Between Models

Counterfactual Frontiers across Causal-Response Experiments
Online Appendix

Chae-Yeon Xon
Yonsei University

September 2026

Abstract

This appendix contains proofs, inference details, simulation design, and additional calculations for the monetary application. The notation and target definitions follow the main text.

JEL Codes: C12, C13, C32, C52, E52

Keywords: structural counterfactuals, partial identification, factorization, weak identification, generated responses, monetary policy

---

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

Contents

Appendix A. Model identification and factorization

A.1. Conventions

All response codomains are finite-dimensional Euclidean spaces with their usual topology. The model envelope is a disjoint union, so a model instance retains its family label. For a map $S : X \to Y$, its fiber at s is $S^{-1}(s)$. A factor map has domain $S(X)$. Each family keeps its own parameterization and manifold chart.

LEMMA A1 (Closure of a singleton image). *If a set* $A \subseteq \mathbb{R}^q$ *is empty, then* $\operatorname{cl} A$ *is empty. If every element of a nonempty* A *equals* a*, then* $\operatorname{cl} A = \{a\}$.

PROOF. The empty set is closed. A singleton in a Euclidean space is closed, and a nonempty set whose elements all equal a is the singleton $\{a\}$. □

A.2. Proof of identification as factorization

The full proof appears in the in-paper appendix.

A.3. Differential factorization

The pointwise and local factorization proofs appear in the in-paper appendix.

A.4. Cross-family agreement

PROOF OF THEOREM 3. Suppose the overlap condition holds. If two models $M_f(\theta)$ and $M_g(\vartheta)$ have the same training response s, then

$$T_{p,f}(\theta) = \lambda_{f,p}(s) = \lambda_{g,p}(s) = T_{p,g}(\vartheta).$$

The same argument covers $f = g$. Hence T_p is constant on every envelope fiber and Theorem 1 applies.

For the reverse implication, if the envelope identifies the target, take any s in a pairwise overlap and choose preimages θ and ϑ. Fiber constancy gives $\lambda_{f,p}(s) = T_{p,f}(\theta) = T_{p,g}(\vartheta) = \lambda_{g,p}(s)$. □

A.5. The three status conditions

PROOF OF PROPOSITION 1. Write $F = \mathfrak{F}_A(s_A^\dagger; \mathfrak{E})$. Coverage failure is $F = \varnothing$, whereas conditional policy-identification failure includes $F \neq \varnothing$, so the two events are mutually exclusive. If $F = \varnothing$, its held-out image is empty and held-out relation failure follows. For a nonempty example with policy-identification and held-out relation failure, take $F = \{M_1\}$ and $s_B^\dagger \neq S_B(M_1)$. For a nonempty example with policy-identification failure and held-out relation failure, take $F = \{M_1, M_2\}$ with $T_p(M_1) \neq T_p(M_2)$ and choose $s_B^\dagger$ outside $\{S_B(M_1), S_B(M_2)\}$. These constructions establish the stated compatibility relations. □

Appendix B. Testing-risk fibers, certification, and approximation

B.1. Local response metric

Proof of Proposition 3. Continuous differentiability gives $S_{A,f}(\theta_f^0 + h/\sqrt{T}) - S_{A,f}(\theta_f^0) = T^{-1/2}D_f h + o(T^{-1/2})$ uniformly for h in a compact set. Substitution of $\Sigma_{A,T} = \Omega_{A,f}/T$ gives (21). For $\theta = \varphi(\eta)$, the chain rule gives $D_{f,\eta} = D_{f,\theta}D\varphi(\eta_0)$ and the congruence in (22). Finally, $h'H_f h = (D_f h)'\Omega_{A,f}^{-1}(D_f h)$, so positive definiteness gives $\ker H_f = \ker D_f$. □

B.2. Testing affinity

For a deterministic test with rejection region A, the sum of conditional errors is $P_0(A) + P_1(A^c) = 1 - \{P_1(A) - P_0(A)\}$. Taking the infimum over randomized tests gives $1 - \sup_A |P_1(A) - P_0(A)|$, which proves the identity in (23).

Proof of Proposition 4. Let the Gaussian means be μ_0 and μ_1, whiten their difference, and write $\delta = \Sigma^{-1/2}(\mu_1 - \mu_0)$ with $\|\delta\| = c$. Projection on $e = \delta/c$ produces unit variance and a mean difference c. The likelihood-ratio threshold is the midpoint, so each conditional error equals $\Phi(-c/2)$. Their sum is $2\Phi(-c/2)$, and division by two gives the Hansen–Sargent detection-error probability. □

B.3. Coverage and sign separation

The complete proof appears immediately after the result in the main text (5).

Proof of Corollary 1. For $\rho \in [0,1]$, every weight ρ^h is nonnegative and the weight at $h = 0$ equals one. A finite weighted sum of negative coordinates is negative. □

Proof of Proposition 6. For each ξ, Proposition 5 gives the common-negative-sign interval $[c_{\text{entry}}(\xi), c_{\text{sign}}(\xi))$. This interval is nonempty exactly when $\Delta(\xi) > 0$. If $\underline{\Delta}(\Xi) > 0$, then $\Delta(\xi) > 0$ for every $\xi \in \Xi$. If some ξ_0 has $\Delta(\xi_0) \le 0$, its interval is empty, so the collection criterion fails. □

B.4. Continuous certificates

Proof of Proposition 7. Write $x_i = \Sigma_{A,T}^{-1/2}\{S_{A,f}(\theta_{f,i}) - s_A^\dagger\}$. Convexity of the norm gives $\|\sum_i \lambda_i x_i\| \le \sum_i \lambda_i \|x_i\| \le c$. Each scalar policy coordinate is affine in the same simplex weights, so its extrema over the simplex occur at vertices. Equation (35) and the margin conclusion follow. □

Proof of Proposition 8. For any $\theta \in \mathcal{N}_f(s_A^\dagger; c)$, choose $\theta_g \in \mathcal{G}_f(c)$ with $\|\theta - \theta_g\| \le \varepsilon_f$. The Lipschitz and numerical bounds give $T_{j,f}(\theta) \le \widehat{T}_{j,f}(\theta_g) + L_{f,j}\varepsilon_f + e_{f,j}$. Taking the supremum over θ and the maximum over families gives (36). □

A boxwise interval implementation can replace the covering calculation. For each parameter box B, it verifies either a distance lower bound above c or a target upper bound below zero. The calculation records closed domains, equilibrium checks, parameter-box widths, response truncation bounds, and solver allowances.

B.5. Envelope monotonicity

The complete proof is given in the main text immediately after the result.

Proof of Proposition 9. Envelope inclusion gives $\mathfrak{F}_A(s;\mathfrak{E}_1) \subseteq \mathfrak{F}_A(s;\mathfrak{E}_2)$. Applying T_p preserves inclusion of raw policy images, and closure is monotone, which yields (37). □

B.6. Approximate factorization

Proof of Lemma 1. Take M and M' in the raw set defining (13). Adding and subtracting the fitted relation values gives

$$\begin{aligned}\|T_p(M) - T_p(M')\|_{W_p} \le \|T_p(M) - \lambda_p\{S_A(M)\}\|_{W_p} + \|\lambda_p\{S_A(M)\} - \lambda_p\{S_A(M')\}\|_{W_p} \\ + \|\lambda_p\{S_A(M')\} - T_p(M')\|_{W_p}.\end{aligned}$$

The first and third terms are at most ϵ_C. The middle term is at most $2L_p\epsilon_A$ by the triangle inequality in response space. Every pair in the raw image is therefore separated by at most $2(\epsilon_C + L_p\epsilon_A)$, and continuity of finite-dimensional norms carries the bound to the closure. □

B.7. Visibility strength

Suppose $C_{f,p} \neq 0$. The affine set $\{h : C_{f,p}h = 1\}$ is nonempty. Under $\ker D_f \subseteq \ker C_{f,p}$, Proposition 2 gives $C_{f,p} = \Lambda D_f$. A sequence in this affine set with $D_f h_n \to 0$ would imply $1 = \Lambda D_f h_n \to 0$. Finite-dimensional norm equivalence and positive definiteness of Σ_A give $\kappa_{f,p} > 0$. Under kernel failure, scaling a direction in $\ker D_f$ with a positive policy loading gives objective value zero in (15).

