Too Big to Aggregate

Corporate Control and Policy-Invariant Aggregation

Online Appendix

Chae-Yeon Xon

Yonsei University

September 2026

Abstract

This appendix contains proofs, corporate-control data construction, additional response estimates, rank and repair calculations, and the numerical state-repair exercise. The sections follow the notation and policy target of the paper.

JEL Codes: C61, C63, E17, E32, F23, L13

Keywords: aggregation, corporate control, multinational firms, policy counterfactuals, state sufficiency

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

Appendix A. Formal Theory and Proofs

A.1. A tractable control-overlap representation

In a quadratic special case,

$$\Pi_i(a_i; p) = b_i(p)' a_i - \frac{d_i}{2} a_i' a_i - \frac{\lambda_i}{2}(s_i' a_i)^2, \qquad d_i > 0, \quad \lambda_i \geq 0, \tag{A1}$$

where $s_i \in \mathbb{R}_+^C$ records the geography of jointly controlled activities. The positive curvature block is the negative Hessian,

$$H_i \equiv -D^2_{a_i} \Pi_i = d_i I + \lambda_i s_i s_i'. \tag{A2}$$

Both d_i and $d_i + \lambda_i \|s_i\|^2$ are positive, so the rank-one inverse is

$$H_i^{-1} = d_i^{-1} I - \kappa_i s_i s_i', \qquad \kappa_i = \frac{\lambda_i}{d_i(d_i + \lambda_i \|s_i\|^2)}. \tag{A3}$$

The pre-policy control-overlap matrix is

$$K_t = \sum_i \omega_{it} \kappa_{it} s_{it} s_{it}', \qquad \omega_{it} \geq 0. \tag{A4}$$

It is positive semidefinite. Its off-diagonal entries increase with the nonnegative activity weights of two jointly controlled countries and the curvature of their shared internal margin. With strategic interaction, the local response has the schematic form

$$R_p = \mathcal{M}(I - \Gamma)^{-1} \mathcal{B}(X, K) p, \tag{A5}$$

where $I - \Gamma$ is nonsingular; $\rho(\Gamma) < 1$ is a sufficient local-stability condition. The special case maps predetermined ownership and geographic-operation data into the theory.

A.2. Employment-transmission signs

Proof of proposition 1. The objective in equation (7) is strictly convex because its second derivative equals $\kappa_H^{-1} > 0$. Its first-order condition gives $\Delta n_{t:t+H} = -\kappa_H[\ell_{0H} + (1 + \alpha_H B^F - \nu_H B^E)\varepsilon_t]$. Differentiating first with respect to ε_t and then with respect to each control coordinate yields equation (8). Linear independence of the two response-loading columns gives a two-dimensional image; collinearity gives a one-dimensional image. □

A.3. Linear factorization

LEMMA A1 (Linear factorization). *Let $A : V \to X$ and $B : V \to Y$ be linear maps between finite-dimensional vector spaces. The following statements are equivalent: (i)* $\ker A \subseteq \ker B$*; (ii) there is a unique linear map $C : \operatorname{im} A \to Y$ such that $B = C \circ A$.*

PROOF. If $B = C \circ A$, then $Av = 0$ implies $Bv = 0$. Conversely, define $C(Av) = Bv$. If $Av = Av'$, then $A(v - v') = 0$, so $B(v - v') = 0$ and $Bv = Bv'$. Hence C is well defined and linear on $\operatorname{im} A$. Uniqueness on the image follows immediately. A linear map on $\operatorname{im} A$ extends to all of X by choosing a complementary subspace. □

A.4. Control transmission

Assumption 1 gives a local equilibrium $a(s, p)$ with $a_p = -J^{-1}F_p$. Along a general state direction h,

$$\begin{aligned} D_h R_p = & -(D_h M)J^{-1}F_p + MJ^{-1}(D_h J)J^{-1}F_p \\ & - MJ^{-1}D_h F_p, \end{aligned} \tag{A6}$$

Under assumptions 2 and 3, the first and third terms vanish. If a target is extended to $Y = Ma + y(X) + G(S, p)$, the right side gains the term $D_h G_p$.

PROOF OF THEOREM 1. Differentiate $J^{-1}J = I$ along h to obtain

$$D_h(J^{-1}) = -J^{-1}(D_h J)J^{-1}. \tag{A7}$$

Since $R_p = -MJ^{-1}F_p$ and M and F_p are invariant along a pure control direction,

$$D_h R_p = -MD_h(J^{-1})F_p = MJ^{-1}(D_h J)J^{-1}F_p. \tag{A8}$$

Substitute equation (6) and use $D_h H^L = 0$. Stacking F_{p_k} gives the finite-policy statement. □

PROOF OF PROPOSITION 2. Each term on the right side of equation (10) is zero under the stated conditions. □

A.5. Primitive response rank

PROOF OF THEOREM 2. Substituting equation (11) into theorem 1 gives

$$D_h R_P = \sum_{r=1}^{R} \ell_r(h) MJ^{-1}B_r J^{-1}P. \tag{A9}$$

Vectorizing and stacking the target gives $\mathcal{H}_{\mathcal{P}} = A_{\mathcal{P}} L_C$. The rank inequality follows from $\operatorname{rank}(AB) \leq \min\{\operatorname{rank} A, \operatorname{rank} B\}$. If L_C is onto $\mathbb{R}^R$ and $A_{\mathcal{P}}$ is injective on $\mathbb{R}^R$, their product has rank R. □

A.6. Local aggregation and minimum repair

PROOF OF THEOREM 3. At s_0, apply lemma A1 with $A = Dq_R(s_0)$ and $B = DT_{\mathcal{P}}(s_0)$. The kernel inclusion is equivalent to a linear map on $\operatorname{im} Dq_R(s_0)$. Choose a complement of that image in $T_{x_0}\mathcal{X}$ and extend the map on the complement. This yields definition 1. For the neighborhood claim, the constant-rank theorem gives coordinates (x, u) in which $q_R(x, u) = x$ after coordinate changes on its image. The kernel condition implies $D_u T_{\mathcal{P}}(x, u) = 0$. Connected open u-sections are joined by piecewise C^1 paths, and integrating the zero directional derivative along each path makes $T_{\mathcal{P}}$ independent of u. This defines the stated C^1 factor on $q_R(U)$. □

PROOF OF THEOREM 4. Let $V = \ker Dq_R(s_0)$, $L = DT_{\mathcal{P}}(s_0)|_V$, and $B = Db(s_0)|_V$ for a d-dimensional repair. Closure requires $\ker B \subseteq \ker L$. Rank-nullity gives

$$\operatorname{rank} B = \dim V - \dim\ker B \geq \dim V - \dim\ker L = \operatorname{rank} L = r_{\mathcal{P}}(s_0). \tag{A10}$$

Since $\operatorname{rank} B \leq d$, the lower bound follows. For the converse, choose $V = \ker L \oplus W$, where $\dim W = r_{\mathcal{P}}$. Let B^* vanish on $\ker L$ and be an isomorphism from W to $\mathbb{R}^{r_{\mathcal{P}}}$. Extend B^* to the full tangent space. A local chart produces a smooth map with that derivative, whose augmented kernel equals $\ker L$. □

PROOF OF COROLLARY 1. Because $\mathcal{V}_C \subseteq \mathcal{V}_R$ and $\mathcal{H}_{\mathcal{P}}$ is a coordinate block of the declared target derivative on $\mathcal{V}_C$, the full hidden-response rank is at least R. Theorem 4 gives the lower bound. Write any $h \in \ker Dq_R \cap \ker Db_0$ uniquely as $h = h_C + h_O$ under the stated direct sum. Since $Db_0 h_O = 0$ and $Db_0 h_C = L_C h_C$, the equality $Db_0 h = 0$ gives $L_C h_C = 0$. The kernel assumptions then give $DT_{\mathcal{P}} h_C = 0$ and $DT_{\mathcal{P}} h_O = 0$, hence $DT_{\mathcal{P}} h = 0$. The augmented kernel criterion holds, and the R-dimensional repair attains the lower bound. □

A.7. Rank-one inverse and approximation loss

For $H = dI + \lambda ss'$ with $d \neq 0$ and $d + \lambda\|s\|^2 \neq 0$, define

$$\widetilde{H}^{-1} = d^{-1}I - \frac{\lambda}{d(d + \lambda\|s\|^2)} ss'. \tag{A11}$$

Using $(ss')^2 = \|s\|^2 ss'$, direct multiplication gives $H\widetilde{H}^{-1} = \widetilde{H}^{-1}H = I$. Under $d > 0$ and $\lambda \geq 0$, both nonzero conditions hold.

For the half-separation bound, let a pooled predictor choose $\widehat{T}$ for two states with $\|T(s) - T(s')\| = \Delta$. The triangle inequality implies

$$\Delta \leq \|T(s) - \widehat{T}\| + \|T(s') - \widehat{T}\|. \tag{A12}$$

At least one error is therefore at least $\Delta/2$.

A.8. Continuous state compression and finite-experiment error

The linear part below specializes the Bernstein/manifold-width comparison in Dahlke et al. (2006, Section 2.2) and the nonlinear-width framework of DeVore et al. (1989). The remaining steps apply uniform Taylor bounds to the policy increment on the national-state fiber.

Proof of Theorem 5. Taylor's formula with an integral remainder in the policy argument gives $G(z,\tau) = \tau R(z) + r_p(z,\tau)$ with $\|r_p(z,\tau)\| \leq K_p \tau^2/2$. Taylor's formula along the line segment from 0 to z gives $R(z) = R(0) + Lz + r_s(z)$ with $\|r_s(z)\| \leq K_s \|z\|^2/2$. Consequently,

$$G(z,\tau) = \tau R(0) + \tau L z + r(z,\tau), \qquad \sup_{\|z\|\leq\rho} \|r(z,\tau)\| \leq \epsilon(\rho,\tau). \tag{A13}$$

Extend the right-singular vectors to an orthonormal basis of $\mathbb{R}^k$. For the stated encoder and decoder, the residual linear operator is $\sum_{j>d} \sigma_j u_j v_j'$, with operator norm σ_{d+1}. The triangle inequality gives the upper bound uniformly over the ball. For the lower bound, fix any continuous b and any decoder g, and restrict b to the radius-ρ sphere in $\operatorname{span}\{v_1, \ldots, v_{d+1}\}$. This sphere is homeomorphic to S^d. The Borsuk–Ulam theorem gives an antipodal pair $z, -z$ with $b(z) = b(-z)$; when $d = 0$, the codomain is a singleton and the same assertion holds directly. Orthogonality implies $\|Lz\| \geq \rho\sigma_{d+1}$. Equation (A13) and the reverse triangle inequality give

$$\|G(z,\tau) - G(-z,\tau)\| \geq 2|\tau|\rho\sigma_{d+1} - 2\epsilon(\rho,\tau). \tag{A14}$$

The decoder assigns the same prediction to this pair. Its maximum error on the pair is at least half their separation. Taking the supremum over the ball and then the infimum over encoders and decoders yields the lower bound, together with nonnegativity of error. If $K_s = K_p = 0$, the bounds agree. For $d \geq k$, the identity encoder recovers z and an exact decoder achieves zero error. □

Primitive curvature bound. Write $A = J^{-1}$, and differentiate the response along unit directions u, v on the fiber. With $J_u = D_z J[u]$ and $J_{uv} = D_z^2 J[u,v]$, the second derivative is

$$D_z^2 R[u, v] = M(-AJ_v AJ_u A + AJ_{uv} A - AJ_u AJ_v A)P. \tag{A15}$$

Submultiplicativity and the triangle inequality give equation (20). If $Y = Ma$ and the scalar-policy equilibrium equation is $F(a, z, \tau) = 0$, the implicit second derivative also gives

$$\|Y_{\tau\tau}\| \le \|M\| \kappa \left(F_{aa} \kappa^2 F_\tau^2 + 2F_{a\tau} \kappa F_\tau + F_{\tau\tau}\right), \tag{A16}$$

where each symbol on the right denotes a uniform norm bound for the indicated derivative of F, and F_τ in the quadratic terms denotes its norm bound. These constants apply on the declared regular neighborhood. A piecewise-smooth cap requires a branch-specific neighborhood or a separate remainder estimate across the boundary. □

Appendix B. Corporate-Control Data Construction

The descriptive inputs apply to the stated control-state construction. Regression tables use the alternative standardized prior-mean shock when their captions identify that reference. Complete-margin diagnostics use the same construction.

B.1. Data architecture

The public ownership data combine GLEIF Level 2 accounting-consolidation relationships and reporting exceptions, ESEF issuer filings, DG Competition merger cases, and SEC submissions (Global Legal Entity Identifier Foundation 2026; XBRL International 2026; European Commission, Directorate-General for Competition 2026; U.S. Securities and Exchange Commission 2026, 2004). GLEIF establishes the reported direct or ultimate accounting-consolidating parent and the relationship period. Transaction announcements, regulatory decisions, and annual reports establish announcement, completion, and effective-control dates. ESEF primary statements supply the euro-area issuer census. SEC Item 2.01 filings establish completed U.S. acquisitions, while SEC Companyfacts and primary 10-Q tables supply acquirer and target accounting levels.

