Counterfactual Sufficiency

State Representation and Response Geometry
Online Appendix

Chae-Yeon Xon
Yonsei University

September 2026

Abstract

This appendix collects supporting geometry, measurement and inference results, numerical-accuracy conditions, diagnostic designs, and the remaining proofs. The proofs of the local spectral-risk expansion, the fiber certificate sandwich, and the feasible-coordinate oracle bound follow their statements in the main paper. The state-design specification and response targets are those of the main paper. The quantitative details describe the supplied model exercises.

JEL Codes: C14, C52, C61, C63, D81

Keywords: counterfactual sufficiency, state representation, response geometry, dimension reduction, active subspaces, structural models

Yonsei University. Email: eunixon@yonsei.ac.kr. This draft is circulated for discussion; comments and suggestions are welcome.

Appendix A. Geometric and function-space extensions

These results record the coordinate and metric conditions used by the main analysis. The finite-dimensional exposition in the paper uses an orthonormal vertical frame; the constructions below give alternative coordinates and a function-space formulation.

REMARK A1 (Observability-Gramian special case). *Consider a linearized state-output system $\delta s_{h+1} = F_h \delta s_h$ and $\delta y_h = C_h \delta s_h$, and let $\Phi_h = F_{h-1} \cdots F_0$ with $\Phi_0 = I$. If the declared targets are the outputs $\{\delta y_h\}_{h \in \mathcal{H}}$ and B is an orthonormal frame for the supplied-state vertical space, then*

$$H_{\mathcal{T}} = B' \left\{ \sum_{h \in \mathcal{H}} \omega_h \Phi_h' C_h' W_h C_h \Phi_h \right\} B. \tag{A1}$$

The matrix in braces is a weighted finite-horizon observability or output-sensitivity Gramian. Equation (A1) *therefore places the linear special case inside observability-based model reduction. Nonlinear empirical-Gramian methods construct related matrices from simulated or experimental state perturbations (Lall, Marsden, and Glavaski 2002; Hahn and Edgar 2002; Hahn, Edgar, and Marquardt 2003).*

A.1. Metric covariance and distribution states

The operator in (13) is written in an orthonormal frame of the vertical space. Let $Z(s) : \mathbb{R}^k \to \mathcal{V}_s$ be an arbitrary local frame, let G_s be the coordinate matrix of the state metric restricted to $\mathcal{V}_s$, and define the target quadratic form

$$Q_{\mathcal{T},s} = \sum_{t \in \mathcal{T}} \omega_t \{D_s R_t(s) Z(s)\}' W_t \{D_s R_t(s) Z(s)\}. \tag{A2}$$

The metric endomorphism $\mathcal{H}_{\mathcal{T},s}$ is characterized by

$$g_s(\mathcal{H}_{\mathcal{T},s} v, w) = \sum_{t \in \mathcal{T}} \omega_t \langle D_s R_t(s) v, D_s R_t(s) w \rangle_{W_t}, \qquad v, w \in \mathcal{V}_s. \tag{A3}$$

Its coordinate matrix is $G_s^{-1} Q_{\mathcal{T},s}$ and its spectral directions solve the generalized eigenproblem

$$Q_{\mathcal{T},s} v = \lambda G_s v. \tag{A4}$$

PROPOSITION A1 (Coordinate covariance). *Let $\tilde{z}$ be another vertical coordinate with nonsingular Jacobian C such that $z = C\tilde{z}$ at s. Then $\widetilde{G}_s = C' G_s C$ and $\widetilde{Q}_{\mathcal{T},s} = C' Q_{\mathcal{T},s} C$. The generalized eigenvalues in* (A4) *are unchanged, and the corresponding tangent subspaces are carried by C. A target-coordinate transformation $y = D\tilde{y}$ with nonsingular D preserves $Q_{\mathcal{T},s}$ when the target metric transforms as $\widetilde{W} = D' W D$.*

The state metric therefore enters the definition of the loss geometry. Units, normalization, and reparameterization are represented through the metric and its pullback. Grey (2026) develops an intrinsic Riemannian formulation of active subspaces with parallel transport across tangent spaces. For a distribution state μ in the interior of the probability simplex, the tangent space is $\{h : \mathbf{1}'h = 0\}$. The Fisher–Rao metric is

$$g_\mu(h, k) = \sum_{i=1}^{n} \frac{h_i k_i}{\mu_i}. \tag{A5}$$

Chentsov's characterization identifies this metric up to scale among Markov-invariant Riemannian metrics on finite probability simplices (Fujiwara 2024). Section 6.2 of the main paper uses (A5) for a distributional state.

A.2. Function-space target operator

Let $\mathbb{V}_s$ be a separable Hilbert vertical space with inner product $\langle \cdot, \cdot \rangle_s$. A target view t has a separable Hilbert output space $\mathbb{Y}_t$, a bounded positive operator W_t on $\mathbb{Y}_t$, and a bounded derivative $A_t : \mathbb{V}_s \to \mathbb{Y}_t$. Let ν be a finite measure on the target family and let ω_t be measurable and positive. Assume

$$\int_{\mathfrak{T}} \omega_t \| W_t^{1/2} A_t \|_{\mathrm{HS}}^2 \, d\nu(t) < \infty. \tag{A6}$$

Define the target operator

$$\mathfrak{H}_{\mathfrak{T},s} = \int_{\mathfrak{T}} \omega_t A_t^* W_t A_t \, d\nu(t), \tag{A7}$$

where the integral is a Bochner integral in the trace-class operators on $\mathbb{V}_s$.

PROPOSITION A2 (Hilbert target operator and sieve limit). *Under* (A6), *$\mathfrak{H}_{\mathfrak{T},s}$ is positive, self-adjoint, trace class, and compact, with*

$$\operatorname{tr} \mathfrak{H}_{\mathfrak{T},s} = \int_{\mathfrak{T}} \omega_t \| W_t^{1/2} A_t \|_{\mathrm{HS}}^2 \, d\nu(t). \tag{A8}$$

If W_t is positive definite on the closure of $\operatorname{im} A_t$ *for ν-almost every t, then*

$$\ker \mathfrak{H}_{\mathfrak{T},s} = \{ v \in \mathbb{V}_s : A_t v = 0 \textit{ for } \nu\textit{-almost every } t \}. \tag{A9}$$

Let Π_m be finite-rank orthogonal projections with $\Pi_m v \to v$ for every $v \in \mathbb{V}_s$, and set $\mathfrak{H}_m = \Pi_m \mathfrak{H}_{\mathfrak{T},s} \Pi_m$. Then

$$\| \mathfrak{H}_m - \mathfrak{H}_{\mathfrak{T},s} \|_1 \to 0, \tag{A10}$$

where $\|\cdot\|_1$ *is the trace norm. If the nonzero eigenvalues are ordered as* $\lambda_1 \geq \lambda_2 \geq \cdots$*, padded by zeros, then*

$$\sum_{j\geq 1} |\lambda_j(\mathfrak{H}_m) - \lambda_j(\mathfrak{H}_{\mathcal{T},s})| \to 0. \tag{A11}$$

A spectral gap at d yields convergence of the leading rank-d spectral projector in operator norm. The rank-d quadratic target loss equals $\sum_{j>d} \lambda_j$.

The direct-gradient special case belongs to infinite-dimensional active-subspace analysis (Kundu and Wycoff 2026). The target operator here acts on the vertical space induced by the supplied representation and stacks the declared structural counterfactual derivatives before the sieve approximation.

COROLLARY A1 (Distributional score representation). *Let* μ *be a positive density and let* $\mathbb{H}_\mu = L^2_0(\mu)$ *be the centered score Hilbert space with inner product* $\langle f, g\rangle_\mu = \mathbb{E}_\mu[fg]$*. Suppose the supplied representation records moments* $\mathbb{E}_\mu[\phi_j]$*,* $j = 1, \ldots, m$*. The vertical score space is*

$$\mathbb{V}_\mu = \{s \in \mathbb{H}_\mu : \mathbb{E}_\mu[\phi_j s] = 0,\ j = 1, \ldots, m\}. \tag{A12}$$

If a scalar target derivative has Riesz representation $DR_t(\mu)[\mu s] = \mathbb{E}_\mu[\psi_t s]$*, let* $\psi_t^{\mathcal{V}}$ *be the orthogonal projection of* ψ_t *onto* $\mathbb{V}_\mu$*. Then*

$$\mathfrak{H}_{\mathcal{T},\mu} s = \int_{\mathcal{T}} \omega_t \psi_t^{\mathcal{V}}\, \mathbb{E}_\mu[\psi_t^{\mathcal{V}} s]\, d\nu(t). \tag{A13}$$

For a finite family of M scalarized target views, $\operatorname{rank} \mathfrak{H}_{\mathcal{T},\mu} \leq M$*. Its nonzero eigenvalues equal those of the* $M \times M$ *view Gram matrix with entries* $\mathbb{E}_\mu[\psi_t^{\mathcal{V}} \psi_{t'}^{\mathcal{V}}]/M$ *under equal weights.*

The score map and Fisher pairing are standard constructions in nonparametric information geometry (Pistone and Sempi 1995; Cheng and Tong 2026). Corollary A1 gives a distributional implementation in which the operator dimension can be indexed by the target family while the distributional state remains function valued.

A.3. Decision-induced target metric

Let a planner choose $a \in \mathbb{R}^m$ after forming a target vector $r \in \mathbb{R}^J$. The twice differentiable loss is $L(a, r)$. At (a_0, r_0), suppose a_0 is an interior minimizer and $H_{aa} = D_{aa}L(a_0, r_0) \succ 0$.

PROPOSITION A3 (Decision metric). *Let* $a^\star(r)$ *denote the local optimizer. Then*

$$Da^\star(r_0) = -H_{aa}^{-1} D_{ar} L(a_0, r_0). \tag{A14}$$

For a target perturbation Δr,

$$L(a^\star(r_0 + \Delta r), r_0) - L(a_0, r_0) = \frac{1}{2}\Delta r' W^{\mathrm{dec}} \Delta r + o(\|\Delta r\|^2), \tag{A15}$$

where

$$W^{\mathrm{dec}} = D_{ra} L(a_0, r_0) H_{aa}^{-1} D_{ar} L(a_0, r_0) \geq 0. \tag{A16}$$

The matrix W^{dec} can be used as the target metric in (13). It has rank at most m, so the decision-relevant quotient can merge target changes that leave the local optimal action unchanged. Decision-focused prediction uses downstream optimization loss to organize prediction tasks (Elmachtoub and Grigas 2022; Lee, Jin, and Lee 2026); equation (A16) supplies the local metric used here. For a finite displacement v in the supplied-fiber chart used in (21), a twice differentiable target response satisfies

$$R_t(\varphi_s(v)) = R_t(s) + A_t(s)v + r_t(v), \qquad \|r_t(v)\| \leq K_t \|v\|^2/2. \tag{A17}$$

The first-order frontier therefore receives a remainder controlled by the declared state radius and target curvature. Policy magnitude can be included as a second expansion coordinate.

A.4. Continuous-encoder approximation

Let $\mathcal{A}_s^{\mathcal{T}} = DF_s(0)$ and write its singular values as $\sigma_1 \geq \cdots \geq \sigma_k$. For $d < k$, define

$$e_d(\rho) = \inf_{\substack{b : B_\rho^k \to \mathbb{R}^d \text{ continuous} \\ g}} \sup_{\|v\| \leq \rho} \|F_s(v) - g(b(v))\|. \tag{A18}$$

PROPOSITION A4 (Continuous-encoder frontier). *Under* (22)*,*

$$\max\{0, \rho\sigma_{d+1}(\mathcal{A}_s^{\mathcal{T}}) - K_s\rho^2/2\} \leq e_d(\rho) \leq \rho\sigma_{d+1}(\mathcal{A}_s^{\mathcal{T}}) + K_s\rho^2/2. \tag{A19}$$

For the affine target $F_s(v) = F_s(0) + \mathcal{A}_s^{\mathcal{T}} v$,

$$e_d(\rho) = \rho\sigma_{d+1}(\mathcal{A}_s^{\mathcal{T}}). \tag{A20}$$

The lower bound uses antipodal coincidence; the upper bound uses the leading right-singular coordinates. The target map, frame, and remainder are those of Section 3.2.

A.5. Coordinate realization

The projector $P_d(s)$ defines retained directions inside the vertical space. Let

$$\mathcal{K}_d(s) = B(s) \ker P_d(s) \subseteq \ker Dq(s) \tag{A21}$$

be the discarded distribution.

PROPOSITION A5 (Local coordinate realization). *Suppose q and the constant-rank distribution $\mathcal{K}_d$ are smooth on a neighborhood of s_0. After restricting the neighborhood, the following statements are equivalent:*

1. *$\mathcal{K}_d$ is involutive;*
2. *there is a smooth map $b : U \to \mathbb{R}^d$ such that (q, b) has constant rank and*

$$\ker D(q, b)(s) = \mathcal{K}_d(s), \qquad s \in U. \tag{A22}$$

The eigengap makes the spectral projector smooth when the target field has corresponding regularity. The eigengap leaves involutivity as a separate condition. A collection of local charts can cover several coordinate neighborhoods. A single global coordinate additionally requires a regular leaf space with a suitable global embedding.