Appendix C. Generated responses and weak-rank inference

C.1. Minimum-distance expansion

Let the minimum-distance criterion for family f be $L_{f,T}(\theta) = \hat{g}_{f,T}(\theta)'\widehat{W}_f\hat{g}_{f,T}(\theta)$. The approximate first-order condition and stochastic differentiability in Assumption A1 imply

$$0 = G_f' W_f \left\{\hat{g}_{f,T}(\theta_f) + G_f(\hat{\theta}_f - \theta_f)\right\} + o_p(T^{-1/2})$$

uniformly on the maintained class. Multiplying by $\sqrt{T}$ and using nonsingularity of $G_f' W_f G_f$ yields

$$\sqrt{T}(\hat{\theta}_f - \theta_f) = -(G_f' W_f G_f)^{-1} G_f' W_f \sqrt{T}\hat{g}_{f,T}(\theta_f) + o_p(1).$$

Substitution of the moment influence representation proves (A6). Uniform consistency keeps the expansion in the interior neighborhood on which the derivative and eigenvalue bounds hold.

C.2. Implicit equilibrium and response derivatives

At a selected equilibrium, total differentiation of $F_f\{z_f(\theta, q), \theta, q\} = 0$ gives

$$D_\theta z_f = -F_{f,z}^{-1} F_{f,\theta}, \qquad D_q z_f = -F_{f,z}^{-1} F_{f,q}.$$

Applying the chain rule to $R_f(\theta, q) = \Phi_f\{z_f(\theta, q), \theta, q\}$ gives (A8). A joint delta method applied to $(\hat{\theta}_f, \hat{q})$ then produces the first two terms in (A7). If a simulated response is an average over S_T centered draws, its stochastic order is $S_T^{-1/2}$. After multiplication by $\sqrt{T}$, its limit scale is $\sqrt{T/S_T} \to \sqrt{\eta}$. Solver and discretization errors vanish from the first-order law by Assumption A1.

When two families use a common estimated treatment path or moments from overlapping samples, the associated influence vectors share components. Stacking before covariance estimation retains terms such as

$$R_{f,q}\mathbb{E}(\psi_q \psi_q') R_{g,q}' \quad \text{and} \quad R_{f,\theta}\mathbb{E}(\psi_{\theta_f} \psi_q') R_{g,q}'.$$

These terms enter the joint covariance through shared inputs and overlapping samples.

Proof of Proposition A1. The minimum-distance expansion proves (A6). Invertibility of $F_{f,z}$ gives the implicit derivatives. The joint functional delta method, centered simulation representation, and rate conditions then give (A7). Stacking the primitive influence vectors and applying fixed linear response derivatives gives the joint family representation and its off-diagonal covariance blocks. $\square$

C.3. Uniform projection inference

The complete proof appears immediately after the result in the main text (4).

The domination uses the value of g_S at a compatible model nuisance. Its derivative rank may change with P or T. Compactness and measurability make the profiled statistic measurable and support numerical approximation of the infimum. The optimizer returns an upper bound from evaluated feasible points and a convergence certificate; ϵ_{num} includes the remaining optimization gap.

C.4. Regular quotient refinement

Proof of Proposition A2. Under constant rank, local coordinates (x, z) can be chosen so that the population moment factors through x. Assumption A2 gives the same factorization for the first-order covariance. After whitening at the truth and writing $u = \sqrt{T}(x - x_0)$, the uniform local quadratic expansion is

$$Q_{S,T}(x,z,b_0) = \|Z_T - Ju\|^2 + o_p(1), \qquad Z_T \Rightarrow N(0, I_{d_S}),$$

where J has column rank r_S. Consistency of the local minimizer permits profiling on an $O_p(1)$ neighborhood. Therefore

$$P^{\mathrm{reg}}_{S,T}(b_0) = \inf_u \|Z_T - Ju\|^2 + o_p(1) = \|(I - \Pi_J)Z_T\|^2 + o_p(1),$$

where $\Pi_J = J(J'J)^{-1}J'$ is the orthogonal projection onto $\operatorname{Im} J$. The matrix $I - \Pi_J$ is idempotent with rank $d_S - r_S$. The continuous mapping theorem gives the limit in (A9). □

An active inequality boundary gives a local cone and a chi-bar-square or another boundary law. Covariance variation along z carries the fiber coordinate into the first-order experiment. The quotient refinement applies to interior cells with covariance factorization and locally constant rank.

C.5. Target-constrained quotient

Proof of Proposition A3. Assumption A3 supplies local coordinates u on the target level set and a uniform expansion

$$Q_{A,f,T}(u,b_0) = \|Z_T - J_{A|b,f}u\|^2 + o_p(1), \qquad Z_T \Rightarrow N(0, I_{d_A}).$$

Consistency permits profiling on this chart. The profiled criterion equals $\|(I - \Pi_{J_{A|b,f}})Z_T\|^2 + o_p(1)$. The projection matrix has rank $r_{f,b}$, which gives (A12). A weak direction normal to the target level set is absent from $B_{f,b}$ and therefore absent from the profiled tangent score. □

Appendix D. Application details and implementation

D.1. Data construction and units

The baseline quarterly sample is 1969Q1–2006Q4. The aggregate series are real output per capita (FRED `A939RX0Q048SBEA`), the GDP deflator (FRED `GDPDEF`), and the federal funds rate (FRED `FEDFUNDS`). Monetary shocks use the Romer and Romer (2004) series distributed with the Ramey data and the natural-language shock of Aruoba and Drechsel (2024). Monthly shocks are aggregated to quarters. The VAR uses two lags, an intercept, a linear trend, and the ordering AD shock, output, inflation, RR shock, and policy rate. Output and inflation enter horizons zero through twelve, and treatment paths are normalized to a 25-basis-point annualized policy-rate peak over horizons zero through twenty-four.

Under this design, $s^{\dagger}_A$ has 52 coordinates, the joint response and policy-rate input has 78 coordinates, and each delayed policy response $\nu_f(h;\theta)$ has 26 coordinates. The local Jacobian $D_{A,f}$

enters tangent calculations, while finite response secants enter the evaluated-design calculations. Inflation and the policy rate are divided by four at the model interface, and persistent interventions use the coefficients ρ^h applied to the delayed response vectors.

D.2. Nonlinear envelope computation

For each family and detection radius, the computation evaluates the equilibrium and the training and policy responses over the declared parameter domain, searches for extrema of training distance and policy targets from multiple starting points, and repeats the calculation under tighter solver tolerances and longer horizons. Family-specific contributions are then combined into the policy image of the compatible set. The numerical comparisons below use the retained baseline support for finite-support frontiers. Native-domain certification requires a separate reproducible enclosure calculation.

At the benchmark radius 3.9199, the retained-design family counts are 613, 561, 664, and 619. Across the eight testing-risk radii in equation (47), the combined counts are $0, 12, 273, 1204$ and $2457, 3399, 3840, 3988$, from smallest to largest radius. The first retained point with a nonnegative delayed or persistent target has distance 4.356667. A local HANK–RE search gives an inflation reversal witness at distance 4.253202.

D.3. Evidence subsets and the CMR expansion

The evidence exercise evaluates 26 response specifications on the same 4,000 structural points. The full-evidence specification reproduces the baseline distances, and the nested horizon and block-deletion results are reported in Tables A1 and A2.