The source collection contains 25,675 ESEF filings and 7,346 issuers. The euro-area linkage census contains 7,309 filings for 1,643 issuers in 18 countries, with 1,166 issuers observed for at least three reporting periods. The repository collection excludes German and Irish national filing collections, so the issuer design defines its population as the observed 18-country repository census. Among the 1,643 issuers, 454 match a GLEIF direct-parent record and 1,277 match a parent-reporting exception. The sample records filing

TABLE A1. Sample Construction and Control-State Coverage

Object	Records	Units	Span	Gate
EA20 ESEF census	7,309 filings	1,643 issuers	18 countries	7,306 valid dates
GLEIF parent linkage	665,510 relations	454 issuers	27.6% linked	502 starts
SEC acquisition sample	24 events	24 acquirers	2012–2018	23 agreements
Public-target perimeter	17 pairs	11 primary events	39 quarters	$G_{\text{eff}} = 8.12$
FATS state	1,056 cells	20 countries	2008–2023	222 country-years
Monthly macro panel	2,614 rows	20 countries	2009–2024	2,394 complete
Quarterly macro panel	888 rows	20 countries	2009–2024	716 employment
EA-MPD sample	315 dates	3 windows	1999–2025	315 complete

The ESEF census is the filings.xbrl.org repository snapshot retrieved on August 25, 2026. Three period ends outside 2018–2026 are removed from the valid-date count. GLEIF reports accounting-consolidation relationships. The SEC sample draws two cohorts from Assets frames dated before their event windows, validates completion in Item 2.01 primary documents, and removes other Item 2.01 filings within eight quarters. The public-target primary panel requires four early-pre, four late-pre, and four post-close observations on a combined perimeter. The monthly panel combines lagged FATS coordinates, manufacturing production, HICP, release-delayed GDP, NFC credit and lending rates, trade openness, manufacturing value-added share, and monthly policy-event sums. Complete denotes availability of the declared national benchmark.

TABLE A2. Summary Statistics

Variable	Mean	SD	P10	Median	P90
Investment growth	1.551	17.909	−5.030	0.500	5.830
Employment growth	−0.058	1.443	−1.314	0.065	1.050
ECB policy exposure	0.021	0.889	−0.905	0.000	0.840
Foreign-control share	0.360	0.181	0.159	0.329	0.614
Extra-EU world share	0.170	0.141	0.068	0.129	0.301
Foreign-control index	0.037	0.907	−0.978	−0.096	1.671
Extra-EU control index	0.006	0.972	−0.735	−0.293	0.789

Investment is quarterly real gross fixed capital formation growth ($N = 888$); employment is quarterly manufacturing log growth in percentage points ($N = 716$). Policy exposure is the monthly rotation-prior monetary shock summed over ECB events, with one observation per month ($N = 192$). Primitive shares use valid unit-interval FATS cells ($N = 1{,}055$ and 895). State indices use 222 country-years. Published aggregate rates enter in raw form.

errors, warnings, period-end validity, source country, registration status, and accounting perimeter before outcome construction.

B.2. European affiliate and industry validation scope

The industry bridge uses the legacy Eurostat table by host country, NACE division, control geography, and year. It constructs the same foreign-over-world and extra-EU-over-world coordinates from employment, tangible investment, turnover, and value added, standardizes each margin within division, averages cells with at least three margins, and dates

TABLE A3. Within-Country Industry Employment Responses

Horizon	Foreign	95% simultaneous interval	Extra-EU	95% simultaneous interval
0	−0.090	[−0.232, 0.052]	0.018	[−0.155, 0.191]
1	−0.115	[−0.284, 0.054]	0.056	[−0.082, 0.195]
2	−0.154	[−0.365, 0.056]	0.013	[−0.179, 0.205]
4	−0.040	[−0.327, 0.248]	−0.034	[−0.264, 0.196]

The outcome is cumulative 100Δ log division employment from quarter t through horizon h. States use the preceding year's legacy FATS cells. The maximum-t critical value is 2.539 across the two monetary interactions and four horizons, using 9,999 calendar-quarter Rademacher score draws and 48 clusters. The four-quarter joint-zero wild value is $p = 0.8725$.

the state one calendar year before the policy quarter. The source layer contains 22,288 margin cells and produces 3,192 country-division-year states across 19 countries and 24 manufacturing divisions. The Eurostat quarterly labor-input table supplies 17,010 seasonally adjusted division-employment observations. Common horizon support yields 9,127 observations from 259 country-division cells in 13 countries over 48 quarters (Eurostat 2026b,c).

For each horizon, the equation includes monetary and information shocks interacted with both industry-state coordinates and absorbs country-by-quarter, division-by-quarter, and country-by-division effects. The first effect set holds fixed the national policy environment and every observed country-quarter condition; the second absorbs common division-specific adjustment; the third absorbs persistent local specialization. Calendar-quarter score wild draws provide the primary inference, with two-way country-division and quarter clustering as a secondary covariance calculation. Table A3 reports the unweighted specification.

The foreign-control path is negative through every horizon and in 11 of 13 four-quarter country leaveouts. The extra-EU path is positive through two quarters and changes sign at four quarters; one of 13 four-quarter leaveouts is positive. These estimates leave the eight simultaneous intervals containing zero. They supply a within-national-environment check on the first corporate coordinate and set the precision boundary for a two-coordinate industry interpretation.

An employment-weighted specification uses the same panel and equation, weighting each observation by the preceding-year WORLD-employment cell and normalizing weights to mean one. Table A4 reports the complete path.

Employment weighting yields the predicted pair at impact and joint evidence through one quarter. At two and four quarters the joint values are 0.1766 and 0.8769, and the four-quarter pair changes direction. The weighted evidence therefore locates an initial within-country response among aggregate-relevant industries and leaves the medium-horizon national loading as the paper's response-representation object.

TABLE A4. Employment-Weighted Within-Country Responses

Horizon	Foreign	95% simultaneous interval	Extra-EU	95% simultaneous interval	Joint p
0	−0.133	[−0.239, −0.028]	0.104	[0.020, 0.187]	0.0001
1	−0.118	[−0.252, 0.015]	0.086	[−0.008, 0.180]	0.0137
2	−0.118	[−0.280, 0.045]	0.045	[−0.081, 0.172]	0.1766
4	0.039	[−0.150, 0.227]	−0.001	[−0.157, 0.155]	0.8769

Weights use preceding-year employment in the FATS WORLD cell, with zero winsorization. The effective sample size is 2,467.1 and the largest normalized weight contributes 0.149 percent of total weight. The maximum-t critical value is 2.554 over all eight loadings. At impact the joint one-sided sign value is 0.0002. Inference uses the same 9,999 calendar-quarter score draws and 48 clusters as Table A3.

A country-history reassignment exercise addresses persistent national types. Five countries provide complete 2008–2019 state histories and common support for every outcome and all eight national controls. The exercise assigns the five two-coordinate histories to countries in all 120 ways and re-estimates the saturated four-horizon system. The observed leading singular value is 2.131, with exact upper-tail value 0.392; the observed oriented four-quarter margin has upper-tail value 0.208. The measured histories yield the pooled sign pair, while their assignment is compatible with many other persistent country-history correspondences in this balanced core.

The affiliate-level contrast calls for local affiliates within the same host country, NACE industry, and policy date, using historical ultimate-control residence and employment on a constant local perimeter. ESEF describes issuer financial reporting: IFRS consolidated primary statements receive detailed tags, while notes receive block tags (European Securities and Markets Authority 2026). GLEIF Level 2 reports direct and ultimate accounting-consolidating parents (Global Legal Entity Identifier Foundation 2026). A scan covers 58,728 tagged facts in 105 issuer filings underlying 172 issuer-years for 29 issuers. Its numeric headcount-tag screen returns zero candidates under the documented concept and unit rules, and the existing balanced financial-outcome screen admits one documented switch. The verified host-affiliate employment and historical-control-by-industry joins each contain zero rows. An affiliate-level design requires stable local identifiers, employment on a consistent local perimeter, NACE industry, historical ultimate-parent residence, and within-country-industry-period overlap across control states. A control-assignment estimand also requires separation from the acquisition package.

B.3. FATS primitives and harmonization

The legacy table covers 2008–2020 under the pre-EBS production system, and the current table covers 2021 onward under the EBS production system. The harmonizer retains source values, status flags, units, control-country categories, industry codes, and vintage

TABLE A5. Persistent-Country Sensitivity at Four Quarters

Specification	Foreign control	Extra-EU control	Family p
Saturated national controls	−0.563 [−0.976, −0.137]	0.646 [0.318, 0.977]	0.0001
Country monetary and information slopes	−0.965 [−2.594, 0.657]	1.028 [−1.498, 3.646]	0.7212
Country slopes and linear trends	−0.725 [−2.089, 0.632]	0.476 [−1.322, 2.277]	0.8221

Entries are four-quarter employment-response interactions under the legacy shock benchmark with pointwise 95 percent basic wild-calendar-quarter intervals. Family values use the all-horizon maximum-absolute-t calibration. All equations include country and calendar-quarter effects, the two state levels, information-shock interactions, and the eight national controls in levels and interacted with both shocks. The latter two specifications use annual within-country FATS index movement after absorbing country-specific policy slopes.

TABLE A6. Control-State Measurement Robustness

Index construction	$\widehat{\delta}_F$	95% interval	$\widehat{\delta}_E$	95% interval
Four-margin baseline	−0.563	[−0.976, −0.137]	0.646	[0.318, 0.977]
Employment omitted	−0.547	[−0.963, −0.122]	0.676	[0.368, 0.987]
Tangible investment omitted	−0.566	[−0.983, −0.139]	0.710	[0.350, 1.073]
Turnover omitted	−0.548	[−0.953, −0.136]	0.609	[0.267, 0.952]
Value added omitted	−0.579	[−0.980, −0.165]	0.585	[0.271, 0.904]

Each row reports the four-quarter employment response. The baseline is recomputed for each index construction on the same declared support. Every row uses the same support rule.

identifiers. For employment, tangible investment, turnover, and value added, it constructs foreign activity over world activity and extra-EU activity over world activity when the published numerator and denominator yield a share in the closed unit interval. The two indices standardize each valid margin within vintage and average country-years with at least three margins.

The harmonized manufacturing panel contains 1,056 country-year-metric rows for 20 countries over 2008–2023. A total of 896 rows contain both intra-EU and extra-EU components, and the maximum absolute gap between reported foreign activity and the sum of these components is one source unit. The state file contains 222 country-years. The estimation sample uses the preceding year's state, and the primary rank window ends before the production-system transition.

B.4. Measurement robustness

The baseline composite uses employment, tangible investment, turnover, and value added. Each leave-one-margin-out index is reconstructed within source vintage, lagged one year, merged into the same saturated national benchmark, and evaluated with 9,999 calendar-quarter wild draws.

The raw pair has variance-inflation factor 2.559 and standardized condition number 2.848. Residualizing both interaction columns on the national benchmark lowers these values to 1.079 and 1.319. Orthogonalizing the extra-EU coordinate against foreign control gives coefficient pair $(-0.023, 0.646)$. The first coefficient follows the residualized foreign-control direction as E/W moves by 0.8365 per foreign-control unit. The declared $(-0.563, 0.646)$ pair holds the other coordinate fixed. The two bases span the same response surface. A separate conditional-geography coordinate based on E/F has correlation 0.285 with B^F and gives pair $(-0.109, 0.437)$ with intervals $[-0.439, 0.227]$ and $[0.237, 0.640]$. This last row uses a different economic construct from the declared E/W state.

An outcome-blind principal-component analysis begins from all eight margin-specific shares. Among 221 country-years with a complete primitive dictionary, its first two components explain 91.34 percent of state variation. The corresponding twelve-row response matrix has normalized singular values 0.309 and 0.159. The covariance-weighted restricted-rank statistic is 17.916 against a 5 percent critical value of 17.013, giving $p = 0.0308$ with the same 9,999 aligned quarter draws. Specifications adding a third or fourth component fail the frozen residualized-interaction support rule because the effective country count falls below four and the largest country share exceeds 0.40. The public data therefore support two measured response directions in this dictionary and provide no basis for a higher-dimensional claim. The replication record retains every basis, completion rule, and coefficient path considered in this analysis.

B.5. Additional channel data

The ECB Bank Lending Survey series measures the national diffusion index for enterprise credit standards, where positive values denote tightening (European Central Bank 2026a). The common four-horizon sample contains 502 country-quarter observations, 16 countries, and 42 calendar-quarter clusters from 2009Q3 through 2019Q4. Eurostat outward FATS supplies manufacturing employment and gross tangible investment for affiliates located inside and outside the EU (Eurostat 2026d). The published two-cell support yields 90 country-year-metric observations over 2021–2023. ECB balance-of-payments statistics supply direct-investment debt assets and liabilities for the internal-finance validation.

