Too Big to Aggregate

Corporate Control and Policy-Invariant Aggregation

Chae-Yeon Xon

Yonsei University

September 2026

Abstract

National accounts locate economic activity, while corporate control links the decisions through which it responds to policy. I ask which control information must accompany a supplied national state to preserve a declared policy response. The rank of the response along omitted state directions gives the minimum local augmentation. Its singular values measure the first-order cost of a smaller augmentation, with curvature bounds for finite experiments. In euro-area manufacturing, ECB monetary responses vary with lagged corporate-control indices. For a quarter-unit index radius and a one-standard-deviation monetary innovation, omitting both measured coordinates gives a fitted worst-case employment error of 0.1125 log percentage points across four horizons, with a simultaneous 95 percent interval of $[0.0312, 0.1939]$. One coordinate lowers the point error to 0.0165. The value of the second coordinate increases when credit and investment enter the target. These estimates measure response heterogeneity across corporate states; identifying the effect of changing control requires separate variation.

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.

1. Introduction

National accounts assign economic activity to the country where it takes place. Corporate groups make decisions across the activities they control. These two boundaries can describe the same current allocation while organizing its response to policy in different ways. Two economies with the same reported national state may therefore react differently to the same intervention. The state directions that generate this difference determine what must be added to the national representation and how costly a smaller addition would be.

The candidate missing state is the corporate decision perimeter. Common control links establishments through internal funding, shared adjustment costs, and joint allocation, while national accounts continue to record their activity by location. The economic unit is the set of activities whose choices enter a common decision problem. The paper asks how much information about monetary responses is left out when a national state omits measured control geography, and how many control coordinates are needed to recover a declared target.

Consider two countries with the same employment, capital, production, and establishments, but with different corporate boundaries. A firm operating across countries can reallocate funds and share adjustment costs across its activities. A monetary contraction can then reach identical local balance sheets through different joint decisions. Common-control curvature, shared inputs, and strategic feedback enter the inverse equilibrium Jacobian that maps policy incidence into national outcomes.

The empirical analysis pairs lagged Eurostat FATS control indices with intraday ECB surprises and conditions on eight national variables and their monetary and information interactions.[1] The central admissible rotation gives four-quarter employment loadings $(-0.466, 0.631)$.[2] The estimand is heterogeneity in monetary responses across measured corporate states. Assignment of corporate control requires a separate design; within-country and issuer evidence describes related variation. Following Chen and Fang (2019), the rank analysis allows for the rank-zero stratum.[3] The comparison across methods yields a repair frontier reported with sampling uncertainty.

For a quarter-unit index radius and a one-event-standard-deviation monetary innovation,[4] the fitted employment RMS frontier is $(0.1125, 0.0165, 0)$ when zero, one, and two coordinates are retained. The simultaneous operator-norm interval for the zero-coordinate loss is $[0.0312, 0.1939]$. Retaining the first coordinate lowers the point loss

[1] FATS assigns control to the ultimate controlling institutional unit. The two indices summarize foreign-controlled and extra-EU-controlled shares of local manufacturing activity.

[2] The central rotation is the midpoint of the rate–equity sign-restricted interval; the empirical analysis also evaluates the full admissible set.

[3] At rank zero, singular-vector directions are unidentified; the reported procedure treats that stratum separately.

[4] The radius is measured in the standardized two-index state metric; the innovation is scaled to unit variance across monetary-policy events.

by 0.0960, or 85.3 percent. Direct evaluation on the aligned score draws gives a 95 percent basic interval of $[0.0286, 0.1598]$ for this reduction and a remaining-error interval of $[0, 0.0979]$. When credit and investment enter with fixed outcome-standard-deviation weights, the share attributed to the second direction rises from 2.11 to 14.37 percent. The usefulness of an additional coordinate therefore depends on the target it is meant to preserve.

The question is one of invariance in the sense of Marschak (1953); Hurwicz (1962); Lucas (1976): which relationships remain usable under a stated class of interventions? Marschak's Maxim directs identification toward the objects needed for the policy question (Heckman 2010). Here the object supplied by existing institutions is a state representation q_R. The question is which intervention-relevant information that representation preserves. The policy, outcomes, horizons, and base state together define the target $T_{\mathcal{P}}$ and, with it, the information required for the exercise.