TABLE A1. Training-horizon frontier on fixed support

Evidence	Retained	$c^{\mathcal{D}}_{\text{entry}}$	$c^{\mathcal{D}}_{\text{sign}}$	$Y < 0$	$\pi < 0$
Joint horizons 0–0	4,000	0.5962	0.5962	17/17	5/17
Joint horizons 0–1	4,000	1.6553	1.6505	17/17	5/17
Joint horizons 0–2	3,996	2.4831	2.4831	17/17	5/17
Joint horizons 0–3	3,884	3.1750	3.5993	17/17	5/17
Joint horizons 0–4	3,724	3.2940	3.6160	17/17	5/17
Joint horizons 0–5	3,653	3.3131	3.7090	17/17	5/17
Joint horizons 0–6	3,497	3.3459	3.8719	17/17	5/17
Joint horizons 0–7	3,396	3.3529	3.9460	17/17	17/17
Joint horizons 0–8	3,268	3.3646	4.0214	17/17	17/17
Joint horizons 0–9	3,078	3.3845	4.1165	17/17	17/17
Joint horizons 0–10	2,841	3.4391	4.2201	17/17	17/17
Joint horizons 0–11	2,588	3.4951	4.3080	17/17	17/17
Joint horizons 0–12	2,457	3.5058	4.3567	17/17	17/17

The support contains the same 4,000 baseline evaluations in every row. Retained counts and negative-cell counts use $c = 3.9199$, $\beta = 0.99$, and delays zero through sixteen. Each distance uses the principal covariance submatrix of its evidence block. $c^{\mathcal{D}}_{\text{sign}}$ is the first evaluated reversal distance; $c^{\mathcal{D}}_{\text{entry}}$ is the largest family-specific minimum distance. The parameter support, fitted shock inputs, and policy targets remain fixed.

TABLE A2. Shock–outcome blocks and deletion experiments

Evidence	Retained	$c^{\mathcal{D}}_{\text{entry}}$	$c^{\mathcal{D}}_{\text{sign}}$	$Y < 0$	$\pi < 0$
RR output	4,000	2.3584	2.3981	17/17	5/17
RR inflation	4,000	2.0878	3.3564	17/17	5/17
AD output	4,000	0.7816	0.7580	17/17	5/17
AD inflation	4,000	1.1097	1.4569	17/17	5/17
All except RR output	3,999	2.6609	3.5921	17/17	5/17
All except RR inflation	3,987	2.8592	3.0165	17/17	5/17
All except AD output	3,354	3.3685	4.2558	17/17	17/17
All except AD inflation	3,428	3.2548	4.2340	17/17	17/17

The support contains the same 4,000 baseline evaluations in every row. Retained counts and negative-cell counts use $c = 3.9199$, $\beta = 0.99$, and delays zero through sixteen. Each distance uses the principal covariance submatrix of its evidence block. $c^{\mathcal{D}}_{\text{sign}}$ is the first evaluated reversal distance; $c^{\mathcal{D}}_{\text{entry}}$ is the largest family-specific minimum distance. The parameter support, fitted shock inputs, and policy targets remain fixed.

The CMR calculation preserves the authors' economic equations and calibration while adding 24 dated news states for the common announced-wedge experiment. The original and augmented contemporaneous responses agree within 6.54×10^{-12}. The empirical fit has rank 50 and condition number 9.927. The paper reports this calibration as one evaluated extension point.

D.4. Native-domain certification gate

The convex response completion uses common simplex weights for each evaluated empirical response and its policy target. At the benchmark, its 2,457 vertices have maximum distance 3.919872 and joint sign margin 0.002193.

The local-box calculation covers the evaluated support and its stated convex completion. A native-domain compatibility or sign conclusion requires a single calculation with a fixed response-column order, shock path, box definition, distance enclosure, and target enclosure.

Horizon checks solve the margin-setting vertices at horizons 200, 225, 250, 275, and 300. At the binding RANK–CD inflation horizon-sixteen cell, the $T = 200$–$T = 300$ difference is 7.843×10^{-7}, or 0.03577 percent of the 0.0021925988 margin. Across the 2,457 compatible vertices, the maximum $T = 200$–$T = 300$ response discrepancy is 2.682×10^{-5} and the maximum relative equilibrium residual is 2.235×10^{-16}.

TABLE A3. Target-unit checks at the binding sign-margin cell

Quantity	Target units	Margin share	Scope
Artifact–$T = 300$ difference	2.987×10^{-10}	0.0000136%	Binding vertex
$T = 200$–$T = 300$ difference	7.843×10^{-7}	0.03577%	Fixed-grid vertex
$T = 275$–$T = 300$ difference	3.728×10^{-8}	0.001700%	Fixed-grid vertex
Local policy secant discrepancy	0.001341	61.16%	Observed vertex secant
Factor-relation secant discrepancy	0.001418	64.69%	Observed vertex secant
All-vertex artifact–$T = 300$ maximum	2.375×10^{-8}	0.001083%	83,538 comparisons
All-vertex $T = 200$–$T = 300$ maximum	2.682×10^{-5}	1.223%	83,538 comparisons
All-vertex local-policy secant maximum	0.3424	15,616%	Baseline design
All-vertex factor-relation secant maximum	0.1658	7,561%	Baseline design
Conditional standard error	0.015442	704.28%	Local influence map
Evaluated 95% projection displacement	0.002716	123.89%	Union upper endpoint

The reference margin is 0.0021925988 at RANK–CD inflation horizon sixteen. Fixed-grid rows use independent equilibrium solves. All-vertex rows cover the 2,457 benchmark-compatible evaluations and 34 targets. Secant maxima measure approximation error on the baseline design. The conditional standard error uses the empirical-response component of the generated-response influence stack.

Four response-center and covariance specifications provide a sensitivity comparison. The baseline admits 2,457 evaluations with margin 0.002193; the current-vintage center admits 2,417 with the same margin; the model-response training covariance admits 3,215 with margin 0.001470;

and the rank-50 pseudoinverse yields all-family entry radius 4.2862. The first evaluated reversal radii are 4.3567, 4.3471, 4.2391, and 5.0199 in the same order.

TABLE A4. Baseline-design response-center and covariance sensitivity

Specification	Rank	Count	Margin	Coverage	Reversal
Baseline, tapered	52	2,457	0.002193	3.5058	4.3567
Current-vintage center, tapered	52	2,417	0.002193	3.5125	4.3471
Baseline, raw pseudoinverse	50	0	n/a	4.2862	5.0199
Model-artifact training-only	52	3,215	0.001470	3.3916	4.2391

Compatible counts use radius 3.9199. Family-entry radius is the largest family minimum over the baseline design. Reversal radius is the first evaluated delayed-target sign reversal. Each populated row and its joint response–policy convex completion retains a negative upper endpoint. The raw rank-50 covariance gives an empty benchmark set.

D.5. Weak-rank simulation design

The local model dimension is six and sample size is $T \in \{150, 300, 600\}$. Let D_T and C_T be the empirical-response and scalar-target Jacobians. The smallest empirical singular value follows

$$\sigma_{\min}(D_T) = \frac{c}{\sqrt{T}}, \qquad c \in \{0, 0.25, 0.5, 1, 2, 5\}, \tag{A1}$$

and the target loading on the associated right singular vector n_T is

$$C_T n_T = \gamma, \qquad \gamma \in \{0, 0.25, 0.5, 1\}. \tag{A2}$$

Generated-response uncertainty varies the correlation between the response and Jacobian influence projections over $\{0, 0.5, 0.9\}$ and their standard-deviation ratio over $\{0, 0.5, 1\}$. Each cell uses 5,000 replications.

Among the cells with $c \leq 1$ and $\gamma > 0$, projection has mean rejection 0.146 percent and mean unbounded-set share 74.318 percent. The target-constrained quotient has mean rejection 4.996 percent and mean unbounded-set share 67.306 percent. In the fixed-positive-singular-value cells, regular quotient rejection ranges from 4.38 to 5.74 percent and the unbounded-set share is zero. These calculations illustrate the behavior of projection and quotient procedures as the response Jacobian approaches weak rank.

