NO TARGET-FREE EQUILIBRIUM LEARNING: ADAPTIVE PHASE TRANSITIONS IN ZERO-SUM GAMES

Anonymous authors
Paper under double-blind review

ABSTRACT

Learning a game has no target-independent statistical complexity. Under the same noisy payoff observations, values and low-exploitability decisions may be accurately learned while point equilibrium recovery remains statistically unresolved. We first prove a dimension-uniform accuracy–sensitivity obstruction for every globally Lipschitz single-valued strategy rule, together with an exact two-by-two separation of value, decision, branch, and point risks. We then introduce TARGET, an implementable procedure that discovers an unknown simple equilibrium interface, estimates its information normal and side targets, and adaptively allocates samples. Uniformly over a fixed quantitative class and compact phase ranges, TARGET recovers the active chart, converges to the oracle Gaussian score, and attains the known-geometry minimax interface-side risk and the least-favorable-pair point and confidence benchmarks. Matching lower bounds show that geometric oracle information has zero first-order value, while an explicit competing-interface family identifies the sharp adaptation boundary. Finally, we derive the exact price of sampling before the downstream equilibrium target is known: one experiment is simultaneously oracle optimal for several targets if and only if their normalized information profiles coincide. A certified generator and a fixed-seed study of 320 heterogeneous games verify implementation and exhibit dimension-free collapse onto the predicted information-normal curves.

1 INTRODUCTION

What does it mean to learn an equilibrium from noisy game data? The usual question hides several different statistical targets. The value, a decision with small exploitability, the side of an equilibrium interface, a point strategy, and a downstream policy need not have the same risk, the same limiting experiment, or the same optimal data-acquisition rule. This paper shows that the target is not a secondary choice of loss: it is part of the statistical problem.

The paper is organized around exactly three claims.

Claim 1. *Stable decisions do not imply stable strategies.* Values and exploitability are globally stable under payoff perturbations, but every accurate single-valued strategy rule must be highly sensitive. A two-by-two family has zero value risk and root-n decision loss while branch and point-strategy risks remain order one.

Claim 2. *Unknown simple interfaces are adaptively learnable at oracle risk.* TARGET uses a vanishing full-support pilot to select a finite equilibrium chart, recover its normal and side targets, and allocate the remaining samples. Uniformly over a fixed quantitative class, its score converges in total variation to the known-geometry Gaussian experiment; its design, side, least-favorable-pair point, and confidence risks attain the matching adaptive minimax bounds. A competing-interface family gives the exact root-sample-size boundary for adaptation.

Claim 3. *There is no target-free experiment.* If the downstream target is revealed only after sampling, the exact price of target ignorance is a convex minimax quantity. One design is simultaneously oracle optimal for all targets if and only if their normalized information profiles coincide.

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

The new adaptive theorem is not an abstract boundary-learning claim. Smooth boundary learning and adaptive experimental design have substantial prior literatures (Locatelli et al., 2018; Kaufmann et al., 2016). Here the unknown boundary, normal, targets, and optimal experiment are all induced by an equilibrium-kernel atlas, and the target-dependent Gaussian transition is attained with the same first-order risk as if that geometry had been supplied. Likewise, generic maximin multi-objective design dualities are known (Tommasi et al., 2023); our object is the post-hoc equilibrium-target protocol and its game-derived profiles.

To test whether the theory describes more than one worked example, we provide a certified support-splitting generator and a fixed-seed study of 320 games in dimensions $3, 5, 10$, and 20. The generator is proved to create strict simple interfaces. After normalization, heterogeneous side, point, and confidence risks collapse onto the exact one-dimensional laws, while the excluded regimes fail in the predicted ways. The computation verifies the implementation; it is not used in place of proof.

All uniform adaptive statements are restricted to one fixed quantitative simple-interface class and a compact phase range. We make no prevalence claim, no exact finite-sample adaptive-honesty claim, and no point-risk claim beyond the declared information-balanced pair. Entropy regularization is an upper construction for Claim 1, not a separate approximation-rate contribution. Non-Gaussian transport, semialgebraic certification, and multi-branch extensions are likewise outside the three headline claims.

2 Claim 1: Stable Decisions Do Not Imply Stable Strategies

Let $A \in \mathbb{R}^{m\times k}$ be the payoff matrix of a two-player zero-sum game, with row player maximizing and column player minimizing. Write $\mathrm{val}(A)$ for its value, $\mathcal{E}(A)$ for its saddle set, and $\mathrm{Gap}_A(z)$ for the saddle-point gap of $z = (x, y)$. We use $d_\Delta(z, z') = \frac{1}{2}(\|x - x'\|_1 + \|y - y'\|_1)$.

Proposition 2.1 (Stable functionals)**.** *For all payoff matrices A, B,*

$$|\,\mathrm{val}(A) - \mathrm{val}(B)| \le \|A - B\|_\infty, \qquad z_B \in \mathcal{E}(B) \Longrightarrow \mathrm{Gap}_A(z_B) \le 2\|A - B\|_\infty. \tag{2.1}$$

Proof. For every x, y, $|x^\top(A-B)y| \le \|A-B\|_\infty$. Taking max–min gives the value bound. If $z_B = (x_B, y_B)$ is a saddle of B, then the best-response payoff against y_B is at most $\mathrm{val}(B) + \|A - B\|_\infty$, whereas the worst response to x_B is at least $\mathrm{val}(B) - \|A - B\|_\infty$; subtraction proves the gap bound. □

Stability of decision quality does not produce a stable point selection. For a single-valued strategy rule S, let $L(S)$ be its global Lipschitz modulus from payoff sup norm to d_Δ, and let $\varepsilon_{\mathrm{eq}}(S) = \sup_A \mathrm{Gap}_A(S(A))$ on the unit payoff cube.

Theorem 2.2 (Accuracy–sensitivity obstruction)**.** *For every $m, k \ge 2$ and every globally Lipschitz single-valued strategy rule,*

$$\varepsilon_{\mathrm{eq}}(S) \ge \begin{cases} 1 - L(S), & 0 \le L(S) \le \frac{1}{2}, \\ 1/(4L(S)), & L(S) > \frac{1}{2}, \end{cases} \qquad \varepsilon_{\mathrm{eq}}(S) \ge \frac{1}{4(1 + L(S))}. \tag{2.2}$$

Hence no continuous exact equilibrium selector exists, and every ε-accurate rule with $\varepsilon < 1/2$ has $L(S) \ge 1/(4\varepsilon)$.

Proof sketch. Embed the path $A_t = \left(\begin{smallmatrix} 0 & -t \\ t & 0 \end{smallmatrix}\right)$ and its sign reversal. Their unique equilibria are opposite pure pairs. The sum of the two output gaps is at least $2t\{1 - L(S)t\}$. Maximizing over $0 < t \le 1$ gives equation (2.2); the complete calculation is theorem C.3. The obstruction applies to every globally Lipschitz single-valued strategy rule and is the substantive claim. Entropy is used only to witness a smooth upper construction with the same inverse-sensitivity dependence, up to its logarithmic dimension factor. Nguyen et al. establish the tight finite-game QRE approximation order (Nguyen et al., 2026); no such approximation-rate novelty is claimed here. □

Proposition 2.3 (Entropy is an upper construction)**.** *Let $S_\tau(A)$ be the unique entropy-regularized saddle rule with regularization level $\tau > 0$. Then, globally over payoff matrices,*

$$\mathrm{Gap}_A(S_\tau(A)) \le \tau(\log m + \log k), \qquad d_\Delta(S_\tau(A), S_\tau(B)) \le \frac{\|A - B\|_\infty}{\tau}. \tag{2.3}$$

Thus entropy realizes the inverse-sensitivity dependence in theorem 2.2, up to its logarithmic entropy range; it is not an additional approximation-rate contribution.

Proof sketch. The first inequality is the saddle-gap bias obtained from $0 \le H(x) \le \log m$ and $0 \le H(y) \le \log k$. For the second, compare the two regularized variational inequalities. Bilinear terms cancel, while the symmetrized entropy Bregman divergence is at least the squared product ℓ_1 distance by Pinsker's inequality. The complete boundary and relative-interior argument is theorem C.8. □

The statistical separation is already exact on the same two-by-two path. If $Z_1, \dots, Z_n \overset{\text{iid}}{\sim} N(\theta, 1)$ and $\theta = \pm c/\sqrt{n}$, then the value is identically zero, the constant uniform strategy pair has gap $c/\sqrt{n}$, and the exact branch error is $\Phi(-c)$. On the two point targets at distance $\Delta = 2$ under d_Δ, the exact minimax squared point risk is

$$\frac{\Delta^2}{4} e^{-c^2/2} \mathbb{E}_{G \sim N(0,1)} [\cosh(cG)^{-1}]. \tag{2.4}$$

Thus the same observations give zero value risk, root-n decision loss, and order-one structural risk. Complete finite-sample calculations are in sections B and G.

3 Claim 2: Unknown Interfaces Are Learnable at Oracle Risk

At round t, the learner chooses a payoff coordinate I_t from the past transcript and observes

$$Y_t = A_{I_t} + \sigma_{I_t} \varepsilon_t, \qquad \varepsilon_t \overset{\text{iid}}{\sim} N(0, 1). \tag{3.1}$$

The learner knows the dimensions, noise bounds, and a finite instance-independent kernel atlas. It is not given the interface payoff B, active chart, incident branches, validity coordinate, normal, side targets, or oracle allocation.

The theorem domain. The result is uniform over one fixed quantitative class $\mathfrak{G}(\Gamma)$. Its assumptions are not unnamed genericity conditions; each is a numerical margin that can be checked from the finite regular-kernel atlas.

Quantitative ingredient	**Role in the adaptive theorem**
Bordered-system conditioning	Uniform analytic branch charts and stable inverses for both incident kernels.
Transversality and curvature	A nonzero information normal and a quadratic Newton/interface remainder.
Inactive-slack and atlas margins	Exact finite chart selection rather than pointwise convergence of an unspecified score.
Separated Lipschitz targets	Stable recovery of the two side strategies and a nondegenerate point loss.
Projection tube and bounded local alternatives	A unique valid interface anchor for every $A_{n,G,h} = B + h/\sqrt{n}$, $\|h\|_F \le H$.

The formal class, including every constant and edge convention, is definition N.1.

TARGET: a fully specified non-oracle procedure. The learner is given the dimensions, noise profile, class constants, and the finite instance-independent atlas, but not the active chart or any instance-specific geometry. It performs five measurable stages:

1. draw a deterministic full-support pilot of size n_0 and form the coordinatewise payoff estimate;
2. minimize certified atlas scores with a deterministic tie rule, then take one Newton correction to the selected interface;

3. evaluate the selected branch maps at the projected anchor to estimate $a_G = \nabla g_G(B)$ and the two side targets;
4. use the forced-exploration design

$$\widehat{q}_{n,i} = (1-\epsilon_n)\frac{\sigma_i|\widehat{a}_i|}{\sum_j \sigma_j|\widehat{a}_j|} + \epsilon_n q_i^0, \tag{3.2}$$

and deterministic largest-remainder rounding for the second-stage counts;
5. construct a one-step interface score from the independent second stage, then report the side, posterior-interpolated point target, or three-report confidence region requested by the loss.

Fixed default outputs cover irregular numerical inputs, so the procedure is globally defined and measurable. Forced exploration protects coordinates whose true or estimated normal component is zero; sample splitting makes the conditional score exactly Gaussian.

Define

$$s_G = \sum_i \sigma_i|a_{G,i}|, \qquad q^\star_{G,i} = \frac{\sigma_i|a_{G,i}|}{s_G}, \qquad \kappa_G(\lambda) = \lambda/s_G. \tag{3.3}$$

For the information-balanced local pair, put

$$v^{\mathrm{lf}}_{G,i} = \frac{\sigma_i \operatorname{sgn}(a_{G,i})}{s_G}\mathbf{1}\{a_{G,i}\neq 0\}, \qquad A^{\mathrm{lf}}_{n,G,\eta,\lambda} = B + \eta\frac{\lambda}{\sqrt{n}}v^{\mathrm{lf}}_G, \tag{3.4}$$

where $\eta \in \{-1,+1\}$. Let $z_{G,\eta,n}(\lambda)$ be the corresponding branch target and define

$$\Delta_{G,n}(\lambda) = d_\Delta(z_{G,+,n}(\lambda), z_{G,-,n}(\lambda))\,. \tag{3.5}$$

Theorem 3.1 (Unknown-interface oracle equivalence)**.** *Fix one class $\mathfrak{G}(\Gamma)$, one local radius H, and one compact phase range $0 \le \lambda \le \Lambda$. Run TARGET with $n_0 = d\lceil n^\rho/d\rceil$, $1/2 < \rho < 1$, and $\epsilon_n = n^{-\beta}$, $0 < \beta < \rho$. Then uniformly over the declared class:*

1. *the active chart is recovered with exponentially small error, and*

$$\|\widehat{B}-B\|_F + \|\widehat{a}-a_G\|_F + \max_{s\in\{+,-\}} d_\Delta(\widehat{z}_s, z_{G,s}) = O_P^{\mathrm{unif}}(n_0^{-1/2}\sqrt{\log n} + n^{-1/2}); \tag{3.6}$$

2. *the one-step statistic obeys*

$$\sup_{G,\|h\|_F\le H} d_{\mathrm{TV}}\left(\mathcal{L}_{B+h/\sqrt{n}}(\widehat{T}_n), N\left(\frac{\langle a_G, h\rangle}{s_G}, 1\right)\right) \longrightarrow 0; \tag{3.7}$$

3. *the random design is oracle efficient, and the composite interface-side risk satisfies*

$$R^{\mathrm{TARGET}}_{\mathrm{side,n}}(G,\lambda) \longrightarrow \Phi(-\kappa_G(\lambda)); \tag{3.8}$$

4. *on the information-balanced least-favorable pair, the squared point-strategy risk satisfies*

$$R^{\mathrm{TARGET}}_{\mathrm{point,n}}(G,\lambda) \longrightarrow \frac{\Delta_{G,n}(\lambda)^2}{4} e^{-\kappa_G(\lambda)^2/2}\mathbb{E}[\cosh(\kappa_G(\lambda)Z)^{-1}]; \tag{3.9}$$

5. *the projected three-report region is uniformly $(1-\alpha-o(1))$-honest and has oracle expected diameter*

$$\Delta_{G,n}(\lambda)[\Phi(z_{1-\alpha} - 2\kappa_G(\lambda)) - \alpha]_+. \tag{3.10}$$

The risk benchmarks in items 3–5 are matching lower bounds for every randomized adaptive procedure, even one supplied the true interface geometry. Thus the value of the geometric oracle is asymptotically zero to first order.

Proof sketch. Pilot concentration and the fixed atlas margin give exact discrete chart recovery. The Newton correction cancels the first-order interface residual, leaving an $O_P(r_n^2)$ term; branch smoothness controls the normal and targets. Conditional on the pilot, the independent second-stage

statistic is exactly Gaussian. Forced exploration and rounding make its random variance converge to s_G^2, including at normal-support changes, and equal-variance Gaussian comparison gives equation (3.7). For the lower bound, choose $v_{G,i}^{\mathrm{lf}} = \sigma_i \operatorname{sgn}(a_{G,i})/s_G$ on the normal support. Every informative coordinate then has the same standardized shift, so a parameter-independent kernel simulates the entire adaptive transcript from iid scalar observations. Neyman–Pearson, posterior interpolation, and the exact three-report confidence calculation give the benchmarks. Full finite-sample bounds are in theorems O.5, P.7, P.11 and Q.4. □

Finite-sample bridge to the oracle experiment. The limiting statement is obtained from a nonasymptotic comparison. On the pilot recovery event, conditioning on the pilot transcript $\mathcal{F}_0$ gives

$$\widehat{T}_n \mid \mathcal{F}_0 \sim N(M_{n,G,h}, 1) \tag{3.11}$$

exactly. If the pilot radius is r, its failure probability is at most α, and the exploration level is ϵ, then

$$\sup_{G,\|h\|_F \le H} d_{\mathrm{TV}}\Big(\mathcal{L}_{B+h/\sqrt{n}}(\widehat{T}_n), N(\langle a_G, h\rangle/s_G, 1)\Big) \le \alpha + \frac{\mathfrak{r}_n(r) + \mathfrak{k}_n(r,\epsilon)}{\sqrt{2\pi}}, \tag{3.12}$$

where $\mathfrak{r}_n(r) = O(\sqrt{n}\, r^2 + n^{-1/2})$ is the one-step linearization error and $\mathfrak{k}_n(r,\epsilon) = O(r + r^2/\epsilon + \epsilon + n_1^{-1})$ collects normal estimation, support protection, forced exploration, and integer rounding. The schedules in theorem 3.1 make every term vanish uniformly. This decomposition is the reason the pilot exponent, forced exploration, and independent refinement stage are all needed; see theorem P.7 for explicit constants.

The margin assumptions are information-theoretic rather than cosmetic.

Theorem 3.2 (Exact competing-interface boundary)**.** *There is an explicit pair of three-by-three games whose active charts, unit normals, and oracle designs are separated by constants, while their atlas margin is $\Theta(\delta_N)$. Every adaptive N-query learner satisfies*

$$\max_j P_j\{\widehat{\xi} \ne \xi_j\} \ge \tfrac{1}{4} e^{-N\delta_N^2/2}, \tag{3.13}$$

and balanced queries attain consistent chart, normal, and design recovery if and only if $N\delta_N^2 \to \infty$ on this family. Hence TARGET's root-n_0 atlas-separation requirement is sharp up to the confidence logarithm.

The proof is the adaptive KL chain rule plus an exact Gaussian sign test; see theorem R.4, proposition R.5, and corollary R.8.

4 Cross-Game Universality

The scalar formulas in theorem 3.1 could still appear tied to one hand-designed game. We therefore construct a heterogeneous certified family. Start from an $s \times s$ core C with a unique interior saddle (x, y), value zero, and rank $s - 1$. Duplicate one active row or column, pad to any dimension $d \ge s + 1$ with strictly dominated actions, and permute the actions. The duplicate mass is the side-target separation; the interface normal is the equilibrium weight vector across the duplicate pair.

Theorem 4.1 (Certified generator and universal collapse)**.** *For every admissible support-splitting construction:*

1. *the center is a strict qualitative simple interface and therefore lies in some quantitative class $\mathfrak{G}(\Gamma_G)$;*
2. *under the oracle normal design and information-balanced local pair, the Gaussian tangent score is exactly $N(\eta\kappa, 1)$, while balanced integer-query scores converge to it in total variation;*
3. *after normalization by the target separation Δ_G, the side, point, and confidence laws are*

$$\Phi(-\kappa), \qquad e^{-\kappa^2/2}\mathbb{E}[\cosh(\kappa Z)^{-1}], \qquad [\Phi(z_{1-\alpha} - 2\kappa) - \alpha]_+, \tag{4.1}$$

independent of dimension, support, orientation, noise, and conditioning.

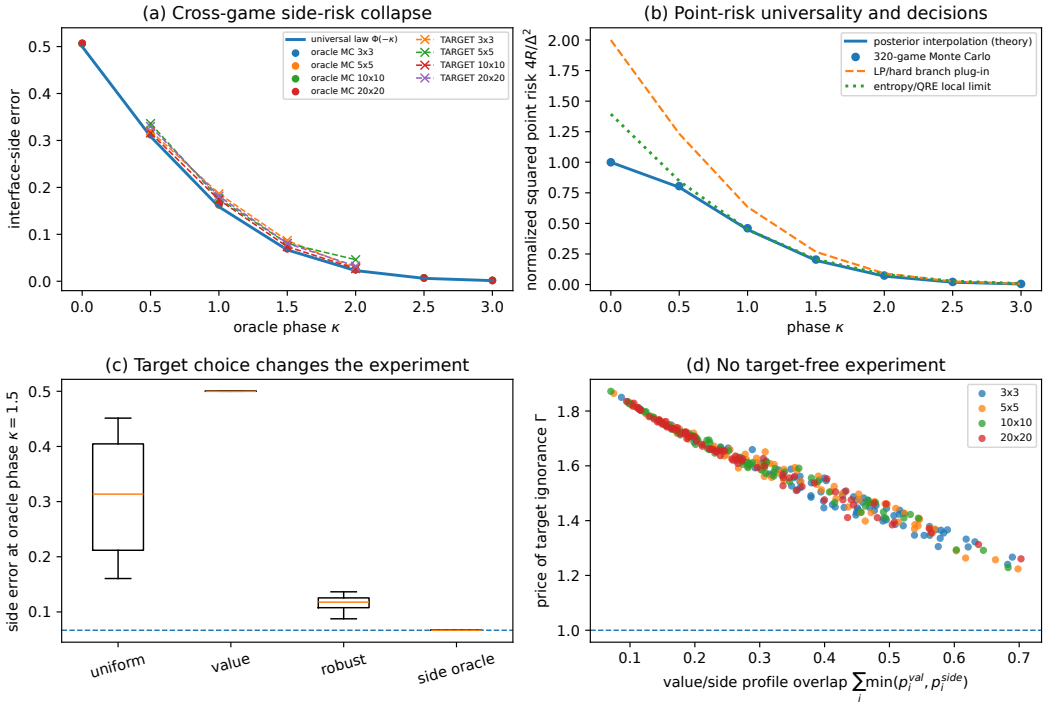

Figure 1: **Cross-game theorem verification.** The fixed-seed corpus contains 320 certified games in dimensions $3, 5, 10, 20$. Oracle and non-oracle side risks collapse onto $\Phi(-\kappa)$; normalized point risks collapse onto the exact posterior law; value, robust post-hoc, and side designs differ; and the target-ignorance price rises as information profiles separate. The markers verify the implementation of proved laws.

Proof sketch. Write $C = U_x M V_y^\top$ with M nonsingular. Positivity of x, y and $\operatorname{rank}(C) = s - 1$ make (x, y) the unique core saddle. At the split center, the equilibrium face consists exactly of all divisions of the duplicated mass. Deleting either duplicate gives the same nonsingular bordered core, so the endpoints are the only incident regular branches. The active slack derivative is the difference of the duplicate rows weighted by y (or duplicate columns weighted by x), and all padding slacks are strict. This proves the simple-interface claim. Sharp Cauchy–Schwarz gives the normal design; the remaining statements reduce to the scalar Gaussian experiment. Complete proofs are in proposition U.2 and theorem U.4. □

The corpus varies active support sizes, row versus column interfaces, equilibrium weights, branch separations, core singular values, coordinate noise, permutations, and information profiles. It contains 2240 oracle curve cells. A separate non-oracle experiment runs TARGET on 12 games per dimension at four phase values. In the fixed run, all adaptive cells selected the correct chart; the median design-variance ratio was 1.0154 and the side-risk root-mean-square error was 0.0354. Across individual oracle cells, the side and normalized point root-mean-square errors were 0.00705 and 0.0111; the largest integer-rounding deviation of the score-variance ratio from one was 5.94×10^{-9}. These are descriptive reproducibility outputs, not replacements for theorem 4.1. Appendix U also supplies a simultaneous finite-sample certificate: for K reported cells and M replications per label, a union of Hoeffding bounds controls every displayed binary mean by $\sqrt{\log(16K/\delta)/(2M)}$ and every normalized point-risk mean by four times that quantity, with probability at least $1 - \delta$.

The negative controls are equally important. Directions on fan faces, higher-order contact, colliding interfaces, nearly singular bordered systems, vanishing side-target separation, and unprotected normal-support changes fail for different reasons. Appendix U proves the generator, states finite Monte Carlo error bounds, reports confidence collapse, and documents all six controls.

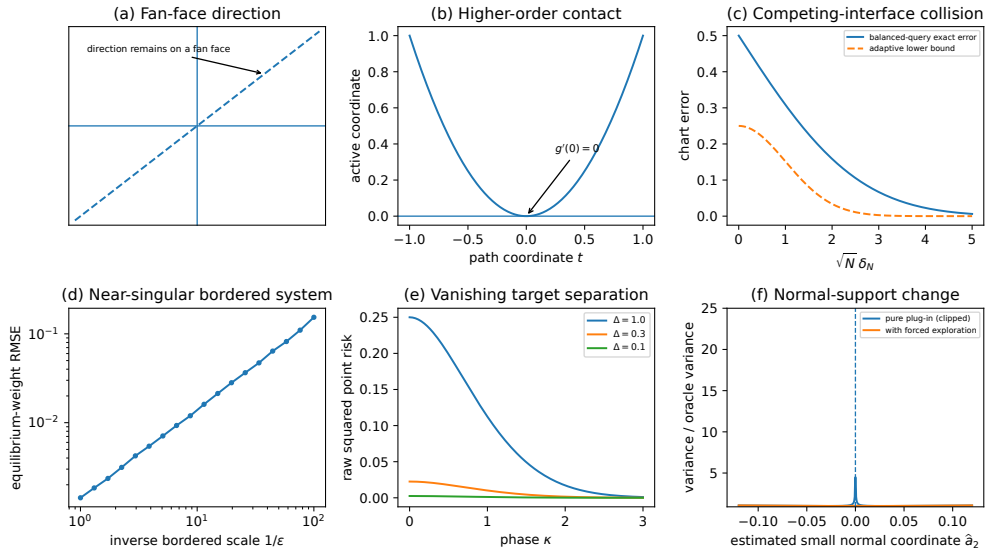

Figure 2: **Negative controls delimit the theorem domain.** Fan-face directions, quadratic contact, competing interfaces, nearly singular bordered systems, vanishing side separation, and unprotected normal-support changes respectively violate the uniqueness, first-order, atlas-margin, conditioning, loss-separation, and exploration hypotheses. The failures are qualitatively different and are not pooled into an unsupported prevalence claim.

5 Claim 3: There Is No Target-Free Experiment

Suppose the sampling design must be chosen before one of R downstream equilibrium targets is revealed. For target r, coordinate derivative $d_{r,i}$, loss matrix $W_r \succeq 0$, and noise σ_i, define

$$c_{r,i} = \sigma_i \|d_{r,i}\|_{W_r}, \qquad p_{r,i} = \frac{c_{r,i}}{\sum_j c_{r,j}}. \tag{5.1}$$

First-order constant targets are removed. Under design q, target r's Gaussian linear risk divided by its oracle minimum is $F_r(q) = \sum_i p_{r,i}^2/q_i$, with the extended convention $c/0 = +\infty$ for $c > 0$. Define

$$\Gamma(p_1, \ldots, p_R) = \min_{q \in \Delta_d} \max_{r \in [R]} F_r(q). \tag{5.2}$$

Theorem 5.1 (Exact price and compatibility)**.** *The primal and dual optima are attained, and*

$$\boxed{\Gamma = \max_{\lambda \in \Delta_R} \left(\sum_i \sqrt{\sum_{r=1}^{R} \lambda_r p_{r,i}^2} \right)^2}. \tag{5.3}$$

For a dual optimizer $\lambda^\star$, an associated robust design is

$$q_i^\star = \frac{\sqrt{\sum_r \lambda_r^\star p_{r,i}^2}}{\sum_j \sqrt{\sum_r \lambda_r^\star p_{r,j}^2}}. \tag{5.4}$$

Complementary slackness identifies the targets that determine the robust experiment:

$$\lambda_r^\star > 0 \quad \Longrightarrow \quad F_r(q^\star) = \Gamma. \tag{5.5}$$

Moreover,

$$1 \le \Gamma \le R, \qquad \Gamma = 1 \iff p_1 = \cdots = p_R, \qquad \Gamma = R \quad \textit{for pairwise disjoint supports.} \tag{5.6}$$

Proof sketch. Express the maximum over targets as a maximum over mixtures $\lambda \in \Delta_R$. Strong duality for the convex epigraph program interchanges the design and mixture optimizations. For fixed λ, sharp Cauchy–Schwarz gives

$$\min_q \sum_i \frac{\sum_r \lambda_r p_{r,i}^2}{q_i} = \left(\sum_i \sqrt{\sum_r \lambda_r p_{r,i}^2} \right)^2,$$

and its equality condition yields equation (5.4). Every $F_r(q) \geq 1$, with equality only at $q = p_r$, proving compatibility. The average-profile design gives the upper bound R; disjoint supports attain it. Full boundary and complementary-slackness arguments are in theorem S.2 and corollary S.3. □

The equilibrium profiles are concrete. At a unique regular equilibrium they have the following coordinate weights before normalization:

Target	**Coordinate weight** c_{ij}	**Geometry encoded**
Value	$\sigma_{ij} x_i^\star y_j^\star$	Equilibrium occupancy of the payoff entry.
Interface side	$\sigma_{ij} \lvert a_{ij} \rvert$	Normal to the active equilibrium interface.
Point strategy	$\sigma_{ij} \lVert d_{ij} \rVert_W$	Local strategy sensitivity under the declared metric.
Policy target	$\sigma_{ij} \lVert d_{ij}^{\mathrm{pol}} \rVert_W$	Statewise or occupancy-weighted policy sensitivity.

One experiment is therefore oracle optimal for all of these objects exactly when occupancy, interface normal, strategy sensitivity, and policy sensitivity induce the same normalized profile. For several interface targets with common oracle phase κ, the robust design has exact worst-case side error $\Phi(-\kappa/\sqrt{\Gamma})$. In the 320-game corpus, $1.2238 \leq \Gamma \leq 1.8720$ for the value–side pair.

The theorem is operational for integer queries. If $q_{\min}^\star = \min_{i:q_i^\star>0} q_i^\star$, largest-remainder rounding to $N \geq 2/q_{\min}^\star$ observations gives proportions $q^{(N)}$ satisfying

$$\max_r F_r(q^{(N)}) \leq \Gamma + \frac{2}{N(q_{\min}^\star)^2} \max_r \sum_i p_{r,i}^2. \tag{5.7}$$

Thus the exact approximate-design penalty is implementable with $O(N^{-1})$ degradation, rather than being only an asymptotic convex program.

6 Scope, Related Work, and Consequences

The three claims have deliberately different scopes. The selection obstruction is global and deterministic. The adaptive theorem is uniform only over a fixed class $\mathfrak{G}(\Gamma)$ and a compact phase interval. Its side statement is composite over bounded half-spaces, whereas its point and confidence equalities concern the declared information-balanced pair. Honest adaptive coverage is asymptotic and uniform, not an exact finite-sample identity. The target-ignorance theorem is an exact finite-dimensional design result once the derivative profiles are declared.

Claim	**Theorem domain**	**Type of conclusion**
1	All finite zero-sum payoff matrices and all globally Lipschitz single-valued strategy rules.	Exact deterministic lower bound; exact finite-sample two-point risks; entropy only as an upper witness.
2	One fixed quantitative simple-interface class, bounded local alternatives, and a compact phase interval.	Finite recovery and total-variation bounds; uniform asymptotic oracle efficiency; exact lower benchmark on the declared least-favorable pair.
3	Any finite collection of nonconstant derivative profiles under the declared quadratic losses.	Exact attained primal–dual identity, equality cases, and integer implementation.

This separation prevents a pairwise point equality, a composite side theorem, and a global selector obstruction from being read as the same quantifier statement.

Adaptive learning of smooth decision boundaries is established in other settings (Locatelli et al., 2018), and general adaptive change-of-measure tools are standard in sequential identification (Kaufmann et al., 2016). The equilibrium-specific contribution is the finite regular-kernel atlas, recovery of game-derived side targets and information normal, exact oracle transition, and matching zero-oracle-value result. Multi-objective optimal-design duality also has antecedents (Tommasi et al., 2023); the new object here is the post-hoc selection among equilibrium targets.

Entropy regularization is not a fourth contribution. Nguyen et al. prove the tight $\Theta(\beta^{-1}(\log m + \log k))$ finite-game QRE approximation order (Nguyen et al., 2026). Our distinct result is the dimension-uniform lower obstruction for every globally Lipschitz single-valued strategy rule; entropy appears only as one upper construction matching the obstruction's inverse-sensitivity dependence up to the logarithmic dimension factor.

No theorem asserts that simple interfaces are globally prevalent. Every strict qualitative simple interface belongs to some quantitative class after localization, but no adaptive result is uniform over the uncontrolled union of all margins. Higher-order contact, simultaneous interfaces, singular-kernel creation, vanishing target separation, and trajectory-constrained exploration are outside the headline theorem.

Statewise equilibrium singularities propagate through finite-horizon Bellman recursion under sectorwise separation; this finite-horizon, generative-oracle, metric-qualified result is retained only in Appendix J, while a trajectory-data theory in which reachability and exploration jointly determine effective information is a separate project.

Proof and implementation map. The logical chain is auditable. First, the finite kernel atlas and the fixed margins convert unknown equilibrium geometry into a separated model-selection problem. Second, the one-step projection and sample split reduce the selected chart to an exactly conditional Gaussian score with an explicit total-variation remainder. Third, the information-balanced local pair makes every informative coordinate carry the same standardized signal, which gives an exact scalar simulation of every adaptive transcript and hence the matching minimax lower bounds. Fourth, convex epigraph duality and sharp Cauchy–Schwarz produce the target-ignorance formula and its equality cases. Finally, the certified support-splitting generator embeds these ingredients across dimensions and noise profiles; the simulations exercise the proved implementation rather than supply evidence for an unproved theorem.

Methodological consequence. A report of "equilibrium-learning complexity" is incomplete unless it names the requested equilibrium object, the loss, and the experiment. When point selection is unresolved, low-gap or set-valued reporting may remain valid; when the downstream target is unknown, the unavoidable price is Γ; and when the interface itself is unknown but quantitatively separated, TARGET pays no first-order geometric-oracle cost. These conclusions are different faces of the same target-indexed decision problem, not separate applications.

Sharp boundaries of the conclusion. The positive and negative theories meet at the scale of the declared margins. On the explicit collision family, chart, normal, and oracle-design recovery are consistent exactly when $N\delta_N^2 \to \infty$; bounded $N\delta_N^2$ leaves a nonvanishing testing error. Vanishing side-target separation can make raw point loss small without making the chart identifiable, while a normal-support change requires forced exploration to preserve the score variance. Finally, $\Gamma = 1$ only under profile alignment; this is an exact compatibility condition, not a heuristic recommendation to use a single design. These boundary cases explain which assumption would have to be strengthened before any broader theorem could be stated.

The resulting message is narrower than a universal theory of games and stronger than an oracle local calculation: a learner can discover a regular unknown equilibrium interface and attain its oracle transition, but neither strategy stability nor experimental optimality exists independently of the requested equilibrium target.

AI USE STATEMENT

Generative AI tools assisted with research ideation, literature discovery, theorem formulation, portions of proof drafting, exposition, LaTeX editing, reference formatting, figure preparation, and code. The authors independently checked the final mathematical statements and proofs, verified every external theorem invocation against its cited primary source, recomputed the closed-form risk and design formulas analytically and numerically, reviewed the supplementary code, and rebuilt the paper and supplement from clean archives. The accompanying verification suite regenerates all figures and reports and checks the declared numerical tolerances without network access. The authors take responsibility for all final claims, proofs, citations, code, and artifacts.

REPRODUCIBILITY STATEMENT

This is a theoretical paper. The appendices state every assumption, provide complete proofs, and give the finite procedures used to verify the branch conditions. Appendices M and U, together with the anonymous `reproducibility/` package, enumerate kernels, verify validity gradients, regenerate the certified 320-game corpus and all displayed curves, solve the design programs, and run the theorem checks from fixed seeds without external data or network access.

REFERENCES

Hirotachi Abo, Irem Portakal, and Luca Sodomaco. A vector bundle approach to Nash equilibria. *Advances in Applied Mathematics*, 175:103028, 2026. doi: 10.1016/j.aam.2025.103028.

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

Timothy B. Armstrong. Asymptotic efficiency bounds for a class of experimental designs, 2022. arXiv:2205.02726; version 3 dated May 18, 2026.

Luc Attia and Miquel Oliu-Barton. Shapley–Snow kernels, multiparameter eigenvalue problems, and stochastic games. *Mathematics of Operations Research*, 46(3):1181–1202, 2021. doi: 10.1287/moor.2020.1104.

Luc Attia, Miquel Oliu-Barton, and Raimundo Saona. Marginal values of a stochastic game. *Mathematics of Operations Research*, 50(1):482–505, 2025. doi: 10.1287/moor.2023.0297.

Yu Bai and Chi Jin. Provable self-play algorithms for competitive Reinforcement Learning. In Hal Daumé III and Aarti Singh (eds.), *Proceedings of the 37th International Conference on Machine Learning*, volume 119 of *Proceedings of Machine Learning Research*, pp. 551–560. PMLR, 2020. URL `https://proceedings.mlr.press/v119/bai20a.html`.

Saugata Basu, Richard Pollack, and Marie-Françoise Roy. *Algorithms in Real Algebraic Geometry*, volume 10 of *Algorithms and Computation in Mathematics*. Springer, Berlin, 2 edition, 2006. doi: 10.1007/3-540-33099-2.

Shicong Cen, Yuting Wei, and Yuejie Chi. Fast policy extragradient methods for competitive games with entropy regularization. In *Advances in Neural Information Processing Systems*, volume 34, pp. 27952–27964, 2021.

Herman Chernoff. Sequential design of experiments. *The Annals of Mathematical Statistics*, 30(3): 755–770, 1959. doi: 10.1214/aoms/1177706205.

Imre Csiszár. Information-type measures of difference of probability distributions and indirect observations. *Studia Scientiarum Mathematicarum Hungarica*, 2:299–318, 1967.

Ruoyu Diao, Yu-Hong Dai, and Liwei Zhang. Stability for Nash equilibrium problems. *Mathematics of Operations Research*, 51(3):2017–2040, 2026. doi: 10.1287/moor.2024.0609. Published online July 2, 2025.

David L. Donoho. Statistical estimation and optimal recovery. *The Annals of Statistics*, 22(1): 238–270, 1994. doi: 10.1214/aos/1176325367.

Asen L. Dontchev and R. Tyrrell Rockafellar. Characterizations of strong regularity for variational inequalities over polyhedral convex sets. *SIAM Journal on Optimization*, 6(4):1087–1105, 1996. doi: 10.1137/S1052623495284029.

Asen L. Dontchev and R. Tyrrell Rockafellar. Ample parameterization of variational inclusions. *SIAM Journal on Optimization*, 12(1):170–187, 2001. doi: 10.1137/S1052623400371016.

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

Zheng Fang and Andres Santos. Inference on directionally differentiable functions. *The Review of Economic Studies*, 86(1):377–412, 2019. doi: 10.1093/restud/rdy049.

Herbert Federer. Curvature measures. *Transactions of the American Mathematical Society*, 93(3): 418–491, 1959. doi: 10.1090/S0002-9947-1959-0110078-1.

Aurélien Garivier and Emilie Kaufmann. Optimal best arm identification with fixed confidence. In *Proceedings of the 29th Annual Conference on Learning Theory*, volume 49 of *Proceedings of Machine Learning Research*, pp. 998–1027. PMLR, 2016. URL https://proceedings.mlr.press/v49/garivier16a.html.

Richard D. Gill and Boris Y. Levit. Applications of the van Trees inequality: A Bayesian Cramér–Rao bound. *Bernoulli*, 1(1–2):59–79, 1995. doi: 10.2307/3318681.

Hédi Hadiji, Sarah Sachs, Tim van Erven, and Wouter M. Koolen. Towards characterizing the first-order query complexity of learning (approximate) Nash equilibria in zero-sum matrix games. In *Advances in Neural Information Processing Systems*, volume 36, pp. 13356–13373, 2023. doi: 10.52202/075280-0586.

Jaroslav Hájek. Local asymptotic minimax and admissibility in estimation. In Lucien M. Le Cam, Jerzy Neyman, and Elizabeth L. Scott (eds.), *Proceedings of the Sixth Berkeley Symposium on Mathematical Statistics and Probability, Volume 1: Theory of Statistics*, pp. 175–194. University of California Press, Berkeley, California, 1972. URL https://projecteuclid.org/euclid.bsmsp/1200514092.

Keisuke Hirano and Jack R. Porter. Asymptotic representations for sequential decisions, adaptive experiments, and batched bandits, 2023. arXiv:2302.03117; version 2 dated February 22, 2025.

Jiashuo Jiang and Mengxiao Zhang. An LP-based approach for bilinear saddle point problem with instance-dependent guarantee and noisy feedback, 2026. arXiv preprint.

Emilie Kaufmann, Olivier Cappé, and Aurélien Garivier. On the complexity of best-arm identification in multi-armed bandit models. *Journal of Machine Learning Research*, 17(1):1–42, 2016. URL https://jmlr.org/papers/v17/kaufman16a.html.

Alexander S. Kechris. *Classical Descriptive Set Theory*, volume 156 of *Graduate Texts in Mathematics*. Springer, New York, 1995. doi: 10.1007/978-1-4612-4190-4.

Jack Kiefer and Jacob Wolfowitz. The equivalence of two extremum problems. *Canadian Journal of Mathematics*, 12:363–366, 1960. doi: 10.4153/CJM-1960-030-4.

Eric B. Laber, Min Qian, Daniel J. Lizotte, William E. Pelham, and Susan A. Murphy. Statistical inference in dynamic treatment regimes, 2010. arXiv preprint.

Lucien Le Cam. *Asymptotic Methods in Statistical Decision Theory*. Springer Series in Statistics. Springer, New York, 1986. doi: 10.1007/978-1-4612-4946-7.

Luofeng Liao and Christian Kroer. Statistical inference and A/B testing in Fisher markets and paced auctions. *Management Science*, 2025. Forthcoming; arXiv:2406.15522.

Michael L. Littman. Markov games as a framework for Multi-Agent Reinforcement Learning. In *Machine Learning Proceedings 1994*, pp. 157–163. Morgan Kaufmann, 1994. doi: 10.1016/B978-1-55860-335-6.50027-1.

Qinghua Liu, Tiancheng Yu, Yu Bai, and Chi Jin. A sharp analysis of model-based Reinforcement Learning with self-play. In Marina Meila and Tong Zhang (eds.), *Proceedings of the 38th International Conference on Machine Learning*, volume 139 of *Proceedings of Machine Learning Research*, pp. 7001–7010. PMLR, 2021. URL https://proceedings.mlr.press/v139/liu21z.html.

Shuyu Liu, Florentina Bunea, and Jonathan Niles-Weed. Asymptotic confidence sets for random linear programs. In Gergely Neu and Lorenzo Rosasco (eds.), *Proceedings of the Thirty Sixth Conference on Learning Theory*, volume 195 of *Proceedings of Machine Learning Research*, pp. 3919–3940. PMLR, 2023. URL https://proceedings.mlr.press/v195/liu23d.html.

Andrea Locatelli and Alexandra Carpentier. Adaptivity to smoothness in x-armed bandits. In Sébastien Bubeck, Vianney Perchet, and Philippe Rigollet (eds.), *Proceedings of the 31st Conference on Learning Theory*, volume 75 of *Proceedings of Machine Learning Research*, pp. 1463–1492. PMLR, 2018. URL https://proceedings.mlr.press/v75/locatelli18a.html.

Andrea Locatelli, Alexandra Carpentier, and Samory Kpotufe. An adaptive strategy for active learning with smooth decision boundary. In Firdaus Janoos, Mehryar Mohri, and Karthik Sridharan (eds.), *Proceedings of Algorithmic Learning Theory*, volume 83 of *Proceedings of Machine Learning Research*, pp. 547–571. PMLR, 2018. URL https://proceedings.mlr.press/v83/locatelli18a.html.

Mark G. Low. On nonparametric confidence intervals. *The Annals of Statistics*, 25(6):2547–2554, 1997. doi: 10.1214/aos/1030741084.

Nan Lu, Ethan Lee, James M. Robins, David Simchi-Levi, and Junwei Lu. Optimal value inference for Reinforcement Learning, 2026. arXiv preprint.

Shu Lu. Symmetric confidence regions and confidence intervals for normal map formulations of stochastic variational inequalities. *SIAM Journal on Optimization*, 24(3):1458–1484, 2014. doi: 10.1137/13090506X.

Arnab Maiti, Kevin Jamieson, and Lillian Ratliff. Instance-dependent sample complexity bounds for zero-sum matrix games. In Francisco Ruiz, Jennifer Dy, and Jan-Willem van de Meent (eds.), *Proceedings of the 26th International Conference on Artificial Intelligence and Statistics*, volume 206 of *Proceedings of Machine Learning Research*, pp. 9429–9469. PMLR, 2023. URL https://proceedings.mlr.press/v206/maiti23a.html.

Arnab Maiti, Ross Boczar, Kevin Jamieson, and Lillian Ratliff. Near-optimal pure exploration in matrix games: A generalization of stochastic bandits and dueling bandits. In Sanjoy Dasgupta, Stephan Mandt, and Yingzhen Li (eds.), *Proceedings of the 27th International Conference on Artificial Intelligence and Statistics*, volume 238 of *Proceedings of Machine Learning Research*, pp. 2602–2610. PMLR, 2024. URL https://proceedings.mlr.press/v238/maiti24a.html.

Arnab Maiti, Ross Boczar, Kevin Jamieson, and Lillian J. Ratliff. Query-efficient algorithm to find all Nash equilibria in a two-player zero-sum matrix game. *ACM Transactions on Economics and Computation*, 13(3):1–18, 2025. doi: 10.1145/3732787. Article 15.

Richard D. McKelvey and Thomas R. Palfrey. Quantal response equilibria for normal form games. *Games and Economic Behavior*, 10(1):6–38, 1995. doi: 10.1006/game.1995.1023.

Mika Meitz and Alexander Shapiro. Minimax asymptotics. *Electronic Journal of Statistics*, 19(2): 5117–5146, 2025. doi: 10.1214/25-EJS2457.

Yurii Nesterov. Smooth minimization of non-smooth functions. *Mathematical Programming*, 103: 127–152, 2005. doi: 10.1007/s10107-004-0552-5.

Jerzy Neyman and Egon Sharpe Pearson. Ix. on the problem of the most efficient tests of statistical hypotheses. *Philosophical Transactions of the Royal Society of London. Series A, Containing Papers of a Mathematical or Physical Character*, 231(694–706):289–337, 1933. doi: 10.1098/rsta.1933.0009.

Khang Nguyen, Valentio Iverson, Sahan Wijetunga, William Chang, and Guillaume Wang. Tight rates of approximation of mixed Nash equilibria by entropy regularization in continuous games. In *Proceedings of the 42nd Conference on Uncertainty in Artificial Intelligence*, volume 337 of *Proceedings of Machine Learning Research*, pp. 4855–4872. PMLR, 2026. URL https://proceedings.mlr.press/v337/nguyen26b.html.

Wenhua Nie, Binhan Luo, Zijie Meng, Jyh-Shing Roger Jang, and Ching-Wen Ma. Equilibrium residuals expose three regimes of matrix-game strategic reasoning in language models, 2026. arXiv preprint.

James M. Robins and Aad W. van der Vaart. Adaptive nonparametric confidence sets. *The Annals of Statistics*, 34(1):229–253, 2006. doi: 10.1214/009053605000000877.

Stephen M. Robinson. Strongly regular generalized equations. *Mathematics of Operations Research*, 5(1):43–62, 1980. doi: 10.1287/moor.5.1.43.

R. Tyrrell Rockafellar and Roger J.-B. Wets. *Variational Analysis*, volume 317 of *Grundlehren der mathematischen Wissenschaften*. Springer, Berlin, 1998. doi: 10.1007/978-3-642-02431-3.

Lloyd S. Shapley. Stochastic games. *Proceedings of the National Academy of Sciences of the United States of America*, 39(10):1095–1100, 1953. doi: 10.1073/pnas.39.10.1095.

Lloyd S. Shapley and Roslyn N. Snow. Basic solutions of discrete games. In Harold W. Kuhn and Albert W. Tucker (eds.), *Contributions to the Theory of Games, Volume I*, volume 24 of *Annals of Mathematics Studies*, pp. 27–35. Princeton University Press, Princeton, NJ, 1950. doi: 10.1515/9781400881727-004.

Masahiro Shiota. Whitney triangulations of semialgebraic sets. *Annales Polonici Mathematici*, 87 (1):237–246, 2005.

Maurice Sion. On general minimax theorems. *Pacific Journal of Mathematics*, 8(1):171–176, 1958. doi: 10.2140/pjm.1958.8.171.

Kenta Takatsu and Arun Kumar Kuchibhotla. Generalized van Trees inequality: Local minimax bounds for non-smooth functionals and irregular statistical models, 2024. arXiv preprint.

Chiara Tommasi, Juan M. Rodríguez-Díaz, and Jesús F. López-Fidalgo. An equivalence theorem for design optimality with respect to a multi-objective criterion. *Statistical Papers*, 64:1041–1056, 2023. doi: 10.1007/s00362-023-01431-2.

Zhixian Zhang, Xiaotian Hou, and Linjun Zhang. Unified inference framework for single and multi-player performative prediction: Method and asymptotic optimality, 2026. arXiv preprint.

A Notation, Standing Assumptions, and Proof Dependencies

This appendix fixes the conventions used throughout the paper and records the assumptions under which the main results are stated. It introduces no new scientific claim. All asymptotics are fixed-dimensional unless explicitly stated otherwise, all local statements are relative to their declared perturbation model, and all uniformity claims name the set over which the remainder or risk is controlled.

A.1 Basic notation

For $r \in \mathbb{N}$, write $[r] = \{1, \dots, r\}$ and

$$\Delta_r = \{p \in \mathbb{R}^r_+ : \mathbf{1}^\mathsf{T} p = 1\}.$$

Vectors are columns, e_i is the ith standard basis vector, and $\langle u, v\rangle = u^\mathsf{T} v$. For matrices, $\|A\|_\infty$ is the entrywise maximum norm and $\|A\|_F$ the Frobenius norm. Column-wise vectorization is $\mathrm{vec}(\cdot)$; $M^\dagger$ is the Moore–Penrose inverse. For $W \succ 0$, $\|u\|_W^2 = u^\mathsf{T} W u$.

For the product strategy space $\mathcal{Z} = \Delta_m \times \Delta_k$, use

$$d_\Delta((x,y),(x',y')) = \frac{1}{2}\big(\|x-x'\|_1 + \|y-y'\|_1\big). \tag{A.1}$$

Its diameter is two. Every point-strategy risk and confidence diameter in the paper uses this normalization. The notation $r_n = O_K(a_n)$ means a remainder bounded by $c_K a_n$ uniformly on K for all sufficiently large n.

A.2 Global terminology and reserved symbols

The following terms are used with fixed meanings throughout the manuscript.

Strategy rule. A measurable single-valued map from observations or payoffs to $\mathcal{Z}$; it need not return an equilibrium.

Approximate-equilibrium rule. A strategy rule accompanied by a stated exploitability guarantee.

Equilibrium selector. A single-valued map $A \mapsto z(A)$ satisfying $z(A) \in \mathcal{E}(A)$ for every payoff in its domain.

Branch label. The label of a local analytic equilibrium germ (after quotienting duplicate kernel representations when this is relevant).

Support label. The pair $(\mathrm{supp}(x), \mathrm{supp}(y))$. A branch label and a support label need not coincide globally and are never used as synonyms.

Point strategy. One element of $\mathcal{Z}$; this is distinct from the set-valued equilibrium target $\mathcal{E}(A)$.

Local-oracle design. A design depending on the declared reference game, perturbation model, target derivative, or bifurcation normal.

Implementable design. Integer query counts or a specified adaptive allocation rule. An approximate design is not called implementable until its rounding or sampling rule is stated.

Strictly valid regular kernel branch. A kernel branch with nonzero bordered determinant and strictly positive validity vector. The shorter phrase "regular equilibrium" is not used for this stronger property.

Target-specific nonregularity. A failure of regular estimation for a named target and loss. The paper never calls a game "nonregular" without identifying the target under discussion.

The following symbols are reserved to avoid collisions:

$$\begin{aligned} \varepsilon_{\mathrm{reg}} \text{ (regularization strength)}, \qquad & s_a \text{ or } s_{a,n} \text{ (information-normal scale)}, \\ p_{\mathrm{reach}} \text{ (reach probability)}, \qquad & r_{\mathrm{noise}} \text{ (variance multiplier)}, \\ \mathfrak{C} \text{ (confidence region)}, \qquad & D_\Psi \text{ (target derivative)}, \\ \dot{M} \text{ (primitive Markov-model direction)}, \qquad & \mathcal{S}^{\mathrm{sol}} \text{ (localized solution map)}. \end{aligned} \tag{A.2}$$

Branch chambers retain the notation $\mathcal{C}$, kernel indices retain β, and validity vectors retain ρ_β.

A.3 Claim-scope conventions

The adjectives used in theorem titles have the following formal meaning.

Exact. A finite-sample equality for the observation model, parameter class, target, loss, and scaling stated in the same theorem. An equality on a two-point subexperiment becomes only a lower bound on a larger class.

Sharp asymptotic. Equality of a declared limiting minimax constant; it is not a finite-sample identity.

Matching order. Equality only in the named rate or exponent. In particular, entropy matches the inverse-sensitivity $1/L$ dependence only up to the factor $\log(mk)$.

Local-oracle optimal. Optimality for a fixed reference model and a specified target/loss criterion. It does not imply global adaptive implementability.

Transport under verified equivalence. Extension to another observation model only after two-sided Le Cam equivalence to the stated Gaussian experiment has been established separately.

Every theorem statement identifies its observation model, parameter class, known quantities, target, loss and scaling, finite-sample or asymptotic status, type of conclusion, uniformity set, and excluded boundary cases either in the statement itself or in an immediately preceding scope paragraph.

A.4 Static zero-sum games and targets

The row player maximizes and the column player minimizes. We write

$$\mathrm{val}(A) = \max_{x \in \Delta_m} \min_{y \in \Delta_k} x^{\mathsf{T}} A y, \tag{A.3}$$

$$\mathrm{Gap}_A(x, y) = \max_{x'} x'^{\mathsf{T}} A y - \min_{y'} x^{\mathsf{T}} A y', \tag{A.4}$$

and $\mathcal{E}(A) = \{z : \mathrm{Gap}_A(z) = 0\}$. In zero-sum games, $\mathcal{E}(A) = \mathcal{X}^\star(A) \times \mathcal{Y}^\star(A)$. The canonical selector is the maximum-entropy equilibrium

$$\mathrm{sel}(A) = \left(\operatorname*{arg\,max}_{x \in \mathcal{X}^\star(A)} \mathsf{Ent}(x), \operatorname*{arg\,max}_{y \in \mathcal{Y}^\star(A)} \mathsf{Ent}(y) \right). \tag{A.5}$$

It is single-valued by strict concavity of entropy, but no continuity is assumed.

The principal targets are distinct: scalar value, exploitability of an output, a point strategy, a branch label, a support label, and the equilibrium set itself. A result for one target never transfers implicitly to another. Regular point strategy risk is n-scaled, jump point-strategy risk is unscaled, branch-label risk uses 0–1 loss, and exploitability is $\sqrt{n}$-scaled in the local separation theorems.

A.5 Variational conventions

For a closed convex set C,

$$T_C(\bar{z}) = \{u : \exists t_n \downarrow 0,\ u_n \to u,\ \bar{z} + t_n u_n \in C\}, \tag{A.6}$$

$$N_C(\bar{z}) = \{w : \langle w, z - \bar{z} \rangle \le 0\ \forall z \in C\}. \tag{A.7}$$

For a set-valued map F, its graphical derivative is

$$DF(\bar{x} \mid \bar{y})(u) = \{v : (u, v) \in T_{\mathrm{gph}\, F}(\bar{x}, \bar{y})\}.$$

Given $\bar{w} \in N_C(\bar{z})$, the critical cone is $T_C(\bar{z}) \cap \bar{w}^{\perp}$. For matrix games,

$$F_A(x, y) = (-Ay, A^{\mathsf{T}} x), \qquad z \in \mathcal{E}(A) \iff 0 \in F_A(z) + N_{\mathcal{Z}}(z).$$

The normal-cone polarity above fixes every sign in Appendix E.

A relative Hadamard derivative is defined by the joint limit $t \downarrow 0$, $h' \to h$ within the declared tangent model. It is stronger than a fixed-direction Gâteaux limit.

A.6 Local Gaussian experiment

Fix a reference game A_0 and a d-dimensional perturbation space $\mathbb{T}$ with Frobenius-orthonormal basis $H^1, \ldots, H^d$. Put

$$\mathcal{U}h = \sum_{\ell=1}^{d} h_\ell H^\ell, \qquad U = [\mathrm{vec}(H^1)\ \cdots\ \mathrm{vec}(H^d)], \qquad A_{n,h} = A_0 + \frac{\mathcal{U}h}{\sqrt{n}}.$$

The reference game is relatively interior along $\mathbb{T}$, so fixed compact sets of h are feasible for all large n. Under deterministic query counts n_{ij} and known positive variances,

$$D_n = \mathrm{diag}\left(\frac{n_{ij}}{n \sigma_{ij}^2} \right), \qquad J_n = U^{\mathsf{T}} D_n U, \qquad J_n \to J \succ 0. \tag{A.8}$$

The exact sufficient statistic satisfies $Z_n \sim \mathcal{N}(h, J_n^{-1})$. The center, tangent model, basis, and design are fixed components of the local experiment. The local theory does not claim global adaptation to an unknown center or unknown branch geometry.

For bounded-loss minimax transfer, the loss is uniformly continuous in its target argument. For quadratic loss, the tangent target is continuous on a compact local set or the theorem uses the expanding-ball order of limits $n \to \infty$ first and radius $R \to \infty$ second.

A.7 Branch and jump assumptions

A regular kernel branch is defined by a nonsingular bordered primal–dual system. Its validity vector contains supported probabilities and off-support reduced slacks. The regular-kernel validity-refinement theory uses exactly:

Regular-kernel completeness (assumption D.5). Every nearby basic equilibrium in $A_0 + \mathbb{T}$ is represented by a kernel already nonsingular at A_0.

First-order transversality (assumption D.7). Every active validity coordinate has a nonzero differential on $\mathbb{T}$.

Uniform expansions hold only on compact normalized direction sets strictly inside full-dimensional validity chambers. No first-order branch assertion is made on the validity-refinement skeleton.

The simple-jump theory additionally assumes assumption G.1: a C^2 defining function separates two locally unique branches with distinct limits. In local coordinates,

$$a_\ell = Dg(A_0)[H^\ell], \qquad s_{a,n}^2 = a^\mathsf{T} J_n^\dagger a, \qquad \kappa_n = \frac{\lambda_n}{s_{a,n}},$$

whenever $a \in \mathrm{range}(J_n)$; otherwise the branch classes are unidentified and we set $s_{a,n} = +\infty$. The least-favorable points induced by J_n are

$$h_{n,\pm}^\star = \pm\lambda_n \frac{J_n^\dagger a}{a^\mathsf{T} J_n^\dagger a}.$$

The fixed parameter classes used to compare experiments are the half-spaces $\{h : \pm a^\mathsf{T} h \geq \lambda_n\}$; the least-favorable pair is a proof object and may depend on the design. The information-normal phase parameter is κ_n.

A.8 Confidence, design, and Markov conventions

An honest point-target region satisfies uniform coverage of the point target; a simultaneous set-valued region covers the entire equilibrium set. The exact confidence-diameter equality in Appendix H is for the declared two-point jump experiment and becomes a lower bound on any larger parameter class.

For experimental design, entry a contributes $J_a = u_a u_a^\mathsf{T}/\sigma_a^2$ and $J(q) = \sum_a q_a J_a$. For derivative D_Ψ and loss W, target identifiability is $\mathrm{range}(D_\Psi^\mathsf{T}) \subseteq \mathrm{range}(J(q))$. The equivalence theorem is stated on the positive definite information region, and the closed-form allocation requires unrestricted entrywise perturbations. The adaptive attainment theorem is pointwise at a fixed regular game, not uniform across a jump.

The Markov extension is finite-horizon, finite-state/action, two-player zero-sum, and generative-oracle based. Primitive transition directions have row sum zero and are relative to a feasible affine transition model. Appendix J proves the directional Bellman recursion and records only a conditional sectorwise branch consequence under explicit statewise chamber separation; no global polyhedral fan is claimed at a dynamically singular state.

A.9 Dependency map and scope

The proof architecture is

$$\begin{aligned} \text{kernel branches} &\Longrightarrow \text{validity geometry} \Longrightarrow \text{target charts}, \\ \text{Gaussian score reduction} &\Longrightarrow \text{target-indexed risks and design}, \\ \text{matrix-game calculus} &\Longrightarrow \text{conditional Bellman pullbacks}. \end{aligned} \tag{A.9}$$

The complete proofs of the main-text results are in Appendices B–J. Appendix K records only proved extensions and Appendix L audits novelty claims.

Unless explicitly stated otherwise, $m, k, d, T, |\mathcal{S}|$, the reference model, and the noise variances are fixed as $n \to \infty$. The paper does not establish general-sum, growing-dimensional, partially observable, continuous-action, higher-order singularity, or trajectory-constrained design theory.

B Exact Analysis of the Two-by-Two Family

Throughout, $[d] \coloneqq \{1, \dots, d\}$. This appendix gives an exact, finite-sample analysis of the one-parameter family

$$A_\theta \coloneqq \begin{pmatrix} 0 & \theta \\ -\theta & 0 \end{pmatrix}, \qquad \theta \in [-1, 1]. \tag{B.1}$$

The family is the canonical local model of a transverse jump bifurcation: the value is constant, the unique equilibrium switches discontinuously at $\theta = 0$, and the complete equilibrium set at the bifurcation is $\Delta_2 \times \Delta_2$. We analyze value estimation, approximate-equilibrium decision loss, equilibrium-branch recovery, and point strategy estimation under the same observation experiment. All minimax values below are exact for every sample size in the stated local parameterization.

B.1 Controlled Gaussian queries reduce to a scalar experiment

At round $t \in [n]$, a possibly randomized and adaptive sampling policy chooses an entry $(I_t, J_t) \in [2] \times [2]$ based on the past transcript and observes

$$Y_t = (A_\theta)_{I_t J_t} + \xi_t, \qquad \xi_t \overset{\text{iid}}{\sim} \mathcal{N}(0, 1). \tag{B.2}$$

A terminal decision rule may use the complete action–observation transcript and additional randomization. The next lemma shows that adaptively choosing entries does not create a hidden information advantage in this family.

Lemma B.1 (Exact reduction of the controlled query experiment)**.** *For every $n \in \mathbb{N}$, the class of controlled experiments generated by* (B.2) *is equivalent, for every statistical decision problem, to observing*

$$X_t = \theta + \varepsilon_t, \qquad \varepsilon_t \overset{\text{iid}}{\sim} \mathcal{N}(0, 1), \qquad t \in [n]. \tag{B.3}$$

More precisely: (i) repeatedly querying entry $(1, 2)$ produces (B.3)*; and (ii) for every adaptive query policy, its complete transcript can be simulated from $(X_1, \dots, X_n)$ by a parameter-independent Markov kernel. Consequently, for every nonnegative loss, the infimum risk over adaptive query policies and terminal decision rules equals the infimum risk over decision rules based on* (B.3)*.*

Proof. The first claim is immediate from $(A_\theta)_{12} = \theta$. For the reverse simulation, let $(X_t)_{t=1}^n$ follow (B.3), and let $(G_t)_{t=1}^n$ be independent standard Gaussians, independent of the X_t's and of the policy's internal randomization. Recursively simulate the policy's action using the previously simulated transcript, and then set

$$\widetilde{Y}_t \coloneqq \begin{cases} X_t, & (I_t, J_t) = (1, 2), \\ -X_t, & (I_t, J_t) = (2, 1), \\ G_t, & (I_t, J_t) \in \{(1, 1), (2, 2)\}. \end{cases} \tag{B.4}$$

Conditional on the simulated past and the chosen action, $\widetilde{Y}_t$ is Gaussian with variance one and mean $(A_\theta)_{I_t J_t}$. Independence of the fresh variables X_t and G_t therefore gives exactly the same conditional law as (B.2). Induction over t proves equality of the joint transcript laws for every θ. Since the simulation kernel does not depend on θ, composing any transcript-based decision rule with this kernel cannot improve upon the scalar experiment. The always-$(1, 2)$ policy gives the opposite reduction, proving equality of the optimal risks. □

B.2 EQUILIBRIUM GEOMETRY AND TARGET DISCONTINUITY

For $x = (p, 1-p) \in \Delta_2$ and $y = (r, 1-r) \in \Delta_2$, the row player's expected payoff is

$$x^\mathsf{T} A_\theta y = \theta(p - r). \tag{B.5}$$

For $u \in \Delta_2$, write $\mathsf{H}(u) \coloneqq -\sum_{i=1}^{2} u_i \log u_i$, with $0 \log 0 \coloneqq 0$. Recall that the saddle-point gap is

$$\mathrm{Gap}_A(x, y) \coloneqq \max_{x' \in \Delta_2} x'^\mathsf{T} A y - \min_{y' \in \Delta_2} x^\mathsf{T} A y'. \tag{B.6}$$

We use the product metric

$$d_\Delta\big((x, y), (x', y')\big) \coloneqq \frac{1}{2}\left(\|x - x'\|_1 + \|y - y'\|_1\right). \tag{B.7}$$

Lemma B.2 (Exact equilibrium geometry)**.** *For every $x = (p, 1-p)$ and $y = (r, 1-r)$,*

$$\mathrm{Gap}_{A_\theta}(x, y) = \begin{cases} \theta(2 - p - r), & \theta > 0, \\ 0, & \theta = 0, \\ |\theta|(p + r), & \theta < 0. \end{cases} \tag{B.8}$$

Moreover,

$$\mathrm{val}(A_\theta) = 0 \qquad \textit{for every } \theta, \tag{B.9}$$

and

$$\mathcal{E}(A_\theta) = \begin{cases} \{(e_1, e_1)\}, & \theta > 0, \\ \Delta_2 \times \Delta_2, & \theta = 0, \\ \{(e_2, e_2)\}, & \theta < 0. \end{cases} \tag{B.10}$$

If sel *denotes the maximum-entropy equilibrium selector, then*

$$\mathrm{sel}(A_\theta) = \begin{cases} (e_1, e_1), & \theta > 0, \\ \big((1/2, 1/2), (1/2, 1/2)\big), & \theta = 0, \\ (e_2, e_2), & \theta < 0. \end{cases} \tag{B.11}$$

In particular, the value is constant and globally regular on this family, whereas every exact point-equilibrium selector is forced to jump between the two unique side equilibria.

Proof. Equation (B.5) follows by direct multiplication. If $\theta > 0$, then

$$\max_{x' \in \Delta_2} x'^\mathsf{T} A_\theta y = \theta(1 - r),$$
$$\min_{y' \in \Delta_2} x^\mathsf{T} A_\theta y' = \theta(p - 1),$$

which gives the first case of (B.8). If $\theta < 0$, writing $\delta = |\theta|$ gives

$$\max_{x' \in \Delta_2} x'^\mathsf{T} A_\theta y = \delta r,$$
$$\min_{y' \in \Delta_2} x^\mathsf{T} A_\theta y' = -\delta p,$$

which gives the third case. At $\theta = 0$, every payoff and every gap is zero.