The analysis begins with the hidden-response operator

$$\mathcal{L}_{\mathcal{P}} = DT_{\mathcal{P}}|_{\ker Dq_R}.$$

Policy-invariant local aggregation holds if and only if the national-state kernel lies inside the policy-target kernel. The rank of the hidden-response operator is the minimum augmentation dimension, and a smooth coordinate map attains that bound locally. Theorem 5 then uses classical manifold-width geometry (DeVore et al. 1989; Dahlke et al. 2006) to obtain the first-order compression cost $|\tau|\rho\sigma_{d+1}$ when d continuous coordinates are retained.[5] The policy displacement τ, state radius ρ, target weights, and state metric give this cost its economic units. Curvature and projection-discrepancy bounds carry the calculation to a finite policy experiment. These results form a characterization, an attained minimum-state theorem, and a finite compression frontier for the same declared target.

The closest comparison comes from the aggregation literature. Hulten's first-order result uses equilibrium sales shares to aggregate microeconomic productivity shocks, and Baqaee and Farhi (2019) show how networks, substitution, returns to scale, and reallocation enter the second-order effect through changing sales shares and general-equilibrium elasticities. Krusell and Smith (1998) establish approximate aggregation using the mean of the wealth distribution, while Den Haan and Rendahl (2010) connect individual policy rules to the aggregate state variables they require. An et al. (2009); Chang et al. (2013) study representative-agent descriptions and policy-dependent aggregate parameters. Koby and Wolf (2020) express general-equilibrium smoothing through the rank deficiency and null space of a dynamic demand–supply elasticity map, which determines which

[5] The singular-value term is local and target specific. The finite-experiment bound also carries the declared state radius, policy displacement, curvature allowance, and projection discrepancy.

distribution-induced demand shifts survive price adjustment. The analysis here starts from an existing statistical representation and restricts attention to the directions in $\ker Dq_R$. The declared response target then determines the dimension and loss frontier of the repair.

A related literature asks which information is sufficient for policy evaluation. Forecasts and impulse responses enter the policy assessment of Barnichon and Mesters (2023), and causal monetary responses support the near-term counterfactuals of Caravello et al. (2026). The present question is which current-state information preserves those response functions. Sequence-space Jacobians provide a macroeconomic language for equilibrium derivatives (Auclert et al. 2021); the concern here is the state information on which those derivatives depend.

Studies of multinational measurement reassign financial positions or profits to trace their aggregate consequences (Coppola et al. 2021; Guvenen et al. 2022). My exercise holds the current allocation fixed and lets the decision perimeter change its response. Evidence on internal funding and international transmission provides the empirical basis for this mechanism (Biermann and Huber 2024; Cetorelli and Goldberg 2012; Zhang et al. 2025). The FATS indices describe the geography of corporate control, while the numerical economy shows how joint decisions generate the hidden-response directions.

2. Corporate Control and the Aggregation Problem

2.1. States, activities, and national statistics

There are C countries and N corporate groups. Group i controls a vector $a_i = (a_{i1}, \dots, a_{iC})'$ of local pricing, production, capacity, investment, financing, or sourcing decisions. The full state $S_t \in \mathcal{S}$ records technologies, local stocks, ownership links, corporate geography, financial positions, shared inputs, and strategic connections. On the neighborhoods considered below, the state space $\mathcal{S}$ and the national-state space $\mathcal{X}$ are finite-dimensional C^2 manifolds. Existing statistical institutions provide the map

$$X_t = q_R(S_t) \in \mathcal{X}. \tag{1}$$

Thus X_t contains the information available in a residence- or location-based macroeconomic representation.

Contemporaneous national outcomes $Y_t \in \mathbb{R}^{d_Y}$ take the local form

$$Y_t = Ma_t + y(X_t), \tag{2}$$

where the stacked action has dimension m and $M \in \mathbb{R}^{d_Y \times m}$ is fixed by the accounting convention. Dynamic targets propagate through

$$X_{t+1} = \mathcal{G}(X_t, Y_t, \xi_{t+1}). \tag{3}$$

For a policy p and horizon H, $T_p(S_t)$ stacks the conditional paths of the declared outcomes. The policy family and target vector are fixed before estimation.