D.6. Local-projection response estimator and structural profiling

For raw quarter t, let w_t multiply its Hamilton-regression contribution and its VAR residual or LP forecast-origin contribution. The estimator $\hat{s}(w)$ refits both stages and peak normalization. With

$D_t = \partial\hat{s}(w)/\partial w_t|_{w=1}$ and bandwidth b, the finite-sample covariance is

$$\widehat{\Sigma}_b = \sum_t D_t D_t' + \sum_{\ell=1}^{b}\left(1 - \frac{\ell}{b+1}\right)\sum_{t=\ell+1}^{275}(D_t D_{t-\ell}' + D_{t-\ell}D_t'). \tag{A3}$$

The LP impact matrix is the Cholesky factor from a VAR with three lags. Direct horizon regressions use current variables, three lags, an intercept, and a trend. Forecast-origin samples shrink with horizon. The main refitting calculation uses bandwidth 28 and reports bandwidths 8 and 16 as sensitivity rows.

For fixed structural parameters, fifty shock coefficients are represented through fifty normalized policy-rate ordinates. Two peak ordinates and two recursive-zero equalities impose the normalization restrictions. After whitening and eliminating the fixed peaks, the inner problem is a convex least-squares criterion on an affine box. A linear program computes a first-order gap bound for each fitted point. The structural search evaluates 102 distinct parameter points across the four families and reprofiles each point under both coordinate choices and all three bandwidths, giving 612 constrained fits on a common support.

TABLE A5. LP-refitted radius pairs across HAC bandwidths

Experiment	Bandwidth	Family entry	First reversal	Count at 3.9199
Training 50	8	14.3999	11.1563	0
Training 50	16	17.5390	14.2231	0
Training 50	28	23.2840	17.1449	0
Full 146	8	204.0180	150.8373	0
Full 146	16	252.6573	211.4665	0
Full 146	28	323.1614	290.0486	0

All six rows use the same 102 evaluated parameter points. Shock paths are reprofiled for each experiment and covariance. The numerical search generated these points under the declared objective and parameter hull; every reported distance is reevaluated at 300 quarters. In every row, reversal precedes all-family entry.

Across the common-support profiles, the maximum normalization equality error is 3.56×10^{-15}, the maximum rate-bound excess is 2.64×10^{-15}, and the largest relative inner-profile gap is 3.74×10^{-10}. The selected entry and reversal points change by at most 0.000126 in distance between 250- and 300-quarter solves, and the corresponding target change is at most 3.59×10^{-6}.

TABLE A6. Horizon checks at LP entry and reversal evaluations

Family	Point	Distance at 300	$\lvert d_{250} - d_{300} \rvert$	Largest target change
RANK–RE	entry	23.2840	1.57e-08	2.01e-09
RANK–RE	first reversal	32.2795	2.06e-07	1.50e-09
HANK–RE	entry	18.2001	3.63e-07	1.41e-06
HANK–RE	first reversal	25.8588	1.26e-04	5.77e-07
RANK–CD	entry	22.3531	9.51e-06	3.59e-06
RANK–CD	first reversal	27.1542	4.54e-09	3.45e-10
HANK–CD	entry	17.0400	1.25e-08	1.63e-08
HANK–CD	first reversal	17.1449	1.54e-06	1.24e-06

Distances use the 50-coordinate LP experiment and bandwidth 28; shock paths are reprofiled at each horizon. Each first-reversal evaluation has a positive target at both horizons. These pointwise comparisons keep the baseline HANK household grid fixed.

D.7. Experiment-gap sampling calculations

For a profiled distance $d_i = \| L^{-1}(s_i - \hat{s}) \|$, the envelope derivative with respect to the empirical center is

$$g_i = -L^{-\prime} \frac{L^{-1}(s_i - \hat{s})}{d_i}. \tag{A4}$$

Gaussian draws update each distance by $d_i^{(b)} = d_i + g_i' Z^{(b)}$ and recompute family minima, all-family entry, and the first sign crossing across the 102 evaluated points. The HAC-8, HAC-16, and HAC-28 upper endpoints of the 95-percent gap ranges are −1.2974, −1.3484, and −4.2955.

For the common-support decomposition, the 102 LP-profiled fitted responses and policy targets remain fixed while response dynamics, peak normalization, and covariance are changed in sequence. Under bandwidth 28, the gaps are −0.0683, −0.9548, −1.0918, and −6.1391. A second decomposition compares the 4,000 baseline points, the common 102-point parameter support under the original shock fit, and the same support under LP shock profiling. Their HAC-28 gaps are 3.3811, 6.2580, and −6.1391.

The initial calendar block-multiplier calculation reruns Hamilton detrending, lag-augmented LP estimation, impact identification, and response construction while updating profiled distances through local envelope derivatives. Under adaptive peak selection, block lengths 8, 16, and 28 give 95-percent ranges $[-5.8329, 0.4248]$, $[-6.4171, 1.0715]$, and $[-8.5502, 0.9632]$. The fixed-multiplier and fixed-peak variants also produce ranges spanning zero. Table A7 records the full grid examined before the selected full-reprofile calculation.

TABLE A7. Calendar-bootstrap specification grid preceding full reprofiling

Block	Normalization	95% gap range	Pr(gap > 0) (%)	Peak switches
8	Fixed multiplier	[-5.6159, 0.3616]	5.503	0
8	Fixed peak cell	[-6.2591, 0.5773]	5.403	0
8	Adaptive peak	[-5.8329, 0.4248]	4.702	1,458
16	Fixed multiplier	[-5.5025, 1.0150]	14.907	0
16	Fixed peak cell	[-6.9476, 1.1636]	8.854	0
16	Adaptive peak	[-6.4171, 1.0715]	7.354	1,433
28	Fixed multiplier	[-7.9836, 0.7461]	5.253	0
28	Fixed peak cell	[-8.7704, 1.2314]	4.752	0
28	Adaptive peak	[-8.5502, 0.9632]	4.302	1,431

Every row uses 1,999 circular calendar block-multiplier draws on the 275-quarter raw-data vintage. Each draw reruns response estimation and then updates the profiled distances through local envelope derivatives on the 102-point support. The HAC metric, structural support, policy targets, and active profile derivative remain fixed. This grid records the bandwidth and normalization specifications examined before full shock-path reprofiling.

The selected calculation uses HAC bandwidth 28. Each of the 1,999 draws recomputes every constrained shock path and both frontier binders. It gives a point gap of -6.1391, bootstrap quantiles $(-10.2187, -4.9865, 2.2214)$ at probabilities $(0.025, 0.5, 0.975)$, and a 6.353-percent positive-gap frequency. All draws produce a frontier. Across 203,898 point profiles, 240 constraint systems are infeasible and the remaining profiles determine every reported frontier. The maximum equality error is 2.49×10^{-14}, the maximum rate-bound excess is 9.92×10^{-15}, and the maximum propagated gap-envelope width is 3.01×10^{-5}.

Appendix E. Supporting conditions and regular refinements

E.1. Primitive conditions for generated responses

The expansion combines minimum-distance estimation, simulation asymptotics, and implicit-function differentiation (Pakes and Pollard 1989; Gouriéroux et al. 1993). Its role here is to carry the shared treatment path and overlapping empirical inputs into every family's response and their off-diagonal covariance blocks.