A.6. Integrable-coordinate frontier

Let μ be a probability measure on $\mathcal{S}$. Let $b : \mathcal{S} \to \mathbb{R}^d$ be a smooth map such that $z = (q, b)$ has constant rank and $Db(s)|_{\mathcal{V}_s}$ has rank d. Assume $\mathbb{E}\,\mathrm{tr}\, H_{\mathcal{T}}(S) < \infty$. In the vertical frame, let $P_b(s)$ be the rank-d projector onto the orthogonal complement of $\ker\{Db(s)|_{\mathcal{V}_s}\}$. Define the orientation loss

$$\Delta_b(s) = \mathrm{tr}\left[\{P_d(s) - P_b(s)\} H_{\mathcal{T}}(s)\right]. \tag{A23}$$

PROPOSITION A6 (Integrable-coordinate loss). *Every admissible b satisfies*

$$\mathrm{tr}\{(I - P_b(s)) H_{\mathcal{T}}(s)\} = \sum_{j>d} \lambda_j(s) + \Delta_b(s), \qquad \Delta_b(s) \geq 0. \tag{A24}$$

When $\gamma_d(s) > 0$, $\Delta_b(s) = 0$ holds exactly when $P_b(s) = P_d(s)$. Moreover,

$$0 \leq \Delta_b(s) \leq \|H_{\mathcal{T}}(s)\|_{\mathrm{op}} \|P_b(s) - P_d(s)\|_{*}. \tag{A25}$$

Hence

$$\text{(A26)}\qquad \inf_b \mathbb{E}\,\mathrm{tr}\{(I - P_b(S))H_{\mathcal{T}}(S)\} \geq \mathbb{E} \sum_{j>d} \lambda_j(S).$$

If the infimum is attained by an admissible $b^\star$ and the eigengap is positive almost everywhere, equality in (A26) *holds if and only if $P_{b^\star} = P_d$ almost everywhere.*

The spectral tail is a pointwise lower envelope. A smooth feature map must also have kernel fields that form level-set foliations. Once an admissible feature is supplied, equation (A24) measures the loss created by this additional constraint.

Appendix B. Measurement and observed design information

The state-coordinate question can be applied to measurement systems and observed response views. The derivative-level results below state the information required by the declared target.

B.1. State measurement completeness

Let $M_{\mathcal{D}} : \mathcal{S} \to \mathbb{R}^r$ collect the state fields supplied by a measurement system, including q when the supplied representation is observed. Let $\boldsymbol{R}_{\mathcal{T}}$ stack the declared counterfactual targets. Work in an orthonormal frame of the supplied-state vertical space and let $P_{\mathcal{D}}(s)$ project onto $\ker DM_{\mathcal{D}}(s)$. Define the residual target operator

$$\text{(A27)}\qquad H_{\mathcal{T}|\mathcal{D}}(s) = P_{\mathcal{D}}(s) H_{\mathcal{T}}(s) P_{\mathcal{D}}(s).$$

PROPOSITION A7 (Counterfactual completeness and measurement repair). *At a state s, define*

$$\text{(A28)}\qquad \delta_{\mathcal{T}|\mathcal{D}}(s) = \mathrm{rank}\left(D\boldsymbol{R}_{\mathcal{T}}(s)|_{\ker DM_{\mathcal{D}}(s)}\right) = \mathrm{rank}\, D(M_{\mathcal{D}}, \boldsymbol{R}_{\mathcal{T}})(s) - \mathrm{rank}\, DM_{\mathcal{D}}(s).$$

First-order target recovery from $M_{\mathcal{D}}$ holds exactly when $\delta_{\mathcal{T}|\mathcal{D}}(s) = 0$. Every collection of k added C^1 scalar measurements that achieves first-order recovery satisfies $k \geq \delta_{\mathcal{T}|\mathcal{D}}(s)$, and a collection with equality exists at the derivative level. Let $A_{\mathcal{T},W,\mathcal{D}}(s)$ stack $\sqrt{\omega_t} W_t^{1/2} DR_t(s) P_{\mathcal{D}}(s)$ and define the loss deficiency $\delta_{\mathcal{T},W|\mathcal{D}}(s) = \mathrm{rank}\, A_{\mathcal{T},W,\mathcal{D}}(s)$. Then

$$\text{(A29)}\qquad \mathrm{rank}\, H_{\mathcal{T}|\mathcal{D}}(s) = \delta_{\mathcal{T},W|\mathcal{D}}(s), \qquad H_{\mathcal{T}|\mathcal{D}}(s) = 0 \iff \delta_{\mathcal{T},W|\mathcal{D}}(s) = 0.$$

Every collection of k added scalar measurements that eliminates first-order declared loss satisfies $k \geq \delta_{\mathcal{T},W|\mathcal{D}}(s)$, and equality is attainable at the derivative level. If the target metric is positive definite on the attainable residual target image, then $\delta_{\mathcal{T},W|\mathcal{D}}(s) = \delta_{\mathcal{T}|\mathcal{D}}(s)$. Let $\mu_1 \geq \mu_2 \geq \cdots$ be *the eigenvalues of $H_{\mathcal{T}|\mathcal{D}}(s)$. Among d additional orthonormal local scalar measurements with derivatives in the unresolved data fiber, the minimum remaining isotropic quadratic target loss is*

$$\sum_{j>d} \mu_j, \tag{A30}$$

attained by measuring the first d eigen-directions of $H_{\mathcal{T}|\mathcal{D}}(s)$. Suppose $M_{\mathcal{D}}$ and $(M_{\mathcal{D}}, \boldsymbol{R}_{\mathcal{T}})$ have constant ranks on a neighborhood with connected local fibers of $M_{\mathcal{D}}$. If $\delta_{\mathcal{T}|\mathcal{D}} = 0$ throughout that neighborhood, then there is a local map g such that $\boldsymbol{R}_{\mathcal{T}} = g \circ M_{\mathcal{D}}$.

The raw deficiency counts repair for exact target recovery. The residual operator measures declared-loss variation inside the state distinctions left unresolved by the measurement system; its rank counts loss repair and its spectrum orders added measurements under a dimension budget.

B.2. Observed design views

Let $\theta_0 \in \Theta \subseteq \mathbb{R}^p$. An observed design view $e \in \mathcal{D}$ supplies a moment vector $m_e(\theta) \in \mathbb{R}^{r_e}$ with

$$J_e = D_\theta m_e(\theta_0). \tag{A31}$$

Stack active views to form $J_{\mathcal{D}}$. With $n_e > 0$ replications and moment covariance $\Sigma_e \succ 0$, define

$$I_{\mathcal{D}} = \sum_{e \in \mathcal{D}} n_e J_e' \Sigma_e^{-1} J_e. \tag{A32}$$

Let $\tau(\theta) \in \mathbb{R}^m$ be a structural target and write $T_\theta = D_\theta \tau(\theta_0)$.

PROPOSITION A8 (Design information and target identification). *The information kernel is*

$$\ker I_{\mathcal{D}} = \bigcap_{e \in \mathcal{D}: n_e > 0} \ker J_e = \ker J_{\mathcal{D}}. \tag{A33}$$

The target is locally identified to first order from the design when

$$\ker J_{\mathcal{D}} \subseteq \ker T_\theta. \tag{A34}$$

Condition (A34) *is equivalent to a linear factorization $T_\theta = L J_{\mathcal{D}}$ for a map L on the image of $J_{\mathcal{D}}$. The residual target-identification dimension is*

$$r_{\mathcal{D}\to\tau} = \operatorname{rank}\left(T_\theta|_{\ker J_{\mathcal{D}}}\right) = \operatorname{rank}\begin{pmatrix} J_{\mathcal{D}} \\ T_\theta \end{pmatrix} - \operatorname{rank} J_{\mathcal{D}}. \tag{A35}$$

Every collection of unrestricted scalar added moments that yields (A34) *has cardinality at least* $r_{\mathcal{D}\to\tau}$*, and a collection of that cardinality exists at the derivative level.*

A positive change in n_e scales one summand of (A32) and preserves the kernel in (A33). Adding a design view $e^\star$ changes the kernel to $\ker I_{\mathcal{D}} \cap \ker J_{e^\star}$. Replication therefore changes information eigenvalues while additional views can change information rank.

COROLLARY A2 (Information accounting). *Let S be a block of additional design moments. Then*

$$r_{\mathcal{D}\to\tau} - r_{\mathcal{D}\cup S\to\tau} = \operatorname{rank} J_{\mathcal{D}\cup S} + \operatorname{rank}\begin{pmatrix} J_{\mathcal{D}} \\ T_\theta \end{pmatrix} - \operatorname{rank} J_{\mathcal{D}} - \operatorname{rank}\begin{pmatrix} J_{\mathcal{D}\cup S} \\ T_\theta \end{pmatrix} \geq 0. \tag{A36}$$

B.3. Target-directed view allocation

Let $M \succ 0$ be the current precision matrix on a local state or parameter coordinate and let $\Sigma = M^{-1}$. Let $H_{\mathcal{T}} \succeq 0$ measure the declared target loss. A candidate view has derivative $J \in \mathbb{R}^{r\times k}$ and noise covariance $R \succ 0$, giving the precision update $M^+ = M + J'R^{-1}J$. Define the target posterior-variance criterion

$$\Psi(M) = \operatorname{tr}(H_{\mathcal{T}} M^{-1}). \tag{A37}$$

PROPOSITION A9 (Target loss from an added view). *The reduction in* (A37) *from the candidate view is*

$$\Psi(M) - \Psi(M^+) = \operatorname{tr}\left[H_{\mathcal{T}}\Sigma J'(R + J\Sigma J')^{-1}J\Sigma\right]. \tag{A38}$$

For a scalar view with precision contribution $\nu aa'$,

$$\Psi(M) - \Psi(M + \nu aa') = \frac{\nu a'\Sigma H_{\mathcal{T}}\Sigma a}{1 + \nu a'\Sigma a}. \tag{A39}$$

A sequential target-directed design selects the candidate view with the largest value in (A38). *The criterion is a local goal-oriented A-optimal design calculation (Attia, Alexanderian, and Saibaba 2018).*

Appendix C. Statistical extensions

C.1. Scattered response-field reconstruction

The local operator can be computed from a solved model or estimated from response data. In the latter case, the first stage must recover the derivative defined by the declared target family. The following reconstruction result conditions on a uniform first-stage rate. Cross-fitted quadratic products remove their noise contribution under the centering conditions in Lemma A1.

Work on a compact state region $\mathcal{S}_0$ contained in one chart. The vertical dimension k is constant. A finite target family $\mathcal{T}$ is used in the main results. Let $s_{1n}, \dots, s_{N_n n}$ be anchor states. For each anchor and target view, obtain two cross-fitted response-derivative estimates $\widehat{A}_{it}^{(1)}$ and $\widehat{A}_{it}^{(2)}$. State-dependent response estimators can supply these local derivatives after their state and shock-size conditions are matched to the declared target; Gonçalves et al. (2024) give such conditions for local projections.

ASSUMPTION A1 (Response field and cover). *The region $\mathcal{S}_0$ is compact. The map q is C^1 with constant rank on a neighborhood of $\mathcal{S}_0$. The operators $A_t(s)$ are bounded uniformly over s and t and satisfy $\|A_t(s) - A_t(s')\|_{\mathrm{op}} \le L\|s - s'\|^{\beta}$ for some $\beta \in (0,1]$. The weights and target metrics are bounded. Partition weights satisfy $\kappa_{ni}(s) \ge 0$, $\sum_i \kappa_{ni}(s) = 1$, and $\kappa_{ni}(s) > 0$ only when $\|s - s_{in}\| \le h_n$, where $h_n \to 0$. The local estimates satisfy*

$$\max_{i,t,b\in\{1,2\}} \|\widehat{A}_{it}^{(b)} - A_t(s_{in})\|_{\mathrm{op}} = O_p(\zeta_n), \qquad \zeta_n \to 0. \tag{A40}$$

Define the cross-fitted local patch

$$\widetilde{H}_i = \sum_{t\in\mathcal{T}} \frac{\omega_t}{2} \left[\widehat{A}_{it}^{(1)\prime} W_t \widehat{A}_{it}^{(2)} + \widehat{A}_{it}^{(2)\prime} W_t \widehat{A}_{it}^{(1)}\right] \tag{A41}$$

and the reconstructed field

$$\widehat{H}_n(s) = \sum_{i=1}^{N_n} \kappa_{ni}(s) \widetilde{H}_i. \tag{A42}$$

LEMMA A1 (Quadratic patch centering). *Suppose $\widehat{A}^{(b)} = A + E^{(b)}$ for $b \in \{1,2\}$, with $\mathbb{E}[E^{(b)} \mid s] = 0$ and $\mathbb{E}[E^{(1)\prime} W E^{(2)} \mid s] = 0$. Then*

$$\mathbb{E}\left[\frac{1}{2}\{\widehat{A}^{(1)\prime} W \widehat{A}^{(2)} + \widehat{A}^{(2)\prime} W \widehat{A}^{(1)}\} \,\Big|\, s\right] = A' W A. \tag{A43}$$

For a same-sample estimate $\widehat{A} = A + E$ with $\mathbb{E}[E \mid s] = 0$,

$$\mathbb{E}[\widehat{A}' W \widehat{A} \mid s] = A' W A + \mathbb{E}[E' W E \mid s]. \tag{A44}$$

The construction can use repeated folds and average the cross products.

PROPOSITION A10 (Scattered reconstruction). *Under Assumption A1,*

$$\sup_{s\in\mathcal{S}_0} \|\widehat{H}_n(s) - H_{\mathcal{T}}(s)\|_{\mathrm{op}} = O_p(r_n), \qquad r_n = \zeta_n + \zeta_n^2 + h_n^{\beta}. \tag{A45}$$

A numerical integration of a continuous target family adds its uniform quadrature error to r_n.

