It’s a Price! It’s a Value!

Economic Measurement Across Environments
Online Appendix

Chae-Yeon Xon
Yonsei University

September 2026

Abstract

This appendix gives the supporting representation results, global quotient conditions, inference procedures, F-MAP retention calculations, and Treasury construction. The source-fiber, local-quotient, and measurement-frontier proofs appear beside their results in the main paper. Section E also derives the fixed-rule evaluation identity.

JEL Codes: C18, C43, C51, D11, D61, E01, G12

Keywords: economic measurement, representation, comparability, transportability, information sufficiency, dimension reduction, measurement design

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

Contents

Appendix A. Foundations and Measurement Representation

A.1. Normalization and attained ranges

Suppose a group G acts on environment descriptions with known transformations ψ_g satisfying $m(x, gz) = \psi_g(m(x,z))$. A declared normalization ρ_z satisfies $\rho_{gz} \circ \psi_g = \rho_z$, making $\bar{m}(x,[z]) = \rho_z(m(x,z))$ well defined on $\mathcal{Z}/G$. The main text uses this normalized map and relabels it m. Working on a declared section of the quotient supplies the same normalized environment domain.

The scalar proposition is set theoretic on attained ranges. Continuous refinements equip I_z with its subspace topology and impose continuity of coordinate maps on the attained latent range. When interval extensions are used, the attained ranges are connected intervals. The order statement imposes strict increase on those ranges.

A.2. Scalar transport on attained ranges

For each environment let $I_z = m(\mathcal{X}, z)$ be its attained range. An object-independent transport is a map $\tau_{z\to z'} : I_z \to I_{z'}$ satisfying

$$m(x, z') = \tau_{z\to z'}(m(x,z)) \quad \text{for every } x. \tag{A1}$$

PROPOSITION A1 (Scalar transport representation). *On a nonempty object region, exact transports exist for every ordered environment pair if and only if there are a scalar state $L : \mathcal{X} \to \mathbb{R}$ and injective coordinates $\phi_z : L(\mathcal{X}) \to I_z$ such that*

$$m(x,z) = \phi_z(L(x)). \tag{A2}$$

The transports are bijections on attained ranges and obey

$$\tau_{z'\to z''} \circ \tau_{z\to z'} = \tau_{z\to z''}. \tag{A3}$$

Two such representations with strictly increasing coordinates have latent states related by a strictly increasing bijection.

To construct the representation, fix z_0 and take $L = m(\cdot, z_0)$. Exact reverse transports make each coordinate injective. The proof below establishes this attained-range representation.

COROLLARY A1 (Ranking reversal). *If two objects exchange their strict ordering between two environments, the class in* (A2) *with all coordinates strictly increasing is empty.*

Path independence here composes exact coordinate changes for a fixed object. In the index-number tradition, path independence of a Divisia integral follows from economic conditions supporting an aggregate (Gorman 1996). The source–target problem in Section 2 allows a vector profile and evaluates sufficiency separately for each declared target family.

A.3. Proof of scalar transport representation

PROOF OF PROPOSITION A1. Fix a reference environment z_0. If object-independent transports exist, define

$$L(x) = m(x, z_0), \qquad \phi_z = \tau_{z_0 \to z}. \tag{A4}$$

Then $m(x, z) = \phi_z(L(x))$ by transport correctness.

It remains to prove injectivity on $I_L = I_{z_0}$. Suppose $\phi_z(a) = \phi_z(b)$ for $a, b \in I_L$. Choose x, x' with $a = m(x, z_0)$ and $b = m(x', z_0)$. Applying the reverse transport gives

$$a = \tau_{z \to z_0}(\phi_z(a)) = \tau_{z \to z_0}(\phi_z(b)) = b.$$

Thus ϕ_z is one-to-one on the attained latent range.

Conversely, suppose $m(x, z) = \phi_z(L(x))$ with each ϕ_z injective on I_L. Since $\phi_z : I_L \to I_z$ is also onto by the definition of I_z, it is a bijection. Define

$$\tau_{z \to z'} = \phi_{z'} \circ \phi_z^{-1}. \tag{A5}$$

The inverse has domain I_z, the attained range used for transport. Equation (A1) follows, and cancellation gives path independence:

$$\tau_{z' \to z''} \circ \tau_{z \to z'} = \phi_{z''} \circ \phi_{z'}^{-1} \circ \phi_{z'} \circ \phi_z^{-1} = \tau_{z \to z''}.$$

Alternatively, if transports are taken as primitive for every ordered pair, both sides of (A3) map an attained value $m(x, z)$ to $m(x, z'')$, so path independence follows directly from correctness.

For uniqueness, let (L, ϕ_z) and $(\widetilde{L}, \widetilde{\phi}_z)$ be two representations whose coordinate families are strictly increasing, and fix z_0. On I_L, set

$$h = \widetilde{\phi}_{z_0}^{-1} \circ \phi_{z_0}.$$

Then $\widetilde{L} = h \circ L$. The map h is a strictly increasing bijection from I_L to $I_{\widetilde{L}}$. Connectedness and continuity enter extensions to intervals and continuous global-coordinate claims. □

A.4. Ranking reversal

PROOF OF COROLLARY A1. If $m(x,z) = \phi_z(L(x))$ and every ϕ_z is strictly increasing, then

$$m(x,z) > m(x',z) \iff L(x) > L(x').$$

The same latent inequality must hold at z', contradicting the reversed ordering. □

A.5. Proof of minimal representation

THEOREM A1 (Minimal profile representation). *Fix an environment neighborhood N_Z. Let $\mathcal{P}$ be a C^1 Banach manifold of real-valued functions on N_Z with C^1 evaluation map. Suppose the profile map $\mathfrak{M} : N \to \mathcal{P}$ is C^1, has constant finite rank r_M, and has an embedded r_M-dimensional image. If $\mathfrak{M} = G \circ s$ with $s : N \to \mathbb{R}^k$ and $G : s(N) \to \mathcal{P}$ extending smoothly to neighborhoods of their images, then $k \geq r_M$. Locally, a chart of the profile image yields a faithful r_M-coordinate state.*

PROOF OF THEOREM A1. On N, the factorization in the profile manifold $\mathcal{P}$ is $\mathfrak{M} = G \circ s$. The Banach-manifold chain rule gives

$$r_M = \operatorname{rank} D\mathfrak{M} = \operatorname{rank}(DG \circ Ds) \leq \operatorname{rank} Ds \leq k, \tag{A6}$$

which proves the lower bound. Faithfulness requires distinct local profile points to map to distinct states. Smooth factorization and constant rank give the dimension inequality.

Conversely, let χ be a chart from a neighborhood in the embedded profile image $\mathcal{M}$ to an open subset of $\mathbb{R}^{r_M}$ and set $s = \chi \circ \mathfrak{M}$. Define the decoder by

$$g(a,z) = \operatorname{ev}(\chi^{-1}(a), z).$$

The evaluation assumption makes g a C^1 map on the chart domain. Then $m(x,z) = g(s(x),z)$. Because χ is one-to-one, s distinguishes precisely the local measurement-equivalence classes. Hence an r_M-dimensional faithful local state exists. □

Appendix B. Interaction Quotients and Active Directions

B.1. Local quotient and functional-decomposition lineage

Section 3 gives the proof of Theorem 2 and its direct coordinate construction. This appendix records the additional global conditions and the relation to functional decomposition.