For family f, let $\hat{g}_{f,T}(\theta)$ be a vector of empirical and calibration moments, let $\widehat{W}_f$ be a minimum-distance weight, and let $\hat{\theta}_f$ minimize the corresponding criterion up to first-order error. Write the equilibrium system as $F_f(z, \theta, q) = 0$, where q is the policy-treatment path estimated from the same or an overlapping sample. A response entering S_A, S_B, or T_p has the form

$$R_f(\theta, q) = \Phi_f\{z_f(\theta, q), \theta, q\}. \tag{A5}$$

ASSUMPTION A1 (Generated-response regularity). *For every family, the following conditions hold uniformly on the declared detectability neighborhoods and over the maintained data-generating class.*

1. *The true coordinate is interior,* $\hat{\theta}_f \to_p \theta_f$ *uniformly, and the sample first-order condition has remainder* $o_p(T^{-1/2})$.
2. *With* $g_f(\theta) = \mathbb{E}\hat{g}_{f,T}(\theta)$ *and* $G_f = Dg_f(\theta_f)$, *the sample moment is stochastically differentiable and*

$$\sqrt{T}\,\hat{g}_{f,T}(\theta_f) = T^{-1/2}\sum_{t=1}^{T}\psi_{g_f,t} + o_p(1)$$

   under a uniform central limit theorem.
3. $\widehat{W}_f \to_p W_f$ *uniformly,* W_f *is symmetric positive definite,* G_f *has full column rank, and the smallest eigenvalue of* $G_f' W_f G_f$ *is bounded away from zero.*
4. *The treatment path and all shared inputs have a joint asymptotic linear representation with the calibration moments; in particular,* $\sqrt{T}(\hat{q} - q) = T^{-1/2}\sum_t \psi_{q,t} + o_p(1)$.
5. *The selected equilibrium has uniformly invertible* $F_{f,z}$*;* F_f *and* Φ_f *have bounded continuous derivatives; solver residuals are* $o_p(T^{-1/2})$*; and, when the estimand is a continuum model, discretization error is* $o(T^{-1/2})$.
6. *Simulation error is centered and has a joint asymptotic linear representation with the preceding inputs. If* S_T *is the number of simulation draws,* $T/S_T \to \eta \in [0, \infty)$.

The full-rank condition in part 3 concerns estimation of the declared finite-dimensional coordinate inside each family. The empirical-response derivative $DS_{A,f}$ can have weak singular values or deficient rank under the identification configurations studied below.

PROPOSITION A1 (Joint representation of generated responses). *Under Assumption A1,*

$$\text{(A6)} \qquad \sqrt{T}(\hat{\theta}_f - \theta_f) = T^{-1/2}\sum_{t=1}^{T}\psi_{\theta_f,t} + o_p(1), \qquad \psi_{\theta_f,t} = -(G_f' W_f G_f)^{-1} G_f' W_f \psi_{g_f,t},$$

and the influence function of (A5) *is*

$$\text{(A7)} \qquad \psi_{R_f} = R_{f,\theta}\psi_{\theta_f} + R_{f,q}\psi_q + \sqrt{\eta}\,\psi_{\mathrm{sim},f},$$

where

$$\text{(A8)} \qquad R_{f,\theta} = \Phi_{f,\theta} - \Phi_{f,z}F_{f,z}^{-1}F_{f,\theta}, \qquad R_{f,q} = \Phi_{f,q} - \Phi_{f,z}F_{f,z}^{-1}F_{f,q}.$$

Stacking families and common inputs yields a joint influence representation with off-diagonal covariance blocks induced by shared samples, treatment paths, and calibration inputs.

Online Appendix C derives (A6) from the minimum-distance first-order condition, applies the implicit-function derivative in (A8), and tracks the simulation rate. The joint expansion includes treatment-path and common-input covariance alongside parameter-estimation covariance.

E.2. A regular quotient refinement

At a regular interior point θ_0 in family f, let $J_S(\theta_0)$ be the derivative of the selected population moment with respect to structural coordinates and let $r_S = \operatorname{rank} J_S(\theta_0)$. Constant-rank coordinates (x, z) can be chosen so that $x \in \mathbb{R}^{r_S}$ indexes the quotient direction and z indexes the local population-moment fiber.

ASSUMPTION A2 (Local quotient law). *The null is correctly specified at an interior point θ_0. Economic and detectability inequalities are inactive in a neighborhood. The rank r_S is locally constant, positive singular values of the quotient Jacobian are bounded away from zero, and the local profiled minimizer is consistent for the true quotient coordinate. On the local moment fiber, the first-order covariance law factors through x, $\Omega_S(x, z) = \widetilde{\Omega}_S(x)$, and the local quadratic expansion is uniform.*

PROPOSITION A2 (Regular quotient-space refinement). *Under Assumptions 2 and A2, the locally profiled regular criterion obeys*

$$P^{\text{reg}}_{S,T}(b_0) \Rightarrow \chi^2_{d_S - r_S}. \tag{A9}$$

The quotient chart makes the population moment a function of x. The covariance restriction gives the same factorization for the first-order law. The resulting quadratic form is the squared norm of the projection of a d_S-dimensional standard normal vector onto the orthogonal complement of an r_S-dimensional score. Cells with inactive inequalities, locally constant rank, and covariance factorization use the quotient refinement. Other reported cells use the $\chi^2_{d_S}$ projection procedure.

E.3. Regular profiling on a target level set

For a scalar target and family f, define the tested level set and its tangent-response Jacobian by

$$\mathcal{H}_f(b, c) = \{\theta \in \mathcal{N}_f(s^\dagger_A; c) : T_{p,f}(\theta) = b\}, \qquad J_{A|b,f}(\theta) = DS_{A,f}(\theta) B_{f,b}(\theta), \tag{A10}$$

where the columns of $B_{f,b}(\theta)$ form an orthonormal basis for $\ker DT_{p,f}(\theta)$. Let $Q_{A,f,T}$ use the training moment g_A and its $d_A \times d_A$ covariance, and set

$$P^{\text{tc}}_{f,T}(b) = \inf_{\theta \in \mathcal{H}_f(b,c)} Q_{A,f,T}(\theta). \tag{A11}$$

ASSUMPTION A3 (Target-constrained quotient law). *The tested level set is a local manifold around a compatible interior point. The rank $r_{f,b} = \operatorname{rank} J_{A|b,f}(\theta)$ is locally constant, its positive singular values have a uniform positive lower bound, the covariance law factors along the local training-moment fiber* *inside the level set, and the profiled minimizer is consistent. Economic and detectability inequalities are inactive on the local chart, and the quadratic expansion is uniform.*

PROPOSITION A3 (Regular target-level profiling). *Under Assumptions 2 and A3,*

$$P^{\mathrm{tc}}_{f,T}(b_0) \Rightarrow \chi^2_{d_A - r_{f,b}}. \tag{A12}$$

The result permits weak singular values of $DS_{A,f}$ in directions normal to the tested policy level set.

Family-specific inversions combine as

$$\mathcal{C}^{\mathrm{tc}}_T = \bigcup_f \{b : P^{\mathrm{tc}}_{f,T}(b) \le \chi^2_{d_A - r_{f,b}, 1-\alpha}\}. \tag{A13}$$

Coverage follows from the component containing a compatible family. Use of the refinement is conditioned on target-constrained Jacobian spectra, distances to inequality boundaries, covariance variation along the level set, and optimizer diagnostics from multiple starts.

Appendix F. Additional evidence and support diagnostics

F.1. Which evidence moves the baseline frontier

On the 4,000-evaluation baseline support, the distance needed for coverage of all four families is 3.5058. The first evaluated sign reversal is at 4.3567. All-family evaluated entry and common negative delayed signs hold on the interval $[3.5058, 4.3567)$. The upper boundary corresponds to symmetric total testing risk 0.02938. A benchmark at 3.9199 leaves distance slack 0.4367 inside this interval; the frontier also gives the conclusion for every other radius.

Table A8 holds the structural evaluations fixed and changes the response coordinates used to filter them. With the AD shock alone, all 4,000 evaluations enter and twelve inflation targets admit positive values. With the RR shock alone, 3,813 enter and every inflation endpoint is negative. Combining both shocks retains 2,457 and moves the inflation reversal distance from 4.1579 to 4.3567. Output endpoints are negative across the initial support, including the empty-filter row. This output result describes a restriction of that support; the inflation comparison measures exclusion by the response evidence.