The main reconstruction result is extended to aligned frames and eigenspace loss. The subsequent criterion incorporates independently recovered response information, and the spectral-tail result states the additional first-stage and sieve assumptions used for inference.

C.2. Reportable coordinate-selection bounds

COROLLARY A3 (Reported oracle bound). *Suppose deterministic sequences $\bar\rho_n, s_n \ge 0$ and a nonnegative estimator $\widehat{\mathfrak{D}}_n$ satisfy*

$$\text{(A46)}\qquad \Pr(\rho_n \le \bar\rho_n) \ge 1 - \beta_{1n}, \qquad \Pr\left(\sup_{\eta\in\Xi_d} |\widehat{\mathfrak{D}}_n(b_\eta) - \mathfrak{D}(b_\eta)| \le s_n\right) \ge 1 - \beta_{2n}.$$

Then, with probability at least $1 - \alpha - \beta_{1n} - \beta_{2n}$,

$$\text{(A47)}\qquad \mathfrak{R}(q, b_{\hat\eta_d}) - \mathfrak{R}_d^\star \le \sqrt{M\left\{\sup_{\eta\in\Xi_d} \widehat{\mathfrak{D}}_n(b_\eta) + s_n\right\}} + 2\bar\rho_n + 2a_n(\alpha) + \eta_n.$$

Under the relative-slack condition in Theorem 5, with probability at least $1 - \alpha - \beta_{1n}$,

$$\text{(A48)}\qquad \mathfrak{R}(q, b_{\hat\eta_d}) \le (1 + \delta_d)\mathfrak{R}_d^\star + 2\bar\rho_n + 2a_n(\alpha) + \eta_n.$$

A uniform rate r_n for all first-stage objects entering the score, including the target derivative, component means, and conditional spectral gap, gives $\bar\rho_n = O(r_n)$ on an event whose probability tends to one when the score functional is uniformly Lipschitz in those objects and the gap stays bounded away from zero. Evaluation of $\widehat{\mathfrak{D}}_n$ additionally uses the conditional density score and target second derivatives. If LD_Ξ/M grows at most polynomially in n, then $r_n = o(1)$, $p_d \log n/n = o(1)$, and $\eta_n = o_p(1)$ make the estimation and optimization terms vanish. The additive generator term also requires $\overline{\mathfrak{D}}_d = o(1)$, or its reported counterpart $\sup_\eta \widehat{\mathfrak{D}}_n(b_\eta) + s_n = o_p(1)$. These conditions state the estimation and approximation burdens created by the generator diagnostic and complete the population sandwich with sampling and approximation requirements.

C.3. Frame-aligned reconstruction

Suppose each anchor uses a local vertical frame. Let $O_{i\to s}$ denote the orthogonal coordinate transport from the frame at s_i to the frame at s, and let $\widehat{O}_{i\to s}$ be an estimate. Define

$$\text{(A49)}\qquad \widehat{H}_n^A(s) = \sum_i \kappa_{ni}(s)\widehat{O}_{i\to s}\widetilde{H}_i\widehat{O}_{i\to s}'.$$

COROLLARY A4 (Aligned local frames). *Suppose Proposition A10's local patch rate holds in each anchor frame,* $\sup_{i,s:\kappa_{ni}(s)>0} \|\widehat{O}_{i\to s} - O_{i\to s}\|_{\mathrm{op}} = O_p(\eta_n)$*, and*

$$\sup_{i,s:\kappa_{ni}(s)>0} \|O_{i\to s} H_{\mathcal{J}}(s_i) O'_{i\to s} - H_{\mathcal{J}}(s)\|_{\mathrm{op}} = O(h_n^{\beta}). \tag{A50}$$

Then

$$\sup_{s\in\mathcal{S}_0} \|\widehat{H}_n^{A}(s) - H_{\mathcal{J}}(s)\|_{\mathrm{op}} = O_p(\zeta_n + \zeta_n^2 + h_n^{\beta} + \eta_n). \tag{A51}$$

Local PCA and frame synchronization provide one route to the transport matrices in point-cloud settings (Singer and Wu 2012). In an observed Euclidean state with a common coordinate system, $O_{i\to s} = I$. A finite-sample positive-semidefinite version replaces $\widehat{H}$ by its positive spectral part $\widehat{H}_+$. Since $H_{\mathcal{J}} \succeq 0$,

$$\|\widehat{H}_+ - H_{\mathcal{J}}\|_{\mathrm{op}} \le 2\|\widehat{H} - H_{\mathcal{J}}\|_{\mathrm{op}}, \tag{A52}$$

so the reconstruction rate is preserved.

C.4. Eigenspaces and loss

ASSUMPTION A2 (Eigenspace separation). *For a retained dimension d,*

$$\inf_{s\in\mathcal{S}_0} \{\lambda_d(s) - \lambda_{d+1}(s)\} \ge \gamma_d > 0. \tag{A53}$$

Let $\widehat{P}_d(s)$ be the leading rank-d projector of $\widehat{H}_n(s)$.

COROLLARY A5 (Synthetic-state recovery). *Under Proposition A10 and Assumption A2,*

$$\sup_{s,j} |\widehat{\lambda}_j(s) - \lambda_j(s)| = O_p(r_n) \tag{A54}$$

and

$$\sup_{s} \|\widehat{P}_d(s) - P_d(s)\|_{\mathrm{op}} = O_p(r_n/\gamma_d). \tag{A55}$$

On events satisfying $\sup_s \|\widehat{H}_n(s) - H_{\mathcal{J}}(s)\|_{\mathrm{op}} < \gamma_d/2$*, the target-loss regret satisfies*

$$0 \le \mathcal{L}_s(\widehat{P}_d) - \mathcal{L}_s(P_d) \le \frac{4d}{\gamma_d} \|\widehat{H}_n(s) - H_{\mathcal{J}}(s)\|_{\mathrm{op}}^2. \tag{A56}$$

Hence the uniform loss regret is $O_p(dr_n^2/\gamma_d)$.

If the target rank is constant on $\mathcal{S}_0$, $\inf_s \lambda_r(s) > 0$, and $\lambda_{r+1}(s) = 0$, a threshold $a_n \to 0$ with $r_n/a_n \to 0$ yields uniform rank recovery through $\widehat{r}(s) = \#\{j : \widehat{\lambda}_j(s) > a_n\}$.

C.5. Structural estimation with response information

Let $Q_B(\theta)$ denote a population criterion based on observed behavior, transitions, or equilibrium moments. The structural model generates a target field $H_{\mathcal{T},\theta}(s)$. The population criterion is

$$Q(\theta) = Q_B(\theta) + \lambda \int_{\mathcal{S}_0} \|H_{\mathcal{T},\theta}(s) - H_0(s)\|_F^2 \, d\nu(s), \tag{A57}$$

where H_0 is the population field recovered from response data. When target responses are generated entirely by the structural model, the field can be computed after behavioral estimation; equation (A57) applies when response information enters estimation.

C.6. Response-data field and model field

At a fixed state, let $Z_i \in \mathbb{R}^{k_m}$ be coordinates of an observed local perturbation in a sieve of the vertical space. For target view t, consider the local response equation

$$Y_{it} = \alpha_t + \beta_t' Z_i + \varepsilon_{it}, \qquad \mathbb{E}[\varepsilon_{it} \mid Z_i] = 0. \tag{A58}$$

The first stage estimates the response derivatives β_t from perturbation coordinates and target responses. Split-sample estimates $\widehat{\beta}_t^{(1)}$ and $\widehat{\beta}_t^{(2)}$ produce the cross-fitted field

$$\widehat{H}_m^D = \sum_t \frac{\omega_t}{2} \left\{ \widehat{\beta}_t^{(1)} \widehat{\beta}_t^{(2)\prime} + \widehat{\beta}_t^{(2)} \widehat{\beta}_t^{(1)\prime} \right\}. \tag{A59}$$

A structural model indexed by θ generates $H_m^M(\theta)$ on the same sieve and state metric. One operator-matching criterion is

$$Q_m^H(\theta) = \|H_m^M(\theta) - H_m^D\|_{\Omega_m}^2. \tag{A60}$$

The structural criterion in (A61) allows (A60) as the field component. The sieve result in Proposition A2 supplies the approximation link when the state is function valued. Let $\widetilde{H}_{\mathcal{T},\theta,n}$ be the numerically computed model field and let ν_n be the sample integration measure. The sample criterion is

$$\widehat{Q}_n(\theta) = \widehat{Q}_{B,n}(\theta) + \lambda \int_{\mathcal{S}_0} \|\widetilde{H}_{\mathcal{T},\theta,n}(s) - \widehat{H}_n(s)\|_F^2 \, d\nu_n(s). \tag{A61}$$

ASSUMPTION A3 (Structural criterion). *The parameter space Θ is compact. The map $(\theta, s) \mapsto H_{\mathcal{T},\theta}(s)$ is jointly continuous and uniformly bounded. The behavioral criterion converges uni-*

formly, $\sup_\theta |\widehat{Q}_{B,n}(\theta) - Q_B(\theta)| = o_p(1)$*. The structural solver satisfies* $\sup_{\theta,s} \|\widetilde{H}_{\mathcal{T},\theta,n}(s) - H_{\mathcal{T},\theta}(s)\|_F = o_p(1)$*. The measure* ν_n *obeys a uniform law of large numbers for the squared field-distance class. The population criterion has unique minimizer* θ_0.

PROPOSITION A11 (Structural consistency). *Under Proposition A10 and Assumption A3, every sequence* $\widehat{\theta}_n \in \operatorname{argmin}_{\theta\in\Theta} \widehat{Q}_n(\theta)$ *satisfies*

$$\widehat{\theta}_n \to_p \theta_0. \tag{A62}$$

C.7. Geometry identification

Write $h_\theta(s) = \operatorname{vech} H_{\mathcal{T},\theta}(s)$ in a fixed local frame and let $G_H(s) = D_\theta h_{\theta_0}(s)$. For a positive weighting field $\Omega(s)$ define

$$\mathcal{I}_H = \int G_H(s)'\Omega(s)G_H(s)\, d\nu(s). \tag{A63}$$

PROPOSITION A12 (First-order information and local identification). *Suppose* $\theta \mapsto h_\theta$ *is Fréchet differentiable at an interior* θ_0 *in the weighted* $L^2(\nu;\Omega)$ *norm, with derivative* $u \mapsto G_H u$*. If* $\mathcal{I}_H$ *is positive definite,* θ_0 *is locally identified by this field in that norm. If* $\operatorname{rank} \mathcal{I}_H = r < p_\theta$*, its range records the* r *first-order parameter directions preserved by the field, while its kernel leaves the field unchanged to first order. For a stack with differentiable behavioral or level-response moments, regular first-order identification means that the stacked derivative has trivial kernel. This property is equivalent to full column rank and is sufficient for local identification under the same norm-differentiability condition.*

The full-rank condition is sufficient for local identification and characterizes its regular first-order form (Chen et al. 2014). Nonlinear curvature can identify directions that have a zero first derivative. Operator matching targets parameter combinations that move the state geometry of the declared counterfactual family. Fixed-sieve minimum-distance specification tests follow from the residual between $\widehat{h}^D$ and $h^M(\widehat{\theta})$ after projecting out the model tangent. Response-matching asymptotics and specification testing supply the surrounding minimum-distance framework (Guerron-Quintana, Inoue, and Kilian 2017; Liu, Plagborg-Moller, and Tan 2026).

C.8. First-order distribution

Suppose the field estimator admits an integrated influence representation

$$\sqrt{n}\, \operatorname{vec}\{\widehat{H}_n(s) - H_0(s)\} = \frac{1}{\sqrt{n}} \sum_{i=1}^{n} \xi_i(s) + o_p(1) \tag{A64}$$

in a function space supporting integration. With $G_\theta(s) = D_\theta \operatorname{vec} H_{\mathcal{J},\theta_0}(s)$, the field term contributes $-2\lambda \int G_\theta(s)'\xi_i(s)\, d\nu(s)$ to the score influence function. Combining this term with the behavioral score gives the sandwich covariance.

C.9. Growing sieve and spectral-tail inference

Let $\mathfrak{H}$ be a positive trace-class target operator, $\tau = \operatorname{tr} \mathfrak{H}$, and $\mathfrak{H}_m = \Pi_m \mathfrak{H} \Pi_m$. For a retained dimension d, define

$$\ell_d(\tau, \mathfrak{H}) = \tau - \kappa_d(\mathfrak{H}), \qquad \kappa_d(\mathfrak{H}) = \sum_{j=1}^{d} \lambda_j(\mathfrak{H}). \tag{A65}$$

The scalar trace is carried separately from the Hilbert–Schmidt operator so that the stochastic limit can be stated in the Hilbert space $\mathbb{R} \times \mathcal{S}_2^{sa}$.

ASSUMPTION A4 (Joint sieve first stage). *For $m_n \to \infty$,*

$$\sqrt{n}\left(\widehat{\tau}_{n,m_n} - \tau_{m_n}, \widehat{\mathfrak{H}}_{n,m_n} - \mathfrak{H}_{m_n}\right) \Rightarrow (Z_\tau, \mathbb{G}) \tag{A66}$$

in $\mathbb{R} \times \mathcal{S}_2^{sa}$, and

$$\sqrt{n}\left\{|\tau_{m_n} - \tau| + \|\mathfrak{H}_{m_n} - \mathfrak{H}\|_2\right\} \to 0. \tag{A67}$$

A sufficient first-stage condition is an asymptotically linear triangular-array expansion in $\mathbb{R} \times \mathcal{S}_2^{sa}$ with an L^2-convergent influence function and an $o_p(n^{-1/2})$ remainder.