Functional-decomposition ancestry. Leontief (1947a, theorem I) characterizes scalar separation through constancy of within-group marginal-derivative ratios across the complementary variables, under his differentiability and nonvanishing-derivative assumptions. His proof constructs a summary by fixing the complementary variables and establishes functional dependence through a rank calculation. The construction above applies the common-first-integral argument to the family of marginal environment responses, producing r coordinates on each side. The cross-difference and reference field then assign these coordinates their roles in the measurement state.

Relation to functional decomposition. Leontief, Nataf, Green, and Goldman–Uzawa provide the functional-decomposition lineage; the constant-rank theorem and Frobenius provide the local geometric tools (Leontief 1947a,b; Nataf 1948; Green 1964; Goldman and Uzawa 1964). Green's Theorem 1 concerns grouping variables within one function, and his Theorem 2 concerns additive replacement of a subset under pairwise marginal-rate restrictions. Theorem 2 applies this lineage to a cardinal object–environment measurement map using a bilateral primitive: the left mixed kernel is indexed by the object and the right mixed kernel by the environment. Its conclusion matches each quotient dimension to the mixed rank, establishes a nonsingular induced mixed derivative, separates the additive baseline fields, and feeds the resulting state into corollary 2. The record frontier uses the resulting profile coordinates to evaluate source-preserving changes in target valuations.

B.2. Scope, baseline values, and cardinal units

The quotient dimensions in theorem 2 equal the mixed rank. The rotating example admits the two-coordinate state x and has mixed rank one. Its failure concerns the rank-matched relational representation. The derivative Gram matrix with a nondegenerate distribution of z has determinant $4\,\mathrm{Var}(z) > 0$ for gradient vector $(1, 2z)'$, so integrated derivative information records both active directions.

For $m(x, z) = x_1 z + x_2$ and reference $z_0 = 0$, the mixed derivative is $(1, 0)'$ and the relational coordinate is $u = x_1$. The full state is (x_2, x_1). A change in x_2 shifts every valuation by the same amount, while preserving its response to z. This example separates interaction information from the baseline field.

The kernel restriction concerns the cardinal measurement map in its stated units. Separate coordinate changes on $\mathcal{X}$ and $\mathcal{Z}$ transform the kernels by the corresponding tangent isomorphisms; adding $a(x) + b(z)$ leaves the mixed derivative unchanged. A monotone transformation of the value can alter its mixed kernels. For $\overline{m}(x, z) = \exp(x_1 z +$

x_2),

$$D_x D_z \overline{m} = e^{x_1 z + x_2} \begin{pmatrix} 1 + x_1 z \\ x_1 \end{pmatrix}. \tag{A7}$$

At $x_1 \neq 0$ this direction rotates with z. Ordinal preference ordering is preserved by the exponential transformation, while the interaction geometry of the cardinal map changes. This specifies the measurement convention under which the separability comparison is made.

B.3. Global descent and topology

DEFINITION A1 (Simple global interaction quotient). *On connected domains N_X, N_Z, let $\mathcal{F}_X, \mathcal{F}_Z$ be the foliations generated by the stable kernels. Their quotients $Q_X = N_X/\mathcal{F}_X$ and $Q_Z = N_Z/\mathcal{F}_Z$ are simple when they are Hausdorff, second-countable C^2 manifolds and their canonical projections q_X, q_Z are surjective submersions with connected fibers.*

THEOREM A2 (Global relational quotient). *Under the smoothness and stable-kernel conditions of the main theorem on connected domains, the cross-difference is constant on every product of leaves and descends uniquely to the set-theoretic leaf spaces:*

$$\Delta_0 m(x,z) = F_Q([x]_{\mathcal{F}_X}, [z]_{\mathcal{F}_Z}). \tag{A8}$$

For a simple quotient, F_Q is smooth and

$$m(x,z) = a(x) + b(z) + F_Q(q_X(x), q_Z(z)). \tag{A9}$$

Conversely, such a representation through surjective submersions with connected fibers onto r-dimensional manifolds and nonsingular induced mixed derivative has stable kernels. Its foliations are simple with leaf spaces diffeomorphic to those manifolds.

PROOF. For $\xi \in K_X(x)$, the function $z \mapsto D_x \Delta_0 m(x,z)[\xi]$ has zero derivative in every environment direction. It is zero at z_0 and hence zero globally by connectedness. Symmetrically, $D_z \Delta_0 m$ annihilates K_Z. Integration along leaves makes $\Delta_0 m$ constant on leaf products and defines F_Q uniquely. For simple foliations, local submersion charts show that the descended function is smooth.

Conversely, differentiating (A9) gives $J_M = Dq_X' D_{q_X} D_{q_Z} F_Q Dq_Z$. The induced derivative is nonsingular, so the kernels are $\ker Dq_X$ and $\ker Dq_Z$. Connected fibers are their leaves. Submersions onto Hausdorff second-countable smooth manifolds yield the claimed simple quotients. □

A simple leaf space can require several Euclidean charts. A single global vector that separates all leaves additionally requires a global chart. The local transport calculation uses the local state image and its tangent spaces.

B.4. Profile dimension and rotating kernels

PROOF OF COROLLARY 1. The relational profile factors through $u(x) \in \mathbb{R}^r$, giving rank at most r. If $Du(x)\xi \neq 0$, nonsingularity of $D_u D_v F_0$ supplies an environment direction whose mixed derivative along ξ is nonzero. Thus the relational profile rank is exactly r. The identity

$$m(x, z) = m(x, z_0) + m(x_0, z) - m(x_0, z_0) + F_0(u(x), v(z))$$

gives the full state $(m(x, z_0), u(x))$. The joint differential has rank r exactly when $D_x m(\cdot, z_0)$ annihilates $\ker Du$. Constant-rank coordinates then make $m(\cdot, z_0)$ a local function of u. At regular points with a nonzero restriction to this kernel, the joint rank is $r + 1$. □

PROOF OF PROPOSITION 1. Direct differentiation gives the displayed rank-one mixed derivative and its rotating kernel. A representation through scalar submersions and a nonzero induced mixed derivative would have left kernel $\ker Du(x)$ for every z, contradicting that rotation. The two coefficient functions $z - z_0$ and $z^2 - z_0^2$ are linearly independent on every nondegenerate interval, so the relational profile rank is two. □

B.5. Directional Gram characterization

Equip tangent spaces with smooth inner products. For a declared environment direction $\eta \in T_z\mathcal{Z}$, let $g_X(x; z, \eta)$ satisfy $\langle g_X(x; z, \eta), \xi\rangle_X = J_M(x, z)[\xi, \eta]$. For a distribution ν_B over environment directions in block B, define

$$G_X(x; B) = \mathbb{E}_{\nu_B}[g_X(x; z, \eta) g_X(x; z, \eta)']. \tag{A10}$$

PROPOSITION A2 (Directional Gram kernel). *Suppose J_M is continuous, ν_B has support equal to the declared direction set, and $\mathbb{E}_{\nu_B}\|g_X\|^2 < \infty$. Then*

$$\ker G_X(x; B) = \{\xi : J_M(x, z)[\xi, \eta] = 0 \textit{ for every } (z, \eta) \in \operatorname{supp} \nu_B\}. \tag{A11}$$

The symmetric result holds on the environment side. Full local support and stable kernels give the same exact kernel in every block.