B.6. Policy shock construction

The EA-MPD monetary-event window supplies 315 policy dates. The short-rate factor is the first principal component of one-, three-, six-, and twelve-month OIS changes, oriented toward a positive twelve-month loading and scaled to unit event-date standard deviation. Admissible rotations assign positive rate and negative EURO STOXX 50 impacts to contractionary policy shocks and positive impacts in both coordinates to information shocks.

The reference uses the midpoint angle of the admissible interval, and the continuous interval defines the shock set. Separate release windows, leave-meeting-out samples, and alternative factor decompositions measure sensitivity.

B.7. SEC acquisition sample and accounting perimeter

The acquisition sample contains 24 isolated transactions, 23 agreement dates, and 17 public-target CIK matches. Eleven pairs satisfy the twelve-observation sample rule: Hologic–Gen-Probe, Stryker–MAKO, AECOM–URS, AT&T–DIRECTV, Pfizer–Hospira, Dentsply–Sirona, Exelon–Pepco, Microsoft–LinkedIn, Amazon–Whole Foods, Concho–RSP Permian, and Tidewater–GulfMark. Transaction filings establish effective dates and public-target matches. Gen-Probe's March 31, 2010 total-assets observation comes from its filed 10-Q consolidated balance sheet (Gen-Probe Incorporated 2010).

The extraction forms levels before transformations. For every quarter before closing, assets equal the sum of acquirer and target values for the same calendar frame and USD unit. For every quarter from closing onward, assets equal the surviving acquirer's consolidated value. The response windows contribute four observations per event in each of the early-pre, late-pre, and post bins.

B.8. Legal-switch estimator and falsification

For acquirer $a(i)$, target $t(i)$, and closing quarter c_i, define the combined asset perimeter

$$A_{iq} = \begin{cases} A_q^{a(i)} + A_q^{t(i)}, & q < c_i, \\ A_q^{a(i)}, & q \geq c_i. \end{cases} \tag{A17}$$

The second branch uses the surviving acquirer's consolidated statement. For horizon h, let $Z_{iq}^h = \log A_{i,q+h} - \log A_{i,q-1}$. Early and late bins each contain four shock quarters whose response endpoints precede the agreement quarter; the post bin contains the four quarters after legal closing. The closing quarter and agreement-to-closing response endpoints leave the estimation set.

Let U_q be the updated Bu–Rogers–Wu FOMC shock summed by quarter and scaled by its published-series quarterly standard deviation (Bu et al. 2021). With D_{iq}^A marking the late or post bins and D_{iq}^P marking the post bin, estimate

$$Z_{iq}^h = \alpha_i + \lambda_q + \delta_1 D_{iq}^A + \delta_2 D_{iq}^P + b_i U_q + \pi_h U_q D_{iq}^A + \tau_h U_q D_{iq}^P + u_{iq}^h. \tag{A18}$$

Event-specific shock slopes are normalized relative to one event because calendar-quarter effects absorb the common shock. The coefficient π_h measures the late-minus-early loading; τ_h measures the post-minus-late loading. Event and quarter fixed effects, bin-level controls, and event-specific shock slopes absorb time-invariant event exposure, common macro conditions, level shifts at closing, and baseline policy sensitivity.

The primary panel contains 11 complete public-target events and 132 response observations over 39 quarters. It passes the declared event-count, quarter-count, residualized-variation, rank, and conditioning gates. The estimated post-minus-late loading at impact is -0.0040, with a two-way-clustered 95 percent interval of $[-0.0588, 0.0507]$; the signed-effect gate is therefore closed. Equation (A18) measures the completed-acquisition package component in equation (A19). The replication record contains the outcome catalogue, event and quarter sign flips, leave-one-event-out paths, future-shock placebo, and agreement-to-closing support assessment.

B.9. Inference

The macro equations use calendar-period clustered inference, country-cluster diagnostics, and Rademacher wild-period bootstrap intervals. Each block statistic covers both control coordinates over all frequency-compatible horizons, and Bonferroni adjustment covers the outcome family. The SEC direct-switch equation uses 11 events and 39 calendar quarters with two-way clustered covariance and finite-sample sign-flip inference.

Appendix C. Corporate-Control Measurement and the Causal Object

Eurostat assigns foreign control through the ultimate controlling institutional unit, so each FATS cell records the residence of the decision perimeter governing host-country activity. The empirical state combines four manufacturing margins and dates them one year before the policy quarter. This timing turns corporate control into a predetermined response characteristic while country and calendar-quarter effects absorb persistent national attributes and common innovations.

The assigned treatment in the macro application is the EA-MPD monetary innovation. The interaction coefficients identify its causal response across predetermined corporate states under the event-timing, sign-decomposition, exclusion, and common-shock conditions in section 5.3. The estimand is a conditional response surface over the observed national and corporate indices.

Let $\mathcal{R}_h(b, w) = \partial_\varepsilon \mathbb{E}[Y_h(\varepsilon; b, w) \mid X]$ denote the policy-response surface, where b is the declared corporate-control coordinate and w collects financing, accounting, integration, technology, and demand primitives. The macro interaction averages response heterogeneity over the companion-primitives distribution conditional on observed b and X; interpreting it as a derivative holding w fixed requires an additional assignment condition.

A legal switch from (b_0, w_0) to (b_1, w_1) obeys the exact decomposition

$$\mathcal{R}_h(b_1, w_1) - \mathcal{R}_h(b_0, w_0) = \underbrace{\mathcal{R}_h(b_1, w_0) - \mathcal{R}_h(b_0, w_0)}_{\Delta_h^B} + \underbrace{\mathcal{R}_h(b_1, w_1) - \mathcal{R}_h(b_1, w_0)}_{\Delta_h^W}. \tag{A19}$$

The first term is the control-coordinate assignment effect at fixed companion primitives. The second term records the companion-primitives component. Designs that hold w fixed identify Δ_h^B; an acquisition contrast identifies their sum. This decomposition fixes the evidentiary role of each design before the estimates are interpreted.

Three data layers trace the corporate interpretation. Inward FATS measures the host-country control state used in estimation. Outward FATS measures the geography of affiliates controlled by euro-area institutional units. ECB balance-of-payments debt flows and the Bank Lending Survey measure internal funding and external credit supply. Each layer retains its own unit, frequency, and inference rule.

C.1. Legal-switch evidence

An auxiliary design follows completed U.S. public acquisitions reported under SEC Form 8-K Item 2.01. It fixes the accounting perimeter by summing acquirer and target levels before closing and following the consolidated acquirer afterward. Equation (A19) assigns this design to the acquisition-package contrast $\Delta_h^B + \Delta_h^W$. Eleven events supply 132 log-asset response observations over 39 quarters. The post-close loading change for a one-standard-deviation Bu–Rogers–Wu contractionary shock is −0.0040, with a 95 percent interval from −0.0588 to 0.0507. The one-quarter future-shock placebo is 0.0341, with an interval from −0.0687 to 0.1368. The estimator uses a fixed support rule, outcome family, temporal falsification, and legal-transfer feasibility conditions. The euro-area FATS interaction corresponds to the two control coordinates.

Appendix D. National-State Sufficiency, Rank, and Response Evidence

Tables and model-comparison figures identify the alternative prior-mean shock benchmark where it is used. The admissible frontier, composite-null inference, and continuous shock set use the declared target metric and rotation rule.

D.1. Alternative shock benchmark

D.2. Quarterly response dimension

Let $\widehat{\Lambda}_Q$ stack the two corporate-state loadings from manufacturing employment, NFC credit growth, and real investment at horizons 0, 1, 2, and 4 quarters. The twelve rows

TABLE A7. State-Dependent Employment Responses

Horizon	$\widehat{\delta}_F$	95% interval	$\widehat{\delta}_E$	95% interval
0	−0.052	[−0.168, 0.068]	0.227	[0.122, 0.331]
1	−0.196	[−0.371, −0.018]	0.228	[0.060, 0.396]
2	−0.199	[−0.448, 0.054]	0.254	[0.048, 0.462]
4	−0.563	[−0.976, −0.137]	0.646	[0.318, 0.977]

The sample contains 462 observations at impact and 441 at four quarters, with 47 common calendar-quarter clusters. Outcomes are percentage-point cumulative log changes. One index unit equals one unit in the average of the available primitive-margin within-vintage z-scores. Coefficient intervals use 9,999 Rademacher wild calendar-quarter-cluster draws. The all-horizon maximum-absolute-t statistic is 4.171 with wild-bootstrap $p = 0.0001$.

share the quarterly frequency, state window, calendar-quarter cluster universe, and joint Rademacher draws. The two declared columns bound this response projection's rank by two. A bound on the full hidden-response rank requires the model's primitive-spanning restrictions. The dimension test is

$$H_0: \quad \operatorname{rank}(\Lambda_Q) \leq 1. \tag{A20}$$

For coefficient vector $\widehat{b} = \operatorname{vec}(\widehat{\Lambda}_Q)$ and joint bootstrap covariance $\widehat{\Sigma}$, the minimum-distance statistic is

$$Q_1 = \min_{u \in \mathbb{R}^{12},\, v \in \mathbb{R}^2} \left[\widehat{b} - \operatorname{vec}(uv')\right]' \widehat{\Sigma}^+ \left[\widehat{b} - \operatorname{vec}(uv')\right]. \tag{A21}$$

Restricted-null bootstrap samples add jointly aligned wild-score perturbations to the fitted rank-one matrix. Transforming outcomes by an invertible row map and transforming $\widehat{\Sigma}$ by the corresponding Kronecker map leaves equation (A21) unchanged.

The composite-null evaluation includes rank zero and several rank procedures. The Chen–Fang procedure retains rank one; the empirical summary therefore uses approximation loss with sampling uncertainty. The full-fiber interpretation invokes the spanning conditions in corollary 1.

D.3. Temporal stability and exclusion tests

The parsimonious reference preserves the foreign-control and extra-EU sign pair in the common-production window, the full sample, and the 2019-forward temporal split. The 2019-forward joint test gives $p = 0.00021$. Re-estimating that reference after excluding each country or each policy year preserves both signs in every run. The saturated benchmark in Table A7 supplies the national-state conditioning result.

FIGURE A1. The Dimension of Quarterly Quantity Responses

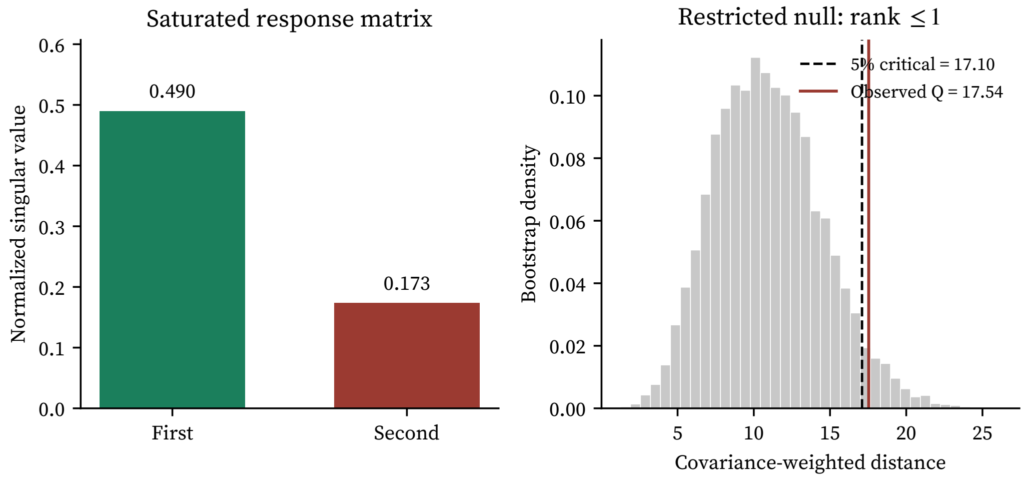

The response matrix stacks employment, NFC credit growth, and real investment at impact and one, two, and four quarters. Panel A gives its two outcome-standard-deviation-normalized singular values. Panel B compares the covariance-weighted minimum-distance statistic with its restricted-rank-one bootstrap distribution. The cumulative, incremental, and raw-unit statistics agree to displayed precision.

TABLE A8. Quarterly Quantity Response Rank

Response-matrix object	Value
First normalized singular value	0.490
Second normalized singular value	0.173
Measured projection	two columns
Fitted rank-one calibration set	$\{2\}$

Restriction	Statistic	5% critical value	p-value	Decision
$\mathrm{rank}(\Lambda_Q) \le 1$	17.538	17.098	0.0407	rejected
Parsimonious reference	28.783	16.806	0.0001	rejected

The sample window is 2009Q2–2020Q4 with 47 common calendar-quarter clusters. Every saturated equation includes the national benchmark in levels and interacted with both shock components. Restricted-null inference uses 9,999 jointly aligned Rademacher draws and the full 24-by-24 coefficient covariance. The covariance has rank 24 and condition number 1,729.5.