FIGURE A1. Evaluated frontiers over training horizons and response radius

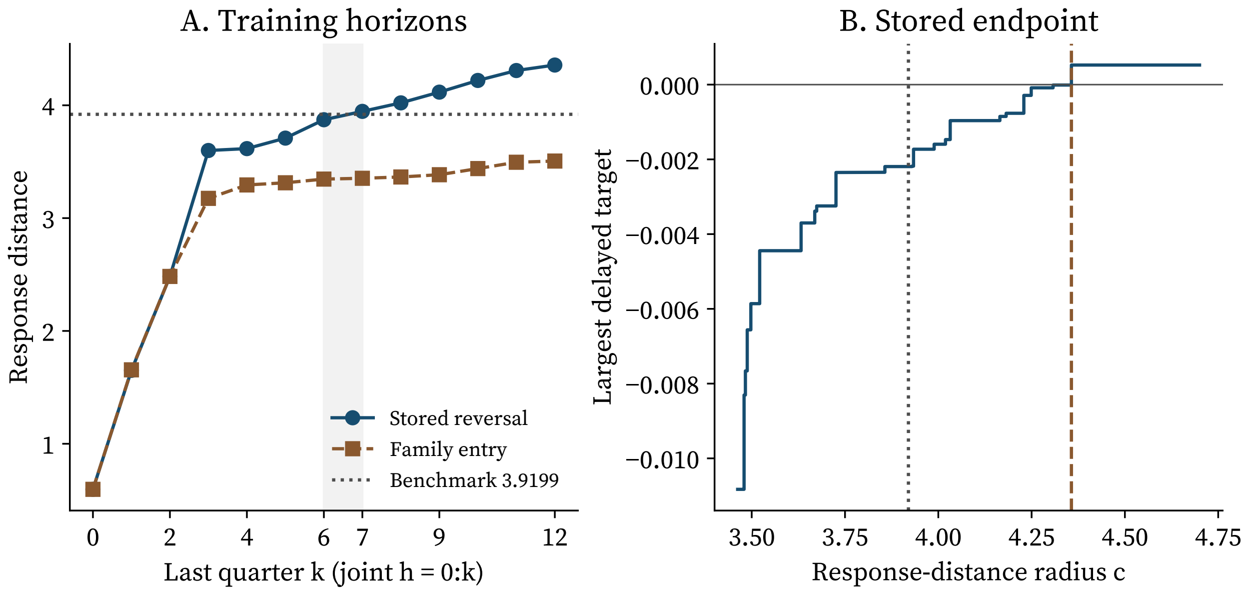

Panel A keeps the 4,000 structural evaluations and fitted inputs fixed. The shaded interval marks the addition of quarter seven. Panel B reports the largest of the 34 delayed target endpoints. Both panels use the baseline covariance and describe structural filtering.

TABLE A8. Evidence and the distance to sign reversal

Evidence	Retained	$c^{\mathcal{D}}_{\text{entry}}$	$c^{\mathcal{D}}_{\text{sign}}$	$Y < 0$	$\pi < 0$
Empty filter	4,000	0.0000	0.0000	17/17	5/17
AD	4,000	1.4476	1.8164	17/17	5/17
RR	3,813	3.0602	4.1579	17/17	17/17
Output, both shocks	3,997	2.5815	2.5815	17/17	5/17
Inflation, both shocks	4,000	2.5263	3.4837	17/17	5/17
Joint horizons 0–3	3,884	3.1750	3.5993	17/17	5/17
Joint horizons 0–6	3,497	3.3459	3.8719	17/17	5/17
Joint horizons 0–7	3,396	3.3529	3.9460	17/17	17/17
Joint horizons 0–12	2,457	3.5058	4.3567	17/17	17/17

The support contains the same 4,000 baseline evaluations in every row. Retained counts and negative-cell counts use $c = 3.9199$, $\beta = 0.99$, and delays zero through sixteen. Each distance uses the principal covariance submatrix of its evidence block. $c^{\mathcal{D}}_{\text{sign}}$ is the first evaluated reversal distance; $c^{\mathcal{D}}_{\text{entry}}$ is the largest family-specific minimum distance. The parameter support, fitted shock inputs, and policy targets remain fixed.

The horizon sequence locates this exclusion within the observed response path. Joint horizons zero–six retain 3,497 evaluations and put the reversing point at distance 3.8719. Adding quarter seven raises its distance to 3.9460, retains 3,396 evaluations, and makes all seventeen inflation endpoints negative. The smallest inflation margin at that step is 0.00008548. Extending the response block through quarter twelve yields margin 0.002193. Online Appendix D reports all thirteen nested horizon designs and the shock–outcome block deletions.

HANK–RE draw 443 is the only baseline evaluation with a nonnegative delayed target. It supplies twelve delayed inflation reversal cells and four persistent inflation reversal cells. Its delay-three inflation target is 0.00052377. The remaining 22 delayed cells and four persistent output cells have zero reversing evaluations in the evaluated support. These records specify the numerical frontier's support. Parameter searches supply additional reversal witnesses under the baseline metric: HANK–RE reaches 4.2077 at inflation delay three, and RANK–RE reaches 4.1847 at inflation delay six. The latter target is 0.00011925 after a 300-quarter solve. Each evaluated witness gives an upper bound on the corresponding parameter-domain infimum.

The conditional-distance decomposition in Lemma 2 identifies the empirical directions that separate draw 443. Its AD distance is 1.8164 and its RR distance is 4.1579. Adding RR conditional on AD contributes 15.6813 to squared distance; adding AD conditional on RR contributes 1.6928. Both blocks together give squared distance 4.3567^2. These increments account for the cross-shock covariance and hold the original parameter support and fitted input paths fixed.

F.2. An executed CMR calibration point

The CMR author-calibration point supplies an output reversal at delay sixteen. Its discounted cumulative output target equals 0.02052. In the first thirteen response quarters, the consumption contribution is −0.12370 and the investment contribution is 0.14422; government expenditure is unchanged. This target measures the response before the delayed tightening is implemented. All seventeen inflation targets and the eight reported persistent output and inflation targets are negative. The extension therefore locates a disagreement in the composition and timing of expenditure before an announced future wedge.

The CMR point's joint training distance is 4.8959. AD evidence admits it at distance 2.4288; RR evidence puts it at 4.6164. Table A9 shows that the AD-filtered expanded support has sixteen negative output cells, and the RR-filtered expanded support has seventeen. The full-evidence benchmark excludes this CMR point. On the 4,001-point expanded support, the output sign frontier becomes 4.8959 and the first joint output–inflation frontier remains 4.3567. These are bounds for the evaluated expansion with a fixed author calibration. Parameter estimation and domain-wide optimization define the remaining CMR-family calculation.

TABLE A9. CMR author-calibration point and evidence filtering

Evidence	CMR distance	Admitted	$Y < 0$	$\pi < 0$
Empty filter	0.0000	1	16/17	5/17
RR	4.6164	0	17/17	17/17
AD	2.4288	1	16/17	5/17
Joint horizons 0–3	3.7651	1	16/17	5/17
Joint horizons 0–7	4.3452	0	17/17	17/17
Joint horizons 0–12	4.8959	0	17/17	17/17

The expanded evaluated support has 4,000 baseline evaluations and one CMR author-calibration point. Negative-cell counts use the union admitted at $c = 3.9199$. The CMR economic parameters are fixed; two 25-quarter monetary-wedge sequences are fitted to the normalized 150-coordinate empirical response. A 24-state news chain implements the delayed interventions under the existing CMR policy rule.

The execution uses the public CMR code with an API port to Dynare 7 and Octave. Its 349-equation author system has maximum static residual 1.29×10^{-12}. The news augmentation has 373 equations and 94 unstable eigenvalues for 94 forward variables. The contemporaneous response agrees with the original author solve within 6.54×10^{-12} per annualized percentage point. Output constructed from expenditure components agrees with accumulated GDP growth within 3.74×10^{-13}. These numerical checks verify the response construction used for the reported CMR point.

F.3. Causal response design

The empirical response vector uses the narrative shock of Romer and Romer (2004) and the text-based shock of Aruoba and Drechsel (2024). Under the maintained shock-exogeneity conditions, the 52 output and inflation responses identify the causal effect of unanticipated policy movements over the first thirteen quarters. The independent response reconstruction reproduces the author OLS impulse responses within 8.66×10^{-11}. Each shock is scaled so that the annualized policy-rate response reaches 25 basis points over horizons zero through twenty-four; the corresponding multipliers are 0.361 for RR and 0.705 for AD.