PROOF. For every ξ, $\langle \xi, G_X \xi\rangle = \mathbb{E}_{\nu_B}[J_M(x, z)[\xi, \eta]^2]$. Nonnegativity gives zero expectation exactly when the form is zero almost everywhere. Continuity and full support extend equality to the declared support. □

For the exact bilinear map $m(x,z) = x'z$, $J_M = I$ and $G_X(B) = \mathbb{E}_B[\eta\eta']$. The exact kernel is the orthogonal complement of the span of the supported directions. Comparisons of leading eigenspaces study the stability of a selected *approximation*; equality of these leading spaces and exact equality of interaction kernels have their respective rank and support conditions. Appendix D states the approximation test and its identification gate.

B.6. Local bilinear representation

The normalization $F_0(u, v_0) = F_0(u_0, v) = 0$ gives

$$F_0(u,v) = (u - u_0)'A_0(v - v_0) + o(\|u - u_0\|\|v - v_0\|), \tag{A12}$$

where $A_0 = D_u D_v F_0(u_0, v_0)$ is nonsingular. An invertible coordinate change absorbs A_0. This is the tangent model for smooth nonlinear systems; the valuation laboratories use exact bilinear primitives.

Appendix C. Residual Information and the Measurement Frontier

C.1. Fibers, local dimension, and certificates

The proof of Theorem 1 is given in Section 2. The remaining arguments connect that result to the common-state coordinates.

For quotient reduction, the cross-difference identity gives $m(x,z) = m(x,z_0) + c(z) + F_0(u(x), v(z))$. At a regular local state image, $Ds_{z_0} : T_x\mathcal{X} \to T_s S$ is surjective. The kernel of this map annihilates every measurement derivative. It follows that Ds_{z_0} maps $\ker DY_O$ onto $\ker D\widetilde{Y}_O$, and the two restricted target images have equal rank. If $z_0 \in O$, the source map includes the coordinate α, giving $\mathrm{d}\alpha = 0$ on its kernel. This proves (18).

Proof of corollary 2. The rank of $D\widetilde{Y}_T$ restricted to $\ker D\widetilde{Y}_O$ is bounded by the dimension of that kernel. Rank–nullity on $T_s S$ gives the equality in (19), and $\dim S = r_M$. When $z_0 \in O$, $D\widetilde{Y}_O$ contains $\mathrm{d}\alpha$, so its kernel is contained in the tangent space generated by the r interaction coordinates. The final bound follows. □

C.2. Decoder scope of the measurement frontier

The complete proof of Theorem 3 is in Section 4.

The average-loss clauses concern the best affine decoder and require finite second moments. For Gaussian exposures, conditional expectation is affine and these clauses also minimize over all measurable decoders. For a general exposure law, unrestricted square-loss prediction has residual loss $\mathbb{E}[\mathrm{Var}(V_T'q \mid V_O'q)]$ in the scalar case and the corresponding conditional covariance trace in the vector case.

C.3. Correspondence with goal-oriented Gaussian design

The average-loss calculation has a Gaussian experimental-design counterpart (Attia et al. 2018). Let $q \sim N(\mu_q, \Sigma_q)$, allow Σ_q to be singular, and observe $y_\varepsilon = V_O' q + \varepsilon e$ with independent $e \sim N(0, I_K)$ and $\varepsilon > 0$. Conditional covariance is

$$\Sigma_\varepsilon = \Sigma_q - \Sigma_q V_O (V_O' \Sigma_q V_O + \varepsilon^2 I_K)^{-1} V_O' \Sigma_q. \tag{A13}$$

For the target $W_T^{1/2} V_T' q$, goal-oriented A-optimal loss is $\operatorname{tr}(M_T \Sigma_\varepsilon)$.

To take the zero-noise limit, write $\Sigma_q = LL'$ on its support and let $A = V_O' L = U_\ell D_\ell Z_\ell'$ be a compact singular-value decomposition, with positive singular values $\sigma_1, \dots, \sigma_\ell$. Then

$$\Sigma_\varepsilon = L \left[P_{\ker A} + Z_\ell \operatorname{diag}\left(\frac{\varepsilon^2}{\sigma_j^2 + \varepsilon^2} \right)_{j=1}^{\ell} Z_\ell' \right] L'. \tag{A14}$$

Consequently $\Sigma_\varepsilon \to L P_{\ker A} L' = \Sigma_{q|O}$ and $\operatorname{tr}(M_T \Sigma_\varepsilon) \to \operatorname{tr}(M_T \Sigma_{q|O})$. The formula also covers $\ell = 0$ through an empty sum. Redundant source measurements and singular exposure supports therefore have the same projection limit. The Gaussian criterion coincides with the retained-record affine-loss criterion at this limit. The frontier additionally evaluates exact recovery on a declared exposure set, smooth or linear completion counts, and available economic record fields. Each of these statements retains the assumptions specified in the main text.

C.4. General bounded sets and acquisition accounting

For any nonempty compact exposure set $\mathcal{U}$ and scalar target,

$$\inf_f \sup_{q \in \mathcal{U}} |t'q - f(V_O' q)| = \frac{1}{2} \sup_{\substack{q, q' \in \mathcal{U} \\ V_O' q = V_O' q'}} |t'(q - q')|. \tag{A15}$$

The midpoint of the smallest interval containing each residual set attains the bound. For convex $\mathcal{U} = -\mathcal{U}$, every half-difference of two same-source objects lies in $\mathcal{U} \cap \ker V_O'$, and every vector in this intersection arises from $(q, -q)$. Hence the error equals the support function $h_{\mathcal{U} \cap \ker V_O'}(t)$.

Let $r_x(A) = \operatorname{rank} DY_A(x)$. Rank increments give the accounting identities

$$d_{O \to (S,T)} = d_{O \to S} + d_{O \cup S \to T}, \tag{A16}$$

$$d_{O \to T} - d_{O \cup S \to T} = r_x(O \cup S) + r_x(O \cup T) - r_x(O) - r_x(O \cup S \cup T) \geq 0. \tag{A17}$$

The inequality is submodularity of linear row-space rank. In a bilinear system, put $\overline{\mathcal{V}}_A = (\mathcal{V}_O + \mathcal{V}_A)/\mathcal{V}_O$. The reduction equals $\dim(\overline{\mathcal{V}}_S \cap \overline{\mathcal{V}}_T)$. Thus acquisition removes the overlap between source innovation and the residual target directions.

PROOF OF PROPOSITION 2. Apply the global source-fiber criterion to (Y_O, P). Under the stated constant ranks, apply its local version to obtain $\ker D(Y_O, P) = \ker DY_O \cap \ker DP \subseteq \ker DY_T$. If $C = c \circ P$, then $\ker DP \subseteq \ker DC$; recovery from (Y_O, C) therefore implies the displayed necessary inclusion. □

For the ideal source-subspace benchmark, squared minimax loss averaged over targets is $B^2\{\mathrm{tr}(\Omega_T) - \mathrm{tr}(P_W \Omega_T)\}$. Ky Fan maximizes the removed trace over $\dim W \leq K$, proving (37). This optimization allows every subspace of the prescribed dimension; restricting measurements to environment loadings gives (35).