D.4. Additional outcome blocks

The mechanism predicts that the employment response travels with corporate financing and changes in the location of foreign-controlled employment. Centered wild-bootstrap block tests place the response in employment, credit and lending rates, and cross-border reallocation. Investment supplies a separate $p = 0.0315$ result and a five-block Holm value of 0.063. Manufacturing production, ECB balance-of-payments internal funding, and ECB Bank Lending Survey credit supply form lower-precision channel checks. This evidence

TABLE A9. Temporal Stability of the Parsimonious Reference

Window	$\widehat{\delta}_F$	95% interval	$\widehat{\delta}_E$	95% interval	Joint p
Full, 2009–2023	−0.367	[−0.573, −0.158]	0.629	[0.415, 0.842]	2.8×10^{-8}
Legacy state, 2009–2020	−0.347	[−0.600, −0.087]	0.600	[0.374, 0.842]	3.8×10^{-8}
2019 forward	−0.247	[−0.480, −0.010]	0.462	[0.200, 0.733]	0.00021

Rows use state main effects, information-shock interactions, country effects, and calendar-quarter effects. The 2019-forward row contains 240 observations and 20 calendar-quarter clusters. Wild calendar-quarter-cluster intervals use 1,999 draws. This appendix reports the saturated benchmark, within-country response-slope estimates, and production-system transition diagnostics.

FIGURE A2. Response Heterogeneity across Outcome Blocks

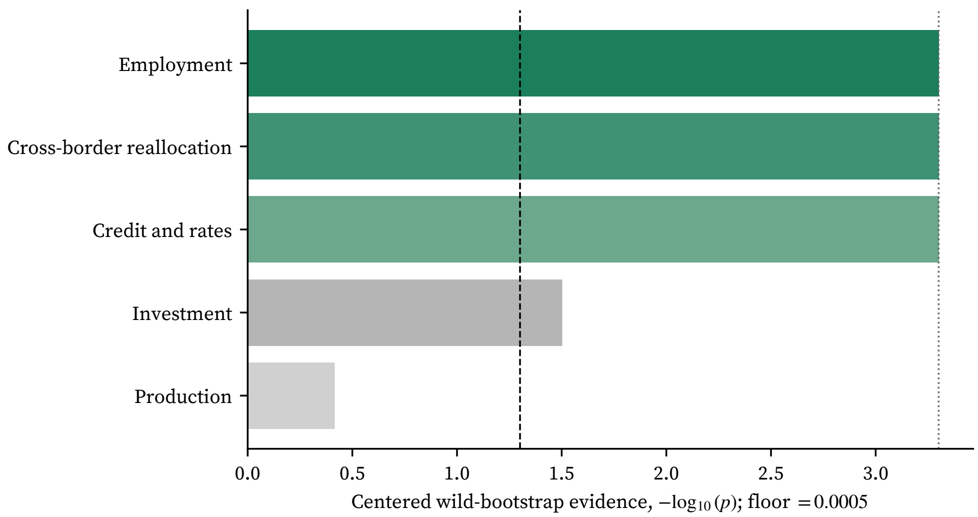

Bars report $-\log_{10}(p)$ from centered joint wild-bootstrap quadratic-form tails. The dashed line marks $p = 0.05$. Each block covers both corporate-control interactions at every declared horizon for that frequency. The employment, credit-rate, and cross-border values equal the 1,999-draw reporting floor of 0.0005.

pattern locates the response in labor, credit, and host-country reallocation while retaining several financing implementations as distinguishable channels.

The two-year foreign-controlled-employment-share equation supplies a reallocation result. Its coefficients on the foreign-control and extra-EU-control interactions are −0.00298 and 0.00372; the joint statistic is 15.73 with centered-bootstrap $p = 0.0005$. The credit-growth block gives statistic 16.77 with $p = 0.0165$, and the combined credit-and-rate block gives statistic 65.32 with $p = 0.0005$.

ECB balance-of-payments data supply an external internal-finance outcome (European Central Bank 2026b). Group net inflow is direct-investment debt liabilities incurred minus debt assets acquired, divided by prior-year GDP; the series stays outside calibration. At one quarter, the foreign-control interaction is 0.984 with interval [0.137, 1.849], and the extra-EU interaction is −1.450 with interval [−3.021, 0.094]; their joint p-value is 0.073. At

TABLE A10. Outcome-Block Tests

Outcome block	Statistic	Degrees of freedom	Bootstrap p	Holm p
Employment	73.65	8	0.0005	0.0025
Credit and rates	65.32	16	0.0005	0.0025
Cross-border reallocation	205.38	14	0.0005	0.0025
Investment	14.98	8	0.0315	0.0630
Manufacturing production	8.64	8	0.3840	0.3840

The employment row uses the full quarterly panel; table A7 reports the primary legacy-state window. Each primary block p-value is the centered wild-bootstrap tail of the displayed quadratic form. The simulation floor is $1/(1{,}999 + 1) = 0.0005$; asymptotic chi-square tails provide a secondary diagnostic. Holm adjustment covers the five displayed blocks.

the declared four-quarter horizon, the coefficients are 0.879 and −1.338, with intervals $[-0.512, 2.200]$ and $[-2.993, 0.309]$ and joint $p = 0.291$. The point estimates describe an inflow toward foreign-controlled balance sheets after contraction and a smaller euro-area funding draw where extra-EU reach is larger. Winsorization yields one-quarter joint $p = 0.074$, and the S1P institutional-sector definition yields $p = 0.148$.

D.5. Rotation equivariance and simultaneous approximation bands

Let Z be the centered, Cholesky-whitened 315×2 rate-equity event matrix, with sample covariance I_2. The alternative prior-mean shock matrix satisfies $S_0 = ZD$ for an invertible 2×2 matrix D. An admissible rotation gives $S_\theta = ZQ_\theta$, where the first column of Q_θ is $(\cos\theta, \sin\theta)'$. For each state or national control, both monetary and information interactions enter the regression. The nuisance column space is consequently invariant to this change of shock basis. If b_0 stacks the two estimated coefficients, then $b_Z = Db_0$ and the monetary coefficient at angle θ is $(\cos\theta, \sin\theta)b_Z$. This uses orthogonal sign-restricted identification in the tradition of Arias et al. (2018), specialized to the two observed event surprises.

Apply this transformation to each of the twelve response equations and to every aligned wild-score draw. Writing the two whitened response matrices as M_0, M_1 and their bootstrap errors as $E_0^{(j)}, E_1^{(j)}$ gives

$$\widehat{\Lambda}_\theta = \cos\theta M_0 + \sin\theta M_1, \qquad E_\theta^{(j)} = \cos\theta E_0^{(j)} + \sin\theta E_1^{(j)}. \tag{A22}$$

The same 9,999 Rademacher quarter weights enter every equation and rotation. Direct four-quarter employment estimates at the first, middle, and last retained rotations agree with the transformed coefficient construction. Row normalizations remain fixed across rotations.

The strict sign restrictions yield an open angle interval with closed limits $[-1.382489, 0]$ radians. Taking its closure preserves the suprema used for the uncertainty envelope. For each bootstrap draw, define $U_j = \sup_\theta \|E_\theta^{(j)}\|_{\mathrm{op}}$. A 2,049-point grid evaluates this supremum with an upper correction

$$U_j \leq \max_{\theta \text{ on grid}} \|E_\theta^{(j)}\|_{\mathrm{op}} + \frac{\Delta\theta}{2}\sqrt{\|E_0^{(j)}\|_F^2 + \|E_1^{(j)}\|_F^2}. \tag{A23}$$

The correction follows from the derivative norm and the distance to the nearest grid point. Its maximum across draws is 0.000316. The 95th percentile of the corrected envelope is $q_{.95} = 0.563745$. Singular-value perturbation inequalities give simultaneous intervals $[\max\{0, \widehat{\sigma}_j(\theta) - q_{.95}\}, \widehat{\sigma}_j(\theta) + q_{.95}]$ for both singular values and every angle. The displayed union band additionally envelopes the estimated singular-value curve, with numerical correction 0.000254. This is a wild-bootstrap sampling approximation conditional on the measured event covariance and FATS indices. Sampling of the event covariance, alternative sign restrictions, and measurement error in control indices define further uncertainty layers.

Coefficient extrema use the endpoints and stationary points of $a\cos\theta + b\sin\theta$. Applied to the four-quarter coefficient errors, this calculation gives a simultaneous pair-and-angle radius of 0.664438. Pointwise reference intervals, simultaneous coefficient envelopes, and spectral envelopes have different coverage targets. The rank-one diagnostic uses a 180-point direction grid followed by bounded minimization around every detected local minimum. All 4,520 covariance matrices have rank 24; their largest condition number is 3,928.6. The reference distance agrees with the restricted-null analysis to displayed precision. The reported envelopes use the same rotation-level coefficients, singular values, and aligned bootstrap draws.

The central-reference twelve-row spectral band is $[0, 0.637]$; the alternative prior-mean band is $[0, 0.585]$. The simultaneous pair-and-angle employment bands are $[-1.271, 0.552]$ and $[-0.207, 1.297]$. Covariance-weighted rank-one distances range from 16.545 to 35.986 across retained rotations; these are observed distances. The prior-mean pair is standardized after averaging and can lie outside the orthogonal family. The main-text reference uses $\theta_c = -0.691244$, with impact matrix $\left(\begin{smallmatrix} 0.770453 & 0.637496 \\ -0.529483 & 0.438110 \end{smallmatrix}\right)$ in rate-factor and equity units. Its signs satisfy all four restrictions and its rotation is orthogonal to numerical precision. Rank tests in Appendix D.8 condition on their declared reference; the spectral envelope covers the entire angle interval conditional on event covariance.

FIGURE A3. Employment Responses and Approximation Cost across Shock Rotations

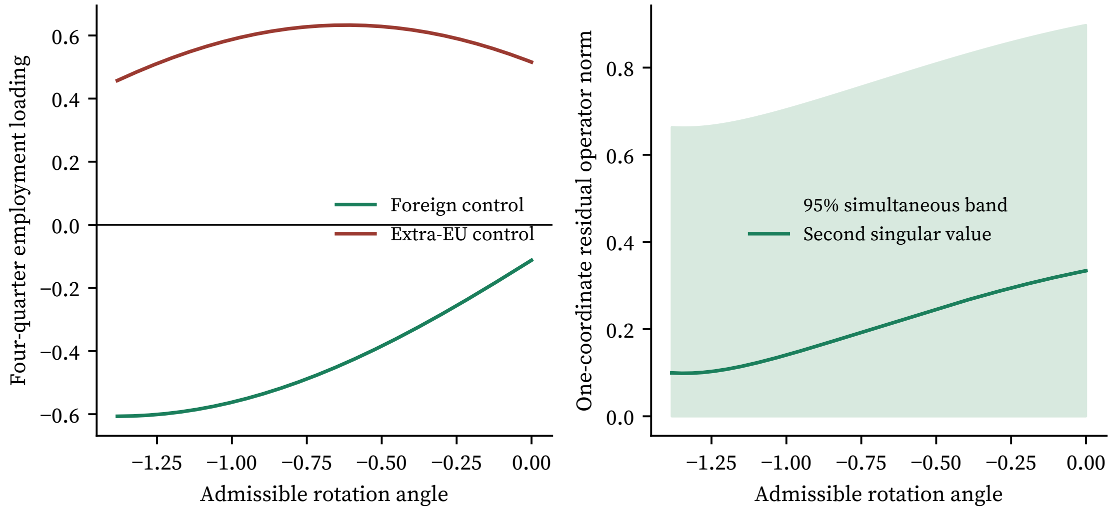

Each point uses the same sample and a one-event-standard-deviation monetary shock under that rotation. The left panel displays coefficient estimates. The right panel adds a 95 percent simultaneous spectral-error band for the twelve-row quantity system covering the continuous admissible angle interval and both singular values, conditional on the event covariance and measured FATS state. The angular grid receives an explicit Lipschitz correction.

D.6. Employment units, metrics, and curvature budgets

The employment outcome is the sum of quarterly $100\Delta \log N$ from t through $t+h$, hence $100 \log(N_{t+h}/N_{t-1})$. The unstandardized four-horizon loading matrix is

$$D_N = \begin{pmatrix} -0.01066 & 0.23821 \\ -0.13593 & 0.24441 \\ -0.13663 & 0.24124 \\ -0.46640 & 0.63131 \end{pmatrix}. \tag{A24}$$

Dividing by two implements RMS loss across four horizons. The calculation restores the raw employment units from the twelve-row system and then fixes the target metric. Policy units retain the event-standard-deviation convention of each shock decomposition. A fixed basis-point rate intervention would require its rotation-specific impact normalization and define a further policy target.