The event-date design separates the policy component from the information-shock component for 319 announcements. Table A10 reports the response of one- and two-year Treasury yields. At the event date, the four max-t intervals lie above zero. The joint impact tests have p-values 2.82×10^{-18} for the one-year yield and 6.29×10^{-14} for the two-year yield. The five-date policy coefficients remain positive under simultaneous coverage.

TABLE A10. H3 event-date Treasury-yield holdout

Yield and window	MP coefficient [simultaneous interval]	CBI coefficient [simultaneous interval]	Joint p
DGS1, event date	0.574 [0.321, 0.827]	0.878 [0.504, 1.253]	2.82×10^{-18}
DGS1, 5 dates	0.886 [0.468, 1.303]	0.647 [−0.213, 1.507]	4.55×10^{-13}
DGS1, 20 dates	0.904 [−0.071, 1.879]	1.340 [−0.253, 2.933]	1.31×10^{-4}
DGS2, event date	0.650 [0.283, 1.016]	0.861 [0.249, 1.474]	6.29×10^{-14}
DGS2, 5 dates	0.705 [0.242, 1.168]	0.440 [−0.304, 1.184]	3.90×10^{-6}
DGS2, 20 dates	0.526 [−0.411, 1.463]	1.126 [−0.469, 2.721]	0.0244

Each regression uses 319 announcement dates and includes a constant, the median-rotation monetary-policy shock, and the median-rotation central-bank-information shock. Coefficients measure percentage-point yield changes per one-unit shock in the public file. Brackets use the 95 percent max-t interval across 12 slope coefficients from 9,999 event-level Rademacher multiplier draws. Joint p-values use the two-slope HC3 Wald statistic.

The expectation block uses individual Survey of Professional Forecasters submissions from 66 survey quarters. A narrative tightening lowers real-growth forecasts three and four horizons ahead. The RR coefficients are −0.622 and −0.713, with simultaneous intervals [−1.198, −0.046] and [−1.294, −0.133]. A survey-quarter mean specification with Newey–West covariance gives −0.633 and interval [−1.255, −0.012] at horizon three.

TABLE A11. H4E Survey of Professional Forecasters holdout

Forecast outcome	h	RR coefficient	AD coefficient	Responses
Real GDP growth	0	0.21 [−2.86, 3.28]	4.30 [−4.21, 12.80]	2,416
	1	−0.08 [−1.80, 1.65]	1.93 [−3.54, 7.40]	2,416
	2	−0.36 [−1.15, 0.43]	1.48 [−2.65, 5.62]	2,412
	3	−0.62 [−1.20, −0.05]	1.45 [−2.54, 5.43]	2,397
	4	−0.71 [−1.29, −0.13]	1.74 [−2.17, 5.65]	2,338
GDP-price growth	0	0.12 [−1.27, 1.51]	2.94 [−0.94, 6.83]	2,323
	1	0.43 [−0.60, 1.45]	2.06 [−1.25, 5.37]	2,322
	2	0.48 [−0.47, 1.42]	1.63 [−1.83, 5.08]	2,319
	3	0.64 [−0.38, 1.66]	1.16 [−2.11, 4.43]	2,302
	4	0.59 [−0.37, 1.54]	0.94 [−2.44, 4.32]	2,250

Each outcome is the annualized log growth forecast between adjacent SPF level forecasts. Regressions use shocks from the preceding calendar quarter. Brackets give 95-percent intervals simultaneous across 20 slopes, formed from common survey-quarter Rademacher multipliers and the coordinatewise larger forecaster- and calendar-clustered standard error. A survey-quarter mean sensitivity uses Newey–West covariance with lag $h + 1$ and Bonferroni coverage; its RR coefficient for real GDP growth at $h = 3$ is −0.633 with interval [−1.255, −0.012].

The yield and forecast estimates locate two links in the transmission sequence. Policy announcements move short Treasury yields immediately, and the narrative tightening lowers expected real growth at longer forecast horizons. The structural envelope uses the aggregate output and inflation responses for estimation and reserves these market and expectation responses for external validation.

F.4. Continuous response completion and local parameter neighborhoods

The 2,457 compatible evaluated joint response–policy pairs generate the convex completion in (34). Proposition 7 gives the same coordinatewise endpoints for this continuum. Its smallest delayed-target margin is 0.002193, its smallest persistent-target margin is 0.002255, and its largest vertex distance is 3.919872, leaving radius slack 5.59×10^{-5}.

TABLE A12. Continuous response completion at radius 3.9199

Family	Entry witness	Convex margin	Certified lower cells	Recorded $\bar{\delta}^{\text{sign}}$
RANK–RE	3.4280	0.002366	0/34	4.1847
HANK–RE	3.4245	0.002255	0/34	4.2077
RANK–CD	3.3857	0.002193	0/34	4.7428
HANK–CD	3.4089	0.002330	0/34	4.9899

The convex margin is the smallest distance below zero among 34 delayed target coordinates over the family-specific convex response completion. Certified lower cells counts targets with validated global sign-reversal lower bounds. Recorded upper values use 300-quarter evaluated reversals and the baseline metric. The RANK–RE and HANK–CD entries incorporate the LP-directed parameter search, followed by original-protocol input fitting. The HANK–RE entry retains the earlier delay-three witness.

The convex-completion endpoints also give an exact policy-target-scale error budget. If a validation exercise supplies a common coordinatewise additive error bound ε, a delayed cell retains its negative sign when its recorded upper endpoint plus ε lies below zero. All 34 cells retain the sign for $\varepsilon < 0.0021925988$; the binding coordinate is inflation at horizon sixteen. Table A13 records the full count at six allowances. Solver residuals, response-space approximation errors, and held-out residuals enter this calculation after conversion into the policy-target units by a validated family-specific bound.

TABLE A13. Delayed signs under a deterministic policy-target error allowance

Allowance ε	Output	Inflation	All cells
0.0000	17/17	17/17	34/34
0.0010	17/17	17/17	34/34
0.0020	17/17	17/17	34/34
0.0025	17/17	15/17	32/34
0.0050	17/17	5/17	22/34
0.0100	17/17	0/17	17/34

A cell is counted when its benchmark convex-completion upper endpoint plus ε remains below zero. The exact common budget for all 34 delayed cells is $\varepsilon < 0.0021925988$, attained by inflation at horizon sixteen. The table is a deterministic perturbation calculation. A statistical or numerical error source enters this budget after a valid policy-target-scale bound has been established.

At the binding RANK–CD inflation cell, the independent $T = 300$ solve differs from the response artifact by 2.99×10^{-10}. The $T = 200$–$T = 300$ change is 7.84×10^{-7}, and the $T = 275$–$T = 300$ change is 3.73×10^{-8}. Across all 2,457 compatible vertices and 34 targets, 4,914 independent solves give maximum artifact–$T = 300$ and $T = 200$–$T = 300$ discrepancies of 2.375×10^{-8} and 2.682×10^{-5}, equal to 0.001083 and 1.223 percent of the margin. The maximum relative equilibrium residual is

2.24×10^{-16}. Across 26 distinct maximizing witnesses, the maximum equilibrium condition number is 2.43×10^{12}. The compatible-design local-policy and factor-relation secant maxima are 0.3424 and 0.1658, equal to 156.16 and 75.61 margins. These secant rows serve as approximation diagnostics; a native certificate uses domain-wide nonlinear enclosures.

The baseline-design sensitivity recomputes response distance under three locally available covariance choices and a current-vintage response center. The current-vintage center admits 2,417 draws and preserves the 0.002193 margin. The model-artifact training-only covariance admits 3,215 draws and yields margin 0.001470. The raw rank-50 covariance reaches all-family evaluated coverage at radius 4.2862, above the benchmark, so its benchmark policy set is empty. The online appendix reports the four-row specification path.