Appendix D. Statistical Estimation and Inference

D.1. Role in the measurement argument

The inference procedure assesses stability of a selected active-space approximation across declared environment blocks. The exact interaction-kernel statement is the Gram characterization in appendix B. A leading-space test adds a statistical approximation dimension r and an eigengap condition. Its rejection concerns persistence of that approximation across blocks. The main measurement frontier uses the source and target maps and their residual variation directly.

For centered direction vectors w_{bt} in block b, define

(A18) $$G_b = \mathbb{E}[w_{bt} w_{bt}'], \qquad H_{0,r}^{SK} : P_{1,r} = \cdots = P_{B,r} =: P_r,$$

where $P_{b,r}$ projects onto the leading r eigenvectors. The relative identification gap and rotation statistic are

(A19) $$\bar{\gamma}_{r,b} = \frac{\lambda_r(G_b) - \lambda_{r+1}(G_b)}{\lambda_r(G_b)},$$

(A20) $$T_{ROT}(r) = \sum_b \frac{n_b}{2r} \|\widehat{P}_{b,r} - \widehat{P}_{0,r}\|_F^2,$$

with $\widehat{P}_{0,r}$ the weighted Grassmann extrinsic mean. The statistical dimension r and the institutional record capacity K are separate inputs.

D.2. Approximation and the identified null

A generic local statistical approximation is

$$Y_{ie} = \alpha_i + \beta_e + u_i'v_e + \varepsilon_{ie}. \tag{A21}$$

The scanner and fixed-income laboratories use bilinear primitives in which environment directions are observed after harmonization. Block-specific eigenvalues and bases remain unrestricted; the null in (A18) imposes equality of leading subspaces.

For the relative gap in (A19), impose the population identification margin

$$\min_b \bar{\gamma}_{r,b} > \underline{\gamma}_r. \tag{A22}$$

The margin in (A22) makes the leading projector locally unique in every block. The identification gate below estimates this margin before the common-projector test receives an inferential interpretation.

D.3. Identification gate

Let $c = 1, \ldots, C$ index predeclared dependence clusters and

$$H_{bc} = \sum_{t \in \mathcal{I}_{bc}} (w_{bt} w_{bt}' - \widehat{G}_b). \tag{A23}$$

For a multiplier draw s, form the unconstrained gap draw

$$\widehat{G}^{*(s)}_{b,\mathrm{gap}} = \widehat{G}_b + \frac{1}{n_b} \sum_{c=1}^{C} \xi_c^{(s)} H_{bc}. \tag{A24}$$

For each draw define the maximum downward gap error

$$E_r^{*(s)} = \max_b \{\widehat{\bar{\gamma}}_{r,b} - \widehat{\bar{\gamma}}^{*(s)}_{r,b}\}, \tag{A25}$$

and let $q^*_{1-\alpha_g}$ be its empirical quantile. Set

$$LCB_{1-\alpha_g}(\bar{\gamma}_{r,b}) = \widehat{\bar{\gamma}}_{r,b} - q^*_{1-\alpha_g}. \tag{A26}$$

Interpretation of the test requires

$$\min_b LCB_{1-\alpha_g}(\bar{\gamma}_{r,b}) > \underline{\gamma}_r. \tag{A27}$$

The primary implementation sets $\alpha_g = 0.05$. The dimension, relative-gap floor, cluster definition, and block length are chosen on training data and fixed for evaluation.

D.4. Null-imposed cluster-multiplier bootstrap

A null draw imposes a common leading subspace while retaining each block's fitted eigenvalues. Let

$$A_b = \widehat{P}_{0,r}\widehat{P}_{b,r} + (I - \widehat{P}_{0,r})(I - \widehat{P}_{b,r}) \tag{A28}$$

and let R_b be the orthogonal polar factor of A_b. On the neighborhood in which corresponding principal angles lie below $\pi/2$, $R_b\widehat{P}_{b,r}R_b' = \widehat{P}_{0,r}$. For centered block rows $\widetilde{w}_{bt} = R_b(w_{bt} - \bar{w}_b)$, define

$$\widetilde{G}_b^0 = R_b\widehat{G}_bR_b'. \tag{A29}$$

Orthogonal similarity preserves the fitted eigenvalues and their order, while every null-center leading projector equals $\widehat{P}_{0,r}$. Define the rotated cluster scores

$$\widetilde{H}_{bc}^0 = \sum_{t\in\mathcal{I}_{bc}} (\widetilde{w}_{bt}\widetilde{w}_{bt}' - \widetilde{G}_b^0). \tag{A30}$$

For $s = 1, \ldots, S$, use shared mean-zero, unit-variance multipliers wherever a cluster affects several blocks and set

$$\widehat{G}_b^{*(s)} = \widetilde{G}_b^0 + \frac{1}{n_b}\sum_{c=1}^{C} \xi_c^{(s)}\widetilde{H}_{bc}^0. \tag{A31}$$

After symmetrization, recompute block projectors, the common projector, and $T_{ROT}^{*(s)}$. The finite-draw p-value is

$$\widehat{p}_{SK,r} = \frac{1 + \sum_{s=1}^{S} \mathbf{1}\{T_{ROT}^{*(s)}(r) \geq T_{ROT}(r)\}}{S+1}. \tag{A32}$$

With zero exceedances and $S = 4{,}999$, equation (A32) gives $\widehat{p}_{SK,r} = 1/5{,}000 = 0.0002$. The primary analysis rejects at 0.05 conditional on passing (A27). Multiple jointly declared dimensions receive an outer family-wise correction; a dimension chosen on an independent training sample is tested once.

D.5. Validity

PROPOSITION A3 (Validity and consistency of the stable-kernel test). *Let $n = \sum_b n_b$. Suppose, as $n \to \infty$, that:*

(a) the number of blocks is fixed and $n_b/n \to \kappa_b \in (0, 1)$;

(b) the joint vector of centered symmetric covariance scores satisfies a Gaussian cluster central limit theorem, has uniformly bounded $4 + \delta$ moments, and has a consistently estimated long-run covariance that is nondegenerate on the relevant projector-contrast tangent space;

(c) the dependence-preserving multiplier array satisfies the corresponding conditional joint multiplier central limit theorem in probability, the number of clusters diverges, and the largest cluster is asymptotically negligible;

(d) under the null, $\lambda_r(G_b)$ is uniformly bounded away from zero and $\min_b \bar{\gamma}_{r,b} \geq \underline{\gamma}_r + c$ for some $c > 0$;

(e) the weighted common projector is locally unique, corresponding block and common projectors have principal angles uniformly below $\pi/2$ with probability tending to one, and the common-projector and polar-rotation null maps are root-n consistent and Hadamard differentiable at the interior null.

Then the polar-rotation multiplier bootstrap consistently estimates the null law of $T_{ROT}(r)$. The identification gate passes with probability tending to one under the interior null. The level-α rule has asymptotic size at most α, with finite-S grid resolution $1/(S + 1)$. Under a fixed alternative, the test is consistent when the block leading projectors have a strictly positive minimum distance from every common rank-r projector.