THEOREM A1 (Spectral-tail limit and numerical directional bootstrap). *Suppose Assumption A4 holds and $\lambda_d(\mathfrak{H}) > 0$. Put $\lambda_\star = \lambda_d(\mathfrak{H})$. Let $P_>$ project onto eigenvalues strictly above $\lambda_\star$, let $P_=$ project onto the $\lambda_\star$ eigenspace, let $r = \operatorname{rank} P_>$, and put $k = d - r$. For a self-adjoint perturbation G, let $\mu_1(G) \geq \cdots \geq \mu_q(G)$ be the eigenvalues of $P_= G P_=$ on the q-dimensional boundary eigenspace. Then*

$$\sqrt{n}\{\widehat{\ell}_{d,n} - \ell_d(\tau, \mathfrak{H})\} \Rightarrow Z_\tau - \operatorname{tr}(P_> \mathbb{G}) - \sum_{j=1}^{k} \mu_j(\mathbb{G}), \tag{A68}$$

where $\widehat{\ell}_{d,n} = \widehat{\tau}_{n,m_n} - \kappa_d(\widehat{\mathfrak{H}}_{n,m_n})$. With an eigengap at d, the limit reduces to $Z_\tau - \operatorname{tr}(P_d \mathbb{G})$. Suppose a bootstrap pair $(Z^*_{\tau,n}, \mathbb{G}^*_n)$ consistently estimates the joint limit in (A66). *If $\epsilon_n \downarrow 0$ and $\sqrt{n}\,\epsilon_n \to \infty$, the conditional law of*

$$T^*_n = Z^*_{\tau,n} - \frac{\kappa_d(\widehat{\mathfrak{H}}_{n,m_n} + \epsilon_n \mathbb{G}^*_n) - \kappa_d(\widehat{\mathfrak{H}}_{n,m_n})}{\epsilon_n} \tag{A69}$$

converges to the law in (A68).

The eigenvalue-sum map is a spectral function with a directional derivative determined by the perturbation restricted to each boundary eigenspace (Lewis 1996). The resampling statement applies the directional delta method and its numerical derivative implementation (Fang and Santos 2019; Hong and Li 2018). Sieve inference for increasing-dimensional functionals provides the surrounding approximation framework (Chen and Pouzo 2015). For a declared loss tolerance ε, define the sufficient dimension

$$d_\varepsilon = \min\{d : \ell_d(\tau, \mathfrak{H}) \le \varepsilon\}. \tag{A70}$$

For the following finite search, assume a terminal dimension D with known tail $\ell_D = 0$. Set $U_{D,n} = 0$ and let the other bounds satisfy $\liminf_n \Pr\{\ell_d \le U_{d,n}\ \forall\, 0 \le d \le D\} \ge 1 - \alpha$. For $\varepsilon \ge 0$, define

$$\widehat{d}_{\varepsilon,1-\alpha} = \min\{d \in \{0, \ldots, D\} : U_{d,n} \le \varepsilon\}. \tag{A71}$$

COROLLARY A6 (Confidence-protected sufficient dimension). *Under the simultaneous coverage condition above,*

$$\liminf_{n\to\infty} \Pr\{\ell_{\widehat{d}_{\varepsilon,1-\alpha}} \le \varepsilon\} \ge 1 - \alpha. \tag{A72}$$

If d_ε *is separated by* $\ell_{d_\varepsilon - 1} > \varepsilon > \ell_{d_\varepsilon}$*, using* $\ell_{-1} = +\infty$*, and the simultaneous bounds converge uniformly to the tail losses, then* $\widehat{d}_{\varepsilon,1-\alpha} \to_p d_\varepsilon$*. For the stochastic cutoffs satisfying Theorem A1, a jointly consistent numerical directional bootstrap provides simultaneous bounds from the* $(1 - \alpha)$ *quantile of* $\max_d(-T^*_{d,n})$ *at a continuity point of its limiting distribution. The cutoff* $d = 0$ *uses the trace limit* Z_τ*, and the known terminal tail uses the fixed bound zero.*

The reported finite-rank Monte Carlo uses eigenvalues $(1, 0.50, 0.25, 0.06, 0.06, 0.02)$. Its tolerance $\varepsilon = 0.075$ gives $d_\varepsilon = 5$ with an eigenvalue multiplicity at the adjacent boundary. With sample size 1,000, 500 replications, and 249 multiplier draws, the plug-in rule selects a sufficient dimension in 0.350 of replications. The simultaneous numerical-directional rule selects a sufficient dimension in 0.996 of replications, has mean selected dimension 5.036, and its joint upper bounds cover all population tail losses in 0.966 of replications.

Appendix D. Numerical approximation

The computational results concern approximation of a declared structural quantity and numerical accuracy of an estimation criterion. The reward and transition closure conditions used by a reduced Bellman problem are stated in the main paper.

D.1. Multilevel corrections

Let $Y_\ell(\theta)$ denote a structural quantity computed at fidelity level ℓ. It may be a model-implied information operator, a moment, or a score contribution. Define

$$\Delta_\ell Y = Y_\ell - Y_{\ell-1}, \qquad \ell \geq 1. \tag{A73}$$

The identity

$$\mathbb{E}[Y_L] = \mathbb{E}[Y_0] + \sum_{\ell=1}^{L} \mathbb{E}[\Delta_\ell Y] \tag{A74}$$

supports a coupled estimator with independent sample averages across levels and common random inputs within each fine-coarse pair. Assume the bias, coupled correction variance, and per-draw cost satisfy

$$|\mathbb{E}[Y - Y_L]| \leq c_1 2^{-\alpha L}, \qquad V_\ell = \mathrm{Var}(\Delta_\ell Y) \leq c_2 2^{-\beta\ell}, \qquad C_\ell \leq c_3 2^{\gamma\ell}. \tag{A75}$$

The variance under level sample sizes N_ℓ is $\sum_\ell V_\ell/N_\ell$ and work is $\sum_\ell N_\ell C_\ell$. The cost-minimizing allocation for a variance budget is proportional to

$$N_\ell \propto \sqrt{V_\ell/C_\ell}. \tag{A76}$$

PROPOSITION A13 (Multilevel work). *Suppose* (A75) *holds and* $\alpha \geq \frac{1}{2}\min\{\beta, \gamma\}$. *A coupled multilevel estimator can attain RMSE* ε *with*

$$\mathrm{Work}_{ML}(\varepsilon) = \begin{cases} O(\varepsilon^{-2}), & \beta > \gamma, \\ O(\varepsilon^{-2}\log^2 \varepsilon^{-1}), & \beta = \gamma, \\ O(\varepsilon^{-2-(\gamma-\beta)/\alpha}), & \beta < \gamma. \end{cases} \tag{A77}$$

The corresponding single-level calculation at the finest required level has work $O(\varepsilon^{-2-\gamma/\alpha})$.

D.2. Multi-index state, angle, and solver refinement

Let $\ell = (\ell_s, \ell_u, \ell_v)$ index state resolution, intervention-angle resolution, and value-function or equilibrium-solver fidelity. In two coordinates,

$$\Delta_s \Delta_u Y_{ij} = Y_{ij} - Y_{i-1,j} - Y_{i,j-1} + Y_{i-1,j-1}. \tag{A78}$$

Sparse index sets retain mixed corrections with state-angle-fidelity contributions selected by their estimated target contribution and cost. The multi-index sampling rates follow from mixed-difference bias, variance, and cost exponents as in Haji-Ali, Nobile, and

Tempone (2016). For a current rank-d state projector, a candidate target view can be ranked by

$$G(t;s) = \mathrm{tr}\left[(I - P_d(s))A_t(s)'W_tA_t(s)\right]. \tag{A79}$$

A computational acquisition rule can integrate (A79) against a state uncertainty weight and choose the next target angle or fidelity block from the feasible menu.

D.3. Equivalence to an exact numerical criterion

Let Q_n^E be the sample structural criterion computed by an exact numerical solver and Q_n^{ML} the criterion computed by the multilevel scheme.

ASSUMPTION A5 (Numerical score and curvature). *The exact estimator is regular at an interior parameter* θ_0*:* $\sqrt{n}\nabla Q_n^E(\theta_0) \Rightarrow N(0,\Omega)$ *and* $\sup_{\theta\in N}\|\nabla^2 Q_n^E(\theta) - J\| \to_p 0$ *on a neighborhood* N*, where* J *is nonsingular. The numerical criterion satisfies*

$$\sup_{\theta\in N}\|\nabla Q_n^{ML}(\theta) - \nabla Q_n^E(\theta)\| = o_p(n^{-1/2}) \tag{A80}$$

and

$$\sup_{\theta\in N}\|\nabla^2 Q_n^{ML}(\theta) - \nabla^2 Q_n^E(\theta)\| = o_p(1). \tag{A81}$$

PROPOSITION A14 (Numerical equivalence). *Let* $\widehat{\theta}_E$ *and* $\widehat{\theta}_{ML}$ *be local minimizers of* Q_n^E *and* Q_n^{ML} *in* N*. Under Assumption A5,*

$$\sqrt{n}(\widehat{\theta}_{ML} - \widehat{\theta}_E) = o_p(1) \tag{A82}$$

and

$$\sqrt{n}(\widehat{\theta}_{ML} - \theta_0) \Rightarrow N(0, J^{-1}\Omega J^{-1\prime}). \tag{A83}$$

Appendix E. Diagnostic structural experiments

The numerical diagnostics examine the representation, estimation, design, and computation steps separately. Appendix G gives their construction.

TABLE A1. Diagnostics for the structural representation

Exercise	Object	Result
Exchangeable-state DDC	Planning closure at $K = 14$	245,760 Bellman states map to 960 count states; mapped values and choice probabilities agree to numerical precision.
Target-dependent DDC	Two-coordinate held-out response fit	Counterfactual coordinates: 0.225 response SD; baseline-value coordinates: 1.002 response SD.
Scattered response field	Cross-fitted local quadratic patches	Projector error falls as anchors increase; same-sample squares retain the quadratic noise component.
Intervention allocation	Goal-oriented posterior target loss	Added policy views reduce the residual information directions selected by the target criterion.
State-angle refinement	Multi-index mixed corrections	State and angle resolution enter separate telescoping axes.
Loss-controlled dimension	$\varepsilon = 0.075$ with an adjacent eigenvalue multiplicity	Plug-in sufficiency: 0.350; simultaneous numerical-directional rule: 0.996; mean protected dimension: 5.036.

The target-dependent DDC sets the dominant baseline value-function directions apart from the finite-counterfactual directions. Figure A1 reports held-out response error as the retained dimension varies. The exact-closure exercise separates target sufficiency from Bellman sufficiency. In that design the reward and pushforward transition law factor through two count states, so the Bellman problem closes exactly. The scattered and boundary-spectrum exercises evaluate the first-stage and inferential conditions used by Proposition A10 and Theorem A1. The multi-index exercise evaluates the numerical decomposition in Online Appendix D.

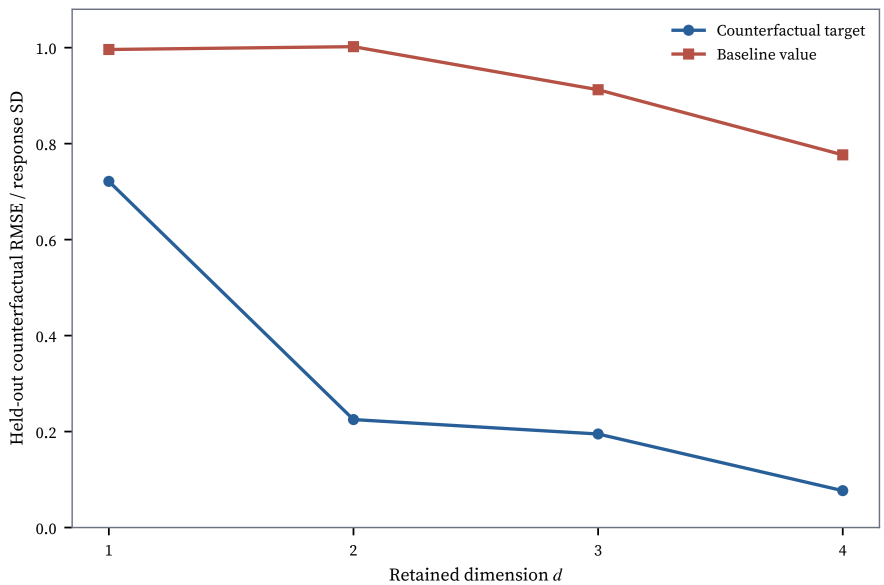

FIGURE A1. Counterfactual response error by retained dimension

Each representation is estimated on a training sample of states and evaluated on held-out finite policy responses. The target coordinates use the counterfactual derivative field. The baseline-value coordinates use the gradient field of the baseline value function. Errors are divided by the held-out response standard deviation.