A strategy pair is a saddle point if and only if its gap is zero. For $\theta > 0$, (B.8) vanishes only when $p = r = 1$; for $\theta < 0$, it vanishes only when $p = r = 0$; and for $\theta = 0$, it vanishes everywhere. This proves (B.10). Every listed saddle point has payoff zero, yielding (B.9). Finally, when $\theta = 0$ the two equilibrium strategy sets are the full simplices, whose unique entropy maximizers are their uniform distributions; away from zero the equilibrium is unique. This proves (B.11). □

B.3 EXACT LOCAL MINIMAX RISKS

Fix $c > 0$ and an integer $n \geq \lceil c^2 \rceil$, and consider the two local parameters

$$\theta_{n,\eta} \coloneqq \eta \frac{c}{\sqrt{n}}, \qquad \eta \in \{-1, +1\}. \tag{B.12}$$

The restriction on n only ensures $\theta_{n,\eta} \in [-1, 1]$. The statistical decision problem is oracle-local in the following explicit sense: the family A_θ, the two-point class $\{-c/\sqrt{n}, +c/\sqrt{n}\}$, the constant c, and the unit observation variance are known; only the sign η is unknown. By lemma B.1, it suffices to use the scalar experiment (B.3). Its normalized sample mean

$$T_n \coloneqq \sqrt{n}\, \overline{X}_n \tag{B.13}$$

obeys

$$T_n \mid \eta \sim \mathcal{N}(\eta c, 1), \qquad \log \frac{dP_+}{dP_-}(T_n) = 2cT_n. \tag{B.14}$$

Thus T_n is sufficient for the two-point experiment. Let $\varphi(t) = (2\pi)^{-1/2} e^{-t^2/2}$ and $\Phi(t) = \int_{-\infty}^{t} \varphi(u)\, du$. The likelihood-ratio calculations below are the symmetric Gaussian specialization of the Neyman–Pearson principle (Neyman & Pearson, 1933); we include complete lower- and upper-bound arguments.

A procedure may use arbitrary randomization. Let $\widehat{v} \in \mathbb{R}$, $\widehat{b} \in \{-1, +1\}$, and $\widehat{z} = (\widehat{x}, \widehat{y}) \in \Delta_2 \times \Delta_2$ denote, respectively, a value estimate, a branch estimate, and a strategy-pair estimate. Define

$$\mathcal{R}^\star_{\mathrm{val},n}(c) \coloneqq \inf_{\widehat{v}} \max_{\eta \in \{-1,+1\}} \mathbb{E}_\eta \left[\left(\widehat{v} - \mathrm{val}(A_{\theta_{n,\eta}}) \right)^2 \right], \tag{B.15}$$

$$\mathcal{R}^\star_{\mathrm{br},n}(c) \coloneqq \inf_{\widehat{b}} \max_{\eta \in \{-1,+1\}} P_\eta(\widehat{b} \neq \eta), \tag{B.16}$$

$$\mathcal{R}^\star_{\mathrm{gap},n}(c) \coloneqq \inf_{\widehat{z}} \max_{\eta \in \{-1,+1\}} \sqrt{n}\, \mathbb{E}_\eta \left[\mathrm{Gap}_{A_{\theta_{n,\eta}}}(\widehat{z}) \right], \tag{B.17}$$

$$\mathcal{R}^\star_{\mathrm{sel},n}(c) \coloneqq \inf_{\widehat{z}} \max_{\eta \in \{-1,+1\}} \mathbb{E}_\eta \left[d_\Delta^2 \left(\widehat{z}, \mathrm{sel}(A_{\theta_{n,\eta}}) \right) \right]. \tag{B.18}$$

The infima range jointly over adaptive sampling policies and terminal decision rules; by lemma B.1, the same values are obtained by rules based only on T_n.

Proposition B.3 (Exact target-dependent minimax separation)**.** *For every $c > 0$ and $n \geq \lceil c^2 \rceil$,*

$$\mathcal{R}^\star_{\mathrm{val},n}(c) = 0, \tag{B.19}$$

$$\mathcal{R}^\star_{\mathrm{br},n}(c) = \Phi(-c), \tag{B.20}$$

$$\mathcal{R}^\star_{\mathrm{gap},n}(c) = 2c\Phi(-c), \tag{B.21}$$

$$\mathcal{R}^\star_{\mathrm{sel},n}(c) = \mathfrak{r}_{\mathrm{pt}}(c) \coloneqq e^{-c^2/2} \mathbb{E}_{G \sim \mathcal{N}(0,1)} \left[\frac{1}{\cosh(cG)} \right]. \tag{B.22}$$

The following rules attain the respective infima:

1. $\widehat{v} \equiv 0$*;*
2. $\widehat{b} = +1$ *if and only if* $T_n \geq 0$*, so the fixed zero-probability tie convention is* $\mathrm{sgn}(0) = +1$*;*
3. $\widehat{x} = \widehat{y} = e_1$ *when* $T_n \geq 0$*, and* $\widehat{x} = \widehat{y} = e_2$ *otherwise;*
4. $$\widehat{x} = \widehat{y} = \big(\pi_c(T_n), 1 - \pi_c(T_n)\big), \qquad \pi_c(t) \coloneqq \frac{1}{1 + e^{-2ct}}. \tag{B.23}$$

The posterior interpolation in item 4 is an estimator of the equilibrium point *under squared d_Δ loss. Except at the two pure endpoints, it need not itself be an equilibrium of the true game. In particular, for every fixed $c > 0$, value risk is zero and optimal expected exploitability is $\Theta(n^{-1/2})$, whereas branch and point-strategy risks remain bounded away from zero as $n \to \infty$.*

Proof. The value identity (B.19) follows from (B.9) by taking $\widehat{v} \equiv 0$.

For equilibrium-branch recovery, let $p_+(t) = \varphi(t-c)$ and $p_-(t) = \varphi(t+c)$ be the densities of T_n. A randomized branch rule can be represented by a measurable function $\delta : \mathbb{R} \to [0,1]$, where $\delta(t)$ is the conditional probability of declaring $+1$. Then

$$\begin{aligned}\max\{\mathbb{E}_+[1-\delta(T_n)], \mathbb{E}_-[\delta(T_n)]\} &\geq \frac{1}{2}\int_{\mathbb{R}} [p_+(t)(1-\delta(t)) + p_-(t)\delta(t)]\, dt \\ &\geq \frac{1}{2}\int_{\mathbb{R}} \min\{p_+(t), p_-(t)\}\, dt. \end{aligned} \tag{B.24}$$

Since $p_+(t) \geq p_-(t)$ if and only if $t \geq 0$,

$$\frac{1}{2}\int_{\mathbb{R}} \min\{p_+, p_-\} = \frac{1}{2}\left[P_+(T_n < 0) + P_-(T_n \geq 0)\right] = \Phi(-c). \tag{B.25}$$

The threshold rule $\delta(t) = \mathbf{1}\{t \geq 0\}$ has both error probabilities equal to $\Phi(-c)$, proving (B.20).

For the gap risk, write $\widehat{x} = (P, 1-P)$ and $\widehat{y} = (R, 1-R)$, and put $U = (P+R)/2 \in [0,1]$. By (B.8),

$$\sqrt{n}\ \mathrm{Gap}_{A_{\theta_{n,+}}}(\widehat{z}) = 2c(1-U), \qquad \sqrt{n}\ \mathrm{Gap}_{A_{\theta_{n,-}}}(\widehat{z}) = 2cU. \tag{B.26}$$

Conditioning on T_n and replacing U by its conditional mean does not change these two risks, so it suffices to consider a measurable $u : \mathbb{R} \to [0,1]$. Then

$$\begin{aligned}\max\left\{2c\int p_+(t)(1-u(t))\, dt, 2c\int p_-(t)u(t)\, dt\right\} &\geq c\int [p_+(t)(1-u(t)) + p_-(t)u(t)]\, dt \\ &\geq c\int \min\{p_+(t), p_-(t)\}\, dt = 2c\Phi(-c), \end{aligned} \tag{B.27}$$

where the last equality uses (B.25). The threshold rule $u(t) = \mathbf{1}\{t \geq 0\}$ attains equality and is implemented by the pure strategy rule in the proposition. This proves (B.21).

For point-strategy estimation, the side equilibria are $z_+ = (e_1, e_1)$ and $z_- = (e_2, e_2)$. With the same variable U,

$$d_\Delta^2(\widehat{z}, z_+) = 4(1-U)^2, \qquad d_\Delta^2(\widehat{z}, z_-) = 4U^2. \tag{B.28}$$

Because the losses are convex in U, conditional Jensen implies that randomization beyond T_n cannot help. For any deterministic $u : \mathbb{R} \to [0,1]$,

$$\begin{aligned}\max\left\{4\int p_+(t)(1-u(t))^2\, dt, 4\int p_-(t)u(t)^2\, dt\right\} &\geq 2\int \left[p_+(t)(1-u(t))^2 + p_-(t)u(t)^2\right] dt. \end{aligned} \tag{B.29}$$

Pointwise in t,

$$p_+(1-u)^2 + p_-u^2 = (p_+ + p_-)\left(u - \frac{p_+}{p_+ + p_-}\right)^2 + \frac{p_+p_-}{p_+ + p_-}. \tag{B.30}$$

Hence the equal-prior Bayes rule is

$$u(t) = \frac{p_+(t)}{p_+(t) + p_-(t)} = \pi_c(t). \tag{B.31}$$

The identities $p_+(-t) = p_-(t)$ and $\pi_c(-t) = 1 - \pi_c(t)$ imply that this rule has equal risks under the two parameters. It therefore attains the lower bound in (B.29), and is minimax. Its common risk is

$$\begin{aligned}2\int_{\mathbb{R}} \frac{p_+(t)p_-(t)}{p_+(t) + p_-(t)}\, dt &= e^{-c^2/2}\int_{\mathbb{R}} \frac{\varphi(t)}{\cosh(ct)}\, dt \\ &= e^{-c^2/2}\mathbb{E}_{G\sim\mathcal{N}(0,1)}\left[\frac{1}{\cosh(cG)}\right], \end{aligned} \tag{B.32}$$

which proves (B.22). □

Corollary B.4 (Four distinct statistical scales)**.** *For each fixed $c > 0$, along the local alternatives $\theta_{n,\pm} = \pm c/\sqrt{n}$,*

$$\mathcal{R}^{\star}_{\mathrm{val},n}(c) = 0, \qquad \inf_{\widehat{z}} \max_{\eta} \mathbb{E}_{\eta}\left[\mathrm{Gap}_{A_{\theta_{n,\eta}}}(\widehat{z})\right] = \frac{2c\Phi(-c)}{\sqrt{n}}, \tag{B.33}$$

whereas

$$\mathcal{R}^{\star}_{\mathrm{br},n}(c) = \Phi(-c) > 0, \qquad \mathcal{R}^{\star}_{\mathrm{sel},n}(c) = \mathfrak{r}_{\mathrm{pt}}(c) > 0. \tag{B.34}$$

Thus value estimation, approximate-equilibrium decision making, equilibrium-branch recovery, and point-strategy estimation occupy different statistical scales under the same sequence of experiments.

Remark B.5 (Endpoint checks). The formulas have the required limiting behavior. Since $\Phi(-c) \to 1/2$ and $e^{-c^2/2}\mathbb{E}[\cosh(cG)^{-1}] \to 1$ as $c \downarrow 0$, indistinguishable alternatives yield random-guess branch error and the minimax squared radius $d^2_{\Delta}(z_+, z_-)/4 = 1$ of the two point targets. As $c \to \infty$, $\Phi(-c) \to 0$ and dominated convergence gives $e^{-c^2/2}\mathbb{E}[\cosh(cG)^{-1}] \to 0$; the two structural targets become perfectly distinguishable. The unscaled optimal exploitability is exactly $2c\Phi(-c)/\sqrt{n}$.

B.4 A FINITE-SAMPLE LOWER BOUND FOR HONEST STRATEGY CONFIDENCE SETS

Let $\mathcal{Z}_2 \coloneqq \Delta_2 \times \Delta_2$, and let $\mathfrak{C}_n$ be a measurable random nonempty compact subset of $\mathcal{Z}_2$, where the hyperspace of nonempty compact subsets is equipped with the Hausdorff-metric Borel σ-field. Write

$$\mathrm{diam}_{d_{\Delta}}(\mathfrak{C}_n) \coloneqq \sup_{z,z' \in \mathfrak{C}_n} d_{\Delta}(z, z'), \tag{B.35}$$

with diameter zero for singletons. We use the convention $d_{\mathrm{TV}}(P, Q) \coloneqq \sup_B |P(B) - Q(B)|$.

Proposition B.6 (Honesty forces nonvanishing diameter)**.** *Suppose that, for some $\alpha \in [0, 1)$,*

$$\min_{\eta \in \{-1,+1\}} P_{\eta}\left(\mathrm{sel}(A_{\theta_{n,\eta}}) \in \mathfrak{C}_n\right) \geq 1 - \alpha. \tag{B.36}$$

Then

$$\max_{\eta \in \{-1,+1\}} \mathbb{E}_{\eta}\left[\mathrm{diam}_{d_{\Delta}}(\mathfrak{C}_n)\right] \geq 4\big[\Phi(-c) - \alpha\big]_+. \tag{B.37}$$

In particular, whenever $\alpha < \Phi(-c)$, no confidence set that is honest at both local alternatives can have expected diameter converging to zero.

Proof. Let $z_+ = (e_1, e_1)$ and $z_- = (e_2, e_2)$. Their distance is

$$d_{\Delta}(z_+, z_-) = 2. \tag{B.38}$$

The total variation distance between the two laws of T_n is

$$d_{\mathrm{TV}}(P_+, P_-) = 1 - 2\Phi(-c), \tag{B.39}$$

which follows because the two densities cross only at zero. Let $A_+ = \{z_+ \in \mathfrak{C}_n\}$ and $A_- = \{z_- \in \mathfrak{C}_n\}$. By honesty and the total variation inequality,

$$P_+(A_-) \geq P_-(A_-) - d_{\mathrm{TV}}(P_+, P_-) \geq 1 - \alpha - d_{\mathrm{TV}}(P_+, P_-). \tag{B.40}$$

Therefore

$$\begin{aligned} P_+(A_+ \cap A_-) &\geq P_+(A_+) + P_+(A_-) - 1 \\ &\geq 1 - 2\alpha - d_{\mathrm{TV}}(P_+, P_-) = 2\Phi(-c) - 2\alpha. \end{aligned} \tag{B.41}$$

On $A_+ \cap A_-$, the confidence-set diameter is at least two. Taking the positive part of the preceding lower bound gives

$$\mathbb{E}_+\left[\mathrm{diam}_{d_{\Delta}}(\mathfrak{C}_n)\right] \geq 2\big[2\Phi(-c) - 2\alpha\big]_+, \tag{B.42}$$

which is (B.37). The same argument can be run under P_-, although the maximum over the two parameters already follows from the displayed P_+ bound. □

Remark B.7 (Preliminary versus exact confidence bounds). The preceding argument is a direct testing-based lower bound. The exact least-favorable two-point minimax expected-diameter curve, together with an attaining measurable three-report region, is derived in theorem H.6.

Remark B.8 (Relation to noisy matrix-game learning). Existing work derives instance-dependent guarantees for approximate equilibria and support identification in noisy zero-sum matrix games (Maiti et al., 2023; Jiang & Zhang, 2026). The purpose of proposition B.3 is different: under a single local experiment, it gives exact minimax risks for four distinct equilibrium targets and shows that their statistical scales separate at a support bifurcation.

C Stable Values and the Selection Uncertainty Principle

This appendix establishes a global asymmetry that drives the statistical phenomena studied in the paper. The value and the decision loss of a zero-sum game are uniformly stable under payoff perturbations. By contrast, every single-valued strategy rule faces a quantitative stability–accuracy tradeoff. The dimension-uniform lower obstruction is the contribution. Entropy is analyzed only as a convenient smooth upper construction that matches the inverse-sensitivity dependence up to its logarithmic entropy-range factor; no new tight QRE approximation theorem is claimed.

Throughout this appendix, let

$$\mathcal{A}_{m,k} \coloneqq [-1,1]^{m\times k}, \qquad \mathcal{Z}_{m,k} \coloneqq \Delta_m \times \Delta_k,$$

where $m, k \geq 2$. For $z = (x, y)$ and $z' = (x', y')$ in $\mathcal{Z}_{m,k}$, recall the product total-variation metric

$$d_\Delta(z, z') \coloneqq \frac{1}{2}\big(\|x - x'\|_1 + \|y - y'\|_1\big). \tag{C.1}$$

All logarithms are natural.

C.1 Uniform stability of value and decision quality

We first record perturbation inequalities that do not require uniqueness, strict complementarity, or any local regularity of the equilibrium correspondence.

Proposition C.1 (Global perturbation bounds)**.** *For arbitrary $A, B \in \mathbb{R}^{m\times k}$, put $\delta \coloneqq \|A - B\|_\infty$. Then*

$$\big|\mathrm{val}(A) - \mathrm{val}(B)\big| \leq \delta, \tag{C.2}$$

$$\big|\mathrm{Gap}_A(z) - \mathrm{Gap}_B(z)\big| \leq 2\delta \qquad \textit{for every } z \in \mathcal{Z}_{m,k}. \tag{C.3}$$

Consequently, every $z_B \in \mathcal{E}(B)$ satisfies

$$\mathrm{Gap}_A(z_B) \leq 2\|A - B\|_\infty. \tag{C.4}$$

The constants in (C.2) *and* (C.3) *are sharp.*

Proof. For every $(x, y) \in \mathcal{Z}_{m,k}$,

$$\big|x^\mathsf{T}(A - B)y\big| \leq \sum_{i,j} x_i y_j |A_{ij} - B_{ij}| \leq \delta. \tag{C.5}$$

Taking first the minimum in y and then the maximum in x gives $\mathrm{val}(A) \leq \mathrm{val}(B) + \delta$; exchanging A and B proves (C.2).

For fixed $z = (x, y)$, define

$$\overline{v}_A(y) \coloneqq \max_{u\in\Delta_m} u^\mathsf{T} A y, \qquad \underline{v}_A(x) \coloneqq \min_{v\in\Delta_k} x^\mathsf{T} A v.$$

By (C.5), $|\overline{v}_A(y) - \overline{v}_B(y)| \leq \delta$ and $|\underline{v}_A(x) - \underline{v}_B(x)| \leq \delta$. Since $\mathrm{Gap}_A(x, y) = \overline{v}_A(y) - \underline{v}_A(x)$, the triangle inequality yields (C.3). Equation (C.4) follows from $\mathrm{Gap}_B(z_B) = 0$.

To see sharpness of (C.2), take two constant payoff matrices differing by δ. For (C.3), take $B = 0$, $A = \left(\begin{smallmatrix} 0 & \delta \\ -\delta & 0 \end{smallmatrix}\right)$, and $z = (e_2, e_2)$; then $\|A - B\|_\infty = \delta$ and $\mathrm{Gap}_A(z) - \mathrm{Gap}_B(z) = 2\delta$. □

C.2 No accurate single-valued strategy rule can be globally stable

A *single-valued strategy rule* is an arbitrary map $S : \mathcal{A}_{m,k} \to \mathcal{Z}_{m,k}$; it need not return an exact equilibrium. Define its global sensitivity and worst-case equilibrium error by

$$\mathrm{Lip}(S) \coloneqq \sup_{A\neq B} \frac{d_\Delta(S(A), S(B))}{\|A - B\|_\infty} \in [0, \infty], \tag{C.6}$$

$$\varepsilon_{\mathrm{eq}}(S) \coloneqq \sup_{A\in\mathcal{A}_{m,k}} \mathrm{Gap}_A(S(A)) \in [0, 2]. \tag{C.7}$$

The next construction embeds the two-branch singularity from section B into every matrix dimension.

For $t \in (0,1]$, define $B_t^+, B_t^- \in \mathcal{A}_{m,k}$ by the block matrices

$$B_t^{\pm} := \begin{pmatrix} A_{\pm t} & \mathbf{1}_2\mathbf{1}_{k-2}^{\mathsf{T}} \\ -\mathbf{1}_{m-2}\mathbf{1}_2^{\mathsf{T}} & 0_{(m-2)\times(k-2)} \end{pmatrix}, \tag{C.8}$$

where empty blocks are omitted when $m = 2$ or $k = 2$, and $A_{\pm t}$ is as in (B.1). Notice that

$$\|B_t^+ - B_t^-\|_\infty = 2t. \tag{C.9}$$

Lemma C.2 (Embedded two-branch gap inequalities)**.** *For every $z = (x,y) \in \mathcal{Z}_{m,k}$ and every $t \in (0,1]$,*

$$\mathrm{Gap}_{B_t^+}(x,y) \geq t(2 - x_1 - y_1), \tag{C.10}$$

$$\mathrm{Gap}_{B_t^-}(x,y) \geq t(x_1 + y_1). \tag{C.11}$$

Moreover, for every $t \in (0,1]$,

$$\mathcal{E}(B_t^+) = \{(e_1, e_1)\}, \qquad \mathcal{E}(B_t^-) = \{(e_2, e_2)\}. \tag{C.12}$$

Proof. For B_t^+, evaluate the row player's best-response term at e_1 and the column player's best-response term at e_1:

$$\max_{u \in \Delta_m} u^{\mathsf{T}} B_t^+ y \geq e_1^{\mathsf{T}} B_t^+ y = t y_2 + \sum_{j>2} y_j,$$

$$\min_{v \in \Delta_k} x^{\mathsf{T}} B_t^+ v \leq x^{\mathsf{T}} B_t^+ e_1 = -t x_2 - \sum_{i>2} x_i.$$

Subtracting and using $\sum_{i>2} x_i = 1 - x_1 - x_2$ and $\sum_{j>2} y_j = 1 - y_1 - y_2$ gives

$$\begin{aligned} \mathrm{Gap}_{B_t^+}(x,y) &\geq 2 - x_1 - y_1 - (1-t)(x_2 + y_2) \\ &= t(2 - x_1 - y_1) + (1-t)\big(2 - x_1 - x_2 - y_1 - y_2\big) \\ &\geq t(2 - x_1 - y_1), \end{aligned}$$

which proves (C.10). Similarly, evaluating both best-response terms for B_t^- at e_2 yields

$$\begin{aligned} \mathrm{Gap}_{B_t^-}(x,y) &\geq 2 - x_1 - x_2 - y_1 - y_2 + t(x_1 + y_1) \\ &\geq t(x_1 + y_1), \end{aligned}$$

proving (C.11).

For uniqueness, every row $i > 2$ is strictly dominated by row 1 in B_t^+ and by row 2 in B_t^-; every column $j > 2$ is strictly dominated, from the minimizing player's perspective, by column 1 in B_t^+ and by column 2 in B_t^-. Eliminating those actions reduces the game to A_t or A_{-t}, whose unique equilibria are given by Lemma B.2. This proves (C.12). □

Define the decreasing function $\psi : [0,\infty) \to (0,1]$ by

$$\psi(L) := \sup_{0<t\leq 1} t(1 - Lt) = \begin{cases} 1 - L, & 0 \leq L \leq \frac{1}{2}, \\ \dfrac{1}{4L}, & L > \frac{1}{2}. \end{cases} \tag{C.13}$$

Theorem C.3 (Global accuracy–sensitivity obstruction for single-valued strategy rules)**.** *For every $m, k \geq 2$ and every strategy rule $S : \mathcal{A}_{m,k} \to \mathcal{Z}_{m,k}$ with $L := \mathrm{Lip}(S) < \infty$,*

$$\boxed{\varepsilon_{\mathrm{eq}}(S) \geq \psi(L) \geq \frac{1}{4(1+L)}.} \tag{C.14}$$

In particular, if $\varepsilon_{\mathrm{eq}}(S) < \frac{1}{2}$, then

$$\mathrm{Lip}(S) \geq \frac{1}{4\varepsilon_{\mathrm{eq}}(S)}. \tag{C.15}$$

No continuous exact equilibrium selector exists on $\mathcal{A}_{m,k}$.

Proof. Fix $t \in (0, 1]$ and write $S(B_t^{\pm}) = (x^{\pm}, y^{\pm})$. By Lemma C.2,

$$\begin{aligned} 2\varepsilon_{\mathrm{eq}}(S) &\geq \mathrm{Gap}_{B_t^+}(S(B_t^+)) + \mathrm{Gap}_{B_t^-}(S(B_t^-)) \\ &\geq t\left[2 - \left(x_1^+ - x_1^-\right) - \left(y_1^+ - y_1^-\right)\right]. \end{aligned} \tag{C.16}$$

For probability vectors p, q, every coordinate satisfies $|p_i - q_i| \leq \frac{1}{2}\|p - q\|_1$. Hence

$$\begin{aligned} \left(x_1^+ - x_1^-\right) + \left(y_1^+ - y_1^-\right) &\leq d_{\Delta}\big(S(B_t^+), S(B_t^-)\big) \\ &\leq L\|B_t^+ - B_t^-\|_{\infty} = 2Lt. \end{aligned}$$

Substitution into (C.16) gives

$$\varepsilon_{\mathrm{eq}}(S) \geq t(1 - Lt) \qquad \text{for every } t \in (0, 1]. \tag{C.17}$$

Optimizing the concave quadratic over $[0, 1]$ proves the first inequality in (C.14). The second follows directly from the two cases in (C.13). If $\varepsilon_{\mathrm{eq}}(S) < 1/2$, the first case is impossible, so $L > 1/2$ and $\varepsilon_{\mathrm{eq}}(S) \geq 1/(4L)$, proving (C.15).

Finally, suppose that S were a continuous exact selector. By (C.12), for every $t \in (0, 1)$, $S(B_t^+) = (e_1, e_1)$ and $S(B_t^-) = (e_2, e_2)$. Both matrix paths converge to the same B_0 as $t \downarrow 0$, contradicting continuity at B_0. □

Remark C.4 (What the lower bound does and does not say). The obstruction is not computational: it applies to every single-valued map, including maps defined by unbounded computation or exact knowledge of A. It is also not an artifact of high dimension, since it is generated by a one-dimensional transverse path. The theorem distinguishes stable *decision quality* from stable *strategy selection*: by Proposition C.1, an equilibrium of a nearby game has small exploitability, even though every accurate single-valued point-strategy rule must vary rapidly or discontinuously.

C.3 Entropy as an upper construction for the selector frontier

For finite $m \times k$ games, Nguyen et al. (2026) establish the tight $\Theta(\beta^{-1}(\log m + \log k))$ QRE approximation order. The analysis below has a different and narrower purpose: it exhibits one globally smooth strategy rule whose bias and sensitivity witness the upper side of the general selector frontier proved above. The QRE approximation rate itself is inherited, not claimed as a contribution.

Let

$$h_d(p) \coloneqq \sum_{i=1}^{d} p_i \log p_i, \qquad 0 \log 0 \coloneqq 0, \tag{C.18}$$

so that $-h_d$ is Shannon entropy. For $\varepsilon_{\mathrm{reg}} > 0$, define the entropy-regularized saddle function

$$\mathcal{L}_{A,\varepsilon_{\mathrm{reg}}}(x, y) \coloneqq x^{\mathsf{T}} A y - \varepsilon_{\mathrm{reg}} h_m(x) + \varepsilon_{\mathrm{reg}} h_k(y), \tag{C.19}$$

and let

$$\mathcal{S}^{\mathrm{sol}}_{\varepsilon_{\mathrm{reg}}}(A) = (x_{\varepsilon_{\mathrm{reg}}}(A), y_{\varepsilon_{\mathrm{reg}}}(A)) \tag{C.20}$$

be its saddle point. Existence is a direct instance of Sion's minimax theorem (Sion, 1958); the logit form below is the zero-sum entropic specialization of quantal-response regularization (McKelvey & Palfrey, 1995).

Proposition C.5 (Well-posedness and logit equations)**.** *For every $A \in \mathbb{R}^{m \times k}$ and $\varepsilon_{\mathrm{reg}} > 0$, $\mathcal{L}_{A,\varepsilon_{\mathrm{reg}}}$ has a unique saddle point $\mathcal{S}^{\mathrm{sol}}_{\varepsilon_{\mathrm{reg}}}(A)$, and both strategy vectors have full support. They satisfy*

$$x_{\varepsilon_{\mathrm{reg}},i}(A) = \frac{\exp((Ay_{\varepsilon_{\mathrm{reg}}}(A))_i/\varepsilon_{\mathrm{reg}})}{\sum_{r=1}^{m} \exp((Ay_{\varepsilon_{\mathrm{reg}}}(A))_r/\varepsilon_{\mathrm{reg}})}, \tag{C.21}$$

$$y_{\varepsilon_{\mathrm{reg}},j}(A) = \frac{\exp(-(A^{\mathsf{T}}x_{\varepsilon_{\mathrm{reg}}}(A))_j/\varepsilon_{\mathrm{reg}})}{\sum_{s=1}^{k} \exp(-(A^{\mathsf{T}}x_{\varepsilon_{\mathrm{reg}}}(A))_s/\varepsilon_{\mathrm{reg}})}. \tag{C.22}$$

Proof. The function in (C.19) is continuous on the compact convex set $\mathcal{Z}_{m,k}$, strictly concave in x, and strictly convex in y. Sion's theorem gives a saddle point. If (x_1, y_1) and (x_2, y_2) were two saddle points, the saddle inequalities would give

$$\mathcal{L}_{A,\varepsilon_{\mathrm{reg}}}(x_2, y_1) \leq \mathcal{L}_{A,\varepsilon_{\mathrm{reg}}}(x_1, y_1) \leq \mathcal{L}_{A,\varepsilon_{\mathrm{reg}}}(x_1, y_2) \leq \mathcal{L}_{A,\varepsilon_{\mathrm{reg}}}(x_2, y_2) \leq \mathcal{L}_{A,\varepsilon_{\mathrm{reg}}}(x_2, y_1).$$

All inequalities are therefore equalities. Strict concavity in x forces $x_1 = x_2$, and strict convexity in y forces $y_1 = y_2$.

For fixed y, put $u = Ay$ and $\pi_i = \exp(u_i/\varepsilon_{\mathrm{reg}})/\sum_r \exp(u_r/\varepsilon_{\mathrm{reg}})$. A direct calculation gives the Gibbs variational identity

$$x^{\mathsf{T}}u - \varepsilon_{\mathrm{reg}} h_m(x) = \varepsilon_{\mathrm{reg}} \log\left(\sum_{r=1}^{m} e^{u_r/\varepsilon_{\mathrm{reg}}}\right) - \varepsilon_{\mathrm{reg}} D_{\mathrm{KL}}(x\|\pi).$$

Hence the unique maximizer is $x = \pi$, which is strictly positive, proving (C.21). Similarly, for fixed x, put $v = A^{\mathsf{T}}x$ and $\pi^y_j = \exp(-v_j/\varepsilon_{\mathrm{reg}})/\sum_s \exp(-v_s/\varepsilon_{\mathrm{reg}})$. Then

$$y^{\mathsf{T}}v + \varepsilon_{\mathrm{reg}} h_k(y) = -\varepsilon_{\mathrm{reg}} \log\left(\sum_{s=1}^{k} e^{-v_s/\varepsilon_{\mathrm{reg}}}\right) + \varepsilon_{\mathrm{reg}} D_{\mathrm{KL}}(y\|\pi^y),$$

so the unique minimizer is $y = \pi^y$, proving (C.22). Evaluating these best-response identities at the saddle point completes the proof. □

Let

$$\mathrm{val}_{\varepsilon_{\mathrm{reg}}}(A) \coloneqq \max_{x\in\Delta_m} \min_{y\in\Delta_k} \mathcal{L}_{A,\varepsilon_{\mathrm{reg}}}(x,y). \tag{C.23}$$

Proposition C.6 (Regularization bias)**.** *For every $A \in \mathbb{R}^{m\times k}$ and $\varepsilon_{\mathrm{reg}} > 0$,*

$$\mathrm{val}(A) - \varepsilon_{\mathrm{reg}} \log k \le \mathrm{val}_{\varepsilon_{\mathrm{reg}}}(A) \le \mathrm{val}(A) + \varepsilon_{\mathrm{reg}} \log m, \tag{C.24}$$

and

$$\mathrm{Gap}_A(\mathcal{S}^{\mathrm{sol}}_{\varepsilon_{\mathrm{reg}}}(A)) \le \varepsilon_{\mathrm{reg}}(\log m + \log k). \tag{C.25}$$

Proof. Since $0 \le -h_m(x) \le \log m$ and $0 \le -h_k(y) \le \log k$,

$$x^{\mathsf{T}}Ay - \varepsilon_{\mathrm{reg}} \log k \le \mathcal{L}_{A,\varepsilon_{\mathrm{reg}}}(x,y) \le x^{\mathsf{T}}Ay + \varepsilon_{\mathrm{reg}} \log m.$$

Taking max–min proves (C.24).

Write $(x_{\varepsilon_{\mathrm{reg}}}, y_{\varepsilon_{\mathrm{reg}}}) = \mathcal{S}^{\mathrm{sol}}_{\varepsilon_{\mathrm{reg}}}(A)$. The row saddle inequality implies, for every $x \in \Delta_m$,

$$x^{\mathsf{T}}Ay_{\varepsilon_{\mathrm{reg}}} \le x^{\mathsf{T}}_{\varepsilon_{\mathrm{reg}}}Ay_{\varepsilon_{\mathrm{reg}}} + \varepsilon_{\mathrm{reg}}\big[-h_m(x_{\varepsilon_{\mathrm{reg}}}) + h_m(x)\big] \le x^{\mathsf{T}}_{\varepsilon_{\mathrm{reg}}}Ay_{\varepsilon_{\mathrm{reg}}} + \varepsilon_{\mathrm{reg}} \log m.$$

The column saddle inequality similarly gives, for every $y \in \Delta_k$,

$$x^{\mathsf{T}}_{\varepsilon_{\mathrm{reg}}}Ay \ge x^{\mathsf{T}}_{\varepsilon_{\mathrm{reg}}}Ay_{\varepsilon_{\mathrm{reg}}} - \varepsilon_{\mathrm{reg}} \log k.$$

Taking the maximum in the first display and the minimum in the second, then subtracting, proves (C.25). □

We next prove a global stability bound with an explicit constant. For strictly positive probability vectors, define

$$D_{\mathrm{KL}}(p\|q) \coloneqq \sum_i p_i \log \frac{p_i}{q_i}.$$

Lemma C.7 (Strong monotonicity of negative entropy)**.** *For $p, q \in \mathrm{ri}(\Delta_d)$,*

$$\langle \nabla h_d(p) - \nabla h_d(q), p - q\rangle = D_{\mathrm{KL}}(p\|q) + D_{\mathrm{KL}}(q\|p) \ge \|p - q\|_1^2. \tag{C.26}$$

Proof. The equality follows by expanding $\nabla h_d(p) = \mathbf{1} + \log p$ coordinatewise. Pinsker's inequality in nats states that $D_{\mathrm{KL}}(p\|q) \ge \frac{1}{2}\|p - q\|_1^2$ (Csiszár, 1967). Applying it in both directions proves the inequality. □

Theorem C.8 (Entropic upper construction for the stability frontier)**.** *For every $A, B \in \mathbb{R}^{m\times k}$ and every $\varepsilon_{\mathrm{reg}} > 0$,*

$$\boxed{d_\Delta(\mathcal{S}^{\mathrm{sol}}_{\varepsilon_{\mathrm{reg}}}(A), \mathcal{S}^{\mathrm{sol}}_{\varepsilon_{\mathrm{reg}}}(B)) \le \frac{\|A-B\|_\infty}{\varepsilon_{\mathrm{reg}}}.} \tag{C.27}$$

Consequently,

$$\mathrm{Lip}(\mathcal{S}^{\mathrm{sol}}_{\varepsilon_{\mathrm{reg}}}) \le \frac{1}{\varepsilon_{\mathrm{reg}}}, \qquad \varepsilon_{\mathrm{eq}}(\mathcal{S}^{\mathrm{sol}}_{\varepsilon_{\mathrm{reg}}}) \le \varepsilon_{\mathrm{reg}}(\log m + \log k). \tag{C.28}$$

More generally, for every estimated matrix $\widehat{A}$,

$$d_\Delta(\mathcal{S}^{\mathrm{sol}}_{\varepsilon_{\mathrm{reg}}}(\widehat{A}), \mathcal{S}^{\mathrm{sol}}_{\varepsilon_{\mathrm{reg}}}(A)) \le \frac{\|\widehat{A}-A\|_\infty}{\varepsilon_{\mathrm{reg}}}, \tag{C.29}$$

$$\mathrm{Gap}_A(\mathcal{S}^{\mathrm{sol}}_{\varepsilon_{\mathrm{reg}}}(\widehat{A})) \le \varepsilon_{\mathrm{reg}}(\log m + \log k) + 2\|\widehat{A}-A\|_\infty. \tag{C.30}$$

Proof. Let $z_A = (x_A, y_A) = \mathcal{S}^{\mathrm{sol}}_{\varepsilon_{\mathrm{reg}}}(A)$ and $z_B = (x_B, y_B) = \mathcal{S}^{\mathrm{sol}}_{\varepsilon_{\mathrm{reg}}}(B)$. By Proposition C.5, all four strategy vectors are strictly positive. Define the regularized saddle operator

$$\mathcal{G}_{A,\varepsilon_{\mathrm{reg}}}(x,y) \coloneqq \begin{pmatrix} -Ay + \varepsilon_{\mathrm{reg}}\nabla h_m(x) \\ A^{\mathsf{T}}x + \varepsilon_{\mathrm{reg}}\nabla h_k(y) \end{pmatrix}. \tag{C.31}$$

The saddle-point first-order conditions are the variational inequalities

$$\big\langle \mathcal{G}_{A,\varepsilon_{\mathrm{reg}}}(z_A), z - z_A\big\rangle \ge 0, \qquad \big\langle \mathcal{G}_{B,\varepsilon_{\mathrm{reg}}}(z_B), z - z_B\big\rangle \ge 0 \quad \text{for every } z \in \mathcal{Z}_{m,k}. \tag{C.32}$$

Insert $z = z_B$ in the first inequality and $z = z_A$ in the second, then add. With $\Delta x = x_A - x_B$ and $\Delta y = y_A - y_B$, this yields

$$\begin{aligned}&\big\langle \mathcal{G}_{A,\varepsilon_{\mathrm{reg}}}(z_A) - \mathcal{G}_{A,\varepsilon_{\mathrm{reg}}}(z_B), z_A - z_B\big\rangle \\ &\qquad \le \big\langle \mathcal{G}_{B,\varepsilon_{\mathrm{reg}}}(z_B) - \mathcal{G}_{A,\varepsilon_{\mathrm{reg}}}(z_B), z_A - z_B\big\rangle.\end{aligned} \tag{C.33}$$

The two bilinear terms on the left cancel: $-\Delta x^{\mathsf{T}} A \Delta y + \Delta y^{\mathsf{T}} A^{\mathsf{T}} \Delta x = 0$. By Lemma C.7, the left-hand side is therefore at least

$$\varepsilon_{\mathrm{reg}}\big(\|\Delta x\|_1^2 + \|\Delta y\|_1^2\big). \tag{C.34}$$

Writing $\delta = \|A-B\|_\infty$, the right-hand side of (C.33) equals

$$\langle (A-B)y_B, \Delta x\rangle + \big\langle (B-A)^{\mathsf{T}}x_B, \Delta y\big\rangle \le \delta\big(\|\Delta x\|_1 + \|\Delta y\|_1\big), \tag{C.35}$$

where the inequality uses $\|(A-B)y_B\|_\infty \le \delta$ and $\|(B-A)^{\mathsf{T}}x_B\|_\infty \le \delta$. Set $a = \|\Delta x\|_1$, $b = \|\Delta y\|_1$, and $d = d_\Delta(z_A, z_B) = (a+b)/2$. Combining (C.34) and (C.35), then using $a^2 + b^2 \ge (a+b)^2/2 = 2d^2$, gives

$$2\varepsilon_{\mathrm{reg}} d^2 \le \varepsilon_{\mathrm{reg}}(a^2+b^2) \le \delta(a+b) = 2\delta d.$$

If $d = 0$ there is nothing to prove; otherwise division by $2d$ yields (C.27). The first two claims in (C.28) follow from (C.27) and Proposition C.6. Equation (C.29) is the same stability estimate with $B = \widehat{A}$. Finally,

$$\begin{aligned}\mathrm{Gap}_A(\mathcal{S}^{\mathrm{sol}}_{\varepsilon_{\mathrm{reg}}}(\widehat{A})) &\le \mathrm{Gap}_{\widehat{A}}(\mathcal{S}^{\mathrm{sol}}_{\varepsilon_{\mathrm{reg}}}(\widehat{A})) + 2\|\widehat{A}-A\|_\infty \\ &\le \varepsilon_{\mathrm{reg}}(\log m + \log k) + 2\|\widehat{A}-A\|_\infty\end{aligned}$$

by Propositions C.1 and C.6, proving (C.30). □

Corollary C.9 (Global accuracy–stability sandwich)**.** *For $L_0 > 0$, let*

$$\mathfrak{S}(L_0) \coloneqq \{S : \mathcal{A}_{m,k} \to \mathcal{Z}_{m,k} : \mathrm{Lip}(S) \le L_0\}.$$

Then

$$\psi(L_0) \le \inf_{S\in\mathfrak{S}(L_0)} \varepsilon_{\mathrm{eq}}(S) \le \min\left\{2, \frac{\log m + \log k}{L_0}\right\}. \tag{C.36}$$

In particular, for $L_0 > 1/2$,

$$\frac{1}{4L_0} \le \inf_{S\in\mathfrak{S}(L_0)} \varepsilon_{\mathrm{eq}}(S) \le \frac{\log(mk)}{L_0}. \tag{C.37}$$

Thus the optimal worst-case equilibrium error scales as L_0^{-1} in the high-accuracy regime, up to the entropy-range factor $\log(mk)$.

Proof. The lower bound follows from theorem C.3 and the fact that ψ is decreasing. For the second upper bound, choose $\varepsilon_{\mathrm{reg}} = L_0^{-1}$ in theorem C.8. The upper bound by 2 follows from any constant strategy rule, because payoffs lie in $[-1, 1]$ and hence every saddle-point gap is at most 2. The final display uses the second case of (C.13) and $\log m + \log k = \log(mk)$. □

Remark C.10 (Exact role of entropy). Theorem C.3 applies to every globally Lipschitz single-valued strategy rule and is dimension uniform. The entropic rule in theorem C.8 is one witness that the L^{-1} dependence can be achieved, with the factor $\log(mk)$ coming from the entropy range. The tight finite-game QRE approximation rate is due to Nguyen et al. (2026); neither proposition C.6 nor corollary C.9 is presented as a new QRE-rate result.

The entropic strategy rule is not merely an arbitrary smooth approximation. At every fixed game, it selects precisely the maximum-entropy equilibrium in the vanishing-regularization limit.

Proposition C.11 (Vanishing regularization selects maximum entropy)**.** *For every fixed $A \in \mathbb{R}^{m\times k}$,*

$$\lim_{\varepsilon_{\mathrm{reg}}\downarrow 0} \mathcal{S}^{\mathrm{sol}}_{\varepsilon_{\mathrm{reg}}}(A) = \mathrm{sel}(A) \qquad \textit{in } d_\Delta. \tag{C.38}$$

Proof. Write $(x_{\varepsilon_{\mathrm{reg}}}, y_{\varepsilon_{\mathrm{reg}}}) = \mathcal{S}^{\mathrm{sol}}_{\varepsilon_{\mathrm{reg}}}(A)$. By (C.25), $\mathrm{Gap}_A(x_{\varepsilon_{\mathrm{reg}}}, y_{\varepsilon_{\mathrm{reg}}}) \to 0$. Compactness of $\mathcal{Z}_{m,k}$ and continuity of the gap therefore imply that every accumulation point of $(x_{\varepsilon_{\mathrm{reg}}}, y_{\varepsilon_{\mathrm{reg}}})$ lies in $\mathcal{E}(A) = \mathcal{X}^\star(A) \times \mathcal{Y}^\star(A)$. The product identity follows from zero-sum duality: if x is maximin-optimal and y is minimax-optimal, then $x^\mathsf{T} A y \geq \mathrm{val}(A)$ and $x^\mathsf{T} A y \leq \mathrm{val}(A)$, so (x, y) is a saddle point.

Let $\bar{x} \in \mathcal{X}^\star(A)$ and $\bar{y} \in \mathcal{Y}^\star(A)$. Since $\bar{x}$ guarantees the value against every column strategy and $\bar{y}$ caps the value against every row strategy,

$$\bar{x}^\mathsf{T} A y_{\varepsilon_{\mathrm{reg}}} \geq \mathrm{val}(A), \qquad x^\mathsf{T}_{\varepsilon_{\mathrm{reg}}} A \bar{y} \leq \mathrm{val}(A).$$

The saddle inequalities for $\mathcal{L}_{A,\varepsilon_{\mathrm{reg}}}$ give

$$\mathcal{L}_{A,\varepsilon_{\mathrm{reg}}}(\bar{x}, y_{\varepsilon_{\mathrm{reg}}}) \leq \mathcal{L}_{A,\varepsilon_{\mathrm{reg}}}(x_{\varepsilon_{\mathrm{reg}}}, y_{\varepsilon_{\mathrm{reg}}}) \leq \mathcal{L}_{A,\varepsilon_{\mathrm{reg}}}(x_{\varepsilon_{\mathrm{reg}}}, \bar{y}).$$

Writing $\mathsf{Ent} = -h$ and combining the last two displays yields

$$\mathsf{Ent}(x_{\varepsilon_{\mathrm{reg}}}) + \mathsf{Ent}(y_{\varepsilon_{\mathrm{reg}}}) \geq \mathsf{Ent}(\bar{x}) + \mathsf{Ent}(\bar{y}) \qquad \text{for every } (\bar{x}, \bar{y}) \in \mathcal{E}(A). \tag{C.39}$$

Take a convergent subsequence with limit $(x_0, y_0) \in \mathcal{E}(A)$ and choose $\bar{x}$ and $\bar{y}$ to be the unique entropy maximizers over $\mathcal{X}^\star(A)$ and $\mathcal{Y}^\star(A)$, respectively. By continuity of entropy, (C.39) implies

$$\mathsf{Ent}(x_0) + \mathsf{Ent}(y_0) \geq \max_{x\in\mathcal{X}^\star(A)} \mathsf{Ent}(x) + \max_{y\in\mathcal{Y}^\star(A)} \mathsf{Ent}(y).$$

The reverse inequality holds because $(x_0, y_0) \in \mathcal{X}^\star(A) \times \mathcal{Y}^\star(A)$. Equality follows, and each component must attain its corresponding entropy maximum. Strict concavity of entropy on each convex equilibrium-strategy set makes those maximizers unique, so every accumulation point equals $\mathrm{sel}(A)$. Compactness then implies convergence of the entire family. □

Remark C.12 (Interpretation). Theorems C.3 and C.8 give a quantitative form of the paper's central distinction. Entropic regularization produces a globally Lipschitz single-valued strategy rule and converges pointwise to the canonical maximum-entropy equilibrium selector, but its Lipschitz modulus must diverge at least inversely with the desired worst-case equilibrium accuracy. Hence regularization cannot eliminate an equilibrium bifurcation; it can only replace a discontinuous exact selector by a controlled bias–sensitivity tradeoff.

D Primal–Dual Equilibrium Branch Geometry

This appendix develops the finite-dimensional branch geometry underlying the local statistical theory. The starting point is the classical representation of extreme optimal strategies by square Shapley–Snow kernels (Shapley & Snow, 1950; Attia & Oliu-Barton, 2021). We recast that representation as a family of bordered linear systems, derive exact differential formulae for every regular kernel branch, and organize the first-order validity inequalities into a finite polyhedral fan. The resulting object is local and directional: it describes which equilibrium branch is selected by a small

payoff perturbation inside a prescribed linear perturbation space. Singular kernels and higher-order contacts are not hidden; they are excluded by explicit assumptions and treated as distinct failure modes.

Throughout, $A \in \mathbb{R}^{m\times k}$ is the row player's payoff matrix, $m, k \geq 2$, and

$$\mathcal{Z} \coloneqq \Delta_m \times \Delta_k.$$

For an index set I, $\mathbf{1}_I$ denotes the all-ones vector in $\mathbb{R}^{|I|}$, and I^c denotes the complement relative to the ambient action set. Submatrices and subvectors are indexed in increasing order. We write $\mathrm{ext}(C)$ for the extreme points of a compact convex set C.

D.1 Linear-programming and basic-solution representations

The row and column optimal-strategy sets are the solution sets of the dual linear programs

$$\mathrm{val}(A) = \max_{x,v} \left\{ v : A^{\mathsf{T}}x \geq v\mathbf{1}_k,\ \mathbf{1}_m^{\mathsf{T}}x = 1,\ x \geq 0 \right\}, \tag{D.1}$$

$$\mathrm{val}(A) = \min_{y,w} \left\{ w : Ay \leq w\mathbf{1}_m,\ \mathbf{1}_k^{\mathsf{T}}y = 1,\ y \geq 0 \right\}. \tag{D.2}$$

Consequently,

$$(x,y) \in \mathcal{E}(A) \quad \Longleftrightarrow \quad Ay \leq \mathrm{val}(A)\mathbf{1}_m, \qquad A^{\mathsf{T}}x \geq \mathrm{val}(A)\mathbf{1}_k. \tag{D.3}$$

Moreover, $\mathcal{E}(A) = \mathcal{X}^\star(A) \times \mathcal{Y}^\star(A)$, and therefore

$$\mathrm{ext}(\mathcal{E}(A)) = \mathrm{ext}(\mathcal{X}^\star(A)) \times \mathrm{ext}(\mathcal{Y}^\star(A)). \tag{D.4}$$

We call the elements of (D.4) *basic equilibria*. This terminology agrees with the basic solutions of Shapley and Snow.

Let

$$\mathfrak{K}_{m,k} \coloneqq \{\beta = (I,J) : \varnothing \neq I \subseteq [m],\ \varnothing \neq J \subseteq [k],\ |I| = |J|\}$$

be the finite collection of square action pairs. For $\beta = (I,J) \in \mathfrak{K}_{m,k}$, put $r = |I| = |J|$ and define the bordered matrices

$$M_\beta^y(A) \coloneqq \begin{pmatrix} A_{I,J} & -\mathbf{1}_I \\ \mathbf{1}_J^{\mathsf{T}} & 0 \end{pmatrix}, \qquad M_\beta^x(A) \coloneqq \begin{pmatrix} A_{I,J}^{\mathsf{T}} & -\mathbf{1}_J \\ \mathbf{1}_I^{\mathsf{T}} & 0 \end{pmatrix}. \tag{D.5}$$

The two determinants coincide. Indeed, if $D_r = \mathrm{diag}(\mathbf{I}_r, -1)$, then

$$M_\beta^x(A) = D_r M_\beta^y(A)^{\mathsf{T}} D_r.$$

We denote their common determinant by

$$\Delta_\beta(A) \coloneqq \det M_\beta^y(A) = \det M_\beta^x(A). \tag{D.6}$$

Expanding the bordered determinant shows that $\Delta_\beta(A)$ is the sum of the cofactors of $A_{I,J}$; hence $\Delta_\beta(A) \neq 0$ is precisely the regularity condition appearing in the Shapley–Snow kernel formula (Shapley & Snow, 1950; Attia & Oliu-Barton, 2021).

Definition D.1 (Regular kernel branch)**.** Fix $\beta = (I,J) \in \mathfrak{K}_{m,k}$. On the open set

$$\mathcal{U}_\beta \coloneqq \{A \in \mathbb{R}^{m\times k} : \Delta_\beta(A) \neq 0\},$$

define

$$\begin{pmatrix} y_{\beta,J}(A) \\ v_\beta^y(A) \end{pmatrix} \coloneqq M_\beta^y(A)^{-1} \begin{pmatrix} 0_r \\ 1 \end{pmatrix}, \qquad \begin{pmatrix} x_{\beta,I}(A) \\ v_\beta^x(A) \end{pmatrix} \coloneqq M_\beta^x(A)^{-1} \begin{pmatrix} 0_r \\ 1 \end{pmatrix}. \tag{D.7}$$

The remaining coordinates of $x_\beta(A)$ and $y_\beta(A)$ are set to zero. We call $z_\beta(A) = (x_\beta(A), y_\beta(A))$ the regular kernel branch associated with β.

The two scalar components in (D.7) are equal. To see this, write $B = A_{I,J}$, $x = x_{\beta,I}(A)$, and $y = y_{\beta,J}(A)$. The bordered systems give

$$By = v_\beta^y \mathbf{1}_I, \qquad B^{\mathsf{T}}x = v_\beta^x \mathbf{1}_J, \qquad \mathbf{1}_I^{\mathsf{T}}x = \mathbf{1}_J^{\mathsf{T}}y = 1.$$

Thus $x^{\mathsf{T}}By = v_\beta^y = v_\beta^x$. We henceforth denote the common value by $v_\beta(A)$.

D.2 EXACT BRANCH-VALIDITY CERTIFICATES

For $A \in \mathcal{U}_\beta$, define the row and column reduced slacks

$$s^{\mathrm{row}}_{\beta,i}(A) \coloneqq v_\beta(A) - A_{i,J}y_{\beta,J}(A), \qquad i \in I^c, \tag{D.8}$$

$$s^{\mathrm{col}}_{\beta,j}(A) \coloneqq x_{\beta,I}(A)^{\mathsf{T}}A_{I,j} - v_\beta(A), \qquad j \in J^c. \tag{D.9}$$

Collect all basic masses and reduced slacks into the analytic validity map

$$\rho_\beta(A) \coloneqq \begin{pmatrix} x_{\beta,I}(A) \\ y_{\beta,J}(A) \\ \left(s^{\mathrm{row}}_{\beta,i}(A)\right)_{i\in I^c} \\ \left(s^{\mathrm{col}}_{\beta,j}(A)\right)_{j\in J^c} \end{pmatrix} \in \mathbb{R}^{m+k}. \tag{D.10}$$

The coordinate ordering in (D.10) is fixed once and for all.

Proposition D.2 (Exact equilibrium certificate)**.** *Let $\beta \in \mathfrak{K}_{m,k}$ and $A \in \mathcal{U}_\beta$. Then*

$$z_\beta(A) \in \mathcal{E}(A) \quad \Longleftrightarrow \quad \rho_\beta(A) \geq 0 \quad \textit{coordinatewise}. \tag{D.11}$$

Whenever (D.11) *holds, $v_\beta(A) = \mathrm{val}(A)$ and $z_\beta(A)$ is a basic equilibrium. If $\rho_\beta(A) > 0$ coordinatewise, then $z_\beta(A)$ is the unique equilibrium of A and*

$$\mathrm{supp}(x_\beta(A)) = I, \qquad \mathrm{supp}(y_\beta(A)) = J.$$

Proof. Suppose first that $\rho_\beta(A) \geq 0$. The bordered systems imply that x_β and y_β are nonnegative and normalized. The equalities on $I \times J$ and the reduced-slack inequalities give

$$Ay_\beta \leq v_\beta \mathbf{1}_m, \qquad A^{\mathsf{T}}x_\beta \geq v_\beta \mathbf{1}_k.$$

Hence, for every $(u, w) \in \mathcal{Z}$,

$$u^{\mathsf{T}}Ay_\beta \leq v_\beta = x_\beta^{\mathsf{T}}Ay_\beta \leq x_\beta^{\mathsf{T}}Aw.$$

Thus $z_\beta(A)$ is a saddle point and $v_\beta(A) = \mathrm{val}(A)$. Conversely, if the branch pair is an equilibrium, then its strategy coordinates are nonnegative and (D.3) gives all reduced-slack inequalities, so $\rho_\beta(A) \geq 0$.

The Shapley–Snow basic-solution theorem implies that a valid square kernel with $\Delta_\beta(A) \neq 0$ represents a basic equilibrium (Shapley & Snow, 1950; Attia & Oliu-Barton, 2021). It remains to prove the uniqueness claim. Assume $\rho_\beta(A) > 0$, and let $(\widetilde{x}, \widetilde{y})$ be any equilibrium. Since Ay_β is strictly below v_β on I^c and $\widetilde{x}^{\mathsf{T}}Ay_\beta = v_\beta$, one has $\widetilde{x}_{I^c} = 0$. Since $y_{\beta,J} > 0$ and $A^{\mathsf{T}}\widetilde{x} \geq v_\beta \mathbf{1}_k$, equality of the weighted average forces $A^{\mathsf{T}}_{I,J}\widetilde{x}_I = v_\beta \mathbf{1}_J$. Together with $\mathbf{1}_I^{\mathsf{T}}\widetilde{x}_I = 1$, the nonsingularity of $M^x_\beta(A)$ gives $\widetilde{x} = x_\beta$. The symmetric argument using the strict column slacks gives $\widetilde{y} = y_\beta$. □

The converse part of Shapley–Snow theory gives a finite completeness statement.

Proposition D.3 (Paired Shapley–Snow coverage)**.** *For every $A \in \mathbb{R}^{m\times k}$ and every*

$$(x, y) \in \mathrm{ext}(\mathcal{X}^\star(A)) \times \mathrm{ext}(\mathcal{Y}^\star(A)),$$

there exists one square action pair $\beta = (I, J)$ such that $\Delta_\beta(A) \neq 0$, $\rho_\beta(A) \geq 0$, and $(x, y) = z_\beta(A)$. This includes rectangular games, zero game value, zero coordinates in the represented strategies, degenerate equilibrium polytopes, singular payoff submatrices with nonsingular bordered systems, and duplicate kernel representations. Consequently,

$$\mathcal{X}^\star(A) = \mathrm{conv}\left\{x_\beta(A) : A \in \mathcal{U}_\beta,\ \rho_\beta(A) \geq 0\right\}, \tag{D.12}$$

$$\mathcal{Y}^\star(A) = \mathrm{conv}\left\{y_\beta(A) : A \in \mathcal{U}_\beta,\ \rho_\beta(A) \geq 0\right\}. \tag{D.13}$$

Proof. For arbitrary convex sets X, Y,

$$\mathrm{ext}(X \times Y) = \mathrm{ext}(X) \times \mathrm{ext}(Y).$$

Indeed, a nonextreme coordinate makes the pair nonextreme, while an extreme- coordinate pair cannot be decomposed nontrivially without decomposing at least one coordinate. Since $\mathcal{E}(A) = \mathcal{X}^\star(A) \times \mathcal{Y}^\star(A)$, the pair (x, y) is therefore a basic solution in the terminology of Shapley–Snow.

Theorem 2.3.1 of Attia & Oliu-Barton (2021), restating the paired Shapley–Snow theorem, yields one square subgame $B = A_{I,J}$ such that

$$c \coloneqq S(\operatorname{co} B) \neq 0, \qquad x_I = \frac{\operatorname{co}(B)\mathbf{1}}{c}, \qquad y_J = \frac{\operatorname{co}(B)^\mathsf{T}\mathbf{1}}{c}, \tag{D.14}$$

with the remaining coordinates padded by zeros, and this completed pair is optimal. For

$$M^y_\beta(A) = \begin{pmatrix} B & -\mathbf{1} \\ \mathbf{1}^\mathsf{T} & 0 \end{pmatrix},$$

Laplace expansion gives

$$\det M^y_\beta(A) = S(\operatorname{co} B) = c \neq 0.$$

The same scalar is, up to the fixed sign convention, the determinant of the transposed bordered system $M^x_\beta(A)$. The vectors in (D.14) solve

$$By = v\mathbf{1}, \qquad B^\mathsf{T} x_I = v\mathbf{1}, \qquad \mathbf{1}^\mathsf{T} x_I = \mathbf{1}^\mathsf{T} y_J = 1, \qquad v = \frac{\det B}{c}.$$

By uniqueness of the bordered-system solutions they equal the branch coordinates from definition D.1. Optimality of the padded pair gives all mass and reduced-slack inequalities, hence $\rho_\beta(A) \geq 0$. This proves the paired common-kernel assertion. The hull identities follow because the two projected equilibrium sets are compact polytopes and therefore the convex hulls of their extreme points. □

D.3 Analyticity and exact directional derivatives

Every regular branch is a rational, hence real analytic, function of the payoff entries on $\mathcal{U}_\beta$. Its derivative has an explicit bordered linear-system representation.

Proposition D.4 (Kernel-branch differential). *Fix $A \in \mathcal{U}_\beta$ and a perturbation $H \in \mathbb{R}^{m\times k}$. Write*

$$x = x_{\beta,I}(A), \qquad y = y_{\beta,J}(A), \qquad v = v_\beta(A).$$

Then the directional derivative $(\dot x, \dot y, \dot v) = D(x_{\beta,I}, y_{\beta,J}, v_\beta)(A)[H]$ is the unique solution of

$$\begin{pmatrix} \dot y \\ \dot v \end{pmatrix} = -M^y_\beta(A)^{-1} \begin{pmatrix} H_{I,J} y \\ 0 \end{pmatrix}, \qquad \begin{pmatrix} \dot x \\ \dot v \end{pmatrix} = -M^x_\beta(A)^{-1} \begin{pmatrix} H^\mathsf{T}_{I,J} x \\ 0 \end{pmatrix}, \tag{D.15}$$

where the scalar derivatives computed by the two systems agree and satisfy

$$\dot v = x^\mathsf{T} H_{I,J} y. \tag{D.16}$$

The derivative of the validity map is

$$D\rho_\beta(A)[H] = \begin{pmatrix} \dot x \\ \dot y \\ \dot v \mathbf{1}_{I^c} - H_{I^c,J} y - A_{I^c,J}\dot y \\ H^\mathsf{T}_{I,J^c} x + A^\mathsf{T}_{I,J^c}\dot x - \dot v \mathbf{1}_{J^c} \end{pmatrix}. \tag{D.17}$$

Proof. Differentiate the two bordered systems in (D.7). Since their right-hand sides are constant, the product rule gives (D.15). Differentiating $A_{I,J} y = v\mathbf{1}_I$ and left-multiplying by x^T yields

$$x^\mathsf{T} H_{I,J} y + x^\mathsf{T} A_{I,J}\dot y = \dot v.$$

Because $x^\mathsf{T} A_{I,J} = v\mathbf{1}^\mathsf{T}_J$ and $\mathbf{1}^\mathsf{T}_J \dot y = 0$, the middle term vanishes, proving (D.16). Differentiating the coordinates of (D.10) gives (D.17). □

In particular, for every compact $K \subset \mathbb{R}^{m\times k}$ and every $A_0 \in \mathcal{U}_\beta$, there exist a neighborhood of A_0 and a finite constant $C_{\beta,K}$ such that

$$\|z_\beta(A_0 + tH) - z_\beta(A_0) - tDz_\beta(A_0)[H]\|_1 \leq C_{\beta,K} t^2 \qquad (H \in K) \tag{D.18}$$

for all sufficiently small $|t|$. The same uniform second-order expansion holds for v_β and ρ_β.

D.4 A LOCAL REGULAR-KERNEL ATLAS

The bordered representation is finite, but a kernel whose bordered determinant vanishes at the reference game can become nonsingular after perturbation. Such *singular-kernel creation* lies outside the regular analytic atlas based at A_0; its limiting behavior is not captured by the present first-order branch construction. We exclude it explicitly rather than silently identifying it with an ordinary basis change.

Fix a reference game A_0 and a linear perturbation space

$$\mathbb{T} \subseteq \mathbb{R}^{m\times k}. \tag{D.19}$$

The full payoff model corresponds to $\mathbb{T} = \mathbb{R}^{m\times k}$; a structured or one-dimensional statistical submodel is represented by a proper subspace. Define the finite collection of regular incident kernels

$$\mathfrak{B}_0 := \{\beta \in \mathfrak{K}_{m,k} : \Delta_\beta(A_0) \neq 0,\ \rho_\beta(A_0) \geq 0\}. \tag{D.20}$$

Assumption D.5 (Regular-kernel completeness)**.** There exists a relative neighborhood $\mathcal{N}$ of A_0 in $A_0 + \mathbb{T}$ such that every basic equilibrium of every $A \in \mathcal{N}$ is represented by a kernel $\beta \in \mathfrak{K}_{m,k}$ satisfying $\Delta_\beta(A_0) \neq 0$.

Lemma D.6 (Local regular-kernel atlas)**.** *Under assumption D.5, after shrinking $\mathcal{N}$ if necessary,*

$$\operatorname{ext}(\mathcal{E}(A)) = \{z_\beta(A) : \beta \in \mathfrak{B}_0,\ \rho_\beta(A) \geq 0\} \qquad (A \in \mathcal{N}), \tag{D.21}$$

where duplicate representations are removed.

Proof. By assumption D.5 and proposition D.3, every basic equilibrium in the original neighborhood is represented by one of the finitely many kernels satisfying $\Delta_\beta(A_0) \neq 0$. If such a kernel does not belong to $\mathfrak{B}_0$, then at least one coordinate of $\rho_\beta(A_0)$ is strictly negative. Continuity keeps that coordinate negative on a relative neighborhood of A_0, so this kernel cannot be valid there. Intersecting these finitely many neighborhoods proves that only indices in $\mathfrak{B}_0$ can occur. The reverse inclusion follows from proposition D.2. □

Thus assumption D.5 is exactly the condition under which the finite family $\{z_\beta : \beta \in \mathfrak{B}_0\}$ is a complete local atlas of basic equilibrium branches. It rules out branch creation through a kernel singular at A_0 but permits degeneracy of the validity inequalities and multiple equilibria at A_0.

For $\beta \in \mathfrak{B}_0$, let

$$\mathcal{A}_\beta := \{a \in [m+k] : \rho_{\beta,a}(A_0) = 0\} \tag{D.22}$$

be the active validity coordinates. For $a \in \mathcal{A}_\beta$, define the linear functional on $\mathbb{T}$

$$\lambda_{\beta,a}(H) := D\rho_{\beta,a}(A_0)[H]. \tag{D.23}$$

Assumption D.7 (First-order transversality)**.** For every $\beta \in \mathfrak{B}_0$ and every $a \in \mathcal{A}_\beta$, the restriction $\lambda_{\beta,a} : \mathbb{T} \to \mathbb{R}$ is nonzero.

This assumption excludes a validity constraint whose first nonzero variation is quadratic or higher along the chosen statistical model. It is weaker than requiring a unique equilibrium at A_0 and allows arbitrarily many incident regular branches.

D.5 A REGULAR-KERNEL VALIDITY REFINEMENT

Let Λ_0 be the finite collection of distinct nonzero linear forms $\lambda_{\beta,a}$ from (D.23), identifying only exact duplicates. The central hyperplane arrangement

$$\mathcal{H}_0 := \{\ker \lambda : \lambda \in \Lambda_0\} \tag{D.24}$$

induces sign cells in $\mathbb{T}$: two perturbations lie in the same cell when every $\lambda \in \Lambda_0$ has the same sign on both.