PROOF. The uniform eigengap makes the leading-projector map Fréchet differentiable. The joint score central limit theorem and the spectral delta method give a joint Gaussian limit for the block projectors and their weighted common projector. At an interior null, the polar transporter converges to the identity and its differential removes the fitted block-to-common rotation from the bootstrap center. The conditional multiplier central limit theorem reproduces the null-tangent expansion. Applying the continuous quadratic map in (A20) yields bootstrap consistency. The positive gap margin makes the gate pass with probability tending to one, and the finite-draw correction gives the stated grid.

Under the stated alternative, T_{ROT} diverges at the sample-size rate because projector separation is bounded away from zero, while the polar-rotation bootstrap remains on the common-projector null scale. □

D.6. Generated states and effective deficiency

Extensions with estimated object or direction states require sample splitting or cross-fitting before second-stage loss comparisons, and inference for smooth loss functionals accounts for generated-state uncertainty. The F-MAP record experiment uses training moments, validation target directions, and a separate evaluation period; its reported paired-loss intervals condition on the trained rules. Near small eigengaps, the identification gate governs rank and principal-angle inference. The regular Gaussian approximation uses the joint score and multiplier limit assumptions as well as the gap condition in proposition A3. Those assumptions identify the population and dependence sequence for the result. Their use with calendar data calls for a corresponding block or cluster justification.

For estimated source and target direction matrices with $\widehat{\Sigma}_x \in \mathbb{R}^{p_x \times p_x}$, form

$$\widehat{R}_{O,T} = P_{\widehat{V}_O^{\perp}} \widehat{V}_T$$

and, using the source-residual covariance

$$\widehat{\Sigma}_{x|O} = \widehat{\Sigma}_x - \widehat{\Sigma}_x \widehat{V}_O (\widehat{V}_O' \widehat{\Sigma}_x \widehat{V}_O)^{\dagger} \widehat{V}_O' \widehat{\Sigma}_x$$

and fixed W_T,

$$\widehat{K}_{O,T}^{(x)} = \widehat{\Sigma}_{x|O}^{1/2} \widehat{R}_{O,T} W_T \widehat{R}_{O,T}' \widehat{\Sigma}_{x|O}^{1/2}. \tag{A33}$$

For $\epsilon \geq 0$, if $\widehat{\lambda}_j$ are its ordered eigenvalues,

$$\widehat{D}_{\epsilon, W_T}(O,T) = \min \left\{ d \in \{0, \ldots, p_x\} : \sum_{j>d} \widehat{\lambda}_j \leq \epsilon \right\}. \tag{A34}$$

At $d = p_x$ the spectral tail is the empty sum, equal to zero, so the set in (A34) is nonempty. This plug-in count is distinct from the simultaneous RMSE capacity rule in the empirical exercise. Bootstrap draws recompute estimated direction spaces and (A33) when required; W_T and ϵ remain fixed.

Appendix E. Measurement-Policy Effects

E.1. Evaluation population and policies

The primary scanner system aggregates 90 F-MAP expenditure food product groups into seven USDA consumption groups. For group j, area g, and month s, group dollars are summed across product groups, group grams are summed across the same records, and the group unit value is dollars times 100 divided by grams. Before normalization, group unit values and quantities reproduce reported expenditure. Each quantity vector is then expressed in 100-gram units and rescaled to sum to 1,000; $p'q$ is the expenditure for that normalized basket.

Training environments cover 2012–2015 and contain $N_0 = 480$ area–month observations. Validation environments cover 2016 and contain $N_v = 120$. Evaluation environments cover 2017–2018 and contain $N_e = 240$. The evaluation array combines each of the 240 basket vectors with each of the 240 target price vectors, giving 57,600 pairs.

The capacity grid is $K = 1, \ldots, 6$, with four retained records as the reference comparison. The budget is read directly from the policy configuration. Separately, the 95-percent training variance floor and 0.10 relative eigengap floor produce a training-price dimension diagnostic of four. The greedy target-directed rule begins with $O = \varnothing$ and at each step selects the feasible training price vector p with the largest validation loss reduction

$$\text{(A35)} \qquad \frac{1}{N_v} \sum_{t \in V} p_t' \left(\Sigma_{q|O} - \Sigma_{q|O \cup \{p\}} \right) p_t .$$

The dispersed rule shares the first selected environment, standardizes each price coordinate by its training root-mean-square deviation, and adds the candidate with the largest minimum squared distance from the acquired set. At each capacity the randomized rule samples K distinct training environments uniformly. The quantity oracle generates the source records and the held-out target values; each decoder is evaluated using its retained records and calibration moments.

E.2. Finite-population effect and repeated-population reference

For trained policies a and b, let $d_{it} = \ell_{it}(b) - \ell_{it}(a)$ and $\bar{d} = (N_x N_t)^{-1} \sum_{i,t} d_{it}$. Every d_{it} is computed from the two policy predictions and the accounting value $p_t' q_i$. Hence $\bar{d}$ is the finite-population acquisition effect.

The repeated-population reference conditions on the trained policies and permits dependence within object and target-environment dimensions. Write $\widetilde{d}_{it} = d_{it} - \bar{d}$, $S_i = \sum_t \widetilde{d}_{it}$, and $C_t = \sum_i \widetilde{d}_{it}$. The implemented two-way cluster variance is the positive part of

$$\text{(A36)} \qquad \widehat{V}(\bar{d}) = \frac{1}{(N_x N_t)^2} \left[\frac{N_x}{N_x - 1} \sum_i S_i^2 + \frac{N_t}{N_t - 1} \sum_t C_t^2 - \frac{N_x N_t}{N_x N_t - 1} \sum_{i,t} \widetilde{d}_{it}^2 \right].$$

The reported interval is $\bar{d} \pm 1.96 \sqrt{\widehat{V}(\bar{d})}$. For randomized acquisition with draws $r = 1, \ldots, R$, the finite-draw lower-tail probability is

$$\text{(A37)} \qquad \widehat{p}_R = \frac{1 + \sum_{r=1}^{R} \mathbb{1}\{\mathrm{MSE}(O_r) \le \mathrm{MSE}(O_a)\}}{R + 1}.$$

TABLE A1. Fixed-rule loss comparisons at four retained fields

Comparator	Comparator MSE	Greedy MSE	MSE reduction	95% effect interval
Dispersed	0.083306	0.023602	71.7%	[0.040935, 0.078474]
Uniform random	0.086360	0.023602	72.7%	[0.052767, 0.072749]

Effects are comparator loss minus implemented greedy target-directed loss over 57,600 pairs. Equation (35) defines the calibration menu optimum; appendix E reports finite-menu gaps for the implemented search rules. Intervals use two-way clustering by basket and target environment conditional on the trained policies. The random row uses Monte Carlo expected loss.

The primary calculation uses $R = 4,999$. Coordinate-system comparisons use $R = 999$.