Appendix F. Proofs

F.1. Proof of Theorem 1

PROOF. If $z(s) = z(s')$, the factorization $q = q_z \circ z$ gives $q(s) = q(s')$ and $\boldsymbol{R}_{\mathcal{T}} = R_z \circ z$ gives $\boldsymbol{R}_{\mathcal{T}}(s) = \boldsymbol{R}_{\mathcal{T}}(s')$. Hence s and s' belong to the same counterfactual-equivalence class. Define $h(z(s)) = \pi^\star(s)$. The preceding implication makes h well defined and gives $\pi^\star = h \circ z$ on the attained image. For the local statement, let $V = \ker Dq(s)$ and $A = D\boldsymbol{R}_{\mathcal{T}}(s)|_V$. Write $C = Db(s)|_V$. First-order closure implies $\ker C \subseteq \ker A$. Rank-nullity gives $\operatorname{rank} C = \dim V - \dim\ker C \geq \dim V - \dim\ker A = \operatorname{rank} A$. Since $\operatorname{rank} C \leq d$, $d \geq \operatorname{rank} A$. For attainment, choose a complement $V = \ker A \oplus W$ with $\dim W = \operatorname{rank} A$. Define a linear map $C^\star : V \to \mathbb{R}^{\operatorname{rank} A}$ that vanishes on $\ker A$ and is an isomorphism on W. Extend $C^\star$ to the full tangent space and realize the extension as the derivative at s of a C^1 map in a local chart. Then $\ker Dq(s) \cap \ker Db(s) = \ker A$, which gives first-order closure. □

F.2. Proof of Proposition 2

PROOF. For $v \in \mathbb{R}^k$,

$$v' H_{\mathcal{T}}(s) v = \sum_{t \in \mathcal{T}} \omega_t \| W_t^{1/2} A_t(s) v \|_2^2. \tag{A84}$$

Every summand is nonnegative and each ω_t is positive. Hence the quadratic form is zero exactly for vectors in every $\ker\{W_t^{1/2} A_t(s)\}$. The stacked weighted map $A_{\mathcal{T},W}$ has Gram operator $H_{\mathcal{T}}$, so they have the same kernel and rank. If W_t is positive definite on $\operatorname{im} A_t(s)$, then $W_t^{1/2} A_t(s) v = 0$ exactly when $A_t(s) v = 0$. Frame equivariance follows by substituting $\widetilde{A}_t = A_t O$. □

F.3. Proof of Proposition A1

PROOF. Write a vertical vector as $Zz = ZC\tilde{z}$. The pullback metric gives $\widetilde{G} = C'GC$. Each target derivative in the new coordinate is $A_t C$, so its quadratic contribution is $C' A_t' W_t A_t C$ and $\widetilde{Q} = C'QC$. If $Qv = \lambda G v$ and $v = C\tilde{v}$, left multiplication by C' gives $\widetilde{Q}\tilde{v} = \lambda \widetilde{G}\tilde{v}$. Hence the generalized spectrum is invariant and the tangent direction represented by v is the same geometric vector as the direction represented by $\tilde{v}$. For target coordinates satisfying $y = D\tilde{y}$, the identity $y' W y = \tilde{y}' D' W D \tilde{y}$ gives the stated metric transformation and leaves the target quadratic form unchanged. □

F.4. Proof of Proposition A2 and Corollary A1

PROOF. For each target view, $L_t = W_t^{1/2} A_t$ is Hilbert–Schmidt and $A_t^* W_t A_t = L_t^* L_t$ is positive trace class with $\|L_t^* L_t\|_1 = \|L_t\|_{\mathrm{HS}}^2$. Condition (A6) therefore gives Bochner integrability in the trace-class Banach space. The integral in (A7) is positive and self-adjoint, and the trace identity (A8) follows from continuity of the trace on positive trace-class operators. Trace-class operators are compact. For $v \in \mathbb{V}_s$,

$$\langle \mathfrak{H}_{\mathfrak{T},s} v, v \rangle_s = \int_{\mathfrak{T}} \omega_t \| W_t^{1/2} A_t v \|_{\mathbb{Y}_t}^2 \, d\nu(t). \tag{A85}$$

The integrand is measurable and nonnegative. The integral is zero exactly when $W_t^{1/2} A_t v = 0$ for ν-almost every target. Positive definiteness of W_t on the target image gives (A9). For the sieve statement, let $\Pi_m \to I$ strongly. A trace-class operator admits finite-rank approximants in trace norm. For a finite-rank operator F, strong convergence of Π_m is uniform on the finite-dimensional range of F and F^*, which gives $\|F - \Pi_m F \Pi_m\|_1 \to 0$. Given $\epsilon > 0$, choose finite-rank F with $\|\mathfrak{H} - F\|_1 < \epsilon$. Since orthogonal projections are contractions,

$$\|\mathfrak{H} - \Pi_m \mathfrak{H} \Pi_m\|_1 \le \|\mathfrak{H} - F\|_1 + \|F - \Pi_m F \Pi_m\|_1 + \|\Pi_m (F - \mathfrak{H}) \Pi_m\|_1 \tag{A86}$$

$$\le 2\epsilon + \|F - \Pi_m F \Pi_m\|_1. \tag{A87}$$

Taking m to infinity and then ϵ to zero gives (A10). Trace-norm perturbation bounds for self-adjoint compact operators give (A11), and operator-norm convergence follows from $\|\cdot\|_{\mathrm{op}} \le \|\cdot\|_1$. A separated spectral projector therefore converges by the Davis–Kahan theorem. Ky Fan's principle gives the eigenvalue-tail loss. For Corollary A1, write an absolute perturbation as $h = \mu s$. The Fisher pairing becomes $\mathbb{E}_\mu[ss']$. The derivative of the supplied moment $\mathbb{E}_\mu[\phi_j]$ in direction s is $\mathbb{E}_\mu[\phi_j s]$, which gives (A12). The Riesz representation of the restricted target derivative is the orthogonal projection $\psi_t^{\mathcal{V}}$. A scalar target therefore contributes the rank-one operator $\psi_t^{\mathcal{V}} \otimes \psi_t^{\mathcal{V}}$, giving (A13). For a finite family, collect the score gradients as columns of an operator $\Psi : \mathbb{R}^M \to \mathbb{V}_\mu$. The nonzero eigenvalues of $M^{-1}\Psi\Psi^*$ equal those of $M^{-1}\Psi^*\Psi$, which is the view Gram matrix. □

F.5. Proof of Theorem 2

PROOF. Let $\mathcal{A}$ be the stacked weighted operator with block rows $\sqrt{\omega_t} W_t^{1/2} A_t$. Then $H_{\mathfrak{T}} = \mathcal{A}'\mathcal{A}$. For a rank-$d$ orthogonal projector P, isotropic quadratic loss from discarding $(I-P)v$ equals $\mathrm{tr}\{\mathcal{A}(I-P)\mathcal{A}'\} = \mathrm{tr}[(I-P)H_{\mathfrak{T}}]$. The Ky Fan variational principle maximizes $\mathrm{tr}(PH_{\mathfrak{T}})$ over rank-d projectors at the leading eigenspace, giving (18). Worst-case norm loss over $\|v\| \le \rho$ equals $\rho\|\mathcal{A}(I-P)\|_{\mathrm{op}}$. The Eckart–Young singular-value characterization gives the minimum $\rho\sigma_{d+1}(\mathcal{A}) = \rho\sqrt{\lambda_{d+1}(H_{\mathfrak{T}})}$. Equation (20) follows from the eigenvalue tail. □

F.6. Proof of Proposition A5

PROOF. Suppose the coordinate map exists. Smooth vector fields X and Y taking values in $\mathcal{K}_d$ annihilate every component f of (q, b). The identity $[X, Y]f = X(Yf) - Y(Xf) = 0$ gives $[X, Y] \in \ker D(q, b) = \mathcal{K}_d$, so $\mathcal{K}_d$ is involutive. Conversely, the constant-rank theorem gives local coordinates (x, y) with $q(x, y) = x$ and vertical space spanned by the y coordinates. Frobenius gives adapted vertical coordinates (b, w) on a smaller neighborhood in which $\mathcal{K}_d$ is spanned by the w directions. In the combined coordinates (x, b, w), the differential kernel of (q, b) is $\mathcal{K}_d$. □

F.7. Proof of Proposition A6

PROOF. Add and subtract the leading spectral projector:

$$\operatorname{tr}\{(I - P_b)H_{\mathcal{J}}\} = \operatorname{tr}\{(I - P_d)H_{\mathcal{J}}\} + \operatorname{tr}\{(P_d - P_b)H_{\mathcal{J}}\}. \tag{A88}$$

The first term is the spectral tail. Ky Fan's principle makes the second term nonnegative. Under an eigengap, equality in the rank-d Ky Fan problem identifies $P_b = P_d$. Trace duality gives (A25). Integration gives (A26). If an admissible $b^\star$ attains the infimum, equality holds exactly when $\mathbb{E}\Delta_{b^\star} = 0$. Nonnegativity implies $\Delta_{b^\star} = 0$ almost everywhere, and the positive eigengap gives alignment there. Conversely, an aligned admissible field attains the spectral lower bound. □

F.8. Proof of Proposition 4

PROOF. Conditional expectation is the orthogonal projection of $\boldsymbol{F}_{\mathcal{J}}(S)$ onto the closed L^2 subspace measurable with respect to $z(S)$. Apply (33) to every coordinate of $\boldsymbol{F}_{\mathcal{J}}$ in an orthonormal output basis and sum. The tangent space of a regular level set is $\ker Dz(s)$. Integration over the disintegration of μ gives (34). For $z = (q, b)$,

$$\|D\boldsymbol{F}_{\mathcal{J}}(s)\Pi_{\ker D(q,b)(s)}\|_{\mathrm{HS}}^2 = \operatorname{tr}\{(I - P_b(s))H_{\mathcal{J}}(s)\}. \tag{A89}$$

Proposition A6 gives (35). Alignment and the uniform bound on C give (36). □

F.9. Proof of Proposition 5 and Proposition 1

PROOF. The Hilbert-space conditional-variance identity gives

$$\mathfrak{R}(z) = \mathbb{E}\|\boldsymbol{F}_{\mathcal{J}} - \mathbb{E}[\boldsymbol{F}_{\mathcal{J}} \mid z, C_z]\|_2^2 \tag{A90}$$

$$+ \mathbb{E}\|\mathbb{E}[\boldsymbol{F}_{\mathcal{J}} \mid z, C_z] - \mathbb{E}[\boldsymbol{F}_{\mathcal{J}} \mid z]\|_2^2. \tag{A91}$$

Apply the componentwise conditional Poincaré inequality to the first term. The second term equals the final term in (39). Vanishing total risk is equivalent to almost-sure factorization through z. For the finite-state result, assign a distinct real number to every class of the finite exact target quotient and define $b(s)$ as the label of the class containing s. Equality of (q, b) implies equality of the target class. □

F.10. Proof of Proposition 3

PROOF. Let $X_1, \ldots, X_d, Y$ be independent and uniform on $[-1, 1]$, let q be constant, put $\phi_k(u) = \cos\{k\pi(u+1)/2\}$, fix $A > 1$, and set

$$F(X_1, \ldots, X_d, Y) = A \sum_{j=1}^{d-1} \phi_2(X_j) + \epsilon\phi_N(X_d) + \phi_1(Y), \tag{A92}$$

where the sum is empty when $d = 1$. Choose an even integer N. Since $\mathbb{E}\phi_k'(U) = 0$ for even k and U uniform on $[-1, 1]$, independence removes every cross term. Apart from the common factor $\pi^2/8$, the integrated derivative operator is diagonal with entries $4A^2$ for $X_1, \ldots, X_{d-1}$, $\epsilon^2 N^2$ for X_d, and 1 for Y. Choose $\epsilon > 0$ such that $1/\epsilon^2 \geq K$, and then choose the even N large enough that $\epsilon^2 N^2 > \max\{1, \rho/(1-\rho)\}$. The leading d-dimensional eigenspace is therefore spanned by $X_1, \ldots, X_d$, and its derivative share is

$$\frac{4(d-1)A^2 + \epsilon^2 N^2}{4(d-1)A^2 + \epsilon^2 N^2 + 1} > \rho. \tag{A93}$$

Compare $b_{\mathrm{AS}} = (X_1, \ldots, X_d)$ with $b_{\mathrm{alt}} = (X_1, \ldots, X_{d-1}, Y)$. Independence gives

$$\mathfrak{R}(q, b_{\mathrm{AS}}) = \mathrm{Var}\{\phi_1(Y)\} = \frac{1}{2}, \qquad \mathfrak{R}(q, b_{\mathrm{alt}}) = \mathrm{Var}\{\epsilon\phi_N(X_d)\} = \frac{\epsilon^2}{2}. \tag{A94}$$

Hence $\mathfrak{R}(q, b_{\mathrm{AS}})/\mathfrak{R}(q, b_{\mathrm{alt}}) = 1/\epsilon^2 \geq K$. □

F.11. Proof of Theorem 4

The proof is given immediately after Theorem 4 in the main paper.