Lemma D.8 (Active-arrangement fan)**.** *The closures of the nonempty sign cells of* (D.24)*, together with all of their faces, form a complete polyhedral fan in $\mathbb{T}$. On the relative interior of each cone, every $\lambda \in \Lambda_0$ has a constant sign in $\{-1, 0, +1\}$.*

Proof. A sign cell is the intersection of finitely many open homogeneous half-spaces and hyperplanes. Its closure is therefore a polyhedral cone. Intersecting two such closures amounts to imposing the common sign restrictions and turning incompatible strict signs into equalities; hence the intersection is a face of each cone. Every $H \in \mathbb{T}$ has a sign vector, so the union of the cones is $\mathbb{T}$. These are precisely the axioms of a complete polyhedral fan. □

Denote this finite central arrangement by

$$\mathfrak{F}(A_0; \mathbb{T}). \tag{D.25}$$

We call (D.25) the *regular-kernel validity refinement*. Its full-dimensional relatively open cells are called *chambers*. The refinement is sufficient for sectorwise branch identification, but is not asserted to be the coarsest or a canonical branch partition: distinct kernel representations of the same analytic branch germ may contribute redundant hyperplanes. Quotienting representations by analytic-germ equality removes duplicate branch labels, but retaining every active representation-level hyperplane is always a safe refinement.

For a chamber C and any $H \in C$, define

$$\mathfrak{B}(C) \coloneqq \{\beta \in \mathfrak{B}_0 : \lambda_{\beta,a}(H) > 0 \text{ for every } a \in \mathcal{A}_\beta\}. \tag{D.26}$$

The definition is independent of the chosen $H \in C$ because all active linear forms have constant nonzero sign on a chamber.

Theorem D.9 (Sectorwise regular-kernel validity refinement away from its skeleton)**.** *Suppose assumptions D.5 and D.7 hold. Let C be a chamber of $\mathfrak{F}(A_0; \mathbb{T})$, and let*

$$\mathbb{S}_\infty(\mathbb{T}) \coloneqq \{H \in \mathbb{T} : \|H\|_\infty = 1\}.$$

For every compact $K \subset C \cap \mathbb{S}_\infty(\mathbb{T})$, there exists $\varepsilon_K > 0$ such that, for every $H \in K$ and every $t \in (0, \varepsilon_K)$,

$$\operatorname{ext}(\mathcal{E}(A_0 + tH)) = \{z_\beta(A_0 + tH) : \beta \in \mathfrak{B}(C)\}, \tag{D.27}$$

with duplicate representations removed. Every branch appearing in (D.27) *satisfies*

$$z_\beta(A_0 + tH) = z_\beta(A_0) + tDz_\beta(A_0)[H] + O_K(t^2), \tag{D.28}$$

uniformly over $H \in K$. Here the payoff path is measured in $\|\cdot\|_\infty$, the strategy remainder is measured in d_Δ (equivalently, in any fixed finite-dimensional norm), and every validity coordinate has a scalar $O_K(t^2)$ remainder uniformly on K.

Moreover, the right-hand side of (D.27) *contains exactly one distinct equilibrium pair. Hence $A_0 + tH$ has a unique equilibrium for all $H \in K$ and sufficiently small $t > 0$; its branch index and first-order derivative are constant throughout the chamber, up to duplicate kernel representations of the same analytic branch.*

Proof. Fix a chamber C and compact $K \subset C \cap \mathbb{S}_\infty(\mathbb{T})$. For every $\beta \in \mathfrak{B}_0$, the validity coordinates inactive at A_0 are strictly positive and remain positive in a common neighborhood by continuity. For an active coordinate $a \in \mathcal{A}_\beta$, the uniform Taylor expansion gives

$$\rho_{\beta,a}(A_0 + tH) = t\lambda_{\beta,a}(H) + O_K(t^2). \tag{D.29}$$

Since K is compact and no active linear form vanishes on the chamber, the absolute value of every $\lambda_{\beta,a}$ is uniformly bounded away from zero on K. Therefore, after decreasing ε_K, the sign of each active validity coordinate in (D.29) equals the sign of its linearization, simultaneously for all incident branches, all active coordinates, and all $H \in K$.

It follows from proposition D.2 that $z_\beta(A_0 + tH)$ is a basic equilibrium precisely when $\beta \in \mathfrak{B}(C)$. The local coverage (D.21) then proves (D.27). The expansion (D.28) follows from (D.18).

Finally, every branch in $\mathfrak{B}(C)$ has all validity coordinates strictly positive for small $t > 0$. By proposition D.2, each such branch would be the unique equilibrium. Hence two distinct branch pairs cannot both occur. Existence of at least one branch follows from the minimax theorem and assumption D.5. Therefore exactly one distinct equilibrium pair occurs on the chamber. □

Corollary D.10 (Sectorwise analytic selector)**.** *Under the hypotheses of theorem D.9, let C be a chamber and let $K \subset C \cap \mathbb{S}_\infty(\mathbb{T})$ be compact. Then there is $\varepsilon_K > 0$ such that the equilibrium map is single-valued on*

$$\{A_0 + tH : H \in K,\ 0 < t < \varepsilon_K\}$$

and is the restriction of one real analytic kernel branch there. Uniformly for $H \in K$, its first-order law is $H \mapsto Dz_{\beta(C)}(A_0)[H]$.

The theorem is deliberately silent on lower-dimensional faces of the fan. If H lies on such a face, at least one active first-order variation vanishes and quadratic or higher-order terms may determine feasibility. This is the precise boundary between the first-order tangent-fan theory and higher-order singularity analysis.

D.6 Regular, kink, and jump interfaces

Let C_+ and C_- be adjacent chambers sharing a codimension-one face, and let z_+ and z_- denote their unique analytic equilibrium branches. Their one-sided limits at A_0 exist by analyticity of the incident kernel branches. We use the following classification.

Definition D.11 (Branch interface type)**.** An adjacent two-chamber interface at A_0 is called

1. *regular* if the two branch germs agree;
2. a *continuous kink* if $z_+(A_0) = z_-(A_0)$ but $Dz_+(A_0) \neq Dz_-(A_0)$ on $\mathbb{T}$;
3. a *jump bifurcation* if $z_+(A_0) \neq z_-(A_0)$.

At every interface, the branch values agree at the reference game:

$$v_+(A_0) = v_-(A_0) = \mathrm{val}(A_0), \tag{D.30}$$

because both limiting branch pairs are equilibria of the same zero-sum game. Thus a jump in equilibrium strategies never forces a jump in the game value. The value may nevertheless have a kink because

$$Dv_+(A_0)[H] = x_+(A_0)^\mathsf{T} H y_+(A_0), \qquad Dv_-(A_0)[H] = x_-(A_0)^\mathsf{T} H y_-(A_0)$$

can differ.

A *simple regular-kernel bifurcation* is an interface for which the local parameter space is divided into exactly two chambers by a single hyperplane $\{H : \lambda(H) = 0\}$, with one distinct equilibrium branch germ on each side, possibly represented by multiple equivalent regular kernels. In later statistical results, the signed normal coordinate is represented by a smooth local defining function g satisfying $g(A_0) = 0$ and $Dg(A_0)|_{\mathbb{T}} = \lambda$. Existence of such a function is an explicit assumption on the interface, not a consequence of first-order transversality alone.

Example D.12 (A codimension-one regular-kernel jump)**.** Consider the one-parameter affine model

$$A(t) = \begin{pmatrix} t & 1 \\ -t & 1 \end{pmatrix}, \qquad |t| < 1. \tag{D.31}$$

For $t > 0$, the unique equilibrium is (e_1, e_1); for $t < 0$, it is (e_2, e_1); and at $t = 0$,

$$\mathcal{E}(A(0)) = \Delta_2 \times \{e_1\}.$$

Both incident branches are 1×1 regular kernels. Their active row slacks are $2t$ and $-2t$, respectively, while all column slacks remain strictly positive. Hence the regular-kernel validity refinement on the one-dimensional perturbation space has the two chambers $t > 0$ and $t < 0$. The equilibrium strategy jumps, whereas

$$\mathrm{val}(A(t)) = |t| \tag{D.32}$$

is continuous and Lipschitz, with a scalar kink at the interface. The maximum-entropy selector at $t = 0$ is $((1/2, 1/2), e_1)$, so the canonical point selector is discontinuous from both sides.

D.7 Analytic branch germs, finite verification, and worked failures

Definition D.13 (Equality of regular analytic branch germs)**.** For regular kernels β_1, β_2, write $\beta_1 \sim \beta_2$ if their rational maps agree on a relative neighborhood of A_0 on which both bordered determinants are nonzero:

$$z_{\beta_1}(A) = z_{\beta_2}(A).$$

Equality only at A_0 is not sufficient. The validity refinement may retain several representations of one germ; doing so can over-refine, but cannot merge, the intrinsic branch regions.

D.7.1 Exact finite verification of the branch assumptions

For each square kernel β, define its valid graph

$$\mathcal{G}_\beta = \{(A, z) : \Delta_\beta(A) \neq 0,\ \rho_\beta(A) \geq 0,\ z = z_\beta(A)\}. \tag{D.33}$$

Let

$$\mathfrak{K}_{\mathrm{reg}}(A_0) = \{\beta : \Delta_\beta(A_0) \neq 0\}, \qquad \mathcal{G}_{\mathrm{reg}} = \bigcup_{\beta \in \mathfrak{K}_{\mathrm{reg}}(A_0)} \mathcal{G}_\beta, \qquad \mathcal{G}_{\mathrm{all}} = \bigcup_{\beta} \mathcal{G}_\beta.$$

Theorem D.14 (Semialgebraic characterization of regular-kernel completeness)**.** *The sets $\mathcal{G}_\beta$, $\mathcal{G}_{\mathrm{reg}}$, and $\mathcal{G}_{\mathrm{all}}$ are semialgebraic. Regular-kernel completeness at $(A_0, \mathbb{T})$ is equivalent to the first-order real-algebraic statement*

$$\begin{aligned} \exists \varepsilon > 0\ \forall A, z: \quad & A - A_0 \in \mathbb{T}, \quad \|A - A_0\|_F^2 < \varepsilon^2, \quad (A, z) \in \mathcal{G}_{\mathrm{all}} \\ & \Longrightarrow (A, z) \in \mathcal{G}_{\mathrm{reg}}. \end{aligned} \tag{D.34}$$

Consequently, the assumption is a decidable first-order property whenever the entries of A_0 and a basis of $\mathbb{T}$ are supplied in an effective real closed field (in particular, for rational or real-algebraic input), by real quantifier elimination. For unrestricted real constants the display remains an exact semialgebraic characterization, but algorithmic decidability requires an effective representation of those constants.

A simpler sufficient certificate is

$$A_0 \notin \mathrm{cl}_{A_0 + \mathbb{T}}\{A : \Delta_\beta(A) \neq 0,\ \rho_\beta(A) \geq 0\} \quad \textit{for every } \beta \textit{ with } \Delta_\beta(A_0) = 0. \tag{D.35}$$

This sufficient condition need not be necessary, because a singular-at-A_0 kernel may duplicate an equilibrium germ already represented by a regular kernel.

Proof. By Cramer's rule, every coordinate of $z_\beta(A)$ and every validity coordinate is a rational function with polynomial numerator and denominator $\Delta_\beta(A)$. On each sign region $\Delta_\beta > 0$ and $\Delta_\beta < 0$, clearing denominators converts the branch equations and validity inequalities into finitely many polynomial equalities and inequalities. Thus $\mathcal{G}_\beta$ is semialgebraic, and finite unions preserve semialgebraicity.

By proposition D.3, the fiber of $\mathcal{G}_{\mathrm{all}}$ over A is exactly the set of basic equilibria, allowing duplicate representations. Therefore (D.34) says precisely that every basic equilibrium of every sufficiently nearby game has at least one representation by a kernel regular at A_0; this is regular-kernel completeness. The formula is first-order over the real field. Quantifier elimination decides it over an effective real closed coefficient field (Basu et al., 2006); without an effective representation of the input constants, it should be read as an exact logical characterization rather than a Turing-algorithm claim.

If (D.35) holds for every singular-at-A_0 kernel, finiteness of the kernel set gives one neighborhood in which none of those kernels is valid. Every nearby basic equilibrium is represented by some kernel, hence by one regular at A_0. This proves sufficiency. It is not necessary when a valid singular kernel represents a branch already represented by a regular kernel, which explains why the exact criterion is the graph inclusion (D.34) rather than separate exclusion of every singular representation. □

Proposition D.15 (Finite transversality and fan-construction certificate)**.** *Let*

$$N_{m,k} = \sum_{r=1}^{\min(m,k)} \binom{m}{r}\binom{k}{r}. \tag{D.36}$$

This is the number of square action pairs. Given A_0 and a basis of $\mathbb{T}$, the following finite procedure either certifies the regular-kernel fan assumptions or returns an explicit failed condition:

1. *enumerate all $N_{m,k}$ kernels;*
2. *compute $\Delta_\beta(A_0)$ and solve every nonsingular bordered system;*
3. *retain the incident valid kernels and compute their active validity coordinates;*
4. *compute the restrictions $D\rho_{\beta,j}(A_0)|_{\mathbb{T}}$ and reject if any is zero;*
5. *decide* (D.34)*, or use* (D.35) *as a sufficient certificate;*
6. *decide equality of regular analytic branch germs, quotient only the* branch labels *by that equivalence, and construct a safe hyperplane refinement using all representation-level active forms.*

Equality of two regular analytic germs on the relative model is decidable by clearing denominators and testing whether the resulting coordinate polynomials vanish identically on $\mathbb{T}$. If two regular kernels define the same germ, then their validity predicates agree on their common neighborhood, because the exact branch certificate says that either predicate is equivalent to the same intrinsic statement $z(A) \in \mathcal{E}(A)$. Thus one may retain a single representative after this equivalence has been verified. Keeping every representation-level active hyperplane instead gives a possibly redundant but always safe chamber refinement.

Proof. Every step except the semialgebraic decision is finite linear algebra or polynomial identity testing. Step 4 is exact because transversality is simply nonvanishing of finitely many linear forms on the supplied basis of $\mathbb{T}$. Step 5 is decidable by theorem D.14. For Step 6, two rational branch maps agree on a relative neighborhood iff, after cross-multiplying their nonzero denominators, every coordinate-difference polynomial vanishes identically on the affine model; polynomial identity testing on a finite-dimensional linear parameterization decides this. By the exact branch certificate, equal germs have equivalent validity predicates on the common regular neighborhood. The active nonzero linear forms then define a finite central hyperplane arrangement; retaining redundant forms only over-refines the same branch geometry and does not invalidate the sectorwise theorem. □

The finite enumeration and bordered solves cost

$$O\left(\sum_{r=1}^{\min(m,k)} \binom{m}{r}\binom{k}{r} r^3\right)$$

arithmetic operations up to validity and derivative evaluation. An arrangement of N active hyperplanes in tangent dimension d has at most $\sum_{j=0}^{d}\binom{N}{j}$ full-dimensional chambers. Exact singular-kernel exclusion is a general semialgebraic decision problem and can be substantially more expensive; the point of theorem D.14 is an exact semialgebraic characterization and, for effectively represented algebraic inputs, a finite decision procedure; no polynomial-time guarantee is claimed.

Property	What is established
Kernel enumeration	finite and explicit
Bordered systems	finite linear algebra
Active-gradient transversality	finite exact nonvanishing test
Duplicate-germ equality	polynomial identity test on the relative model
Regular-kernel completeness	first-order semialgebraic decision problem
Completeness complexity	potentially very high; no polynomial guarantee

D.8 From a validity interface to a simple jump

Proposition D.16 (A regular-kernel interface induces a simple jump chart)**.** *Let $A_0 + \mathbb{T}$ be the relative perturbation model. Suppose two incident regular branches z_+ and z_- are analytic near A_0, and there is a real analytic validity coordinate r such that*

$$r(A_0) = 0, \qquad Dr(A_0)|_{\mathbb{T}} \neq 0.$$

Assume that in some relative neighborhood $\mathcal{N}$ of A_0:

1. *$z_+(A)$ is the unique equilibrium for $r(A) > 0$;*
2. *$z_-(A)$ is the unique equilibrium for $r(A) < 0$;*
3. *no third branch is valid in $\mathcal{N}$.*

Then $\{A \in \mathcal{N} : r(A) = 0\}$ is a real analytic codimension-one relative submanifold, $g = r$ is a local defining function for the two sides, and if $z_+(A_0) \neq z_-(A_0)$ then the simple transverse jump assumption holds.

Proof. The derivative of r on the model is nonzero, so the relative analytic implicit-function theorem makes $r^{-1}(0)$ a codimension-one analytic submanifold after shrinking $\mathcal{N}$. The two connected local sides are $r > 0$ and $r < 0$. By assumption, the equilibrium selector equals the analytic branch z_+ on the positive side and z_- on the negative side. Analyticity of the regular kernel branches gives the one-sided limits $z_+(A_0)$ and $z_-(A_0)$. If these differ, the selector has a transverse jump across the smooth boundary, exactly as required by the simple-jump assumption. □

Lemma D.17 (Normalization invariance of the jump experiment)**.** *Replace a defining function g by $\widetilde{g} = cg$ with $c \neq 0$. If $a = Dg(A_0)|_{\mathbb{T}}$, $\widetilde{a} = D\widetilde{g}(A_0)|_{\mathbb{T}}$, and the signed normal margin is rescaled as $\widetilde{\lambda} = |c|\lambda$, then*

$$\widetilde{a} = ca, \qquad \widetilde{s_a} = |c|s_a, \qquad \widetilde{\kappa} = \frac{\widetilde{\lambda}}{\widetilde{s_a}} = \frac{\lambda}{s_a} = \kappa.$$

The local-oracle branch allocation $q_i^\star \propto \sigma_i|a_i|$ is unchanged.

Proof. Linearity of differentiation gives $\widetilde{a} = ca$. For every information matrix J,

$$\widetilde{s_a}^2 = \widetilde{a}^\mathsf{T} J^\dagger \widetilde{a} = c^2 a^\mathsf{T} J^\dagger a = c^2 s_a^2.$$

The displayed invariance of κ follows. Multiplication of every design weight numerator $\sigma_i|a_i|$ by the same factor $|c|$ cancels under normalization. □

D.9 A THREE-WAY WORKED VALIDITY REFINEMENT

For $(u, v) \in \mathbb{R}^2$, define

$$A(u, v) := \begin{pmatrix} u & 1 & 1 \\ v & 2 & 1 \\ 0 & 1 & 2 \end{pmatrix}, \qquad A_0 := A(0, 0), \tag{D.37}$$

and restrict attention to the relative neighborhood $\max\{|u|, |v|\} < 1/2$. Let

$$H_1 = e_1 e_1^\mathsf{T}, \qquad H_2 = e_2 e_1^\mathsf{T}, \qquad \mathbb{T} = \operatorname{span}\{H_1, H_2\}.$$

For $i \in \{1, 2, 3\}$, let $\beta_i = (\{i\}, \{1\})$.

Proposition D.18 (Exact three-way validity refinement)**.** *For every $\max\{|u|, |v|\} < 1/2$:*

1. *the unique column equilibrium is e_1;*
2. *the row equilibrium set is*

$$\operatorname{conv}\{e_i : i \in \arg\max\{u, v, 0\}\}; \tag{D.38}$$

3. *the value is*

$$\operatorname{val}(A(u, v)) = \max\{u, v, 0\}; \tag{D.39}$$

4. *every extreme equilibrium pair is represented by one of the three regular kernels $\beta_1, \beta_2, \beta_3$;*
5. *the active branch-validity forms at A_0 are*

$$\beta_1 : \quad u - v,\ u, \tag{D.40}$$
$$\beta_2 : \quad v - u,\ v, \tag{D.41}$$
$$\beta_3 : \quad -u,\ -v. \tag{D.42}$$

Consequently, regular-kernel completeness and first-order transversality hold relative to $\mathbb{T}$. The hyperplanes $u = 0$, $v = 0$, and $u = v$ form a finite sufficient validity refinement. The coarsest branch regions are

$$\mathcal{C}_1 = \{u > \max\{v, 0\}\}, \tag{D.43}$$
$$\mathcal{C}_2 = \{v > \max\{u, 0\}\}, \tag{D.44}$$
$$\mathcal{C}_3 = \{0 > \max\{u, v\}\}, \tag{D.45}$$

with selected equilibria (e_i, e_1) and zero branch derivatives. The hyperplane arrangement strictly over-refines this three-region partition.

Proof. Fix $x \in \Delta_3$. The expected payoffs of the three columns are

$$x^\mathsf{T} A(u, v)e_1 = ux_1 + vx_2,$$
$$x^\mathsf{T} A(u, v)e_2 = x_1 + 2x_2 + x_3 = 1 + x_2,$$
$$x^\mathsf{T} A(u, v)e_3 = x_1 + x_2 + 2x_3 = 1 + x_3.$$

The first quantity belongs to $(-1/2, 1/2)$, whereas the latter two are at least one. Hence e_1 is the unique minimizing column against every row strategy. The row player therefore maximizes

$$x^\mathsf{T} A(u, v)e_1 = ux_1 + vx_2,$$

whose maximizers are precisely (D.38); this proves the first three claims. The extreme row equilibria are the pure maximizers, while the column equilibrium is uniquely e_1, proving the kernel-coverage claim.

Each 1×1 bordered kernel matrix has determinant one, so all three branches are regular at A_0. For branch (e_1, e_1), row optimality against column one requires $u \geq v$ and $u \geq 0$, while the two column slacks are $1 - u$ and $1 - u$, strictly positive in the stated neighborhood. This gives (D.40). The other two displays follow in the same way. Every active differential is nonzero on $\mathbb{T}$, proving transversality. Since every nearby extreme equilibrium is one of these three pure pairs, regular-kernel completeness also holds. The branch regions and zero derivatives follow directly from the fact that the equilibrium strategies are constant on each strict argmax region. Finally, for example, $\mathcal{C}_1$ intersects both $v > 0$ and $v < 0$, so the arrangement hyperplane $v = 0$ splits one branch region and is therefore a strict refinement. □

D.10 Failure modes excluded by the first-order theory

Example D.19 (A branch created through a kernel singular at the reference game)**.** Let

$$R = \begin{pmatrix} 0 & 1 & -1 \\ -1 & 0 & 1 \\ 1 & -1 & 0 \end{pmatrix}, \qquad A(t) = tR.$$

For every $t \neq 0$, the game has the unique equilibrium $(\mathbf{1}_3/3, \mathbf{1}_3/3)$. Indeed, $R\mathbf{1}_3 = 0$, and if a row-optimal x satisfies $R^\mathsf{T} x \geq 0$, then the components sum to zero, so $R^\mathsf{T} x = 0$ and $x = \mathbf{1}_3/3$; similarly for the column player. At $t = 0$, the full-support bordered kernel is singular because

$$\begin{pmatrix} 0 & -\mathbf{1}_3 \\ \mathbf{1}_3^\mathsf{T} & 0 \end{pmatrix}$$

has every $(p, 0)$ with $\mathbf{1}_3^\mathsf{T} p = 0$ in its null space. Thus a nearby unique equilibrium germ is generated by a kernel singular at the reference game and is outside the regular analytic atlas used by the validity fan.

Example D.20 (Quadratic contact invisible to the first-order validity forms)**.** Consider the 3×2 game

$$A(t) = \begin{pmatrix} 1 + t & -1 \\ -1 & 1 \\ 1 & -1 + t \end{pmatrix}. \tag{D.46}$$

The branch supported on the first two rows and both columns has

$$x(t) = y(t) = \left(\frac{2}{4 + t}, \frac{2 + t}{4 + t}\right), \qquad v(t) = \frac{t}{4 + t}.$$

The off-support row-three slack is

$$v(t) - A_3(t)y(t) = -\frac{t^2}{4+t} = -\frac{1}{4}t^2 + O(t^3). \tag{D.47}$$

It is active at zero, its first derivative vanishes, and its quadratic term is strictly negative. Hence the branch is invalid for every sufficiently small $t \neq 0$, although first-order signs alone do not detect this.

Example D.21 (Fan faces and a three-way intersection)**.** In proposition D.18, the direction $u = v > 0$ lies on the active face $u - v = 0$ and leaves rows one and two tied; no unique branch is selected. At $(u, v) = (0, 0)$, all three pure row branches meet. The first-order fan therefore supplies chamberwise branches but no two-branch or unique-branch conclusion on these lower-dimensional sets.

D.11 Relation to existing stability theory

Shapley and Snow characterize basic equilibria through square kernels; modern accounts make the same cofactor and kernel structure explicit (Shapley & Snow, 1950; Attia & Oliu-Barton, 2021). Robinson's strong-regularity theory provides existence, uniqueness, and Lipschitz localization for regular generalized equations (Robinson, 1980), and recent Nash-equilibrium stability work derives differentiable single-valued localizations of KKT solution mappings under appropriate regularity hypotheses (Diao et al., 2026). The role of this appendix is different and more specialized: for finite zero-sum games, it constructs all regular basic branches explicitly, records their exact validity inequalities and derivatives, and assembles those inequalities into a perturbation-direction fan. The next appendix proves that the bordered-branch derivative agrees with the critical-cone linearization of the saddle-point generalized equation.

E Generalized Equations and Critical Cones

This appendix gives a variational characterization of first-order equilibrium motions. The strategy space is polyhedral, so the graphical derivative of its normal-cone mapping reduces exactly to the normal cone of a critical cone (Rockafellar & Wets, 1998). We specialize this reduction to zero-sum matrix games, derive an explicit product-simplex critical cone, formulate the resulting auxiliary saddle problem, and connect it to the regular kernel branches of section D. The canonical perturbation theory is classical (Dontchev & Rockafellar, 2001; 1996); the game-specific contribution here is the explicit identification of the tangent game and its exact agreement with the branch calculus used later in the statistical theory.

Throughout, equip matrix spaces with the Frobenius inner product and vector spaces with the Euclidean inner product. For a nonempty closed convex set $C \subseteq \mathbb{R}^d$, use the normal-cone convention

$$N_C(\bar{z}) \coloneqq \{w : \langle w, z - \bar{z}\rangle \leq 0 \text{ for every } z \in C\}, \qquad \bar{z} \in C, \tag{E.1}$$

and set $N_C(\bar{z}) = \varnothing$ when $\bar{z} \notin C$. The Bouligand tangent cone is

$$T_C(\bar{z}) \coloneqq \left\{u : \exists\, t_\nu \downarrow 0,\ z_\nu \in C,\ \frac{z_\nu - \bar{z}}{t_\nu} \to u\right\}. \tag{E.2}$$

For a set-valued map M, its graphical derivative at $(\bar{z}, \bar{w}) \in \operatorname{gph} M$ is

$$DM(\bar{z} \mid \bar{w})(u) \coloneqq \{q : (u, q) \in T_{\operatorname{gph} M}(\bar{z}, \bar{w})\}. \tag{E.3}$$

We call M *proto-differentiable* at $\bar{z}$ for $\bar{w}$ when $\operatorname{gph} M$ is geometrically derivable at $(\bar{z}, \bar{w})$; in that case (E.3) is the exact first-order graph of M, rather than only an outer tangent approximation.

E.1 The saddle-point generalized equation

Let

$$\mathcal{Z} \coloneqq \Delta_m \times \Delta_k, \qquad z = (x, y),$$

and define the affine saddle operator

$$F_A(x, y) \coloneqq \begin{pmatrix} -Ay \\ A^{\mathsf{T}}x \end{pmatrix}. \tag{E.4}$$

Its Jacobian in the strategy variable is the skew-symmetric linear operator

$$\mathcal{L}_A(p,q) \coloneqq \begin{pmatrix} -Aq \\ A^\mathsf{T} p \end{pmatrix}, \qquad \langle \mathcal{L}_A u, u \rangle = 0. \tag{E.5}$$

For a payoff perturbation H, its parameter derivative at $z = (x, y)$ is

$$\mathcal{B}_z(H) \coloneqq \begin{pmatrix} -Hy \\ H^\mathsf{T} x \end{pmatrix}. \tag{E.6}$$

Proposition E.1 (Saddle points as a generalized equation)**.** *For every payoff matrix A and every $z \in \mathcal{Z}$,*

$$z \in \mathcal{E}(A) \quad \Longleftrightarrow \quad 0 \in F_A(z) + N_{\mathcal{Z}}(z). \tag{E.7}$$

If $z_0 = (x_0, y_0) \in \mathcal{E}(A_0)$ and $v_0 = \mathrm{val}(A_0)$, then

$$w_0 \coloneqq -F_{A_0}(z_0) = \begin{pmatrix} A_0 y_0 \\ -A_0^\mathsf{T} x_0 \end{pmatrix} \in N_{\mathcal{Z}}(z_0). \tag{E.8}$$

Proof. Because normal cones commute with Cartesian products, (E.7) is equivalent to

$$Ay \in N_{\Delta_m}(x), \qquad -A^\mathsf{T} x \in N_{\Delta_k}(y).$$

The first inclusion says $\langle Ay, x' - x \rangle \le 0$ for every $x' \in \Delta_m$, so x maximizes $x' \mapsto x'^\mathsf{T} Ay$. The second says $\langle A^\mathsf{T} x, y' - y \rangle \ge 0$ for every $y' \in \Delta_k$, so y minimizes $y' \mapsto x^\mathsf{T} A y'$. These are precisely the saddle inequalities. The final statement follows by substituting A_0 and z_0. □

E.2 Explicit critical cones for the product simplex

Fix A_0 and $z_0 = (x_0, y_0) \in \mathcal{E}(A_0)$. Write

$$I_0 \coloneqq \mathrm{supp}(x_0), \qquad J_0 \coloneqq \mathrm{supp}(y_0), \tag{E.9}$$

and define the sets of tied pure best responses

$$\widehat{I}_0 \coloneqq \{i : (A_0 y_0)_i = v_0\}, \qquad \widehat{J}_0 \coloneqq \left\{j : (A_0^\mathsf{T} x_0)_j = v_0\right\}. \tag{E.10}$$

The saddle inequalities imply $I_0 \subseteq \widehat{I}_0$ and $J_0 \subseteq \widehat{J}_0$.

For $x \in \Delta_d$ with support I, the tangent and normal cones are

$$T_{\Delta_d}(x) = \left\{p : \mathbf{1}_d^\mathsf{T} p = 0,\ p_i \ge 0 \text{ for } i \notin I\right\}, \tag{E.11}$$

$$N_{\Delta_d}(x) = \{r : \exists \alpha \in \mathbb{R},\ r_i = \alpha \text{ for } i \in I,\ r_i \le \alpha \text{ for } i \notin I\}. \tag{E.12}$$

These formulas follow directly from the linear description of the simplex; more generally, tangent and normal cones to a polyhedral set are determined by its active linear constraints (Rockafellar & Wets, 1998).

Definition E.2 (Equilibrium critical cone)**.** The critical cone of the equilibrium generalized equation at (A_0, z_0) is

$$\mathcal{K}_0 \coloneqq T_{\mathcal{Z}}(z_0) \cap w_0^{\perp}, \tag{E.13}$$

where w_0 is given by (E.8).

Proposition E.3 (Product and support formula for the critical cone)**.** *The equilibrium critical cone factors as*

$$\mathcal{K}_0 = \mathcal{K}_x \times \mathcal{K}_y, \tag{E.14}$$

where

$$\mathcal{K}_x = \left\{p : \mathbf{1}_m^\mathsf{T} p = 0,\ p_i = 0 \textit{ for } i \notin \widehat{I}_0,\ p_i \ge 0 \textit{ for } i \in \widehat{I}_0 \setminus I_0\right\}, \tag{E.15}$$

$$\mathcal{K}_y = \left\{q : \mathbf{1}_k^\mathsf{T} q = 0,\ q_j = 0 \textit{ for } j \notin \widehat{J}_0,\ q_j \ge 0 \textit{ for } j \in \widehat{J}_0 \setminus J_0\right\}. \tag{E.16}$$

Thus a strictly suboptimal pure action receives zero first-order mass, whereas an unused action tied at the equilibrium value can receive nonnegative first-order mass.

Proof. The product formulas for tangent and normal cones give $T_{\mathcal{Z}}(z_0) = T_{\Delta_m}(x_0) \times T_{\Delta_k}(y_0)$. For every tangent pair (p, q), normality of the two components of w_0 implies

$$\langle A_0 y_0, p\rangle \le 0, \qquad \langle -A_0^{\mathsf{T}} x_0, q\rangle \le 0.$$

Their sum is zero if and only if both terms are zero, which proves (E.14).

For the row component, use $\mathbf{1}_m^{\mathsf{T}} p = 0$ to write

$$\langle A_0 y_0, p\rangle = \sum_{i=1}^{m} \big((A_0 y_0)_i - v_0\big) p_i.$$

The coefficient is zero on I_0, nonpositive off I_0, and every off-support coordinate of p is nonnegative. Hence the sum vanishes exactly when $p_i = 0$ at every strictly suboptimal row, yielding (E.15). The column formula follows similarly from

$$\langle -A_0^{\mathsf{T}} x_0, q\rangle = \sum_{j=1}^{k} \big(v_0 - (A_0^{\mathsf{T}} x_0)_j\big) q_j.$$

□

E.3 Polyhedral normal-cone differentiation

The graph of a normal-cone mapping is generally nonconvex even when the underlying set is convex. Polyhedrality nevertheless yields an exact local reduction.

Theorem E.4 (Graphical derivative of a polyhedral normal cone)**.** *Let $C \subseteq \mathbb{R}^d$ be nonempty and polyhedral, $\bar{z} \in C$, and $\bar{w} \in N_C(\bar{z})$. Put*

$$K_C(\bar{z}, \bar{w}) \coloneqq T_C(\bar{z}) \cap \bar{w}^{\perp}.$$

Then N_C is proto-differentiable at $\bar{z}$ for $\bar{w}$ and

$$DN_C(\bar{z} \mid \bar{w})(u) = N_{K_C(\bar{z},\bar{w})}(u) \qquad \textit{for every } u \in \mathbb{R}^d. \tag{E.17}$$

In particular,

$$DN_{\mathcal{Z}}(z_0 \mid w_0)(u) = N_{\mathcal{K}_0}(u). \tag{E.18}$$

Proof. For a polyhedral set, the reduction lemma identifies the graph of N_C around $(\bar{z}, \bar{w})$, after translation, with the graph of $N_{K_C(\bar{z},\bar{w})}$ around $(0, 0)$. Since the latter graph is a closed cone, its tangent cone at the origin is the graph itself. Therefore the tangent to $\operatorname{gph} N_C$ is exactly $\operatorname{gph} N_K$, which is (E.17). This is the polyhedral normal-cone formula of Rockafellar & Wets (1998, Example 13.44). Applying it to $C = \mathcal{Z}$ proves (E.18). □

E.4 Canonical perturbations and the tangent equilibrium system

To separate payoff perturbations from canonical right-hand perturbations, define

$$\mathcal{S}^{\mathrm{sol}}(A, r) \coloneqq \{z \in \mathcal{Z} : r \in F_A(z) + N_{\mathcal{Z}}(z)\}, \qquad r \in \mathbb{R}^{m+k}. \tag{E.19}$$

Thus $\mathcal{S}^{\mathrm{sol}}(A, 0) = \mathcal{E}(A)$. Write $r = (r_x, r_y)$ and a perturbation direction as $h = (h_x, h_y)$. The canonical parameter is important: the derivative with respect to r is $-\mathbf{I}_{m+k}$, so the parameterization is ample in the sense of Dontchev & Rockafellar (2001).

Theorem E.5 (Canonical equilibrium proto-derivative: exact ample-parameterization formula)**.** *The map $\mathcal{S}^{\mathrm{sol}}$ is proto-differentiable at $(A_0, 0)$ for every $z_0 \in \mathcal{E}(A_0)$. For every payoff direction H and canonical direction $h \in \mathbb{R}^{m+k}$,*

$$D\mathcal{S}^{\mathrm{sol}}((A_0, 0) \mid z_0)(H, h) = \Big\{u : \ 0 \in \mathcal{L}_{A_0} u + \mathcal{B}_{z_0}(H) - h + N_{\mathcal{K}_0}(u)\Big\}. \tag{E.20}$$

The graph of the derivative map in (E.20) *is a finite union of polyhedral cones; in particular, the derivative is positively homogeneous and piecewise polyhedral.*

Proof. We give a direct graph-coordinate proof, which also verifies the ample- parameterization invocation of Dontchev & Rockafellar (2001, Theorem 6.1). Put $N = N_{\mathcal{Z}}$ and consider the ambient C^1 bijection

$$\Theta(A, z, w) = \big((A, F_A(z) + w), z\big),$$

with inverse

$$\Theta^{-1}((A, r), z) = \big(A, z, r - F_A(z)\big).$$

Both maps are polynomial. Moreover,

$$\operatorname{gph} \mathcal{S}^{\mathrm{sol}} = \Theta\big(\mathbb{R}^{m\times k} \times \operatorname{gph} N\big).$$

Thus the graph of the solution map is a smooth change of coordinates of the product of payoff space with $\operatorname{gph} N$. Let $w_0 = -F_{A_0}(z_0) \in N(z_0)$. Differentiating Θ at (A_0, z_0, w_0) gives

$$D\Theta(A_0, z_0, w_0)[H, u, \omega] = \big((H, \mathcal{B}_{z_0}(H) + \mathcal{L}_{A_0} u + \omega), u\big).$$

Because Θ is a C^1 diffeomorphism, tangent cones transform exactly. Consequently,

$$\begin{aligned} u \in D\mathcal{S}^{\mathrm{sol}}((A_0, 0) \mid z_0)(H, h) \quad &\Longleftrightarrow \quad \big(u, h - \mathcal{B}_{z_0}(H) - \mathcal{L}_{A_0} u\big) \in T_{\operatorname{gph} N}(z_0, w_0) \\ &\Longleftrightarrow \quad h - \mathcal{B}_{z_0}(H) - \mathcal{L}_{A_0} u \in DN(z_0 \mid w_0)(u). \end{aligned}$$

By theorem E.4,

$$DN(z_0 \mid w_0)(u) = N_{\mathcal{K}_0}(u).$$

Rearranging yields exactly (E.20). This direct coordinate-change argument proves equality rather than an outer inclusion and also establishes proto-differentiability of $\mathcal{S}^{\mathrm{sol}}$.

For comparison with the general theory, the smooth inclusion map is $f((A, r), z) = F_A(z) - r$ and its parameter derivative is $\mathcal{B}_{z_0}(H) - h$. The h-block is $-\mathbf{I}_{m+k}$, so the full-row- rank ample condition of Dontchev & Rockafellar (2001, Theorem 6.1) holds. Finally, $\mathcal{K}_0$ is polyhedral, hence the graph of its normal-cone map is a finite union of polyhedral cones; the displayed linear preimage has the same property. □

For later use, denote the right-hand side of (E.20) by

$$\mathfrak{T}_{A_0, z_0}(H, h). \tag{E.21}$$

This is the *critical-cone tangent equilibrium correspondence*.

E.5 The tangent game

The inclusion in (E.20) is itself a zero-sum saddle problem, but on the critical cones rather than the original simplices. For $u = (p, q) \in \mathcal{K}_x \times \mathcal{K}_y$, define

$$\Gamma_{H,h}(p, q) \coloneqq p^{\mathsf{T}} A_0 q + p^{\mathsf{T}} H y_0 + x_0^{\mathsf{T}} H q + h_x^{\mathsf{T}} p - h_y^{\mathsf{T}} q. \tag{E.22}$$

Proposition E.6 (Equivalent tangent-game formulations)**.** *For $u = (p, q)$, the following are equivalent:*

1. $u \in \mathfrak{T}_{A_0, z_0}(H, h)$*;*
2. $u \in \mathcal{K}_0$ *and, for every* $(p', q') \in \mathcal{K}_x \times \mathcal{K}_y$*,*

$$\Gamma_{H,h}(p', q) \le \Gamma_{H,h}(p, q) \le \Gamma_{H,h}(p, q'); \tag{E.23}$$

3. *with*

$$g = \mathcal{L}_{A_0} u + \mathcal{B}_{z_0}(H) - h,$$

one has the conic complementarity conditions

$$u \in \mathcal{K}_0, \qquad g \in \mathcal{K}_0^*, \qquad \langle g, u\rangle = 0, \tag{E.24}$$

where $\mathcal{K}_0^* = \{g : \langle g, v\rangle \ge 0 \ \forall v \in \mathcal{K}_0\}$.

The tangent game may have no saddle point; emptiness records the absence of a finite first-order equilibrium motion based at z_0 in the specified direction.

Proof. The row and column gradients of (E.22) are

$$\nabla_p \Gamma_{H,h} = A_0 q + H y_0 + h_x, \qquad \nabla_q \Gamma_{H,h} = A_0^{\mathsf{T}} p + H^{\mathsf{T}} x_0 - h_y.$$

Therefore the saddle-point generalized equation for $\Gamma_{H,h}$ on $\mathcal{K}_x \times \mathcal{K}_y$ is exactly (E.20), proving the equivalence of (i) and (ii). Since $\mathcal{K}_0$ is a closed convex cone,

$$N_{\mathcal{K}_0}(u) = \mathcal{K}_0^\circ \cap u^\perp \qquad (u \in \mathcal{K}_0),$$

where the polar cone uses the inequality $\langle a, v\rangle \le 0$. Thus $-g \in N_{\mathcal{K}_0}(u)$ is equivalent to (E.24). □

$$\boxed{\text{bordered branch derivative} \iff \text{critical-cone tangent solution} \implies \text{statistical target chart.}} \tag{E.25}$$

E.6 Payoff-only derivatives and kernel-branch equivalence

Let

$$\mathfrak{E}(A) \coloneqq \mathcal{E}(A) \tag{E.26}$$

be the payoff-only equilibrium correspondence. Given the perturbation subspace $\mathbb{T}$ fixed in section D, also define its relative restriction at A_0 by

$$\mathfrak{E}_{\mathbb{T}}(G) \coloneqq \mathcal{E}(A_0 + G), \qquad G \in \mathbb{T}. \tag{E.27}$$

The unrestricted payoff parameterization is not ample in general: for fixed $z_0 = (x_0, y_0)$, the range of $H \mapsto \mathcal{B}_{z_0}(H)$ satisfies the linear identity

$$x_0^{\mathsf{T}}(-H y_0) + y_0^{\mathsf{T}}(H^{\mathsf{T}} x_0) = 0. \tag{E.28}$$

Accordingly, we first state an unconditional necessary condition and then use the branch atlas of section D to obtain exact payoff-only derivatives on fan chambers.

Proposition E.7 (Critical-cone necessity for payoff perturbations)**.** *For every H,*

$$D\mathfrak{E}(A_0 \mid z_0)(H) \subseteq \mathfrak{T}_{A_0, z_0}(H, 0). \tag{E.29}$$

Proof. Let $u \in D\mathfrak{E}(A_0 \mid z_0)(H)$. By definition there exist $t_\nu \downarrow 0$, $H_\nu \to H$, and $z_\nu \in \mathcal{E}(A_0 + t_\nu H_\nu)$ such that $z_\nu \to z_0$ and $(z_\nu - z_0)/t_\nu \to u$. Put $w_\nu = -F_{A_0 + t_\nu H_\nu}(z_\nu)$, so $w_\nu \in N_{\mathcal{Z}}(z_\nu)$. Smoothness of F gives

$$\frac{w_\nu - w_0}{t_\nu} \longrightarrow -\mathcal{L}_{A_0} u - \mathcal{B}_{z_0}(H).$$

Hence $-\mathcal{L}_{A_0} u - \mathcal{B}_{z_0}(H) \in DN_{\mathcal{Z}}(z_0 \mid w_0)(u)$. Apply (E.18). □

Theorem E.8 (Equivalence with regular kernel branches)**.** *Suppose assumptions D.5 and D.7 hold at A_0 relative to the perturbation space $\mathbb{T}$. Let C be a chamber of $\mathfrak{F}(A_0; \mathbb{T})$, choose a representative regular kernel $\beta(C)$ of its unique equilibrium branch, and put*

$$z_C^0 \coloneqq z_{\beta(C)}(A_0).$$

Then, for every nonzero $H \in C$ and every $z_0 \in \mathcal{E}(A_0)$, the graphical derivative of the restricted *equilibrium map in* (E.27) *is*

$$D\mathfrak{E}_{\mathbb{T}}(0 \mid z_0)(H) = \begin{cases} \{D z_{\beta(C)}(A_0)[H]\}, & z_0 = z_C^0, \\ \varnothing, & z_0 \neq z_C^0. \end{cases} \tag{E.30}$$

In particular,

$$D z_{\beta(C)}(A_0)[H] \in \mathfrak{T}_{A_0, z_C^0}(H, 0), \tag{E.31}$$

so the bordered-system derivative in proposition D.4 solves the critical-cone tangent game. If $\mathbb{T} = \mathbb{R}^{m \times k}$, then (E.30) *is the unrestricted graphical derivative of the payoff-equilibrium correspondence.*

Proof. Fix $H \in C$. A tangent to the restricted map is generated by sequences $t_\nu \downarrow 0$, $H_\nu \in \mathbb{T}$ with $H_\nu \to H$, and $z_\nu \in \mathcal{E}(A_0 + t_\nu H_\nu)$. Since C is relatively open in $\mathbb{T}$, one has $H_\nu \in C$ eventually. Put $r_\nu = \|H_\nu\|_\infty$, $\widetilde{H}_\nu = H_\nu / r_\nu$, and $s_\nu = t_\nu r_\nu$. Because $H \neq 0$, the normalized directions $\widetilde{H}_\nu$ eventually lie in a compact subset of $C \cap \mathbb{S}_\infty(\mathbb{T})$. Hence theorem D.9 implies that, for all large ν,

$$z_\nu = z_{\beta(C)}(A_0 + s_\nu \widetilde{H}_\nu) = z_{\beta(C)}(A_0 + t_\nu H_\nu).$$

Using the uniform analytic branch expansion from theorem D.9, together with linearity and continuity of $Dz_{\beta(C)}(A_0)$, gives

$$\frac{z_\nu - z^0_C}{t_\nu} = Dz_{\beta(C)}(A_0)[H_\nu] + o(1) \longrightarrow Dz_{\beta(C)}(A_0)[H].$$

Conversely, the constant direction sequence $H_\nu = H$ and the analytic branch itself realize this derivative. This proves the nonempty case in (E.30); uniqueness of the perturbed equilibrium also shows that no tangent based at a different z_0 can exist. Finally, (E.31) follows from proposition E.7. □

At a strictly valid regular kernel branch, the critical system is not merely necessary: it is uniquely solvable for every canonical perturbation.

Proposition E.9 (Strict regular branches)**.** *Let $\beta = (I, J)$ satisfy $\Delta_\beta(A_0) \neq 0$ and $\rho_\beta(A_0) > 0$, and put $z_0 = z_\beta(A_0)$. Then $\widehat{I}_0 = I$, $\widehat{J}_0 = J$, and*

$$\mathcal{K}_x = \left\{p : \mathrm{supp}(p) \subseteq I,\ \mathbf{1}_I^\mathsf{T} p_I = 0\right\}, \qquad \mathcal{K}_y = \left\{q : \mathrm{supp}(q) \subseteq J,\ \mathbf{1}_J^\mathsf{T} q_J = 0\right\} \tag{E.32}$$

are linear subspaces. For every (H, h), the canonical tangent system has a unique solution $u = (p, q)$, given on the supports by

$$\begin{pmatrix} q_J \\ \lambda_x \end{pmatrix} = M^y_\beta(A_0)^{-1} \begin{pmatrix} -H_{I,J} y_{0,J} - (h_x)_I \\ 0 \end{pmatrix}, \tag{E.33}$$

$$\begin{pmatrix} p_I \\ \lambda_y \end{pmatrix} = M^x_\beta(A_0)^{-1} \begin{pmatrix} -H^\mathsf{T}_{I,J} x_{0,I} + (h_y)_J \\ 0 \end{pmatrix}, \tag{E.34}$$

with all off-support coordinates equal to zero. For $h = 0$, $\lambda_x = \lambda_y = x_0^\mathsf{T} H y_0$ and $(p, q) = Dz_\beta(A_0)[H]$. Consequently, the canonical generalized equation is strongly regular at the strict branch in the sense of Robinson.

Proof. Strict positivity of the off-support reduced slacks implies that no unused pure action is tied at the value, giving $\widehat{I}_0 = I$ and $\widehat{J}_0 = J$. Hence (E.32) follows from proposition E.3. On a linear subspace K, $N_K(u) = K^\perp$ for every $u \in K$. Therefore the tangent inclusion is equivalent to requiring

$$A_{0,I,J} q_J + H_{I,J} y_{0,J} + (h_x)_I = \lambda_x \mathbf{1}_I,$$
$$A^\mathsf{T}_{0,I,J} p_I + H^\mathsf{T}_{I,J} x_{0,I} - (h_y)_J = \lambda_y \mathbf{1}_J,$$

together with the two zero-sum constraints. These are exactly (E.33)– (E.34). The bordered matrices are nonsingular, proving uniqueness. When $h = 0$, multiplying the first equation by $x^\mathsf{T}_{0,I}$ and the second by $y^\mathsf{T}_{0,J}$ shows that both scalars equal $x_0^\mathsf{T} H y_0$; comparison with (D.15) gives the final identity. It remains to justify the strong-regularity conclusion rather than merely the uniqueness of the tangent equation. Strict positivity of every supported mass and every off-support slack persists in a neighborhood of $(A_0, z_0, 0)$. Hence the active faces of the two simplices are fixed there, and the canonical right-hand-side generalized equation reduces exactly to the two bordered systems obtained above, with coefficients depending continuously (indeed, analytically) on the payoff and right-hand side. Their determinants are nonzero at the reference point and therefore remain nonzero locally. The implicit-function theorem yields a unique analytic solution map $(A, r) \mapsto z(A, r)$ on this neighborhood. On a smaller neighborhood its Jacobian is bounded, so the solution map is locally Lipschitz. Thus the localized inverse of $z \mapsto F_{A_0}(z) + N_{\mathcal{Z}}(z)$ is single-valued and Lipschitz around $(0, z_0)$, which is precisely Robinson strong regularity. This direct verification is consistent with the classical polyhedral characterizations in Robinson (1980); Dontchev & Rockafellar (1996). □

E.7 THE VALUE REMAINS DIRECTIONALLY REGULAR

The equilibrium correspondence may have an empty or multivalued graphical derivative at a singular game. The scalar value nevertheless has a single-valued Hadamard directional derivative everywhere.

Proposition E.10 (Hadamard directional derivative of the game value)**.** *For every A_0 and H,*

$$\begin{aligned}\mathrm{val}'(A_0; H) &:= \lim_{\substack{t\downarrow 0\\ H'\to H}} \frac{\mathrm{val}(A_0 + tH') - \mathrm{val}(A_0)}{t}\\ &= \max_{x\in\mathcal{X}^\star(A_0)} \min_{y\in\mathcal{Y}^\star(A_0)} x^\mathsf{T} H y\\ &= \min_{y\in\mathcal{Y}^\star(A_0)} \max_{x\in\mathcal{X}^\star(A_0)} x^\mathsf{T} H y.\end{aligned} \tag{E.35}$$

Thus the value derivative is the value of the perturbation game H restricted to the original equilibrium-strategy polytopes.

Proof. The last two expressions in (E.35) agree by the minimax theorem because the two equilibrium-strategy sets are nonempty, compact, and convex. Let $t_\nu \downarrow 0$ and $H_\nu \to H$, and choose $(x_\nu, y_\nu) \in \mathcal{E}(A_0 + t_\nu H_\nu)$. The value perturbation bound in proposition C.1 keeps the difference quotients bounded. Along an arbitrary convergent subsequence, write $(x_\nu, y_\nu) \to (\bar{x}, \bar{y})$. Continuity of the value and the saddle inequalities imply $\bar{x} \in \mathcal{X}^\star(A_0)$ and $\bar{y} \in \mathcal{Y}^\star(A_0)$.

For every $x \in \mathcal{X}^\star(A_0)$,

$$\mathrm{val}(A_0 + t_\nu H_\nu) \geq x^\mathsf{T}(A_0 + t_\nu H_\nu) y_\nu \geq \mathrm{val}(A_0) + t_\nu x^\mathsf{T} H_\nu y_\nu.$$

Hence every subsequential lower limit of the difference quotient is at least $\max_{x\in\mathcal{X}^\star(A_0)} x^\mathsf{T} H \bar{y}$, which is no smaller than the restricted-game value in (E.35). Similarly, for every $y \in \mathcal{Y}^\star(A_0)$,

$$\mathrm{val}(A_0 + t_\nu H_\nu) \leq x_\nu^\mathsf{T}(A_0 + t_\nu H_\nu) y \leq \mathrm{val}(A_0) + t_\nu x_\nu^\mathsf{T} H_\nu y.$$

Every subsequential upper limit is therefore at most $\min_{y\in\mathcal{Y}^\star(A_0)} \bar{x}^\mathsf{T} H y$, which is no larger than the same restricted-game value. All convergent subsequences have the same limit, and boundedness of the quotient then proves the Hadamard derivative formula. □

E.8 CRITICAL-CONE DIAGNOSIS OF THE TWO-BY-TWO JUMP

Example E.11 (Finite derivatives on one side and no derivative on the other)**.** Let

$$H_\star = \begin{pmatrix} 0 & 1\\ -1 & 0\end{pmatrix}, \qquad A_\theta = \theta H_\star,$$

as in equation (B.1), and take $A_0 = 0$. At $z_+ = (e_1, e_1)$, both critical-cone factors equal the ray

$$R_+ := \{a(-1, 1) : a \geq 0\}.$$

Since $\mathcal{L}_{A_0} = 0$ and

$$\mathcal{B}_{z_+}(H_\star) = \big((0, 1), (0, 1)\big),$$

the tangent system has the unique solution $u = 0$. This is the derivative of the constant equilibrium branch for $\theta > 0$. In the opposite direction,

$$\mathcal{B}_{z_+}(-H_\star) = \big((0, -1), (0, -1)\big),$$

and the tangent system has no solution: for each ray factor, the vector $(0, 1)$ belongs neither to the normal cone at the origin nor to the normal space at a nonzero ray point. Hence

$$D\mathfrak{E}(0 \mid z_+)(H_\star) = \{0\}, \qquad D\mathfrak{E}(0 \mid z_+)(-H_\star) = \varnothing.$$

The equilibrium therefore leaves every $O(t)$-neighborhood of z_+ under the negative perturbation and jumps to (e_2, e_2).

At the maximum-entropy equilibrium $\bar{z} = ((1/2, 1/2), (1/2, 1/2))$, both critical-cone factors are the zero-sum line $\{a(-1, 1) : a \in \mathbb{R}\}$. The vector $\mathcal{B}_{\bar{z}}(\pm H_\star)$ has a nonzero component in this line and therefore cannot be canceled by its orthogonal normal space. Thus

$$D\mathfrak{E}(0 \mid \bar{z})(H_\star) = D\mathfrak{E}(0 \mid \bar{z})(-H_\star) = \varnothing,$$

which is the first-order signature of the point-selector discontinuity. By contrast, proposition E.10 gives $\mathrm{val}'(0; \pm H_\star) = 0$.

E.9 Relation to existing variational stability theory

The normal-cone derivative $DN_C = N_K$ for polyhedral C and the associated auxiliary variational inequality are classical consequences of variational analysis (Rockafellar & Wets, 1998). Canonical and ample parameterizations of such inclusions yield exact proto-derivative formulas (Dontchev & Rockafellar, 2001), while strong regularity characterizes locally single-valued Lipschitz solution mappings over polyhedral sets (Dontchev & Rockafellar, 1996; Robinson, 1980). Recent work develops related regularity criteria for Nash KKT mappings (Diao et al., 2026).

The role of this appendix is narrower and game-specific. It identifies the critical cone in terms of supported and tied pure best responses, interprets the linearized inclusion as a zero-sum tangent game, and proves exact agreement with the regular kernel branches of section D. This is the bridge from deterministic support geometry to the loss-dependent Gaussian limit experiments developed in the next appendix.

F Gaussian Local Experiments and Local Minimax Theory

This appendix couples the deterministic tangent geometry of sections D and E to statistical decision theory. For a deterministic asymptotic sampling design, the local payoff experiment is not merely locally asymptotically normal: after reduction to a sufficient statistic, it is an *exact* finite-dimensional Gaussian shift for every sample size. We exploit this structure to prove a target-agnostic local minimax transfer theorem. Any equilibrium functional admitting a uniform local tangent chart induces a Gaussian decision problem whose target is that tangent chart. The value yields a nonlinear restricted-game target at singular points; regular strategy branches yield linear targets; and continuous support kinks yield piecewise-linear targets. The classical local asymptotic minimax principle goes back to Hájek (1972) and Le Cam (1986); our contribution is the equilibrium-specific target geometry and an exact comparison argument that avoids an appeal to an abstract convergence-of-experiments theorem. The relation to inference for directionally differentiable maps is discussed after proposition F.13; see also Fang & Santos (2019). Throughout this appendix, the center A_0, perturbation space $\mathbb{T}$, and deterministic sampling design are fixed components of the local experiment; procedures may depend on them. Constructing a globally adaptive localization scheme is a separate problem and is not claimed here.

F.1 Structured local payoff models and exact Gaussian reduction

Let $p := mk$, identify $\mathbb{R}^{m\times k}$ with $\mathbb{R}^p$ by column-wise vectorization, and fix a d-dimensional linear perturbation space $\mathbb{T} \subseteq \mathbb{R}^{m\times k}$. Choose a Frobenius-orthonormal basis $H^1, \dots, H^d$ of $\mathbb{T}$ and define the linear isometry

$$\mathcal{U} : \mathbb{R}^d \to \mathbb{T}, \qquad \mathcal{U}h := \sum_{\ell=1}^{d} h_\ell H^\ell. \tag{F.1}$$

Let $U \in \mathbb{R}^{p\times d}$ have columns $\mathrm{vec}(H^1), \dots, \mathrm{vec}(H^d)$, so that $\mathrm{vec}(\mathcal{U}h) = Uh$ and $U^\mathsf{T}U = \mathbf{I}_d$. We assume that A_0 is relatively interior to the payoff model along $\mathbb{T}$: for some $\epsilon_0 > 0$,

$$A_0 + \{H \in \mathbb{T} : \|H\|_F < \epsilon_0\} \subseteq [-1, 1]^{m\times k}. \tag{F.2}$$

Consequently, every bounded local parameter set is feasible for all sufficiently large n.

For each n, let $n_{ij} \in \mathbb{N}_0$ be deterministic query counts with $\sum_{i,j} n_{ij} = n$. Conditional on the unknown payoff matrix A, the data are independent and satisfy

$$Y_{ij,r} \sim \mathcal{N}(A_{ij}, \sigma_{ij}^2), \qquad 1 \le r \le n_{ij}, \tag{F.3}$$

where $0 < \sigma_{ij}^2 < \infty$ are known. Write

$$D_n := \mathrm{diag}\left(\frac{n_{ij}}{n\sigma_{ij}^2}\right)_{i,j}, \qquad J_n := U^\mathsf{T} D_n U. \tag{F.4}$$

If $n_{ij} = 0$, set $\overline{Y}_{ij} = (A_0)_{ij}$ by convention; its value is irrelevant because the corresponding diagonal entry of D_n is zero.

Assumption F.1 (Identifiable deterministic local design)**.** There exists n_0 such that $J_n \succ 0$ for all $n \geq n_0$, and

$$J_n \longrightarrow J \succ 0. \tag{F.5}$$

We write $\Sigma_n \coloneqq J_n^{-1}$ and $\Sigma \coloneqq J^{-1}$.

For $h \in \mathbb{R}^d$, define the local payoff matrix

$$A_{n,h} \coloneqq A_0 + \frac{\mathcal{U}h}{\sqrt{n}}. \tag{F.6}$$

For experiment-comparison statements we use the algebraic Gaussian extension of this family to all $h \in \mathbb{R}^d$. By (F.2), on every fixed compact local set this extension coincides with the original bounded-payoff model for all sufficiently large n. Define the normalized score and its information-normalized version by

$$\Delta_n \coloneqq U^{\mathsf{T}} D_n \sqrt{n} \left(\mathrm{vec}(\overline{Y}_n) - \mathrm{vec}(A_0) \right), \tag{F.7}$$

$$Z_n \coloneqq J_n^{-1} \Delta_n. \tag{F.8}$$

Proposition F.2 (Exact semidefinite Gaussian score reduction)**.** *Let $F \in \mathbb{R}^{N \times d}$ be arbitrary and suppose*

$$Y \sim \mathcal{N}(Fh, \mathbf{I}_N), \qquad h \in \mathbb{R}^d. \tag{F.9}$$

Define

$$\Delta = F^{\mathsf{T}} Y, \qquad J = F^{\mathsf{T}} F \succeq 0. \tag{F.10}$$

Then the experiment in (F.9) *is Blackwell-equivalent to*

$$\Delta \sim \mathcal{N}(Jh, J). \tag{F.11}$$

The equivalence is exact even when J is singular.

For deterministic Gaussian entry counts n_s, the normalized whitened design has the row $u_s^{\mathsf{T}}/(\sigma_s\sqrt{n})$ repeated n_s times. Hence

$$J_n = \sum_s \frac{n_s}{n\sigma_s^2} u_s u_s^{\mathsf{T}}, \tag{F.12}$$

and the complete local payoff experiment is exactly equivalent to $\Delta_n \sim \mathcal{N}(J_n h, J_n)$.

Proof. The statistic Δ has the law in (F.11). Let P_F denote the orthogonal projector onto $\mathrm{range}(F)$ and write $Y = P_F Y + (I - P_F)Y$. The residual $(I - P_F)Y$ is standard Gaussian on $\mathrm{range}(F)^{\perp}$, is independent of $P_F Y$, and has a law independent of h. Moreover,

$$P_F Y = F J^{\dagger} F^{\mathsf{T}} Y = F J^{\dagger} \Delta. \tag{F.13}$$

Thus Y is reconstructed from Δ by the parameter-independent kernel that returns $F J^{\dagger}\Delta + R$, where $R \sim \mathcal{N}(0, I - P_F)$ is independent. The reverse map $Y \mapsto F^{\mathsf{T}} Y$ is deterministic. Therefore both deficiencies are zero. The entrywise specialization follows after whitening each observed sample and collecting the repeated feature rows. □

Remark F.3 (Singular score reduction versus full-model transfer). Proposition F.2 is exact for every $J_n \succeq 0$ and therefore applies to sparse interface-side designs. The Gaussian-shift target transfer and the regular information theorem below use assumption F.1, namely $J_n \to J \succ 0$. They must not be applied to a design that identifies only one side functional while leaving other directions of the declared perturbation model unidentified.

Proposition F.4 (Exact Gaussian reduction)**.** *Under Assumption F.1, the experiment generated by all observations in* (F.3)*, under the local model $\{A_{n,h} : h \in \mathbb{R}^d\}$, is statistically equivalent to observing*

$$Z_n \sim \mathcal{N}(h, \Sigma_n). \tag{F.14}$$

More precisely, Z_n is sufficient for h, and the conditional law of the full data given Z_n is independent of h.

Proof. Relative to $h = 0$, the Gaussian log-likelihood ratio equals

$$\log \frac{dP_{n,h}}{dP_{n,0}} = \sum_{i,j} \sum_{r=1}^{n_{ij}} \left[\frac{Y_{ij,r} - (A_0)_{ij}}{\sigma_{ij}^2} \frac{(\mathcal{U}h)_{ij}}{\sqrt{n}} - \frac{(\mathcal{U}h)_{ij}^2}{2n\sigma_{ij}^2} \right]$$
$$= h^\mathsf{T} \Delta_n - \frac{1}{2} h^\mathsf{T} J_n h.$$

The factorization criterion therefore makes Δ_n, and equivalently Z_n, sufficient. Under $P_{n,h}$,

$$\Delta_n \sim \mathcal{N}(J_n h, J_n),$$

which proves (F.14). In this dominated Gaussian family, sufficiency also yields a parameter-independent regular conditional law of the full observation given Z_n; hence the two experiments simulate one another by Markov kernels and are statistically equivalent. □

We use total variation in the convention

$$d_{\mathrm{TV}}(P, Q) \coloneqq \sup_B |P(B) - Q(B)|. \tag{F.15}$$

Lemma F.5 (Uniform convergence of the Gaussian shifts)**.** *Under Assumption F.1,*

$$\eta_n \coloneqq \sup_{h \in \mathbb{R}^d} d_{\mathrm{TV}} \left(\mathcal{N}(h, \Sigma_n), \mathcal{N}(h, \Sigma) \right) = d_{\mathrm{TV}} \left(\mathcal{N}(0, \Sigma_n), \mathcal{N}(0, \Sigma) \right) \longrightarrow 0. \tag{F.16}$$

Proof. Translation invariance of total variation gives the equality. Moreover,

$$D_{\mathrm{KL}} \left(\mathcal{N}(0, \Sigma_n) \| \mathcal{N}(0, \Sigma) \right) = \frac{1}{2} \left[\mathrm{tr}(\Sigma^{-1} \Sigma_n) - d - \log \det(\Sigma^{-1} \Sigma_n) \right] \longrightarrow 0$$

by $\Sigma_n \to \Sigma \succ 0$. Pinsker's inequality then implies (F.16). □

F.2 Target charts and a local minimax transfer theorem

The statistical observation experiment is common to all targets; only the local target chart changes. This is the source of the target-dependent statistical complexity developed in the paper.

Let $(\mathcal{Y}, d_{\mathcal{Y}})$ be a metric target space and $\mathcal{D}$ a standard Borel decision space. Let

$$\ell : \mathcal{D} \times \mathcal{Y} \to [0, M]$$

be measurable and uniformly continuous in its target argument, uniformly over decisions, in the sense that

$$\omega_\ell(r) \coloneqq \sup_{a \in \mathcal{D}} \sup_{d_{\mathcal{Y}}(y, y') \le r} |\ell(a, y) - \ell(a, y')| \longrightarrow 0 \qquad (r \downarrow 0). \tag{F.17}$$

Randomized procedures are allowed and are represented by Markov kernels from the observation space into $\mathcal{D}$.