The complete evaluation array uses the same 240 area–month identifiers in basket and target roles. A node jackknife allows each identifier to induce dependence through either role. For node g, let $\bar{d}_{(-g)}$ remove row g and column g, retaining $(N-1)^2$ cells, and let $\bar{d}_{(-\cdot)} = N^{-1}\sum_g \bar{d}_{(-g)}$. The variance reference is

$$\widehat{V}_{\text{node}}(\bar{d}) = \frac{N-1}{N}\sum_{g=1}^{N}\left(\bar{d}_{(-g)} - \bar{d}_{(-\cdot)}\right)^2. \tag{A38}$$

The forward 2017-basket by 2018-price cell places disjoint area–month observations in the two roles and uses (A36). Both references condition on the trained acquisition rules; the complete-array finite-population mean remains the directly observed estimand. A repeated-population interpretation additionally requires a sampling or weak-dependence approximation for the area–month units. Removing an identifier from both roles addresses its shared occurrence in the array; serial and cross-area dependence require their own cluster or block specification. The source-array mean is unchanged by the choice of variance reference.

Across all evaluation pairs, dispersed MSE is 0.083306 and greedy target-directed MSE is 0.023602. The effect is 0.059704 with row–column cluster standard error 0.009576 and interval [0.040935, 0.078474]. The node-jackknife standard error is 0.010558 and its interval is [0.039011, 0.080398]. Expected random-policy MSE is 0.086360. Ninety-two of 4,999 random source sets attain MSE weakly below greedy target-directed MSE, yielding $\widehat{p}_R = 0.0186$. The disjoint 2017-basket by 2018-price cell gives an effect of 0.04999 and interval [0.02708, 0.07289].

E.3. Separate row–column variance reference

The row–column intervals hold the policies fixed and use equation (A36). Because the complete array reuses area–month identifiers in both roles, the main text instead emphasizes equation (A38) and its simultaneous extension. For the first residual certificate, the row–column reference gives the loss-reduction interval $[0.01323, 0.02251]$. This interval is conditional on that variance reference.

E.4. Fixed-rule evaluation identity

DERIVATION OF (42). For the fixed affine decoder in (39), valuation error equals $p_t' B_O(q_i - \mu_q)$. Averaging its square over all $N_x N_t$ pairs and using cyclicity of trace gives

$$\bar{\ell}_E(O) = \mathrm{tr}\left[M_{\mathrm{eval}} B_O \left\{ \frac{1}{N_x} \sum_i (q_i - \mu_q)(q_i - \mu_q)' \right\} B_O' \right]. \tag{A39}$$

Write $q_i - \mu_q = (q_i - \bar{q}_E) + \delta_E$. The cross terms average to zero, and the matrix in braces is $\Sigma_E + \delta_E \delta_E'$. Substitution yields (42). All averages are finite-array moments with denominator N_x or N_t; the result is an algebraic identity for the fixed rule. □

The derivation separates a change in the averaging environment from re-estimation of the rule. It also supplies a code-level check: direct errors on the Cartesian array and the moment expansion should agree. It leaves the sampling assumptions for intervals unchanged. A population version with independently sampled object and target draws uses their product distribution; dependence or selective pairing requires the corresponding joint law.

E.5. Simultaneous bounds and policy reselection

For simultaneous evaluation, index the 240 held-out area–month nodes by g and let $\widehat{\theta}_{(-g)}$ delete node g from both the basket and target dimensions. With $G = 240$ and $\widehat{\theta}$ stacking the reported policy effects and losses, the pseudovalue $\xi_g = G\widehat{\theta} - (G-1)\widehat{\theta}_{(-g)}$ supplies a node-level linearization. One common set of 4,999 Gaussian multipliers acts on the pseudovalues for exchange-policy loss and its effect against dispersed acquisition at capacities one through five. The 95-percent critical value is the 0.95 quantile of the maximum studentized perturbation. Capacity six is reported as algebraic closure. A confidence-protected capacity is the smallest K whose simultaneous RMSE upper bound lies below a supplied basis-point threshold.

The loading and selection-aware exercises below use the public standard errors and declared cross-group covariance specifications. Their percentile ranges summarize those sensitivity calculations, while the simultaneous band conditions on the previously selected policies.

TABLE A2. Certificate-count closure across capacities and calibration samples

Calibration	Cases	Predicted d	Observed d	Maximum error
Both rules, K = 1	2	5	5	1.7×10^{-13}
Both rules, K = 2	2	4	4	1.7×10^{-13}
Both rules, K = 3	2	3	3	2.3×10^{-13}
Both rules, K = 4	2	2	2	2.3×10^{-13}
Both rules, K = 5	2	1	1	5.7×10^{-13}
Both rules, K = 6	2	0	0	7.3×10^{-12}
Leave-area-out, K = 4	10	2	2	3.4×10^{-13}

A case is a source rule and capacity, or one excluded metropolitan calibration area. Each case is checked at all three relative rank tolerances. The completing dimension is the first tested certificate count with maximum absolute error at most 10^{-9}. Main-sample cases evaluate 57,600 pairs; each leave-area-out case evaluates 24 excluded-area baskets against 240 targets.

E.6. Residual information and coordinate systems

The centered seven-group quantity covariance has numerical rank six at relative tolerance 10^{-10}. Projection onto this exposure support gives source-loading rank four and evaluation-target-loading rank six. Their conditional residual covariance has rank two, matching the support-relative identity in (31). Certificate directions diagonalize the validation target-loss operator on this residual covariance. The first direction reduces test MSE from 0.023602 to 0.005732; the second reduces it to 2.63×10^{-27} with the support-root SVD decoder.

The covariance formula in (39) is evaluated through $\Sigma_q = LL'$ and the pseudoinverse of $P_O L$. This avoids squaring the condition number of the source information matrix. The numerical closure criterion is maximum absolute held-out value error at most 10^{-9}. Relative loading-rank tolerances are 10^{-8}, 10^{-10}, and 10^{-12}.

The leave-area-out exercise excludes one metropolitan area from both the training moments and the validation target distribution, reselects four feasible source environments, and forms the certificates in that calibration sample. Evaluation uses the excluded area's later baskets. The comparison also records the loss obtained with one fewer certificate at each tolerance.

E.7. Public loading and policy-selection uncertainty

F-MAP reports the standard error s_{jeg} of each weighted EFPG unit value, calculated by ERS from 200 replicate-weight estimates (U.S. Department of Agriculture, Economic Research Service 2024). For consumption group g in environment e, let G_{jeg} be weighted grams and

TABLE A3. Policy reselection and held-out node uncertainty

Uncertainty propagated	Median reduction (%)	95% interval (%)	Pr($\Delta L > 0$)
Policy reselection from published unit-value SEs	75.7	[44.3, 89.7]	1.000
Reselection and held-out node uncertainty	75.9	[44.2, 90.1]	1.000

The first row summarizes 499 re-extractions using published marginal EFPG unit-value standard errors and draw-specific exchange and dispersed policies. The second crosses each re-extraction with ten common node-multiplier perturbations of its exchange and dispersed held-out loss matrices. Its 4,990 draws propagate policy reselection and evaluation-array uncertainty under the diagonal unit-value covariance specification. ΔL is dispersed loss minus exchange loss.

let p_{jeg} be the EFPG unit value. The aggregation used for each draw is

$$p_{eg} = \frac{\sum_{j \in g} G_{jeg} p_{jeg}}{\sum_{j \in g} G_{jeg}}, \qquad s_{eg}^2 = \frac{\sum_{j \in g} G_{jeg}^2 s_{jeg}^2}{(\sum_{j \in g} G_{jeg})^2}, \qquad p_{eg}^{(b)} = p_{eg} + s_{eg} E_{eg}^{(b)}. \tag{A40}$$

Published F-MAP fields supply marginal EFPG standard errors. Equation (A40) uses diagonal EFPG covariance and independent standard-normal draws across group–environment cells. Purchase-gram weights remain fixed. Each of 499 draws reselects the four greedy and dispersed environments from perturbed training and validation loadings. Exchange search then starts from the draw-specific greedy set and the published exchange set and retains the lower calibration loss.