The retained-support coordinate boxes are the closed coordinatewise hulls of each family's 1,000 retained parameter draws. They have five active coordinates in the RE families and six in the CD families; coordinates with retained standard deviation at most 10^{-8} are fixed. Horizon-300 searches give feasible sign reversals for RANK–RE inflation at delay zero with distance 4.7738, HANK–RE inflation at delay zero with distance 4.2532, and RANK–CD output at delay zero with distance 4.7428. The equilibrium relative residuals are at most 2.03×10^{-17} across these witnesses. Each distance is an upper bound on its family sign-reversal distance.

The HANK–CD search over the retained draws and additional candidate points supplies an output reversal at delay four; reevaluation with the original shock fitting and baseline metric gives distance 4.9899 and target 0.00001217. The RANK–RE witness at 4.1847 and this HANK–CD witness retain positive targets at horizons 250 and 300. Their 300-quarter equilibrium condition numbers are 1.39×10^{13} and 6.68×10^{9}, respectively.

The corrected AD/RR response-column calculation leaves local-box response compatibility open. The frontier results retain their evaluated-support and convex-completion domains. A local or global native-domain compatibility claim requires the additional gate specified in Appendix D.

F.5. Delayed and persistent interventions

For each compatible retained point, the model maps a one-quarter 25-basis-point tightening delayed by h quarters into discounted cumulative output and inflation responses. The evaluated set and its convex response completion share the endpoints in Table A14.

TABLE A14. Cumulative responses to delayed 25-basis-point tightenings

Delay h	Output endpoint interval	Inflation endpoint interval
0	$[-0.110, -0.031]$	$[-0.019, -0.002]$
4	$[-0.179, -0.054]$	$[-0.032, -0.005]$
8	$[-0.211, -0.063]$	$[-0.037, -0.005]$
12	$[-0.183, -0.044]$	$[-0.033, -0.004]$
16	$[-0.120, -0.011]$	$[-0.026, -0.002]$

Entries are the union endpoints across 2,457 compatible retained baseline evaluations from four model families. Proposition 7 gives the same endpoints for their continuous joint response–policy convex completion. Targets discount quarterly responses at $\beta = 0.99$.

Both endpoints are negative at every delay from zero through sixteen quarters. Across the 34 target–delay cells, output endpoints range from -0.211 to -0.011 and inflation endpoints range from -0.037 to -0.002. At delay eight, output lies in $[-0.211, -0.063]$ and inflation lies in $[-0.037, -0.005]$. The smallest margin occurs for inflation at delay sixteen. The output profile reaches its largest absolute lower endpoint near delay eight.

Persistent interventions combine the delayed basis policies according to $u_t = 0.25\rho^t$. Table A15 reports the corresponding endpoints.

TABLE A15. Cumulative responses to persistent tightening paths

Persistence ρ	Output endpoint interval	Inflation endpoint interval
0.00	$[-0.110, -0.031]$	$[-0.019, -0.002]$
0.50	$[-0.254, -0.074]$	$[-0.045, -0.006]$
0.75	$[-0.594, -0.184]$	$[-0.105, -0.016]$
0.90	$[-1.383, -0.438]$	$[-0.240, -0.038]$

The path is $u_t = 0.25\rho^t$ for $t = 0, \dots, 16$. Linearity combines the compatible delayed-intervention responses before endpoints are taken. The four rows use the same 2,457 retained evaluations and continuous convex completion as Table A14.

At $\rho = 0.50$, output lies in $[-0.254, -0.074]$ and inflation in $[-0.045, -0.006]$. At $\rho = 0.90$, the intervals are $[-1.383, -0.438]$ and $[-0.240, -0.038]$. These endpoints apply to the compatible retained evaluations and their convex response completion. The SPF response supplies an observable counterpart at horizons where a persistent intervention continues to load on the economy.

F.6. Weak-rank inference diagnostics

The Gaussian design results are reported next to the projection argument in Section 5.3. The application modes have Jacobian ranks 7, 7, 9, and 9, and their reported local visibility measures range from 1.417 to 98.064 across 136 family–target cells. These are local calculations at the selected modes.

F.7. Primitive response uncertainty and re-estimation

The raw-data reconstruction uses 275 dated quarters from 1954Q3 through 2023Q1. It refits the Hamilton regression on its full lead–lag footprint, estimates responses over 1969Q1–2006Q4, and recomputes each shock's 25-basis-point peak normalization. Calendar weights enter Hamilton dependent-variable dates and VAR residual or LP forecast-origin dates. Numerical derivatives of this two-step estimator supply a 275×150 primitive influence array. The dependence calculation applies a Bartlett HAC sandwich on that common calendar with bandwidths 8, 16, and 28. Recursive RR impact restrictions and fixed peak-normalized rate ordinates have zero derivative.

The VAR point reproduces the baseline normalized response within 1.28×10^{-11}. Its primitive derivative has numerical rank 40. The selected 50-coordinate training covariance is singular, so the positive-definite distance in (16) requires a different response coordinate system or a singular-experiment extension. The LP estimator uses three lags and horizon-specific forecast-origin samples. Its 50-coordinate training covariance is positive definite under all three reported bandwidths. Halving the case-weight differentiation step changes the derivative arrays by relative norms below 1.22×10^{-7}, with zero peak-index switches.

TABLE A16. Response estimator and fixed-input compatibility

Experiment	Coordinates	Entry	Reversal	Count
Baseline taper	52	3.51	4.36	2,457
VAR, HAC 28	50	–	–	–
LP, HAC 8	50	27.22	29.49	0
LP, HAC 16	50	36.11	38.27	0
LP, HAC 28	50	46.43	49.82	0

Entry and reversal are baseline-design radii. Count is the number of evaluations at radius 3.9199. New rows use the original parameter support and fitted inputs, rescaled to the estimator's peak normalization. The VAR derivative has rank 40 across 150 primitive responses; its 50-coordinate training covariance is singular, so the full-rank distance is withheld. LP uses three lags. All HAC rows use the 275-quarter calendar and refit preprocessing and normalization.

Table A16 keeps the original structural parameters and fitted shock inputs fixed, converting their normalization to the new estimator's scale. Every LP row has an empty compatible set at 3.9199. Family entry occurs at radii 27.22, 36.11, and 46.43 for bandwidths 8, 16, and 28. The corresponding evaluated reversal distances are 29.49, 38.27, and 49.82. These rows measure estimator and covariance sensitivity at the original fits. The following refitting experiment varies both structural parameters and shock paths.

The primitive covariance estimates use the stated raw-data vintage and shock-series construction and provide the empirical metric for the subsequent structural refitting exercise.

F.8. Additional response checks

Family-held-out prediction evaluates 68 family–horizon cells, with minimum predictive coverage 0.863 and zero five-percent generated-loss rejections. Treasury-yield, SPF, funding-spread, QFR, SCF, and CEX estimates provide additional transmission evidence. These response checks are reported on their own empirical scales and complement the training-response frontier.

Under the baseline response experiment, family entry and common negative signs hold on the compatible retained design and its convex response–policy completion. On the refitted LP support, the first sign crossing precedes all-family entry and every positive target is below 0.001. The central full-reprofile resampling range spans zero, with 6.353 percent of draws producing a positive gap. Native-domain compatibility remains an open computational gate.

References

Aruoba, S. Boragan and Thomas Drechsel, "Identifying Monetary Policy Shocks: A Natural Language Approach," 2024. NBER Working Paper 32417; forthcoming, American Economic Journal: Macroeconomics.

Gouriéroux, Christian, Alain Monfort, and Eric Renault, "Indirect Inference," *Journal of Applied Econometrics*, 1993, *8*, S85–S118.

Pakes, Ariel and David Pollard, "Simulation and the Asymptotics of Optimization Estimators," *Econometrica*, 1989, *57* (5), 1027–1057.

Romer, Christina D. and David H. Romer, "A New Measure of Monetary Shocks: Derivation and Implications," *American Economic Review*, 2004, *94* (4), 1055–1084.