Definition F.6 (Uniform local tangent chart)**.** Let $\mathcal{H} \subseteq \mathbb{R}^d$ be compact. A sequence of measurable target maps $\psi_n : \mathcal{H} \to \mathcal{Y}$ admits the measurable tangent chart $\phi : \mathcal{H} \to \mathcal{Y}$ when

$$\varepsilon_n(\mathcal{H}) \coloneqq \sup_{h \in \mathcal{H}} d_{\mathcal{Y}}(\psi_n(h), \phi(h)) \longrightarrow 0. \tag{F.18}$$

Define the finite-sample local minimax risk

$$\mathfrak{R}_n(\mathcal{H}; \psi_n, \ell) \coloneqq \inf_{\delta_n} \sup_{h \in \mathcal{H}} \mathbb{E}_{n,h} \left[\ell(\delta_n, \psi_n(h)) \right], \tag{F.19}$$

where the infimum ranges over all procedures based on the complete data, and define the Gaussian tangent risk

$$\mathfrak{R}_G(\mathcal{H}; \phi, \ell) \coloneqq \inf_{\delta} \sup_{h \in \mathcal{H}} \mathbb{E} \left[\ell \left(\delta(h + \Sigma^{1/2} G), \phi(h) \right) \right], \qquad G \sim \mathcal{N}(0, \mathbf{I}_d). \tag{F.20}$$

Theorem F.7 (Target-chart minimax transfer in an exact Gaussian score experiment)**.** *Suppose Assumption F.1 and Definition F.6 hold. Then*

$$|\mathfrak{R}_n(\mathcal{H};\psi_n,\ell)-\mathfrak{R}_G(\mathcal{H};\phi,\ell)|\le M\eta_n+\omega_\ell\left(\varepsilon_n(\mathcal{H})\right), \tag{F.21}$$

where η_n is defined in (F.16)*. Consequently,*

$$\mathfrak{R}_n(\mathcal{H};\psi_n,\ell)\longrightarrow\mathfrak{R}_G(\mathcal{H};\phi,\ell). \tag{F.22}$$

Proof. By proposition F.4, it is without loss to restrict every finite-sample procedure to a Markov kernel applied to Z_n. Fix such a kernel δ. For each $h\in\mathcal{H}$, contraction of total variation under a Markov kernel and (F.15) give

$$\begin{aligned}&\left|\mathbb{E}_{\mathcal{N}(h,\Sigma_n)}\ell(\delta,\psi_n(h))-\mathbb{E}_{\mathcal{N}(h,\Sigma)}\ell(\delta,\phi(h))\right|\\&\quad\le Md_{\mathrm{TV}}\left(\mathcal{N}(h,\Sigma_n),\mathcal{N}(h,\Sigma)\right)+\omega_\ell\left(d_{\mathcal{Y}}(\psi_n(h),\phi(h))\right)\\&\quad\le M\eta_n+\omega_\ell(\varepsilon_n(\mathcal{H})).\end{aligned}$$

Taking the supremum over h preserves this bound. Taking the infimum first over kernels in the finite experiment and then over the same kernels in the limit experiment proves one inequality in (F.21); the reverse comparison is identical. The limit follows from lemma F.5 and (F.18). □

The bounded-loss assumption is not restrictive on a fixed compact local set for quadratic loss, because decisions can be projected onto a compact convex set containing every local target.

Corollary F.8 (Quadratic loss on compact local sets)**.** *Let $\mathcal{Y}=\mathcal{D}=\mathbb{R}^r$, let $W\succ 0$, and write $\|u\|_W^2=u^{\mathsf{T}}Wu$. If ϕ is continuous on compact $\mathcal{H}$ and $\psi_n\to\phi$ uniformly there, then*

$$\inf_{\delta_n}\sup_{h\in\mathcal{H}}\mathbb{E}_{n,h}\|\delta_n-\psi_n(h)\|_W^2\longrightarrow\inf_{\delta}\sup_{h\in\mathcal{H}}\mathbb{E}\left[\|\delta(h+\Sigma^{1/2}G)-\phi(h)\|_W^2\right]. \tag{F.23}$$

Proof. Uniform convergence and compactness imply that, for all sufficiently large n, the union of $\psi_n(\mathcal{H})$ and $\phi(\mathcal{H})$ lies in one closed W-ball C. The metric projection onto C cannot increase squared W-distance to any point of C. Hence both minimax problems may be restricted to decisions in C, on which quadratic loss is bounded and uniformly continuous. Apply theorem F.7. □

F.3 Equilibrium target charts

We now verify the uniform tangent condition for the two continuous equilibrium targets used later: the scalar value and a continuous point-selector branch.

F.3.1 The value target

Define, for $h\in\mathbb{R}^d$,

$$\psi_{\mathrm{val},n}(h)\coloneqq\sqrt{n}\left[\mathrm{val}(A_{n,h})-\mathrm{val}(A_0)\right], \tag{F.24}$$

$$\phi_{\mathrm{val}}(h)\coloneqq\max_{x\in\mathcal{X}^\star(A_0)}\min_{y\in\mathcal{Y}^\star(A_0)}x^{\mathsf{T}}(\mathcal{U}h)y. \tag{F.25}$$

Proposition F.9 (Uniform value tangent chart)**.** *For every compact $\mathcal{H}\subseteq\mathbb{R}^d$,*

$$\sup_{h\in\mathcal{H}}|\psi_{\mathrm{val},n}(h)-\phi_{\mathrm{val}}(h)|\longrightarrow 0. \tag{F.26}$$

Consequently, under every bounded loss satisfying (F.17)*, and under quadratic loss on compact local sets, the local minimax risk for value estimation converges to the Gaussian shift problem with target* (F.25)*.*

Proof. By proposition E.10, the value is Hadamard directionally differentiable at A_0, with derivative

$$H\longmapsto\max_{x\in\mathcal{X}^\star(A_0)}\min_{y\in\mathcal{Y}^\star(A_0)}x^{\mathsf{T}}Hy.$$

This derivative is continuous, indeed 1-Lipschitz in $\|\cdot\|_\infty$. If (F.26) failed, compactness would provide a subsequence $h_n\to h\in\mathcal{H}$ along which the error stays bounded away from zero. Hadamard directional differentiability applied with $t_n=n^{-1/2}$ and directions $\mathcal{U}h_n\to\mathcal{U}h$, together with continuity of the derivative, gives a contradiction. □

At a singular game, ϕ_{val} may be nonlinear, but it is always single-valued. Thus value estimation can have a non-Gaussian tangent target without suffering the discontinuity of a strategy branch.

F.3.2 Continuous strategy kinks

Let $\mathcal{C}_1, \dots, \mathcal{C}_J$ be full-dimensional chambers of the regular-kernel validity refinement in theorem D.9, with branch limits z_j^0 and branch derivatives L_j.

Definition F.10 (Sector-separated local set)**.** A compact set $\mathcal{H} \subseteq \mathbb{R}^d$ is *sector-separated* relative to $\mathcal{C}_1, \dots, \mathcal{C}_J$ when every nonzero $h \in \mathcal{H}$ satisfies $\mathcal{U}h \in \bigcup_{j=1}^{J} \mathcal{C}_j$, and, for each j, the set

$$K_j \coloneqq \left\{ \frac{\mathcal{U}h}{\|\mathcal{U}h\|_\infty} : h \in \mathcal{H} \setminus \{0\},\ \mathcal{U}h \in \mathcal{C}_j \right\} \tag{F.27}$$

is a compact subset of $\mathcal{C}_j \cap \{H : \|H\|_\infty = 1\}$.

The compact-containment requirement is exactly an angular separation from the fan skeleton. It ensures that the chamberwise expansions of theorem D.9 hold uniformly over the local parameter set.

Proposition F.11 (Uniform selector chart at a continuous kink)**.** *Assume the hypotheses of theorem D.9. Suppose the chambers $\mathcal{C}_1, \dots, \mathcal{C}_J$ have the common branch limit*

$$z_1^0 = \cdots = z_J^0 = z_0 = \mathrm{sel}(A_0), \tag{F.28}$$

and let $\mathcal{H}$ be sector-separated relative to these chambers. Define

$$\psi_{\mathrm{sel},n}(h) \coloneqq \sqrt{n}\,[\mathrm{sel}(A_{n,h}) - z_0]\,, \tag{F.29}$$

$$\phi_{\mathrm{sel}}(h) \coloneqq \begin{cases} L_j[\mathcal{U}h], & \mathcal{U}h \in \mathcal{C}_j,\ h \neq 0, \\ 0, & h = 0. \end{cases} \tag{F.30}$$

Then ϕ_{sel} is well defined on $\mathcal{H}$ and

$$\sup_{h \in \mathcal{H}} \|\psi_{\mathrm{sel},n}(h) - \phi_{\mathrm{sel}}(h)\|_2 \longrightarrow 0. \tag{F.31}$$

Consequently, the local minimax problem converges to the Gaussian shift experiment with the piecewise-linear target (F.30)*, for bounded losses and for quadratic loss on $\mathcal{H}$.*

Proof. For nonzero h with $\mathcal{U}h \in \mathcal{C}_j$, write $t = \|\mathcal{U}h\|_\infty/\sqrt{n}$ and $H = \mathcal{U}h/\|\mathcal{U}h\|_\infty$. Then $H \in K_j$, and the uniform branch expansion in theorem D.9 gives

$$\mathrm{sel}(A_{n,h}) = z_0 + \frac{1}{\sqrt{n}} L_j[\mathcal{U}h] + O\left(\frac{\|\mathcal{U}h\|_\infty^2}{n}\right),$$

uniformly over $h \in \mathcal{H} \setminus \{0\}$. The selector equals the unique equilibrium in every full-dimensional chamber, and (F.28) handles $h = 0$. Compactness of $\mathcal{H}$ yields (F.31). The map ϕ_{sel} is continuous on $\mathcal{H}$: it is linear on each sector, distinct sector pieces cannot accumulate at a nonzero point because of sector separation, and every piece converges to 0 as $h \to 0$. The minimax claims follow from theorem F.7 and corollary F.8. □

The sector separation is substantive. On a nonzero fan face, the first-order validity signs do not determine the selected branch, and a higher-order chart would be required. The reduction itself is not the end of the decision analysis: proposition F.12 solves a nontrivial continuous-kink example exactly, showing that the plug-in rule can also be the global minimax Gaussian decision for a one-sided piecewise-linear target. At a jump bifurcation, the common-limit assumption (F.28) fails and root-n point-strategy localization is impossible; that regime is treated by testing in section G.

F.3.3 An exactly solved continuous-kink Gaussian problem

On the one-dimensional slice $v = 0$ of (D.37), the value target is

$$\phi(u) = u_+ \coloneqq \max\{u, 0\}.$$

Consider the Gaussian shift $Z \sim \mathcal{N}(u, 1)$ under squared error.

Proposition F.12 (Exact minimax rule for the positive-part kink)**.** *The estimator*

$$\delta^\star(Z) = Z_+ \tag{F.32}$$

has risk

$$R(u) = \begin{cases} (u^2+1)\Phi(u) + u\varphi(u), & u < 0, \\ \Phi(u) - u\varphi(u) + u^2\Phi(-u), & u \geq 0, \end{cases} \tag{F.33}$$

where φ and Φ are the standard normal density and distribution function. Moreover,

$$\boxed{\inf_{\delta} \sup_{u\in\mathbb{R}} \mathbb{E}_u\big(\delta(Z) - u_+\big)^2 = 1,} \tag{F.34}$$

and $\delta^\star$ attains the minimax value.

Proof. Write $Z = u + G$, with $G \sim \mathcal{N}(0,1)$. If $u < 0$, then the target is zero and

$$R(u) = \mathbb{E}[(u+G)^2 \mathbf{1}\{G > -u\}].$$

Using

$$\int_a^\infty g\varphi(g)\,dg = \varphi(a), \qquad \int_a^\infty g^2\varphi(g)\,dg = a\varphi(a) + \Phi(-a),$$

with $a = -u$ gives the first line of (F.33). If $u \geq 0$, then the estimation error equals G on $\{G > -u\}$ and $-u$ otherwise, giving the second line.

For $u \geq 0$,

$$R'(u) = 2u\Phi(-u) \geq 0,$$

so R increases from $R(0) = 1/2$ to one. For $u < 0$,

$$R'(u) = 2\{\varphi(u) + u\Phi(u)\} = 2\int_{-\infty}^u (u-t)\varphi(t)\,dt > 0,$$

so R increases from zero to $1/2$. Therefore $\sup_u R(u) = 1$, proving the upper bound in (F.34).

For the lower bound, fix $R > 0$ and restrict the parameter to $u \in [R, 3R]$. On this interval $u_+ = u$. Translating the experiment and the decision by $2R$ shows that its minimax risk equals the minimax risk for estimating a Gaussian mean over $[-R, R]$. By lemma F.15, these compact-interval risks converge to one as $R \to \infty$. The full-line positive-part problem contains every such restricted problem, hence its minimax risk is at least one. This matches the risk of Z_+. □

F.4 Plug-in limits versus minimax limits

The preceding results characterize the best attainable local decision risk, not merely the distribution of one plug-in estimator. For comparison, the following proposition records the directional delta-method limit.

Proposition F.13 (Directional plug-in limit)**.** *Suppose $\psi_n(h)$ is a Euclidean local target chart satisfying $\psi_n \to \phi$ locally uniformly on $\mathbb{R}^d$, where ϕ is continuous. Let $\widehat{h}_n = Z_n$ and use the local plug-in decision $\widehat{\psi}_n = \psi_n(\widehat{h}_n)$. For every fixed h,*

$$\widehat{\psi}_n - \psi_n(h) \implies \phi(h + \Sigma^{1/2}G) - \phi(h). \tag{F.35}$$

If ϕ is nonlinear, the right-hand side is generally non-Gaussian and the plug-in decision need not solve the Gaussian minimax problem in (F.20)*.*

Proof. Under $A_{n,h}$, $Z_n = h + \Sigma_n^{1/2}G_n$ with $G_n \sim \mathcal{N}(0, \mathbf{I}_d)$, and hence $Z_n \Rightarrow h + \Sigma^{1/2}G$. Tightness permits restriction to an arbitrarily large compact set with arbitrarily high probability. Local uniform convergence of ψ_n to ϕ, followed by the continuous mapping theorem, gives $\psi_n(Z_n) \Rightarrow \phi(h + \Sigma^{1/2}G)$; subtracting $\psi_n(h) \to \phi(h)$ proves (F.35). □

This is the familiar directional delta-method phenomenon studied in general form by Fang & Santos (2019). Our minimax transfer theorem is stronger in a different direction: it identifies the limiting decision problem over all procedures, not only the weak limit of the natural plug-in rule.

F.4.1 Linear Hadamard derivatives are Fréchet derivatives in finite dimensions

Lemma F.14 (Hadamard-to-Fréchet upgrade)**.** *Let E, F be finite-dimensional normed spaces, let $U \subseteq E$ be a neighborhood of x_0, and let $f : U \to F$ be locally Lipschitz near x_0. Suppose that f is Hadamard directionally differentiable at x_0 and that its directional derivative is a linear map $L : E \to F$. Then f is Fréchet differentiable at x_0 with derivative L:*

$$\lim_{h \to 0} \frac{\|f(x_0 + h) - f(x_0) - Lh\|}{\|h\|} = 0. \tag{F.36}$$

Proof. If the conclusion failed, there would be $\varepsilon > 0$ and $h_n \to 0$ such that

$$\|f(x_0 + h_n) - f(x_0) - Lh_n\| \geq \varepsilon \|h_n\|.$$

Set $t_n = \|h_n\|$ and $u_n = h_n / t_n$. Compactness of the unit sphere gives a subsequence, not relabeled, with $u_n \to u$. Hadamard directional differentiability gives

$$\frac{f(x_0 + t_n u_n) - f(x_0)}{t_n} \to Lu.$$

Linearity and continuity of L give $Lu_n \to Lu$, so

$$\frac{f(x_0 + h_n) - f(x_0) - Lh_n}{\|h_n\|} \to 0,$$

a contradiction. □

F.5 Regular targets and sharp information bounds

We next recover the sharp regular local asymptotic minimax constant as a special case. The proof uses a finite-dimensional Gaussian location lemma and a multivariate van Trees lower bound. The latter is standard; see Gill & Levit (1995) for the form used below.

Lemma F.15 (Gaussian linear minimax risk)**.** *Let*

$$Z = h + \Sigma^{1/2} G, \qquad h \in \mathbb{R}^d, \qquad G \sim \mathcal{N}(0, \mathbf{I}_d),$$

where $\Sigma \succ 0$. Let $D \in \mathbb{R}^{r \times d}$ and $W \succ 0$. Then

$$\inf_{\delta} \sup_{h \in \mathbb{R}^d} \mathbb{E}_h \|\delta(Z) - Dh\|_W^2 = \operatorname{tr}(WD\Sigma D^{\mathsf{T}}), \tag{F.37}$$

and $\delta^\star(Z) = DZ$ attains the bound. Moreover, if

$$r_R \coloneqq \inf_{\delta} \sup_{\|h\|_2 \leq R} \mathbb{E}_h \|\delta(Z) - Dh\|_W^2,$$

then

$$\lim_{R \to \infty} r_R = \operatorname{tr}(WD\Sigma D^{\mathsf{T}}). \tag{F.38}$$

Proof. The estimator DZ has constant risk $\operatorname{tr}(WD\Sigma D^{\mathsf{T}})$, giving the upper bounds.

For the lower bound on expanding balls, choose a continuously differentiable density π supported in the open unit ball, vanishing at its boundary, and having finite Fisher-information matrix

$$I_\pi = \int \nabla \log \pi(u) \nabla \log \pi(u)^{\mathsf{T}} \pi(u)\, du.$$

Let $\pi_R(h) = R^{-d} \pi(h/R)$, which is supported in the ball of radius R and has prior information $R^{-2} I_\pi$. The multivariate van Trees inequality applied to the Gaussian location model and the linear estimand Dh yields, for every decision rule,

$$\int \mathbb{E}_h \|\delta(Z) - Dh\|_W^2 \, \pi_R(h)\, dh \geq \operatorname{tr}\left[WD \left(\Sigma^{-1} + R^{-2} I_\pi \right)^{-1} D^{\mathsf{T}} \right].$$

The supremum risk on the radius-R ball dominates this Bayes risk. Letting $R \to \infty$ gives the right-hand side of (F.38). Since r_R is nondecreasing and is bounded above by the constant risk of DZ, equality follows. The full-space identity (F.37) is immediate. □

Let Ψ be a Euclidean target defined near A_0 on $A_0 + \mathbb{T}$. Suppose it is Fréchet differentiable along $\mathbb{T}$: for some $D \in \mathbb{R}^{r\times d}$,

$$\sqrt{n}\,[\Psi(A_{n,h}) - \Psi(A_0)] = Dh + r_n(h), \qquad \sup_{\|h\|_2\le R} \|r_n(h)\|_2 \to 0 \quad \text{for every } R < \infty. \tag{F.39}$$

Define

$$\mathfrak{R}^{(2)}_{n,R}(\Psi, W) \coloneqq \inf_{\delta_n} \sup_{\|h\|_2\le R} \mathbb{E}_{n,h} \left\|\delta_n - \sqrt{n}\,[\Psi(A_{n,h}) - \Psi(A_0)]\right\|^2_W. \tag{F.40}$$

Theorem F.16 (Sharp regular local asymptotic minimax constant)**.** *Under Assumption F.1 and* (F.39)*, for every $W \succ 0$,*

$$\boxed{\lim_{R\to\infty} \lim_{n\to\infty} \mathfrak{R}^{(2)}_{n,R}(\Psi, W) = \operatorname{tr}(WD\Sigma D^{\mathsf{T}}).} \tag{F.41}$$

The tangent decision $\delta_n = DZ_n$ has worst-case risk converging to the right-hand side uniformly on every fixed radius-R local set and therefore attains the double-limit bound in (F.41)*.*

Proof. For each fixed R, (F.39) and corollary F.8 reduce the finite-sample risk to the Gaussian problem

$$\inf_{\delta} \sup_{\|h\|_2\le R} \mathbb{E}\|\delta(h + \Sigma^{1/2}G) - Dh\|^2_W.$$

Letting $R \to \infty$ and applying lemma F.15 proves (F.41). For $\delta_n = DZ_n$,

$$\delta_n - \sqrt{n}[\Psi(A_{n,h}) - \Psi(A_0)] = D(Z_n - h) - r_n(h).$$

The covariance term converges to $\operatorname{tr}(WD\Sigma D^{\mathsf{T}})$, the remainder is uniformly negligible on each fixed ball, and the cross term vanishes by Cauchy–Schwarz. □

The order of limits in (F.41) is the classical local asymptotic minimax order: first take the sampling limit on a fixed local neighborhood, then expand that neighborhood. This prevents boundary information from artificially lowering the minimax risk (Hájek, 1972; Le Cam, 1986).

Corollary F.17 (Regular value efficiency)**.** *Suppose A_0 has a unique equilibrium $z_0 = (x_0, y_0)$. The value is locally Lipschitz and its Hadamard derivative is linear; therefore lemma F.14 upgrades it to a relative Fr'echet derivative. Define $g_{\mathrm{val}} \in \mathbb{R}^d$ by*

$$(g_{\mathrm{val}})_\ell \coloneqq x_0^{\mathsf{T}} H^\ell y_0, \qquad \ell = 1, \dots, d. \tag{F.42}$$

Then the sharp regular local asymptotic minimax constant for squared value error is

$$\boxed{g_{\mathrm{val}}^{\mathsf{T}} \Sigma g_{\mathrm{val}}.} \tag{F.43}$$

Proof. Uniqueness gives $\mathcal{X}^\star(A_0) = \{x_0\}$ and $\mathcal{Y}^\star(A_0) = \{y_0\}$, so (F.25) is the linear functional $h \mapsto g_{\mathrm{val}}^{\mathsf{T}} h$. Apply theorem F.16 with $r = 1$, $D = g_{\mathrm{val}}^{\mathsf{T}}$, and $W = 1$. □

Corollary F.18 (Regular strategy efficiency)**.** *Suppose A_0 lies on a strictly valid regular kernel branch β, so the equilibrium is locally unique and differentiable. Let $D_{\mathrm{sel}} \in \mathbb{R}^{(m+k)\times d}$ have columns*

$$(D_{\mathrm{sel}})_{:\ell} \coloneqq Dz_\beta(A_0)[H^\ell]. \tag{F.44}$$

For every $W \succ 0$, the sharp regular local asymptotic minimax constant for the scaled point-strategy target is

$$\boxed{\operatorname{tr}\left(WD_{\mathrm{sel}}\Sigma D_{\mathrm{sel}}^{\mathsf{T}}\right).} \tag{F.45}$$

Proof. The analytic branch expansion in proposition D.4 gives (F.39) with $D = D_{\mathrm{sel}}$. Apply theorem F.16. □

F.6 What the reduction does and does not cover

The preceding results establish three distinct local experiments.

1. At a strictly valid regular kernel branch, the target chart is linear and the exact asymptotic risk is an information-covariance trace.
2. At a continuous kink, the chart is piecewise linear on the tangent fan. The local minimax problem remains a finite-dimensional Gaussian decision problem, but neither the optimal decision nor its distribution need be linear or Gaussian.
3. At a jump bifurcation, root-n point-strategy centering fails because the incident branch limits are separated by a constant. Uniform target convergence of the form (F.18) cannot hold on a local set crossing the bifurcation. The correct limit problem is branch testing with an unscaled structural target, developed in section G.

Thus the Gaussian observation law alone does not determine the statistical difficulty. The target tangent map—linear, piecewise linear, or discontinuous—determines the local minimax universality class.

G Composite Interface-Side Testing and Least-Favorable-Pair Risks

This appendix treats the discontinuous regime excluded from the root-n strategy chart in section F. The parameter class used to compare experiments is fixed independently of the design: it consists of two half-spaces separated by the linearized equilibrium interface. The exact composite target is the interface-side label. Within the local two-branch neighborhood this side label equals the equilibrium branch label; outside that neighborhood no global branch interpretation is asserted. The least-favorable pair is induced by the information matrix and is used only in the lower-bound proof and the two-point point-strategy calculation. The Gaussian information matrix may be singular.

G.1 Simple transverse jump charts

Retain the relative local payoff model

$$A_{n,h} = A_0 + \frac{\mathcal{U}h}{\sqrt{n}}, \qquad h \in \mathbb{R}^d. \tag{G.1}$$

Assumption G.1 (Simple transverse jump bifurcation)**.** There exist a relatively open neighborhood $\mathcal{N}$ of A_0 in $A_0 + \mathbb{T}$, a C^2 function $g : \mathcal{N} \to \mathbb{R}$, and continuous branch maps z_+ and z_- on the closures of the two sides such that:

1. $g(A_0) = 0$ and $Dg(A_0)|_{\mathbb{T}} \neq 0$;
2. $\mathcal{E}(A) = \{z_+(A)\}$ for $g(A) > 0$ and $\mathcal{E}(A) = \{z_-(A)\}$ for $g(A) < 0$;
3. the one-sided limits $z_+^0, z_-^0 \in \mathcal{E}(A_0)$ satisfy

$$\Delta_0 := d_\Delta(z_+^0, z_-^0) > 0. \tag{G.2}$$

In the local basis, define the boundary covector

$$a_\ell = Dg(A_0)[H^\ell], \qquad a = (a_1, \dots, a_d)^\mathsf{T} \neq 0. \tag{G.3}$$

The bridge from a single active regular-kernel interface to this assumption is proposition D.16. Replacing g by cg, $c \neq 0$, does not change the transition parameter or the local-oracle branch allocation; see lemma D.17.

G.2 Semidefinite information and fixed half-space classes

For each n, consider the exact Gaussian score experiment

$$\Delta_n \sim \mathcal{N}(J_n h, J_n), \qquad J_n \succeq 0, \tag{G.4}$$

and, for a margin $\lambda_n > 0$, the design-independent classes

$$\mathcal{H}_{n,+}(\lambda_n) = \{h : a^\mathsf{T} h \geq \lambda_n\}, \qquad \mathcal{H}_{n,-}(\lambda_n) = \{h : a^\mathsf{T} h \leq -\lambda_n\}. \tag{G.5}$$

Define the linearized interface-side label by

$$b_a(h) = \begin{cases} +, & h \in \mathcal{H}_{n,+}(\lambda_n), \\ -, & h \in \mathcal{H}_{n,-}(\lambda_n). \end{cases} \tag{G.6}$$

Theorem G.2 (Semidefinite Gaussian interface-side minimax theorem)**.** *For the experiment* (G.4)*, let*

$$R^\star_{\mathrm{side}}(J; a, \lambda) = \inf_{\widehat{b}} \max \left\{ \sup_{h \in \mathcal{H}_+(\lambda)} P_h(\widehat{b} = -), \sup_{h \in \mathcal{H}_-(\lambda)} P_h(\widehat{b} = +) \right\}.$$

Then:

1. *If $a \notin \mathrm{range}(J)$, then*

$$R^\star_{\mathrm{side}}(J; a, \lambda) = \frac{1}{2}. \tag{G.7}$$

2. *If $a \in \mathrm{range}(J)$, put*

$$s_a^2 \coloneqq a^\mathsf{T} J^\dagger a > 0. \tag{G.8}$$

Then

$$\boxed{R^\star_{\mathrm{side}}(J; a, \lambda) = \Phi\left(-\frac{\lambda}{s_a}\right).} \tag{G.9}$$

An optimal test is

$$\widehat{b}(\Delta) = \begin{cases} +, & a^\mathsf{T} J^\dagger \Delta \geq 0, \\ -, & a^\mathsf{T} J^\dagger \Delta < 0, \end{cases} \tag{G.10}$$

and a least-favorable pair is

$$h^\star_\pm = \pm\lambda \frac{J^\dagger a}{a^\mathsf{T} J^\dagger a}. \tag{G.11}$$

Proof. Because J is symmetric, $\mathrm{range}(J) = \ker(J)^\perp$.

Assume first that $a \notin \mathrm{range}(J)$. Then there exists $v \in \ker(J)$ with $a^\mathsf{T} v \neq 0$. Replacing v by $-v$ if necessary, assume $a^\mathsf{T} v > 0$, and set

$$h_+ = \frac{\lambda}{a^\mathsf{T} v} v, \qquad h_- = -h_+.$$

Then $h_+ \in \mathcal{H}_+(\lambda)$, $h_- \in \mathcal{H}_-(\lambda)$, and $Jh_+ = Jh_- = 0$. Thus the two observations have exactly the same law $\mathcal{N}(0, J)$. Every test has maximal error at least $1/2$, while an independent fair coin attains $1/2$, proving (G.7).

Now assume $a \in \mathrm{range}(J)$. Since $J^\dagger$ is positive definite on $\mathrm{range}(J)$ and $a \neq 0$, $s_a^2 > 0$. Define

$$T = \frac{a^\mathsf{T} J^\dagger \Delta}{s_a}.$$

Because $J^\dagger J$ is the orthogonal projector onto $\mathrm{range}(J)$ and a belongs to this range,

$$\mathbb{E}_h[T] = \frac{a^\mathsf{T} J^\dagger J h}{s_a} = \frac{a^\mathsf{T} h}{s_a},$$
$$\mathrm{Var}_h(T) = \frac{a^\mathsf{T} J^\dagger J J^\dagger a}{s_a^2} = 1.$$

Hence $T \sim \mathcal{N}(a^\mathsf{T} h / s_a, 1)$. The sign test (G.10) therefore has error at most $\Phi(-\lambda/s_a)$ on each half-space.

For the lower bound, the points in (G.11) satisfy $a^\mathsf{T} h^\star_\pm = \pm\lambda$. Their observation means are

$$Jh^\star_\pm = \pm\frac{\lambda}{s_a^2} J J^\dagger a = \pm\frac{\lambda}{s_a^2} a.$$

The equal-covariance singular-Gaussian likelihood ratio on $\mathrm{range}(J)$ is

$$\log \frac{dP_+}{dP_-}(\Delta) = \frac{2\lambda}{s_a^2} a^\mathsf{T} J^\dagger \Delta.$$

Thus the Neyman–Pearson test is exactly (G.10), and its two errors are both $\Phi(-\lambda/s_a)$. Restricting the composite problem to this pair proves the corresponding lower bound. □

The half-space classes in (G.5) are held fixed when J_n or a sampling design is varied. The pair (G.11) is induced by the design; it is not the parameter class used to compare designs.

G.3 Least-favorable normal-pair equilibrium risks

Assume henceforth that $a \in \mathrm{range}(J_n)$, $\lambda_n = o(\sqrt{n})$, and the normalized least-favorable directions lie in compact angular subsets of the two strict local side chambers. Put

$$v_n = \frac{J_n^\dagger a}{a^\mathsf{T} J_n^\dagger a}, \qquad h_{n,\pm}^\star = \pm \lambda_n v_n. \tag{G.12}$$

Then $a^\mathsf{T} v_n = 1$. Since $v_n/\|v_n\|$ remains in a compact angular subset of a strict side chamber, continuity of $u \mapsto a^\mathsf{T} u$ and separation from $a^\perp$ imply

$$\inf_n \left| a^\mathsf{T} \frac{v_n}{\|v_n\|} \right| > 0, \qquad \sup_n \|v_n\| < \infty.$$

Consequently $\|h_{n,\pm}^\star\| = O(\lambda_n)$. Define

$$A_{n,\pm} = A_0 + \frac{\mathcal{U} h_{n,\pm}^\star}{\sqrt{n}}, \qquad z_{n,\pm} = z_\pm(A_{n,\pm}), \qquad \Delta_{z,n} = d_\Delta(z_{n,+}, z_{n,-}). \tag{G.13}$$

Because g is C^2 on the relative model and $Dg(A_0)[\mathcal{U}h] = a^\mathsf{T} h$,

$$g(A_{n,\pm}) = \frac{a^\mathsf{T} h_{n,\pm}^\star}{\sqrt{n}} + O\left(\frac{\|h_{n,\pm}^\star\|^2}{n}\right),$$

$$\sqrt{n}\, g(A_{n,\pm}) = \pm\lambda_n + O\left(\frac{\lambda_n^2}{\sqrt{n}}\right). \tag{G.14}$$

Thus the points lie on opposite local equilibrium branches for all sufficiently large n, and $\Delta_{z,n} \to \Delta_0$.

Theorem G.3 (Exact equilibrium-branch and point-strategy risks on the least-favorable pair)**.** *On the declared pair $\{h_{n,+}^\star, h_{n,-}^\star\}$, which lies on the two local equilibrium branches, the oracle decision problem is given the two point targets $z_{n,+}$ and $z_{n,-}$. The exact local equilibrium-branch-germ error is $\Phi(-\kappa_n)$ and the exact squared point-strategy risk is*

$$\boxed{R_{\mathrm{sel},n}^\star = \frac{\Delta_{z,n}^2}{4} e^{-\kappa_n^2/2} \mathbb{E}_{G\sim\mathcal{N}(0,1)}[\cosh(\kappa_n G)^{-1}].} \tag{G.15}$$

The attaining point estimator is the componentwise posterior interpolation

$$\widehat{z}_n = \big(\pi_n x_{n,+} + (1-\pi_n) x_{n,-}, \pi_n y_{n,+} + (1-\pi_n) y_{n,-}\big), \quad \pi_n = (1 + e^{-2\kappa_n T_n})^{-1}, \tag{G.16}$$

where $T_n = a^\mathsf{T} J_n^\dagger \Delta_n / s_{a,n}$. It satisfies

$$d_\Delta(\widehat{z}_n, z_{n,+}) = (1-\pi_n)\Delta_{z,n}, \qquad d_\Delta(\widehat{z}_n, z_{n,-}) = \pi_n \Delta_{z,n}. \tag{G.17}$$

Every larger parameter class containing the least-favorable pair inherits the two displayed risks as minimax lower bounds.

Proof. The equilibrium-branch claim on the declared least-favorable pair follows from theorem G.2, because the two points lie in the local strict branch chambers. Under the pair, $T_n \sim \mathcal{N}(\pm\kappa_n, 1)$ and the likelihood ratio is $e^{2\kappa_n T_n}$. Conditional on T_n, the posterior positive probability is π_n. If $d_\pm = d_\Delta(z, z_{n,\pm})$, then $d_+ + d_- \geq \Delta_{z,n}$, and minimizing $\pi_n d_+^2 + (1-\pi_n) d_-^2$ gives the lower bound $\pi_n(1-\pi_n)\Delta_{z,n}^2$. Convexity of the product simplex and the norm-induced metric give the exact identities (G.17), so (G.16) attains the bound pointwise. Reflection symmetry makes it equalizer and hence minimax. Integrating against the equal mixture density $\varphi(t) e^{-\kappa_n^2/2} \cosh(\kappa_n t)$ proves (G.15). Restricting a larger minimax problem to the pair proves the final lower-bound statement. □

Corollary G.4 (Information-normal transition regimes)**.** *If $\Delta_{z,n} \to \Delta_0 > 0$, then:*

1. *$\kappa_n \to \infty$ implies both structural risks vanish;*
2. *$\kappa_n \to \kappa \in (0, \infty)$ gives the finite limits obtained by replacing κ_n by κ in the exact formulae;*
3. *$\kappa_n \to 0$ gives equilibrium-branch error $1/2$ on the least-favorable pair and point-strategy risk $\Delta_0^2/4$.*

Writing the signed physical boundary distance as $\delta_n = \lambda_n/\sqrt{n}$, the phase parameter is $\sqrt{n}\,\delta_n/s_{a,n}$. If $\lambda_n \to \infty$ while $\lambda_n = o(\sqrt{n})$, the sequence is moderate-local rather than contained in one fixed compact LAN neighborhood.

Corollary G.5 (No common root-n point-strategy chart)**.** *For every center $c_n \in \mathcal{Z}$,*

$$\max_{\eta \in \{-1,+1\}} d_\Delta(c_n, z_{n,\eta}) \geq \frac{\Delta_{z,n}}{2}. \tag{G.18}$$

Thus no centering sequence makes both root-n point-strategy errors bounded.

Proof. Apply the triangle inequality to $z_{n,+}, c_n, z_{n,-}$. □

G.4 Parametric stability under an identifiable payoff design

The preceding structural theorem allows singular information. To estimate the entire payoff matrix by plug-in, assume additionally the identifiable fixed design of assumption F.1, so $Z_n \sim \mathcal{N}(h, \Sigma_n)$ with $\sup_n \|\Sigma_n\|_{\mathrm{op}} < \infty$. Define

$$\widehat{A}_n = A_0 + \frac{\mathcal{U} Z_n}{\sqrt{n}}, \qquad \widehat{v}_n = \mathrm{val}(\widehat{A}_n), \qquad \widehat{z}_n = \mathrm{sel}(\widehat{A}_n). \tag{G.19}$$

By lemma H.2, these are globally defined Borel procedures. Put

$$C_r = \sup_n \mathbb{E}\|\mathcal{U}\Sigma_n^{1/2} G\|_\infty^r < \infty, \qquad r \in \{1, 2\}. \tag{G.20}$$

Proposition G.6 (Uniform root-n value and gap bounds)**.** *For every local parameter h,*

$$\mathbb{E}_{n,h} |\widehat{v}_n - \mathrm{val}(A_{n,h})|^2 \leq C_2/n, \tag{G.21}$$

$$\mathbb{E}_{n,h}\, \mathrm{Gap}_{A_{n,h}}(\widehat{z}_n) \leq 2C_1/\sqrt{n}, \tag{G.22}$$

$$\mathbb{E}_{n,h}\, \mathrm{Gap}^2_{A_{n,h}}(\widehat{z}_n) \leq 4C_2/n. \tag{G.23}$$

Proof. Under $A_{n,h}$, $\widehat{A}_n - A_{n,h} \overset{d}{=} \mathcal{U}\Sigma_n^{1/2} G/\sqrt{n}$. Apply the global value and gap transfer bounds from proposition C.1 and take moments. □

Thus the exact structural transition and the stable scalar/decision bounds can be observed in one full-information local experiment. Sparse local-oracle interface-side designs need not identify the entire payoff matrix; their exact composite side risk is instead governed by the semidefinite theorem above.

H Honest Confidence Regions

This appendix turns the bifurcation risks from section G into a theory of honest uncertainty quantification. The central distinction is between pointwise coverage and *honesty*: a region must cover the requested equilibrium target uniformly over the stated parameter set, including parameters inside the statistical bifurcation layer. Classical confidence-set theory relates honest diameter to testing and to local moduli of continuity (Donoho, 1994; Low, 1997; Robins & van der Vaart, 2006); recent optimization work constructs confidence sets for random linear-program solution sets (Liu et al., 2023). Here the target is the equilibrium geometry itself. We prove an exact finite-sample two-point

diameter frontier at a simple jump, construct an exact Gaussian payoff region whose projection simultaneously covers every compatible equilibrium, and show that the same construction contracts at the parametric rate on regular branches but cannot do so through a jump.

For the projected-payoff construction and the regular-branch contraction result, we retain the identifiable local Gaussian shift from section F:

$$A_{n,h} = A_0 + \frac{\mathcal{U}h}{\sqrt{n}}, \qquad Z_n \sim \mathcal{N}(h, \Sigma_n), \qquad J_n = \Sigma_n^{-1} \to J \succ 0. \tag{H.1}$$

The exact two-point jump frontier in section H.2 instead uses the semidefinite score experiment of section G.2; it does not require invertibility in nuisance directions. For $u \in \mathbb{R}^d$, write

$$\|u\|_{J_n} \coloneqq (u^{\mathsf{T}} J_n u)^{1/2}. \tag{H.2}$$

The target strategy space $\mathcal{Z} = \Delta_m \times \Delta_k$ is equipped with d_Δ from section 2; its diameter is two. We distinguish three coverage notions throughout:

Point target	$P\{\Psi(h) \in \mathfrak{C}_n\} \geq 1 - \alpha.$
Selected equilibrium	$P\{\mathrm{sel}(A_{n,h}) \in \mathfrak{C}_n\} \geq 1 - \alpha.$
Entire equilibrium set	$P\{\mathcal{E}(A_{n,h}) \subseteq \mathfrak{C}_n\} \geq 1 - \alpha.$

Only the last notion covers every equilibrium simultaneously.

H.1 Honesty and the testing–diameter principle

Let $(\mathcal{Y}, d_{\mathcal{Y}})$ be a compact metric space, and let $\mathcal{K}(\mathcal{Y})$ denote its nonempty compact subsets, equipped with the Hausdorff Borel structure. A confidence region is a measurable map $\mathfrak{C}_n : \mathbb{R}^d \to \mathcal{K}(\mathcal{Y})$. For a target $\Psi_n : \Theta \to \mathcal{Y}$, where $\Theta \subseteq \mathbb{R}^d$, define

$$\mathrm{diam}_{\mathcal{Y}}(C) \coloneqq \sup_{y,y' \in C} d_{\mathcal{Y}}(y, y'). \tag{H.3}$$

Definition H.1 (Honest target confidence region)**.** For $\alpha \in (0,1)$, $\mathfrak{C}_n$ is $(1-\alpha)$-honest for Ψ_n over Θ if

$$\inf_{h \in \Theta} P_{n,h}\{\Psi_n(h) \in \mathfrak{C}_n(Z_n)\} \geq 1 - \alpha. \tag{H.4}$$

If $\Gamma_n : \Theta \rightrightarrows \mathcal{Y}$ is set-valued, we call $\mathfrak{C}_n$ simultaneously honest for Γ_n if

$$\inf_{h \in \Theta} P_{n,h}\{\Gamma_n(h) \subseteq \mathfrak{C}_n(Z_n)\} \geq 1 - \alpha. \tag{H.5}$$

Lemma H.2 (Borel measurability of the canonical equilibrium selector)**.** *The maps*

$$A \mapsto \mathcal{X}^\star(A), \qquad A \mapsto \mathcal{Y}^\star(A)$$

are nonempty compact-valued Borel correspondences. Their unique entropy maximizers are Borel measurable. Consequently,

$$\mathrm{sel} : \mathbb{R}^{m \times k} \to \mathcal{Z}$$

is Borel measurable.

Proof. The value map is globally 1-Lipschitz in the entrywise sup norm. The optimal strategy sets admit the closed descriptions

$$\mathcal{X}^\star(A) = \{x \in \Delta_m : A^{\mathsf{T}} x \geq \mathrm{val}(A)\mathbf{1}_k\},$$
$$\mathcal{Y}^\star(A) = \{y \in \Delta_k : Ay \leq \mathrm{val}(A)\mathbf{1}_m\}.$$

Hence their graphs are closed in the product of Euclidean spaces, and their values are nonempty compact sets. A Borel graph with compact sections in Polish spaces is a Borel correspondence by the Arsenin–Kunugui theorem; see Theorem 18.18 of Kechris (1995).

Entropy is continuous on each simplex. The measurable maximum theorem (Aliprantis–Border, Theorem 18.19) implies that the entropy-argmax correspondences are measurable and compact-valued (Aliprantis & Border, 2006). Each optimal-strategy set is convex and entropy is strictly concave, so each argmax is a singleton. A singleton-valued Borel correspondence is a Borel function, proving the claim. □

Lemma H.3 (Measurable reduction to the three two-point confidence reports)**.** *Let $(\mathcal{Y}, d)$ be compact metric and fix distinct $y_+, y_- \in \mathcal{Y}$. For $C \in \mathcal{K}(\mathcal{Y})$ define*

$$\mathfrak{R}(C) = \begin{cases} \{y_+, y_-\}, & y_+, y_- \in C, \\ \{y_+\}, & y_+ \in C,\ y_- \notin C, \\ \{y_-\}, & y_- \in C,\ y_+ \notin C, \\ \{y_+\}, & y_+, y_- \notin C. \end{cases}$$

Then $\mathfrak{R} : \mathcal{K}(\mathcal{Y}) \to \mathcal{K}(\mathcal{Y})$ is Borel, never decreases coverage of either point, and satisfies $\operatorname{diam} \mathfrak{R}(C) \le \operatorname{diam} C$.

Proof. Under the Hausdorff topology, $C \mapsto d(y, C)$ is continuous. Thus each event $\{C : y_\pm \in C\} = \{C : d(y_\pm, C) = 0\}$ is closed, and the four cases defining $\mathfrak{R}$ are Borel. If C contains one or both targets, the reduced set contains exactly those targets and is a subset of C; if it contains neither, coverage already fails at both points and the fixed singleton cannot reduce coverage. The diameter claim is immediate. □

Proposition H.4 (Measurability of the projected equilibrium confidence set)**.** *Let $z \mapsto \mathcal{Q}(z) \subset \mathbb{R}^d$ be a nonempty compact ellipsoid whose graph is closed and which varies continuously in the Hausdorff metric. Define*

$$\Gamma(z) = \bigcup_{u \in \mathcal{Q}(z)} \mathcal{E}\left(A_0 + \frac{\mathcal{U}u}{\sqrt{n}}\right).$$

Then $\Gamma(z)$ is nonempty compact for every z, its graph is closed, and $z \mapsto \Gamma(z)$ is a Borel random compact set.

Proof. The set

$$\mathcal{G} = \left\{(z, u, w) : u \in \mathcal{Q}(z),\ w \in \mathcal{Z}, \operatorname{Gap}_{A_0 + \mathcal{U}u/\sqrt{n}}(w) = 0\right\}$$

is closed because the ellipsoid graph and the gap map are closed/continuous. For fixed z, its (u, w)-section is compact. Its projection onto (z, w) is therefore closed: if $(z_j, w_j) \to (z, w)$, choose witnesses $u_j \in \mathcal{Q}(z_j)$; local boundedness of the ellipsoids gives a convergent subsequence $u_j \to u \in \mathcal{Q}(z)$, and continuity of the gap gives $w \in \Gamma(z)$. Hence Γ has closed graph and nonempty compact values. The Arsenin–Kunugui theorem then implies weak Borel measurability; for compact metric range this is equivalent to Borel measurability as a map into the Hausdorff hyperspace $\mathcal{K}(\mathcal{Z})$. □

The next lemma is purely decision-theoretic and does not use Gaussianity. It is the confidence-region analogue of the standard reduction from estimation to testing.

Lemma H.5 (Testing forces confidence diameter)**.** *Let P_0, P_1 be two probability laws, let $y_0, y_1 \in \mathcal{Y}$, and put $\Delta = d_{\mathcal{Y}}(y_0, y_1)$. Suppose a random compact set C satisfies*

$$P_i\{y_i \in C\} \ge 1 - \alpha_i, \qquad i \in \{0, 1\}. \tag{H.6}$$

Then, for each $i \in \{0, 1\}$,

$$P_i\{\operatorname{diam}_{\mathcal{Y}}(C) \ge \Delta\} \ge [1 - \alpha_0 - \alpha_1 - d_{\mathrm{TV}}(P_0, P_1)]_+, \tag{H.7}$$

$$\mathbb{E}_i[\operatorname{diam}_{\mathcal{Y}}(C)] \ge \Delta\,[1 - \alpha_0 - \alpha_1 - d_{\mathrm{TV}}(P_0, P_1)]_+. \tag{H.8}$$

Proof. Let $E_j = \{y_j \in C\}$. By the definition of total variation, $P_i(E_{1-i}) \ge P_{1-i}(E_{1-i}) - d_{\mathrm{TV}}(P_0, P_1)$. Hence

$$\begin{aligned} P_i(E_0 \cap E_1) &\ge 1 - P_i(E_i^c) - P_i(E_{1-i}^c) \\ &\ge 1 - \alpha_i - \alpha_{1-i} - d_{\mathrm{TV}}(P_0, P_1). \end{aligned}$$

On $E_0 \cap E_1$, the set C contains both target points and therefore has diameter at least Δ. This proves (H.7); multiplying the event indicator by Δ gives (H.8). □

For equal-covariance Gaussian shifts, if $\delta = \|h_1 - h_0\|_{J_n}$, whitening and projection onto the mean-difference direction give

$$d_{\mathrm{TV}}(\mathcal{N}(h_0, \Sigma_n), \mathcal{N}(h_1, \Sigma_n)) = 1 - 2\Phi(-\delta/2). \tag{H.9}$$

Thus nearby parameters with separated targets force wide honest regions.

H.2 THE EXACT TWO-POINT CONFIDENCE-DIAMETER FRONTIER AT A JUMP

Retain the semidefinite score experiment and least-favorable pair from sections G.2 and G.3. Thus $a \in \mathrm{range}(J_n)$,

$$s_{a,n}^2 = a^\mathsf{T} J_n^\dagger a, \qquad \kappa_n = \lambda_n / s_{a,n},$$

and under the two points $h_{n,\pm}^\star$ the scalar sufficient statistic

$$T_n = \frac{a^\mathsf{T} J_n^\dagger \Delta_n}{s_{a,n}}$$

has law $\mathcal{N}(\pm\kappa_n, 1)$. Let $z_{n,\pm}$ be the two oracle point-strategy targets on local branches and $\Delta_{z,n} = d_\Delta(z_{n,+}, z_{n,-})$. For $\alpha \in (0, 1/2)$, put

$$q_\alpha(\kappa) \coloneqq [\Phi(z_{1-\alpha} - 2\kappa) - \alpha]_+, \qquad z_{1-\alpha} = \Phi^{-1}(1-\alpha). \tag{H.10}$$

Let $\mathfrak{C}_{\mathrm{conf},n,\alpha}$ be the nonempty random compact regions satisfying

$$P_{n,\eta}\{z_{n,\eta} \in \mathfrak{C}_n\} \geq 1 - \alpha, \qquad \eta \in \{-1, +1\}.$$

Theorem H.6 (Exact two-point honest-diameter frontier at a simple jump)**.** *Suppose the local procedure is given the two oracle targets $z_{n,+}, z_{n,-}$. Then*

$$\boxed{\inf_{\mathfrak{C}_n \in \mathfrak{C}_{\mathrm{conf},n,\alpha}} \max_{\eta \in \{-1,+1\}} \mathbb{E}_{n,\eta} \operatorname{diam}_{d_\Delta}(\mathfrak{C}_n) = \Delta_{z,n} q_\alpha(\kappa_n).} \tag{H.11}$$

An attaining Borel rule is

$$\mathfrak{C}_n^\star(T_n) = \begin{cases} \{z_{n,+}\}, & T_n > c_{n,\alpha}, \\ \{z_{n,-}\}, & T_n < -c_{n,\alpha}, \\ \{z_{n,-}, z_{n,+}\}, & |T_n| \leq c_{n,\alpha}, \end{cases} \qquad c_{n,\alpha} = [z_{1-\alpha} - \kappa_n]_+. \tag{H.12}$$

Threshold ties are assigned to the two-target report. Every larger honest parameter class containing the two least-favorable points has minimax expected diameter at least the quantity in (H.11)*.*

Proof. By lemma H.3, any random compact region can be reduced, without decreasing either coverage indicator or increasing diameter, to the three Borel reports $\{z_{n,+}\}$, $\{z_{n,-}\}$, and $\{z_{n,+}, z_{n,-}\}$. Write $\delta_+(t), \delta_-(t), \delta_0(t)$ for the conditional probabilities of these reports. Then $\delta_+ + \delta_- + \delta_0 = 1$ and expected diameter is $\Delta_{z,n}\mathbb{E}_\eta \delta_0(T_n)$.

Given any feasible rule, average it with its reflected, label-swapped rule. Coverage remains feasible, and the larger of the two expected diameters cannot increase. Hence it suffices to consider symmetric rules satisfying $\delta_-(t) = \delta_+(-t)$ and $\delta_0(t) = \delta_0(-t)$. Their two expected diameters are equal, and the two coverage constraints reduce to

$$\int \delta_+(t) p_-(t)\, dt \leq \alpha, \qquad \delta_+(t) + \delta_+(-t) \leq 1 \quad \text{a.e.}$$

The common undecided probability is

$$1 - \int \delta_+(t)\{p_+(t) + p_-(t)\}\, dt.$$

Thus we maximize the last integral under the displayed constraints. The benefit-to-cost ratio is

$$\frac{p_+(t) + p_-(t)}{p_-(t)} = 1 + e^{2\kappa_n t},$$

which is strictly increasing. The Neyman–Pearson rearrangement therefore places the positive singleton on an upper tail and, by symmetry, the negative singleton on its reflected lower tail. This argument already allows randomization: the feasible set is convex and the objective is linear, while the monotone likelihood ratio makes an extreme threshold rule optimal.

If $c_{n,\alpha} = z_{1-\alpha} - \kappa_n > 0$, the negative-law error of the upper-tail rule $\{T_n > c_{n,\alpha}\}$ is exactly α, the two singleton tails are disjoint, and under the positive law

$$P_+\{|T_n| \leq c_{n,\alpha}\} = \Phi(z_{1-\alpha} - 2\kappa_n) - \alpha.$$

If $c_{n,\alpha} \leq 0$, the sign rule uses the disjoint half-lines separated at zero and has error $\Phi(-\kappa_n) \leq \alpha$, so the optimum undecided probability is zero. Combining the cases proves the positive-part formula. Restricting any larger-class honest problem to the pair gives the final lower bound. □

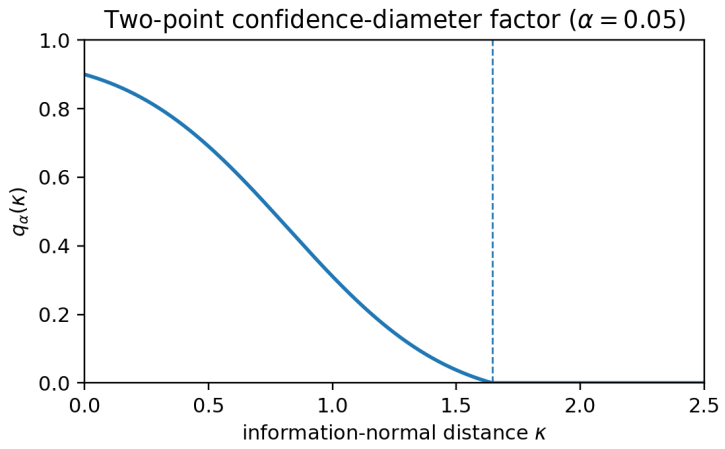

Figure 3: The exact two-point diameter factor $q_\alpha(\kappa) = [\Phi(z_{1-\alpha} - 2\kappa) - \alpha]_+$ for $\alpha = 0.05$. Once $\kappa \geq z_{1-\alpha}$, the sign rule satisfies both coverage constraints and the optimum diameter is zero.

Corollary H.7 (Three confidence-diameter regimes)**.** *If $\Delta_{z,n} \to \Delta_0 > 0$, the optimum in* (H.11) *tends to zero when $\kappa_n \to \infty$, to $\Delta_0 q_\alpha(\kappa)$ when $\kappa_n \to \kappa \in (0, \infty)$, and to $\Delta_0(1-2\alpha)$ when $\kappa_n \to 0$.*

H.3 Exact payoff ellipsoids and equilibrium projection

The lower bounds above apply to every honest procedure. We now construct a single confidence region that is exactly honest for the payoff parameter and, by projection, simultaneously covers all equilibria compatible with that uncertainty. This is a statistical existence construction: evaluating the union of equilibrium sets over an ellipsoid can be computationally difficult, and no polynomial-time projection algorithm is claimed. Let

$$\left(r^{\mathrm{conf}}_{\alpha,d}\right)^2 := F^{-1}_{\chi^2_d}(1-\alpha) \tag{H.13}$$

be the $(1-\alpha)$ quantile of a chi-square random variable with d degrees of freedom, and define the local-parameter ellipsoid

$$\mathbb{B}_{n,\alpha}(z) := \left\{u \in \mathbb{R}^d : \|u - z\|_{J_n} \leq r^{\mathrm{conf}}_{\alpha,d}\right\}. \tag{H.14}$$

Because $\|Z_n - h\|^2_{J_n} \sim \chi^2_d$,

$$P_{n,h}\{h \in \mathbb{B}_{n,\alpha}(Z_n)\} = 1 - \alpha \qquad \text{for every } h \in \mathbb{R}^d. \tag{H.15}$$

Define the projected value interval and equilibrium region by

$$\mathfrak{C}^{\mathrm{val}}_{n,\alpha}(z) := \{\mathrm{val}(A_{n,u}) : u \in \mathbb{B}_{n,\alpha}(z)\}, \tag{H.16}$$

$$\mathfrak{C}^{\mathrm{eq}}_{n,\alpha}(z) := \bigcup_{u \in \mathbb{B}_{n,\alpha}(z)} \mathcal{E}(A_{n,u}). \tag{H.17}$$

Although the local Gaussian family is algebraically extended beyond the bounded payoff cube in section F, both $\mathrm{val}(A)$ and $\mathcal{E}(A)$ are well-defined for every real payoff matrix, so these projections require no truncation.

Proposition H.8 (Exact projected confidence regions)**.** *Under Assumption F.1, the following hold.*

1. *$\mathfrak{C}^{\mathrm{val}}_{n,\alpha}(z)$ is a compact interval, and $\mathfrak{C}^{\mathrm{eq}}_{n,\alpha}(z)$ is a nonempty compact subset of $\mathcal{Z}$. The map $z \mapsto \mathfrak{C}^{\mathrm{eq}}_{n,\alpha}(z)$ is a measurable random compact set.*

2. *For every $h \in \mathbb{R}^d$,*

$$P_{n,h}\left\{\mathrm{val}(A_{n,h}) \in \mathfrak{C}^{\mathrm{val}}_{n,\alpha}(Z_n)\right\} \geq 1 - \alpha, \tag{H.18}$$

$$P_{n,h}\left\{\mathcal{E}(A_{n,h}) \subseteq \mathfrak{C}^{\mathrm{eq}}_{n,\alpha}(Z_n)\right\} \geq 1 - \alpha. \tag{H.19}$$

In particular, the equilibrium region is honest for every single-valued equilibrium selector.

3. Let

$$\Lambda_n \coloneqq \sup_{\|u\|_{J_n} \le 1} \|\mathcal{U}u\|_\infty. \tag{H.20}$$

Then, deterministically for every z,

$$\operatorname{diam}\left(\mathfrak{C}^{\mathrm{val}}_{n,\alpha}(z)\right) \le \frac{2r^{\mathrm{conf}}_{\alpha,d}\Lambda_n}{\sqrt{n}}. \tag{H.21}$$

Proof. The ellipsoid $\mathbb{B}_{n,\alpha}(z)$ is nonempty, compact, and connected. The value map is continuous, so its image is a compact connected subset of $\mathbb{R}$, hence an interval. The graph of the equilibrium correspondence is closed because

$$\operatorname{gph}(\mathcal{E}) = \{(A, w) : \operatorname{Gap}_A(w) = 0\} \tag{H.22}$$

with $(A, w) \mapsto \operatorname{Gap}_A(w)$ continuous. Therefore

$$\{(u, w) : u \in \mathbb{B}_{n,\alpha}(z),\ w \in \mathcal{E}(A_{n,u})\} \tag{H.23}$$

is compact, and its projection onto w is exactly $\mathfrak{C}^{\mathrm{eq}}_{n,\alpha}(z)$. The graph of the map $z \mapsto \mathfrak{C}^{\mathrm{eq}}_{n,\alpha}(z)$ is closed by the same subsequence argument. The precise hyperspace conclusion is supplied by proposition H.4: the compact-valued map is Borel measurable into $\mathcal{K}(\mathcal{Z})$ with the Hausdorff Borel structure.

On the event $h \in \mathbb{B}_{n,\alpha}(Z_n)$, the true value and the entire true equilibrium set are included in their corresponding projections. Equation (H.15) proves the coverage claims. Finally, if u, v belong to the same ellipsoid, then $\|u - v\|_{J_n} \le 2r^{\mathrm{conf}}_{\alpha,d}$ and

$$\begin{aligned} |\operatorname{val}(A_{n,u}) - \operatorname{val}(A_{n,v})| &\le \|A_{n,u} - A_{n,v}\|_\infty \\ &= \frac{\|\mathcal{U}(u - v)\|_\infty}{\sqrt{n}} \le \frac{2r^{\mathrm{conf}}_{\alpha,d}\Lambda_n}{\sqrt{n}}, \end{aligned}$$

using proposition C.1. □

H.4 Local oscillation and diameter adaptation

The correct size of a projected equilibrium region is controlled by a local oscillation of the equilibrium correspondence in the information metric. Define

$$\Omega^{\mathrm{eq}}_n(r; h) \coloneqq \sup\Big\{ d_\Delta(w, w') : \|u - h\|_{J_n} \le r,\ \|v - h\|_{J_n} \le r, \tag{H.24}$$

$$w \in \mathcal{E}(A_{n,u}),\ w' \in \mathcal{E}(A_{n,v}) \Big\}. \tag{H.25}$$

This is the equilibrium analogue of the local moduli used in minimax estimation and honest-confidence theory (Donoho, 1994; Low, 1997; Robins & van der Vaart, 2006).

Proposition H.9 (Modulus control of projected diameter)**.** *For every $h \in \mathbb{R}^d$,*

$$P_{n,h}\left\{\operatorname{diam}_{d_\Delta}\left(\mathfrak{C}^{\mathrm{eq}}_{n,\alpha}(Z_n)\right) \le \Omega^{\mathrm{eq}}_n(2r^{\mathrm{conf}}_{\alpha,d}; h)\right\} \ge 1 - \alpha, \tag{H.26}$$

while

$$\mathbb{E}_{n,h} \operatorname{diam}_{d_\Delta}\left(\mathfrak{C}^{\mathrm{eq}}_{n,\alpha}(Z_n)\right) \le \Omega^{\mathrm{eq}}_n(2r^{\mathrm{conf}}_{\alpha,d}; h) + 2\alpha. \tag{H.27}$$

Conversely, if a confidence region is honest over a parameter set $\Theta \subseteq \mathbb{R}^d$ for a single-valued target $\Psi_n : \Theta \to \mathcal{Y}$, then

$$\sup_{h \in \Theta} \mathbb{E}_{n,h} \operatorname{diam}_{\mathcal{Y}}(\mathfrak{C}_n) \ge \sup_{u,v \in \Theta} d_{\mathcal{Y}}(\Psi_n(u), \Psi_n(v)) \left[2\Phi\left(-\frac{\|u - v\|_{J_n}}{2} \right) - 2\alpha \right]_+. \tag{H.28}$$

Proof. On the event $h \in \mathbb{B}_{n,\alpha}(Z_n)$, every $u \in \mathbb{B}_{n,\alpha}(Z_n)$ satisfies

$$\|u - h\|_{J_n} \le \|u - Z_n\|_{J_n} + \|Z_n - h\|_{J_n} \le 2r^{\mathrm{conf}}_{\alpha,d}. \tag{H.29}$$

The definition of Ω^{eq}_n therefore gives (H.26). Since $\operatorname{diam}_{d_\Delta}(\mathcal{Z}) = 2$, splitting the expectation according to this event proves (H.27). For the converse, apply lemma H.5 to every pair $u, v \in \Theta$ and use the Gaussian total-variation formula (H.9); then take the supremum. □

The upper and lower bounds use different but operationally matched objects: the upper bound is the actual equilibrium oscillation inside a confidence ellipsoid, while the lower bound weights target separation by the exact Gaussian testing overlap. At a jump, target separation stays positive as information distance tends to zero, forcing order-one diameter. On a regular branch, the oscillation is linear and the region contracts at rate $n^{-1/2}$.

Theorem H.10 (Sharp first-order contraction on a strictly valid regular kernel branch)**.** *Suppose A_0 has a strictly valid regular kernel branch z_β and there is a neighborhood $\mathcal{N}$ of A_0 on which $\mathcal{E}(A) = \{z_\beta(A)\}$. Define the coordinate derivative*

$$D_\beta h := Dz_\beta(A_0)[\mathcal{U}h], \qquad h \in \mathbb{R}^d, \tag{H.30}$$

and, for a tangent vector $w = (p, q)$,

$$\|w\|_\Delta := \frac{1}{2}(\|p\|_1 + \|q\|_1). \tag{H.31}$$

Then, uniformly over every compact set of fixed local parameters $\mathcal{H} \subset \mathbb{R}^d$, the following convergence in probability holds:

$$\boxed{\sqrt{n}\,\operatorname{diam}_{d_\Delta}\left(\mathfrak{C}^{\mathrm{eq}}_{n,\alpha}(Z_n)\right) \xrightarrow{P_{n,h}} 2r^{\mathrm{conf}}_{\alpha,d}\left\|D_\beta J^{-1/2}\right\|_{2\to\Delta},} \tag{H.32}$$

where

$$\left\|D_\beta J^{-1/2}\right\|_{2\to\Delta} := \sup_{\|g\|_2\le 1} \|D_\beta J^{-1/2} g\|_\Delta, \tag{H.33}$$

If $D_\beta = 0$, the displayed limit constant is zero; the statement then asserts $o_P(n^{-1/2})$ diameter on the strictly valid regular kernel branch. The convergence is uniform in the sense that, for every $\varepsilon > 0$,

$$\sup_{h\in\mathcal{H}} P_{n,h}\left\{\left|\sqrt{n}\operatorname{diam}_{d_\Delta}(\mathfrak{C}^{\mathrm{eq}}_{n,\alpha}(Z_n)) - 2r^{\mathrm{conf}}_{\alpha,d}\|D_\beta J^{-1/2}\|_{2\to\Delta}\right| > \varepsilon\right\} \longrightarrow 0. \tag{H.34}$$

Proof. Uniformly over h in a fixed compact set, $Z_n = h + \Sigma_n^{1/2} G$ is tight. Hence, with probability tending to one uniformly in h, the entire ellipsoid $\mathbb{B}_{n,\alpha}(Z_n)$ lies in a fixed bounded subset of $\mathbb{R}^d$. On that event, all matrices $A_{n,u}$ in the ellipsoid belong to $\mathcal{N}$ for sufficiently large n, so the projected equilibrium region is the image of the single analytic branch z_β.

The uniform second-order expansion from proposition D.4 gives, on every bounded coordinate set,

$$z_\beta(A_{n,u}) = z_\beta(A_0) + \frac{D_\beta u}{\sqrt{n}} + O(n^{-1}). \tag{H.35}$$

Therefore the scaled diameter differs by $o_{P_{n,h}}(1)$, uniformly in $h \in \mathcal{H}$, from the diameter of the linear image

$$\left\{D_\beta u : \|u - Z_n\|_{J_n} \le r^{\mathrm{conf}}_{\alpha,d}\right\}. \tag{H.36}$$

Translation by $D_\beta Z_n$ does not affect diameter, and $u - Z_n = J_n^{-1/2} g$ with $\|g\|_2 \le r^{\mathrm{conf}}_{\alpha,d}$. Since the image is centrally symmetric, its diameter is exactly

$$2r^{\mathrm{conf}}_{\alpha,d}\|D_\beta J_n^{-1/2}\|_{2\to\Delta}. \tag{H.37}$$

Continuity of the induced norm and $J_n^{-1/2} \to J^{-1/2}$ prove (H.32)– (H.34). □

At a strictly valid regular kernel branch, the projected region therefore has a sharp root-n diameter constant. At a simple jump, by contrast, theorem H.6 shows that no honest region can contract uniformly through the information-normal transition layer. Meanwhile, the projected value interval always satisfies the deterministic root-n width bound (H.21). Hence the same payoff ellipsoid induces fundamentally different confidence geometries:

$$\boxed{\begin{array}{c}\text{root-}n\text{ value intervals and regular strategy regions,}\\ \text{but order-one honest strategy regions at a jump.}\end{array}} \tag{H.38}$$

I TARGET-DEPENDENT LOCAL-ORACLE EXPERIMENTAL DESIGN

This appendix characterizes the observation designs induced by the local minimax theory of section F. The principal point is that an "optimal design for the game" is not well-defined: the criterion depends on whether the inferential target is the value, a smooth equilibrium strategy, or the side of an equilibrium bifurcation. We first treat arbitrary structured local models, then derive closed forms for unrestricted entrywise perturbations, and finally exhibit a family in which local-oracle value and interface-side designs are exactly incompatible.