F.12. Proof of Proposition 6

PROOF. Fix a cell g and suppress the candidate and cell indices. Because the graph is connected and L is nonnegative and self-adjoint in the empirical inner product, it has an orthonormal eigenbasis $\{\phi_k\}_{k=0}^{n_g-1}$ with $L\phi_k = \nu_k\phi_k$ and $0 = \nu_0 < \nu_1 \leq \cdots$. Each centered residual coordinate has an expansion $r_\ell^\circ = \sum_{k\geq 1} a_{\ell k}\phi_k$. Its contribution to certificate minus risk and to the diagnostic is

$$\sum_{\ell,k\geq 1}\left(\frac{\nu_k}{\nu_1}-1\right)a_{\ell k}^2 \quad \text{and} \quad \sum_{\ell,k\geq 1}\left(\frac{\nu_k}{\nu_1}-1\right)^2 a_{\ell k}^2. \tag{A95}$$

The first expression is nonnegative. Cauchy–Schwarz bounds it by the square root of the centered residual norm times the diagnostic. Weighting by π_g, summing across cells, and applying Cauchy–Schwarz once more gives (67)–(68); the cell-mean term enters risk and certificate equally. Conditional on the auxiliary sample, Hoeffding's inequality bounds the deviation for one candidate by $M_G\sqrt{\log(2|\mathfrak{B}_{d,G}|/\alpha)/(2n_G)}$. A union bound over the finite menu gives (69). □

F.13. Proof of Theorem 5

The proof is given immediately after Theorem 5 in the main paper.

F.14. Proof of Corollary A3

PROOF. Intersect the empirical-process event in (56) with the two events in (A46). A union bound gives probability at least $1-\alpha-\beta_{1n}-\beta_{2n}$. On this intersection, Theorem 5 and $\mathfrak{R}_d^\star \leq M$ give

$$\sqrt{\{\mathfrak{R}_d^\star - \Gamma(b_d^\star)\}\mathfrak{D}(b_d^\star)} \leq \sqrt{M\left\{\sup_{\eta\in\Xi_d}\widehat{\mathfrak{D}}_n(b_\eta)+s_n\right\}}, \tag{A96}$$

which proves (A47). Intersecting only (56) and the first-stage score event, then applying the relative-slack condition, proves (A48). □

F.15. Proof of Proposition A8 and Corollary A2

PROOF. For any $v\in\mathbb{R}^p$,

$$v' I_{\mathfrak{D}} v = \sum_{e:n_e>0} n_e\|\Sigma_e^{-1/2}J_e v\|_2^2. \tag{A97}$$

Hence (A33) follows. The linear-factorization lemma states that linear maps J and T satisfy $\ker J\subseteq\ker T$ exactly when $T=LJ$ for a linear map L on $\operatorname{im} J$. Restrict T_θ to $V=\ker J_{\mathfrak{D}}$. If C is the derivative of d added scalar moments and target identification holds after augmentation, then $\ker(C|_V)\subseteq\ker(T_\theta|_V)$. Rank-nullity gives $d\geq\operatorname{rank}(T_\theta|_V)$. A map with kernel equal to $\ker(T_\theta|_V)$ and image dimension $r_{\mathfrak{D}\to\tau}$ attains the count. The rank formula in (A35) is the rank increment from stacking T_θ below $J_{\mathfrak{D}}$. Corollary A2 subtracts the two rank increments. Nonnegativity follows from submodularity of row-space dimension. □

F.16. Proof of Lemma A1

PROOF. Expand $\widehat{A}^{(1)\prime}W\widehat{A}^{(2)}$ as $A'WA + E^{(1)\prime}WA + A'WE^{(2)} + E^{(1)\prime}WE^{(2)}$ and take the conditional expectation. The stated moment conditions remove the three error terms. The same-sample identity follows from the expansion $\widehat{A}'W\widehat{A} = A'WA+E'WA+A'WE+E'WE$. □

F.17. Proof of Proposition A10

PROOF. Write $E_{it}^{(b)} = \widehat{A}_{it}^{(b)} - A_t(s_i)$. For one target view,

$$\frac{1}{2}\{\widehat{A}^{(1)\prime}W\widehat{A}^{(2)} + \widehat{A}^{(2)\prime}W\widehat{A}^{(1)}\} - A'WA = \frac{1}{2}\{E^{(1)\prime}WA + A'WE^{(2)} + E^{(2)\prime}WA + A'WE^{(1)}\} \tag{A98}$$

$$+ \frac{1}{2}\{E^{(1)\prime}WE^{(2)} + E^{(2)\prime}WE^{(1)}\}. \tag{A99}$$

Bounded A_t, W_t, and weights together with (A40) give

$$\max_i \|\widetilde{H}_i - H_{\mathfrak{J}}(s_i)\|_{\mathrm{op}} = O_p(\zeta_n + \zeta_n^2). \tag{A100}$$

Next decompose

$$\widehat{H}_n(s) - H_{\mathfrak{J}}(s) = \sum_i \kappa_{ni}(s)\{\widetilde{H}_i - H_{\mathfrak{J}}(s_i)\} + \sum_i \kappa_{ni}(s)\{H_{\mathfrak{J}}(s_i) - H_{\mathfrak{J}}(s)\}. \tag{A101}$$

The first term is uniformly $O_p(\zeta_n + \zeta_n^2)$. Hölder continuity of A_t and boundedness imply Hölder continuity of $H_{\mathfrak{J}}$ with the same exponent, so the second term is uniformly $O(h_n^{\beta})$ on the support of the partition weights. Combining the bounds gives (A45). □

F.18. Proof of Proposition A9

PROOF. The Woodbury identity gives

$$(M + J'R^{-1}J)^{-1} = \Sigma - \Sigma J'(R + J\Sigma J')^{-1}J\Sigma. \tag{A102}$$

Taking the trace against $H_{\mathfrak{J}}$ yields (A38). For a scalar view, write $J = a'$ and $R = \nu^{-1}$. The middle inverse becomes $(\nu^{-1} + a'\Sigma a)^{-1}$, which gives (A39). □

F.19. Proof of Corollary A4

PROOF. Write $\widehat{O}_i = \widehat{O}_{i\to s}$ and $O_i = O_{i\to s}$. For each active anchor,

$$\widehat{O}_i\widetilde{H}_i\widehat{O}_i' - H_{\mathfrak{J}}(s) = \widehat{O}_i\{\widetilde{H}_i - H_{\mathfrak{J}}(s_i)\}\widehat{O}_i' \tag{A103}$$

$$+\{\widehat{O}_i H_{\mathfrak{J}}(s_i)\widehat{O}_i' - O_i H_{\mathfrak{J}}(s_i) O_i'\} \tag{A104}$$

$$+\{O_i H_{\mathfrak{J}}(s_i) O_i' - H_{\mathfrak{J}}(s)\}. \tag{A105}$$

The first term has operator norm $O_p(\zeta_n + \zeta_n^2)$ uniformly over active anchors. The second term is bounded by a constant times $\|\widehat{O}_i - O_i\|_{\mathrm{op}}$ because the target operators and orthogonal transports are uniformly bounded. The third term is $O(h_n^{\beta})$ by (A50). Convex averaging with the partition weights gives (A51). □

F.20. Proof of Corollary A5

PROOF. Equation (A54) follows from Weyl's eigenvalue inequality. Davis–Kahan applied under Assumption A2 gives (A55). Let $E = \widehat{H} - H_{\mathfrak{J}}$. Since $\widehat{P}_d$ maximizes $\mathrm{tr}(P\widehat{H})$ over rank-d projectors,

$$\mathrm{tr}\{(P_d - \widehat{P}_d)\widehat{H}\} \le 0. \tag{A106}$$

Therefore

$$\mathcal{L}_s(\widehat{P}_d) - \mathcal{L}_s(P_d) = \mathrm{tr}\{(P_d - \widehat{P}_d) H_{\mathfrak{J}}\} \tag{A107}$$

$$\le \mathrm{tr}\{(P_d - \widehat{P}_d)(H_{\mathfrak{J}} - \widehat{H})\} \tag{A108}$$

$$\le \|P_d - \widehat{P}_d\|_* \|E\|_{\mathrm{op}}. \tag{A109}$$

The difference of two rank-d projectors has rank at most $2d$, so $\|P_d - \widehat{P}_d\|_* \le 2d\|P_d - \widehat{P}_d\|_{\mathrm{op}}$. On $\|E\|_{\mathrm{op}} < \gamma_d/2$, the operator-norm Davis–Kahan bound gives $\|P_d - \widehat{P}_d\|_{\mathrm{op}} \le 2\|E\|_{\mathrm{op}}/\gamma_d$. Combining the inequalities gives (A56). For fixed s and separated eigenvalues, the Fréchet derivative of the rank-d spectral projector in direction E is

$$DP_H[E] = \sum_{j \le d} \sum_{k > d} \frac{v_k v_k' E v_j v_j' + v_j v_j' E v_k v_k'}{\lambda_j - \lambda_k}. \tag{A110}$$

An asymptotically linear expansion for $\widehat{H}(s)$ and the functional delta method applied to (A110) give a Gaussian limit for $\mathrm{vec}\,\widehat{P}_d(s)$. □

F.21. Proof of Proposition A12

PROOF. Write $\|\cdot\|_{\Omega}$ for the weighted field norm. The derivative obeys $\|G_H u\|_{\Omega}^2 = u' \mathcal{I}_H u$. If $\lambda_{\min}(\mathcal{I}_H) = a^2 > 0$, then $\|G_H u\|_{\Omega} \ge a\|u\|$. Fréchet differentiability gives

$$\|h_{\theta_0+u} - h_{\theta_0}\|_{\Omega} \ge a\|u\| - o(\|u\|) > 0$$

for sufficiently small nonzero u. This proves local identification. At deficient rank, the Gram identity identifies the kernel and range of the first-order field map. Stacking differentiable moments gives another Gram identity. A trivial derivative kernel is equivalent to full column rank in the finite-dimensional parameter space; applying the same lower-bound argument proves sufficiency for local identification. □

F.22. Proof of Proposition A11

PROOF. Proposition A10 and the structural-solver condition imply

$$\sup_{\theta,s} \left\| \{\widetilde{H}_{\mathcal{T},\theta,n}(s) - \widehat{H}_n(s)\} - \{H_{\mathcal{T},\theta}(s) - H_0(s)\} \right\|_F = o_p(1). \tag{A111}$$

Uniform boundedness and the uniform law of large numbers then give uniform convergence of the field-distance term in (A61) to its population counterpart in (A57). Adding uniform convergence of the behavioral criterion yields $\sup_{\theta\in\Theta} |\widehat{Q}_n(\theta) - Q(\theta)| = o_p(1)$. Compactness, continuity, and the unique minimizer condition give (A62) by the argmin theorem. □

F.23. Proof of Theorem A1

PROOF. Assumption A4 and Slutsky's theorem give joint convergence of $(\widehat{\tau}_{n,m_n}, \widehat{\mathfrak{H}}_{n,m_n})$ at the root-n scale around $(\tau, \mathfrak{H})$. The map $\kappa_d(H) = \sum_{j=1}^{d} \lambda_j(H)$ is Hadamard directionally differentiable on self-adjoint Hilbert–Schmidt perturbations at compact H when $\lambda_d(H) > 0$. Let $\lambda_\star = \lambda_d(H)$ and decompose the leading spectral space into the eigenspaces strictly above $\lambda_\star$ and the boundary eigenspace. Fan's variational formula gives

$$\kappa'_{d,H}(G) = \operatorname{tr}(P_{>}G) + \sum_{j=1}^{d-r} \lambda_j(P_{=}GP_{=}|_{\operatorname{im} P_{=}}). \tag{A112}$$

The directional delta method applied to $(\tau, H) \mapsto \tau - \kappa_d(H)$ yields (A68). When $\lambda_d > \lambda_{d+1}$, the boundary eigenspace is fully retained and the displayed derivative equals $\operatorname{tr}(P_d G)$. For the bootstrap, define the numerical derivative of κ_d at $\widehat{H}$ in direction $\mathbb{G}_n^*$ with step ϵ_n. The conditions $\epsilon_n \to 0$ and $\sqrt{n}\,\epsilon_n \to \infty$ place the base-point estimation error below the numerical derivative step. Conditional weak convergence of the bootstrap pair and the numerical directional delta method give conditional convergence of (A69) to the directional derivative law. □

F.24. Proof of Corollary A6

PROOF. The choice $U_{D,n} = 0$ makes the search set nonempty for every sample and every $\varepsilon \geq 0$. On the simultaneous coverage event, $U_{\widehat{d},n} \leq \varepsilon$ implies $\ell_{\widehat{d}} \leq U_{\widehat{d},n} \leq \varepsilon$. The probability of this event has limit inferior at least $1 - \alpha$, which gives (A72). Under separation and uniform convergence of the bounds, the inequalities defining d_ε hold with the same ordering for the sample bounds with probability approaching one. □

F.25. Proof of Proposition 7

PROOF. For any bounded function $\bar{V}$ on the reduced state, define $V(s) = \bar{V}(z(s))$. Equations (70) and (71) imply that each action-specific Bellman update of V depends on s through $z(s)$. Hence the Bellman operator maps functions that factor through z into the same class. Contraction of the discounted log-sum Bellman operator gives a fixed point in that class, establishing $V_\theta(s, p) = \bar{V}_\theta(z(s), p)$. The logit conditional choice probabilities are functions of the action-specific value indices and therefore also factor through z. □

F.26. Proof of Proposition A13

PROOF. Choose L so that $2^{-\alpha L}$ is of order ε. For fixed L, minimize $\sum_{\ell=0}^{L} N_\ell C_\ell$ subject to $\sum_{\ell=0}^{L} V_\ell / N_\ell \leq \varepsilon^2/2$. The Lagrange first-order conditions give N_ℓ proportional to $\sqrt{V_\ell / C_\ell}$ and minimum stochastic work proportional to

$$\varepsilon^{-2}\left(\sum_{\ell=0}^{L} \sqrt{V_\ell C_\ell}\right)^2. \tag{A113}$$