The two composite indices average standardized FATS ratios. Their units differ from percentage-point changes in F/W or E/W. Across the 222 country-years, the covariance

TABLE A11. Repair Error in Employment Units

State metric	Radius	$d = 0$	$d = 1$	Point gain	$d = 2$
Fixed index, Euclidean	0.25	0.1125	0.0165	0.0960	0
Fixed index, Euclidean	0.50	0.2251	0.0330	0.1920	0
Fixed index, Euclidean	1.00	0.4502	0.0661	0.3841	0
Within-vintage covariance ellipse	0.25	0.0603	0.0170	0.0433	0

Entries are worst-case errors in the fitted two-index conditional response projection, measured in employment RMS log percentage points at one policy unit. Point gain is the reduction from $d = 0$ to $d = 1$. At the primary Euclidean row, simultaneous operator-norm intervals are $[0.0312, 0.1939]$ for $d = 0$ and $[0, 0.0979]$ for $d = 1$; direct evaluation on aligned score draws gives $[0.0286, 0.1598]$ for the point gain. Policy sizes 0.25 and 0.5 multiply every entry by the respective factor. The $d = 2$ zero uses the identity encoder on the measured two-dimensional state. The covariance ellipse changes the allowed state directions and their scale.

TABLE A12. External Scale of the Measured Employment Frontier

Object	Log percentage points	Employment equivalents per million	95% interval
Zero-coordinate error	0.1125	1,125	[312, 1,939]
One-coordinate error	0.0165	165	[0, 979]
Zero-to-one point gain	0.0960	960	[286, 1,598]

Employment equivalents use the local derivative $N|e|/100$ at $N = 1{,}000{,}000$ manufacturing workers. Eurostat annual national-accounts employment for manufacturing supplies the external benchmark (Eurostat 2026a). The central unit monetary innovation moves the mean of the four-maturity OIS factor by 2.789 basis points across 315 complete EA-MPD events. The first two intervals are simultaneous operator-norm envelopes; the gain interval evaluates the aligned-score functional. Entries scale a fitted RMS response error and describe local response units.

after subtracting vintage means is

$$\Sigma_B = \begin{pmatrix} 0.820707 & 0.687423 \\ 0.687423 & 0.944130 \end{pmatrix}. \tag{A25}$$

The ellipsoid $z = \Sigma_B^{1/2} u$, $\|u\| \le \rho$, uses singular values of $L_N \Sigma_B^{1/2}$, which are 0.2411 and 0.0678. The corresponding twelve-row values are 0.3595 and 0.1573. Thus a singular value acquires economic meaning together with its target weights and state metric. The 178 consecutive within-country, same-vintage annual index changes have median Euclidean length 0.0896 and 90th percentile 0.2436. These distances describe observed movements, which also change national conditions; the radius grid specifies conditional index-space experiments.

For each $\rho \in \{0.25, 0.5, 1\}$ and $|\tau| \in \{0.25, 0.5, 1\}$, the release evaluates all pairs $\eta_s, \eta_p \in \{0, 0.1, 0.25, 0.5, 1\}$ with $K_s = 2\eta_s \widehat{\sigma}_{N1}$ and $K_p = 2\eta_p \widehat{\sigma}_{N1}$. At unit radius and policy size, each η is its Taylor-budget fraction of the leading projected response norm. These are sensitivity calibrations. At $(\rho, \tau) = (0.25, 1)$, state curvature $\eta_s = 0.1$ contributes 0.00281 log percentage points to the remainder; policy curvature $\eta_p = 0.1$ contributes 0.04502. Each budget enters both sides of the bound. The zero-curvature calculation is exact for the fitted linear response surface. The full-economy extension uses the following additional link.

COROLLARY A1 (Metric and projection transfer). *Let $z \in \mathbb{R}^k$ be the fiber chart of theorem 5 and let $A : \mathbb{R}^k \to \mathbb{R}^2$ be the derivative of the measured-index map on that chart. Suppose A is onto, $m^2 I_2 \le AA' \le M^2 I_2$ with $0 < m \le M$, and $\|L - \Lambda A\|_{\mathrm{op}} \le \delta$ in the stated target units. For $d < k$, with singular values of Λ padded by zeros,*

$$\text{(A26)} \qquad \begin{aligned} \max\{0, |\tau|\rho[m\sigma_{d+1}(\Lambda) - \delta] - \epsilon\} &\le e_d(\rho, \tau) \\ &\le |\tau|\rho[M\sigma_{d+1}(\Lambda) + \delta] + \epsilon. \end{aligned}$$

PROOF. The Loewner bounds imply $m^2\Lambda\Lambda' \le \Lambda AA'\Lambda' \le M^2\Lambda\Lambda'$, so eigenvalue monotonicity gives $m\sigma_j(\Lambda) \le \sigma_j(\Lambda A) \le M\sigma_j(\Lambda)$. Singular-value perturbation by a matrix of operator norm at most δ changes each singular value by at most δ. Substituting these inequalities into theorem 5 gives the result. □

The main-text frontier uses $A = I_2$ and $\delta = 0$ for its measured projection. Applying it to an economic fiber additionally assigns A, δ, K_s, and K_p; changes in the index map enter the curvature of $R(z)$ as well as its derivative. When a sampling band satisfies $\|\widehat{\Lambda} - \Lambda\|_{\mathrm{op}} \le q$, the corollary applies with the estimate and discrepancy $\delta + Mq$. If the full fiber contains further directions, the two-coordinate upper bound is $|\tau|\rho\delta + \epsilon$. Identity recovery at $d = 2$ concerns a response surface whose state is the measured two-index vector.

For the primary employment calculation, $q = 0.32537$. The measured identity projection gives the sampling-adjusted leading-response bound $\widehat{\sigma}_1 - q = 0.12478$, which becomes 0.03120 error units after multiplication by $|\tau|\rho = 0.25$. In the full-fiber form of corollary A1, positivity requires $m\widehat{\sigma}_1 - Mq - \delta > \epsilon/(|\tau|\rho)$. This inequality records the lift, sampling, projection, and curvature quantities separately. The 0.09602 point gain from one coordinate has a basic aligned-score functional interval $[0.02863, 0.15981]$ and percentile interval $[0.03222, 0.16340]$. The aligned-score functional gives the gain interval; the operator-norm interval supplies the simultaneous loss statement.

D.7. Tests of uniform signs and economically specified loss

Write $m_F(\theta) = -\delta_F(\theta)$ and $m_E(\theta) = \delta_E(\theta)$ for the four-quarter employment coefficients. The joint alternative is $\inf_{j\in\{F,E\},\theta\in\Theta} m_j(\theta) > 0$ on the closed admissible interval. Its null is the union of component nulls $m_j(\theta) \le 0$. With oriented centered wild-score errors $e_j^{(b)}(\theta)$, define

$$p_j(\theta) = \frac{1 + \sum_{b=1}^{9999} \mathbf{1}\{e_j^{(b)}(\theta) \geq \widehat{m}_j(\theta)\}}{10000}, \qquad p_{\mathrm{IUT}} = \sup_{j,\theta} p_j(\theta). \tag{A27}$$

This is an intersection–union test (Berger and Hsu 1996). If a component null holds at (j_0, θ_0), the event $p_{\mathrm{IUT}} \leq \alpha$ implies $p_{j_0}(\theta_0) \leq \alpha$. Pointwise level control therefore bounds the probability of rejecting the union null. The bootstrap approximation supplies the pointwise calibration conditional on the observed event covariance, sample, FATS indices, and regression design; the intersection–union argument preserves that calibration for the conjunction.

Each error-minus-estimate curve is $a_b \cos\theta + b_b \sin\theta$. On the interval of length 1.382489 radians, each nonzero curve has at most one interior crossing. Sweeping its crossing angles, endpoints, and ties computes the supremum of the empirical tail probability over the continuum. The foreign-negative and extra-EU-positive component suprema are 0.3679 and 0.0401, both attained at the zero-angle limit; the joint value is 0.3679. Separate component reporting gives Holm values of 0.3679 and 0.0802. These targets differ from the two-sided coefficient-and-angle confidence band in Appendix D.5.

For a projected loss tolerance $\ell \geq 0$ and $c = \ell/(|\tau|\rho)$, the null $\sigma_2(\Lambda_\theta) \leq c$ implies $\widehat{\sigma}_2(\theta) - c \leq \|\widehat{\Lambda}_\theta - \Lambda_\theta\|_{\mathrm{op}}$. The bootstrap tail of the right side yields a conservative one-sided test of loss above ℓ. This avoids regularity assumptions on a derivative of σ_2 at rank one. For the continuous-angle intersection–union value, the 2,049-point grid is supplemented, for each draw, by half the grid spacing times $[\|M_0\|_F^2 + \|M_1\|_F^2]^{1/2} + [\|E_0^{(b)}\|_F^2 + \|E_1^{(b)}\|_F^2]^{1/2}$. This bounds the angular change in the error-norm-minus-estimated-singular-value process. For the alternative prior-mean benchmark at $(\rho, \tau) = (0.5, 1)$ and $\ell = 0$, the value is 0.9775 and the continuous supremum lies between 0.9891 and 0.9894. At $\ell = 0.01$, the values are 0.9986 and 1; at $\ell = 0.05$ and 0.10, both equal 1. These projection-loss tests and the curvature scenarios answer separate statistical and finite-experiment questions.

D.8. Composite-null rank inference

The null is $\mathrm{rank}(\Lambda) \leq 1$ for the 12×2 response projection. The comparison uses the common 47-quarter support, fixed outcome scales, and 9,999 aligned quarter-wild coefficient errors. The central reference is the midpoint of the event-covariance sign interval, chosen independently of outcome coefficients.

Write $\phi_1(A) = \sigma_2(A)^2$ and let $\widehat{r}$ estimate the null rank. The analytic Chen–Fang calibration uses singular values of the bootstrap error projected onto the estimated left and right null spaces (Chen and Fang 2019). At $\widehat{r} = 0$, the draw statistic is $\sigma_2(E^{(b)})^2$; at $\widehat{r} = 1$, it is $\|U_2' E^{(b)} v_2\|^2$. This distinction carries the rank-zero stratum into the null law.

TABLE A13. Rank Inference across Methods

Calibration	Tail probability
Chen–Fang two-step, $\beta = .005$	0.3848
Chen–Fang analytic, $c = .5, 1$	0.3798
Chen–Fang analytic, $c = 2$	0.2988
Chen–Fang numerical, $c = .5$	0.3192
Chen–Fang numerical, $c = 1$	0.2852
Chen–Fang numerical, $c = 2$	0.2919
Robin–Smith, exact rank one	0.3521
KP, exact rank one	0.0132
KP intersection–union, rank at most one	0.0132

Entries use the central admissible shock reference. The first row adds the first-step allowance to the second-step probability. Robin–Smith uses Gaussian weighted-chi-square draws. All procedures condition on the event covariance and state construction.

Estimator-unit scaling cancels the common convergence-rate square between the statistic and critical values.

The two-step procedure first applies sequential KP tests at $\beta = 0.005$ and then uses the selected null-space statistic at level $0.05 - \beta$. The central reference selects rank one. Its second-step tail probability is 0.3798; adding the first-step error allowance gives 0.3848. The analytic threshold is $\kappa = cG^{-1/4}$ for $c \in \{0.5, 1, 2\}$ and $G = 47$. The numerical derivative evaluates $[\phi_1(\widehat{\Lambda} + \kappa\sqrt{G}E) - \phi_1(\widehat{\Lambda})]/(\kappa^2 G)$.

For comparison, the identity-weight characteristic-root statistic of Robin and Smith (2000) uses the eigenvalues of the projected covariance as chi-square weights. The SVD Wald statistic of Kleibergen and Paap (2006) weights the same residual by its full projected covariance and uses 11 degrees of freedom at exact rank one. An intersection–union version takes the maximum of the rank-zero and rank-one tail probabilities. These alternatives address different weighting and null-calibration choices. The primary result retains rank one; the model-independent empirical summary is a target- and metric-specific spectrum with sampling uncertainty.

At the central rotation, allowance values β of 0.001, 0.005, and 0.01 give adjusted bounds 0.3808, 0.3848, and 0.3898.

D.9. Control-index internal consistency

The joint complete-four-margin sample contains 220 country-years. Splitting employment, tangible investment, turnover, and value added into the three possible two-margin partitions gives foreign-control correlations 0.9558, 0.9585, and 0.9602. The extra-EU correlations are 0.9645, 0.9544, and 0.9680. First principal components account for 94.00 and 94.55 percent of within-coordinate margin variance. The equal-weight composite correlations with raw-share means are 0.9989 and 0.9945. The release also reports parallel-item reliability formulas as descriptive diagnostics. Common data-source errors and latent-state validity remain outside their identifying content.