The joint calculation replays every draw-specific policy and loss from the loading calculation to a maximum absolute error of 1.9×10^{-14}. The selection-and-evaluation distribution has a 75.9-percent median reduction and central range $[44.2, 90.1]$ percent. Its selected-policy MSE has median 0.02158 and central range $[0.01639, 0.02741]$. Every one of the 4,990 perturbed loss effects is positive. These percentiles cover the declared diagonal loading model and node perturbation; the covariance grid below records the separate dependence sensitivity.

The missing covariance enters a second perturbation. Define

$$\bar{s}_{eg} = \frac{\sum_{j \in g} G_{jeg} s_{jeg}}{\sum_{j \in g} G_{jeg}}. \tag{A41}$$

For $\rho \in \{0, 0.25, 0.5, 0.75, 1\}$, the sensitivity calculation draws

$$p_{eg}^{(b)}(\rho) = p_{eg} + \sqrt{1-\rho}\, s_{eg} E_{eg}^{(b)} + \sqrt{\rho}\, \bar{s}_{eg} U_e^{(b)}, \tag{A42}$$

where $U_e^{(b)}$ is common across consumption groups in environment e. This construction gives the underlying EFPG errors a common equicorrelation component while preserving

TABLE A4. Composition-index dictionary

Consumption group	Index 1 weight	Index 2 weight
Dairy	-0.103	-0.147
Fruit	1.000	-0.114
Grains	0.219	1.000
Meat and Protein Foods	0.040	-0.022
Other foods	-0.423	0.064
Prepared meals, sides, and salads	-0.040	-0.262
Vegetables	-0.694	-0.519

The weights define (43); the displayed coefficients are rounded to three decimals and the calculations use full precision. Directions use the four-record calibration residual covariance and 2016 validation targets. Each index is stored in share-index points.

their published marginal scales. The policy is reselected in each of 199 draws per grid point.

Across the five covariance specifications in (A42), the exchange median lies between 72.8 and 76.1 percent and the 2.5th percentile lies between 31.5 and 46.0 percent. The replication output contains the draw-level covariance, weighting, information-span, coordinate-resolution, and target-moment diagnostics.

E.8. Composition indices as retained records

For an original certificate loading a_j, let $\bar{a}_j$ be its mean across the seven groups, $b_j = a_j - \bar{a}_j\mathbf{1}$, and $h_j = \max_g |b_{jg}|$. Choose the sign σ_j making the largest-magnitude coefficient positive and set $w_j = \sigma_j b_j / h_j$. Since every basket satisfies $\mathbf{1}' q_i = 1,000$, the stored composition index obeys

$$C_j(q_i) = \frac{\sigma_j}{10h_j} a_j' q_i - \frac{100\sigma_j \bar{a}_j}{h_j}. \tag{A43}$$

The coefficient on $a_j' q_i$ is nonzero, so this affine recoding and its inverse preserve the information available to an affine decoder. The record is one additional scalar composition statistic for each basket.

The first index places weight one on fruit and weight −0.694 on vegetables; the second places weight one on grains and weight −0.519 on vegetables. The table supplies the contributions of all seven groups. These are retained exposure summaries for composition changes. Their common dictionary specifies the units, source-price rows, normalization, mean basket, covariance, and weights once for the recorded population.

The retained-field representation contains four source expenditure fields and two composition-index fields. The decoder combines these fields with the common dictionary

TABLE A5. Greedy and exhaustive selection on the same candidate menu

Menu	K	Menus	Mean gap (%)	Max. gap (%)	Greedy MSE	Opt. rule MSE
Full	1	1	0.00	0.00	0.462314	0.462314
Full	2	1	8.65	8.65	0.175843	0.157026
Reduced	1	12	0.00	0.00	0.479202	0.479202
Reduced	2	12	11.74	40.49	0.194527	0.176382
Reduced	3	12	31.11	69.40	0.084878	0.073720
Reduced	4	12	41.17	101.11	0.035249	0.020527

Gaps use the calibration objective in (A44). Reduced-menu entries average across 12 menus; maximum gaps describe the same menus. The last two columns give held-out MSE, averaged across menus, for the greedy and calibration-optimal rules. The full menu has 480 candidates, and each reduced menu has 20. Every held-out comparison uses 57,600 basket–target pairs.

and the requested target price vector. The first index removes 75.7132 percent of source-only MSE; both indices reduce the maximum absolute held-out value error below 4×10^{-13}. Quantity vectors serve as the evaluation oracle.

E.9. Exhaustive menu comparison

Write $L_{\mathrm{cal}}(O) = \mathrm{tr}(M_{\mathrm{validation}} \Sigma_{q|O})$, using training basket covariance and the equal-weight 2016 target-price second moment. For each supplied menu $\mathcal{Z}$, the exhaustive rule chooses $O_K^\star \in \arg\min_{O \subseteq \mathcal{Z},\ |O|=K} L_{\mathrm{cal}}(O)$. Its comparison with the greedy set O_K^g is

$$G_K = 100 \left\{ \frac{L_{\mathrm{cal}}(O_K^g)}{L_{\mathrm{cal}}(O_K^\star)} - 1 \right\}. \tag{A44}$$

Held-out loss is evaluated after each set has been selected under this calibration criterion.

The full menu has 480 training environments. The comparison evaluates all 480 singleton sets and all $\binom{480}{2} = 114,960$ pairs. Twelve reduced menus each contain 20 candidates, with two training environments drawn from each metropolitan area. Enumerating capacities $K = 1, 2, 3, 4$ adds 74,340 subset evaluations, for 189,780 objectives in total.

For the full-menu two-record problem, calibration losses are 0.122619 for greedy and 0.112860 at the exhaustive optimum. At four records across reduced menus, mean held-out loss is 0.035249 for greedy and 0.020527 for the calibration-optimal rule. Mean and maximum calibration gaps are 41.17 and 101.11 percent. The full 480-candidate menu is enumerated through capacity two.

The calculation factors $\Sigma_q = LL'$ on its six-dimensional support and computes each source row space from $P_O L$ by singular-value decomposition. Its residual projector acts on the validation target moment to produce $L_{\mathrm{cal}}(O)$. The comparison retains the primary greedy rule and calendar split.