Under (A75), the summand is of order $2^{(\gamma-\beta)\ell/2}$. The geometric sum is bounded when $\beta > \gamma$, grows at rate L when $\beta = \gamma$, and grows at rate $2^{(\gamma-\beta)L/2}$ when $\beta < \gamma$. Substituting $L \asymp \alpha^{-1} \log_2 \varepsilon^{-1}$ gives (A77). A single-level estimator uses $O(\varepsilon^{-2})$ draws at level L, each with cost $O(2^{\gamma L})$, giving the stated work rate. □

F.27. Proof of Proposition A14

PROOF. The first-order condition for the multilevel estimator and a mean-value expansion around the exact estimator give

$$0 = \nabla Q_n^{ML}(\widehat{\theta}_E) + \widetilde{J}_n(\widehat{\theta}_{ML} - \widehat{\theta}_E), \tag{A114}$$

where $\widetilde{J}_n$ is a Hessian evaluated along the segment between the two estimators. The exact first-order condition gives $\nabla Q_n^E(\widehat{\theta}_E) = 0$. Assumption A5 implies $\nabla Q_n^{ML}(\widehat{\theta}_E) = o_p(n^{-1/2})$

and $\widetilde{J}_n \to_p J$. Invertibility of J gives (A82). Adding the exact-estimator expansion gives (A83). □

F.28. Representation, finite-radius, decision, and measurement results

Proof of the universal property in Theorem 1. The maps $\bar{q}$ and $\bar{R}_{\mathcal{T}}$ are well defined because both objects are constant on each equivalence class. Let z carry q and $\boldsymbol{R}_{\mathcal{T}}$. Define $h(z(s)) = \pi^\star(s)$. Equality $z(s) = z(s')$ implies equality of q and $\boldsymbol{R}_{\mathcal{T}}$, hence equality of their quotient classes. This makes h well defined on $z(\mathcal{S})$. The identity $\pi^\star = h \circ z$ follows by construction, and surjectivity of $z : \mathcal{S} \to z(\mathcal{S})$ gives uniqueness. □

Proof of Proposition A4. For the upper bound, let $v_1, \ldots, v_d$ be the leading right singular vectors of $\mathcal{A}_s^{\mathcal{T}}$ and set $b(v) = (v_1'v, \ldots, v_d'v)$. The truncated singular expansion gives linear approximation error at most $\rho\sigma_{d+1}$, and (22) adds $K_s\rho^2/2$. For the lower bound, restrict any continuous encoder to the radius-ρ sphere in the span of the first $d+1$ right singular vectors. Borsuk–Ulam gives an antipodal pair $v, -v$ with $b(v) = b(-v)$ (Borsuk 1933). The linear target difference has norm at least $2\rho\sigma_{d+1}$. The two remainders reduce the pairwise separation by at most $K_s\rho^2$. A common decoder value for the pair has maximum error at least one half of their separation. The affine case sets the remainder to zero. □

The proof is given immediately after Theorem 3 in the main paper.

Proof of Proposition A3. The first-order condition is $D_a L(a^\star(r), r) = 0$. Differentiation at r_0 gives (A14). Taylor expansion of $a \mapsto L(a, r_0)$ around a_0 gives one half of the action-Hessian quadratic form plus a smaller-order term. Substitution of (A14) yields (A15). □

Proof of Proposition A7. On a fiber of $M_{\mathcal{D}}$, the raw target derivative has the rank in (A28). The joint-rank identity follows from rank-nullity. First-order exact recovery is equivalent to inclusion of the data-fiber tangent space in the raw target kernel. Any k added scalar fields have derivative rank at most k on the data fiber, which gives the exact-recovery lower bound. A basis of the raw target image on the data fiber gives equality at the derivative level. The residual operator is $A'_{\mathcal{T},W,\mathcal{D}} A_{\mathcal{T},W,\mathcal{D}}$, so its rank equals $\delta_{\mathcal{T},W|\mathcal{D}}$ and (A29) follows. The same rank argument applied to the weighted target map gives the loss-repair lower bound and derivative-level attainment. Positive definiteness on the attainable residual image equates the raw and weighted kernels and therefore the two deficiencies. After d orthonormal measurement directions are added inside the unresolved fiber, the remaining isotropic quadratic target loss is the trace of $H_{\mathcal{T}|\mathcal{D}}$ on their orthogonal complement. The Ky Fan variational principle gives (A30). Under the neighborhood rank conditions, constant-rank coordinates write $M_{\mathcal{D}}(x, u) = x$ and zero raw deficiency gives $D_u \boldsymbol{R}_{\mathcal{T}} = 0$. Integration along connected local u-fibers defines the local factor. □

Appendix G. Simulation and quantitative design

G.1. Exchangeable-state DDC

Each hidden binary state in group $g \in \{1, 2\}$ follows a two-state Markov chain. Group 1 uses $\Pr(1|0) = 0.12$ and $\Pr(1|1) = 0.88$; group 2 uses 0.18 and 0.82. Conditional independence and exchangeability imply that the next-period count distribution depends on the current hidden configuration through the two current counts. The mileage transition is deterministic: keep advances one bin up to the maximum and replacement resets mileage to zero. This establishes the pushforward condition in Proposition 7.

G.2. Six-dimensional DDC

Each coordinate takes values in $\{-1, 0, 1\}$. Baseline action payoffs are

$$u_0(z) = -0.10 - 0.45z_1^2 + 0.18z_4, \tag{A115}$$

$$u_1(z) = -1.15 + 1.35z_1 + 0.80z_4 - 0.18z_4^2. \tag{A116}$$

Action-dependent transitions shift the first and fourth coordinates. The first policy view adds $0.85[0.95z_2 + 0.18z_2z_1 + 0.12z_5^2]$ to u_1 and the second adds $0.85[0.95z_3 - 0.16z_3z_4 + 0.10z_6^2]$. The discount factor is 0.9. Gradients use second-order finite differences on the three-point tensor grid. Response surfaces use polynomial features through degree four and a ridge constant 10^{-8}.

G.3. Scattered recovery

For each selected anchor and each of the two counterfactual gradients, two independent Gaussian error vectors with coordinate standard deviation 0.04 are added. Sample sizes are 50, 100, 200, and 400. Each case uses 300 replications. The same-sample comparison uses $\widehat{A}'\widehat{A}$ and the cross-fitted construction uses the symmetrized product of two noisy estimates.

G.4. Two-dimensional multi-index DDC

The state lies on $[0, 1]^2$. Each action has a product transition kernel with action-dependent persistence and drift. The intervention direction is an angle on $[0, 2\pi)$. A finite intervention of magnitude 0.75 enters replacement utility through two state functions and a second-order angular interaction. State grids contain 5, 9, 17, and 33 points per coordinate. Angle grids contain 4, 8, 16, and 64 equally spaced points. The target operator averages outer products of finite-policy CCP-response gradients across the angle grid. Each Bellman solve uses tolerance 10^{-10}.

G.5. Rotating local geometry

The persistent state uses 61 equally spaced points on $[0,1]$. Four hidden coordinates take values in $\{-1,0,1\}$ and remain fixed across Bellman transitions. Keep advances the persistent state by one grid point and replacement resets it to zero. The two finite policy views load on (z_1, z_2) and (z_3, z_4) through angles $1.15\pi x - 0.2\pi$ and $0.85\pi x + 0.1\pi$ with state-dependent amplitudes. Local hidden-state derivatives use second-order finite differences. The scattered experiment samples 40 hidden configurations per anchor and view, adds independent Gaussian derivative noise with standard deviation 0.015, and uses Gaussian state weights. Anchor counts are 6, 10, 16, 24, and 36. Each case uses 200 replications.

G.6. Target-directed angle acquisition

Three hidden coordinates take values in $\{-1,0,1\}$ and remain fixed over the Bellman transition. The persistent mileage state has 15 bins. Twenty-four candidate angles use $a(\vartheta) \propto (\cos\vartheta, \sin\vartheta, 0.75\cos 2\vartheta)'$ with policy magnitude 0.8. Each angle-specific operator is replaced by its leading rank-one component for the design calculation. The target operator is the equal-weight average across the 24 components and is normalized to unit trace. The target posterior-variance criterion uses prior precision $0.2I$ and per-view information scale 3. The random comparison contains 1,000 acquisition sequences.

G.7. One-period heterogeneous-household diagnostic

The household chooses next-period assets on a 31-point grid over $[0,15]$. Income takes values $(0.5,1,1.5)$ and follows (72). Preferences have discount factor 0.96 and coefficient of relative risk aversion 2. The savings choice receives type-I extreme-value smoothing with scale 0.08. The baseline real interest rate is 0.02, the wage is one, and the transfer is zero. The stationary distribution is computed from the baseline transition matrix. The supplied representation records mass, mean assets, and mean income. The vertical basis is obtained from the null space of these three linear functionals. The Fisher metric is represented in that basis by $B' \operatorname{diag}(\mu^{-1})B$. The five policy views are one-period changes to the real interest rate, transfer, or wage. Policy reverts to baseline from period one onward. Target derivatives are computed from the impact transition and powers of the baseline transition. The target family contains three outcomes at five horizons. Each outcome is divided by its baseline aggregate level before target operators are formed.

G.8. General-equilibrium transition experiment

The general-equilibrium transition routine uses 15 equally spaced asset points over $[0,20]$; income follows (72); and preferences use $\beta = 0.94$, relative risk aversion 2, and extreme-value scale 0.8. Production has capital share 0.36 and depreciation 0.08. The stationary capital stock is found by bisection on $[4, 5]$ until 36 bisection steps have been completed. The resulting household transition matrix determines the stationary distribution. A transition path uses a ten-period terminal horizon. A backward household recursion produces period-specific choice probabilities at a proposed capital path. Forward propagation of the distribution produces an implied capital path, and damped fixed-point updates with weight 0.5 continue until the maximum capital-path discrepancy is below 10^{-8} or 100 iterations have been reached. The capital-income package has initial tax rate 0.25 and the labor-income package has initial rate 0.30; both decay by 0.5 per period. Current tax revenue is rebated through a lump-sum transfer. Responses are divided by the baseline level of each outcome before the target vector is formed. The supplied distributional functionals are total mass, aggregate assets, and the masses of the first two income states. If X stacks these four functions and μ_0 is the stationary mass vector, the construction takes the null space of $(\sqrt{\mu_0}X)'$ and rescales it by $\mu_0^{-1/2}$ to obtain the 41-dimensional Fisher-orthonormal basis B. For each standard-normal direction z_j normalized to unit length, the score is Bz_j and the finite-difference radius is

$$\epsilon_j = 0.08/\max_i |(Bz_j)_i|. \tag{A117}$$

This guarantees nonnegative perturbed masses $\mu_0(1 \pm \epsilon_j Bz_j)$. The calculation uses 46 normalized directions. Central differences produce the 40 target derivatives, and least squares on the normalized directions reconstructs the derivative map. The decision experiment uses 35 training directions and 11 test directions. The two package-response columns form $G(s)$. The disturbance equals negative one half of the sum of the two baseline package responses, and the action penalty is 2×10^{-7}. A finite-difference derivative of the optimal action with respect to vec G constructs the decision metric. Linear response decoders are fit on the training directions and evaluated on the held-out directions. The greedy dictionary contains asset powers two through four, asset-income interactions, three wealth cells, and income-specific low-wealth cells.

G.9. KW97 fixed-model stress test

The KW97 stress test uses eight model periods, 31 common solution draws, and full state-space solution without interpolation. The two parameterizations are the distributed `kw_97_extended_respy` and `kw_97_extended` specifications. Five policies change the schooling reward, white-collar wage, blue-collar wage, military wage, and both civilian wages. Common random numbers and aligned solved state spaces are used for every policy. The 6,520 core states each carry four latent types. A fixed split allocates 70 percent of core states to training and 30 percent to testing. For each supplied representation, two training folds estimate local linear policy-response derivatives. Symmetrized cross-products form the cross-fitted operators. The common target metric and state standardization are estimated on training states. Equal-budget methods then use the same quadratic ridge decoder to predict the 55-component finite-policy response vector on held-out core states. Reported values average two derivative-estimation runs for each parameterization and cover small, primary, and large supplied representations.

The finite-graph calculation uses the primary supplied representation and the same training normalization, split, targets, coordinate maps, and ridge penalties. A second two-fold split of the training core states forms out-of-fold decoder residuals. Let g index an exact supplied-state cell and let b_i denote the added coordinates of training state i in that cell. For distinct states in a cell, the graph weight is

$$w_{ij} = \left(1 + \frac{\|b_i - b_j\|_2^2}{h_g^2}\right)^{-1}, \tag{A118}$$

where h_g is the median positive pairwise distance in the cell. The diagonal weight is zero and $L_g = \operatorname{diag}(W_g\mathbf{1}) - W_g$. The complete positive-weight graph is connected whenever the cell contains at least two distinct coordinates. Singleton cells contribute their residual-mean term. The calculation evaluates the first positive eigenvalue, risk, certificate, diagnostic, and both sides of (67)–(68).

Candidate selection minimizes the training graph certificate separately by model, estimation run, and dimension. The held-out comparison is evaluated afterward. For $d = 3$, the selected baseline-value candidate is compared with the counterfactual-fiber candidate by resampling the 1,956 test core states with replacement 1,999 times. The percentile intervals in Table A2 condition on the fitted candidates, selected graph rule, structural model, and realized split. The resulting candidate diagnostics and bootstrap summary generate Figure A2.