D.10. Pre-announcement public-information conditioning

Following the information-conditioning argument of Bauer and Swanson (2023), the analysis projects both central-reference event shocks on a declared euro-area information set. The six predictors are the latest published monthly industrial-production growth rate, its 30-day update, the latest published unemployment rate, its 30-day update, the 30-day EUR/USD log return, and the 30-day change in the AAA ten-year spot yield. Each macro update compares the latest value available at the event cutoff with the latest value available 30 calendar days earlier; its units are percentage points. The exchange-rate change is $100\Delta\log(\mathrm{USD/EUR})$ and the yield change is in percentage points. These variables describe levels and updates of public information.

Eurostat's vintage tables record the historical website-publication date in the revision-date dimension (Eurostat 2026e). For each cutoff, the algorithm selects the latest reference month and vintage in the eligible archive. Eighteen original Eurostat releases supplement unemployment-vintage gaps during 2011, 2012, and 2014. Their release dates, headline values, reference months, and then-published euro-area compositions enter the historical information set. The maximum reference-month ages are 134 days for production and 113 days for unemployment, measured from the reference month's start. Values released on the policy-event date enter the next event's information set. Exchange-rate observations precede the event date; yield observations are dated at least three calendar days earlier, allowing for subsequent daily publication (European Central Bank 2026c,d). The observed lags are 1–2 days for exchange rates and 3–7 days for yields. Both 30-day anchors follow the same timing rule. Market histories use contemporaneous historical values; source coverage and revisions remain measurement conditions.

Every event within a retained quarter must pass the six-variable availability gate. This retains all 117 policy dates and all 47 quarters of the original 2009Q2–2020Q4 response window. Each retained event uses the corresponding publication cutoff. Same-day releases, survey-consensus forecast errors, inflation releases, and information outside these six predictors have separate coverage requirements. The conditioning set therefore defines the scope of this exercise.

Let $s_e = (\varepsilon_e^P, \varepsilon_e^I)'$ denote the fixed central admissible decomposition in its original event-standard-deviation units, and let z_e contain an intercept and the six standardized predictors. Estimate $s_e = \widehat{\Gamma}' z_e + \widetilde{s}_e$ by OLS. The residual shocks retain the original scale. For quarter t, let $Z_t = \sum_{e \in t} z_e$; its first component is the event count. Extend equation (22) by replacing both shocks with their residualized sums and adding Z_t interacted with both

TABLE A14. Public-Information Conditioning on the Common Sample

Four-quarter employment	Raw reference	News conditioned
Foreign-control loading	−0.4664	−0.3650
Quarter-wild 95% interval	[−0.945, 0.001]	[−0.676, −0.062]
HC1-scaled interval	[−1.006, 0.061]	[−0.757, 0.017]
Bartlett-4, HC1 interval	[−1.146, 0.216]	[−0.782, 0.056]
Extra-EU-control loading	0.6313	0.6639
Quarter-wild 95% interval	[0.261, 1.010]	[0.280, 1.047]
HC1-scaled interval	[0.213, 1.059]	[0.180, 1.147]
Bartlett-4, HC1 interval	[0.250, 1.018]	[0.119, 1.203]
Country-quarter observations	441	441
Regression columns	88	158

Both columns use 117 events, 47 quarters, and 12 countries at the fixed central rotation. Outcome units are cumulative employment log percentage points per original event-standard-deviation innovation and index unit. The news-conditioned column residualizes both shock components and includes all news interactions, including event count. An aligned set of 9,999 quarter-wild draws supports the common-sample comparison. HC1 multiplies each equation's score errors by $[G(n-1)/\{(G-1)(n-k)\}]^{1/2}$. The serial sensitivity uses Gaussian five-quarter moving-average multipliers, giving Bartlett covariance weights through lag four, with HC1 scaling. All intervals are pointwise and conditional on source measurement and the fixed shock decomposition.

corporate indices and all eight standardized national controls. Quarter effects absorb common news levels. For each modifier ν_{ct},

$$\nu_{ct} \sum_{e \in t} s_e = \nu_{ct} \sum_{e \in t} \widetilde{s}_e + \widehat{\Gamma}'(\nu_{ct} Z_t). \tag{A28}$$

The complete interaction span makes this a reparameterization of the regression using raw shocks and the same news controls. The first-stage projection coefficients cancel from the target estimator. All outcome equations retain the original horizon definitions and fixed outcome scales.

The monetary first stage has $R^2 = 0.0904$ and a quarter-wild joint predictor test of $p = 0.1828$; the information component gives $R^2 = 0.0413$ and $p = 0.4991$. Leave-one-calendar-year-out prediction raises monetary and information MSE by 11.61 and 11.20 percent relative to a training-sample intercept. The sample orthogonality identity is an algebraic property of projection. Its causal content depends on the specified information set and the maintained exclusion.

The paired changes in the four-quarter loadings are 0.1014 and 0.0326, with aligned quarter-wild intervals [−0.252, 0.461] and [−0.407, 0.470]. News conditioning retains the negative and positive point-estimate pair. Degrees-of-freedom and serial-correlation adjustments leave the extra-EU interval above zero; the corresponding foreign-control intervals include zero. These adjustments are reported alongside the declared quarter-wild calculation.

The news-conditioned employment singular values are 0.4819 and 0.0641 in the RMS metric, giving the radius-0.25 fitted frontier $(0.1205, 0.0160, 0)$. Its second direction accounts for 1.74 percent of squared response norm. The one-coordinate loss has a quarter-wild spectral interval $[0, 0.0873]$ and a Bartlett-4, HC1 interval $[0, 0.1137]$. Employment-only Chen–Fang adjusted bounds are 0.5059 and 0.7251 under these two calibrations. These quantities describe the measured two-index projection with $A = I_2$ and discrepancy $\delta = 0$; the full-fiber transfer conditions in Appendix D.6 remain separate.

The twelve-row quantity projection gives singular values 0.6069 and 0.3689, and a second-direction share of 26.98 percent. For this conditional system, the Chen–Fang two-step procedure exits at the preliminary rank-above-one branch at allowance $\beta = 0.005$, under both score calibrations. The 0.005 bound denotes that branch. The associated unweighted one-coordinate loss intervals include zero. Covariance weighting, target choice, and sampling calibration continue to enter the rank interpretation.

D.11. Announcement windows and policy composition

EA-MPD measures the press release from a 13:25–13:35 lower median to a 14:00–14:15 upper median, the conference from a 14:15–14:25 lower median to a 15:40–15:50 upper median, and the monetary event from the first lower median to the final upper median (Altavilla et al. 2019). The specification uses all 117 events and 47 quarters in the response sample. For each window, it forms covariance principal components from the four-maturity basis $(1, 3, 6, 12)$ months and the five-maturity basis that adds the two-year OIS. Each rate factor and STOXX50 surprise receives the same two-dimensional sign decomposition as the main shock. The midpoint of each admissible interval is determined from event covariance and impact signs.

The joint-window equation includes release monetary and information shocks and conference monetary and information shocks. Every shock is interacted with both corporate indices and the eight national controls. Corporate-state levels, standardized national-control levels, country effects, and calendar-quarter effects complete the equation. The policy-composition equation uses the five-maturity monetary-event shock and adds three event-level components. The path component is the first principal component of two- and three-year OIS changes. The long-rate component uses five- and ten-year German yields. The sovereign-spread component uses Italian and Spanish spreads over German yields at two, five, and ten years. Sequential projections on the monetary and information shocks and preceding components produce unit-event-standard-deviation residual components. A final equation residualizes all five shocks on the six pre-announcement predictors and includes their quarterly interactions with corporate and national states.

TABLE A15. ECB Treatment Composition and Four-Quarter Employment

Shock and conditioning set	Foreign control	Extra-EU control	Wild/Bartlett p
Monetary event, OIS through one year	−0.466 [−0.934, −0.008]	0.631 [0.251, 1.010]	0.0018/0.0033
Monetary event, OIS through two years	−0.469 [−0.926, −0.014]	0.614 [0.230, 0.998]	0.0023/0.0055
Policy release, joint-window equation	−0.460 [−0.832, −0.096]	0.586 [0.028, 1.161]	0.0117/0.0082
Conference, joint-window equation	−0.426 [−0.849, 0.002]	0.518 [−0.100, 1.137]	0.0885/0.1466
Monetary event plus three components	−0.705 [−1.207, −0.215]	0.494 [0.058, 0.920]	0.0084/0.0593
Components plus pre-announcement news	−0.404 [−0.795, −0.017]	0.628 [0.147, 1.107]	0.0028/0.0612

Coefficient brackets are pointwise 95 percent basic intervals from 9,999 aligned Rademacher calendar-quarter score draws. The final column reports quarter-wild/Bartlett-4 joint two-coordinate quadratic-form tails. Bartlett-4 values use five-quarter Gaussian moving-average multipliers and HC1 scaling. For the policy-release row, Bartlett-4 coefficient intervals are [−0.874, −0.038] and [0.005, 1.170]. For the final row they are [−1.024, 0.224] and [−0.059, 1.323]. Every equation contains shock interactions with the two FATS indices and eight national controls.

The four- and five-maturity monetary-event estimates are nearly unchanged. In the joint equation, the release component carries the more precise response pair. Path, long-rate, and sovereign-spread interactions have intervals spanning zero in the final equation, with respective Bartlett-4 joint values 0.2364, 0.8343, and 0.7996. The point pair remains in the declared quadrants after the full conditioning set, while serial-dependence calibration governs its joint precision. These results refine the assigned policy innovation; the corporate indices continue to parameterize its conditional response surface.

Appendix E. Numerical Economy and Calibration

The empirical comparison vectors in the numerical appendices retain their stated legacy-shock benchmarks. All collision, loss, and regret values refer to the model-generated economy.

E.1. Households and production

Country c has a representative household with preferences

$$\mathbb{E}_0 \sum_{t=0}^{\infty} \beta^t [u(C_{ct}) - v(N_{ct})], \qquad 0 < \beta < 1. \tag{A29}$$

Utility is increasing and concave, labor disutility is increasing and convex, and admissible paths satisfy the integrability and transversality conditions required for a finite objective. Markets clear at the national and international margins specified in the baseline.

Firm i operates local units with production

$$y_{ict} = A_{ict} z_{it}^{\nu} k_{ict}^{\alpha} n_{ict}^{1-\alpha}, \qquad 0 < \alpha < 1, \quad \nu \geq 0, \tag{A30}$$

where z_{it} is a firm-wide intangible or organizational state. Local tangible capital and the shared state evolve according to

$$\begin{aligned} k_{ic,t+1} &= (1-\delta_k)k_{ict} + I_{ict}, \\ z_{i,t+1} &= (1-\delta_z)z_{it} + \Phi\left(\sum_c RD_{ict}, \chi_{it}\right), \end{aligned} \tag{A31}$$

with $0 \le \delta_k, \delta_z \le 1$ and nonnegative continuously differentiable Φ.

E.2. Internal finance and corporate perimeter

Each group can transfer funds τ_{ict} across controlled affiliates subject to

$$\sum_c \tau_{ict} = 0, \qquad \Psi_i(\tau_{it}) = \frac{\psi_i}{2}\sum_c \tau_{ict}^2, \qquad \psi_i \ge 0. \tag{A32}$$

External borrowing pays

$$r_{ict}^b = r_t^{ECB} + \phi_c(\ell_{ict}, X_{ct}). \tag{A33}$$

The ownership map Ω_t assigns local units to corporate groups. Changing Ω_t changes which liquidity constraints, shared states, and product interactions are internalized. Two economies can therefore share the same national distributions of k, n, y, and I and have different control states,

$$B_t = B(\Omega_t, z_t, \Gamma_t). \tag{A34}$$

E.3. Strategic interaction and equilibrium

The industry block is chosen before quantitative policy validation. Differentiated-product demand provides strategic pricing when product-market data support it; investment or capacity interaction provides a dynamic alternative. The selected block maps the empirical common-control, shared-resource, and strategic coordinates into the Jacobian terms of equation (6).

The reported public-data exercise uses two actions and maps the fixed country-year indices into $f = \{1+\exp(-B^F)\}^{-1}$ and $e = \{1+\exp(-B^E)\}^{-1}$. For $\lambda, \psi, \gamma \ge 0$, define

$$\begin{aligned} d_1(f) &= 1 + \lambda(0.5+f), \qquad d_2(e) = 1 + \psi(0.5+e), \\ q(f,e,p) &= \min\{\lambda f + \gamma(1+0.08|p|), 0.9\sqrt{d_1(f)d_2(e)}\}, \\ J(f,e,p) &= \begin{bmatrix} d_1(f) & -q(f,e,p) \\ -q(f,e,p) & d_2(e) \end{bmatrix}. \end{aligned} \tag{A35}$$

TABLE A16. Moment Mapping and Normalizations

Parameter block	Target moments	Source	Status
Production and capital	$\beta = 0.96$, $\alpha = 0.33$, $\delta_k = 0.08$	Declared normalizations	fixed
Corporate geography	$\lambda = 0.360$; mean foreign-control share	FATS, 1,055 cells	calibrated
Internal finance	$\psi = 0.442$; conditional geography $g^{E\|F} = E/F$	FATS, 895 cells	calibrated
Strategic interaction	$\gamma = 0.156$; $0.2\|\operatorname{corr}(B^F, B^E)\|$	FATS state, 222 rows	calibrated
Policy shock process	315 event dates and three windows	ECB EA-MPD	evaluation

The curvature targets use unit-interval FATS cells. The two state coordinates have correlation 0.780. Preference, production, and depreciation entries are declared normalizations for the broader dynamic environment; equation (A35) supplies the reported collision computations. Policy-response estimates remain in the validation set.