I.1 DESIGN GEOMETRY FOR A SMOOTH TARGET

Let

$$\mathcal{A} \coloneqq [m] \times [k]$$

be the set of queryable payoff entries. In the local coordinates of section F, let $u_a^\mathsf{T} \in \mathbb{R}^{1\times d}$ be the row of U corresponding to $a \in \mathcal{A}$. One observation of entry a contributes Fisher information

$$J_a \coloneqq \frac{u_a u_a^\mathsf{T}}{\sigma_a^2}. \tag{I.1}$$

For a design $q \in \Delta_{|\mathcal{A}|}$, define

$$J(q) \coloneqq \sum_{a\in\mathcal{A}} q_a J_a. \tag{I.2}$$

Let a regular target have derivative $D_\Psi \in \mathbb{R}^{r\times d}$, and let $W \in \mathbb{R}^{r\times r}$ be positive definite. Put

$$C \coloneqq D_\Psi^\mathsf{T} W D_\Psi \succeq 0.$$

The target-specific design criterion is

$$\Phi_{D_\Psi,W}(q) \coloneqq \begin{cases} \operatorname{tr}\big(W D_\Psi J(q)^\dagger D_\Psi^\mathsf{T}\big), & \operatorname{range}(D_\Psi^\mathsf{T}) \subseteq \operatorname{range}(J(q)), \\ +\infty, & \text{otherwise}, \end{cases} \tag{I.3}$$

where $J(q)^\dagger$ is the Moore–Penrose inverse. The range condition says exactly that every locally relevant target direction is identifiable, even if nuisance directions are not.

Proposition I.1 (Target-identifiable Gaussian risk)**.** *Consider the Gaussian score experiment*

$$\Delta \sim \mathcal{N}(J(q)h, J(q)), \qquad h \in \mathbb{R}^d, \tag{I.4}$$

and the linear target $D_\Psi h$. If $\operatorname{range}(D_\Psi^\mathsf{T}) \subseteq \operatorname{range}(J(q))$, then

$$\widehat{\psi}_q \coloneqq D_\Psi J(q)^\dagger \Delta \tag{I.5}$$

has constant weighted quadratic risk $\Phi_{D_\Psi,W}(q)$ and is minimax over $h \in \mathbb{R}^d$. If the range condition fails, the global minimax risk is infinite. Moreover, $q \mapsto \Phi_{D_\Psi,W}(q)$ is convex and lower semicontinuous on the design simplex; hence it attains its minimum whenever it is finite for at least one design.

Proof. Let $P_q = J(q)^\dagger J(q)$, the orthogonal projector onto $\operatorname{range}(J(q))$. The range condition is equivalent to $D_\Psi = D_\Psi P_q$. Hence (I.5) is unbiased and has covariance

$$D_\Psi J(q)^\dagger J(q) J(q)^\dagger D_\Psi^\mathsf{T} = D_\Psi J(q)^\dagger D_\Psi^\mathsf{T}.$$

The minimax lower bound follows by restricting h to $\operatorname{range}(J(q))$ and applying lemma F.15 on that subspace. If the range condition fails, there exists $v \in \ker J(q)$ with $D_\Psi v \neq 0$. The laws at h and $h + tv$ are identical for every t, while the targets separate without bound; the global minimax quadratic risk is therefore infinite.

For $b \in \mathbb{R}^d$ and $M \succeq 0$, define

$$\varphi_b(M) \coloneqq \begin{cases} b^\mathsf{T} M^\dagger b, & b \in \operatorname{range}(M), \\ +\infty, & \text{otherwise}. \end{cases}$$

Completing the square on $\mathrm{range}(M)$, and scaling a null-space direction otherwise, gives the variational representation

$$\varphi_b(M) = \sup_{z\in\mathbb{R}^d} \left\{2b^\mathsf{T} z - z^\mathsf{T} M z\right\}. \tag{I.6}$$

Writing the rows of $W^{1/2}D_\Psi$ as b_s^T yields $\Phi_{D_\Psi,W}(q) = \sum_s \varphi_{b_s}(J(q))$. Since $J(q)$ is affine in q, (I.6) makes each summand a supremum of continuous affine functions of q. Thus the criterion is convex and lower semicontinuous. Compactness of the simplex gives existence. □

The next result is a target-specific general equivalence theorem. It is in the classical optimal-design lineage of Kiefer & Wolfowitz (1960), but the criterion here is the equilibrium target risk induced by D_Ψ.

Theorem I.2 (Target-equivalence condition)**.** *Suppose $q^\star$ satisfies $J(q^\star) \succ 0$. Define the sensitivity of entry a by*

$$\chi_a(q) \coloneqq \frac{1}{\sigma_a^2} u_a^\mathsf{T} J(q)^{-1} C J(q)^{-1} u_a. \tag{I.7}$$

Then the following are equivalent:

1. *$q^\star$ minimizes $\Phi_{D_\Psi,W}$ over all designs with nonsingular information;*
2. $$\max_{a\in\mathcal{A}} \chi_a(q^\star) \le \Phi_{D_\Psi,W}(q^\star). \tag{I.8}$$

This equivalence condition is deliberately restricted to the positive definite information region. Singular boundary optima are covered by the convex lower-semicontinuous criterion in proposition I.1 and by the semidefinite closed forms below, not by differentiating an inverse at the boundary.

At every such optimum,

$$\sum_{a\in\mathcal{A}} q_a^\star \chi_a(q^\star) = \Phi_{D_\Psi,W}(q^\star), \qquad q_a^\star > 0 \Longrightarrow \chi_a(q^\star) = \Phi_{D_\Psi,W}(q^\star). \tag{I.9}$$

Proof. On the positive definite information region, $\Phi_{D_\Psi,W}(q) = \mathrm{tr}(CJ(q)^{-1})$ is convex and differentiable. Since $D(M^{-1})[H] = -M^{-1}HM^{-1}$,

$$D\Phi_{D_\Psi,W}(q)[e_a - q] = \Phi_{D_\Psi,W}(q) - \chi_a(q).$$

Optimality over the simplex is equivalent to nonnegative directional derivatives toward every vertex, which is precisely (I.8). Finally,

$$\sum_a q_a \chi_a(q) = \mathrm{tr}\big(J(q)J(q)^{-1}CJ(q)^{-1}\big) = \Phi_{D_\Psi,W}(q).$$

If every sensitivity is at most this weighted average, equality must hold at every positively weighted design point. □

I.2 Closed forms under unrestricted entrywise perturbations

Assume now that $\mathbb{T} = \mathbb{R}^{m\times k}$ and use the canonical entry basis, so $d = mk$ and $U = \mathbf{I}_{mk}$. Let $d_a \in \mathbb{R}^r$ be the column of D_Ψ corresponding to payoff entry a, and write $\|v\|_W = (v^\mathsf{T} W v)^{1/2}$.

If $D_\Psi = 0$, then the first-order target is constant, the Gaussian linear minimax risk is zero, and every design is optimal for this first-order criterion. Hence the normalized allocation below is stated only for $D_\Psi \neq 0$. Coordinates with $d_a = 0$ may receive zero mass, whereas $d_a \neq 0$ requires positive mass in the unrestricted diagonal model. For structured atomic information, several distinct designs can generate the same optimal information matrix, so uniqueness is not asserted outside the unrestricted coordinate case.

Theorem I.3 (Closed-form local-oracle full-coordinate design)**.** *If the target derivative is not identically zero, the unique minimizer of* (I.3) *is*

$$\boxed{q_a^{\star,\Psi} = \frac{\sigma_a \|d_a\|_W}{\sum_{b\in\mathcal{A}} \sigma_b \|d_b\|_W},} \tag{I.10}$$

with the convention that entries with $d_a = 0$ receive zero mass. The optimal risk is

$$\boxed{\Phi^\star_{D_\Psi,W} = \left(\sum_{a\in\mathcal{A}} \sigma_a \|d_a\|_W\right)^2.} \tag{I.11}$$

Proof. Here $J(q) = \mathrm{diag}(q_a/\sigma_a^2)_a$, and therefore

$$\Phi_{D_\Psi,W}(q) = \sum_{a:d_a\neq 0} \frac{\sigma_a^2 \|d_a\|_W^2}{q_a}, \tag{I.12}$$

with value $+\infty$ if $q_a = 0$ for some $d_a \neq 0$. Cauchy–Schwarz gives

$$\left(\sum_a \sigma_a \|d_a\|_W\right)^2 \leq \left(\sum_a \frac{\sigma_a^2\|d_a\|_W^2}{q_a}\right)\left(\sum_a q_a\right).$$

Equality holds exactly at (I.10). Any mass assigned to a zero-derivative coordinate would strictly reduce the mass on a positive-derivative coordinate and increase the criterion, proving uniqueness. □

A useful extension is immediate. For targets (D_{Ψ_s}, W_s) and nonnegative weights ω_s, the design minimizing the weighted aggregate $\sum_s \omega_s \Phi_{D_{\Psi_s},W_s}(q)$ is

$$q_a^\star \propto \sigma_a \left(\sum_s \omega_s \|d_{s,a}\|_{W_s}^2\right)^{1/2}. \tag{I.13}$$

Thus changing the loss weights moves the optimal experiment along an explicit multi-target Pareto path.

Local-oracle value design. At a game with a unique equilibrium $(x^\star, y^\star)$, proposition E.10 gives

$$D\,\mathrm{val}(A)[H] = x^{\star\mathsf{T}} H y^\star, \qquad \frac{\partial\,\mathrm{val}(A)}{\partial A_{ij}} = x_i^\star y_j^\star. \tag{I.14}$$

Hence the local-oracle value design is

$$\boxed{q_{ij}^{\star,\mathrm{val}} = \frac{\sigma_{ij} x_i^\star y_j^\star}{\sum_{a,b}\sigma_{ab}x_a^\star y_b^\star}, \qquad \Phi^\star_{\mathrm{val}} = \left(\sum_{i,j}\sigma_{ij}x_i^\star y_j^\star\right)^2.} \tag{I.15}$$

Under homoscedastic noise, this reduces to

$$q_{ij}^{\star,\mathrm{val}} = x_i^\star y_j^\star, \qquad \Phi^\star_{\mathrm{val}} = \sigma^2. \tag{I.16}$$

Thus the local-oracle value experiment samples action pairs according to the equilibrium product distribution.

Local-oracle interface-side design. At a simple bifurcation, let

$$a_{ij} \coloneqq Dg(A_0)[E^{ij}], \tag{I.17}$$

where E^{ij} is the coordinate matrix. The interface-side parameter class is fixed independently of the design:

$$\mathcal{H}_\pm(\lambda) = \{h : \pm a^\mathsf{T} h \geq \lambda\}. \tag{I.18}$$

Under design q, let

$$s_a^2(q) \coloneqq a^\mathsf{T} J(q)^\dagger a, \tag{I.19}$$

with $s_a(q) = +\infty$ when $a \notin \mathrm{range}(J(q))$. By theorem G.2, the exact composite interface-side risk in the Gaussian score experiment is

$$R^\star_{\mathrm{side}}(q;\lambda) = \begin{cases} \Phi(-\lambda/s_a(q)), & a \in \mathrm{range}(J(q)), \\ 1/2, & a \notin \mathrm{range}(J(q)). \end{cases} \tag{I.20}$$

Thus a local-oracle interface-side design minimizes $s_a^2(q)$ over one fixed pair of half-space classes. In full coordinates,

$$\boxed{q_{ij}^{\star,\mathrm{side}} = \frac{\sigma_{ij}|a_{ij}|}{\sum_{r,s}\sigma_{rs}|a_{rs}|}, \qquad s_{a,\star}^2 = \left(\sum_{i,j}\sigma_{ij}|a_{ij}|\right)^2,} \tag{I.21}$$

with zero mass allowed exactly on coordinates with $a_{ij} = 0$. This is the finite-entry analogue of c-optimal design (Elfving, 1952). This interface-side design is local-oracle because the boundary covector a is unknown. Equations (I.15) and (I.21) depend on different geometric objects: equilibrium occupancy versus bifurcation normal.

Corollary I.4 (Fixed-class local-oracle interface-side entry design)**.** *In the unrestricted entrywise model, let*

$$J(q) = \mathrm{diag}\left(\frac{q_i}{\sigma_i^2}\right)_{i=1}^p, \qquad q \in \Delta_p.$$

For the fixed classes (I.18)*, the interface-side problem is identifiable iff* $q_i > 0$ *whenever* $a_i \neq 0$*. On this domain,*

$$s_a^2(q) = a^\mathsf{T} J(q)^\dagger a = \sum_{i:a_i\neq 0}\frac{\sigma_i^2 a_i^2}{q_i}. \tag{I.22}$$

The unique approximate design minimizing interface-side risk is

$$\boxed{q_i^{\star,\mathrm{side}} = \frac{\sigma_i|a_i|}{\sum_j \sigma_j|a_j|},} \tag{I.23}$$

with $q_i^{\star,\mathrm{side}} = 0$ *when* $a_i = 0$*, and*

$$\inf_q s_a^2(q) = \left(\sum_i \sigma_i|a_i|\right)^2. \tag{I.24}$$

Consequently,

$$\inf_q R^\star_{\mathrm{side}}(J(q); a, \lambda) = \Phi\left(-\frac{\lambda}{\sum_i \sigma_i|a_i|}\right).$$

The parameter classes being compared do not depend on q*; the pair* (G.11) *is the least-favorable pair induced by the chosen design.*

Proof. For diagonal $J(q)$, $a \in \mathrm{range}(J(q))$ iff every nonzero coordinate of a receives positive mass, and the pseudoinverse formula gives (I.22). For every finite-risk design, Cauchy–Schwarz yields

$$\left(\sum_i \sigma_i|a_i|\right)^2 = \left(\sum_{i:a_i\neq 0}\frac{\sigma_i|a_i|}{\sqrt{q_i}}\sqrt{q_i}\right)^2 \leq \left(\sum_{i:a_i\neq 0}\frac{\sigma_i^2 a_i^2}{q_i}\right)\left(\sum_{i:a_i\neq 0} q_i\right) \leq s_a^2(q).$$

Equality requires all mass on the nonzero coordinates and $q_i \propto \sigma_i|a_i|$, proving (I.23)– (I.24). The interface-side error is strictly increasing in $s_a(q)$, so the same design minimizes the exact composite half-space risk. □

Approximate versus integer designs. The simplex variable q is an approximate design. With a sample budget n, choose integer counts by largest-remainder rounding. If the oracle design is interior on the coordinates relevant to the target, then proposition I.5 gives $\|n^{-1}N_n - q\|_\infty \leq n^{-1}$ and changes every smooth information criterion by $O(n^{-1})$. Exact finite-n formulae below apply when the displayed allocation is exactly realizable; otherwise they are the corresponding approximate-design formulae with an $O(n^{-1})$ information perturbation.

I.3 Integer implementation of approximate designs

Proposition I.5 (Rounding an interior approximate design)**.** *Let $q \in \Delta_p$ have support S and satisfy $q_i \geq q_{\min} > 0$ for $i \in S$. There exist integer counts $N_i^{(n)} \geq 0$, $\sum_i N_i^{(n)} = n$, such that $N_i^{(n)} = 0$ for $i \notin S$ and*

$$\left| \frac{N_i^{(n)}}{n} - q_i \right| \leq \frac{1}{n} \qquad (i \in S). \tag{I.25}$$

For all $n \geq 2/q_{\min}$, let $q_i^{(n)} = N_i^{(n)}/n$. Then, for every vector a supported on S,

$$\left| \sum_{i \in S} \frac{\sigma_i^2 a_i^2}{q_i^{(n)}} - \sum_{i \in S} \frac{\sigma_i^2 a_i^2}{q_i} \right| \leq \frac{2}{n q_{\min}^2} \sum_{i \in S} \sigma_i^2 a_i^2. \tag{I.26}$$

Hence every interior approximate-design interface-side risk is attained by integer counts up to an $O(n^{-1})$ information-criterion error. Boundary coordinates with $q_i = 0$ are implementable exactly by assigning zero counts.

Proof. For $i \in S$, write $nq_i = \lfloor nq_i \rfloor + r_i$ with $r_i \in [0, 1)$. Put

$$R_n = n - \sum_{i \in S} \lfloor nq_i \rfloor = \sum_{i \in S} r_i,$$

which is an integer in $\{0, \ldots, |S| - 1\}$. Set $N_i^{(n)} = \lfloor nq_i \rfloor + 1$ for the R_n largest fractional parts r_i and $N_i^{(n)} = \lfloor nq_i \rfloor$ for the remaining support coordinates; set $N_i^{(n)} = 0$ outside S. Then the counts sum to n and $|N_i^{(n)} - nq_i| < 1$, proving (I.25). For $n \geq 2/q_{\min}$, $q_i^{(n)} \geq q_i - 1/n \geq q_{\min}/2$. Therefore

$$\left| \frac{1}{q_i^{(n)}} - \frac{1}{q_i} \right| = \frac{|q_i^{(n)} - q_i|}{q_i q_i^{(n)}} \leq \frac{2}{n q_{\min}^2}.$$

Multiplying by $\sigma_i^2 a_i^2$ and summing proves (I.26). □

I.4 A declared-family local-oracle incompatibility theorem

Consider the two-parameter family

$$A_{\mu,\theta} = \mu \mathbf{1}\mathbf{1}^\mathsf{T} + \begin{pmatrix} 0 & \theta \\ -\theta & 0 \end{pmatrix} = \begin{pmatrix} \mu & \mu + \theta \\ \mu - \theta & \mu \end{pmatrix}. \tag{I.27}$$

Then $\mathrm{val}(A_{\mu,\theta}) = \mu$, while the unique equilibrium branch is (e_1, e_1) for $\theta > 0$ and (e_2, e_2) for $\theta < 0$. Let diagonal queries have variance 1 and off-diagonal queries variance multiplier $r_{\mathrm{noise}} > 1$.

For an arbitrary design q, set

$$\alpha = q_{12} + q_{21}, \qquad \xi = q_{12} - q_{21}.$$

The per-observation Fisher information for (μ, θ) is

$$J(q) = \begin{pmatrix} 1 - \alpha + \alpha/r_{\mathrm{noise}} & \xi/r_{\mathrm{noise}} \\ \xi/r_{\mathrm{noise}} & \alpha/r_{\mathrm{noise}} \end{pmatrix}. \tag{I.28}$$

For fixed α, replacing (q_{12}, q_{21}) by $(\alpha/2, \alpha/2)$ sets $\xi = 0$ and weakly decreases both diagonal entries of $J(q)^{-1}$. Hence every Pareto-efficient design is equivalent to one with equal off-diagonal allocation. Write that design as q_α; its diagonal split is immaterial and

$$J(q_\alpha) = \mathrm{diag}(1 - \alpha + \alpha/r_{\mathrm{noise}}, \alpha/r_{\mathrm{noise}}). \tag{I.29}$$

Theorem I.6 (Exact local-oracle value–interface-side conflict in the declared family)**.** *Fix a variance multiplier $r_{\mathrm{noise}} > 1$ and local interface-side alternatives $\theta_{n,\eta} = \eta c/\sqrt{n}$, $c > 0$. In the Gaussian approximate-design experiment with information $nJ(q_\alpha)$—and exactly for deterministic counts whenever nq_α is integral—the scaled value variance and minimax interface-side error are*

$$V_{r_{\mathrm{noise}}}(\alpha) = \frac{1}{1 - \alpha + \alpha/r_{\mathrm{noise}}}, \qquad B_{r_{\mathrm{noise}},c}(\alpha) = \Phi\left(-c \sqrt{\frac{\alpha}{r_{\mathrm{noise}}}} \right). \tag{I.30}$$

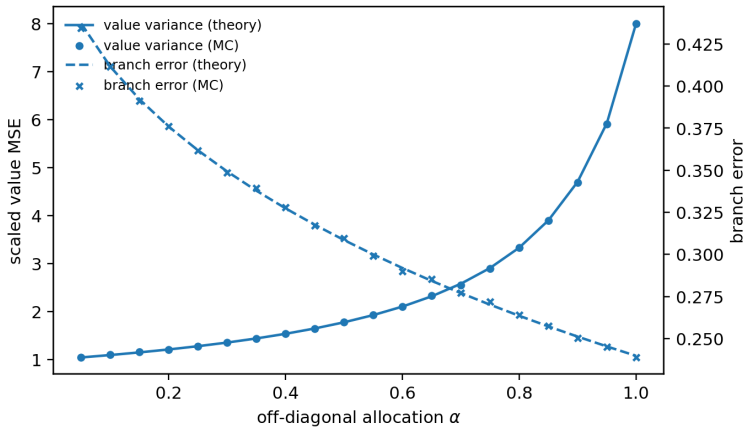

Figure 4: Finite-sample verification of the value–interface-side conflict. Solid and dashed curves are the exact approximate-design predictions; markers are seeded Monte Carlo estimates.

The function $V_{r_{\mathrm{noise}}}$ is strictly increasing and $B_{r_{\mathrm{noise}},c}$ is strictly decreasing. These formulas are exact for the continuous Gaussian approximate-design score experiment. For deterministic integer counts they are evaluated exactly at the realized proportions; the continuous allocation is exactly realizable only at compatible rational budgets, and largest-remainder rounding approaches it with the $O(n^{-1})$ criterion error proved below. Consequently, every $\alpha \in [0,1]$ is Pareto efficient,

$$\alpha^{\star}_{\mathrm{val}} = 0, \qquad \alpha^{\star}_{\mathrm{side}} = 1, \tag{I.31}$$

and no design minimizes both targets. Eliminating α, the exact Pareto frontier is

$$\boxed{V = \left[1 - \frac{r_{\mathrm{noise}} - 1}{c^2}\left\{-\Phi^{-1}(B)\right\}^2\right]^{-1}, \quad B \in \left[\Phi(-c/\sqrt{r_{\mathrm{noise}}}), \frac{1}{2}\right].} \tag{I.32}$$

In particular, the local-oracle value design has scaled value variance 1 but no interface-side information, whereas the local-oracle interface-side design has interface-side error $\Phi(-c/\sqrt{r_{\mathrm{noise}}})$ and inflates the value variance by the factor r_{noise}.

Proof. For (I.28), the feature vectors are $(1,0)$ on the diagonal and $(1,1)$, $(1,-1)$ off the diagonal. If $J = \left(\begin{smallmatrix} s & r \\ r & t \end{smallmatrix}\right)$, then $(J^{-1})_{11} = t/(st - r^2) \geq 1/s$ and $(J^{-1})_{22} = s/(st - r^2) \geq 1/t$, proving the symmetrization claim. Under (I.29), efficient estimation of μ has asymptotic variance $V_{r_{\mathrm{noise}}}(\alpha)/n$. The normalized sufficient statistic for θ has variance $r_{\mathrm{noise}}/\alpha$, so the exact symmetric Gaussian testing error is $B_{r_{\mathrm{noise}},c}(\alpha)$, with the convention $B_{r_{\mathrm{noise}},c}(0) = 1/2$. Differentiation gives strict monotonicity for $r_{\mathrm{noise}} > 1$. Finally, $-\Phi^{-1}(B) = c\sqrt{\alpha/r_{\mathrm{noise}}}$, and substitution into $V_{r_{\mathrm{noise}}}(\alpha)$ gives (I.32). □

Numerical verification. The supplementary script `verify_design_conflict.py` independently solves the full-coordinate convex design problem, implements largest-remainder integer rounding, and compares the exact curves in (I.30) with seeded finite-sample Monte Carlo estimates. For the displayed configuration, the numerical optimizer agrees with the closed form to below 10^{-8} in maximum coordinate error; the simulation is used only as a reproducibility check, not as part of the proof.

I.5 Fixed-confidence characteristic time

The same target dependence appears in sequential identification. Let $\mathcal{M}$ be a payoff model, let $\Psi : \mathcal{M} \to (\mathcal{Y}, d)$ be a measurable target, and define

$$\mathrm{Alt}_{\Psi,\varepsilon}(A) := \{B \in \mathcal{M} : d(\Psi(B), \Psi(A)) > 2\varepsilon\}. \tag{I.33}$$

A sequential procedure chooses entries predictably from the past transcript, stops at an almost surely finite stopping time T, and returns a measurable $\widehat{\Psi}$. It is (ε, δ)-correct when

$$P_A\{d(\widehat{\Psi}, \Psi(A)) \le \varepsilon\} \ge 1 - \delta \qquad \forall A \in \mathcal{M}.$$

Define

$$\Gamma^\star_{\Psi,\varepsilon}(A) \coloneqq \sup_{q \in \Delta_{|\mathcal{A}|}} \inf_{B \in \mathrm{Alt}_{\Psi,\varepsilon}(A)} \frac{1}{2} \sum_{a \in \mathcal{A}} q_a \frac{(A_a - B_a)^2}{\sigma_a^2}. \tag{I.34}$$

Proposition I.7 (Target-specific sequential lower bound)**.** *Let $\delta \in (0, 1/2)$ and suppose the procedure is (ε, δ)-correct and $\mathbb{E}_A \mathsf{T} < \infty$. Then*

$$\boxed{\mathbb{E}_A \mathsf{T} \ge \frac{\mathrm{kl}(1-\delta, \delta)}{\Gamma^\star_{\Psi,\varepsilon}(A)},} \tag{I.35}$$

with the convention $1/0 = +\infty$.

Proof. Fix $B \in \mathrm{Alt}_{\Psi,\varepsilon}(A)$ and let $E_A = \{d(\widehat{\Psi}, \Psi(A)) \le \varepsilon\}$. The event $E_B = \{d(\widehat{\Psi}, \Psi(B)) \le \varepsilon\}$ is disjoint from E_A by the strict 2ε target separation. Hence correctness gives $P_A(E_A) \ge 1 - \delta$ and $P_B(E_A) \le \delta$. Data processing for the stopped transcript therefore yields

$$D_{\mathrm{KL}}(P_A^{\mathsf{T}} \| P_B^{\mathsf{T}}) \ge \mathrm{kl}(1-\delta, \delta). \tag{I.36}$$

For the truncated time $\mathsf{T}_N = \mathsf{T} \wedge N$, the chain rule and predictability of the queried entry give

$$D_{\mathrm{KL}}(P_A^{\mathsf{T}_N} \| P_B^{\mathsf{T}_N}) = \frac{1}{2} \sum_a \mathbb{E}_A N_a(\mathsf{T}_N) \frac{(A_a - B_a)^2}{\sigma_a^2}.$$

All one-step Gaussian KL divergences are finite and $\mathbb{E}_A \mathsf{T} < \infty$. The truncated transcript sigma-fields increase to the stopped transcript sigma-field; relative entropy is continuous from below along increasing sigma-fields. Therefore $D_{\mathrm{KL}}(P_A^{\mathsf{T}_N} \| P_B^{\mathsf{T}_N}) \uparrow D_{\mathrm{KL}}(P_A^{\mathsf{T}} \| P_B^{\mathsf{T}})$, while monotone convergence gives $\mathbb{E}_A N_a(\mathsf{T}_N) \uparrow \mathbb{E}_A N_a(\mathsf{T})$. Hence the same decomposition holds at T; equivalently, this is the stopped change-of-measure identity of Kaufmann et al. (2016, Lemma 1). Let $\bar{q}_a = \mathbb{E}_A N_a(\mathsf{T}) / \mathbb{E}_A \mathsf{T}$. Combining the identity with (I.36), taking the infimum over B, and then using the supremum in (I.34) proves the result. This is the standard stopped change-of-measure argument of sequential experimental design and best-arm identification (Chernoff, 1959; Kaufmann et al., 2016; Garivier & Kaufmann, 2016). □

At a locally linearized simple boundary, the fixed-confidence class is again a fixed half-space. Theorem G.2 shows that its information exponent is governed by $a^\mathsf{T} J(q)^\dagger a$; hence the same local-oracle interface-side design appears in fixed-budget and fixed-confidence formulations.

I.6 Two-stage attainment at a fixed regular game

The following result is deliberately pointwise and interior. It does not adapt uniformly through a jump and does not cover sparse boundary optima without an additional forced-exploration analysis.

Proposition I.8 (Two-stage oracle attainment for an interior oracle)**.** *Assume the unrestricted entrywise model. Let K be a compact convex neighborhood of a fixed game A, contained in an open set on which Ψ is twice continuously differentiable. Suppose its oracle design $q^\star(B) \coloneqq q^{\star, D_\Psi(B)}$ from* (I.10) *is unique and continuous on K, and*

$$\inf_{B \in K} \min_a q_a^\star(B) \ge q_{\min} > 0. \tag{I.37}$$

Let $m_n / \log n \to \infty$ and $m_n / n \to 0$. Use m_n uniformly allocated pilot observations to form $\widetilde{A}_n$. On the event $\widetilde{A}_n \in K$, use largest-remainder integer allocation of the second stage according to $q^\star(\widetilde{A}_n)$; otherwise use the uniform allocation. Let $\overline{A}_n^{(2)}$ be the vector of second-stage means, let Π_K be Euclidean projection onto K, and define the estimator globally by

$$\widehat{\Psi}_n \coloneqq \Psi(\Pi_K \overline{A}_n^{(2)}). \tag{I.38}$$

Then

$$\sqrt{n}\,[\widehat{\Psi}_n - \Psi(A)] \Longrightarrow \mathcal{N}\left(0, D_\Psi(A)\,\mathrm{diag}\left(\frac{\sigma_a^2}{q_a^\star(A)}\right) D_\Psi(A)^\mathsf{T}\right), \tag{I.39}$$

and

$$n\mathbb{E}_A\|\widehat{\Psi}_n - \Psi(A)\|_W^2 \longrightarrow \Phi^\star_{D_\Psi(A),W}. \tag{I.40}$$

Proof. Choose $r > 0$ so that the closed Frobenius ball $\overline{B}_F(A, 2r)$ lies in the interior of K. Uniform Gaussian pilot sampling gives, for constants $c_1, c_2 > 0$,

$$P_A\{\|\widetilde{A}_n - A\|_F > r\} \le c_1 e^{-c_2 m_n}.$$

Thus $nP_A\{\widetilde{A}_n \notin K\} \to 0$. On the pilot-success event, continuity and uniqueness give $q^\star(\widetilde{A}_n) \to q^\star(A)$ in probability. Largest-remainder rounding changes every second-stage proportion by at most $1/(n - m_n)$, while (I.37) keeps all counts of order n.

Conditional on the pilot, the second-stage error $E_n = \overline{A}_n^{(2)} - A$ is centered Gaussian and

$$n\,\mathrm{Cov}(E_n \mid \widetilde{A}_n) \longrightarrow \mathrm{diag}\left(\frac{\sigma_a^2}{q_a^\star(A)}\right)$$

in probability. Its conditional fourth moments are uniformly bounded after multiplication by n^2. Projection is inactive with probability $1 - o(n^{-1})$: on pilot success, Gaussian concentration and the positive distance from A to K^c give an exponentially small probability that $\overline{A}_n^{(2)} \notin K$; pilot failure already has probability $o(n^{-1})$.

Since Ψ is C^2 on a neighborhood of K, there is $L < \infty$ such that

$$\|\Psi(A + u) - \Psi(A) - D_\Psi(A)u\|_W \le L\|u\|_F^2 \tag{I.41}$$

whenever $A + u \in K$. Consequently $n\mathbb{E}\|\text{remainder}\|_W^2 = O(n^{-1})$, and the cross term with the linear part vanishes by Cauchy–Schwarz. Conditional Slutsky gives the CLT; the uniform fourth-moment bound gives uniform integrability, so expectations converge and yield (I.40). The limit is the oracle criterion from theorem I.3. □

The design conclusions can be summarized as

$$\boxed{\text{target tangent} \longmapsto \text{information criterion} \longmapsto \text{optimal observation allocation.}} \tag{I.42}$$

For values, the tangent is the equilibrium occupancy; for local interface-side recovery, it is the bifurcation normal. These objects need not align, and theorem I.6 shows that their local-oracle objectives can select opposite endpoints of the explicit allocation frontier in theorem I.6.

J Bellman Pullbacks in Finite-Horizon Zero-Sum Markov Games

This appendix propagates the matrix-game geometry of sections D and E through finite-horizon dynamic programming. We first establish a model-perturbation bound that is uniform over policies, then prove a Hadamard directional Bellman recursion, and finally give a fully conditional sectorwise policy-branch pullback under explicit statewise separation conditions. We do not claim an unconditional global Bellman fan. The last subsection gives an exact statistical separation in a two-stage Markov game: initial value and exploitability remain parametric while recovery of a downstream equilibrium branch has nonvanishing risk.

Finite stochastic games originate with Shapley (1953); their use as a reinforcement-learning model was developed by Littman (1994). Modern finite-horizon competitive-RL algorithms also rely on the statewise minimax Bellman recursion used below (Bai & Jin, 2020; Liu et al., 2021). Directional derivatives of stochastic game values are already established by Attia et al. (2025); our narrower claim couples a finite-horizon recursion to statewise equilibrium-branch geometry and target-specific statistical separation.

J.1 Finite-horizon model and Bellman saddle principle

Fix a horizon $T \in \mathbb{N}$, a finite state space $\mathcal{S}$, and finite action sets $\mathcal{A} = [m]$ and $\mathcal{B} = [k]$. A two-player zero-sum Markov game is

$$\mathcal{M} = \big(\{R_t\}_{t=1}^{T}, \{P_t\}_{t=1}^{T-1}\big),$$

where

$$R_t(s,a,b) \in \mathbb{R}, \qquad P_t(\cdot \mid s,a,b) \in \Delta_{|\mathcal{S}|}.$$

The row player receives R_t and the column player receives $-R_t$. A nonstationary Markov policy pair is denoted (π, ν), with $\pi_t(\cdot \mid s) \in \Delta_m$ and $\nu_t(\cdot \mid s) \in \Delta_k$. Its continuation value is

$$V_{\mathcal{M},t}^{\pi,\nu}(s) \coloneqq \mathbb{E}_{\mathcal{M}}^{\pi,\nu}\left[\sum_{u=t}^{T} R_u(S_u, A_u, B_u) \,\middle|\, S_t = s\right]. \tag{J.1}$$

For a policy pair, define the dynamic saddle gap

$$\mathrm{Gap}_{\mathcal{M},t}^{\mathrm{dyn}}(\pi, \nu; s) \coloneqq \sup_{\pi'} V_{\mathcal{M},t}^{\pi',\nu}(s) - \inf_{\nu'} V_{\mathcal{M},t}^{\pi,\nu'}(s), \tag{J.2}$$

where the deviations may be arbitrary behavioral policies. Against a fixed Markov opponent, a Markov best response exists, so the extrema in (J.2) are unchanged if restricted to nonstationary Markov policies.

Set $V_{\mathcal{M},T+1} \equiv 0$ and define backward

$$Q_{\mathcal{M},t}(s,a,b) \coloneqq R_t(s,a,b) + \sum_{s' \in \mathcal{S}} P_t(s' \mid s,a,b) V_{\mathcal{M},t+1}(s'), \tag{J.3}$$

$$V_{\mathcal{M},t}(s) \coloneqq \mathrm{val}\big(Q_{\mathcal{M},t}(s)\big), \tag{J.4}$$

where the transition term is absent at $t = T$. Write

$$\mathcal{E}_{\mathcal{M},t}(s) \coloneqq \mathcal{E}\big(Q_{\mathcal{M},t}(s)\big).$$

Proposition J.1 (Bellman saddle principle)**.** *For every finite-horizon zero-sum Markov game,* (J.3)–(J.4) *equal the upper and lower subgame values. If*

$$(\pi_t(\cdot \mid s), \nu_t(\cdot \mid s)) \in \mathcal{E}_{\mathcal{M},t}(s) \qquad \textit{for every } (t,s), \tag{J.5}$$

then (π, ν) *is a subgame-perfect saddle policy and*

$$\mathrm{Gap}_{\mathcal{M},t}^{\mathrm{dyn}}(\pi, \nu; s) = 0 \qquad \textit{for every } (t,s).$$

Proof. The claim follows by backward induction. It is the minimax theorem for the terminal matrix game at $t = T$. Suppose the claim holds from $t+1$ onward. Condition on the first state and action pair at time t. The continuation under the statewise saddle policies is exactly $V_{\mathcal{M},t+1}(S_{t+1})$, while either player's arbitrary continuation deviation is bounded by the corresponding induction hypothesis. Hence the one-stage game at state s has payoff matrix $Q_{\mathcal{M},t}(s)$, and any saddle point of this matrix, followed by the continuation saddle policies, is a saddle for the entire subgame. This proves both equality of the upper and lower values and (J.5). □

J.2 Uniform dynamic stability

The matrix perturbation bounds of proposition C.1 admit a dynamic analogue. Let $\widetilde{\mathcal{M}}$ have the same state and action spaces and horizon as $\mathcal{M}$. Suppose both games satisfy

$$\max_{t,s,a,b} \max\{|R_t(s,a,b)|, |\widetilde{R}_t(s,a,b)|\} \le R_{\max}.$$

Define

$$\delta_R \coloneqq \max_{t,s,a,b} |R_t(s,a,b) - \widetilde{R}_t(s,a,b)|, \qquad \delta_P \coloneqq \max_{t<T,s,a,b} \|P_t(\cdot \mid s,a,b) - \widetilde{P}_t(\cdot \mid s,a,b)\|_1. \tag{J.6}$$

For $L_t \coloneqq T - t + 1$, put

$$B_t(\delta_R, \delta_P) \coloneqq L_t \delta_R + \frac{L_t(L_t - 1)}{2} R_{\max} \delta_P. \tag{J.7}$$

Proposition J.2 (Value and exploitability stability)**.** *For every policy pair (π, ν) and every t,*

$$\max_s \left| V^{\pi,\nu}_{\mathcal{M},t}(s) - V^{\pi,\nu}_{\widetilde{\mathcal{M}},t}(s) \right| \le B_t(\delta_R, \delta_P). \tag{J.8}$$

Consequently,

$$\|V_{\mathcal{M},t} - V_{\widetilde{\mathcal{M}},t}\|_\infty \le B_t(\delta_R, \delta_P). \tag{J.9}$$

If $(\widetilde{\pi}, \widetilde{\nu})$ is a saddle policy for $\widetilde{\mathcal{M}}$, then

$$\sup_s \mathrm{Gap}^{\mathrm{dyn}}_{\mathcal{M},t}(\widetilde{\pi}, \widetilde{\nu}; s) \le 2B_t(\delta_R, \delta_P). \tag{J.10}$$

Proof. For a fixed policy pair, let the left side of (J.8) be d_t. Since every continuation value is bounded by $(T-t)R_{\max}$, the policy Bellman equations give

$$d_t \le \delta_R + d_{t+1} + (T-t)R_{\max}\delta_P, \qquad d_{T+1} = 0.$$

Summation yields (J.7). The argument is uniform over policies; taking max–min values proves (J.9). Finally,

$$\sup_{\pi'} V^{\pi',\widetilde{\nu}}_{\mathcal{M},t}(s) \le \sup_{\pi'} V^{\pi',\widetilde{\nu}}_{\widetilde{\mathcal{M}},t}(s) + B_t,$$
$$\inf_{\nu'} V^{\widetilde{\pi},\nu'}_{\mathcal{M},t}(s) \ge \inf_{\nu'} V^{\widetilde{\pi},\nu'}_{\widetilde{\mathcal{M}},t}(s) - B_t.$$

The two middle extrema are equal because $(\widetilde{\pi}, \widetilde{\nu})$ is a saddle policy. Subtraction proves (J.10). □

J.3 Directional Bellman calculus

We now parameterize the model smoothly. Let $\Theta \subseteq \mathbb{R}^p$ contain an open neighborhood of ϑ_0. For every (t, s, a, b), assume

$$\vartheta \longmapsto R^\vartheta_t(s,a,b), \qquad \vartheta \longmapsto P^\vartheta_t(\cdot \mid s,a,b) \tag{J.11}$$

are twice continuously differentiable near ϑ_0, with the latter map valued in the probability simplex. Write the base model as $\mathcal{M}_0 = \mathcal{M}_{\vartheta_0}$ and, for $h \in \mathbb{R}^p$,

$$\dot{R}^h_t := DR^{\vartheta_0}_t[h], \qquad \dot{P}^h_t := DP^{\vartheta_0}_t[h]. \tag{J.12}$$

Each transition derivative has row sum zero. All Hadamard sequences below are relative to Θ; no arbitrary two-sided perturbation of a zero transition probability is allowed. Put $\dot{V}^h_{T+1} \equiv 0$ and define recursively

$$\dot{Q}^h_t(s,a,b) := \dot{R}^h_t(s,a,b) + \sum_{s'} \dot{P}^h_t(s' \mid s,a,b) V^0_{t+1}(s') + \sum_{s'} P^0_t(s' \mid s,a,b) \dot{V}^h_{t+1}(s'), \tag{J.13}$$

$$\dot{V}^h_t(s) := \max_{x \in \mathcal{X}^\star(Q^0_t(s))} \min_{y \in \mathcal{Y}^\star(Q^0_t(s))} x^\mathsf{T} \dot{Q}^h_t(s) y. \tag{J.14}$$

Theorem J.3 (Hadamard directional Bellman recursion)**.** *For every (t, s), the map $\vartheta \mapsto V^\vartheta_t(s)$ is locally Lipschitz and Hadamard directionally differentiable at ϑ_0. Its derivative in direction h is the quantity defined by* (J.13)–(J.14)*. More precisely, if $\epsilon_n \downarrow 0$ and $h_n \to h$, then*

$$\frac{Q^{\vartheta_0+\epsilon_n h_n}_t(s) - Q^0_t(s)}{\epsilon_n} \longrightarrow \dot{Q}^h_t(s), \qquad \frac{V^{\vartheta_0+\epsilon_n h_n}_t(s) - V^0_t(s)}{\epsilon_n} \longrightarrow \dot{V}^h_t(s). \tag{J.15}$$

For any base statewise equilibrium $z^0_{t,s} \in \mathcal{E}(Q^0_t(s))$, every graphical derivative of the statewise equilibrium correspondence in model direction h satisfies the critical-cone system

$$0 \in \mathcal{L}_{Q^0_t(s)} u + \mathcal{B}_{z^0_{t,s}}\big(\dot{Q}^h_t(s)\big) + N_{\mathcal{K}_{t,s}}(u), \tag{J.16}$$

where

$$\mathcal{K}_{t,s} = T_{\mathcal{Z}}(z^0_{t,s}) \cap [-F_{Q^0_t(s)}(z^0_{t,s})]^\perp.$$

Proof. Proceed backward in t. The terminal claim is trivial. Suppose the value claim holds at $t+1$. Proposition J.2 and smoothness of the model imply $\|V_{t+1}^{\vartheta_0+\epsilon h} - V_{t+1}^0\|_\infty = O(\epsilon)$ uniformly for h in compact sets. Expanding the transition-value product,

$$\begin{aligned}\frac{P_t^{\vartheta_0+\epsilon_n h_n} V_{t+1}^{\vartheta_0+\epsilon_n h_n} - P_t^0 V_{t+1}^0}{\epsilon_n} &= \frac{P_t^{\vartheta_0+\epsilon_n h_n} - P_t^0}{\epsilon_n} V_{t+1}^0 + P_t^0 \frac{V_{t+1}^{\vartheta_0+\epsilon_n h_n} - V_{t+1}^0}{\epsilon_n} \\ &\quad + \frac{P_t^{\vartheta_0+\epsilon_n h_n} - P_t^0}{\epsilon_n} \big(V_{t+1}^{\vartheta_0+\epsilon_n h_n} - V_{t+1}^0\big),\end{aligned}$$

where the last term converges to zero. This proves the Q-limit in (J.15). The value limit then follows from the Hadamard matrix-game derivative proposition E.10, giving (J.14). Local Lipschitz continuity follows from the same induction and the matrix value perturbation bound in proposition C.1. Finally, (J.16) is proposition E.7 applied to the induced payoff direction $\dot{Q}_t^h(s)$. □

J.4 Quadratic-remainder robustness of statewise chambers

Lemma J.4 (Quadratic-remainder robustness of a strict validity chamber)**.** *Fix a regular kernel branch β at A_0, so $\Delta_\beta(A_0) \neq 0$ and $\rho_\beta(A_0) \geq 0$. Let $\mathcal{A}_\beta = \{j : \rho_{\beta,j}(A_0) = 0\}$ be its active validity coordinates. Let K be compact and suppose*

$$\gamma := \min_{H\in K} \min_{j\in\mathcal{A}_\beta} D\rho_{\beta,j}(A_0)[H] > 0, \tag{J.17}$$

when $\mathcal{A}_\beta \neq \varnothing$; when $\mathcal{A}_\beta = \varnothing$, the margin condition is vacuous. Assume

$$A_{\varepsilon,H} = A_0 + \varepsilon H + R_{\varepsilon,H}, \qquad \sup_{H\in K} \|R_{\varepsilon,H}\|_\infty \leq C\varepsilon^2.$$

Then there exists $\varepsilon_0 > 0$ such that for every $0 < \varepsilon < \varepsilon_0$ and $H \in K$:

1. *$\Delta_\beta(A_{\varepsilon,H}) \neq 0$ and $\rho_\beta(A_{\varepsilon,H}) > 0$;*
2. *$z_\beta(A_{\varepsilon,H})$ is the unique equilibrium;*
3. *uniformly in $H \in K$,*

$$z_\beta(A_{\varepsilon,H}) = z_\beta(A_0) + \varepsilon Dz_\beta(A_0)[H] + O_K(\varepsilon^2). \tag{J.18}$$

Proof. The branch map, determinant, and validity map are analytic on a neighborhood of A_0. If $\mathcal{A}_\beta \neq \varnothing$, uniform Taylor expansion on the compact set K gives, for each active coordinate,

$$\begin{aligned}\rho_{\beta,j}(A_{\varepsilon,H}) &= D\rho_{\beta,j}(A_0)[\varepsilon H + R_{\varepsilon,H}] + O_K(\|\varepsilon H + R_{\varepsilon,H}\|_\infty^2) \\ &= \varepsilon D\rho_{\beta,j}(A_0)[H] + O_K(\varepsilon^2) \geq \varepsilon\gamma + O_K(\varepsilon^2).\end{aligned}$$

This is at least $\varepsilon\gamma/2$ for sufficiently small ε, uniformly in H. If $\mathcal{A}_\beta = \varnothing$, there are no active coordinates to check. Every inactive coordinate is strictly positive at A_0 and therefore remains positive by uniform continuity. Likewise, $\Delta_\beta(A_{\varepsilon,H})$ remains separated from zero. Hence $\rho_\beta(A_{\varepsilon,H}) > 0$, and the exact branch certificate proposition D.2 gives validity and uniqueness.

Analytic Taylor expansion of z_β gives

$$\begin{aligned}z_\beta(A_{\varepsilon,H}) &= z_\beta(A_0) + Dz_\beta(A_0)[\varepsilon H + R_{\varepsilon,H}] + O_K(\varepsilon^2) \\ &= z_\beta(A_0) + \varepsilon Dz_\beta(A_0)[H] + O_K(\varepsilon^2),\end{aligned}$$

which proves (J.18). □

If an induced Bellman direction is zero, this lemma applies only when the reference statewise branch is already strict, in which case the branch persists and its first derivative is zero. At a singular reference game, the zero direction lies on the fan skeleton and no first-order branch-selection claim is available.

J.5 SECTORWISE STATEWISE BRANCH PULLBACKS

The directional recursion above is unconditional within the declared smooth finite-dimensional model. We next state the strongest branch result used by the paper. Equip the primitive parameter space with the Euclidean norm and let K be a compact subset of its unit sphere. The first-order quantities $\dot{Q}_t^h(s)$ are computed by the backward recursion in (J.13), so the conditions below are finite statewise checks once the static validity refinements are known.

Proposition J.5 (Conditional sectorwise Bellman branch pullback)**.** *Assume the model maps in* (J.11) *are C^2 on a relative neighborhood of the reference parameter. Suppose that, for every stage–state pair (t, s), one of the following holds:*

1. *$Q_t^0(s)$ has a strictly valid regular kernel branch; or*
2. *there is a regular-kernel validity chamber $C_{t,s}$ such that*

$$\inf_{h\in K} \|\dot{Q}_t^h(s)\|_\infty > 0, \qquad \left\{ \frac{\dot{Q}_t^h(s)}{\|\dot{Q}_t^h(s)\|_\infty} : h \in K \right\} \Subset C_{t,s}. \tag{J.19}$$

Then there are statewise branches $\beta_{t,s}$ such that, uniformly over $h \in K$,

$$Q_t^{\vartheta_0+\varepsilon h}(s) = Q_t^0(s) + \varepsilon \dot{Q}_t^h(s) + O_K(\varepsilon^2), \tag{J.20}$$

$$V_t^{\vartheta_0+\varepsilon h}(s) = V_t^0(s) + \varepsilon \dot{V}_t^h(s) + O_K(\varepsilon^2), \tag{J.21}$$

$$z_{t,s}^{\varepsilon,h} = z_{t,s}^0 + \varepsilon D z_{\beta_{t,s}}(Q_t^0(s))[\dot{Q}_t^h(s)] + O_K(\varepsilon^2). \tag{J.22}$$

In case (1), a zero induced direction is allowed and gives zero first-order policy motion. In case (2), the nonvanishing condition in (J.19) *is essential. No conclusion is made for a zero direction at a singular statewise game.*

Proof. Proceed backward in t. At $t = T$, the model's C^2 expansion gives (J.20). In case (1), strict validity persists and the analytic branch expansion applies directly. In case (2), the normalized induced directions form a compact subset of one chamber and their norms are bounded below, so every active validity differential has a uniform positive margin on the compact set $\{\dot{Q}_T^h(s) : h \in K\}$. Lemma J.4 therefore gives (J.22). The branch envelope formula gives the value expansion.

Assume the statements hold at $t + 1$. The C^2 primitive model expansion, together with the uniform second-order value expansion at $t + 1$, yields

$$Q_t^{\vartheta_0+\varepsilon h}(s) = Q_t^0(s) + \varepsilon \dot{Q}_t^h(s) + O_K(\varepsilon^2)$$

uniformly in $h \in K$; the mixed product of the transition and continuation remainders is $O_K(\varepsilon^2)$. Apply the same strict-branch or chamber-separated argument at (t, s) to obtain the policy expansion, and then use the analytic branch envelope identity for the value expansion. Finiteness of the horizon and state space makes all constants uniform. This closes the backward induction. □

The proposition is deliberately sectorwise. Downstream value derivatives may be nonlinear across sectors, and we do not claim that the resulting primitive partition is a globally polyhedral fan. Checking the hypothesis requires only the backward derivative recursion, finite statewise kernel enumeration, and the chamber-separation tests in (J.19); the main paper uses only the unconditional directional recursion and the exact statistical separation below.

J.6 A ZERO INDUCED BELLMAN DIRECTION

Example J.6 (A nonzero primitive Markov perturbation with zero induced Bellman direction)**.** Consider a two-stage model in which the stage-one transition, independently of actions, reaches terminal states s_+ and s_- with probabilities $1/2$ and $1/2$. Let the primitive reward direction add $+1$ at s_+ and -1 at s_-. Then

$$\dot{V}_2(s_+) = 1, \qquad \dot{V}_2(s_-) = -1,$$

but every stage-one action pair has

$$\dot{Q}_1(a,b) = \frac{1}{2}\dot{V}_2(s_+) + \frac{1}{2}\dot{V}_2(s_-) = 0.$$

Thus a nonzero primitive direction can induce the zero statewise payoff direction. At a singular stage-one game, zero lies on the fan skeleton and no first-order policy branch is selected.

J.7 An exact dynamic statistical separation

We now embed the two-by-two singularity into a genuinely dynamic game. Fix $p_{\mathrm{reach}} \in (0,1]$ and consider a two-stage game with initial state s_0. At stage one, both players have a single action, receive zero reward, and the game moves to a decision state $s_\star$ with probability p_{reach} and to a null state $s_\perp$ with probability $1 - p_{\mathrm{reach}}$. At stage two, the null state has a single zero-reward action, while the decision state has payoff matrix

$$A_{\mu,\theta} = \mu \mathbf{1}_2 \mathbf{1}_2^\top + \begin{pmatrix} 0 & \theta \\ -\theta & 0 \end{pmatrix}. \tag{J.23}$$

Denote this Markov game by $\mathcal{M}_{\mu,\theta}^{(p_{\mathrm{reach}})}$. Then

$$V_2(s_\star) = \mu, \qquad V_1(s_0) = p_{\mathrm{reach}}\mu, \tag{J.24}$$

and the statewise equilibrium branch at $s_\star$ is determined by the sign of θ. Moreover, for any strategy pair z used at $s_\star$,

$$\mathrm{Gap}^{\mathrm{dyn}}_{\mathcal{M}_{\mu,\theta}^{(p_{\mathrm{reach}})},1}(z; s_0) = p_{\mathrm{reach}}\,\mathrm{Gap}_{A_{\mu,\theta}}(z). \tag{J.25}$$

Let n be even. A generative reward oracle supplies $n/2$ independent unit variance observations of each off-diagonal entry:

$$Y_i^+ \sim \mathcal{N}(\mu+\theta, 1), \qquad Y_i^- \sim \mathcal{N}(\mu-\theta, 1), \qquad i = 1, \dots, n/2. \tag{J.26}$$

Put

$$\widehat{\mu}_n = \frac{\overline{Y}^+ + \overline{Y}^-}{2}, \qquad \widehat{\theta}_n = \frac{\overline{Y}^+ - \overline{Y}^-}{2}. \tag{J.27}$$

Then exactly

$$\sqrt{n}\begin{pmatrix} \widehat{\mu}_n - \mu \\ \widehat{\theta}_n - \theta \end{pmatrix} \sim \mathcal{N}(0, \mathbf{I}_2), \tag{J.28}$$

and the two coordinates are independent.

For a Markov policy pair, use the sup-state metric

$$d_{\mathrm{pol}}((\pi,\nu),(\pi',\nu')) \coloneqq \max_{t,s} d_\Delta\big((\pi_t(\cdot \mid s), \nu_t(\cdot \mid s)), (\pi'_t(\cdot \mid s), \nu'_t(\cdot \mid s))\big). \tag{J.29}$$

The two pure equilibrium branches in (J.23) are separated by $d_{\mathrm{pol}} = 2$.

The sup-state metric gives a low-reach state full weight: it asks for the entire subgame-perfect policy. Under an occupancy-weighted policy metric, the contribution of $s_\star$ would instead be multiplied by its reachability and would generally shrink with p_{reach}.

Corollary J.7 (Exact zero-value dynamic separation)**.** *Fix $c > 0$, let $n \geq \lceil c^2 \rceil$, and set $\theta_{n,\eta} = \eta c/\sqrt{n}$ for $\eta \in \{-1,+1\}$. Then the initial and strategic-state values are identically zero and*

$$\inf_{\widehat{V}_n} \max_\eta \mathbb{E}_\eta\Big[\widehat{V}_n - V_1^{\theta_{n,\eta}}(s_0)\Big]^2 = 0, \tag{J.30}$$

$$\inf_{\widehat{b}_n} \max_\eta P_\eta(\widehat{b}_n \neq \eta) = \Phi(-c), \tag{J.31}$$

$$\inf_{(\widehat{\pi}_n,\widehat{\nu}_n)} \max_\eta \sqrt{n}\,\mathbb{E}_\eta\,\mathrm{Gap}^{\mathrm{dyn}}_{\mathcal{M}_{\theta_{n,\eta}},1}(\widehat{\pi}_n,\widehat{\nu}_n; s_0) = 2p_{\mathrm{reach}}\,c\Phi(-c), \tag{J.32}$$

$$\inf_{(\widehat{\pi}_n,\widehat{\nu}_n)} \max_\eta \mathbb{E}_\eta d^2_{\mathrm{pol}}\big((\widehat{\pi}_n,\widehat{\nu}_n),(\pi^\star_\eta,\nu^\star_\eta)\big) = e^{-c^2/2}\mathbb{E}[\cosh(cG)^{-1}], \tag{J.33}$$

where $G \sim \mathcal{N}(0,1)$. The exact branch and point-policy risks do not depend on p_{reach}, whereas the scaled initial exploitability is multiplied by p_{reach}.

Proof. The terminal matrix game has value zero for every θ by lemma B.2. Hence

$$V_2^\theta(s_\star) = 0, \qquad V_1^\theta(s_0) = p_{\mathrm{reach}}\, V_2^\theta(s_\star) = 0,$$

which proves (J.30).

By lemma B.1, the complete adaptive reward-query experiment is decision-theoretically equivalent to n independent $\mathcal{N}(\theta, 1)$ observations. Therefore the exact branch and point-strategy minimax risks are the quantities in proposition B.3, proving (J.31) and (J.33); the sup-state policy metric equals d_Δ at the sole strategic state.

For every terminal strategy pair z, conditioning on whether $s_\star$ is reached gives the pathwise identity

$$\mathrm{Gap}^{\mathrm{dyn}}_{\mathcal{M}_\theta,1}(z; s_0) = p_{\mathrm{reach}}\, \mathrm{Gap}_{A_\theta}(z).$$

Multiplying the exact scaled matrix-gap risk from proposition B.3 by p_{reach} proves (J.32). □

Corollary J.8 (Declared two-stage risks with an unknown additive value coordinate)**.** *Fix $c > 0$ and consider the local alternatives*

$$\mu_n = \frac{h}{\sqrt{n}}, \qquad \theta_{n,\eta} = \eta\frac{c}{\sqrt{n}}, \qquad \eta \in \{-1, +1\}. \tag{J.34}$$

Then the following hold.

1. *The initial value is regular, with exact local asymptotic minimax risk*

$$\lim_{R\to\infty} \lim_{n\to\infty} \inf_{\widehat{V}_n} \sup_{|h|\le R,\, \eta\in\{-1,+1\}} n\mathbb{E}_{h,\eta}\left[\widehat{V}_n - V_1^{\mu_n,\theta_{n,\eta}}(s_0)\right]^2 = p_{\mathrm{reach}}^2. \tag{J.35}$$

   The estimator $p_{\mathrm{reach}}\,\widehat{\mu}_n$ attains the bound.

2. *The finite-sample minimax branch error at $s_\star$ is*

$$\inf_{\widehat{b}_n} \max_{\eta} P_{h,\eta}(\widehat{b}_n \neq \eta) = \Phi(-c). \tag{J.36}$$

3. *The finite-sample minimax scaled initial exploitability is*

$$\inf_{(\widehat{\pi}_n,\widehat{\nu}_n)} \max_{\eta} \sqrt{n}\, \mathbb{E}_{h,\eta}\, \mathrm{Gap}^{\mathrm{dyn}}_{\mathcal{M}^{(p_{\mathrm{reach}})}_{\mu_n,\theta_{n,\eta}},1}(\widehat{\pi}_n, \widehat{\nu}_n; s_0) = 2p_{\mathrm{reach}}\, c\Phi(-c). \tag{J.37}$$

4. *For squared d_{pol} loss, the finite-sample point-policy minimax risk is*

$$\inf_{(\widehat{\pi}_n,\widehat{\nu}_n)} \max_{\eta} \mathbb{E}_{h,\eta} d_{\mathrm{pol}}^2\left((\widehat{\pi}_n, \widehat{\nu}_n), (\pi^\star_{n,\eta}, \nu^\star_{n,\eta})\right) = e^{-c^2/2}\mathbb{E}_{G\sim\mathcal{N}(0,1)}[\cosh(cG)^{-1}]. \tag{J.38}$$

Consequently, as $p_{\mathrm{reach}} \downarrow 0$, the initial value risk and the optimal initial exploitability shrink as p_{reach}^2 and p_{reach}, respectively, while branch and point-policy risks are unchanged.

Proof. Equation (J.24) follows from the Bellman recursion, and adding $\mu\mathbf{1}_2\mathbf{1}_2^\mathsf{T}$ does not alter the statewise equilibrium strategies. The initial state has no strategic choice, so (J.25) follows by conditioning on whether $s_\star$ is reached.

The statistic in (J.28) splits into two independent Gaussian location experiments. The value target is $\sqrt{n}V_1(s_0) = p_{\mathrm{reach}}\, h$, so lemma F.15 gives (J.35). The branch target depends only on the sign of the second coordinate, giving (J.36) by the exact Gaussian test in theorem G.2. The additive parameter μ cancels from all saddle gaps. Therefore Appendix B, together with (J.25), gives (J.37). Finally, the two Markov policy branches differ only at $s_\star$ and have d_{pol} separation two; the exact posterior-barycenter calculation in theorem G.3 with $\Delta_n = 2$ proves (J.38). □

Remark J.9 (Generative versus trajectory data). The separation above uses a generative reward oracle at $s_\star$. Under on-policy trajectory data, the information available about the downstream branch would itself depend on the reachability probability p_{reach}. The result isolates a different phenomenon: even when the singular state can be queried directly, its contribution to initial value and exploitability can be arbitrarily small while branch and point-policy targets remain nonregular under the stated squared sup-state loss.

The Bellman conclusions can be summarized as

$$\boxed{\text{model direction} \longmapsto \text{Bellman payoff direction} \longmapsto \begin{cases}\text{single-valued value tangent,}\\ \text{statewise policy branch or jump.}\end{cases}} \tag{J.39}$$

Under the squared uniform sup-state policy metric used above and the stated smooth sectorwise conditions, the dynamic program preserves the central separation of the paper: value perturbations compose through a scalar Hadamard recursion, whereas point-policy perturbations inherit the relevant statewise branch geometry. The transition-feedback calculation in proposition M.4 independently verifies all reward, transition, and downstream-value derivative terms. This conclusion is metric dependent; occupancy-weighted policy loss would generally attenuate low-reach states.

K Verified Extensions under Explicit Conditions

This appendix records three extensions that follow from the preceding theory without changing its geometric core. First, the exact Gaussian observation model can be replaced by any sequence of local experiments that is asymptotically equivalent, in Le Cam distance, to the same Gaussian shift. Second, for strongly convex tie-breaking regularizers that are finite on ambient neighborhoods of the simplices, we prove a bias–sensitivity theorem and a corresponding regularizer-selected equilibrium. Negative entropy is treated separately in Appendix C. Third, a finite collection of jump branches reduces to a Gaussian multi-hypothesis decision problem. We state only extensions for which complete proofs are given below; higher-order branch contacts, growing game dimensions, and trajectory-constrained designs are not covered.

K.1 Transport through asymptotically equivalent experiments

Let $\mathcal{H}$ be a compact local parameter set. A statistical experiment is denoted

$$\mathcal{E}_n = \big(\Omega_n, \mathcal{A}_n, \{P_{n,h} : h \in \mathcal{H}\}\big).$$

For experiments $\mathcal{E} = (P_h)_{h\in\mathcal{H}}$ and $\mathcal{F} = (Q_h)_{h\in\mathcal{H}}$, define the directed deficiency

$$\delta_{\mathcal{H}}(\mathcal{E},\mathcal{F}) \coloneqq \inf_K \sup_{h\in\mathcal{H}} d_{\mathrm{TV}}(KP_h, Q_h), \tag{K.1}$$

where the infimum is over Markov kernels from the sample space of $\mathcal{E}$ to that of $\mathcal{F}$, and define the Le Cam distance

$$\Delta_{\mathcal{H}}(\mathcal{E},\mathcal{F}) \coloneqq \max\big\{\delta_{\mathcal{H}}(\mathcal{E},\mathcal{F}), \delta_{\mathcal{H}}(\mathcal{F},\mathcal{E})\big\}. \tag{K.2}$$

This is the operational comparison of statistical experiments used throughout Le Cam's decision theory (Le Cam, 1986).

Let $\mathcal{G}_n = \{Q_{n,h} : h \in \mathcal{H}\}$ be a Gaussian shift experiment. Let $\psi_n, \phi : \mathcal{H} \to (\mathcal{Y}, d_{\mathcal{Y}})$ be finite-sample and limiting target charts, and set

$$\epsilon_n \coloneqq \sup_{h\in\mathcal{H}} d_{\mathcal{Y}}(\psi_n(h), \phi(h)).$$

For a bounded loss $\ell : \mathcal{D}\times\mathcal{Y} \to [0, M]$, uniformly continuous in its target argument with modulus ω_ℓ, write

$$\mathfrak{R}(\mathcal{E}_n, \psi_n) \coloneqq \inf_\delta \sup_{h\in\mathcal{H}} \mathbb{E}_{P_{n,h}}[\ell(\delta, \psi_n(h))],$$
$$\mathfrak{R}(\mathcal{G}_n, \phi) \coloneqq \inf_\delta \sup_{h\in\mathcal{H}} \mathbb{E}_{Q_{n,h}}[\ell(\delta, \phi(h))],$$

where randomized rules are allowed.

Theorem K.1 (Decision-theoretic transport by Le Cam distance)**.** *For every n,*

$$\boxed{|\mathfrak{R}(\mathcal{E}_n, \psi_n) - \mathfrak{R}(\mathcal{G}_n, \phi)| \le M\Delta_{\mathcal{H}}(\mathcal{E}_n, \mathcal{G}_n) + \omega_\ell(\epsilon_n).} \tag{K.3}$$

Consequently, if $\Delta_{\mathcal{H}}(\mathcal{E}_n, \mathcal{G}_n) \to 0$ and $\epsilon_n \to 0$, then the two minimax risks have the same limit. The same conclusion holds for quadratic loss whenever the finite and limiting target ranges are contained in a common compact convex set, after projection onto that set in the loss norm.

Proof. Fix $\varepsilon > 0$. Choose a Markov kernel $K_{E\to G}$ from the sample space of $\mathcal{E}_n$ to that of $\mathcal{G}_n$ such that

$$\sup_{h\in\mathcal{H}} d_{\mathrm{TV}}(K_{E\to G}P_{n,h}, Q_{n,h}) \le \delta_{\mathcal{H}}(\mathcal{E}_n, \mathcal{G}_n) + \varepsilon. \tag{K.4}$$

Let δ_G be any randomized decision rule in the Gaussian experiment, and let $\delta_E = \delta_G \circ K_{E\to G}$. For every h, boundedness of ℓ and contraction of total variation through the decision kernel give

$$\begin{aligned}\mathbb{E}_{P_{n,h}}\ell(\delta_E, \psi_n(h)) &\le \mathbb{E}_{Q_{n,h}}\ell(\delta_G, \psi_n(h)) + M d_{\mathrm{TV}}(K_{E\to G}P_{n,h}, Q_{n,h}) \\ &\le \mathbb{E}_{Q_{n,h}}\ell(\delta_G, \psi_n(h)) + M\{\delta_{\mathcal{H}}(\mathcal{E}_n, \mathcal{G}_n) + \varepsilon\}. \end{aligned}\tag{K.5}$$

Taking the supremum over h, then the infimum over δ_G, and finally letting $\varepsilon \downarrow 0$ yields

$$\mathfrak{R}(\mathcal{E}_n, \psi_n) \le \mathfrak{R}(\mathcal{G}_n, \psi_n) + M\delta_{\mathcal{H}}(\mathcal{E}_n, \mathcal{G}_n). \tag{K.6}$$

Repeating the argument with a near-optimal kernel $K_{G\to E}$ gives the reverse directed inequality

$$\mathfrak{R}(\mathcal{G}_n, \psi_n) \le \mathfrak{R}(\mathcal{E}_n, \psi_n) + M\delta_{\mathcal{H}}(\mathcal{G}_n, \mathcal{E}_n). \tag{K.7}$$

Therefore

$$|\mathfrak{R}(\mathcal{E}_n, \psi_n) - \mathfrak{R}(\mathcal{G}_n, \psi_n)| \le M\Delta_{\mathcal{H}}(\mathcal{E}_n, \mathcal{G}_n). \tag{K.8}$$

For every decision d and parameter h, uniform continuity in the target argument gives

$$|\ell(d, \psi_n(h)) - \ell(d, \phi(h))| \le \omega_\ell(\epsilon_n).$$

Taking expectations, suprema, and infima shows

$$|\mathfrak{R}(\mathcal{G}_n, \psi_n) - \mathfrak{R}(\mathcal{G}_n, \phi)| \le \omega_\ell(\epsilon_n). \tag{K.9}$$

The triangle inequality proves (K.3).