The repeated-split comparison fixes the 55-target normalization and active mask at the original design before varying the core-state split over 20 seeds. Within each split, the state standardization, local operators, ridge penalties, out-of-fold decoder risks, graph objects, candidate rankings, and held-out losses are re-estimated from the assigned training states. All four latent types of a core state remain in the same partition, leaving 1,956 held-out core states per comparison. Minimum out-of-fold decoder risk matches the held-out oracle in all 80 model–run–split comparisons. It selects the baseline-value family in 75 comparisons, and every corresponding paired 1,999-draw interval excludes zero. The graph certificate matches 63 of 80 held-out oracles. The maximum graph-identity discrepancy across the repeated-split comparison is 1.78×10^{-15}.

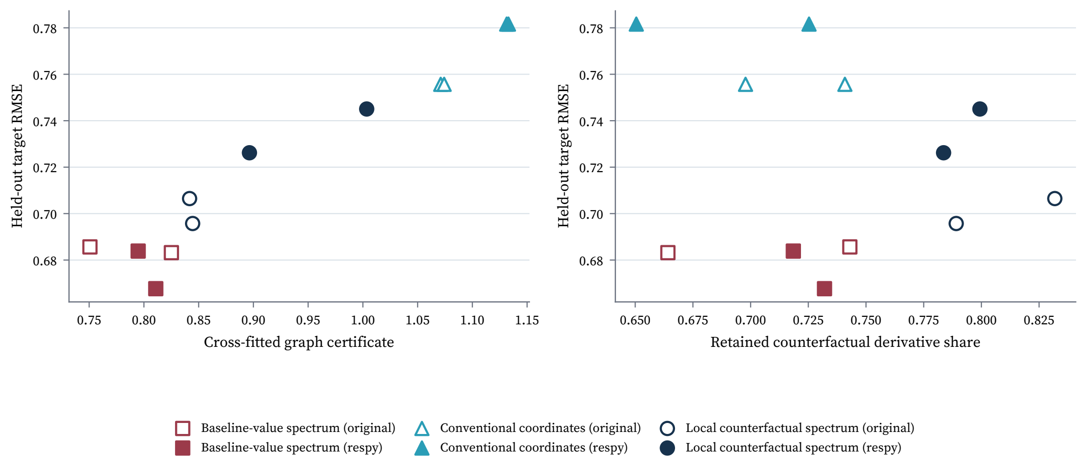

FIGURE A2. Generator certificate, derivative share, and held-out target risk

Each point is one candidate under one parameterization and estimation run at $d = 3$. Squares denote the baseline-value spectrum, triangles conventional coordinates, and circles the local counterfactual spectrum. Filled markers denote the distributed respy parameterization and hollow markers the distributed original parameterization. The left panel compares the training-sample graph certificate with held-out RMSE. The right panel compares retained counterfactual derivative share with the same held-out outcome.

G.10. Original-split selection diagnostics

At $d = 3$, certificate minimization selects the baseline-value spectrum under both parameterizations and both derivative-estimation runs. That candidate is also the held-out risk minimizer in all four comparisons. Relative to the counterfactual-fiber spectrum, it lowers held-out target RMSE by 1.80–8.21 percent. A paired bootstrap over the 1,956 held-out core states gives intervals that exclude zero in every comparison. Table A2 reports the selection calculation.

TABLE A2. Training-selected graph certificate and held-out KW97 risk, $d = 3$

Model / run	$\widehat{\mathfrak{C}}_G$(CF)	$\widehat{\mathfrak{C}}_G$(selected)	RMSE(CF)	RMSE(selected)	Reduction [95% CI]
respy / 1	0.896	0.811	0.7262	0.6677	8.06 [7.37, 8.78]
respy / 2	1.003	0.795	0.7450	0.6839	8.21 [7.56, 8.93]
original / 1	0.842	0.751	0.7065	0.6856	2.95 [2.60, 3.34]
original / 2	0.845	0.825	0.6957	0.6832	1.80 [1.53, 2.09]

CF denotes the local counterfactual-fiber spectrum. The selected candidate is the baseline-value spectrum in every row. Certificates use two-fold out-of-fold training residuals. Intervals are percentile intervals from 1,999 paired core-state bootstrap draws of the held-out RMSE reduction, in percent.

Across the twelve $d = 3$ candidate–model–run observations in the original split, the Spearman correlation between the graph certificate and held-out RMSE is 0.926; the correlation between retained counterfactual derivative share and held-out RMSE is −0.112. Across all budgets in that split, certificate minimization selects the held-out best candidate in six of twelve model–run–budget comparisons, while direct out-of-fold risk selection does so in twelve of twelve.

References

Athey, Susan, and Stefan Wager. 2021. “Policy Learning With Observational Data.” *Econometrica* 89(1):133–161.

Attia, Ahmed, Alen Alexanderian, and Arvind K. Saibaba. 2018. “Goal-Oriented Optimal Design of Experiments for Large-Scale Bayesian Linear Inverse Problems.” *Inverse Problems* 34(9):095009.

Auclert, Adrien, Bence Bardoczy, Matthew Rognlie, and Ludwig Straub. 2021. “Using the Sequence-Space Jacobian to Solve and Estimate Heterogeneous-Agent Models.” *Econometrica* 89(5):2375–2408.

Bai, Jushan, and Serena Ng. 2002. “Determining the Number of Factors in Approximate Factor Models.” *Econometrica* 70(1):191–221.

Bajari, Patrick, C. Lanier Benkard, and Jonathan Levin. 2007. “Estimating Dynamic Models of Imperfect Competition.” *Econometrica* 75(5):1331–1370.

Bakry, Dominique, Ivan Gentil, and Michel Ledoux. 2014. *Analysis and Geometry of Markov Diffusion Operators*. Cham: Springer.

Beraja, Martin. 2023. “A Semistructural Methodology for Policy Counterfactuals.” *Journal of Political Economy* 131(1):190–201.

Bigoni, Daniele, Youssef Marzouk, Clémentine Prieur, and Olivier Zahm. 2022. “Nonlinear Dimension Reduction for Surrogate Modeling Using Gradient Information.” *Information and Inference: A Journal of the IMA* 11(4):1597–1639.

Borsuk, Karol. 1933. “Drei Sätze über die *n*-dimensionale euklidische Sphäre.” *Fundamenta Mathematicae* 20:177–190.

Chen, Xiaohong, and Demian Pouzo. 2015. “Sieve Wald and QLR Inferences on Semi/Nonparametric Conditional Moment Models.” *Econometrica* 83(3):1013–1079.

Chen, Xiaohong, Victor Chernozhukov, Sokbae Lee, and Whitney K. Newey. 2014. “Local Identification of Nonparametric and Semiparametric Models.” *Econometrica* 82(2):785–809.

Cheng, Bing, and Howell Tong. 2026. “An Approach to Fisher-Rao Metric for Infinite Dimensional Non-Parametric Information Geometry.” *Entropy* 28(4):374.

Chernozhukov, Victor, Iván Fernández-Val, and Blaise Melly. 2013. “Inference on Counterfactual Distributions.” *Econometrica* 81(6):2205–2268.

Christensen, Timothy, and Benjamin Connault. 2023. “Counterfactual Sensitivity and Robustness.” *Econometrica* 91(1):263–298.

Chung, Fan R. K. 1997. *Spectral Graph Theory*. Providence: American Mathematical Society.

Constantine, Paul G., Armin Eftekhari, Jeffrey Hokanson, and Rachel A. Ward. 2017. “A Near-Stationary Subspace for Ridge Approximation.” *Computer Methods in Applied Mechanics and Engineering* 326:402–421.

Constantine, Paul G., Eric Dow, and Qiqi Wang. 2014. “Active Subspace Methods in Theory and Practice: Applications to Kriging Surfaces.” *SIAM Journal on Scientific Computing* 36(4):A1500–A1524.

Elmachtoub, Adam N., and Paul Grigas. 2022. "Smart Predict, then Optimize." *Management Science* 68(1):9–26.

Fang, Zheng, and Andres Santos. 2019. "Inference on Directionally Differentiable Functions." *Review of Economic Studies* 86(1):377–412.

Fujiwara, Akio. 2024. "Hommage to Chentsov's Theorem." *Information Geometry* 7:79–98.

Givan, Robert, Thomas Dean, and Matthew Greig. 2003. "Equivalence Notions and Model Minimization in Markov Decision Processes." *Artificial Intelligence* 147(1–2):163–223.

Gonçalves, Sílvia, Ana María Herrera, Lutz Kilian, and Elena Pesavento. 2024. "State-Dependent Local Projections." *Journal of Econometrics* 244(2):105702.

Grey, Zachary. 2026. "A Riemannian View on Active Subspaces." arXiv:2607.25163.

Guerron-Quintana, Pablo A., Atsushi Inoue, and Lutz Kilian. 2017. "Impulse Response Matching Estimators for DSGE Models." *Journal of Econometrics* 196(1):144–155.

Hahn, Juergen, and Thomas F. Edgar. 2002. "An Improved Method for Nonlinear Model Reduction Using Balancing of Empirical Gramians." *Computers & Chemical Engineering* 26(10):1379–1397.

Hahn, Juergen, Thomas F. Edgar, and Wolfgang Marquardt. 2003. "Controllability and Observability Covariance Matrices for the Analysis and Order Reduction of Stable Nonlinear Systems." *Journal of Process Control* 13(2):115–127.

Haji-Ali, Abdul-Lateef, Fabio Nobile, and Raul Tempone. 2016. "Multi-Index Monte Carlo: When Sparsity Meets Sampling." *Numerische Mathematik* 132(4):767–806.

Hebden, James, and Fabian Winkler. 2026. "Computation of Policy Counterfactuals in Sequence Space." *Journal of Economic Dynamics and Control* 182:105228.

Hong, Han, and Jessie Li. 2018. "The Numerical Delta Method." *Journal of Econometrics* 206(2):379–394.

Kotta, Ülle, Claude H. Moog, and Maris Tõnso. 2018. "Minimal Realizations of Nonlinear Systems." *Automatica* 95:207–212.

Kundu, Poorbita, and Nathan Wycoff. 2026. "Active Subspaces in Infinite Dimension." *Proceedings of Machine Learning Research* 300:2089–2097.

Lall, Sanjay, Jerrold E. Marsden, and Sonja Glavaski. 2002. "A Subspace Approach to Balanced Truncation for Model Reduction of Nonlinear Control Systems." *International Journal of Robust and Nonlinear Control* 12(6):519–535.

Lee, Junhyeong, Sangjin Jin, and Yongjae Lee. 2026. "Decision-Focused Learning via Tangent-Space Projection of Prediction Error." arXiv:2605.01361.

Lewis, A. S. 1996. "Derivatives of Spectral Functions." *Mathematics of Operations Research* 21(3):576–588.

Kalouptsidi, Myrto, Paul T. Scott, and Eduardo Souza-Rodrigues. 2021. "Identification of Counterfactuals in Dynamic Discrete Choice Models." *Quantitative Economics* 12(2):351–403.

Liu, Laura, and Mikkel Plagborg-Moller. 2023. "Full-Information Estimation of Heterogeneous Agent Models Using Macro and Micro Data." *Quantitative Economics* 14(1):1–35.

Liu, Laura, Mikkel Plagborg-Moller, and Nelson Matthew P. Tan. 2026. "Limited-Information Estimation of Heterogeneous Agent Models." arXiv:2608.13953.

Magnac, Thierry, and David Thesmar. 2002. "Identifying Dynamic Discrete Decision Processes." *Econometrica* 70(2):801–816.

Moore, B. C. 1981. “Principal Component Analysis in Linear Systems: Controllability, Observability, and Model Reduction.” *IEEE Transactions on Automatic Control* 26(1):17–32.

Pistone, Giovanni, and Carlo Sempi. 1995. “An Infinite-Dimensional Geometric Structure on the Space of All the Probability Measures Equivalent to a Given One.” *Annals of Statistics* 23(5):1543–1561.

Reiter, Michael. 2009. “Solving Heterogeneous-Agent Models by Projection and Perturbation.” *Journal of Economic Dynamics and Control* 33(3):649–665.

Rust, John. 1987. “Optimal Replacement of GMC Bus Engines: An Empirical Model of Harold Zurcher.” *Econometrica* 55(5):999–1033.

Singer, Amit, and Hau-Tieng Wu. 2012. “Vector Diffusion Maps and the Connection Laplacian.” *Communications on Pure and Applied Mathematics* 65(8):1067–1144.

Sussmann, Hector J. 1976. “Existence and Uniqueness of Minimal Realizations of Nonlinear Systems.” *Mathematical Systems Theory* 10(1):263–284.

Verdière, Romain, Clémentine Prieur, and Olivier Zahm. 2025. “Diffeomorphism-Based Feature Learning Using Poincaré Inequalities on Augmented Input Space.” *Journal of Machine Learning Research* 26(139):1–31.

Winberry, Thomas. 2018. “A Method for Solving and Estimating Heterogeneous Agent Macro Models.” *Quantitative Economics* 9(3):1123–1151.

Zahm, Olivier, Paul G. Constantine, Clementine Prieur, and Youssef M. Marzouk. 2020. “Gradient-Based Dimension Reduction of Multivariate Vector-Valued Functions.” *SIAM Journal on Scientific Computing* 42(1):A534–A558.