Let $Open_+ = \max\{Open, 0\}$, $MVA_+ = \max\{MVA, 0\}$, $c = [1 + 0.05\log(1 + Open_+),\ 0.5 + 0.5e]'$, and $w = MVA_+/(MVA_+ + 0.15)$. The equilibrium solves $F(a; f, e, p) = J(f, e, p)a + pc = 0$, and the target is $T(f, e, p) = [w, 1 - w]a + \psi ep$. The final term maps extra-EU internal-finance capacity into the geographic buffer of equation (7). Since $d_1, d_2 \geq 1$ and $q^2 \leq 0.81 d_1 d_2$, $\det J = d_1 d_2 - q^2 \geq 0.19 d_1 d_2 > 0$. Sylvester's criterion gives positive definiteness, so the equilibrium is unique and equals $a = -J^{-1}pc$.

Equilibrium is computed under the same ownership map used to construct the state. The determinant bound establishes uniqueness for every evaluated state and policy. At the reported response point $p = 1$, the uncapped branch lies below the cap by at least 0.830246 across all 222 states, so the evaluated derivatives stay within one C^1 branch.

E.4. Calibration and validation

The reported two-action calibration uses three FATS moments: the mean foreign-control share, the conditional extra-EU share $g^{E|F} = E/F$ among valid foreign-controlled cells, and the cross-country correlation of the two state indices. The conditional geography moment calibrates an internal-finance channel parameter; the empirical state coordinate remains the extra-EU share E/W standardized within vintage. Preference, production, and depreciation values normalize the surrounding dynamic environment. The employment loading pair, policy-invariance loss, collision responses, ECB direct-investment flows, Bank Lending Survey responses, and outward-FATS affiliate geography stay outside the numerical objective and retain their source-specific coverage.

The saturated empirical vector has norm 0.8571 and the model derivative vector has norm 0.0842, giving a norm ratio of 0.0983. The least-squares common scale mapping from model units to the saturated empirical vector is 9.443, after which the directional

TABLE A17. Response Geometry Outside the Calibration Objective

Object	Foreign-control loading	Extra-EU loading	Direction ratio
Saturated employment estimate	−0.563	0.646	1.147
Parsimonious employment reference	−0.347	0.600	1.729
Calibrated-economy mean derivative	−0.0277	0.0795	2.870

Direction ratio is the extra-EU loading divided by the absolute foreign-control loading. The cosine between the saturated empirical and model loading vectors is 0.928; the parsimonious-reference cosine is 0.982. Negative foreign-control and positive extra-EU derivatives occur in all 222 model states. FATS cross-sectional moments determine the numerical parameters; employment-response coefficients serve as evaluation moments.

residual equals 37.3 percent of the empirical norm. These quantities define the magnitude discrepancy alongside the directional comparison. Collision, error, repair, and regret calculations below use outcomes generated by equation (A35) and retain model-objective units.

E.5. Aggregate collisions

Within the calibrated stationary distribution or simulated state space, choose pairs (S^A, S^B) satisfying

$$d_X(X(S^A), X(S^B)) \leq \epsilon_X \tag{A36}$$

while $d_B(B(S^A), B(S^B))$ exceeds a declared control-distance threshold. The matching metric, tolerance, and control-distance threshold are declared before policy outcomes are computed. Each pair receives the same policy intervention. Collision search uses states and predetermined model primitives; target divergence is computed afterward.

Appendix F. Model-Internal Counterfactual Error, State Repair, and Policy Loss

F.1. Counterfactual divergence

For collision pair (A, B) and horizon h, define

$$CD_h(p; A, B) = T_{p,h}(S^A) - T_{p,h}(S^B), \qquad CDI_h = \frac{|CD_h|}{\sigma(Y_h)}, \tag{A37}$$

For $\sigma(Y_h) > 0$, CDI_h expresses the response gap in standard deviations of Y_h; CD_h reports the same gap in the units of Y_h.

FIGURE A4. Aggregate Collisions and Hidden Counterfactual Divergence

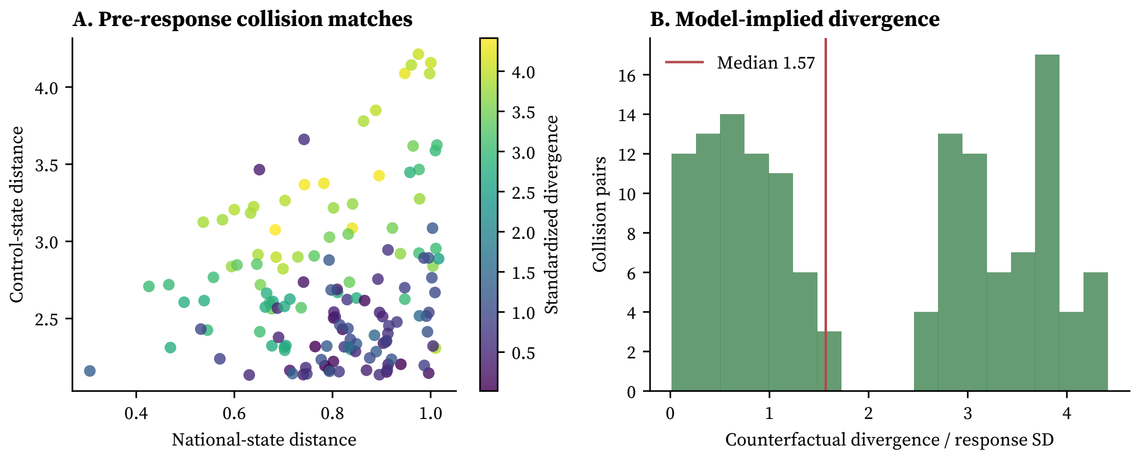

Collision matching uses trade openness, manufacturing value-added share, and log GDP in the national metric and the two FATS indices in the control metric. The rule retains national distances below 1.022 and control distances above 2.131 before computing the policy response.

F.2. Counterfactual aggregation error

Let T_p^{full} be the model counterfactual, T_p^X the macro-state prediction, and $T_p^{X,B}$ the control-augmented prediction. For symmetric positive-definite W, define

$$
\begin{aligned}
CAE_X &= \|T_p^{full} - T_p^X\|_W, \\
CAE_{X,B} &= \|T_p^{full} - T_p^{X,B}\|_W. \qquad \text{(A38)}
\end{aligned}
$$

When $CAE_X > 0$, the repair share is

$$RepairShare = 1 - \frac{CAE_{X,B}}{CAE_X}. \qquad \text{(A39)}$$

The interval $[0, 1]$ corresponds to $0 \le CAE_{X,B} \le CAE_X$, and values below zero correspond to $CAE_{X,B} > CAE_X$. The reporting sample requires baseline error above a declared threshold.

F.3. Policy invariance

Let $\widehat{T}_1^X(S)$ and $\widehat{T}_1^{X,B}(S)$ be response predictors estimated at the unit policy magnitude. For a validation set $\mathcal{V}$ and a policy-magnitude set $\mathcal{P}_V$, define

$$
\begin{aligned}
PI_X &= \frac{1}{|\mathcal{V}||\mathcal{P}_V|} \sum_{S\in\mathcal{V}} \sum_{p\in\mathcal{P}_V} \left| T_p^{full}(S) - p\widehat{T}_1^X(S) \right|, \\
PI_{X,B} &= \frac{1}{|\mathcal{V}||\mathcal{P}_V|} \sum_{S\in\mathcal{V}} \sum_{p\in\mathcal{P}_V} \left| T_p^{full}(S) - p\widehat{T}_1^{X,B}(S) \right|. \qquad \text{(A40)}
\end{aligned}
$$

FIGURE A5. State Repair, Policy Invariance, and Policy Loss

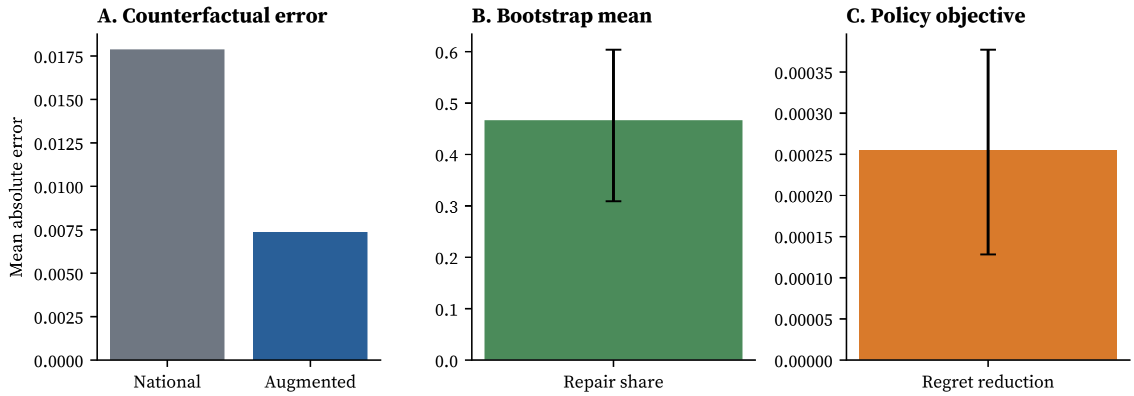

Panel A evaluates 81 states from 2019 onward. Panels B and C report means and 95 percent state-resampling intervals from 999 draws. The repair-share reporting sample contains 72 states above the tenth percentile of national-state error.

The difference $PI_X - PI_{X,B}$ measures the reduction in response-prediction loss when one mapping trained at $p = 1$ travels to the declared policy magnitudes. The reported set is $\mathcal{P}_V = \{0.75, 1.25\}$.

F.4. Policy regret

For a nonempty feasible policy set with an attained optimum, let $p^*(S) \in \arg\max_p V(p; S)$. Let $\widehat{p}_X$ and $\widehat{p}_{X,B}$ be policies chosen from the macro-state and augmented states. Regret is

$$(A41)\qquad Regret_X = V(p^*; S) - V(\widehat{p}_X; S), \qquad Regret_{X,B} = V(p^*; S) - V(\widehat{p}_{X,B}; S).$$

Optimality makes both objects nonnegative. The welfare reporting sample comprises policy applications with an explicitly specified objective and aggregation error above a declared threshold.

The collision pairs have median national distance 0.818 and median control distance 2.608. Their median model response gap is 1.568 response standard deviations. For the 2019–2023 validation states, adding both control coordinates lowers mean model-internal counterfactual error from 0.0179 to 0.0073. The mean repair share is 0.466 with a state-resampling interval of $[0.309, 0.604]$, while its median is 0.651. Applying the unit-policy mapping at policy magnitudes 0.75 and 1.25 lowers mean absolute response error by 0.0104. Mean regret falls by 0.000255 with interval $[0.000128, 0.000377]$, equal to 74.4 percent of mean national-state regret and 0.0453 percent of the squared response target. The oracle-policy match rate rises from 51.9 to 84.0 percent; regret reductions are positive in 35.8 percent, zero in 59.3 percent, and negative in 4.9 percent of model states.

TABLE A18. Model-Internal Aggregation and Repair

Object	Mean	Median	P90	Mean 95% interval
National-state distance d_X	0.805	0.818	0.996	[0.783, 0.828]
Control-state distance d_B	2.713	2.608	3.428	[2.636, 2.798]
Counterfactual divergence *CDI*	2.038	1.568	3.896	[1.818, 2.274]
CAE_X	0.0179	0.0141	0.0370	[0.0149, 0.0208]
$CAE_{X,B}$	0.0073	0.0055	0.0140	[0.0055, 0.0098]
Repair share	0.466	0.651	0.916	[0.309, 0.604]
Policy-invariance loss reduction	0.0104	0.0104	0.0104	[0.0077, 0.0132]
Regret reduction	0.000255	0.000000	0.001126	[0.000128, 0.000377]

Distances and divergence use 140 collision pairs. Counterfactual errors and regret use 81 validation states. Mean intervals resample the corresponding state or collision rows. Policy-invariance loss applies the $p = 1$ mapping at policy magnitudes 0.75 and 1.25. The regret objective targets an outcome response of −0.75 with policy cost 0.08 over 61 policy values from −1.5 to 1.5.