For quadratic loss, let C be a common closed convex compact set containing the finite and limiting target ranges, and let Π_C^W be metric projection in the loss norm. The projection inequality $\|\Pi_C^W(d) - y\|_W \le \|d - y\|_W$ for every $y \in C$ permits restriction to C. The squared loss is bounded and uniformly continuous on $C \times C$, so the preceding argument applies. □

Corollary K.2 (Conditional transfer under verified two-sided Le Cam equivalence)**.** *Suppose two-sided Le Cam equivalence has been independently verified for a sequence of local observation experiments, uniformly on every compact local parameter set, to the Gaussian shift appearing in section 3. Then all bounded-loss local minimax conclusions in Appendices F–H transfer, including continuous-kink decision limits, jump-branch classification, and unscaled point-strategy risks. The regular quadratic-risk bounds remain valid under the compact-projection or expanding-ball conditions used in corollary F.8 and theorem F.16.*

Remark K.3 (Scope of the non-Gaussian extension). The corollary is an experiment-comparison statement, not an unqualified claim that every non-Gaussian observation model is locally Gaussian. Verifying the Le Cam-distance condition requires model-specific regularity.

K.2 Strongly convex regularized strategy rules

Let $\|\cdot\|_X$ and $\|\cdot\|_Y$ be norms on $\mathbb{R}^m$ and $\mathbb{R}^k$, with dual norms $\|\cdot\|_{X,*}$ and $\|\cdot\|_{Y,*}$. Let $R_X : \mathbb{R}^m \to (-\infty, +\infty]$ and $R_Y : \mathbb{R}^k \to (-\infty, +\infty]$ be proper closed convex functions, finite on neighborhoods of the corresponding simplices, and respectively μ_X- and μ_Y-strongly convex. Put

$$\Omega_X := \max_{x\in\Delta_m} R_X(x) - \min_{x\in\Delta_m} R_X(x), \tag{K.10}$$

$$\Omega_Y := \max_{y\in\Delta_k} R_Y(y) - \min_{y\in\Delta_k} R_Y(y), \tag{K.11}$$

and $\mu := \min\{\mu_X, \mu_Y\}$. For a matrix M, define

$$\|M\|_\diamond := \max\left\{\sup_{y\in\Delta_k}\|My\|_{X,*}, \sup_{x\in\Delta_m}\|M^{\mathsf{T}}x\|_{Y,*}\right\}. \tag{K.12}$$

For $\varepsilon_{\mathrm{reg}} > 0$, let $\mathcal{S}^{\mathrm{sol}}_{\varepsilon_{\mathrm{reg}},R}(A) = (x_{\varepsilon_{\mathrm{reg}},R}(A), y_{\varepsilon_{\mathrm{reg}},R}(A))$ be the saddle point of

$$\mathcal{L}_{A,\varepsilon_{\mathrm{reg}},R}(x, y) = x^{\mathsf{T}}Ay - \varepsilon_{\mathrm{reg}}R_X(x) + \varepsilon_{\mathrm{reg}}R_Y(y). \tag{K.13}$$

Strongly convex prox functions of this type are standard in smoothing frameworks (Nesterov, 2005). Theorem C.8 treats negative entropy separately; it is not a literal specialization of the ambient-neighborhood-finite assumptions used here.

Theorem K.4 (Strongly convex regularization bias–sensitivity theorem)**.** *For every A and $\varepsilon_{\mathrm{reg}} > 0$, the regularized saddle point $\mathcal{S}^{\mathrm{sol}}_{\varepsilon_{\mathrm{reg}},R}(A)$ exists and is unique. For all payoff matrices A, B,*

$$\mathrm{Gap}_A(\mathcal{S}^{\mathrm{sol}}_{\varepsilon_{\mathrm{reg}},R}(A)) \le \varepsilon_{\mathrm{reg}}(\Omega_X + \Omega_Y), \tag{K.14}$$

$$\frac{1}{2}\left(\|x_{\varepsilon_{\mathrm{reg}},R}(A) - x_{\varepsilon_{\mathrm{reg}},R}(B)\|_X + \|y_{\varepsilon_{\mathrm{reg}},R}(A) - y_{\varepsilon_{\mathrm{reg}},R}(B)\|_Y\right) \le \frac{\|A-B\|_\diamond}{\varepsilon_{\mathrm{reg}}\mu}. \tag{K.15}$$

Furthermore, define

$$\mathcal{S}^{\mathrm{sol}}_R(A) \coloneqq \left(\operatorname*{arg\,min}_{x\in\mathcal{X}^\star(A)} R_X(x), \operatorname*{arg\,min}_{y\in\mathcal{Y}^\star(A)} R_Y(y)\right). \tag{K.16}$$

Both minimizers are unique, and

$$\boxed{\mathcal{S}^{\mathrm{sol}}_{\varepsilon_{\mathrm{reg}},R}(A) \longrightarrow \mathcal{S}^{\mathrm{sol}}_R(A) \quad \textit{as } \varepsilon_{\mathrm{reg}} \downarrow 0.} \tag{K.17}$$

Proof. Write $X = \Delta_m$ and $Y = \Delta_k$. Because X and Y are compact convex sets and the objective is continuous, concave in x, and convex in y, Sion's minimax theorem gives a saddle point (Sion, 1958). For fixed y, the function $x \mapsto x^\mathsf{T}Ay - \varepsilon_{\mathrm{reg}}R_X(x)$ is $\varepsilon_{\mathrm{reg}}\mu_X$-strongly concave on X; for fixed x, the function $y \mapsto x^\mathsf{T}Ay + \varepsilon_{\mathrm{reg}}R_Y(y)$ is $\varepsilon_{\mathrm{reg}}\mu_Y$-strongly convex on Y. Two distinct saddle points would contradict strict concavity in the row component or strict convexity in the column component, so the saddle point is unique.

Let $(x_A, y_A) = \mathcal{S}^{\mathrm{sol}}_{\varepsilon_{\mathrm{reg}},R}(A)$. The saddle inequalities say that for every $x \in X$ and $y \in Y$,

$$x^\mathsf{T}Ay_A - \varepsilon_{\mathrm{reg}}R_X(x) \le x_A^\mathsf{T}Ay_A - \varepsilon_{\mathrm{reg}}R_X(x_A), \tag{K.18}$$

$$x_A^\mathsf{T}Ay_A + \varepsilon_{\mathrm{reg}}R_Y(y_A) \le x_A^\mathsf{T}Ay + \varepsilon_{\mathrm{reg}}R_Y(y). \tag{K.19}$$

Rearranging gives

$$x^\mathsf{T}Ay_A - x_A^\mathsf{T}Ay_A \le \varepsilon_{\mathrm{reg}}\{R_X(x) - R_X(x_A)\} \le \varepsilon_{\mathrm{reg}}\Omega_X,$$
$$x_A^\mathsf{T}Ay_A - x_A^\mathsf{T}Ay \le \varepsilon_{\mathrm{reg}}\{R_Y(y) - R_Y(y_A)\} \le \varepsilon_{\mathrm{reg}}\Omega_Y.$$

Taking the maximum over x in the first display and the maximum over the negative column-deviation term in the second proves (K.14).

For stability, extend each regularizer by its simplex indicator: $\overline{R}_X = R_X + \iota_X$ and $\overline{R}_Y = R_Y + \iota_Y$. Fermat optimality gives subgradients

$$g_X^A \in \partial\overline{R}_X(x_A), \qquad Ay_A = \varepsilon_{\mathrm{reg}}g_X^A, \tag{K.20}$$

$$g_Y^A \in \partial\overline{R}_Y(y_A), \qquad -A^\mathsf{T}x_A = \varepsilon_{\mathrm{reg}}g_Y^A, \tag{K.21}$$

and analogous relations for (x_B, y_B). Strong convexity of the extended functions on their domains implies

$$\langle g_X^A - g_X^B, x_A - x_B\rangle \ge \mu_X\|x_A - x_B\|_X^2, \tag{K.22}$$

$$\langle g_Y^A - g_Y^B, y_A - y_B\rangle \ge \mu_Y\|y_A - y_B\|_Y^2. \tag{K.23}$$

Put $\Delta x = x_A - x_B$ and $\Delta y = y_A - y_B$. Pair the differences of (K.20) and (K.21) with Δx and Δy, respectively, and add. The terms $\langle A\Delta y, \Delta x\rangle$ and $-\langle A^\mathsf{T}\Delta x, \Delta y\rangle$ cancel, leaving

$$\begin{aligned}\varepsilon_{\mathrm{reg}}\mu_X\|\Delta x\|_X^2 + \varepsilon_{\mathrm{reg}}\mu_Y\|\Delta y\|_Y^2 &\le \langle (A-B)y_B, \Delta x\rangle - \langle (A-B)^\mathsf{T}x_B, \Delta y\rangle \\ &\le \|A-B\|_\diamond\{\|\Delta x\|_X + \|\Delta y\|_Y\}.\end{aligned} \tag{K.24}$$

With $a = \|\Delta x\|_X$, $b = \|\Delta y\|_Y$, and $\mu = \min\{\mu_X, \mu_Y\}$, the left side is at least $\varepsilon_{\mathrm{reg}}\mu(a^2 + b^2) \ge \varepsilon_{\mathrm{reg}}\mu(a+b)^2/2$. If $a + b > 0$, division proves (K.15); otherwise the claim is immediate.

It remains to identify the zero-temperature limit. Compactness gives accumulation points as $\varepsilon_{\mathrm{reg}} \downarrow 0$, and (K.14) implies that every such point belongs to $\mathcal{E}(A)$. Fix any equilibrium $(x^\star, y^\star) \in \mathcal{E}(A)$. Evaluating (K.18) at $x^\star$ gives

$$\mathrm{val}(A) - \varepsilon_{\mathrm{reg}}R_X(x^\star) \le x_A^\mathsf{T}Ay_A - \varepsilon_{\mathrm{reg}}R_X(x_A), \tag{K.25}$$

because $x^{\star\mathsf{T}} A y_A \geq \mathrm{val}(A)$. Evaluating (K.19) at $y^\star$ gives

$$x_A^\mathsf{T} A y_A + \varepsilon_{\mathrm{reg}} R_Y(y_A) \leq \mathrm{val}(A) + \varepsilon_{\mathrm{reg}} R_Y(y^\star), \tag{K.26}$$

because $x_A^\mathsf{T} A y^\star \leq \mathrm{val}(A)$. Combining the two displays yields

$$R_X(x_A) + R_Y(y_A) \leq R_X(x^\star) + R_Y(y^\star). \tag{K.27}$$

Passing to an accumulation point $(\bar{x}, \bar{y})$ shows that it minimizes $R_X(x) + R_Y(y)$ over $\mathcal{E}(A) = \mathcal{X}^\star(A) \times \mathcal{Y}^\star(A)$. The product structure makes this sum minimization equivalent to the two separate minimizations in (K.16). Strong convexity makes each minimizer unique. Hence every accumulation point equals $\mathcal{S}_R^{\mathrm{sol}}(A)$, which proves convergence of the entire family. □

Remark K.5 (Relation to the entropy theorem). Theorem C.8 treats $R(p) = \sum_i p_i \log p_i$ by a separate argument using interior Gibbs solutions and Pinsker's inequality. Negative entropy, under its standard convex extension, is not finite on an ambient open neighborhood of the simplex and has no finite boundary subgradient at a zero coordinate. Therefore it does not satisfy the neighborhood-finiteness hypothesis of theorem K.4; no literal special-case claim is made.

K.3 Finite multi-branch jump experiments

Let $\mathcal{H}_1, \ldots, \mathcal{H}_J \subset \mathbb{R}^d$ be pairwise disjoint, nonempty compact sets, with $0 \notin \mathcal{H}_j$, such that each normalized payoff direction lies in a compact subset of one full-dimensional fan chamber $\mathcal{C}_j$. Suppose the corresponding branch limits $z_1^0, \ldots, z_J^0$ are pairwise distinct. Put $\mathcal{H} = \bigcup_j \mathcal{H}_j$ and define $b(h) = j$ on $\mathcal{H}_j$. For $Z = h + \Sigma^{1/2} G$, define

$$\mathfrak{R}_{\mathrm{br},G}^{(J)} := \inf_{\delta : \mathbb{R}^d \rightsquigarrow [J]} \sup_{h \in \mathcal{H}} \Pr_h\{\delta(Z) \neq b(h)\}, \tag{K.28}$$

$$\mathfrak{R}_{\mathrm{sel},G}^{(J)} := \inf_{\delta : \mathbb{R}^d \rightsquigarrow \mathcal{Z}} \sup_{h \in \mathcal{H}} \mathbb{E}_h d_\Delta^2(\delta(Z), z_{b(h)}^0). \tag{K.29}$$

Proposition K.6 (Gaussian reduction and pairwise lower bounds for finite multi-branch jumps)**.** *Under the tangent-fan assumptions and the identifiable Gaussian design of assumption F.1, the finite-sample branch and unscaled point-strategy minimax risks over $\mathcal{H}$ converge to the Gaussian risks in* (K.29)*. If*

$$d_{j\ell} = \inf_{h \in \mathcal{H}_j,\, h' \in \mathcal{H}_\ell} \|h - h'\|_J, \qquad \Delta_{j\ell} = d_\Delta(z_j^0, z_\ell^0),$$

then

$$\mathfrak{R}_{\mathrm{br},G}^{(J)} \geq \max_{j \neq \ell} \Phi\left(-\frac{d_{j\ell}}{2}\right), \tag{K.30}$$

$$\mathfrak{R}_{\mathrm{sel},G}^{(J)} \geq \max_{j \neq \ell} \frac{\Delta_{j\ell}^2}{4} \Phi\left(-\frac{d_{j\ell}}{2}\right). \tag{K.31}$$

Proof. The uniform branch expansion from theorem D.9 makes each branch label eventually exact and the point target uniformly close to the piecewise constant map $h \mapsto z_{b(h)}^0$. Bounded-loss minimax convergence follows from theorem F.7. Fix $j \neq \ell$. Compactness implies that the distance $d_{j\ell}$ is attained at points $h_j^\star \in \mathcal{H}_j$ and $h_\ell^\star \in \mathcal{H}_\ell$. Restricting the full multiclass problem to the simple two-point experiment $\{h_j^\star, h_\ell^\star\}$ gives the exact equal-covariance Gaussian error $\Phi(-d_{j\ell}/2)$ and hence the stated lower bound. For strategy loss, convert any estimator to the nearer-target classifier between z_j^0 and z_ℓ^0. On a classification error the triangle inequality forces distance at least $\Delta_{j\ell}/2$, and squared loss is therefore at least $\Delta_{j\ell}^2/4$. Maximizing over pairs proves both displays. No exact composite-class testing formula is asserted. □

Remark K.7 (What remains outside the completed extension theory). The results above do not classify directions on tangent-fan faces, where the first nonzero validity term may be quadratic or of higher order. Such contacts can induce different localization scales—for example, a quadratic validity term balanced against root-n noise naturally leads to an $n^{-1/4}$ parameter scale—and require a separate higher-order theory. They also do not establish uniformity when m, k, or the local dimension grow with the sample size, and they do not replace the generative observation model in the Markov-game separation by dependent trajectory data.

L Detailed Related Work and Novelty Audit

This appendix audits the novelty claims of the manuscript theorem by theorem. The audit was conducted against publicly available literature through 25 September 2026. It is necessarily a scoped review rather than a proof of absolute priority: terminology differs across game theory, variational analysis, statistics, stochastic programming, and reinforcement learning, and no search can exclude unpublished or concurrent work. We therefore avoid blanket "first" claims. The defensible contribution is the three-claim conjunction

$$\boxed{\begin{array}{c}\text{dimension-uniform obstruction to stable accurate point selection}\\ +\ \text{data-driven recovery of unknown simple interfaces with oracle-risk attainment}\\ +\ \text{exact post-hoc target-ignorance duality and compatibility.}\end{array}} \tag{L.1}$$

The certified cross-game generator tests the predicted normalized laws within the theorem domain; it supports Claim 2 but is not a fourth claim. The Bellman material is likewise appendix-level corroboration rather than part of the headline conjunction. No individual ingredient in (L.1) is presented as a new general theory in isolation.

L.1 Audit protocol

For each headline result we distinguish three layers: inherited machinery, adjacent results, and the manuscript-specific statement. Inherited machinery includes Shapley–Snow kernels, strong regularity of generalized equations, polyhedral graphical derivatives, Le Cam comparison, and classical optimal design. Adjacent work proves related sensitivity, inference, query-complexity, or confidence results for different targets or models. The word *singular* is target dependent here: a game may be regular for value and singular for a point strategy or branch under the same experiment. This use is distinct from the algebraic Nash discriminant of Abo et al. (2026).

L.2 Deterministic equilibrium geometry and selection

The branch representation in Appendix D rests on Shapley–Snow basic solutions and modern kernel formulations (Shapley & Snow, 1950; Attia & Oliu-Barton, 2021). The generalized-equation analysis uses Robinson strong regularity and polyhedral proto-derivatives (Robinson, 1980; Rockafellar & Wets, 1998; Dontchev & Rockafellar, 1996; 2001). Recent Nash-equilibrium stability work characterizes strong regularity, differentiable single-valued localization, and isolated calmness for Nash KKT mappings (Diao et al., 2026). We therefore do not claim to introduce strong regularity or differentiable equilibrium localization.

The manuscript-specific use of this machinery is statistical. Theorem D.9 builds an explicit finite branch atlas with exact validity inequalities and a conical first-order partition of payoff directions; Theorem E.8 identifies the chamber derivative with the critical-cone system used as a target map in the local experiments of Sections 3– 3.

The value and fixed-strategy exploitability Lipschitz bounds are elementary and not claimed as new. Recent work also contrasts Lipschitz exploitability with discontinuous vertex-returning LP selectors (Nie et al., 2026). QRE and entropy-regularized algorithms are established (McKelvey & Palfrey, 1995; Cen et al., 2021), and Nguyen et al. (2026) prove the tight $\Theta(\beta^{-1}(\log m + \log k))$ finite-game QRE approximation order. We do not claim that rate. The retained theorem is instead the dimension-uniform lower bound on the worst-case equilibrium error of every globally Lipschitz single-valued strategy rule. Entropy supplies only one upper construction with the same inverse-sensitivity dependence, up to its logarithmic dimension factor.

L.3 Nonsmooth inference and random optimization

The regular covariance bound belongs to the Hájek–Le Cam local asymptotic minimax tradition (Hájek, 1972; Le Cam, 1986). Directional delta methods and bootstrap validity for nonsmooth maps are developed by Fang & Santos (2019); generalized van Trees inequalities yield sharp bounds for irregular functionals (Takatsu & Kuchibhotla, 2024); and recent minimax-optimization asymptotics obtain non-Gaussian limits for optimal values and solution estimators (Meitz & Shapiro, 2025).

For random linear programs, Liu et al. (2023) derive solution-set limits and confidence sets, while Lu (2014) develops confidence regions for stochastic variational inequalities. The novelty claim is therefore not the existence of a piecewise-normal or non-Gaussian optimizer limit.

Theorem F.7 instead characterizes the limiting minimax decision problem for several equilibrium targets under one common experiment:

$$\begin{gathered}\text{value}, \quad \text{low exploitability}, \quad \text{point strategy}, \\ \text{interface-side label / local branch germ}, \quad \text{support label}, \quad \text{equilibrium set}.\end{gathered} \tag{L.2}$$

Theorem G.3 gives an exact composite minimax risk for the linearized interface-side label over fixed half-spaces. On the induced least-favorable two-point normal experiment it gives exact least-favorable-pair local-branch-germ and point-strategy risks, while Corollary H.7 gives the exact two-point minimax expected diameter of an honest strategy region. Over larger equilibrium-learning neighborhoods, the two-point quantities are lower bounds. The claim is thus a loss-conditioned equilibrium phase diagram, not a general claim to have discovered nonsmooth inference.

L.4 Inference for equilibrium models and noisy games

Statistical inference is already available for important equilibrium models. For Fisher markets and pacing equilibria, Liao & Kroer (2025) derive asymptotic distributions, confidence intervals, local minimax optimality, and A/B-testing procedures. In performative prediction, Zhang et al. (2026) study asymptotic normality and efficiency in single- and multiplayer models. Hence equilibrium inference itself is not new. The distinction here is the analysis of zero-sum saddle targets exactly where point strategies or supports cease to be continuous, together with a comparison of several targets under the same observation law.

Noisy zero-sum game learning also has a substantial literature. Maiti et al. (2023) prove instance-dependent bounds for approximate equilibria; Maiti et al. (2024) study pure-equilibrium identification; and Jiang & Zhang (2026) combine support identification with restricted LP estimation in potentially nonunique games. Hadiji et al. (2023) study first-order payoff-query complexity for exact and approximate equilibria, whereas Maiti et al. (2025) study global noiseless query complexity for recovering the entire equilibrium set. Our observation model is noisy entrywise sampling and our question is local and loss indexed. Proposition B.3 is not presented as the first noisy-game result, nor is Theorem G.3 a new global query-complexity theorem.

L.5 Confidence regions and experimental design

Testing, local moduli, and honest confidence-set diameter are classically linked (Donoho, 1994; Low, 1997; Robins & van der Vaart, 2006). Confidence regions for stochastic VIs and random linear programs are also established (Lu, 2014; Liu et al., 2023). Appendix H therefore does not claim the first equilibrium confidence region. Its narrower conclusions are an exact finite-sample expected-diameter frontier for the declared two-point Gaussian jump experiment, an attaining three-decision region, simultaneous projection-based coverage of the whole zero-sum equilibrium set, and the direct contrast between root-n value intervals and order-one honest point-strategy regions.

The convex design criterion and sensitivity equalization in Section 5 descend from Kiefer–Wolfowitz and Elfving theory (Kiefer & Wolfowitz, 1960; Elfving, 1952); sequential information allocation has a long history (Chernoff, 1959; Kaufmann et al., 2016; Garivier & Kaufmann, 2016). The equilibrium-specific claim is that different targets expose different design directions: equilibrium occupancy for value and the bifurcation normal for interface-side classification. Proposition I.6 gives one declared family in which the local-oracle value and interface-side objectives select opposite endpoints of the sampling frontier.

L.6 Supporting Bellman appendix, not a headline contribution

Target-specific point-policy nonregularity is not unique to games. Dynamic-treatment-regime parameters are nonsmooth near treatment ties (Laber et al., 2010), recent reinforcement-learning work gives debiased asymptotically normal inference for the optimal value (Lu et al., 2026), and directional derivatives of stochastic-game values are studied by Attia et al. (2025). We therefore make no broad priority claim for dynamic nonregularity.

Appendix J proves one supporting statement: statewise equilibrium singularities propagate through finite-horizon Bellman recursion under sectorwise separation. The result is finite-horizon, finite-state/action, generative-oracle based, and metric qualified. Its low-reach example has zero initial-value risk, $O(p_{\mathrm{reach}} n^{-1/2})$ exploitability, and $\Theta(1)$ point-policy risk only under the declared squared uniform sup-state policy metric; under an occupancy-weighted metric the low-reach contribution generally shrinks with p_{reach}. This result appears neither in the title nor the abstract and is not one of the paper's three claims. A trajectory-data theorem coupling reachability, exploration, and target-specific information remains outside the paper.

L.7 Claim-by-claim summary

Claim 1. *Stable decisions do not imply stable strategies.* Relative to residual stability, discontinuous LP selectors, and QRE, the retained claim is a dimension-uniform lower obstruction for every globally Lipschitz single-valued strategy rule, together with the exact value/gap versus branch/point separation. Entropy is only an upper witness for the inverse-sensitivity dependence and is not claimed as a new QRE-rate result.

Claim 2. *Unknown simple interfaces are adaptively learnable at oracle risk.* Relative to smooth boundary learning, adaptive design, local asymptotic minimax theory, and noisy-game identification, the retained claim is an implementable finite-atlas procedure with uniform geometry recovery, uniform Gaussian-score convergence, oracle-efficient side and declared least-favorable-pair point/confidence risks, matching lower bounds, and an explicit adaptation-impossibility boundary. The certified corpus checks the dimension-free normalized laws inside this fixed quantitative domain; it is not evidence of global prevalence.

Claim 3. *There is no target-free experiment.* Relative to classical and multi-objective optimal design, the retained claim is the post-hoc equilibrium-target protocol, its exact attained minimax dual, the compatibility equivalence $\Gamma = 1$, and the sharp disjoint-support value $\Gamma = R$ for game-derived target profiles.

L.8 Safe novelty language and excluded claims

Supported formulations are:

1. we prove a dimension-uniform accuracy–sensitivity obstruction for globally Lipschitz single-valued strategy rules and an exact separation between stable decision targets and singular point/branch targets;
2. we give an implementable procedure that learns an unknown simple equilibrium interface and uniformly attains the declared known-geometry local minimax benchmarks, with matching adaptive lower bounds and an exact competing-interface boundary;
3. we derive the exact price of choosing an experiment before the downstream equilibrium target is known, including the compatibility and disjoint-support characterizations; and
4. a certified generator and fixed-seed study verify implementation of the predicted cross-game normalized laws within the proved quantitative domain.

The manuscript should not claim a new tight QRE approximation rate, global prevalence of simple interfaces, exact finite-sample adaptive honesty, a point minimax theorem beyond the declared information-balanced pair, a general dynamic equilibrium fan, or a trajectory-data Markov-game theory. It also does not provide a complete theory of general-sum games, higher-order singularities, growing dimensions, or partial observation.

Audit conclusion. The literature contains close precedents for the individual tools. The headline novelty is therefore exactly the three-claim structure above: a general point-selection obstruction, a non-oracle adaptive-interface theorem with matching minimax limits, and an exact post-hoc target-ignorance theorem. Entropy and the finite-horizon Bellman result are supporting constructions, not coequal contributions. Any change to a headline theorem, loss, observation model, or uniformity domain should trigger a new literature search and an update to this audit.

M WORKED EXAMPLES AND COMPUTATIONAL VERIFICATION

This appendix supplies the examples, figures, and anonymous verification code used to stress-test the theorem package. The computations are not substitutes for the proofs in Appendices B–K; each example below is solved analytically first, and the code independently checks the resulting formulas by linear programming, kernel enumeration, convex design optimization, Monte Carlo integration, and finite differences. The complete package is in the supplementary directory `reproducibility/`.

Verification categories. The supplement separates (i) exact symbolic or rational identities and finite kernel counts, (ii) deterministic numerical optimization and linear-program cross-checks with declared tolerances, and (iii) fixed-seed Monte Carlo sampling used only as a diagnostic. None of these computations replaces an analytic proof, and complex-step differentiation is applied only to smooth branch formulas away from validity tests, max/min operations, and rank-change boundaries.

M.1 PRACTICAL FINITE VERIFICATION ON THE WORKED EXAMPLES

Table 1 records exact finite enumeration sizes and a reference implementation runtime. Candidate and incident-kernel counts, active-form counts, and chamber counts are mathematical properties of the declared examples; runtime is only a machine-dependent diagnostic. The reference audit used Python 3.13.5, NumPy 2.3.5, SciPy 1.17.0 with the HiGHS linear-program solver, one process on a five-core Intel Xeon Platinum 8573C, and reports the median of 200 runs after eight warm-up runs.

Table 1: Practical finite verification of the worked examples. "Cells/germs" reports strict arrangement cells followed by distinct branch regions; for the Bellman example it reports one strict statewise cell at each of three games. No displayed example requires general real quantifier elimination.

Example	Cand.	Inc.	Forms	Cells/germs	Completeness method	ms
2×2 jump	5	4	1	$2/2$	direct exact family reduction	4.68
3×3 three-way	19	3	3	$6/3$	exact column-dominance reduction	15.61
3×3 four-way	19	4	2	$4/4$	direct duplicate-action atlas	10.48
Bellman feedback	15	3	0	$3/3$	statewise strict regularity	2.40

The candidate counts use (D.36). The incident-kernel, active-form, and chamber counts follow respectively from lemma B.2 and propositions M.1 and D.18 and the strict statewise systems in Appendix J. The first three rows enumerate all square support pairs of the static game. The Bellman row aggregates the three statewise 2×2 games. The largest exact input/displayed rational denominators are respectively $1/1$, $1/1$, $2/16$, and $100/1600$. The first two families are certified by exact direct reductions, the four-way family by an exact analytic branch atlas, and the Bellman example by strict statewise bordered systems. General completeness is still a semialgebraic decision problem; the table shows that the finite procedure is operational on the examples claimed, not that it has polynomial complexity. Exact counts and regenerated timings are written to `outputs/checkability_report.json` and `outputs/checkability_report.md`.

M.2 A FOUR-BRANCH MIXED-EQUILIBRIUM 3×3 FAMILY

The pure three-way example in proposition D.18 verifies the branch assumptions but has zero branch derivatives. The following family has four incident regular kernels, nonzero strategy derivatives, a four-class statistical problem, and incompatible value and branch designs. Define

$$A_0 = \begin{pmatrix} 1 & 1 & -1 \\ 1 & 1 & -1 \\ -1 & -1 & 1 \end{pmatrix}, \tag{M.1}$$

$$H_u = \frac{1}{2}\begin{pmatrix} 1 & 1 & 1 \\ -1 & -1 & -1 \\ 0 & 0 & 0 \end{pmatrix}, \tag{M.2}$$

$$H_v = \frac{1}{2}\begin{pmatrix} 1 & -1 & 0 \\ 1 & -1 & 0 \\ 1 & -1 & 0 \end{pmatrix}, \tag{M.3}$$

with

$$A(u,v) = A_0 + uH_u + vH_v. \tag{M.4}$$

For $s,t \in \{-1,+1\}$, let

$$i_s = \begin{cases} 1, & s=+1, \\ 2, & s=-1, \end{cases} \qquad j_t = \begin{cases} 2, & t=+1, \\ 1, & t=-1, \end{cases}$$

and let $\beta_{s,t}$ be the square kernel with row support $\{i_s,3\}$ and column support $\{j_t,3\}$.

Proposition M.1 (Exact four-branch mixed-equilibrium geometry)**.** *Assume $|u|,|v|<1$.*

1. *If $uv \neq 0$, $A(u,v)$ has a unique equilibrium. With $s=\operatorname{sgn}(u)$ and $t=\operatorname{sgn}(v)$, it is the branch $\beta_{s,t}$ and equals*

$$x_{i_s}(u,v) = \frac{1}{2} + \frac{tv}{8}, \qquad x_3(u,v) = \frac{1}{2} - \frac{tv}{8}, \tag{M.5}$$

$$y_{j_t}(u,v) = \frac{1}{2} - \frac{su}{8}, \qquad y_3(u,v) = \frac{1}{2} + \frac{su}{8}, \tag{M.6}$$

   with all other coordinates zero, and

$$\operatorname{val}(A(u,v)) = \frac{|u|}{4} - \frac{|v|}{4} + \frac{|u||v|}{16}. \tag{M.7}$$

2. *At $u=0$, the two duplicate row branches remain tied; at $v=0$, the two duplicate column branches remain tied. At the origin,*

$$\mathcal{X}^\star(A_0) = \left\{ x \in \Delta_3 : x_1 + x_2 = \frac{1}{2},\ x_3 = \frac{1}{2} \right\}, \tag{M.8}$$

$$\mathcal{Y}^\star(A_0) = \left\{ y \in \Delta_3 : y_1 + y_2 = \frac{1}{2},\ y_3 = \frac{1}{2} \right\}. \tag{M.9}$$

3. *The four incident extreme equilibrium pairs are represented by the four kernels $\beta_{s,t}$. Every bordered determinant equals four. The only active validity coordinates at A_0 are the off-support duplicate slacks, whose exact values on branch $\beta_{s,t}$ are*

$$\rho^{\mathrm{row}}_{s,t}(u,v) = su, \qquad \rho^{\mathrm{col}}_{s,t}(u,v) = tv. \tag{M.10}$$

   Hence the regular-kernel validity fan consists of the four quadrants

$$\mathcal{C}_{s,t} = \{(u,v) : su > 0,\ tv > 0\}. \tag{M.11}$$

4. *On $\mathcal{C}_{s,t}$, the branch derivative at the origin is*

$$Dx_{s,t}(0)[\dot u,\dot v] = \frac{t\dot v}{8}(e_{i_s} - e_3), \tag{M.12}$$

$$Dy_{s,t}(0)[\dot u,\dot v] = \frac{s\dot u}{8}(e_3 - e_{j_t}), \tag{M.13}$$

$$Dv_{s,t}(0)[\dot u,\dot v] = \frac{s\dot u - t\dot v}{4}. \tag{M.14}$$

Thus the theory is not restricted to a single 2×2 pure-action jump: the origin has four mixed-strategy branch germs, nonzero sectorwise derivatives, and two independent interface normals. The four-label Gaussian experiment below nevertheless factorizes into two symmetric sign coordinates.

Proof. The first two rows satisfy

$$A(u,v)_{1,\cdot} - A(u,v)_{2,\cdot} = u\mathbf{1}_3^{\mathsf{T}}.$$

Hence row one strictly dominates row two when $u>0$, and row two strictly dominates row one when $u<0$. Similarly,

$$A(u,v)_{\cdot,1} - A(u,v)_{\cdot,2} = v\mathbf{1}_3,$$

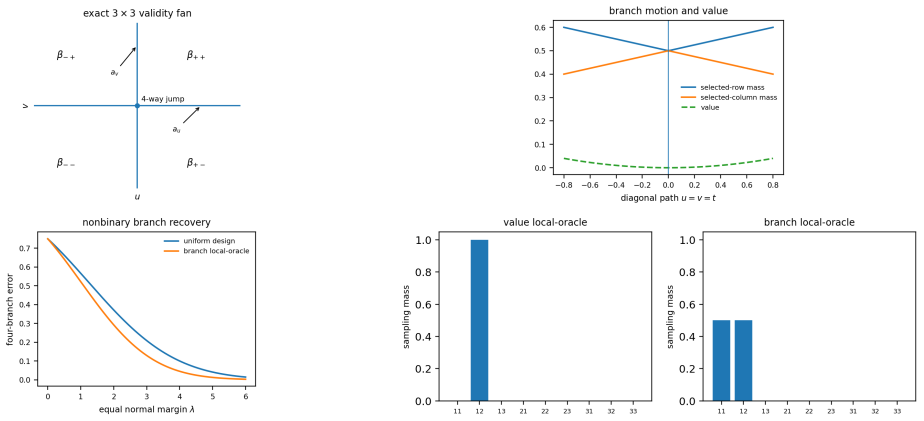

Figure 5: The exact four-branch 3×3 example. The validity fan has four chambers and mixed branch limits. Along $u = v = t$, both strategy probabilities move at first order while the value has a quadratic leading term. The exact four-class branch error improves under the branch local-oracle design. The value local-oracle samples one payoff entry, whereas the branch local-oracle splits mass across two orthogonal entry features.

so the minimizing player excludes column one when $v > 0$ and excludes column two when $v < 0$. After deleting the dominated duplicate row and column, the selected 2×2 subgame is

$$B_{s,t}(u, v) = \begin{pmatrix} 1 + r + c & -1 + r \\ -1 + c & 1 \end{pmatrix}, \qquad r = \frac{su}{2} = \frac{|u|}{2}, \quad c = -\frac{tv}{2} = -\frac{|v|}{2}. \tag{M.15}$$

Its 2×2 saddle denominator is identically four. The standard formulas therefore give

$$x_{i_s} = \frac{1 - (-1 + c)}{4} = \frac{1}{2} - \frac{c}{4}, \qquad y_{j_t} = \frac{1 - (-1 + r)}{4} = \frac{1}{2} - \frac{r}{4},$$

which are (M.5)– (M.6). The value is

$$\frac{(1 + r + c) - (-1 + r)(-1 + c)}{4} = \frac{r + c}{2} - \frac{rc}{4},$$

which is (M.7). The remaining duplicate row differs by exactly $|u|$ against every column mixture, and the remaining duplicate column differs by exactly $|v|$ against every row mixture. This proves (M.10), uniqueness in the strict quadrants, and the axis/origin equilibrium sets. The bordered determinant equals the saddle denominator four. Differentiating the closed forms proves (M.14). □

M.3 Four-branch local risks and incompatible local-oracle designs

Assume unit observation variance and let $f_{ij} = (H_u(i, j), H_v(i, j))^\mathsf{T}$ be the entry feature. For an approximate design q, its local information is

$$J(q) = \sum_{i,j} q_{ij} f_{ij} f_{ij}^\mathsf{T}. \tag{M.16}$$

Uniform sampling gives

$$J(q^{\text{unif}}) = \frac{1}{6} \mathbf{I}_2. \tag{M.17}$$

Proposition M.2 (Exact four-branch classification risk)**.** *Let an observation be*

$$Z = (Z_u, Z_v), \qquad Z_u \sim \mathcal{N}(s\lambda_u, s_u^2), \quad Z_v \sim \mathcal{N}(t\lambda_v, s_v^2),$$

independently, with $(s, t) \in \{-1, +1\}^2$*. Under* 0–1 *loss for the full four-branch label, the coordinatewise sign rule is minimax and*

$$\boxed{R_4^\star = 1 - \Phi\left(\frac{\lambda_u}{s_u}\right) \Phi\left(\frac{\lambda_v}{s_v}\right).} \tag{M.18}$$

Consequently, for equal margins $\lambda_u = \lambda_v = \lambda$,

$$R^\star_{4,\mathrm{unif}}(\lambda) = 1 - \Phi(\lambda/\sqrt{6})^2, \tag{M.19}$$

$$R^\star_{4,\mathrm{br}}(\lambda) = 1 - \Phi(\lambda/2)^2 \tag{M.20}$$

under the uniform and branch local-oracle designs below.

Proof. Under the uniform prior on the four labels, the likelihood factorizes into two independent symmetric Gaussian binary experiments. The Bayes classifier is therefore the coordinatewise sign rule, with correct-classification probability equal to the product of the two one-coordinate correct probabilities. The sign rule has the same risk under all four labels. By the standard equalizer-Bayes argument, it is minimax, proving (M.18). The two displayed specializations follow from (M.17) and (M.24) below. □

Fix the chamber $u > 0, v > 0$. The value derivative is

$$g_{\mathrm{val}} = \left(\frac{1}{4}, -\frac{1}{4}\right)^{\mathsf{T}}. \tag{M.21}$$

The two branch normals are e_1 and e_2.

Proposition M.3 (Exact value-versus-branch design conflict in the 3×3 family)**.** *For the score experiment induced by* (M.16)*:*

1. *A value local-oracle design puts unit mass on entry $(1, 2)$, whose feature is $(1/2, -1/2)^{\mathsf{T}}$. Its target variance is*
$$\inf_q g_{\mathrm{val}}^{\mathsf{T}} J(q)^\dagger g_{\mathrm{val}} = \frac{1}{4}. \tag{M.22}$$

2. *For the symmetric branch criterion*
$$\Psi_{\mathrm{br}}(q) = \max\{e_1^{\mathsf{T}} J(q)^\dagger e_1, e_2^{\mathsf{T}} J(q)^\dagger e_2\}, \tag{M.23}$$
with value $+\infty$ when either normal is unidentified, one optimum puts mass $1/2$ on entries $(1, 1)$ and $(1, 2)$. It has
$$J(q^{\star,\mathrm{br}}) = \frac{1}{4}\mathbf{I}_2, \qquad \inf_q \Psi_{\mathrm{br}}(q) = 4. \tag{M.24}$$

3. *The value design has rank-one information and does not identify both branch normals. Thus it has infinite value under* (M.23)*.*

Proof. Let f_{ij} denote the nine entry features. For the value lower bound, set $z = (1, -1)^{\mathsf{T}}$. Whenever g_{val} is identifiable, generalized Cauchy–Schwarz gives

$$(z^{\mathsf{T}} g_{\mathrm{val}})^2 \le (g_{\mathrm{val}}^{\mathsf{T}} J(q)^\dagger g_{\mathrm{val}})(z^{\mathsf{T}} J(q) z).$$

Now $z^{\mathsf{T}} g_{\mathrm{val}} = 1/2$ and $|z^{\mathsf{T}} f_{ij}| \le 1$, so $z^{\mathsf{T}} J(q) z \le 1$. Hence every design has value variance at least $1/4$. Entry $(1, 2)$ has $f_{12} = (1/2, -1/2)^{\mathsf{T}} = 2g_{\mathrm{val}}$, and observing only that entry attains $1/4$.

For the branch lower bound, a singular $J(q)$ cannot contain both independent normals in its range. If $J(q) \succ 0$, then

$$\max_i e_i^{\mathsf{T}} J(q)^{-1} e_i \ge \frac{1}{2}\operatorname{tr}(J(q)^{-1}) \ge \frac{2}{\operatorname{tr} J(q)}.$$

Every atomic feature has squared Euclidean norm at most $1/2$, so $\operatorname{tr} J(q) \le 1/2$, proving the lower bound four. The features of entries $(1, 1)$ and $(1, 2)$ are $(1/2, 1/2)$ and $(1/2, -1/2)$; equal mass on them gives $J = \mathbf{I}_2/4$, attaining the bound. The value design has range $\operatorname{span}\{(1, -1)\}$ and therefore fails to identify e_1 and e_2 separately. □

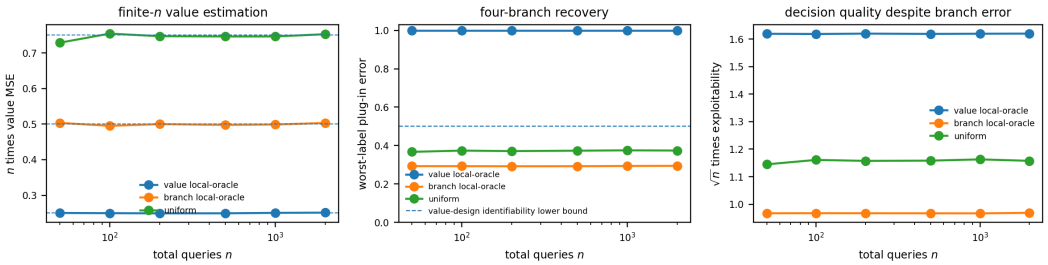

Figure 6: Finite-n verification in the four-branch 3×3 family. Dashed horizontal lines in the value panel are the exact asymptotic constants in (M.25). The middle panel compares branch recovery under the three designs; the dashed line is the nonidentifiability lower bound for the value design. The right panel shows root-n decision quality even when point-branch recovery fails. Each point uses 60,000 Gaussian score replications and deterministic integer rounding of the approximate design.

M.4 Finite-n target comparison in the 3×3 game

To verify that the asymptotic constants are visible at moderate sample sizes, the code simulates the exact Gaussian score experiment associated with integer-rounded query counts. For value estimation, the true local parameter is $h = (2, 1)$ in the chamber $u > 0, v > 0$. The predicted scaled mean-squared-error constants are

$$\frac{1}{4} \quad \text{(value local-oracle)}, \qquad \frac{1}{2} \quad \text{(branch local-oracle)}, \qquad \frac{3}{4} \quad \text{(uniform)}. \tag{M.25}$$

For branch recovery, the four local parameters are $h_{s,t} = (2s, 2t)$, and the displayed statistic is the worst-label error of the minimum-norm score plug-in rule. The branch and uniform designs have exact minimax sign-rule errors

$$1 - \Phi(1)^2 \approx 0.2921, \qquad 1 - \Phi(2/\sqrt{6})^2 \approx 0.3713, \tag{M.26}$$

respectively. The value design has rank-one information and does not identify both normals; its four-label plug-in error is one in this symmetric experiment, while every uniformly valid branch procedure has worst-case error at least $1/2$ on a pair of statistically indistinguishable labels.

For decision quality, the score plug-in is converted to the exact branch strategy of the estimated local game and evaluated under the true payoff matrix. The right panel of figure 6 shows that $\sqrt{n}$ times expected exploitability stabilizes even when branch recovery does not. This plot is a fixed-seed Monte Carlo diagnostic, not an additional minimax theorem.

M.5 Excluded mechanisms as an exact failure suite

The five panels of figure 7 correspond exactly to the examples already proved in examples J.6 and D.19 to D.21. For completeness, their scalar diagnostics are:

$$\det M_{\mathrm{RPS}}(t) = 9t^2, \tag{M.27}$$

$$\rho_{\mathrm{quad}}(t) = -\frac{t^2}{4 + t}, \tag{M.28}$$

$$\dot{Q}_1 = \tfrac{1}{2}(1) + \tfrac{1}{2}(-1) = 0. \tag{M.29}$$

The fan-face and multi-branch panels use the exact arrangements $u = 0$, $v = 0$, and $u = v$ from proposition D.18. Thus every exclusion in the first-order theory is accompanied by a concrete family rather than only a verbal limitation.

M.6 A Bellman finite-difference verification with transition feedback

The main Markov corollary deliberately uses a transparent downstream embedding. The following separate check verifies the full directional recursion when both transition derivatives and downstream value derivatives feed into an upstream matrix game.

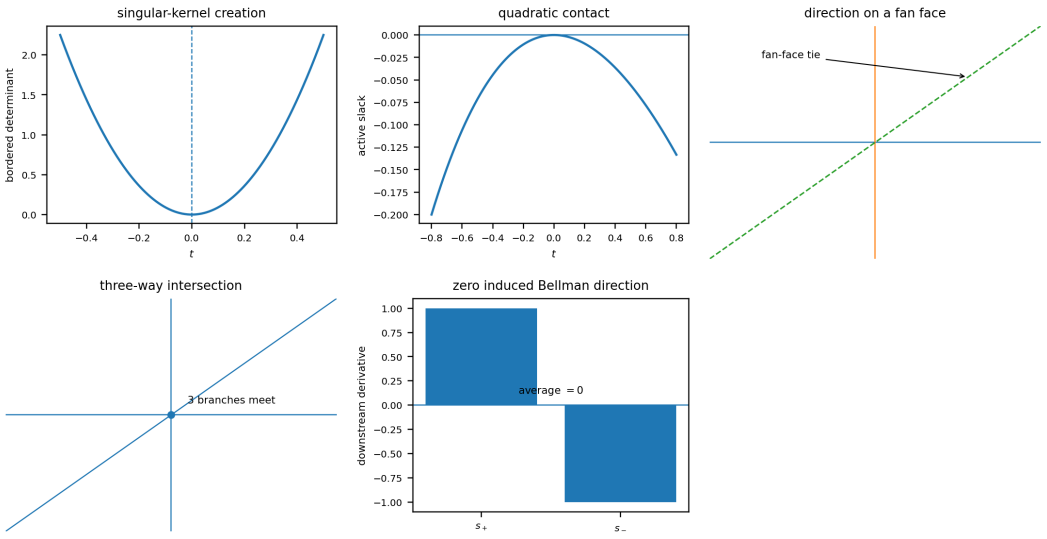

Figure 7: Five excluded mechanisms, each verified by a closed-form family: a bordered determinant that vanishes at the reference game but is nonzero nearby; a validity slack with zero first derivative and negative quadratic term; a direction on a fan face; a three-way branch intersection; and a nonzero primitive Markov perturbation whose induced upstream Bellman direction is zero.

At the two terminal states, define

$$A_A(\theta) = \frac{1}{2}\mathbf{1}_2\mathbf{1}_2^\top + \begin{pmatrix} 1+\theta & -1 \\ -1 & 1 \end{pmatrix}, \tag{M.30}$$

$$A_B(\theta) = -\frac{1}{4}\mathbf{1}_2\mathbf{1}_2^\top + \begin{pmatrix} 1 & -1 \\ -1 & 1+2\theta \end{pmatrix}. \tag{M.31}$$

Their values are

$$V_A(\theta) = \frac{1}{2} + \frac{\theta}{4+\theta}, \qquad V_B(\theta) = -\frac{1}{4} + \frac{\theta}{2+\theta}, \tag{M.32}$$

so $\dot{V}_A(0) = 1/4$ and $\dot{V}_B(0) = 1/2$. Let the action-dependent probability of reaching state A be

$$P_A(\theta) = \begin{pmatrix} 4/5 & 3/10 \\ 2/5 & 7/10 \end{pmatrix} + \theta \begin{pmatrix} 1/20 & -1/25 \\ 3/100 & -1/50 \end{pmatrix}, \tag{M.33}$$

with the complementary probability assigned to state B. Choose the stage-one reward so that $Q_1(0)$ is matching pennies. The Bellman derivative is then

$$\dot{Q}_1(0) = \begin{pmatrix} 27/80 & 79/200 \\ 169/400 & 31/100 \end{pmatrix}. \tag{M.34}$$

Proposition M.4 (Exact derivative in the transition-feedback example)**.** *At the stage-one reference game,*

$$\dot{V}_1(0) = \frac{293}{800}, \qquad \dot{x}_1(0) = \left(-\frac{11}{1600}, \frac{11}{1600}\right), \qquad \dot{y}_1(0) = (0,0). \tag{M.35}$$

The centered finite-difference errors for value and strategy derivatives are $O(\epsilon^2)$.

Proof. Differentiate the Bellman product:

$$\dot{Q}_1 = \dot{P}_A(V_A(0) - V_B(0)) + P_A(0)\dot{V}_A(0) + (1 - P_A(0))\dot{V}_B(0),$$

which gives (M.34). Apply the elementary 2×2 interior-equilibrium formulas to $Q_1(0) + \theta\dot{Q}_1(0)$. Direct differentiation yields (M.35). All primitive maps and the regular 2×2 branch are analytic near zero, so centered differences have the standard $O(\epsilon^2)$ remainder. □

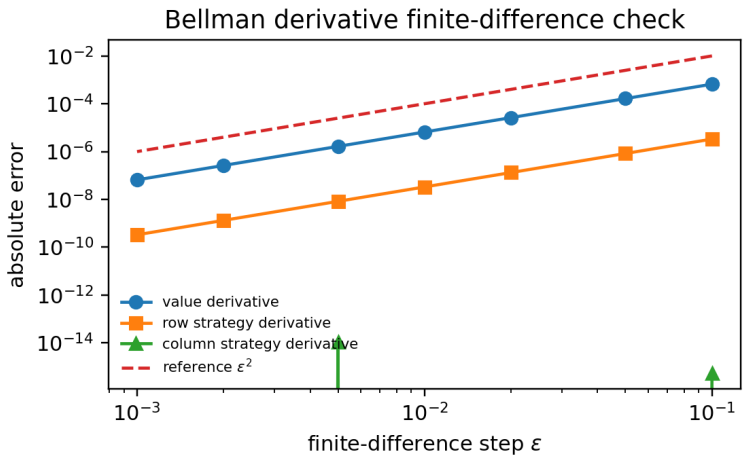

Figure 8: Centered finite-difference verification of proposition M.4. The slopes agree with the predicted quadratic truncation error.

M.7 Executable verification protocol

The anonymous package contains the following independently executable modules.

`kernels.py`. Enumerates square kernels, solves both bordered systems, computes every mass/slack validity coordinate, differentiates validity by complex step, and enumerates strict hyperplane chambers.

`zero_sum.py`. Solves independent primal and dual linear programs and evaluates exploitability.

`risk.py`. Reproduces interface-side, point-strategy, scaled-gap, confidence, continuous-kink, and four-class risk formulas.

`design.py`. Builds semidefinite information matrices, solves target-design programs, and implements deterministic largest-remainder rounding.

`bellman.py`. Evaluates the transition-feedback Markov example and its finite differences.

`checkability.py`. Enumerates the worked examples, counts incident kernels and active forms, constructs their strict chambers, and records environment-specific reference timings.

`verification.py`. Runs all symbolic, LP, convex-optimization, Monte Carlo, failure-mode, and Bellman checks and writes machine-readable reports.

From the supplement directory, the single command `python verify_package.py` recreates a clean output directory, runs the full deterministic test suite, regenerates every figure, checks the recorded tolerances, and writes a machine-readable clean-room report. The archived offline run exits successfully. The archived reference clean-room environment uses Python 3.13.5 on Linux x86-64 with pinned dependencies from `requirements.txt`. On the stated five-core Xeon reference instance, 17 tests pass and the complete offline run takes 20.32 seconds; both quantities are regenerated and may vary with hardware and solver build. Runtime, hardware, fixed seeds, expected output files, and numerical tolerances are recorded in `outputs/clean_room_report.json` and `EXPECTED_TOLERANCES.json`. Complex-step differentiation is applied only within a fixed regular analytic kernel chart, never through a nonsmooth branch decision. Selected numerical residuals are

Check	Maximum observed discrepancy
3×3 closed form versus independent LP strategies	1.67×10^{-16}
3×3 closed form versus LP value	1.71×10^{-16}
Quadratic-contact slack identity	8.24×10^{-17}
Numerical value-design optimum versus $1/4$	1.49×10^{-8}
Numerical branch-design optimum versus 4	1.20×10^{-7}
Two-point interface-side Monte Carlo versus exact curve	7.65×10^{-4}
Two-point point-strategy Monte Carlo versus exact curve	2.46×10^{-3}
Smallest-step Bellman value-derivative error	6.58×10^{-8}

These calculations are diagnostics for signs, constants, and implementation; they are not used as evidence in place of any theorem proof.

Boundary and rank-change unit tests. The clean-room suite contains explicit regression tests for every boundary mechanism requested in the proof audit:

1. a bordered determinant changing sign and a sequence of nearly singular bordered systems;
2. an active validity coordinate with zero first derivative at quadratic contact;
3. two distinct regular kernels representing one analytic equilibrium germ;
4. pseudoinverse range conditions before and after an information-rank change;
5. empirical coverage and expected diameter of the measurable three-report confidence rule;
6. fan-face, singular-kernel, and zero-induced-Bellman-direction failure cases.

Complex-step differentiation is confined to a fixed nonsingular bordered system and differentiates only its analytic mass/slack formulas. No complex step is taken through kernel enumeration, determinant thresholding, max/min, branch selection, or pseudoinverse rank changes.

N Quantitative non-oracle equilibrium-interface classes

This appendix replaces the instance-specific geometry oracle by a finite, instance-independent branch atlas. The learner knows the dimensions, the entrywise noise variances, the atlas construction, and the numerical class constants below. It is not given the true interface game, the active chart, the incident equilibrium branches, the interface normal, the side targets, or the target-optimal design.

Identify $\mathbb{R}^{m\times k}$ with $\mathbb{R}^d$, $d = mk$, through a fixed vectorization, and use the Frobenius inner product. Let $\mathcal{I} = [m] \times [k]$ denote the payoff coordinates. Assume the known noise profile satisfies

$$0 < \underline{\sigma} \le \sigma_i \le \overline{\sigma} < \infty, \qquad i \in \mathcal{I}. \tag{N.1}$$

N.1 Finite analytic atlas and observable diagnostic scores

For fixed m, k, enumerate every ordered tuple

$$\xi = (\beta_+, \ell_+, \beta_-, \ell_-) \tag{N.2}$$

consisting of two canonically represented regular-kernel germs and one coordinate from each validity vector. Duplicate analytic branch germs are identified as in section D. Let Ξ denote the resulting finite set. Every $\xi \in \Xi$ has an open regularity domain $\mathcal{O}_\xi$, analytic branch maps $z_{\xi,+}, z_{\xi,-}$, and active validity coordinates

$$r_{\xi,+}(A) = [\rho_{\beta_+}(A)]_{\ell_+}, \qquad r_{\xi,-}(A) = [\rho_{\beta_-}(A)]_{\ell_-}. \tag{N.3}$$

The canonical orientation is fixed once and for all by the deterministic order on germ classes and validity coordinates.

For each ξ, let $\mathcal{C}_\xi$ be the closed semialgebraic set of bounded payoff matrices at which both selected active coordinates vanish, all other validity coordinates of the two branches are nonnegative, and the

two bordered systems have smallest singular value at least K_{ker}^{-1}. The determinant-sign regions are split before denominators are cleared, as in section D. Define the observable diagnostic score

$$\mathsf{S}_\xi(A) \coloneqq \mathrm{dist}_F(A, \mathcal{C}_\xi), \tag{N.4}$$

with distance to the empty set interpreted as $+\infty$. For nonempty $\mathcal{C}_\xi$, this score is globally one-Lipschitz and can be evaluated to arbitrary certified precision by finite semialgebraic optimization. The quantitative class below separates the true zero score from all false candidate sets.

N.2 Quantitative simple-interface class

Let

$$\Gamma = (M, r_0, \tau, K_{\mathrm{ker}}, \gamma_{\mathrm{det}}, \gamma_{\mathrm{slack}}, \gamma_{\mathrm{atlas}}, \gamma_{\mathrm{trans}}, L_0, L_2, \underline{\chi}, \overline{\chi}, \gamma_{\mathrm{sep}}, L_z, H) \tag{N.5}$$

collect positive finite constants. The determinant margin below is retained because it is an easily audited algebraic certificate, whereas the inverse bound is the quantity used directly by the branch-smoothness estimates.

Definition N.1 (Quantitatively learnable simple interface)**.** An instance is a pair $G = (B, \xi_\star)$, with $\|B\|_\infty \le M$ and $\xi_\star \in \Xi$. Write

$$g_G \coloneqq r_{\xi_\star,+}, \qquad z_{G,\pm} \coloneqq z_{\xi_\star,\pm}, \qquad a_G(A) \coloneqq \nabla g_G(A). \tag{N.6}$$

The instance belongs to $\mathfrak{G}(\Gamma)$ when the following conditions hold on the Frobenius ball $\mathbb{B}_G \coloneqq \overline{B}_F(B, r_0)$.

(Q1) **Uniform regularity and conditioning.** $\mathbb{B}_G \subseteq \mathcal{O}_{\xi_\star}$, and for both incident kernels,

$$|\Delta_{\beta_s}(A)| \ge \gamma_{\mathrm{det}}, \qquad \max\{\|M^x_{\beta_s}(A)^{-1}\|_{\mathrm{op}}, \|M^y_{\beta_s}(A)^{-1}\|_{\mathrm{op}}\} \le K_{\mathrm{ker}} \tag{N.7}$$

for every $A \in \mathbb{B}_G$ and $s \in \{+,-\}$.

(Q2) **One transverse active coordinate.**

$$g_G(B) = 0, \qquad \gamma_{\mathrm{trans}} \le \|a_G(A)\|_F \le L_0, \qquad \|D^2 g_G(A)\|_{\mathrm{op}} \le L_2 \tag{N.8}$$

throughout $\mathbb{B}_G$.

(Q3) **Strict inactive margins and coherent sides.** Every inactive validity coordinate of either incident branch is at least γ_{slack} on $\mathbb{B}_G$. Moreover,

$$r_{\xi_\star,-}(A) = -\chi_G(A) g_G(A), \qquad 0 < \underline{\chi} \le \chi_G(A) \le \overline{\chi} < \infty, \tag{N.9}$$

for an analytic χ_G. If $g_G(A) > 0$, then $z_{G,+}(A)$ is the unique equilibrium; if $g_G(A) < 0$, then $z_{G,-}(A)$ is the unique equilibrium.

(Q4) **Atlas isolation.**

$$\mathsf{S}_{\xi_\star}(B) = 0, \qquad \inf_{\xi \ne \xi_\star} \mathsf{S}_\xi(B) \ge \gamma_{\mathrm{atlas}}. \tag{N.10}$$

Thus competing regular interfaces, fan-face intersections, and nearby singular-kernel creations are excluded by a quantitative model-selection margin. Because distance to a closed set is one-Lipschitz, no additional score-regularity assumption is required.

(Q5) **Smooth and separated side targets.** The two branch maps are L_z-Lipschitz from Frobenius payoff distance to d_Δ, and

$$d_\Delta(z_{G,+}(B), z_{G,-}(B)) \ge \gamma_{\mathrm{sep}}. \tag{N.11}$$

(Q6) **Bounded local alternatives.** The unknown payoff belongs to

$$A_{n,G,h} = B + \frac{h}{\sqrt{n}}, \qquad \|h\|_F \le H. \tag{N.12}$$

(Q7) **Uniform projection tube.** Put

$$\mathcal{M}_G \coloneqq \{A \in \mathcal{C}_{\xi_\star} : \|A - B\|_F \le r_0\}, \qquad \mathcal{M}_G^{\mathrm{in}} \coloneqq \mathcal{M}_G \cap \overline{B}_F(B, r_0/2). \tag{N.13}$$

Every payoff at distance less than τ from $\mathcal{M}_G^{\mathrm{in}}$ has a unique nearest point in the entire candidate set $\mathcal{C}_{\xi_\star}$, and that point belongs to $\mathcal{M}_G$. We assume $0 < \tau \le r_0/8$.

The learner is given Γ, the finite atlas Ξ, and the noise profile. It is not given G, B, $\xi_\star$, g_G, $a_G(B)$, or $z_{G,\pm}(B)$.

The oracle interface-side allocation associated with a nonzero normal $a \in \mathbb{R}^d$ is

$$Q_i(a) := \frac{\sigma_i |a_i|}{\sum_{j=1}^d \sigma_j |a_j|}, \qquad i \in [d]. \tag{N.14}$$

It is invariant under nonzero rescaling of a. We write

$$q_G^\star := Q(a_G(B)). \tag{N.15}$$

Lemma N.2 (Lipschitz normal and design maps)**.** *Suppose $u, v \in \mathbb{R}^d$ satisfy $\|u\|_2, \|v\|_2 \geq \gamma > 0$. Then*

$$\left\| \frac{u}{\|u\|_2} - \frac{v}{\|v\|_2} \right\|_2 \leq \frac{2}{\gamma} \|u - v\|_2. \tag{N.16}$$

For unit vectors ν, ν',

$$\|Q(\nu) - Q(\nu')\|_1 \leq \frac{2\overline{\sigma}\sqrt{d}}{\underline{\sigma}} \|\nu - \nu'\|_2. \tag{N.17}$$

Proof. For nonzero u, v, add and subtract $u/\|v\|_2$ to obtain

$$\left\| \frac{u}{\|u\|_2} - \frac{v}{\|v\|_2} \right\|_2 \leq \frac{|\|u\|_2 - \|v\|_2|}{\|v\|_2} + \frac{\|u - v\|_2}{\|v\|_2} \leq \frac{2\|u - v\|_2}{\gamma}.$$

For $S(\nu) = \sum_i \sigma_i |\nu_i|$, unit length implies $S(\nu) \geq \underline{\sigma} \|\nu\|_1 \geq \underline{\sigma}$. Hence

$$\begin{aligned} \|Q(\nu) - Q(\nu')\|_1 &\leq \frac{\overline{\sigma}}{S(\nu)} \|\nu - \nu'\|_1 + \frac{|S(\nu) - S(\nu')|}{S(\nu)} \\ &\leq \frac{2\overline{\sigma}}{\underline{\sigma}} \|\nu - \nu'\|_1 \leq \frac{2\overline{\sigma}\sqrt{d}}{\underline{\sigma}} \|\nu - \nu'\|_2. \end{aligned}$$

□

Proposition N.3 (Nonemptiness)**.** *The class $\mathfrak{G}(\Gamma)$ is nonempty for strictly positive constants. For example, let*

$$B = \begin{pmatrix} 0 & 1 & 1 \\ -1/4 & 2 & 1 \\ 0 & 1 & 2 \end{pmatrix}. \tag{N.18}$$

In a sufficiently small Frobenius neighborhood of B, column one is the unique minimizer's action, row two is uniformly dominated, and the two branch germs (e_1, e_1) and (e_3, e_1) meet along $g(A) = A_{11} - A_{31} = 0$. All conditions in definition N.1 hold with positive margins.

Proof. At B, the maximum entry of column one is zero, whereas columns two and three have maximum two; thus column one remains uniquely minimizing under a small perturbation. Within column one, rows one and three tie at zero and row two has payoff $-1/4$, so row two remains uniformly suboptimal. Therefore row one is the unique maximizing action when $A_{11} > A_{31}$, row three is unique when $A_{31} > A_{11}$, and the two branches meet exactly on the affine hyperplane $A_{11} = A_{31}$. Their one-by-one bordered matrices have determinant one and uniformly bounded inverses. The active gradient is $E_{11} - E_{31}$, the Hessian vanishes, the branch maps are constant, the interface is an affine hyperplane and hence has an arbitrarily large local projection tube, and the strategy separation is $d_\Delta((e_1, e_1), (e_3, e_1)) = 1$. By continuity, all determinant, slack, conditioning, atlas-separation, and target-separation margins remain positive on a sufficiently small closed ball. □

O TARGET: AN IMPLEMENTABLE ADAPTIVE INTERFACE-LEARNING ALGORITHM

This appendix gives a fully specified procedure and proves finite-sample, uniform recovery of the unknown interface chart, normal, side targets, and oracle interface-side allocation. Adaptive learning of unknown smooth boundaries is studied in other statistical models (Locatelli et al., 2018); the game-specific structure here is a finite kernel-validity atlas with equilibrium and design targets. The uniform adaptive Gaussian score and oracle-efficiency theorems are proved in section P.

O.1 Adaptive observation protocol

At round t, the learner chooses a payoff coordinate $I_t \in [d]$, measurably from the past transcript and private randomization, and observes

$$Y_t = A_{I_t} + \sigma_{I_t}\varepsilon_t, \qquad \varepsilon_t \overset{\text{iid}}{\sim} \mathcal{N}(0,1). \tag{O.1}$$

The true payoff is $A = A_{n,G,h}$ from equation (N.12), with unknown $G \in \mathfrak{G}(\Gamma)$ and unknown h. Geometry recovery is proved uniformly over the entire ball $\|h\|_F \le H$. The side and posterior point outputs are intended for the declared information-balanced least-favorable two-point submodel

$$A^{\text{pair}}_{n,G,\eta} = B + \eta\frac{\lambda_n}{\sqrt{n}}v^{\text{lf}}_G, \qquad \eta \in \{-1,+1\}, \tag{O.2}$$

where

$$v^{\text{lf}}_{G,i} = \begin{cases} \dfrac{\sigma_i \operatorname{sgn}(a_{G,i}(B))}{\sum_j \sigma_j |a_{G,j}(B)|}, & a_{G,i}(B) \neq 0, \\ 0, & a_{G,i}(B) = 0. \end{cases} \tag{O.3}$$

The nonnegative separation schedule λ_n is known, but $G, B, v^{\text{lf}}_G, \eta$ are not. This information-balanced direction satisfies $Dg_G(B)[v^{\text{lf}}_G] = 1$; moreover every informative payoff coordinate has the same standardized local shift. It is the design-independent least-favorable direction used by the matching lower bound in section Q.

Fix a full-support pilot design $q^0 \in \Delta_d$, and put

$$q^0_{\min} \coloneqq \min_i q^0_i > 0. \tag{O.4}$$

Choose a pilot budget $n_0 < n$, let $n_1 = n - n_0$, and choose a forced exploration level $\epsilon_n \in (0,1)$.