TABLE A6. Full-menu one-exchange comparison

K	Exchanges	Greedy cal.	Exchange cal.	Greedy test	Exchange test	vs. dispersed (%)
1	0	0.545895	0.545895	0.462314	0.462314	0.0
2	2	0.122619	0.112860	0.175843	0.157026	62.7
3	4	0.047130	0.041666	0.067755	0.067946	80.0
4	6	0.016848	0.014863	0.023602	0.018594	77.7
5	4	0.004022	0.003015	0.005737	0.004691	80.2

Calibration columns use training basket moments and 2016 target prices. Test columns use the 57,600 pair evaluation array. The final column reports the exchange rule's MSE reduction relative to dispersed acquisition. Capacity six reaches numerical closure and is omitted from the displayed loss comparison.

E.10. One-exchange search and forward evaluation

Starting from the greedy set, the exchange routine evaluates every subset obtained by replacing one selected environment with one unselected environment, accepts the largest strict calibration improvement, and repeats until every one-exchange neighbor has weakly higher calibration loss. The rule uses the same training exposure moments, 2016 validation targets, and 480-environment menu as the acquisition criterion. This is a Fedorov-style local exchange routine; Lau and Zhou (2022) analyzes local search for experimental-design criteria under explicit conditioning assumptions.

At two records, exchange search reaches the enumerated full-menu optimum. In the 12 reduced menus, its mean gaps from the enumerated optimum are 0.00, 0.19, and 3.08 percent at capacities two, three, and four; the corresponding greedy gaps are 11.74, 31.11, and 41.17 percent. The maximum four-record exchange gap is 15.14 percent, compared with 101.11 percent for greedy. On the full menu at four records, six exchanges select training rows 382, 402, 418, and 468 and reduce complete-array MSE to 0.018594.

For the disjoint calendar evaluation, 2017 supplies 120 baskets and 2018 supplies 120 target price vectors. The four-record exchange rule has MSE 0.025115, compared with 0.083410 for dispersed acquisition, a 69.9-percent reduction. Its paired effect is 0.058295, with two-way cluster standard error 0.011378 and interval $[0.035993, 0.080596]$. Selection remains anchored to 2012–2016 calibration data.

E.11. Capacity units and marginal loss

The zero-record predictor is $p_t'\mu_q$. Its held-out MSE is 40.889540, and the greedy target-directed capacities one through six yield the sequence in main table 4. The benchmark expenditure $\bar{v}_{\text{cal}} = \bar{p}'_{\text{validation}}\mu_q$ is 330.145092 dollars for the normalized 100-kilogram basket. Normalized RMSE is $10,000\sqrt{L_{\text{eval}}(K)}/\bar{v}_{\text{cal}}$ basis points. Marginal saved loss is measured per additional scalar field and per evaluated basket–target pair.

For the greedy sequence $O_K = O_{K-1} \cup \{a_K\}$, the calibration increments satisfy

$$L_{\mathrm{cal}}(O_{K-1}) - L_{\mathrm{cal}}(O_K) = VM(a_K \mid O_{K-1}, T_{\mathrm{validation}}). \tag{A45}$$

The identity holds at every step, and held-out increments telescope from capacity zero to capacity six. Calibration and held-out gains are reported separately. The full capacity schedule supplies each candidate value of a cost-penalized objective $L(K) + \lambda K$ under an institution-supplied per-field penalty λ. At four fields, exchange and dispersed RMSE are 0.136360 and 0.288628 dollars. Their RMSE difference is 0.152268 dollars, or 4.612 basis points of benchmark expenditure. Their MSE difference is 0.064712 squared dollars per evaluated pair. With N future revaluations and conversion factor κ from squared-dollar loss to the institution's cost unit, exchange is selected whenever its incremental search and implementation cost is below $N\kappa(0.064712)$. The four-field storage charge is common to the two policies.

Appendix F. Treasury Construction and Tail Performance

The fixed-income calculation combines Treasury auction records with Federal Reserve zero-coupon yield curves. Fixed-rate nominal notes and bonds are converted into dated promised cash-flow vectors, and the yield curves are converted into discount vectors on a semiannual maturity grid. The environment split contains 216 training month ends, 96 validation month ends, and 120 test month ends. Cash-flow objects comprise 521 training securities and 581 test securities. Duration, convexity, key-rate durations, and residual-aligned cash-flow projections are evaluated as additional retained exposure fields.

The published Svensson curves form the discount representation. Source selection and rank calculations use training and validation inputs; the displayed revaluation losses use the held-out cash-flow and curve sample.

TABLE A7. Treasury certificate benchmark

Certificate	Dimension	RMSE	Bias	Tail RMSE	Relative loss
Source prices only	0	14.563	7.303	35.426	1.000
Duration	1	0.789	0.050	1.841	0.002932
Duration and convexity	2	0.631	0.038	1.667	0.001878
Key-rate durations	5	0.212	−0.077	0.481	0.000212
Validation-optimal residual	1	2.234	0.464	6.286	0.023538
Full cash-flow oracle	full	0	0	0	0

Source-only RMSE, bias, and tail RMSE are 14.563, 7.303, and 35.426. Tail RMSE uses the top decile of Macaulay duration under the source discount vector. The validation-optimal residual certificate uses 1990–2007 cash-flow moments and 2008–2015 discount curves. Calibration MSE is 1.878 for this field and 2.468 for duration. The optimum is for that calibration criterion over one added linear field. The table uses 2016–2025 cash flows and curves, under which duration attains the lower RMSE.

References

Ahmed Attia, Alen Alexanderian, and Arvind K. Saibaba. Goal-oriented optimal design of experiments for large-scale bayesian linear inverse problems. *Inverse Problems*, 34(9):095009, 2018. doi: 10.1088/1361-6420/aad210.

Steven M. Goldman and Hirofumi Uzawa. A note on separability in demand analysis. *Econometrica*, 32(3):387–398, 1964. doi: 10.2307/1913043.

W. M. Gorman. Notes on Divisia indices. In Charles Blackorby and Anthony F. Shorrocks, editors, *Separability and Aggregation: The Collected Works of W. M. Gorman, Volume I*, pages 49–60. Oxford University Press, 1996. doi: 10.1093/0198285213.003.0004. Originally circulated as a 1970 typescript.

H. A. John Green. *Aggregation in Economic Analysis: An Introductory Survey*. Princeton University Press, Princeton, NJ, 1964.

Lap Chi Lau and Hong Zhou. A local search framework for experimental design. *SIAM Journal on Computing*, 51(4):900–951, 2022. doi: 10.1137/20M1386542.

Wassily Leontief. A note on the interrelation of subsets of independent variables of a continuous function with continuous first derivatives. *Bulletin of the American Mathematical Society*, 53(4): 343–350, 1947a. doi: 10.1090/S0002-9904-1947-08796-6.

Wassily Leontief. Introduction to a theory of the internal structure of functional relationships. *Econometrica*, 15(4):361–373, 1947b. doi: 10.2307/1905335.

André Nataf. Sur la possibilité de construction de certains macromodèles. *Econometrica*, 16(3): 232–244, 1948. doi: 10.2307/1907277.

U.S. Department of Agriculture, Economic Research Service. Food-at-home monthly area prices, 2012–2018, 2024. URL `https://www.ers.usda.gov/data-products/food-at-home-monthly-area-prices/documentation`. Data product; documentation updated April 2026 and accessed August 2026.