Appendix G. Quantitative Details

G.1. Equilibrium definition

The computation maps the F/W and E/W control indices into $f = \{1 + \exp(-B^F)\}^{-1}$ and $e = \{1+\exp(-B^E)\}^{-1}$. The equilibrium is $Ja + pc = 0$, with J from equation (A35). Its target is $T = [w, 1-w]a + \psi ep$. The calibrated values are $\lambda = 0.359977$, $\psi = 0.441927$, and $\gamma = 0.156098$; their targets use 1,055 valid foreign-control cells, 895 conditional geography cells satisfying $W > 0$, $F > 0$, and $0 \le E \le F$, and the absolute correlation 0.780488 between the two indices. Across the 895 cells, the identity $E/F = (E/W)/(F/W)$ has maximum absolute numerical error 1.06×10^{-15}. Policy-response estimates enter the response-geometry evaluation in table A17.

G.2. Collision algorithm

The collision search standardizes the national-state metric using pre-policy scales, imposes common support, and matches on state coordinates and predetermined primitives. Counterfactual targets are computed after pair construction. Tolerances ϵ_X and the minimum d_B are set before policy simulation. Results are reported over all admissible pairs and under tighter tolerances.

G.3. Channel and validation moments

The numerical calibration uses the three FATS moments listed in table A16. Policy-response coefficients, balance-of-payments flows, lending standards, and outward-affiliate geography form evaluation objects. Each object keeps its native unit and coverage,

TABLE A19. Channel and Validation Evidence

Validation block	Open-data object	Result
Internal group funding	ECB BPS debt assets and liabilities	$h = 4$: $(0.879, -1.338)$; $p = 0.291$
Enterprise credit supply	ECB Bank Lending Survey	block statistic 12.46; $p = 0.0701$
Credit quantities and rates	ECB BSI and MIR	statistic 65.32; $p = 0.0005$
Host-country reallocation	Eurostat inward FATS	statistic 205.38; $p = 0.0005$
Affiliate geography	Eurostat outward FATS	2023 medians: $(0.593, 0.493)$
Employment response direction	Eurostat employment and EA-MPD	saturated cosine 0.928

The BPS row uses 619 country-quarter observations, 19 countries, and 43 calendar-quarter clusters; at one quarter its loading pair is $(0.984, -1.450)$ with joint $p = 0.073$. The BLS row uses 502 observations, 16 countries, 42 calendar-quarter clusters, and 9,999 aligned wild-cluster draws. Credit-rate and reallocation p-values use 1,999 aligned draws and equal the 0.0005 reporting floor. The outward-FATS pair reports the extra-EU shares of manufacturing employment and gross tangible investment; published-cell counts are 16 and 13. Every row remains outside the numerical calibration objective.

so the evidence enters as coefficient vectors, joint response tests, or distributional shares according to the source design.

Direct-investment group funding is $F^{grp}_{ct} = (DI^{debt,L}_{ct} - DI^{debt,A}_{ct})/GDP_{c,t-1}$. The BPS sample begins in 2013Q1 and reports the all-entity series, winsorized flows, the S1P sector scope, and the all-entity series net of resident special-purpose entities. Enterprise credit supply is the ECB Bank Lending Survey diffusion index. Its outcome is the change in credit standards from $t-1$ through $t+h$, with positive values denoting tightening; the regression uses the state and shock interactions from equation (22), a lagged standards level, country effects, and calendar-quarter effects. At one quarter, the BLS loading pair is $(1.830, -1.791)$ and the horizon-specific centered-bootstrap p-value is 0.1327. The four-horizon common-sample statistic is 12.46 with centered-bootstrap $p = 0.0701$ and asymptotic diagnostic 0.1319.

For outward-FATS metric k, affiliate geography is $s^{k,ext}_{cy} = k^{ext}_{cy}/(k^{int}_{cy} + k^{ext}_{cy})$. In 2023, the median extra-EU share of manufacturing employment is 0.593 across 16 euro-area controlling countries, and the median share of gross tangible investment is 0.493 across 13 countries. These outward shares measure the geographic span of decision perimeters, while the inward-FATS response block measures reallocation among host-country control categories.

G.4. Finite-policy curvature in the numerical economy

The example uses the cross-effect $q(p) = \lambda f + \gamma(1 + 0.08p)$ on its uncapped positive-policy branch. Hence J_p has off-diagonal entries $-\eta$, where $\eta = 0.08\gamma = 0.012488$, and the equilibrium response has policy curvature. For $p \in [0.75, 1.25]$, $f, e \in (0,1)$, and trade openness up to the largest observed state value, set $\bar{q} = \lambda + 1.1\gamma$ and $\kappa = [1 + \frac{1}{2}\min\{\lambda, \psi\} - \bar{q}]^{-1}$. Diagonal dominance gives $\|J^{-1}\| \leq \kappa = 1.542487$. The lower cap margin is 0.548586, so this whole domain stays inside the smooth branch.

Writing $F(a, p) = J(p)a + pc$, the second policy derivative is $a_{pp} = 2J^{-1}J_pJ^{-1}(J_pa + c)$. The aggregation row has norm at most one, and $\|c\| \leq \bar{c} = 1.465877$ on the stated domain. Therefore an admissible bound is $K_p = 2\kappa^2\eta(1 + 1.25\eta\kappa)\bar{c}$, evaluated with unrounded constants and rounded upward as 0.08920502. For displacement $\tau = p - 1$ with $|\tau| \leq 0.25$, the policy-only Taylor remainder is at most 0.002788 model units. Across the evaluated states, the maximum observed policy-only remainder is 0.000737.

For state curvature, fix trade openness and aggregation weights and use Euclidean changes in the two indices. A policy-independent level offset can hold the baseline allocation fixed while preserving every policy increment. Write $V = J^{-1}$ and $J_p = \left(\begin{smallmatrix} 0 & -\eta \\ -\eta & 0 \end{smallmatrix}\right)$. At $p_0 = 1$, the scalar policy response is $R = -MVc + MVJ_pVc + \psi e$. This expression includes the control derivatives of incidence $c = (1 + 0.05\log(1 + \text{trade}), 0.5 + 0.5e)'$. Logistic first and second derivatives have absolute bounds 1/4 and $1/(6\sqrt{3})$. With $\varphi = (1 + \sqrt{5})/2$, valid derivative bounds are

$$J_1 = \sqrt{(\varphi\lambda/4)^2 + (\psi/4)^2}, \quad J_2 = \frac{\max\{\varphi\lambda, \psi\}}{6\sqrt{3}}, \quad V_1 = \kappa^2 J_1, \quad V_2 = 2\kappa^3 J_1^2 + \kappa^2 J_2. \tag{A42}$$

The first bound uses the norm φ of $\left(\begin{smallmatrix} 1 & -1 \\ -1 & 0 \end{smallmatrix}\right)$ and Cauchy–Schwarz over the two index directions; the second uses their separate logistic Hessians. Set $c_1 = 1/8$ and $c_2 = 1/(12\sqrt{3})$. Applying the product rule twice to R gives

$$\begin{aligned} K_s = {} & V_2\bar{c} + 2V_1c_1 + \kappa c_2 + \frac{\psi}{6\sqrt{3}} \\ & + \eta\{(2\kappa V_2 + 2V_1^2)\bar{c} + 4\kappa V_1 c_1 + \kappa^2 c_2\}. \end{aligned} \tag{A43}$$

The unrounded calculation is 0.81432458, so $K_s = 0.814325$ is a uniform upper bound on this fiber. The largest Hessian norm at the 222 observed states is 0.035183. Across the evaluated state-policy changes with radii 0.05, 0.1, 0.25, and 0.5 and $\tau = \pm 0.25$, the largest joint Taylor remainder is 24.24 percent of $(|\tau|K_s\rho^2 + K_p\tau^2)/2$. Both constants remain in the model's outcome and policy units; the employment-unit sensitivity analysis supplies its own budgets.

References

Carlo Altavilla, Luca Brugnolini, Refet S. Gürkaynak, Roberto Motto, and Giuseppe Ragusa. Measuring euro area monetary policy. *Journal of Monetary Economics*, 108:162–179, 2019.

Jonas E. Arias, Juan F. Rubio-Ramírez, and Daniel F. Waggoner. Inference based on structural vector autoregressions identified with sign and zero restrictions: Theory and applications. *Econometrica*, 86(2):685–720, 2018. doi: 10.3982/ECTA14468.

Michael D. Bauer and Eric T. Swanson. A reassessment of monetary policy surprises and high-frequency identification. *NBER Macroeconomics Annual*, 37:87–155, 2023. doi: 10.1086/723574.

Roger L. Berger and Jason C. Hsu. Bioequivalence trials, intersection–union tests and equivalence confidence sets. *Statistical Science*, 11(4):283–302, 1996. doi: 10.1214/ss/1032280304.

Chunya Bu, John Rogers, and Wenbin Wu. A unified measure of fed monetary policy shocks. *Journal of Monetary Economics*, 118:331–349, 2021.

Qihui Chen and Zheng Fang. Improved inference on the rank of a matrix. *Quantitative Economics*, 10(4):1787–1824, 2019. doi: 10.3982/QE1139.

Stephan Dahlke, Erich Novak, and Winfried Sickel. Optimal approximation of elliptic problems by linear and nonlinear mappings I. *Journal of Complexity*, 22(1):29–49, 2006. doi: 10.1016/j.jco.2005.06.005.

Ronald A. DeVore, Ralph Howard, and Charles Micchelli. Optimal nonlinear approximation. *Manuscripta Mathematica*, 63:469–478, 1989.

European Central Bank. Bank lending survey statistics: Enterprise credit standards, 2026a. ECB Data Portal BLS national diffusion indices and survey methodology, accessed August 26, 2026.

European Central Bank. Balance of payments and international investment position data: Direct-investment debt instruments, 2026b. ECB Data Portal BPS country-quarter transactions and methodological definitions, accessed August 26, 2026.

European Central Bank. Euro foreign exchange reference rates, 2026c. URL `https://www.ecb.europa.eu/stats/policy_and_exchange_rates/euro_reference_exchange_rates/html/index.en.html`. Daily rates and publication schedule, accessed August 28.

European Central Bank. Euro area yield curves, 2026d. URL `https://www.ecb.europa.eu/stats/financial_markets_and_interest_rates/euro_area_yield_curves/html/index.en.html`. AAA spot rates and publication schedule, accessed August 28.

European Commission, Directorate-General for Competition. Merger cases search and merger statistics, 2026. Competition Policy database, accessed August 2026.

European Securities and Markets Authority. Electronic reporting, 2026. URL `https://www.esma.europa.eu/issuer-disclosure/electronic-reporting`. ESEF scope and requirements, accessed August 28, 2026.

Eurostat. National accounts employment data by industry, 2026a. URL `https://ec.europa.eu/eurostat/databrowser/view/nama_10_a10_e/default/table`. Dataset `nama_10_a10_e`, domestic-concept manufacturing employment in thousands of persons, API snapshot retrieved August 30, 2026.

Eurostat. Foreign affiliates statistics: Inward and outward fats, 2026b. Legacy FATS 2008–2020 and European Business Statistics FATS from 2021 onward, API snapshots retrieved August 25, 2026.

Eurostat. Labour input in industry: Quarterly data, 2026c. URL https://ec.europa.eu/eurostat/databrowser/view/sts_inlb_q/default/table. Dataset sts_inlb_q, NACE Rev. 2 division employment indices, API snapshots retrieved August 30, 2026.

Eurostat. Outward foreign affiliates statistics by controlling and affiliate economy, 2026d. Legacy table fats_out2_r2 and EBS table fats_out_activ, accessed August 26, 2026.

Eurostat. Principal european economic indicators: Vintage data, 2026e. URL https://ec.europa.eu/eurostat/cache/metadata/en/euroind_vtg_esms.htm. Metadata and historical publication vintages, accessed August 28.

Gen-Probe Incorporated. Quarterly report for the period ended march 31, 2010, 2010. Form 10-Q, SEC accession 0000950123-10-044493, filed May 5, 2010.

Global Legal Entity Identifier Foundation. Level 2 data: Who owns whom, 2026. Relationship records and reporting exceptions, snapshot retrieved August 25, 2026.

Frank Kleibergen and Richard Paap. Generalized reduced rank tests using the singular value decomposition. *Journal of Econometrics*, 133(1):97–126, 2006. doi: 10.1016/j.jeconom.2005.02.011.

Jean-Marc Robin and Richard J. Smith. Tests of rank. *Econometric Theory*, 16(2):151–175, 2000. doi: 10.1017/S0266466600162012.

U.S. Securities and Exchange Commission. Additional form 8-k disclosure requirements and acceleration of filing date, 2004. Item 2.01 completion disclosure and significance rule.

U.S. Securities and Exchange Commission. Edgar application programming interfaces, 2026. Submissions, Companyfacts, and XBRL Frames interfaces, accessed August 26, 2026.

XBRL International. Filings.xbrl.org: Esef filing repository and api, 2026. Repository snapshot retrieved August 25, 2026.