Lemma O.1 (Balanced integer rounding)**.** *For every $q \in \Delta_d$ and integer $N \ge 1$, there is a deterministic map $\operatorname{Round}_N(q) = (N_1, \dots, N_d) \in \mathbb{N}_0^d$ satisfying*

$$\sum_i N_i = N, \qquad \left|\frac{N_i}{N} - q_i\right| < \frac{1}{N}, \qquad \left\|\frac{N_\cdot}{N} - q\right\|_1 < \frac{d}{N}. \tag{O.5}$$

If $Nq_i \ge 1$, then $N_i \ge 1$.

Proof. Let $b_i = \lfloor Nq_i \rfloor$, and distribute the remaining $N - \sum_i b_i$ units to the coordinates with the largest fractional parts, using the fixed coordinate order to break ties. Every coordinate changes by at most one from its floor, proving the coordinate and ℓ_1 bounds. If $Nq_i \ge 1$, then either $b_i \ge 1$ or $Nq_i = 1$, in which case $b_i = 1$. □

O.2 The TARGET algorithm

Definition O.2 (TARGET)**.** The *Target-Adaptive Geometry Estimation and Refinement* procedure has the following steps.

1. Full-support pilot. Set $\big(N^{(0)}_i\big)_{i=1}^d = \operatorname{Round}_{n_0}(q^0)$. Query each coordinate $N^{(0)}_i$ times and form its sample mean $\widetilde{A}_i$.

2. Finite-atlas selection. Using certified semialgebraic optimization, compute $\widetilde{\mathsf{S}}_\xi(\widetilde{A})$ with

$$|\widetilde{\mathsf{S}}_\xi(\widetilde{A}) - \mathsf{S}_\xi(\widetilde{A})| \le \eta_{\text{num}}, \qquad 0 \le \eta_{\text{num}} \le \gamma_{\text{atlas}}/16, \tag{O.6}$$

and choose

$$\widehat{\xi} = \min_{\prec} \operatorname*{argmin}_{\xi \in \Xi} \widetilde{\mathsf{S}}_\xi(\widetilde{A}), \tag{O.7}$$

where $\prec$ is the fixed atlas order. Empty candidate sets are removed offline. If the selected chart is not regular at $\widetilde{A}$, return the declared default outputs and use q^0 in stage two.

3. One-step interface anchor and geometry estimates. Write $\widehat{g} = g_{\widehat{\xi}}$ and $\widetilde{a} = \nabla\widehat{g}(\widetilde{A})$. If $\|\widetilde{a}\|_F = 0$, use the defaults. Otherwise set

$$\widehat{B} \coloneqq \widetilde{A} - \frac{\widehat{g}(\widetilde{A})}{\|\widetilde{a}\|_F^2}\widetilde{a}, \tag{O.8}$$

$$\widehat{a} \coloneqq \nabla\widehat{g}(\widehat{B}), \qquad \widehat{\nu} \coloneqq \frac{\widehat{a}}{\|\widehat{a}\|_F}, \tag{O.9}$$

$$\widehat{z}_\pm \coloneqq z_{\widehat{\xi},\pm}(\widehat{B}). \tag{O.10}$$

If $\widehat{B} \notin \mathcal{O}_{\widehat{\xi}}$, $\|\widehat{a}\|_F = 0$, or any displayed quantity is undefined, use the fixed defaults. This clause affects only off-model or pilot-failure events.

4. Estimated target-optimal design with forced exploration. Let

$$\widehat{q}_i^{\,\mathrm{or}} \coloneqq Q_i(\widehat{a}) = \frac{\sigma_i|\widehat{a}_i|}{\sum_j \sigma_j|\widehat{a}_j|}, \tag{O.11}$$

and mix it with the pilot design:

$$\overline{q}_n = (1-\epsilon_n)\widehat{q}^{\,\mathrm{or}} + \epsilon_n q^0. \tag{O.12}$$

Set $(N_i^{(1)})_{i=1}^d = \mathrm{Round}_{n_1}(\overline{q}_n)$ and $\widehat{q}_{n,i} = N_i^{(1)}/n_1$. Query coordinate i another $N_i^{(1)}$ times and form the independent stage-two mean $\widehat{A}_{2,i}$.

5. One-step interface statistic. Define

$$\widehat{G}_n \coloneqq \widehat{g}(\widetilde{A}) + \langle\nabla\widehat{g}(\widetilde{A}), \widehat{A}_2 - \widetilde{A}\rangle, \tag{O.13}$$

$$\widehat{s}_n^2 \coloneqq \sum_{i=1}^d \frac{\sigma_i^2[\nabla\widehat{g}(\widetilde{A})]_i^2}{\widehat{q}_{n,i}}, \qquad \widehat{T}_n \coloneqq \frac{\sqrt{n_1}\,\widehat{G}_n}{\widehat{s}_n}. \tag{O.14}$$

The side output is $\widehat{b}_n = \mathrm{sgn}(\widehat{T}_n)$, with $\mathrm{sgn}(0) = +1$.

For a declared two-point localization radius λ_n, put

$$\widehat{\kappa}_n = \sqrt{\frac{n_1}{n}}\frac{\lambda_n}{\widehat{s}_n}, \qquad \widehat{\pi}_n = \frac{1}{1+\exp(-2\widehat{\kappa}_n\widehat{T}_n)}, \tag{O.15}$$

and return the point estimate

$$\widehat{z}_n = \big(\widehat{\pi}_n\widehat{x}_+ + (1-\widehat{\pi}_n)\widehat{x}_-, \widehat{\pi}_n\widehat{y}_+ + (1-\widehat{\pi}_n)\widehat{y}_-\big). \tag{O.16}$$

All fallback outputs and tie rules are fixed in advance, so the procedure is defined on the entire transcript space.

Every operation in definition O.2 is finite: the atlas is finite, bordered systems are solved finitely many times, branch and validity formulas are explicit rational-analytic maps, and the rounding rule is deterministic. Certified distance-to-set evaluation can be implemented by semialgebraic optimization and bisection to the tolerance in equation (O.6). This establishes implementability in the exact finite-dimensional model; no polynomial-time complexity claim is made. The point output is included to make the procedure complete; its oracle-risk analysis is a separate theorem.

Proposition O.3 (Measurability and full-support second stage)**.** *TARGET is a Borel Markov decision rule. If*

$$n_1\epsilon_n q^0_{\min} \geq 1, \tag{O.17}$$

then $N_i^{(1)} \geq 1$ *for every coordinate, and*

$$\|\widehat{q}_n - \overline{q}_n\|_1 < \frac{d}{n_1}. \tag{O.18}$$

Proof. Finite minimization with deterministic tie breaking is Borel. On every regularity domain, all branch, validity, gradient, and normal maps are continuous; extension by fixed defaults preserves Borel measurability. The remaining operations are continuous arithmetic followed by deterministic rounding. Since $\overline{q}_{n,i} \geq \epsilon_n q^0_{\min}$, equation (O.17) and lemma O.1 give positive counts and the stated rounding bound. □

O.3 FINITE-SAMPLE UNIFORM GEOMETRY RECOVERY

Define

$$C_B \coloneqq 1 + \frac{L_0}{\gamma_{\mathrm{trans}}}, \qquad C_\nu \coloneqq \frac{2L_2 C_B}{\gamma_{\mathrm{trans}}}, \qquad C_q \coloneqq \frac{4\overline{\sigma}\sqrt{d}\, L_2 C_B}{\underline{\sigma}\gamma_{\mathrm{trans}}}. \tag{O.19}$$

Let

$$\delta_\Gamma \coloneqq \min\left\{ \frac{r_0}{2C_B}, \frac{\gamma_{\mathrm{atlas}}}{4}, \frac{\gamma_{\mathrm{sep}}}{4L_z C_B} \right\}. \tag{O.20}$$

For $\alpha \in (0,1)$, define

$$r_{n,n_0}(\alpha) \coloneqq \frac{H}{\sqrt{n}} + \overline{\sigma}\sqrt{\frac{2d\log(2d/\alpha)}{\underline{q}_0 n_0}}, \qquad \underline{q}_0 \coloneqq \frac{q^0_{\min}}{2}, \tag{O.21}$$

where $n_0 \geq 2/q^0_{\min}$ ensures $N_i^{(0)} \geq \underline{q}_0 n_0$.

Lemma O.4 (Uniform pilot concentration)**.** *Uniformly over $G \in \mathfrak{G}(\Gamma)$ and $\|h\|_F \leq H$,*

$$\Pr_{A_{n,G,h}}\left\{ \|\widetilde{A} - B\|_F > r_{n,n_0}(\alpha) \right\} \leq \alpha. \tag{O.22}$$

More generally, for every $t > 0$,

$$\Pr_{A_{n,G,h}}\left\{ \|\widetilde{A} - B\|_F > \frac{H}{\sqrt{n}} + t \right\} \leq 2d\exp\left(-\frac{\underline{q}_0 n_0 t^2}{2d\overline{\sigma}^2} \right). \tag{O.23}$$

Proof. Coordinate i has pilot variance $\sigma_i^2/N_i^{(0)} \leq \overline{\sigma}^2/(\underline{q}_0 n_0)$. If the Frobenius error exceeds t, some coordinate error exceeds $t/\sqrt{d}$. A Gaussian tail bound and a union bound therefore give

$$\Pr\{\|\widetilde{A} - A_{n,G,h}\|_F > t\} \leq 2d\exp\left(-\frac{\underline{q}_0 n_0 t^2}{2d\overline{\sigma}^2} \right).$$

Since $\|A_{n,G,h} - B\|_F \leq H/\sqrt{n}$, the two displays follow. □

Theorem O.5 (Finite-sample uniform recovery of unknown equilibrium geometry)**.** *Suppose*

$$r_{n,n_0}(\alpha) \leq \delta_\Gamma. \tag{O.24}$$

Then, uniformly over $G \in \mathfrak{G}(\Gamma)$ and $\|h\|_F \leq H$, with probability at least $1-\alpha$, all of the following statements hold simultaneously.

1. ***Exact chart and branch-pair recovery:***

$$\widehat{\xi} = \xi_\star. \tag{O.25}$$

2. ***Interface-anchor error:***

$$\|\widehat{B} - B\|_F \leq C_B r_{n,n_0}(\alpha). \tag{O.26}$$

3. ***Second-order interface residual:***

$$|g_G(\widehat{B})| \leq \frac{L_2 L_0^2}{2\gamma_{\mathrm{trans}}^2} r_{n,n_0}(\alpha)^2. \tag{O.27}$$

4. ***Raw and normalized normal recovery:***

$$\|\widehat{a} - a_G(B)\|_F \leq L_2 C_B r_{n,n_0}(\alpha), \tag{O.28}$$

$$\|\widehat{\nu} - \nu_G\|_F \leq C_\nu r_{n,n_0}(\alpha), \qquad \nu_G \coloneqq \frac{a_G(B)}{\|a_G(B)\|_F}. \tag{O.29}$$

5. ***Side-target recovery and retained separation:***

$$d_\Delta(\widehat{z}_\pm, z_{G,\pm}(B)) \le L_z C_B r_{n,n_0}(\alpha), \tag{O.30}$$

$$d_\Delta(\widehat{z}_\pm, z_{G,\pm}(A_{n,G,h})) \le L_z \left(C_B r_{n,n_0}(\alpha) + \frac{H}{\sqrt{n}} \right), \tag{O.31}$$

$$d_\Delta(\widehat{z}_+, \widehat{z}_-) \ge \frac{\gamma_{\mathrm{sep}}}{2}. \tag{O.32}$$

6. ***Oracle-design recovery:***

$$\|\widehat{q}^{\,\mathrm{or}} - q^\star_G\|_1 \le C_q r_{n,n_0}(\alpha), \tag{O.33}$$

$$\|\widehat{q}_n - q^\star_G\|_1 < C_q r_{n,n_0}(\alpha) + 2\epsilon_n + \frac{d}{n_1}. \tag{O.34}$$

Consequently,

$$\sup_{G \in \mathfrak{G}(\Gamma),\ \|h\|_F \le H} \Pr_{A_{n,G,h}} \{\widehat{\xi} \ne \xi_\star\} \le \alpha. \tag{O.35}$$

Proof. Work on the event $\mathcal{E} = \{\|\widetilde{A} - B\|_F \le r\}$ with $r = r_{n,n_0}(\alpha)$. By lemma O.4, $\Pr(\mathcal{E}) \ge 1 - \alpha$.

Distance-to-set scores are one-Lipschitz. Since $\mathsf{S}_{\xi_\star}(B) = 0$, $\mathsf{S}_{\xi_\star}(\widetilde{A}) \le r$. For every false chart with nonempty candidate set,

$$\mathsf{S}_\xi(\widetilde{A}) \ge \mathsf{S}_\xi(B) - r \ge \gamma_{\mathrm{atlas}} - r.$$

An empty candidate set has score $+\infty$ everywhere and is irrelevant. Including the certified error from equation (O.6), the approximate true score is at most $r + \eta_{\mathrm{num}} \le 5\gamma_{\mathrm{atlas}}/16$, whereas every approximate false score is at least $\gamma_{\mathrm{atlas}} - r - \eta_{\mathrm{num}} \ge 11\gamma_{\mathrm{atlas}}/16$. This proves equation (O.25).

Henceforth write $g = g_G$. The segment between B and $\widetilde{A}$ lies in $\mathbb{B}_G$, and the gradient upper bound gives

$$|g(\widetilde{A})| = |g(\widetilde{A}) - g(B)| \le L_0 \|\widetilde{A} - B\|_F \le L_0 r.$$

The Newton displacement in equation (O.8) therefore has norm at most $L_0 r/\gamma_{\mathrm{trans}}$, proving equation (O.26). The choice $r \le r_0/(2C_B)$ keeps both $\widetilde{A}$ and $\widehat{B}$ inside $\mathbb{B}_G$.

Let $s = \widehat{B} - \widetilde{A}$. Taylor expansion at $\widetilde{A}$ gives

$$g(\widehat{B}) = g(\widetilde{A}) + \langle \nabla g(\widetilde{A}), s \rangle + \frac{1}{2} D^2 g(\bar{A})[s, s]$$

for a point $\bar{A}$ on the segment. By construction, the first two terms cancel. Since $\|s\|_F \le L_0 r/\gamma_{\mathrm{trans}}$, equation (O.27) follows.

The Hessian bound and equation (O.26) imply

$$\|\widehat{a} - a_G(B)\|_F \le L_2 \|\widehat{B} - B\|_F \le L_2 C_B r,$$

which proves equation (O.28). Apply lemma N.2 with $\gamma = \gamma_{\mathrm{trans}}$ to obtain equation (O.29).

The branch Lipschitz bound gives equation (O.30). Applying it once more between B and $A_{n,G,h}$ proves equation (O.31). The triangle inequality and equation (N.11) yield

$$d_\Delta(\widehat{z}_+, \widehat{z}_-) \ge \gamma_{\mathrm{sep}} - 2L_z C_B r \ge \gamma_{\mathrm{sep}}/2,$$

proving equation (O.32).

Since $Q(a) = Q(a/\|a\|)$, combine equation (N.17) and equation (O.29) to obtain equation (O.33). Finally,

$$\|\overline{q}_n - q^\star_G\|_1 \le (1 - \epsilon_n)\|\widehat{q}^{\,\mathrm{or}} - q^\star_G\|_1 + \epsilon_n \|q^0 - q^\star_G\|_1 \le C_q r + 2\epsilon_n.$$

Add the rounding error from equation (O.18) to prove equation (O.34). All conclusions hold on $\mathcal{E}$, completing the proof. □

Corollary O.6 (Explicit uniform chart-selection tail)**.** *If $H/\sqrt{n} < \delta_\Gamma$, then*

$$\sup_{G\in\mathfrak{G}(\Gamma),\ \|h\|_F\le H} \Pr_{A_{n,G,h}}\{\widehat{\xi}\neq\xi_\star\} \le 2d\exp\left\{-\frac{\underline{q}_0 n_0(\delta_\Gamma - H/\sqrt{n})^2}{2d\overline{\sigma}^2}\right\}. \tag{O.36}$$

Proof. Apply equation (O.23) with $t = \delta_\Gamma - H/\sqrt{n}$. On the complementary event, the hypotheses of theorem O.5 hold with recovery radius at most δ_Γ, so the chart is selected exactly. □

Corollary O.7 (Uniform consistency with a vanishing pilot fraction)**.** *Let*

$$n_0 = d\left\lceil \frac{n^\rho}{d}\right\rceil, \qquad \rho\in(1/2,1), \qquad \epsilon_n = n^{-\beta}, \quad \beta\in(0,1), \qquad \alpha_n = n^{-2}. \tag{O.37}$$

Then $n_0/n\to 0$, $n_1\epsilon_n q^0_{\min}\to\infty$, and uniformly over $G\in\mathfrak{G}(\Gamma)$, $\|h\|_F\le H$,

$$\Pr\{\widehat{\xi}\neq\xi_\star\} \le n^{-2} \quad \textit{for all sufficiently large } n, \tag{O.38}$$

$$\|\widehat{B}-B\|_F + \|\widehat{a}-a_G(B)\|_F = O^{\mathrm{unif}}_{\Pr}\left(n^{-\rho/2}\sqrt{\log n}+n^{-1/2}\right), \tag{O.39}$$

$$\max_{s\in\{+,-\}} d_\Delta(\widehat{z}_s, z_{G,s}(B)) = O^{\mathrm{unif}}_{\Pr}\left(n^{-\rho/2}\sqrt{\log n}+n^{-1/2}\right), \tag{O.40}$$

$$\|\widehat{q}_n - q^\star_G\|_1 = O^{\mathrm{unif}}_{\Pr}\left(n^{-\rho/2}\sqrt{\log n}+n^{-\beta}+n^{-1}\right). \tag{O.41}$$

Moreover,

$$\sqrt{n}\,|g_G(\widehat{B})| = o^{\mathrm{unif}}_{\Pr}(1), \qquad \sqrt{n}\,\|\widehat{B}-B\|_F^2 = o^{\mathrm{unif}}_{\Pr}(1). \tag{O.42}$$

Proof. With $\alpha_n = n^{-2}$, the stochastic radius in equation (O.21) is

$$O\left(n^{-1/2}+n^{-\rho/2}\sqrt{\log n}\right),$$

which eventually lies below δ_Γ. Apply theorem O.5. The design statement follows by adding $2n^{-\beta}+d/n_1$. Finally,

$$\sqrt{n}\left(n^{-1/2}+n^{-\rho/2}\sqrt{\log n}\right)^2 = O\left(n^{-1/2}+n^{1/2-\rho}\log n+n^{-\rho/2}\sqrt{\log n}\right)\longrightarrow 0$$

because $\rho>1/2$. Combine this with equations (O.26) and (O.27). □

Remark O.8 (Exact boundary of the result). theorem O.5 is a complete finite-sample uniform recovery theorem. It proves that the geometry required by the oracle phase transition can be learned with a vanishing pilot fraction and that the estimated second-stage design converges to the oracle interface-side design. The total-variation Gaussian limit, oracle side and point risks, and adaptive confidence result are proved in section P. The matching adaptive minimax lower bound remains a distinct subsequent theorem; no lower-bound claim is used in the recovery proof.

P Uniform Adaptive Gaussian Limits and Oracle Efficiency

This appendix proves the two results left open by remark O.8: a uniform total-variation Gaussian limit for the data-driven one-step statistic and uniform attainment of the corresponding oracle design, side-classification, point-strategy, and two-point confidence benchmarks. General adaptive-experiment limit theories study which asymptotic laws are attainable under sequential information constraints (Hirano & Porter, 2023); here the Gaussian coordinate model and sample splitting yield a stronger problem-specific statement: conditional on the pilot, the second-stage statistic is exactly Gaussian, so the proof reduces to uniform control of its random mean and variance. Adaptive-design efficiency bounds have also been developed for other estimands and experiment classes (Armstrong, 2022); our claim is narrower and problem-specific, concerning equilibrium-interface geometry, the normal design, and the associated side, point, and confidence targets.

Throughout this appendix, retain definition N.1 and the TARGET procedure of definition O.2. Fix a compact phase range

$$0\le\lambda\le\Lambda, \qquad \Lambda\,\overline{\sigma}\sqrt{d}\le H\,\underline{\sigma}\gamma_{\mathrm{trans}}, \tag{P.1}$$

so that the information-balanced least-favorable pair below is contained in the bounded local class. Put

$$a_G \coloneqq a_G(B), \qquad s_G \coloneqq \sum_{i=1}^{d} \sigma_i |a_{G,i}|, \qquad q_G^\star = Q(a_G). \tag{P.2}$$

Define the information-balanced normal direction coordinatewise by

$$v_{G,i}^{\mathrm{lf}} = \begin{cases} \sigma_i \operatorname{sgn}(a_{G,i})/s_G, & a_{G,i} \neq 0, \\ 0, & a_{G,i} = 0. \end{cases} \tag{P.3}$$

Then

$$\langle a_G, v_G^{\mathrm{lf}} \rangle = 1, \qquad \|v_G^{\mathrm{lf}}\|_2 \leq v_+ \coloneqq \frac{\overline{\sigma}\sqrt{d}}{\underline{\sigma}\gamma_{\mathrm{trans}}}. \tag{P.4}$$

By equations (N.1) and (N.8),

$$s_- \coloneqq \underline{\sigma}\gamma_{\mathrm{trans}} \leq s_G \leq s_+ \coloneqq \overline{\sigma}\sqrt{d}\, L_0. \tag{P.5}$$

For a full-support design q, define the normal-variance functional

$$\mathcal{V}_G(q) \coloneqq \sum_{i=1}^{d} \frac{\sigma_i^2 a_{G,i}^2}{q_i}. \tag{P.6}$$

At a boundary coordinate we use the extended convention $0/0 = 0$ and $c/0 = +\infty$ for $c > 0$.

P.1 Oracle benchmarks and a projected target refinement

Lemma P.1 (Oracle normal design)**.** *For every nonzero normal a_G,*

$$\inf_{q \in \Delta_d} \mathcal{V}_G(q) = s_G^2, \qquad q_G^\star \in \arg\min_q \mathcal{V}_G(q). \tag{P.7}$$

The minimizer is unique on the support of a_G; coordinates with $a_{G,i} = 0$ may receive zero mass.

Proof. Cauchy–Schwarz gives

$$s_G^2 = \left(\sum_i \frac{\sigma_i |a_{G,i}|}{\sqrt{q_i}} \sqrt{q_i} \right)^2 \leq \mathcal{V}_G(q).$$

Equality holds exactly when $q_i \propto \sigma_i |a_{G,i}|$ on the nonzero support, which is equation (N.14). □

For point and confidence decisions, both estimated branch targets must lie in the strategy simplex. The one-step anchor in equation (O.8) has only an $O(r^2)$ interface residual, so we add the following exact projection postprocessing. It does not alter the one-step statistic.

Definition P.2 (Projected TARGET refinement)**.** After selecting $\widehat{\xi}$, let

$$\check{B} \in \operatorname*{argmin}_{C \in \mathcal{C}_{\widehat{\xi}}} \|C - \widetilde{A}\|_F, \tag{P.8}$$

using the fixed lexicographic tie rule outside the unique-projection event. If the candidate set is empty, use the declared defaults. Define

$$\check{z}_\pm = z_{\widehat{\xi},\pm}(\check{B}), \qquad \check{\Delta} = d_\Delta(\check{z}_+, \check{z}_-). \tag{P.9}$$

The point output in this appendix is

$$\check{z}_n = \big(\widehat{\pi}_n \check{x}_+ + (1 - \widehat{\pi}_n)\check{x}_-, \widehat{\pi}_n \check{y}_+ + (1 - \widehat{\pi}_n)\check{y}_-\big). \tag{P.10}$$

Lemma P.3 (Uniform projected-target recovery)**.** *On the recovery event $\mathcal{E}_n = \{\|\widetilde{A} - B\|_F \leq r\}$ with*

$$r \leq \min\{\delta_\Gamma, \tau/2\}, \tag{P.11}$$

we have $\widehat{\xi} = \xi_\star$, the minimizer in equation (P.8) *is unique, and*

$$\|\check{B} - B\|_F \le 2r, \tag{P.12}$$

$$d_\Delta(\check{z}_\pm, z_{G,\pm}(B)) \le 2L_z r. \tag{P.13}$$

For the information-balanced least-favorable pair

$$A^{\mathrm{lf}}_{n,G,\eta,\lambda} = B + \eta \frac{\lambda}{\sqrt{n}} v^{\mathrm{lf}}_G, \qquad \eta \in \{-1, +1\}, \tag{P.14}$$

put

$$z_{G,\eta,n}(\lambda) \coloneqq z_{G,\eta}(A^{\mathrm{lf}}_{n,G,\eta,\lambda}), \qquad \Delta_{G,n}(\lambda) \coloneqq d_\Delta(z_{G,+,n}(\lambda), z_{G,-,n}(\lambda)). \tag{P.15}$$

Then

$$\max_{\eta \in \{-1,+1\}} d_\Delta(\check{z}_\eta, z_{G,\eta,n}(\lambda)) \le e_n(r) \coloneqq L_z \left(2r + \frac{\Lambda v_+}{\sqrt{n}}\right). \tag{P.16}$$

Proof. On $\mathcal{E}_n$, exact chart recovery follows from theorem O.5. Since $B \in \mathcal{C}_{\xi_\star}$, the projection distance is at most $\|\widetilde{A} - B\|_F \le r$. Condition (Q7) gives uniqueness and locality, and the triangle inequality gives $\|\check{B} - B\|_F \le \|\check{B} - \widetilde{A}\|_F + \|\widetilde{A} - B\|_F \le 2r$. The branch Lipschitz bound proves equation (P.13). Finally, $\|A^{\mathrm{lf}}_{n,G,\eta,\lambda} - B\|_F = \lambda \|v^{\mathrm{lf}}_G\|_2 / \sqrt{n} \le \Lambda v_+ / \sqrt{n}$, and another application of branch Lipschitzness proves equation (P.16). □

P.2 Uniform stability of the random normal variance

Let

$$b_n \coloneqq \nabla g_G(\widetilde{A}), \qquad c_n \coloneqq \widehat{a} = \nabla g_G(\widehat{B}) \tag{P.17}$$

when the true chart is selected. The second-stage design is based on c_n, whereas the score variance uses b_n. This distinction is the only place where forced exploration needs a rate condition.

Put

$$\overline{S} \coloneqq \overline{\sigma}\sqrt{d}\, L_0, \qquad L_\Delta \coloneqq L_2(1 + C_B), \tag{P.18}$$

and define

$$C_{V,1} \coloneqq 10\overline{S}\,\overline{\sigma}\sqrt{d}\, L_\Delta, \qquad C_{V,2} \coloneqq \frac{2\overline{\sigma}^2 L_\Delta^2}{q^0_{\min}},$$

$$C_{V,3} \coloneqq 2\overline{S}^2, \qquad C_{V,4} \coloneqq 8d\overline{S}^2, \tag{P.19}$$

with

$$\omega_n(r, \epsilon) \coloneqq C_{V,1} r + C_{V,2} \frac{r^2}{\epsilon} + C_{V,3}\epsilon + \frac{C_{V,4}}{n_1}. \tag{P.20}$$

Lemma P.4 (Uniform random-variance and design-efficiency bound)**.** *Suppose $\mathcal{E}_n$ holds,*

$$0 < \epsilon_n \le \frac{1}{2}, \qquad n_1 \epsilon_n q^0_{\min} \ge 2. \tag{P.21}$$

Then

$$|\widehat{s}_n^2 - s_G^2| \le \omega_n(r, \epsilon_n), \tag{P.22}$$

$$|\mathcal{V}_G(\widehat{q}_n) - s_G^2| \le \omega_n(r, \epsilon_n). \tag{P.23}$$

Consequently, if

$$\omega_n(r, \epsilon_n) \le s_-^2/2, \tag{P.24}$$

then

$$\widehat{s}_n \ge s_-/\sqrt{2}, \qquad |\widehat{s}_n - s_G| \le \frac{\omega_n(r, \epsilon_n)}{s_-}. \tag{P.25}$$

Moreover,

$$0 \le \frac{\mathcal{V}_G(\widehat{q}_n)}{\inf_q \mathcal{V}_G(q)} - 1 \le \frac{\omega_n(r, \epsilon_n)}{s_-^2}. \tag{P.26}$$

Proof. Write $S(u) = \sum_i \sigma_i |u_i|$, $q^c = Q(c_n)$, and $\overline{q} = (1-\epsilon_n)q^c + \epsilon_n q^0$. On $\mathcal{E}_n$,

$$\|b_n - c_n\|_2 \le L_\Delta r, \qquad \|c_n - a_G\|_2 \le L_\Delta r. \tag{P.27}$$

Also $S(b_n), S(c_n), S(a_G) \le \overline{S}$.

By lemma P.1,

$$S(c_n)^2 \le \mathcal{V}(c_n, \overline{q}) \le \frac{S(c_n)^2}{1-\epsilon_n}, \tag{P.28}$$

so the forced-exploration error is at most $2\epsilon_n \overline{S}^2$. Because $\overline{q}_i \ge \epsilon_n q^0_{\min}$ and $|\widehat{q}_{n,i} - \overline{q}_i| < 1/n_1$, condition equation (P.21) implies

$$\widehat{q}_{n,i} \ge \overline{q}_i/2. \tag{P.29}$$

For coordinates with $c_{n,i} \neq 0$, $\overline{q}_i \ge (1-\epsilon_n)\sigma_i |c_{n,i}|/S(c_n)$. Hence

$$|\mathcal{V}(c_n, \widehat{q}_n) - \mathcal{V}(c_n, \overline{q})| \le \frac{8d\overline{S}^2}{n_1}. \tag{P.30}$$

Let $u \in \{b_n, a_G\}$ and $e = u - c_n$. From equation (P.29),

$$\begin{aligned} |\mathcal{V}(u, \widehat{q}_n) - \mathcal{V}(c_n, \widehat{q}_n)| &\le 4\sum_i \frac{\sigma_i^2 |c_{n,i}||e_i|}{\overline{q}_i} + 2\sum_i \frac{\sigma_i^2 e_i^2}{\overline{q}_i} \\ &\le 8\overline{S}\,\overline{\sigma}\sqrt{d}\,\|e\|_2 + \frac{2\overline{\sigma}^2}{\epsilon_n q^0_{\min}}\|e\|_2^2. \end{aligned}$$

Finally,

$$|S(c_n)^2 - S(a_G)^2| \le 2\overline{S}\,\overline{\sigma}\sqrt{d}\,\|c_n - a_G\|_2. \tag{P.31}$$

Combining the last four displays with equation (P.27) proves equations (P.22) and (P.23). The scale bound follows from $|\sqrt{x} - \sqrt{y}| = |x-y|/(\sqrt{x}+\sqrt{y})$, and equation (P.26) follows from lemma P.1. □

P.3 Uniform adaptive Gaussian approximation

For $G \in \mathfrak{G}(\Gamma)$ and $\|h\|_F \le H$, define the oracle second-stage mean

$$\mu_{G,n}(h) \coloneqq \sqrt{\frac{n_1}{n}}\frac{\langle a_G, h\rangle}{s_G}. \tag{P.32}$$

For the information-balanced least-favorable pair, define

$$\kappa_G(\lambda) \coloneqq \frac{\lambda}{s_G}, \qquad \kappa^{(1)}_{G,n}(\lambda) \coloneqq \sqrt{\frac{n_1}{n}}\kappa_G(\lambda). \tag{P.33}$$

Lemma P.5 (Conditional exact normality)**.** *Conditional on the pilot transcript and on every event on which TARGET uses a regular selected chart and positive second-stage counts,*

$$\widehat{T}_n \mid \mathcal{F}_0 \sim \mathcal{N}(M_{n,G,h}, 1), \tag{P.34}$$

where

$$M_{n,G,h} = \frac{\sqrt{n_1}}{\widehat{s}_n}\left\{ g_G(\widetilde{A}) + \langle \nabla g_G(\widetilde{A}), A_{n,G,h} - \widetilde{A}\rangle \right\}, \tag{P.35}$$

Proof. Conditional on the pilot, the design counts and the vector $b_n = \nabla g_G(\widetilde{A})$ are fixed, while the stage-two coordinate means are independent Gaussians with $\widehat{A}_{2,i} - A_{n,G,h} \sim \mathcal{N}(0, \sigma_i^2/N_i^{(1)})$. Thus the random part of $\widehat{G}_n$ is Gaussian with variance $\sum_i b_{n,i}^2 \sigma_i^2/N_i^{(1)} = \widehat{s}_n^2/n_1$, proving the claim. □

Lemma P.6 (Gaussian-mixture total-variation bound)**.** *Suppose $X \mid \mathcal{F} \sim \mathcal{N}(M, 1)$ on an event $\mathcal{E}$, and let X be arbitrary on $\mathcal{E}^c$. Then for every $\mu \in \mathbb{R}$,*

$$d_{\mathrm{TV}}(\mathcal{L}(X), \mathcal{N}(\mu, 1)) \le \Pr(\mathcal{E}^c) + \frac{1}{\sqrt{2\pi}}\operatorname*{ess\,sup}_{\mathcal{E}} |M - \mu|. \tag{P.36}$$

Proof. Total variation is convex under mixing. For equal-variance normals,

$$d_{\mathrm{TV}}(\mathcal{N}(m,1),\mathcal{N}(\mu,1)) = 2\Phi(|m-\mu|/2) - 1 \le |m-\mu|/\sqrt{2\pi}.$$

Average this conditional bound on $\mathcal{E}$, and bound the complement by one. □

Define

$$\mathfrak{r}_n(r) := \frac{\sqrt{2}}{s_-}\left(2L_2\sqrt{n}\,r^2 + \frac{L_2H^2}{2\sqrt{n}}\right), \tag{P.37}$$

and

$$\mathfrak{k}_n(r,\epsilon) := \frac{\sqrt{2}L_0H}{s_-^3}\omega_n(r,\epsilon). \tag{P.38}$$

Theorem P.7 (Finite-sample uniform adaptive Gaussian approximation)**.** *Let $\alpha \in (0,1)$, set $r = r_{n,n_0}(\alpha)$, and suppose*

$$r \le \min\{\delta_\Gamma, \tau/2\}, \quad 0 < \epsilon_n \le 1/2, \quad n_1\epsilon_n q^0_{\min} \ge 2, \quad \omega_n(r,\epsilon_n) \le s_-^2/2. \tag{P.39}$$

Then

$$\sup_{G\in\mathfrak{G}(\Gamma),\ \|h\|_F\le H} d_{\mathrm{TV}}\left(\mathcal{L}_{A_{n,G,h}}(\widehat{T}_n), \mathcal{N}(\mu_{G,n}(h),1)\right) \le \alpha + \frac{\mathfrak{r}_n(r)+\mathfrak{k}_n(r,\epsilon_n)}{\sqrt{2\pi}}. \tag{P.40}$$

For the information-balanced least-favorable pair, uniformly over $G \in \mathfrak{G}(\Gamma)$, $\eta \in \{-1,+1\}$, and $0 \le \lambda \le \Lambda$,

$$d_{\mathrm{TV}}\left(\mathcal{L}_{A^{\mathrm{lf}}_{n,G,\eta,\lambda}}(\widehat{T}_n), \mathcal{N}(\eta\kappa_G(\lambda),1)\right) \le \alpha + \frac{\mathfrak{r}_n(r)+\mathfrak{k}_n(r,\epsilon_n)}{\sqrt{2\pi}} + \frac{\Lambda}{s_-\sqrt{2\pi}}\frac{n_0}{n}. \tag{P.41}$$

Proof. Work on $\mathcal{E}_n$. Exact chart recovery holds. Taylor expansion of g_G at $\widetilde{A}$ gives

$$\begin{aligned}&\left|g_G(\widetilde{A}) + \langle b_n, A_{n,G,h} - \widetilde{A}\rangle - g_G(A_{n,G,h})\right| \\ &\qquad \le \frac{L_2}{2}\|A_{n,G,h} - \widetilde{A}\|_F^2 \le 2L_2r^2.\end{aligned}$$

A second Taylor expansion at B gives

$$\left|g_G(A_{n,G,h}) - \frac{\langle a_G, h\rangle}{\sqrt{n}}\right| \le \frac{L_2H^2}{2n}. \tag{P.42}$$

Using $\widehat{s}_n \ge s_-/\sqrt{2}$ and equation (P.25), the difference between the conditional mean in equation (P.35) and $\mu_{G,n}(h)$ is at most $\mathfrak{r}_n(r)+\mathfrak{k}_n(r,\epsilon_n)$. Now apply lemma P.6 and $\Pr(\mathcal{E}_n^c) \le \alpha$.

On the information-balanced least-favorable pair,

$$\mu_{G,n}(\eta\lambda v_G^{\mathrm{lf}}) = \eta\kappa^{(1)}_{G,n}(\lambda). \tag{P.43}$$

The inequality $1 - \sqrt{1-n_0/n} \le n_0/n$ and the equal-variance normal bound give the final term in equation (P.41). □

Corollary P.8 (Uniform total-variation limit under power schedules)**.** *Let*

$$n_0 = d\left\lceil n^\rho/d\right\rceil, \qquad \epsilon_n = n^{-\beta}, \qquad \alpha_n = n^{-2}, \qquad \frac{1}{2} < \rho < 1, \quad 0 < \beta < \rho. \tag{P.44}$$

Then

$$\sup_{G\in\mathfrak{G}(\Gamma),\ \|h\|_F\le H} d_{\mathrm{TV}}\left(\mathcal{L}_{A_{n,G,h}}(\widehat{T}_n), \mathcal{N}\left(\frac{\langle a_G,h\rangle}{s_G},1\right)\right) \longrightarrow 0. \tag{P.45}$$

In particular,

$$\sup_{G\in\mathfrak{G}(\Gamma),\eta\in\{-1,+1\},\lambda\in[0,\Lambda]} d_{\mathrm{TV}}\left(\mathcal{L}_{A^{\mathrm{lf}}_{n,G,\eta,\lambda}}(\widehat{T}_n), \mathcal{N}(\eta\kappa_G(\lambda),1)\right) \longrightarrow 0. \tag{P.46}$$

Proof. From equation (O.21),

$$r = O\left(n^{-1/2} + n^{-\rho/2}\sqrt{\log n}\right).$$

Hence $\sqrt{n}\, r^2 \to 0$, $r^2/\epsilon_n \to 0$ because $\beta < \rho$, $\epsilon_n \to 0$, $n_1^{-1} \to 0$, and $n_0/n \to 0$. Thus $\omega_n(r, \epsilon_n), \mathfrak{r}_n(r), \mathfrak{k}_n(r, \epsilon_n) \to 0$. Apply theorem P.7; the additional replacement of $\sqrt{n_1/n}$ by one costs at most a constant multiple of n_0/n in total variation. □

Lemma P.9 (Uniform phase-parameter consistency)**.** *Under the conditions of lemma P.4, for every* $0 \le \lambda \le \Lambda$*,*

$$|\widehat{\kappa}_n - \kappa_G(\lambda)| \le C_\kappa \left\{\omega_n(r, \epsilon_n) + \frac{n_0}{n}\right\}, \tag{P.47}$$

where $C_\kappa < \infty$ *depends only on* Λ *and* s_-*. On the information-balanced least-favorable pair,*

$$|M_{n,G,\eta,\lambda} - \eta\widehat{\kappa}_n| \le \mathfrak{r}_n(r). \tag{P.48}$$

Proof. Using $\widehat{s}_n \ge s_-/\sqrt{2}$, $s_G \ge s_-$, and equation (P.25),

$$\begin{aligned} |\widehat{\kappa}_n - \kappa_G(\lambda)| &\le \lambda \left| \frac{\sqrt{n_1/n}}{\widehat{s}_n} - \frac{1}{s_G} \right| \\ &\le \frac{\sqrt{2}\Lambda}{s_-}\frac{n_0}{n} + \frac{\sqrt{2}\Lambda}{s_-^3}\omega_n(r, \epsilon_n). \end{aligned}$$

This proves the first display. For the second, return to the proof of theorem P.7. On the information-balanced least-favorable pair, the target linear term is exactly $\eta\lambda/\sqrt{n}$, so after division by $\widehat{s}_n$ only the two Taylor remainders collected in $\mathfrak{r}_n(r)$ remain. □

P.4 Uniform oracle efficiency

For $0 \le \lambda \le \Lambda$, define the bounded composite side classes

$$\mathcal{H}^H_{G,+}(\lambda) = \{h : \|h\|_F \le H,\ \langle a_G, h\rangle \ge \lambda\}, \tag{P.49}$$

$$\mathcal{H}^H_{G,-}(\lambda) = \{h : \|h\|_F \le H,\ \langle a_G, h\rangle \le -\lambda\}. \tag{P.50}$$

The phase-range condition in equation (P.1) ensures that the least-favorable pair in equation (P.14) belongs to these classes. Define the adaptive side risk

$$R^{\mathrm{ad}}_{\mathrm{side,n}}(G, \lambda) = \max\left\{ \sup_{h \in \mathcal{H}^H_{G,+}(\lambda)} P_{n,G,h}\{\widehat{b}_n = -\}, \sup_{h \in \mathcal{H}^H_{G,-}(\lambda)} P_{n,G,h}\{\widehat{b}_n = +\} \right\}. \tag{P.51}$$

Let

$$r_{\mathrm{pt}}(\kappa) \coloneqq e^{-\kappa^2/2}\mathbb{E}_{Z\sim\mathcal{N}(0,1)}[\cosh(\kappa Z)^{-1}]. \tag{P.52}$$

For the information-balanced least-favorable pair, define

$$R^{\mathrm{ad}}_{\mathrm{pt,n}}(G, \lambda) = \max_{\eta \in \{-1,+1\}} \mathbb{E}_{n,G,\eta,\lambda} d^2_\Delta(\check{z}_n, z_{G,\eta,n}(\lambda)). \tag{P.53}$$

Lemma P.10 (Lipschitz benchmark functions)**.** *On* $[0, K]$*,*

$$|r'_{\mathrm{pt}}(\kappa)| \le K + \sqrt{2/\pi}. \tag{P.54}$$

For every fixed $\alpha \in (0, 1/2)$*,*

$$q_\alpha(\kappa) = [\Phi(z_{1-\alpha} - 2\kappa) - \alpha]_+ \tag{P.55}$$

is $\sqrt{2/\pi}$*-Lipschitz.*

Proof. Differentiate under the Gaussian expectation:

$$r'_{\mathrm{pt}}(\kappa) = -\kappa r_{\mathrm{pt}}(\kappa) - e^{-\kappa^2/2}\mathbb{E}[Z\ \mathrm{sech}(\kappa Z)\tanh(\kappa Z)].$$

Use $0 \le r_{\mathrm{pt}} \le 1$, $|\,\mathrm{sech}\tanh\,| \le 1$, and $\mathbb{E}|Z| = \sqrt{2/\pi}$. The positive-part map is one-Lipschitz and $|2\Phi'(x)| \le \sqrt{2/\pi}$, proving the second claim. □

Define the deterministic efficiency modulus

$$\mathfrak{E}_n \coloneqq \alpha_n + \sqrt{n}\, r_n^2 + n^{-1/2} + \omega_n(r_n, \epsilon_n) + \frac{n_0}{n} + e_n(r_n), \tag{P.56}$$

where $r_n = r_{n,n_0}(\alpha_n)$.

Theorem P.11 (Uniform oracle efficiency of TARGET)**.** *Assume*

$$\alpha_n \to 0, \quad r_n \to 0, \quad \sqrt{n}\, r_n^2 \to 0, \quad \epsilon_n \to 0, \quad \frac{r_n^2}{\epsilon_n} \to 0, \quad n_1 \epsilon_n q_{\min}^0 \to \infty, \quad \frac{n_0}{n} \to 0. \tag{P.57}$$

Then $\mathfrak{E}_n \to 0$. There is a finite constant C_{eff}, depending only on $\Gamma, H, \Lambda, \alpha$, such that, for all sufficiently large n, the bounds below hold uniformly over $G \in \mathfrak{G}(\Gamma)$ and $0 \le \lambda \le \Lambda$.

1. ***Design efficiency.***

$$\frac{\mathcal{V}_G(\widehat{q}_n)}{\inf_q \mathcal{V}_G(q)} \xrightarrow{P} 1 \quad \textit{uniformly in } G. \tag{P.58}$$

More quantitatively,

$$\sup_{G \in \mathfrak{G}(\Gamma),\, \|h\|_F \le H} P_{n,G,h}\left\{ \frac{\mathcal{V}_G(\widehat{q}_n)}{\inf_q \mathcal{V}_G(q)} - 1 > \frac{\omega_n(r_n, \epsilon_n)}{s_-^2} \right\} \le \alpha_n. \tag{P.59}$$

2. ***Composite interface-side oracle risk.***

$$\sup_{G \in \mathfrak{G}(\Gamma),\, 0 \le \lambda \le \Lambda} \left| R^{\mathrm{ad}}_{\mathrm{side,n}}(G, \lambda) - \Phi(-\kappa_G(\lambda)) \right| \longrightarrow 0. \tag{P.60}$$

More precisely, the absolute difference is at most $C_{\mathrm{eff}}\, \mathfrak{E}_n$.

3. ***Least-favorable-pair point risk.***

$$\sup_{G \in \mathfrak{G}(\Gamma),\, 0 \le \lambda \le \Lambda} \left| R^{\mathrm{ad}}_{\mathrm{pt,n}}(G, \lambda) - \frac{\Delta_{G,n}(\lambda)^2}{4} r_{\mathrm{pt}}(\kappa_G(\lambda)) \right| \longrightarrow 0. \tag{P.61}$$

More precisely, the absolute difference is at most $C_{\mathrm{eff}}\, \mathfrak{E}_n$.

4. ***Adaptive honest confidence and oracle diameter.*** *Fix $\alpha \in (0, 1/2)$, and let*

$$\widehat{c}_{n,\alpha} = [z_{1-\alpha} - \widehat{\kappa}_n]_+. \tag{P.62}$$

Let $\mathbb{B}_\Delta(z, e) = \{w \in \mathcal{Z} : d_\Delta(w, z) \le e\}$, and define

$$\widehat{\mathfrak{C}}_{n,\alpha} = \begin{cases} \mathbb{B}_\Delta(\check{z}_+, e_n(r_n)), & \widehat{T}_n > \widehat{c}_{n,\alpha}, \\ \mathbb{B}_\Delta(\check{z}_-, e_n(r_n)), & \widehat{T}_n < -\widehat{c}_{n,\alpha}, \\ \mathbb{B}_\Delta(\check{z}_+, e_n(r_n)) \cup \mathbb{B}_\Delta(\check{z}_-, e_n(r_n)), & |\widehat{T}_n| \le \widehat{c}_{n,\alpha}. \end{cases} \tag{P.63}$$

The map from the transcript to $\widehat{\mathfrak{C}}_{n,\alpha}$ is Borel measurable as a random compact set in the Hausdorff hyperspace: its centers, radius, threshold, and three report events are Borel, and finite unions of closed metric balls are compact in $\mathcal{Z}$. Then

$$\inf_{G \in \mathfrak{G}(\Gamma),\, 0 \le \lambda \le \Lambda,\, \eta \in \{-1, +1\}} P_{n,G,\eta,\lambda}\{ z_{G,\eta,n}(\lambda) \in \widehat{\mathfrak{C}}_{n,\alpha} \} \ge 1 - \alpha - o(1), \tag{P.64}$$

In fact the left side is at least $1 - \alpha - C_{\mathrm{eff}}\, \mathfrak{E}_n$. Moreover,

$$\sup_{G \in \mathfrak{G}(\Gamma),\, 0 \le \lambda \le \Lambda} \left| \max_\eta \mathbb{E}_{n,G,\eta,\lambda} \operatorname{diam}_{d_\Delta}(\widehat{\mathfrak{C}}_{n,\alpha}) - \Delta_{G,n}(\lambda) q_\alpha(\kappa_G(\lambda)) \right| \longrightarrow 0. \tag{P.65}$$

and the displayed absolute difference is at most $C_{\mathrm{eff}}\, \mathfrak{E}_n$.

Under the power schedules in equation (P.44)*, all assumptions hold.*

Proof. The design statement follows immediately from lemma P.4 and $\Pr(\mathcal{E}_n^c) \le \alpha_n$.

For the side risk, let δ_n^{TV} denote the right-hand side of equation (P.40). If $h \in \mathcal{H}_{G,+}^{H}(\lambda)$, then $\mu_{G,n}(h) \ge \kappa_{G,n}^{(1)}(\lambda)$, so

$$P_{n,G,h}\{\widehat{b}_n = -\} \le \Phi(-\kappa_{G,n}^{(1)}(\lambda)) + \delta_n^{\mathrm{TV}}.$$

The negative class is symmetric. Conversely, the least-favorable pair belongs to the classes, and equation (P.40) gives the reverse inequality at those two points. Thus

$$\left| R_{\mathrm{side,n}}^{\mathrm{ad}}(G,\lambda) - \Phi(-\kappa_{G,n}^{(1)}(\lambda)) \right| \le \delta_n^{\mathrm{TV}}.$$

Since Φ is $1/\sqrt{2\pi}$-Lipschitz and $|\kappa_{G,n}^{(1)} - \kappa_G| \le (\Lambda/s_-)(n_0/n)$, the side result follows.

For point risk, work first on $\mathcal{E}_n$ and condition on the pilot. By lemma P.5, under the pair $\widehat{T}_n \sim \mathcal{N}(M_{n,\eta}, 1)$ conditionally. By lemma P.9,

$$|M_{n,\eta} - \eta\widehat{\kappa}_n| \le \mathfrak{r}_n(r_n). \tag{P.66}$$

In the ideal conditional experiment $T \sim \mathcal{N}(\eta\widehat{\kappa}_n, 1)$, posterior interpolation between the fixed estimated targets $\check{z}_+, \check{z}_-$ has exact equalizer risk

$$\frac{\check{\Delta}^2}{4} r_{\mathrm{pt}}(\widehat{\kappa}_n). \tag{P.67}$$

The squared strategy loss is bounded by four, so replacing the ideal law by the actual conditional law changes risk by at most $4\mathfrak{r}_n(r_n)/\sqrt{2\pi}$. Replacing the estimated target by the true target changes squared loss by at most $4e_n(r_n)$, because the diameter of $\mathcal{Z}$ is two. On $\mathcal{E}_n^c$, the loss is at most four.

Furthermore, $|\check{\Delta} - \Delta_{G,n}(\lambda)| \le 2e_n(r_n)$, and lemma P.9 controls $|\widehat{\kappa}_n - \kappa_G(\lambda)|$ uniformly. Apply lemma P.10 to equation (P.67); every error is bounded by a constant multiple of $\mathfrak{E}_n$, proving equation (P.61).

For confidence, on $\mathcal{E}_n$ the true target lies in the corresponding ball by lemma P.3. In the ideal conditional experiment $N(\eta\widehat{\kappa}_n, 1)$, the three-report threshold in equation (P.62) includes the correct estimated target with probability at least $1-\alpha$. Replacing the ideal law by the actual conditional law costs at most $\mathfrak{r}_n(r_n)/\sqrt{2\pi}$, and the complement of $\mathcal{E}_n$ costs at most α_n. This proves equation (P.64).

For the diameter, on $\mathcal{E}_n$,

$$\left| \mathrm{diam}(\widehat{\mathfrak{C}}_{n,\alpha}) - \check{\Delta}\,\mathbf{1}\{|\widehat{T}_n| \le \widehat{c}_{n,\alpha}\} \right| \le 2e_n(r_n). \tag{P.68}$$

The ideal conditional probability of the middle report is $q_\alpha(\widehat{\kappa}_n)$. Total variation controls the actual middle probability, lemma P.10 controls the random phase parameter, and $|\check{\Delta} - \Delta_{G,n}| \le 2e_n(r_n)$. The failure event contributes at most two, the diameter of $\mathcal{Z}$. Hence the diameter error is also bounded by a constant multiple of $\mathfrak{E}_n$, proving the final claim. □

Remark P.12 (Meaning of oracle efficiency). The oracle benchmarks above are the exact Gaussian approximate-design benchmarks for a learner that knows B, the active interface normal, and the two side targets. theorem P.11 proves that estimating those objects has asymptotically zero first-order cost on the quantitative class. The matching lower bound against every adaptive procedure, even one given the geometry, is proved in theorem Q.4.

Q Matching Adaptive Minimax Lower Bounds

This appendix proves that the benchmarks attained in theorem P.11 are unavoidable even if the learner is given the true interface game, active chart, normal, and side targets. Thus removing the geometry oracle has asymptotically zero first-order cost on the quantitative class. The proof uses an information-balanced least-favorable pair whose standardized shift is identical in every informative payoff coordinate. This choice is essential: a Euclidean-normal pair need not be least favorable for an adaptive coordinate-query experiment.

Throughout, retain the quantitative class of definition N.1. For $G = (B, \xi_\star)$, write

$$a_G = \nabla g_G(B), \qquad s_G = \sum_{i=1}^{d} \sigma_i |a_{G,i}|, \qquad \kappa_G(\lambda) = \frac{\lambda}{s_G}. \tag{Q.1}$$

Let

$$v_{G,i}^{\mathrm{lf}} = \begin{cases} \sigma_i \operatorname{sgn}(a_{G,i})/s_G, & a_{G,i} \neq 0, \\ 0, & a_{G,i} = 0, \end{cases} \tag{Q.2}$$

and define the pair

$$A_{n,G,\eta,\lambda}^{\mathrm{lf}} = B + \eta \frac{\lambda}{\sqrt{n}} v_G^{\mathrm{lf}}, \qquad \eta \in \{-1, +1\}. \tag{Q.3}$$

Then

$$\langle a_G, v_G^{\mathrm{lf}} \rangle = 1, \qquad \|v_G^{\mathrm{lf}}\|_2 \leq \frac{\overline{\sigma}\sqrt{d}}{\underline{\sigma}\gamma_{\mathrm{trans}}}. \tag{Q.4}$$

The phase-range condition equation (P.1) therefore places the pair inside the bounded local class.

Lemma Q.1 (Uniform branch validity of the information-balanced pair)**.** *There exists $n_{\mathrm{lf}} < \infty$, depending only on Γ and Λ, such that for every $n \geq n_{\mathrm{lf}}$, every $G \in \mathfrak{G}(\Gamma)$, and every $0 < \lambda \leq \Lambda$,*

$$\operatorname{sgn} g_G(A_{n,G,\eta,\lambda}^{\mathrm{lf}}) = \eta. \tag{Q.5}$$

At $\lambda = 0$, both incident branch germs are valid at B.

Proof. Taylor expansion at B, together with $\langle a_G, v_G^{\mathrm{lf}} \rangle = 1$, gives

$$g_G(A_{n,G,\eta,\lambda}^{\mathrm{lf}}) = \eta \frac{\lambda}{\sqrt{n}} + R_{n,G,\eta,\lambda}, \qquad |R_{n,G,\eta,\lambda}| \leq \frac{L_2 \lambda^2 v_+^2}{2n}.$$

For $0 < \lambda \leq \Lambda$, the remainder divided by the leading magnitude is at most $L_2 \Lambda v_+^2/(2\sqrt{n})$, uniformly in G. This is less than one for all sufficiently large n. Condition (Q3) then identifies the corresponding strict equilibrium branch. At $\lambda = 0$, the definition of the interface chart gives simultaneous validity of the two incident germs. □

Put

$$z_{G,\eta,n}^{\mathrm{lf}}(\lambda) = z_{G,\eta}(A_{n,G,\eta,\lambda}^{\mathrm{lf}}), \qquad \Delta_{G,n}^{\mathrm{lf}}(\lambda) = d_\Delta(z_{G,+,n}^{\mathrm{lf}}(\lambda), z_{G,-,n}^{\mathrm{lf}}(\lambda)). \tag{Q.6}$$

These are the same targets and separation denoted by $z_{G,\eta,n}(\lambda)$ and $\Delta_{G,n}(\lambda)$ in equation (P.15).

Q.1 Adaptive transcript domination by one scalar Gaussian experiment

An *oracle-assisted adaptive procedure* may use $G, B, a_G, z_{G,+}, z_{G,-}$ when selecting coordinates and forming its terminal decision. This only enlarges the class of procedures and therefore produces a valid lower bound for every non-oracle method.

Lemma Q.2 (Exact scalar domination of every adaptive transcript)**.** *Fix G, n, λ, and let π be any possibly randomized adaptive n-query policy in equation* (O.1)*, even one given the full geometry G. Consider the scalar experiment*

$$X_t \overset{\mathrm{iid}}{\sim} \mathcal{N}\left(\eta \frac{\kappa_G(\lambda)}{\sqrt{n}}, 1\right), \qquad t = 1, \ldots, n. \tag{Q.7}$$

There is a Markov kernel, independent of η, mapping $X_1, \ldots, X_n$ to a transcript having exactly the law generated by π under equation (Q.3)*. Conversely, for any coordinate i_0 with $a_{G,i_0} \neq 0$, the policy that always queries i_0 is Blackwell-equivalent to equation* (Q.7)*. Hence the most informative oracle-assisted adaptive query experiment on the pair is exactly the scalar Gaussian experiment*

$$T_n = \frac{1}{\sqrt{n}} \sum_{t=1}^{n} X_t \sim \mathcal{N}(\eta \kappa_G(\lambda), 1). \tag{Q.8}$$

Proof. Let U be the private random seed of π, independent of the scalar observations, and let $(Z_t)_{t=1}^n$ be iid standard Gaussians independent of everything else. Construct the transcript recursively. Given the simulated past and U, run the policy to obtain I_t. If $a_{G,I_t} \neq 0$, set

$$Y_t = B_{I_t} + \sigma_{I_t} \operatorname{sgn}(a_{G,I_t}) X_t. \tag{Q.9}$$

If $a_{G,I_t} = 0$, set

$$Y_t = B_{I_t} + \sigma_{I_t} Z_t. \tag{Q.10}$$

For an informative coordinate, the conditional mean under η is

$$B_{I_t} + \eta \frac{\lambda}{\sqrt{n}} \frac{\sigma_{I_t} \operatorname{sgn}(a_{G,I_t})}{s_G} = [A^{\mathrm{lf}}_{n,G,\eta,\lambda}]_{I_t},$$

and its conditional variance is $\sigma^2_{I_t}$. For a null-normal coordinate, the two pair points have the same coordinate, so equation (Q.10) is also exact. Freshness of X_t, Z_t yields the required conditional independence. Induction on t proves equality of transcript laws. The simulation kernel depends on G and π, which are granted to the oracle, but not on η.

Conversely, fix i_0 with $a_{G,i_0} \neq 0$, query it at every round, and transform

$$X_t = \operatorname{sgn}(a_{G,i_0}) \frac{Y_t - B_{i_0}}{\sigma_{i_0}}.$$

This gives exactly equation (Q.7). The two experiments are therefore Blackwell-equivalent for that policy, and every other adaptive transcript is a garbling of the same scalar sequence. □

The transcript simulation is a problem-specific exact version of the adaptive change-of-measure decomposition commonly used in pure-exploration lower bounds; see, for example, Kaufmann et al. (2016, Lemma 1). No asymptotic LAN argument is needed here.

Q.2 Exact oracle lower bounds

For a side rule $\widehat{b} \in \{-1, +1\}$, define its pair risk by

$$R^{\mathrm{pair}}_{\mathrm{side,n}}(\widehat{b}; G, \lambda) = \max_{\eta \in \{-1,+1\}} P^{\mathrm{lf}}_{n,G,\eta,\lambda}\{\widehat{b} \neq \eta\}. \tag{Q.11}$$

For a point estimator $\widehat{z}$, define

$$R^{\mathrm{pair}}_{\mathrm{pt,n}}(\widehat{z}; G, \lambda) = \max_{\eta} \mathbb{E}^{\mathrm{lf}}_{n,G,\eta,\lambda} d^2_\Delta(\widehat{z}, z^{\mathrm{lf}}_{G,\eta,n}(\lambda)). \tag{Q.12}$$

For $\alpha \in (0, 1/2)$, let an oracle two-point confidence rule be a Borel random compact set $\mathfrak{C} \subseteq \mathcal{Z}$ satisfying

$$\inf_{\eta \in \{-1,+1\}} P^{\mathrm{lf}}_{n,G,\eta,\lambda}\{z^{\mathrm{lf}}_{G,\eta,n}(\lambda) \in \mathfrak{C}\} \geq 1 - \alpha. \tag{Q.13}$$

Theorem Q.3 (Exact oracle benchmarks under arbitrary adaptive querying)**.** *For every $G \in \mathfrak{G}(\Gamma)$, every $n \geq n_{\mathrm{lf}}$, and every $0 \leq \lambda \leq \Lambda$, even if the procedure is given the true geometry,*

$$\inf_{\pi, \widehat{b}} R^{\mathrm{pair}}_{\mathrm{side,n}}(\widehat{b}; G, \lambda) = \Phi(-\kappa_G(\lambda)), \tag{Q.14}$$

$$\inf_{\pi, \widehat{z}} R^{\mathrm{pair}}_{\mathrm{pt,n}}(\widehat{z}; G, \lambda) = \frac{\Delta^{\mathrm{lf}}_{G,n}(\lambda)^2}{4} r_{\mathrm{pt}}(\kappa_G(\lambda)), \tag{Q.15}$$

$$\inf_{\pi, \mathfrak{C}:\,\text{(Q.13)}} \max_{\eta} \mathbb{E}^{\mathrm{lf}}_{n,G,\eta,\lambda} \operatorname{diam}_{d_\Delta}(\mathfrak{C}) = \Delta^{\mathrm{lf}}_{G,n}(\lambda) q_\alpha(\kappa_G(\lambda)). \tag{Q.16}$$

The infima range over all randomized adaptive policies and terminal decisions. Moreover, every side rule on the bounded composite classes in equation (P.50) *has risk at least the right-hand side of equation* (Q.14)*.*

Proof. By lemma Q.2, every adaptive transcript is a garbling of $T_n \sim \mathcal{N}(\eta\kappa_G, 1)$; querying one informative coordinate at every round attains that scalar experiment exactly. Neyman–Pearson therefore gives equation (Q.14).

For point loss, condition on $T_n = t$ under the equal prior and let $\pi(t) = \{1 + \exp(-2\kappa_G t)\}^{-1}$. For any decision z, put $d_\pm = d_\Delta(z, z^{\mathrm{lf}}_{G,\pm,n})$. The triangle inequality gives $d_+ + d_- \geq \Delta^{\mathrm{lf}}_{G,n}$, and minimizing $\pi(t)d_+^2 + (1-\pi(t))d_-^2$ over nonnegative d_+, d_- subject to this inequality gives

$$\pi(t)(1-\pi(t))\Delta^{\mathrm{lf}}_{G,n}(\lambda)^2.$$

Componentwise posterior interpolation between the two strategy targets lies in the product simplex and attains this lower bound. Reflection symmetry makes the rule equalizer. Integrating against the equal-mixture density yields equation (Q.15).

The confidence equality is the exact two-point Gaussian diameter theorem of theorem H.6, applied after the exact Blackwell reduction. Its attaining three-report rule is available to the oracle because the two point targets are known.

Finally, the pair directions satisfy $\langle a_G, \pm\lambda v^{\mathrm{lf}}_G\rangle = \pm\lambda$, so they belong to the two fixed composite halfspace classes. Restricting the composite side problem to this pair proves the final lower bound. □

Q.3 Matching adaptive minimax theorem

For a fixed G, let $R^{\star,\mathrm{or}}_{\mathrm{side,n}}(G,\lambda)$ denote the infimum of the composite side risk over all adaptive procedures allowed to know G. Define $R^{\star,\mathrm{or}}_{\mathrm{pt,n}}(G,\lambda)$ analogously on the pair, with the two targets known. For confidence, let $D^{\star,\mathrm{or}}_{n,\alpha}(G,\lambda)$ be the exact oracle optimum in equation (Q.16). Allowing geometry information can only lower risk, so these quantities are lower bounds for the corresponding non-oracle problems.

Theorem Q.4 (Uniform matching lower bound and zero first-order oracle value)**.** *Run TARGET under the schedules and assumptions of theorem P.11, with the least-favorable pair equation* (Q.3)*. Then uniformly over $G \in \mathfrak{G}(\Gamma)$ and $0 \leq \lambda \leq \Lambda$,*

$$0 \leq R^{\mathrm{ad}}_{\mathrm{side,n}}(G,\lambda) - R^{\star,\mathrm{or}}_{\mathrm{side,n}}(G,\lambda) \longrightarrow 0, \tag{Q.17}$$

$$0 \leq R^{\mathrm{ad}}_{\mathrm{pt,n}}(G,\lambda) - R^{\star,\mathrm{or}}_{\mathrm{pt,n}}(G,\lambda) \longrightarrow 0, \tag{Q.18}$$

$$0 \leq \max_{\eta} \mathbb{E}^{\mathrm{lf}}_{n,G,\eta,\lambda} \operatorname{diam}_{d_\Delta}(\widehat{\mathfrak{E}}_{n,\alpha}) - D^{\star,\mathrm{or}}_{n,\alpha+\varepsilon^{\mathrm{cov}}_n}(G,\lambda) \longrightarrow 0, \tag{Q.19}$$

where $\varepsilon^{\mathrm{cov}}_n = C_{\mathrm{eff}}\mathfrak{E}_n \to 0$. In particular,

$$R^{\star,\mathrm{or}}_{\mathrm{side,n}}(G,\lambda) \geq \Phi(-\kappa_G(\lambda)), \tag{Q.20}$$

$$R^{\star,\mathrm{or}}_{\mathrm{pt,n}}(G,\lambda) = \frac{\Delta^{\mathrm{lf}}_{G,n}(\lambda)^2}{4} r_{\mathrm{pt}}(\kappa_G(\lambda)), \tag{Q.21}$$

$$D^{\star,\mathrm{or}}_{n,\alpha}(G,\lambda) = \Delta^{\mathrm{lf}}_{G,n}(\lambda) q_\alpha(\kappa_G(\lambda)). \tag{Q.22}$$

Thus a learner that is not supplied the interface geometry attains the same first-order side, point, and confidence limits as a learner given the full geometry.

Proof. The lower bounds and equalities in equations (Q.20) to (Q.22) follow from theorem Q.3. The TARGET side and point upper bounds in theorem P.11, stated for equation (Q.3), converge uniformly to the same benchmarks. Since an infimum over oracle-assisted procedures is no larger than the risk of TARGET and no smaller than the displayed oracle lower benchmark, the first two claims follow by squeezing.

For confidence, TARGET has coverage at least $1 - \alpha - \varepsilon^{\mathrm{cov}}_n$. The exact two-point lower bound at that coverage level is

$$\Delta^{\mathrm{lf}}_{G,n}(\lambda) q_{\alpha+\varepsilon^{\mathrm{cov}}_n}(\kappa_G(\lambda)).$$

The function $(\alpha,\kappa) \mapsto q_\alpha(\kappa)$ is uniformly continuous on every compact subset of $(0,1/2) \times [0, \Lambda/s_-])$. The confidence upper bound in theorem P.11 converges uniformly to the same quantity with α, proving equation (Q.19). □

Corollary Q.5 (Global adaptive minimax equality)**.** *Let $\mathfrak{P}_n$ be the class of all non-oracle adaptive procedures that know only the atlas, noise profile, and quantitative class constants. Then*

$$\inf_{\Pi_n \in \mathfrak{P}_n} \sup_{G \in \mathfrak{G}(\Gamma),\, 0 \le \lambda \le \Lambda} R_{\mathrm{side,n}}(\Pi_n; G, \lambda) = \sup_{G \in \mathfrak{G}(\Gamma),\, 0 \le \lambda \le \Lambda} \Phi(-\kappa_G(\lambda)) + o(1), \tag{Q.23}$$

$$\inf_{\Pi_n \in \mathfrak{P}_n} \sup_{G \in \mathfrak{G}(\Gamma),\, 0 \le \lambda \le \Lambda} R_{\mathrm{pt,n}}(\Pi_n; G, \lambda) = \sup_{G \in \mathfrak{G}(\Gamma),\, 0 \le \lambda \le \Lambda} \frac{\Delta^{\mathrm{lf}}_{G,n}(\lambda)^2}{4} r_{\mathrm{pt}}(\kappa_G(\lambda)) + o(1). \tag{Q.24}$$

TARGET attains both displays.

Proof. For every fixed G, λ, a non-oracle procedure is also an admissible procedure in the oracle-assisted lower-bound comparison, so the supremum of the pointwise oracle benchmarks lower-bounds each global minimax value. TARGET is one procedure in $\mathfrak{P}_n$, and the uniform upper bounds in theorem Q.4 give the reverse inequalities up to $o(1)$. □

R An Exact Adaptation-Impossibility Boundary

The positive theory assumes a fixed atlas-isolation margin. This appendix shows that the root-sample-size scale is intrinsic: if two competing simple interfaces are separated by only order $N^{-1/2}$ after N adaptive queries, no method can uniformly recover the active chart, its normal, or its oracle design. The result is proved on an explicit three-by-three zero-sum family. Kernel conditioning, transversality, curvature, ambient interface reach, and point-target separation remain fixed, while the local isolation radius, admissible projection tube, inactive slack, and atlas-isolation margin shrink at the common collision scale.

R.1 Two competing simple interfaces

Assume unit noise in every payoff coordinate. For $0 < \delta \le 1/4$, define

$$B^{(2)}_\delta = \begin{pmatrix} 0 & 2 & 2 \\ 0 & 2 & 2 \\ -\delta & 2 & 2 \end{pmatrix}, \qquad B^{(3)}_\delta = \begin{pmatrix} 0 & 2 & 2 \\ -\delta & 2 & 2 \\ 0 & 2 & 2 \end{pmatrix}. \tag{R.1}$$

In both games, column one is the unique minimizing action. At $B^{(2)}_\delta$, rows one and two tie and row three is lower by δ; at $B^{(3)}_\delta$, rows one and three tie and row two is lower by δ. Let

$$\xi_2 : \text{the interface between } (e_1, e_1) \text{ and } (e_2, e_1), \qquad a_2 = E_{11} - E_{21}, \tag{R.2}$$

$$\xi_3 : \text{the interface between } (e_1, e_1) \text{ and } (e_3, e_1), \qquad a_3 = E_{11} - E_{31}. \tag{R.3}$$

Their unit normals and oracle designs are

$$\nu_2 = (E_{11} - E_{21})/\sqrt{2}, \qquad q^\star_2 = \tfrac{1}{2}(e_{11} + e_{21}), \tag{R.4}$$

$$\nu_3 = (E_{11} - E_{31})/\sqrt{2}, \qquad q^\star_3 = \tfrac{1}{2}(e_{11} + e_{31}). \tag{R.5}$$

Thus

$$\|\nu_2 - \nu_3\|_F = 1, \qquad \|q^\star_2 - q^\star_3\|_1 = 1. \tag{R.6}$$

Lemma R.1 (Quantitative geometry of the competing-interface family)**.** *The two games belong to a sequence of quantitative classes with*

$$r_0(\delta) = \delta/8, \qquad \tau(\delta) = \delta/64, \qquad \gamma_{\mathrm{slack}}(\delta) \ge 3\delta/4, \qquad \gamma_{\mathrm{atlas}}(\delta) = \sqrt{\frac{2}{3}}\,\delta. \tag{R.7}$$

All remaining constants can be chosen independently of δ: the bordered determinants equal one, the active gradients have norm $\sqrt{2}$, the active Hessians vanish, the ambient interfaces have infinite reach, and the two side targets have d_Δ-separation one.

Proof. Columns two and three have payoff two in every row, whereas the maximum entry of column one is zero. Hence column one is the unique minimizer's action with a fixed margin. The row problem is therefore the maximum of the three entries in column one. At $B^{(2)}_\delta$, rows one and two

form the unique active tie and row three has inactive slack δ; the other game is symmetric. On the Frobenius ball of radius $\delta/8$, any payoff difference changes by at most $\delta/4$, so the inactive row slack remains at least $3\delta/4$. The relevant one-by-one bordered matrices have determinant one and uniformly bounded inverses. The active functions are $A_{11} - A_{21}$ and $A_{11} - A_{31}$, so the gradient norm is $\sqrt{2}$, the Hessian is zero, and the interfaces are affine hyperplanes of infinite reach. The pure side targets have product-simplex distance one. Since the active interfaces are affine hyperplanes, metric projection is globally unique; choosing $\tau = \delta/64 \le r_0/8$ verifies the declared local tube condition.

It remains to compute the closest competing candidate interface. For $B_\delta^{(2)}$, the interface between rows one and three requires $x_1 = x_3 \ge x_2$ in the column-one payoff vector $(0, 0, -\delta)$. Its Euclidean projection is $(-\delta/3, -\delta/3, -\delta/3)$, whose squared distance is $2\delta^2/3$. The row-two/row-three interface is identical by symmetry, and no other candidate is closer. Hence the false-chart distance is $\sqrt{2/3}\,\delta$. The calculation for $B_\delta^{(3)}$ is the same. □

R.2 Sequential information bound

Let $P_{2,\delta}^{\pi,N}$ and $P_{3,\delta}^{\pi,N}$ be the transcript laws generated by an arbitrary randomized adaptive policy π using N coordinate queries under the two games. Let $N_{ij}(N)$ denote the number of queries to coordinate (i, j).

Lemma R.2 (Exact adaptive transcript KL)**.** *For every adaptive policy,*

$$D_{\mathrm{KL}}(P_{2,\delta}^{\pi,N} \| P_{3,\delta}^{\pi,N}) = \frac{\delta^2}{2} \mathbb{E}_{2,\delta}^{\pi}\{N_{21}(N) + N_{31}(N)\} \tag{R.8}$$

$$\le \frac{N\delta^2}{2}. \tag{R.9}$$

Proof. Condition on the past transcript and the private random seed. The policy uses the same action kernel under both games. Given the selected coordinate, the conditional laws are unit-variance Gaussians, so their conditional KL is one half of the squared mean difference. The games differ by δ only at coordinates $(2, 1)$ and $(3, 1)$. The chain rule for relative entropy and the tower property give equation (R.8); the count sum is at most N. This is the fixed-horizon Gaussian specialization of the adaptive change-of-measure identity used in bandit lower bounds (Kaufmann et al., 2016). □

Lemma R.3 (Two-experiment testing bound)**.** *For any two laws P, Q and any test $\phi \in \{0, 1\}$,*

$$P\{\phi = 1\} + Q\{\phi = 0\} \ge \frac{1}{2} \exp\{-D_{\mathrm{KL}}(P\|Q)\}. \tag{R.10}$$

Consequently the larger of the two error probabilities is at least $\frac{1}{4} e^{-D_{\mathrm{KL}}(P\|Q)}$.

Proof. If $P \not\ll Q$, the right-hand side is zero and the result is immediate. Otherwise let $L = dP/dQ$ and put $\mathcal{A}(P, Q) = \int \sqrt{dP\, dQ}$, the Hellinger affinity. Jensen's inequality under P gives

$$\mathcal{A}(P, Q) = \mathbb{E}_P[L^{-1/2}] \ge \exp\{-D_{\mathrm{KL}}(P\|Q)/2\}.$$

Cauchy–Schwarz applied to $|\sqrt{dP} - \sqrt{dQ}|(\sqrt{dP} + \sqrt{dQ})$ yields $\|P - Q\|_{\mathrm{TV}} \le \sqrt{1 - \mathcal{A}(P, Q)^2}$. The minimum sum of the two testing errors is $1 - \|P - Q\|_{\mathrm{TV}}$. Since $1 - \sqrt{1 - x} \ge x/2$ for $x \in [0, 1]$, we obtain

$$1 - \|P - Q\|_{\mathrm{TV}} \ge \frac{1}{2}\mathcal{A}(P, Q)^2 \ge \frac{1}{2} e^{-D_{\mathrm{KL}}(P\|Q)}.$$

□

R.3 The root-sample-size adaptation boundary

Theorem R.4 (Competing-interface impossibility)**.** *Let $\widehat{\xi}_N$ be any estimator of the active chart based on an arbitrary adaptive N-query transcript. Then*

$$\inf_{\pi, \widehat{\xi}_N} \max_{j \in \{2,3\}} P_{j,\delta}^{\pi,N}\{\widehat{\xi}_N \ne \xi_j\} \ge \frac{1}{4} \exp\left(-\frac{N\delta^2}{2}\right). \tag{R.11}$$

The same lower bound holds for every unit-normal estimator in the form

$$\inf_{\pi,\widehat{\nu}_N} \max_{j\in\{2,3\}} P^{\pi,N}_{j,\delta}\left\{\|\widehat{\nu}_N - \nu_j\|_F \geq \frac{1}{2}\right\} \geq \frac{1}{4}\exp\left(-\frac{N\delta^2}{2}\right), \tag{R.12}$$

and for every oracle-design estimator,

$$\inf_{\pi,\widehat{q}_N} \max_{j\in\{2,3\}} P^{\pi,N}_{j,\delta}\left\{\|\widehat{q}_N - q^\star_j\|_1 \geq \frac{1}{2}\right\} \geq \frac{1}{4}\exp\left(-\frac{N\delta^2}{2}\right). \tag{R.13}$$

The infima range jointly over all adaptive policies and terminal estimators.

Proof. Apply lemma R.3 with $P = P^{\pi,N}_{2,\delta}$, $Q = P^{\pi,N}_{3,\delta}$, and use equation (R.9). This proves the chart claim for every fixed policy and hence after optimizing over policies.

For normal estimation, the two true unit normals are distance one apart. If $\|\widehat{\nu}_N - \nu_j\|_F < 1/2$, nearest-normal decoding recovers j uniquely. Thus an estimator violating equation (R.12) would induce a chart test violating equation (R.11). The oracle designs are also distance one apart in ℓ_1, so the identical reduction proves equation (R.13). □

The theorem is about arbitrary adaptive querying, not only the two-stage TARGET procedure. It says that a vanishing atlas margin cannot be hidden inside a model-selection proof: root-N separation is information-theoretically necessary. Adaptation lower bounds are known in other sequential learning problems, including adaptation to unknown smoothness (Locatelli & Carpentier, 2018); the present boundary is specific to collision of game-derived equilibrium-interface charts and their oracle designs.

Proposition R.5 (Balanced-query attainment on the collision family)**.** *Let $N \geq 2$, and put $m_2 = \lfloor N/2 \rfloor$ and $m_3 = N - m_2$, and query coordinates $(2,1)$ and $(3,1)$ exactly m_2 and m_3 times. Let $\overline{Y}_{21}$ and $\overline{Y}_{31}$ be the two sample means and select chart two when $\overline{Y}_{21} - \overline{Y}_{31} \geq 0$, chart three otherwise. Its exact maximal chart error is*

$$\Phi\left(-\frac{\delta}{\sqrt{m_2^{-1} + m_3^{-1}}}\right). \tag{R.14}$$

For even N, this is $\Phi(-\delta\sqrt{N}/2)$. Outputting the normal and oracle design associated with the selected chart has the same failure probability. Consequently, on the explicit collision family, chart, normal, and oracle-design consistency are attainable if and only if

$$N\delta_N^2 \longrightarrow \infty. \tag{R.15}$$

Proof. The two sample means are independent. Under game two, $\overline{Y}_{21} - \overline{Y}_{31}$ is Gaussian with mean δ and variance $m_2^{-1} + m_3^{-1}$; under game three its mean is $-\delta$ with the same variance. The sign rule is therefore the symmetric likelihood-ratio test and has the displayed error. Since $m_2^{-1} + m_3^{-1} = 4/N + O(N^{-2})$, the error tends to zero exactly when $N\delta_N^2 \to \infty$. The normal and design are deterministic functions of the selected chart, so their recovery errors coincide with the chart error. Necessity follows from theorem R.4. □

R.4 A POSITIVE ORACLE-DESIGN EFFICIENCY GAP

For a design q, let

$$V_2(q) = \frac{1}{q_{11}} + \frac{1}{q_{21}}, \qquad V_3(q) = \frac{1}{q_{11}} + \frac{1}{q_{31}}, \tag{R.16}$$

with the extended-value convention at zero mass. Each oracle minimum equals four.

Lemma R.6 (Exact robust-design penalty)**.** *One design chosen without knowing which interface is active satisfies*

$$\Gamma_0 := \min_{q\in\Delta_9} \max_{j\in\{2,3\}} \frac{V_j(q)}{4} = \frac{3+2\sqrt{2}}{4} > 1. \tag{R.17}$$

The minimizer has

$$q_{11} = \frac{1}{1+\sqrt{2}}, \qquad q_{21} = q_{31} = \frac{1}{2+\sqrt{2}}, \tag{R.18}$$

and zero mass on irrelevant coordinates.

Proof. Mass on any irrelevant coordinate can be moved to a relevant coordinate without increasing either variance. Write $x = q_{11}$, $y = q_{21}$, $z = q_{31}$, so $x + y + z = 1$. Then

$$\max\{V_2, V_3\} = \frac{1}{x} + \frac{1}{\min\{y, z\}} \geq \frac{1}{x} + \frac{2}{1-x},$$

with equality at $y = z = (1-x)/2$. Differentiation gives $x = (1+\sqrt{2})^{-1}$, and substitution yields $3 + 2\sqrt{2}$. Divide by the oracle value four. □

Theorem R.7 (No uniform oracle-design adaptation below the atlas threshold)**.** *Fix any*

$$0 < \varepsilon < \Gamma_0 - 1. \tag{R.19}$$

For every random design $\widehat{q}_N$ produced from an arbitrary adaptive N-query transcript,

$$\max_{j \in \{2,3\}} P^{\pi,N}_{j,\delta} \left\{ \frac{V_j(\widehat{q}_N)}{4} - 1 > \varepsilon \right\} \geq \frac{1}{4} \exp\left(-\frac{N\delta^2}{2} \right). \tag{R.20}$$

Consequently, if $\sup_N N\delta_N^2 < \infty$, no procedure can have

$$\frac{V_j(\widehat{q}_N)}{\inf_q V_j(q)} \xrightarrow{P_{j,\delta_N}} 1 \quad \textit{uniformly over } j \in \{2,3\}. \tag{R.21}$$

Proof. For $j \in \{2,3\}$, let

$$\mathcal{Q}_j(\varepsilon) = \{q : V_j(q)/4 \leq 1 + \varepsilon\}.$$

By lemma R.6 and $1 + \varepsilon < \Gamma_0$, the two sets are disjoint. Define a chart test that outputs two when $\widehat{q}_N \in \mathcal{Q}_2(\varepsilon)$ and three otherwise. Under game two its error is at most $P_2\{\widehat{q}_N \notin \mathcal{Q}_2(\varepsilon)\}$. Under game three, disjointness implies $\{\widehat{q}_N \in \mathcal{Q}_2(\varepsilon)\} \subseteq \{\widehat{q}_N \notin \mathcal{Q}_3(\varepsilon)\}$, so its error is bounded by the corresponding inefficiency probability. Apply equation (R.11). The final assertion follows along any subsequence on which $N\delta_N^2$ is bounded. □

Corollary R.8 (Exact collision threshold and class-wide necessary condition)**.** *On the explicit two-game collision family, chart, unit-normal, and oracle-design consistency are possible if and only if $N\delta_N^2 \to \infty$. For uniform recovery or oracle-design efficiency on any class containing this family, it is therefore necessary that*

$$N\delta_N^2 \longrightarrow \infty. \tag{R.22}$$

Since $\gamma_{\mathrm{atlas},N} = \sqrt{2/3}\,\delta_N$ in lemma R.1, this is equivalent to

$$N\gamma_{\mathrm{atlas},N}^2 \longrightarrow \infty. \tag{R.23}$$

For the TARGET pilot, $N = n_0$; thus $n_0\gamma_{\mathrm{atlas},n}^2 \to \infty$ is information-theoretically necessary for uniform geometry recovery. The positive recovery theorem has the same root-n_0 separation scale, up to the confidence logarithm used to drive the failure probability to zero.

Remark R.9 (What the boundary does and does not say). The impossibility is caused by two competing regular simple interfaces colliding and becoming statistically indistinguishable. The local isolation radius and inactive margin necessarily shrink in this construction, but the bordered matrices remain uniformly regular, transversality stays nonzero, contact remains first order, and point-target separation stays fixed. Those mechanisms can generate additional adaptation boundaries, but they are not needed for the root-N atlas threshold proved here.

S THE EXACT PRICE OF POST-HOC TARGET IGNORANCE

The local-oracle conflict in theorem I.6 compares two particular equilibrium targets. This appendix replaces that example-level message by a general theorem. The protocol is deliberately post hoc: the learner chooses the sampling design before learning which downstream target will be requested. Generic maxi–min and multi-objective design criteria have an established literature (Tommasi et al., 2023); the manuscript-specific objects below are the equilibrium-derived information profiles and their consequences for value, interface, strategy, and policy targets.

S.1 Coordinate Gaussian targets and normalized information profiles

Work in the unrestricted entrywise Gaussian score experiment from theorem I.3. There are d payoff coordinates, coordinate i has noise standard deviation $\sigma_i > 0$, and $q \in \Delta_d$ is an approximate design. Let $r \in [R]$ index possible post-hoc targets. Target r has derivative columns $d_{r,i}$ and positive semidefinite loss matrix W_r. Define

$$c_{r,i} \coloneqq \sigma_i \|d_{r,i}\|_{W_r}, \qquad C_r \coloneqq \sum_{i=1}^{d} c_{r,i}. \tag{S.1}$$

A target with $C_r = 0$ is first-order constant and has zero Gaussian linear risk under every design. Such a target imposes no post-hoc design constraint. Delete all zero targets and relabel the remaining set so that $C_r > 0$ for every $r \in [R]$. Put

$$p_{r,i} \coloneqq \frac{c_{r,i}}{C_r}, \qquad p_r \in \Delta_d. \tag{S.2}$$

The exact target-r Gaussian linear minimax risk is

$$\Phi_r(q) = \sum_{i=1}^{d} \frac{c_{r,i}^2}{q_i}, \tag{S.3}$$

with the extended conventions $0/0 = 0$ and $c/0 = +\infty$ for $c > 0$. By theorem I.3,

$$\Phi_r^\star = \inf_{q \in \Delta_d} \Phi_r(q) = C_r^2, \qquad q_r^\star = p_r. \tag{S.4}$$

Thus the normalized inefficiency of q for target r is

$$F_r(q) \coloneqq \frac{\Phi_r(q)}{\Phi_r^\star} = \sum_{i=1}^{d} \frac{p_{r,i}^2}{q_i}. \tag{S.5}$$

Definition S.1 (Price of target ignorance)**.** The learner chooses q before an adversary reveals $r \in [R]$. The exact relative price of target ignorance is

$$\Gamma(p_1, \dots, p_R) \coloneqq \min_{q \in \Delta_d} \max_{r \in [R]} F_r(q). \tag{S.6}$$

Let

$$U \coloneqq \{i : \max_r p_{r,i} > 0\} \tag{S.7}$$

be the union of target-relevant coordinates.

S.2 Exact minimax dual and robust design

Theorem S.2 (Exact target-ignorance dual)**.** *For arbitrary profiles $p_1, \dots, p_R \in \Delta_d$,*

$$\boxed{\Gamma(p_1, \dots, p_R) = \max_{\lambda \in \Delta_R} \left(\sum_{i=1}^{d} \sqrt{\sum_{r=1}^{R} \lambda_r p_{r,i}^2} \right)^2.} \tag{S.8}$$

The primal minimum and dual maximum are attained. Every primal optimizer puts zero mass outside U and strictly positive mass on every coordinate in U. There exists a dual optimizer $\lambda^\star \in \Delta_R$ associated with a primal optimizer $q^\star$ such that, with

$$b_i^\star \coloneqq \sum_{r=1}^{R} \lambda_r^\star p_{r,i}^2, \tag{S.9}$$

one has $b_i^\star > 0$ for every $i \in U$ and

$$\boxed{q_i^\star = \frac{\sqrt{b_i^\star}}{\sum_{j \in U} \sqrt{b_j^\star}} \quad (i \in U), \qquad q_i^\star = 0 \quad (i \notin U).} \tag{S.10}$$

Moreover,

$$\lambda_r^\star > 0 \quad \Longrightarrow \quad F_r(q^\star) = \Gamma. \tag{S.11}$$

Thus every target receiving positive dual weight is exactly worst case at the robust design.

Proof. Mass outside U contributes to no target and can be moved to a relevant coordinate without increasing any F_r, so the problem reduces to Δ_U. If $q_i \downarrow 0$ for any $i \in U$, then some target has $p_{r,i} > 0$ and hence $F_r(q) \to +\infty$. Therefore the primal objective is coercive at the boundary of Δ_U and attains its minimum at an interior point.

Write the epigraph form

$$\min_{q \in \mathrm{ri}\,\Delta_U,\ t \in \mathbb{R}} t \quad \text{subject to} \quad F_r(q) \le t, \quad r \in [R]. \tag{S.12}$$

It is a convex program and satisfies Slater's condition: take the uniform design on U and choose t strictly larger than every finite F_r. Its Lagrangian is

$$t\left(1 - \sum_r \lambda_r\right) + \sum_{i \in U} \frac{b_i(\lambda)}{q_i}, \qquad b_i(\lambda) = \sum_r \lambda_r p_{r,i}^2,$$

with multipliers $\lambda_r \ge 0$. Infimization over t is finite exactly when $\lambda \in \Delta_R$. For fixed such λ, sharp Cauchy–Schwarz gives

$$\sum_{i \in U} \frac{b_i(\lambda)}{q_i} \ge \left(\sum_{i \in U} \sqrt{b_i(\lambda)}\right)^2, \tag{S.13}$$

with equality, in the extended infimum sense, at $q_i \propto \sqrt{b_i(\lambda)}$. Strong convex duality proves equation (S.8). The dual objective is continuous on the compact simplex Δ_R, so a maximizer exists.

Let $(q^\star, t^\star)$ be primal optimal. KKT gives a dual optimal $\lambda^\star$ and a multiplier ν for $\sum_i q_i = 1$ such that

$$\frac{b_i^\star}{(q_i^\star)^2} = \nu, \qquad i \in U. \tag{S.14}$$

Because $q_i^\star > 0$, every $b_i^\star$ is positive. Normalizing equation (S.14) yields equation (S.10). Complementary slackness gives equation (S.11). □

Corollary S.3 (Sharp compatibility and worst-case price)**.** *For every collection of nonzero target profiles,*

$$\boxed{1 \le \Gamma(p_1, \dots, p_R) \le R.} \tag{S.15}$$

Furthermore,

$$\boxed{\Gamma(p_1, \dots, p_R) = 1 \iff p_1 = p_2 = \cdots = p_R.} \tag{S.16}$$

If the supports of $p_1, \dots, p_R$ are pairwise disjoint, then

$$\boxed{\Gamma(p_1, \dots, p_R) = R.} \tag{S.17}$$

Proof. For every r and q, Cauchy–Schwarz gives

$$F_r(q) = \sum_i \frac{p_{r,i}^2}{q_i} \ge \left(\sum_i p_{r,i}\right)^2 = 1.$$

Equality holds if and only if $q = p_r$, including zero coordinates. Thus $\Gamma \ge 1$, and $\Gamma = 1$ exactly when one design equals every profile.

For the upper bound, use $\bar{q} = R^{-1} \sum_{s=1}^R p_s$. Whenever $p_{r,i} > 0$, $\bar{q}_i \ge p_{r,i}/R$, so

$$F_r(\bar{q}) \le R \sum_i p_{r,i} = R.$$

Finally, under pairwise disjoint supports, choose $\lambda_r = 1/R$ in equation (S.8). Then

$$\sum_i \sqrt{\sum_r \lambda_r p_{r,i}^2} = \frac{1}{\sqrt{R}} \sum_r \sum_{i \in \mathrm{supp}(p_r)} p_{r,i} = \sqrt{R}.$$

The dual lower bound is therefore R, matching the general upper bound. □

Proposition S.4 (Integer implementation)**.** *Fix the profiles and let $q^\star$ be any robust optimizer. Put $q^\star_{\min} = \min_{i\in U} q^\star_i > 0$. Largest-remainder rounding to N queries gives proportions $q^{(N)}$ with*

$$\max_r F_r(q^{(N)}) \le \Gamma + \frac{2}{N(q^\star_{\min})^2} \max_r \sum_{i\in U} p_{r,i}^2 \tag{S.18}$$

for every $N \ge 2/q^\star_{\min}$. Hence the post-hoc target-ignorance penalty is implementable to $O(N^{-1})$.

Proof. Largest-remainder rounding gives $|q_i^{(N)} - q_i^\star| < N^{-1}$ and $q_i^{(N)} \ge q^\star_{\min}/2$ on U. Thus

$$\left| \frac{1}{q_i^{(N)}} - \frac{1}{q_i^\star} \right| \le \frac{2}{N(q^\star_{\min})^2}.$$

Multiply by $p_{r,i}^2$, sum, and take the maximum over r. □

S.3 Weighted targets and exact classification consequences

For weights $w_r > 0$, define

$$\Gamma_w \coloneqq \min_{q\in\Delta_d} \max_r w_r F_r(q). \tag{S.19}$$

The same proof gives

$$\boxed{\Gamma_w = \max_{\lambda\in\Delta_R} \left(\sum_i \sqrt{\sum_r \lambda_r w_r p_{r,i}^2} \right)^2.} \tag{S.20}$$

Corollary S.5 (Multiple interface-side targets)**.** *For interface target r, let $a_r \ne 0$ be its normal, let*

$$C_r = \sum_i \sigma_i |a_{r,i}|, \qquad p_{r,i} = \frac{\sigma_i |a_{r,i}|}{C_r}, \qquad \kappa_r^\star = \frac{\lambda_r}{C_r} > 0. \tag{S.21}$$

Under design q, the exact fixed-halfspace error is

$$B_r(q) = \Phi\left(-\frac{\kappa_r^\star}{\sqrt{F_r(q)}} \right). \tag{S.22}$$

Consequently, with $w_r = (\kappa_r^\star)^{-2}$,

$$\boxed{\min_q \max_r B_r(q) = \Phi\left(-\frac{1}{\sqrt{\Gamma_w}} \right).} \tag{S.23}$$

If every target has the same oracle phase $\kappa_r^\star = \kappa$, the robust design is exactly the optimizer in theorem S.2 and

$$\min_q \max_r B_r(q) = \Phi\left(-\frac{\kappa}{\sqrt{\Gamma}} \right). \tag{S.24}$$

For unequal phases, the unweighted robust design still guarantees

$$B_r(q^\star) \le \Phi\left(-\frac{\kappa_r^\star}{\sqrt{\Gamma}} \right) \qquad \textit{for every } r. \tag{S.25}$$

Proof. The semidefinite halfspace theorem and diagonal information formula give $V_r(q) = C_r^2 F_r(q)$, hence equation (S.22). Because $x \mapsto \Phi(-1/\sqrt{x})$ is increasing on $(0,\infty)$,

$$\max_r B_r(q) = \Phi\left(-\frac{1}{\sqrt{\max_r w_r F_r(q)}} \right).$$

Minimization and equation (S.20) prove the exact identity. The remaining statements are immediate. □

S.4 EQUILIBRIUM-DERIVED PROFILES

Proposition S.6 (Value, interface, strategy, and policy profiles)**.** *In the unrestricted entrywise Gaussian experiment, the following are exact instances of theorem S.2.*

1. *At a unique equilibrium $(x^\star, y^\star)$, the value profile is*
$$p_{ij}^{\mathrm{val}} = \frac{\sigma_{ij} x_i^\star y_j^\star}{\sum_{a,b} \sigma_{ab} x_a^\star y_b^\star}. \tag{S.26}$$
2. *At a simple interface with normal a, the side profile is*
$$p_{ij}^{\mathrm{side}} = \frac{\sigma_{ij} |a_{ij}|}{\sum_{a,b} \sigma_{ab} |a_{ab}|}. \tag{S.27}$$
3. *For a regular point-strategy target with derivative columns d_{ij} and quadratic metric W, the point profile is*
$$p_{ij}^{\mathrm{pt}} = \frac{\sigma_{ij} \|d_{ij}\|_W}{\sum_{a,b} \sigma_{ab} \|d_{ab}\|_W}. \tag{S.28}$$
4. *Any differentiable finite-horizon value or policy target in a declared local coordinate model has the same form after replacing d_{ij} by the corresponding derivative column and W by the declared policy metric.*

A single entrywise experiment is simultaneously oracle optimal for every listed nonzero target if and only if their normalized profiles coincide. Pairwise disjoint profiles incur the maximal relative penalty equal to the number of targets.

Proof. The value derivative is $D\,\mathrm{val}(A)[H] = x^{\star\mathsf{T}} H y^\star$, so its coordinate columns are $x_i^\star y_j^\star$. The side variance uses the scalar derivative a_{ij}. The point and policy statements are direct applications of equation (S.1). Apply corollary S.3. □

S.5 THE EXISTING DESIGN-CONFLICT FAMILY AS A COROLLARY

The exact profile dual above uses unrestricted direct coordinate information. For completeness, the same target-ignorance principle extends to arbitrary compact design families and target criteria.

Proposition S.7 (General compatibility principle)**.** *Let $\mathcal{Q}$ be compact and let $R_r : \mathcal{Q} \to (0, \infty]$ be lower-semicontinuous target criteria with finite positive oracle values $R_r^\star = \min_{q \in \mathcal{Q}} R_r(q)$. Define*
$$\Gamma_{\mathcal{Q}} = \min_{q \in \mathcal{Q}} \max_r \frac{R_r(q)}{R_r^\star}. \tag{S.29}$$
Then
$$\Gamma_{\mathcal{Q}} \geq 1, \qquad \Gamma_{\mathcal{Q}} = 1 \Longleftrightarrow \bigcap_r \operatorname*{argmin}_{q \in \mathcal{Q}} R_r(q) \neq \varnothing. \tag{S.30}$$

Proof. Every normalized criterion is at least one. Equality of their maximum to one holds exactly when one design attains every individual minimum. □

Corollary S.8 (Robust post-hoc design in the declared value–side family)**.** *In theorem I.6, normalize the value criterion by its oracle minimum and normalize the side normal variance by its oracle minimum:*
$$E_{\mathrm{val}}(\alpha) = \frac{1}{1 - \alpha + \alpha / r_{\mathrm{noise}}}, \qquad E_{\mathrm{side}}(\alpha) = \frac{1}{\alpha}, \tag{S.31}$$
with $E_{\mathrm{side}}(0) = +\infty$. If the target is revealed after the design is chosen, then
$$\boxed{\min_{\alpha \in [0,1]} \max\{E_{\mathrm{val}}(\alpha), E_{\mathrm{side}}(\alpha)\} = 2 - \frac{1}{r_{\mathrm{noise}}}.} \tag{S.32}$$

The unique robust allocation is

$$\boxed{\alpha^\star = \frac{r_{\mathrm{noise}}}{2r_{\mathrm{noise}} - 1}.} \tag{S.33}$$

At this allocation the two normalized criteria are equal. The corresponding side-classification error at separation c is

$$\Phi\left(-\frac{c}{\sqrt{2r_{\mathrm{noise}} - 1}}\right). \tag{S.34}$$

Thus the former conflict example is a concrete corollary of the general post-hoc target protocol rather than the conceptual endpoint of the design theory.

Proof. The value and side oracle minima are 1 and r_{noise}, respectively, so the normalized criteria are exactly those in equation (S.31). The first is strictly increasing and the second strictly decreasing. Their minimax intersection solves

$$\frac{1}{1 - \alpha + \alpha/r_{\mathrm{noise}}} = \frac{1}{\alpha},$$

which gives equation (S.33); the common value is $1/\alpha^\star = 2 - 1/r_{\mathrm{noise}}$. Substitution into equation (I.30) gives equation (S.34). □

S.6 Novelty boundary

The duality in theorem S.2 is a closed-form instance of maxi–min multi-objective design and should not be advertised as the first general equivalence theorem of that kind. The equilibrium-specific content is that value occupancy, interface normals, strategy derivatives, and policy metrics generate distinct profiles under the same payoff-query experiment. The compatibility theorem then gives an exact statement of when "learning the game" admits one target-free experiment and when it does not.

T Exact Theorem Domain: Quantitative Simple Interfaces

This appendix resolves the prevalence question conservatively. We do not claim that simple transverse equilibrium interfaces are Lebesgue-generic, Zariski-generic, or Baire-generic in the full payoff space. Instead, every adaptive theorem is stated on the explicit quantitative class $\mathfrak{G}(\Gamma)$ from definition N.1. We prove that this domain is nonvacuous, locally robust, and exactly the fixed-margin envelope of the qualitative simple-interface instances used by the geometry theory. This is sufficient for the uniform adaptive upper and lower bounds, and avoids silently replacing a quantitative theorem by an unproved genericity claim.

T.1 Qualitative simple interfaces

Definition T.1 (Strict qualitative simple-interface instance)**.** A pair $G = (B, \xi_\star)$ is a strict qualitative simple-interface instance if there exists a neighborhood $\mathcal{O}$ of B such that:

1. the two incident kernel branches in $\xi_\star$ are regular and analytic on $\mathcal{O}$;
2. the active coordinate g_G satisfies $g_G(B) = 0$ and $\nabla g_G(B) \neq 0$;
3. all inactive validity coordinates of both branches are strictly positive at B, and the opposite active coordinate is $-\chi_G g_G$ for an analytic function with $\chi_G(B) > 0$;
4. the two branches are the unique equilibria on the two open sides in a neighborhood of B;
5. no false candidate-interface set $\mathcal{C}_\xi$, $\xi \neq \xi_\star$, contains B;
6. the two side targets at B are distinct.

The definition is qualitative: it asserts strictness at one instance but does not fix uniform numerical margins across a class. Let $\mathfrak{G}_{\mathrm{simple}}$ denote the collection of all such instances for the fixed game dimensions, atlas, and noise bounds.

Theorem T.2 (Quantitative envelope theorem)**.** *Fix any finite local-radius bound $H < \infty$.*

1. *Every strict qualitative simple-interface instance belongs to $\mathfrak{G}(\Gamma_G)$ for some tuple of strictly positive finite constants Γ_G after shrinking its neighborhood.*
2. *Conversely, every $G \in \mathfrak{G}(\Gamma)$ is a strict qualitative simple-interface instance.*

Consequently,

$$\boxed{\mathfrak{G}_{\mathrm{simple}} = \bigcup_{\Gamma \in \mathfrak{P}} \mathfrak{G}(\Gamma),} \tag{T.1}$$

where $\mathfrak{P}$ ranges over all admissible positive class-constant tuples with the fixed dimensions and noise bounds.

Proof. Fix a qualitative instance. Regularity of the two bordered systems at B and continuity give a closed ball on which their determinants remain nonzero and their inverses remain bounded. The nonzero gradient condition and continuity give positive lower and finite upper bounds on $\|\nabla g_G\|_F$; analyticity gives a finite Hessian bound on a smaller closed ball.

Every inactive validity coordinate is positive at B, so finitely many continuity arguments give one common positive slack margin. Likewise, $\chi_G(B) > 0$ yields positive lower and finite upper bounds for χ_G. The finite atlas and the assumption $B \notin \mathcal{C}_\xi$ for every false chart imply

$$\min_{\xi \neq \xi_\star} \mathrm{dist}_F(B, \mathcal{C}_\xi) > 0,$$

because each $\mathcal{C}_\xi$ is closed and there are finitely many of them. This gives a positive atlas margin. Distinct side targets and continuity give a positive target-separation margin and a finite branch Lipschitz constant on a sufficiently small closed ball.

Finally, the regular level-set theorem makes $g_G^{-1}(0)$ a C^2 embedded hypersurface near B. A sufficiently small compact patch of a C^2 embedded hypersurface has a tubular neighborhood with unique nearest-point projection; equivalently it has positive local reach in the sense of Federer (1959). Shrinking once more gives the projection-tube condition in definition N.1. Choose $M > \|B\|_\infty$ and collect all resulting bounds into Γ_G. This proves the first direction. The converse follows immediately from conditions (Q1)–(Q7). □

T.2 Local robustness of a fixed-margin class

Theorem T.3 (Relative openness along an interface patch)**.** *Let $G = (B, \xi_\star) \in \mathfrak{G}(\Gamma)$. There exists*

$$\varepsilon_\Gamma = \min\left\{ \frac{r_0}{4}, \frac{\gamma_{\mathrm{atlas}}}{2}, \frac{\gamma_{\mathrm{sep}}}{4L_z} \right\} > 0 \tag{T.2}$$

such that every $\widetilde{B} \in \mathcal{M}_G^{\mathrm{in}}$ with $\|\widetilde{B} - B\|_F < \varepsilon_\Gamma$ defines $\widetilde{G} = (\widetilde{B}, \xi_\star)$ belonging to a quantitative class $\mathfrak{G}(\Gamma^-)$ with

$$r_0^- = r_0/2, \quad \tau^- = \tau/2, \quad \gamma_{\mathrm{atlas}}^- = \gamma_{\mathrm{atlas}}/2, \quad \gamma_{\mathrm{sep}}^- = \gamma_{\mathrm{sep}}/2, \tag{T.3}$$

and all remaining lower bounds no smaller, upper bounds no larger, except that the payoff bound may be enlarged from M to $M + r_0$. Thus fixed-margin simple interfaces form a relatively open family along every certified interface patch.

Proof. The ball $\overline{B}_F(\widetilde{B}, r_0/2)$ is contained in $\overline{B}_F(B, r_0)$ when $\|\widetilde{B} - B\|_F < r_0/4$. Hence all determinant, inverse, gradient, Hessian, inactive-slack, coherent-side, and branch-Lipschitz bounds remain valid.

Because $\widetilde{B} \in \mathcal{C}_{\xi_\star}$, $\mathsf{S}_{\xi_\star}(\widetilde{B}) = 0$. For every false chart, the one-Lipschitz property of distance to a closed set gives

$$\mathsf{S}_\xi(\widetilde{B}) \geq \mathsf{S}_\xi(B) - \|\widetilde{B} - B\|_F \geq \gamma_{\mathrm{atlas}}/2.$$

The target separation satisfies

$$d_\Delta(z_{G,+}(\widetilde{B}), z_{G,-}(\widetilde{B})) \geq \gamma_{\mathrm{sep}} - 2L_z\|\widetilde{B} - B\|_F \geq \gamma_{\mathrm{sep}}/2.$$

It remains to verify the reduced projection tube. The inner patch for $\widetilde{G}$ with radius $r_0^-/2 = r_0/4$ lies inside $\mathcal{M}_G^{\mathrm{in}}$. Hence every point within distance $\tau/2$ of that patch has the unique original projection. If C is that projection and C_0 is a point of the reduced inner patch witnessing the tube distance, then

$$\|C - \widetilde{B}\|_F \le \|C - x\|_F + \|x - C_0\|_F + \|C_0 - \widetilde{B}\|_F < \tau + r_0/4 \le 3r_0/8 < r_0/2.$$

Thus the projection belongs to the reduced patch $\overline{B}_F(\widetilde{B}, r_0/2)$, proving (Q7) for Γ^-. The payoff bound follows from $\|\widetilde{B}\|_\infty \le M + r_0$. □

T.3 Exact domain of the adaptive theorems

Theorem T.4 (Fixed-margin theorem domain)**.** *Fix one admissible tuple Γ, one local-radius bound H, and one compact phase range $[0, \Lambda]$. Then:*

1. *the geometry-recovery theorem in Appendix O is uniform over exactly $G \in \mathfrak{G}(\Gamma)$ and $\|h\|_F \le H$;*
2. *the adaptive Gaussian, design, side-risk, point-risk, and confidence theorems in Appendix P are uniform over the same fixed class and the displayed compact phase range;*
3. *the matching lower bounds in Appendix Q use the same class and range;*
4. *no one of these statements asserts uniformity over the uncontrolled union $\bigcup_\Gamma \mathfrak{G}(\Gamma)$ or over sequences whose margins vanish without satisfying the displayed error conditions.*

In particular, for a sequence of classes with atlas margin $\gamma_{\mathrm{atlas},n}$, uniform chart recovery requires at least

$$n_0 \gamma_{\mathrm{atlas},n}^2 \longrightarrow \infty \tag{T.4}$$

on the explicit collision family of Appendix R. Therefore a theorem uniform over all qualitative simple interfaces, with no quantitative margin, is false.

Proof. The first three items are the quantifiers in theorems O.5, P.7, P.11 and Q.4. Their deterministic error moduli depend on the fixed constants in Γ. The competing-interface construction in Appendix R keeps determinant conditioning, transversality, reach, and target separation nondegenerate while letting the atlas margin shrink proportionally to δ_n. Its exact lower bound shows that chart recovery is impossible when $n_0\delta_n^2$ remains bounded, and $\gamma_{\mathrm{atlas},n} \asymp \delta_n$. This proves equation (T.4) and rules out margin-free uniformity. □

Remark T.5 (No unproved prevalence claim). Semialgebraic sets admit finite Whitney stratifications (Shiota, 2005), and generic transversality results are available in definable categories. Those facts do not by themselves prove that the particular regular-kernel, no-singular-branch, two-germ equilibrium interfaces required here have full measure or form a Zariski-open set. The paper therefore uses theorems T.2 and T.4 instead of claiming a global prevalence theorem. Any future genericity result must specify the ambient parameter space, the meaning of genericity, and the codimension of every excluded singular-kernel and multi-interface event.

U Cross-Game Universality: A Certified Generator and Study

This appendix supplies the theorem-verification study used in the main paper. The study is not a benchmark on external data. We first construct a high-dimensional family of certified equilibrium interfaces, prove that every member has the same normalized Gaussian decision experiment, and only then use Monte Carlo calculations to check the implementation of the exact formulas. The construction varies dimension, active support, branch separation, noise, interface orientation, and information profiles while preserving the theorem scope of theorem T.4.

U.1 A support-splitting interface generator

Fix integers $s \ge 2$ and $d \ge s + 1$. Let $x, y \in \mathrm{ri}(\Delta_s)$, and let $C \in \mathbb{R}^{s\times s}$ satisfy

$$Cy = 0, \qquad x^\top C = 0, \qquad \mathrm{rank}(C) = s - 1. \tag{U.1}$$

A convenient parametrization is

$$C = U_x M V_y^\top, \tag{U.2}$$

where the columns of U_x and V_y are orthonormal bases of $x^\perp$ and $y^\perp$, respectively, and $M \in \mathbb{R}^{(s-1)\times(s-1)}$ is nonsingular. Thus the nonzero singular values of C can be prescribed through M.

Choose an index $p \in [s]$. In the *row-split construction*, replace row p of C by two identical rows, denoted r_+ and r_-, and retain the remaining $s-1$ rows. This gives an $(s+1)\times s$ active block. Embed it in a $d \times d$ matrix B by assigning a constant $P > \|C\|_\infty$ to every active-row/inactive-column entry, assigning $-P$ to every inactive-row/active-column entry, and assigning zero to the inactive–inactive block. Arbitrary row and column permutations are allowed. Define

$$\begin{aligned} z_+ &= (x_+, y), & z_- &= (x_-, y), \\ (x_+)_{r_+} &= x_p, & (x_-)_{r_-} &= x_p, \end{aligned} \tag{U.3}$$

with the remaining active row masses inherited from x and all omitted masses zero. The column-split construction is obtained symmetrically by duplicating one active column; its two side targets share row strategy x and place the mass y_p on opposite duplicate columns.

Lemma U.1 (Unique interior core saddle)**.** *Under equation* (U.1)*, (x, y) is the unique saddle point of C, and its value is zero.*

Proof. The identities $Cy = 0$ and $x^\top C = 0$ show that (x, y) is a saddle point of value zero. If $u \in \Delta_s$ is row-optimal, then $u^\top C e_j \geq 0$ for every j. Since $y_j > 0$ and

$$\sum_j y_j u^\top C e_j = u^\top C y = 0,$$

every inequality is an equality, so $u^\top C = 0$. The left nullspace of C is one-dimensional because $\operatorname{rank}(C) = s-1$; normalization gives $u = x$. The same argument, with the roles reversed, gives the uniqueness of y. □

Proposition U.2 (Certified split interface)**.** *Assume that the rows and columns of C are pairwise distinct. The row-split construction has value zero and equilibrium set*

$$\mathcal{E}(B) = \left\{ \left(x^{(\theta)}, y\right) : 0 \leq \theta \leq 1 \right\}, \qquad x^{(\theta)}_{r_+} = \theta x_p, \quad x^{(\theta)}_{r_-} = (1-\theta) x_p, \tag{U.4}$$

with all other active masses fixed by x. The two endpoint kernels define regular analytic branches with limiting targets z_+ and z_-. A canonical validity coordinate g can be chosen so that

$$Dg(B)[H] = \sum_{j \in J} y_j (H_{r_+,j} - H_{r_-,j}) = \langle a, H \rangle_F, \tag{U.5}$$

where J is the active column set. In particular,

$$\|a\|_F > 0, \qquad d_\Delta(z_+, z_-) = x_p > 0. \tag{U.6}$$

After shrinking a full-dimensional neighborhood of B, the positive side has unique equilibrium branch $z_+(A)$ and the negative side has unique branch $z_-(A)$; no third regular branch is incident. Hence $(B, \xi_\star)$ is a strict qualitative simple-interface instance in the sense of definition T.1. The column-split construction has normal

$$Dg(B)[H] = \sum_{i \in I} x_i (H_{i,c_-} - H_{i,c_+}) \tag{U.7}$$

and target separation y_p.

Proof. We prove the row statement; the column statement follows by applying the same argument to $-B^\top$.

First, distribute the mass x_p arbitrarily across the two duplicate rows. Against the active columns, the resulting row mixture is exactly $x^\top C = 0$; against an inactive column it earns P. Thus it guarantees value zero. Conversely, the active column strategy y gives payoff zero to every active

row and payoff $-P$ to every inactive row, so it holds the value to zero. Therefore every point in equation (U.4) is an equilibrium.

Now let u be any row-optimal strategy. Against y, its payoff is $-P$ times its inactive-row mass. Row optimality requires every pure-column payoff to be nonnegative, so the y-average is nonnegative; hence u puts no mass on inactive rows. The positive y-weighted average of its active-column payoffs is then zero, so all active-column payoffs are zero. Collapsing the two duplicate coordinates gives a vector in the left nullspace of C, and lemma U.1 forces the collapsed vector to be x. This proves the row part of equation (U.4). For any column-optimal strategy v, use the equilibrium row strategy that splits x_p equally across the duplicates. Its expected payoff against an inactive column is P, so column optimality forces zero inactive-column mass. All active row payoffs must then be zero, and the right-nullspace argument gives $v = y$.

For either endpoint branch, deleting the unused duplicate row leaves a square active matrix equal to C up to permutation. Its two bordered matrices are nonsingular. Indeed, if $Cu - \alpha\mathbf{1} = 0$ and $\mathbf{1}^\top u = 0$, then left multiplication by $x^\top$ gives $\alpha = 0$; hence u is proportional to y, and the normalization gives $u = 0$. The transposed bordered matrix is handled identically. The implicit-function theorem therefore produces two regular analytic kernel branches.

On the plus branch, the only non-strict active validity inequality at B is the excluded-row slack

$$g(A) = v_+(A) - A_{r_-,:}y_+(A).$$

Because $v_+(A) = A_{r_+,:}y_+(A)$ and the duplicate rows coincide at B, differentiation cancels the derivative of y_+ and gives equation (U.5). Its gradient is nonzero because $y \in \mathrm{ri}(\Delta_s)$. Let $h(A)$ be the excluded-r_+ validity coordinate of the minus branch, oriented to be positive on the minus side. Direct differentiation at B gives

$$Dh(B)[H] = \sum_{j \in J} y_j(H_{r_-,j} - H_{r_+,j}) = -Dg(B)[H].$$

We next prove the coherent identity $h = -\chi g$. If $g(A) = 0$, the plus branch makes both duplicate rows best responses. Its column vector and value therefore solve the bordered column system for the minus kernel; nonsingularity makes that solution unique. The corresponding bordered row solution varies continuously from the strictly positive vector at B, so it is feasible after shrinking the neighborhood. Thus $h(A) = 0$. The converse follows by exchanging the two duplicates, and the two analytic zero sets agree locally. Since $Dg(B) \neq 0$, analytic coordinates (g, w) are valid near B. Choose the g-coordinate so that its coordinate vector e_g satisfies $Dg(B)[e_g] = 1$. The derivative identity above gives $\partial_g h(B) = -1$. The identity $h(0, w) = 0$ and the fundamental theorem of calculus give

$$h(g, w) = g \int_0^1 \partial_g h(tg, w)\, dt.$$

At B the integral equals -1; by continuity it remains negative and bounded away from zero after localization. Writing its negative as $\chi(A)$ proves $h(A) = -\chi(A)g(A)$ with analytic $0 < \underline{\chi} \leq \chi(A) \leq \overline{\chi} < \infty$. Every remaining support mass is strictly positive, and every padding slack equals P at B.

Finally, any regular kernel branch incident at B must use all core actions by lemma U.1, cannot use a padding action because its slack is strict, and must choose exactly one of the two duplicate rows because a regular Shapley–Snow kernel is square. Hence the two endpoint branches are the only incident regular germs. Pairwise distinct core rows and columns exclude an additional duplicate pivot. Continuity now yields a neighborhood in which the sign of g selects exactly one endpoint branch. The metric identity in equation (U.6) follows because the row vectors differ by moving mass x_p between two coordinates, so their ℓ_1 distance is $2x_p$, while the column vectors agree. □

Corollary U.3 (Quantitative theorem membership)**.** *Every certified split interface belongs, after shrinking its neighborhood, to some fixed-margin class $\mathfrak{G}(\Gamma_G)$. Any finite corpus of such interfaces is contained in one common class $\mathfrak{G}(\Gamma_{\mathrm{corp}})$ obtained by taking the minimum of its positive margins and the maximum of its finite upper bounds.*

Proof. The first statement is theorem T.2 applied to proposition U.2. A finite minimum of strictly positive margins is positive, and a finite maximum of finite bounds is finite, which proves the corpus statement. □

U.2 Exact game-invariant laws

For either orientation, vectorize the normal a, let the coordinate noise standard deviations be $\sigma_i > 0$, and define

$$s_G = \sum_i \sigma_i |a_i|, \qquad q^\star_{G,i} = \frac{\sigma_i |a_i|}{s_G}, \qquad v^{\mathrm{lf}}_{G,i} = \frac{\sigma_i \operatorname{sgn}(a_i)}{s_G} \mathbf{1}\{a_i \neq 0\}. \tag{U.8}$$

The usual extended-value convention is used when $q_i = 0$.

Theorem U.4 (Cross-game universality on the information-balanced pair)**.** *For every certified split interface and every $\kappa \geq 0$, consider the information-balanced local pair*

$$A_{n,\eta} = B + \eta \frac{\kappa s_G}{\sqrt{n}} v^{\mathrm{lf}}_G, \qquad \eta \in \{-1, +1\}. \tag{U.9}$$

The ideal approximate design $q^\star_G$ has normal variance s^2_G, and its Gaussian tangent score is exactly

$$T \sim \mathcal{N}(\eta\kappa, 1). \tag{U.10}$$

If $N^{(n)}_i$ are balanced integer counts with $N^{(n)}_i/n \to q^\star_{G,i}$ on the informative support, then the corresponding standardized coordinate-query score converges to equation (U.10) *in total variation. Moreover, $g(A_{n,\eta}) = \eta\kappa s_G/\sqrt{n} + O(n^{-1})$ and the two branch targets converge to z_η. Consequently, writing $\Delta_G = d_\Delta(z_+, z_-)$, the exact Gaussian-limit benchmarks are*

$$R^\star_{\mathrm{side}}(\kappa) = \Phi(-\kappa), \tag{U.11}$$

$$\frac{4R^\star_{\mathrm{point}}(\kappa)}{\Delta^2_G} = r_{\mathrm{pt}}(\kappa) = e^{-\kappa^2/2} \mathbb{E}_{Z\sim\mathcal{N}(0,1)}[\cosh(\kappa Z)^{-1}], \tag{U.12}$$

$$\frac{D^\star_\alpha(\kappa)}{\Delta_G} = q_\alpha(\kappa) = [\Phi(z_{1-\alpha} - 2\kappa) - \alpha]_+. \tag{U.13}$$

After normalization by κ and Δ_G, these laws contain no game dimension, support, orientation, noise profile, branch separation, or conditioning parameter.

Proof. Sharp Cauchy–Schwarz gives, for every design q,

$$\sum_i \frac{\sigma_i^2 a_i^2}{q_i} \geq \left(\sum_i \sigma_i |a_i| \right)^2 = s^2_G,$$

with equality at $q^\star_G$. Moreover, $\langle a, v^{\mathrm{lf}}_G \rangle = 1$. Under the ideal approximate design, the standardized linear normal score has mean $\eta\kappa$ and variance one, proving equation (U.10). Under balanced integer counts its mean is unchanged and its variance converges to one; the total-variation distance between equal-mean Gaussian laws therefore tends to zero. Taylor expansion of g and branch continuity give the asserted finite-game target convergence. Neyman–Pearson then gives equation (U.11). Conditional on $T = t$, the positive posterior weight is $(1 + e^{-2\kappa t})^{-1}$. The metric segment joining the two simplex targets is available componentwise, so posterior interpolation has conditional risk $\pi(t)(1 - \pi(t))\Delta^2_G$. Integrating under the equal mixture density gives equation (U.12), exactly as in theorem G.3. The three-report confidence calculation gives equation (U.13); see theorem Q.3. No generator-specific quantity remains after the displayed normalization. □

Proposition U.5 (Decision-rule limits at a split interface)**.** *On the experiment equation* (U.10)*:*

1. *the hard branch, and hence the local limit of an LP plug-in rule, has normalized squared point risk $4\Phi(-\kappa)$;*
2. *posterior interpolation attains equation* (U.12)*;*
3. *an entropy-regularized saddle rule with $\tau_n = s_G/(2\sqrt{n})$ has duplicate-action odds converging to $\exp(2T_n)$ and hence normalized risk*

$$4 \max_{\eta\in\{-1,+1\}} \mathbb{E}_{T\sim\mathcal{N}(\eta\kappa,1)} \left[\mathbf{1}\{\eta = +1\} - \frac{1}{1 + e^{-2T}} \right]^2. \tag{U.14}$$

This generally differs from the minimax curve because the Bayes-optimal odds are $\exp(2\kappa T)$, not $\exp(2T)$.

The entropy comparison is an algorithmic negative control, not part of the universality theorem or any headline claim, and not a claim of a new QRE approximation rate.

Proof. A hard rule returns the wrong endpoint with probability $\Phi(-\kappa)$; its squared loss is then Δ_G^2, proving the first claim. The second was proved in theorem U.4. For the third, the KKT equations of the entropy-regularized saddle problem give

$$\log \frac{x_{r_+}^{(\tau)}}{x_{r_-}^{(\tau)}} = \frac{(A_{r_+,:} - A_{r_-,:})y^{(\tau)}}{\tau}.$$

Regularity of the collapsed core implies $y^{(\tau_n)} \to y$ uniformly on the local pair. The numerator, evaluated on the Gaussian payoff estimate, equals $s_G T_n/\sqrt{n} + o_P(n^{-1/2})$. Substitution of $\tau_n = s_G/(2\sqrt{n})$ gives log odds $2T_n + o_P(1)$. The duplicate mass converges to x_p, so the normalized squared target loss converges to equation (U.14) by bounded convergence. Nguyen et al. (2026) establish the tight $\Theta(\beta^{-1}(\log m + \log k))$ finite-game QRE approximation order; here entropy is used only as a smooth selection control against which the minimax posterior rule can be compared. □

U.3 The finite certified corpus

The released script `run_universality_study.py` in the `reproducibility/` directory uses seed 20260925 and creates 80 games in each dimension $d \in \{3, 5, 10, 20\}$, for a total of 320. It cycles the active support through $2, \ldots, \min\{d-1, 6\}$, alternates row and column splits, samples interior weights with a deterministic positive floor, uses equation (U.2) with controlled nonzero singular values, rejects cores with nearly coincident rows or columns, applies random permutations, and draws heterogeneous standard deviations log-uniformly from $[0.5, 2]$. The exact certification tests equation (U.1), target separation, normal normalization, padding slacks, and the full-support pilot recovery map.

For every game and $\kappa \in \{0, 0.5, 1, 1.5, 2, 2.5, 3\}$, the script draws 2048 independent scores under each label from the integer-rounded coordinate-query law. A separate non-oracle study uses 12 games per dimension, four phase values, and 96 balanced-label replications per cell. It runs a specialized implementation of TARGET that estimates the active block, split orientation, duplicate pair, equilibrium weights, normal, side targets, and target-optimal allocation from an independent full-support pilot. This specialization is a finite-atlas implementation of definition O.2; it does not receive any instance-specific geometry.

Proposition U.6 (Simultaneous Monte Carlo certification)**.** *Suppose each of K reported cells uses M independent replications under each label. With probability at least $1-\delta$, every label-specific side-error, coverage, and normalized-diameter mean differs from its expectation by at most*

$$\sqrt{\frac{\log(16K/\delta)}{2M}}, \tag{U.15}$$

and every label-specific normalized posterior point-risk mean differs by at most

$$4\sqrt{\frac{\log(16K/\delta)}{2M}}. \tag{U.16}$$

The same bounds hold after taking the maximum or minimum over the two labels. Thus the simulations provide a finite, auditable implementation check; the validity of equations (U.11) *to* (U.13) *does not depend on Monte Carlo evidence.*

Proof. Side errors, coverage indicators, and normalized diameters are Bernoulli variables; there are six label-specific Bernoulli means per cell. There are two label-specific point-loss means per cell, and each normalized point loss lies in $[0, 4]$. Two-sided Hoeffding bounds followed by a union bound over all $8K$ means give the displayed constants. Maxima and minima are one-Lipschitz in the sup norm of the two label-specific means. □

The run produced 2240 oracle game–phase rows and 192 adaptive summary rows. Across individual oracle game–phase cells, the side-risk root-mean-square error was 7.05×10^{-3} and the normalized

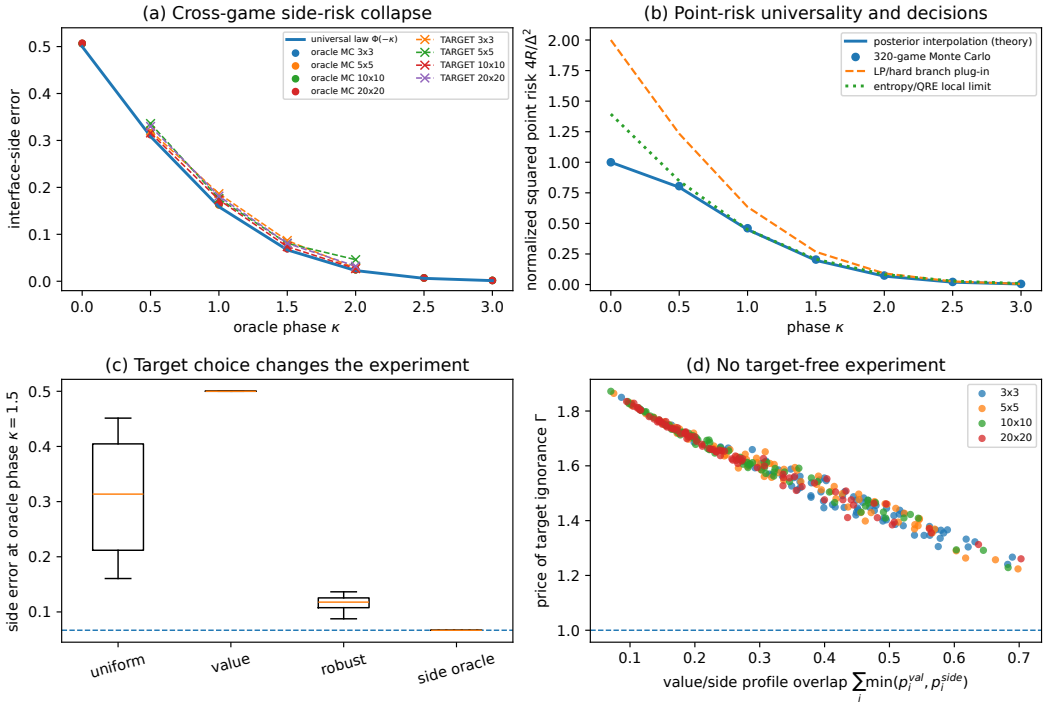

Figure 9: **Certified cross-game universality.** (a) Oracle Monte Carlo and non-oracle TARGET side errors collapse onto $\Phi(-\kappa)$ across dimensions and heterogeneous interfaces. (b) Normalized point risks collapse onto the posterior-interpolation law; hard LP selection and the entropy/QRE local limit are controls. (c) Uniform, value-optimal, robust post-hoc, and side-optimal designs are compared on the same side target. (d) The exact target-ignorance price increases as the value and side information profiles separate. Markers are finite-sample diagnostics of proved formulas, not empirical premises of the theory.

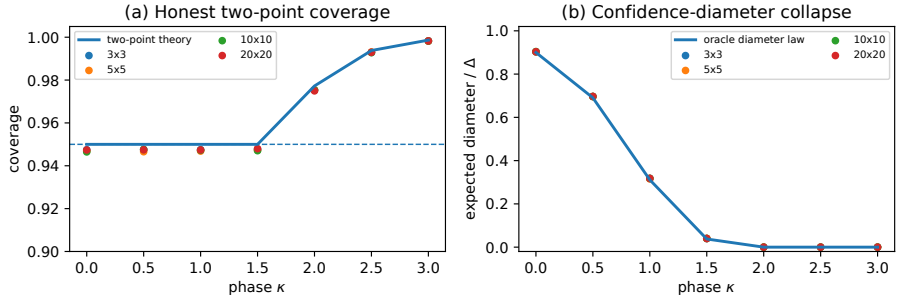

Figure 10: **Confidence universality.** Coverage and expected diameter for the exact two-point three-report region collapse onto the oracle laws. Coverage is exactly $1-\alpha$ before the one-report threshold and becomes conservative afterward.

point-risk root-mean-square error was 1.11×10^{-2}. The largest integer-rounding deviation of the score-variance ratio from one was 5.94×10^{-9}. Every adaptive cell recovered the correct chart in the fixed seeded run; the median variance ratio to the oracle design was 1.0154, its 95th percentile across summary cells was 1.0209, and the adaptive side-risk root-mean-square error was 3.54×10^{-2}. These numbers are descriptive outputs of the released computation and are not promoted to probabilistic theorems about future random corpora.

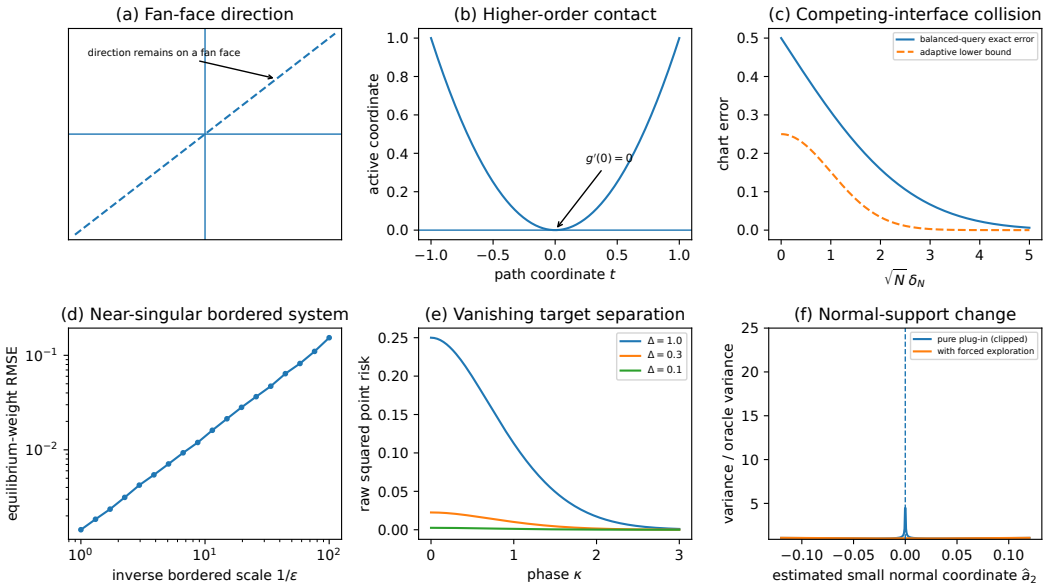

Figure 11: **Negative controls.** The theorem excludes directions on fan faces, higher-order contact, unresolved competing interfaces, nearly singular bordered systems, vanishing target separation, and unprotected normal-support changes. The last panel shows why TARGET mixes its estimated oracle design with forced exploration.

For each generated game, the value profile $p_{ij}^{\mathrm{val}} \propto \sigma_{ij} x_i^\star y_j^\star$ differs from the interface profile $p_{ij}^{\mathrm{side}} \propto \sigma_{ij}|a_{ij}|$. Solving the exact post-hoc design dual of theorem S.2 gave

$$1.2238 \le \Gamma \le 1.8720, \qquad \mathrm{median}(\Gamma) = 1.5990, \tag{U.17}$$

for the fixed corpus. The strict inequality $\Gamma > 1$ is also structural: by corollary S.3, equality would require identical profiles, whereas a split interface assigns positive side information to both duplicate actions and a one-sided value derivative assigns mass to only the active endpoint.

U.4 Negative controls and exact scope

The six panels in figure 11 correspond to six distinct failures, not to one generic "hard game" label. A fan-face direction does not select a unique relative-interior cone. Higher-order contact violates first-order transversality. The collision family of theorem R.4 has exact attainable threshold $N\delta_N^2 \to \infty$. A vanishing bordered singular value amplifies equilibrium-weight error. Raw point risk disappears with Δ_G^2, so a fixed positive target-separation margin is necessary for an order-one point-risk statement. Finally, a pure plug-in allocation can assign zero mass to a coordinate whose true normal component is small but nonzero; forced exploration removes this pole, exactly as required by equation (P.57).

No conclusion in this appendix extends the adaptive theorem beyond a fixed quantitative class. The corpus deliberately samples only certified simple interfaces. It does not establish prevalence among arbitrary games, does not cover higher-order or multi-interface collisions, and does not turn a finite-sample plot into a uniform theorem. Those exclusions are the content, not a defect, of theorem T.4.