2.2. The corporate decision problem

Firm i maximizes

$$\Pi_i(a_i, a_{-i}; S, p) = R_i(a_i, a_{-i}; S, p) - \frac{1}{2}a_i' D_i a_i - \frac{1}{2}\sum_{k=1}^{r_i} \lambda_{ik}(g_{ik}' a_i)^2. \tag{4}$$

The matrix D_i is symmetric positive semidefinite and $\lambda_{ik} \geq 0$. Each loading g_{ik} joins activities that use a common intangible, managerial capacity, internal liquidity, product architecture, data infrastructure, or another firm-wide resource. Reassigning a local activity changes which cross-country terms enter the owner's decision problem, even when current local capital, employment, output, and sales remain fixed.

Stack equilibrium conditions as

$$F(a; S, p) = 0, \qquad J(S, p) = D_a F(a; S, p). \tag{5}$$

The equilibrium Jacobian can be written as

$$J = H^L + H^C + H^G - \Gamma. \tag{6}$$

The block H^L collects local curvature, H^C common-control effects, H^G shared-input and reallocation effects, and Γ strategic cross-firm feedback. These blocks are determined by ownership, operating, and network objects measured before the policy exercise.

2.3. Three organizational channels

Common control. An owner internalizes the effect of one affiliate's decision on the returns to other controlled activities. A control change alters the off-diagonal decision curvature even when every activity remains in the same physical location.

Shared inputs and reallocation. A firm can move liquidity, managerial attention, intangible capital, R&D knowledge, capacity, or product design across its controlled locations. A policy change can therefore move local activities through an internal shadow price omitted from the national-accounts state.

Strategic propagation. Competitors respond to the combined product and geographic footprint of the corporate group. A change in control changes the best-response system faced by rivals and can propagate an initially local policy exposure through markets or networks.

These three channels describe how corporate control can change the equilibrium response operator. For the declared policy family, the national aggregate remains sufficient when the derivatives of all three channel blocks vanish along the hidden control directions.

2.4. Currency-area exposure and geographic buffers

Let B^F measure the foreign-controlled share of world-classified local activity and let B^E measure the extra-EU-controlled share of world-classified local activity. Both enter the empirical implementation as within-vintage standardized indices averaged across four activity margins. A contractionary monetary innovation $\varepsilon_t > 0$ raises the affiliate financing wedge by $(1 + \alpha_H B^F - \nu_H B^E)\varepsilon_t$, where $\alpha_H > 0$ captures transmission through a group-wide euro balance sheet and $\nu_H > 0$ captures access to cash flow and financing outside the currency area. Let cumulative labor adjustment solve

$$\min_{\Delta n_{t:t+H}} \left\{ \frac{(\Delta n_{t:t+H})^2}{2\kappa_H} + \left[\ell_{0H} + (1 + \alpha_H B^F - \nu_H B^E)\varepsilon_t\right] \Delta n_{t:t+H} \right\}, \qquad \kappa_H > 0. \tag{7}$$

PROPOSITION 1 (Two-coordinate employment transmission). *The solution to equation* (7) *satisfies*

$$\frac{\partial^2 \Delta n_{t:t+H}}{\partial \varepsilon_t \partial B^F} = -\kappa_H \alpha_H < 0, \qquad \frac{\partial^2 \Delta n_{t:t+H}}{\partial \varepsilon_t \partial B^E} = \kappa_H \nu_H > 0. \tag{8}$$

The pair (B^F, B^E) *spans the financing block whenever the two empirical loadings are linearly independent across declared outcomes or horizons.*

The first derivative captures amplification through the part of the corporate perimeter exposed to the common policy innovation. The second captures insulation through resources whose cash flows and financing conditions extend beyond the currency area. Equation (8) gives sign restrictions for the primary employment response. The general Jacobian system allows common-input and strategic channels to operate alongside this financing mechanism.

FIGURE 1. Same National State, Different Corporate Control

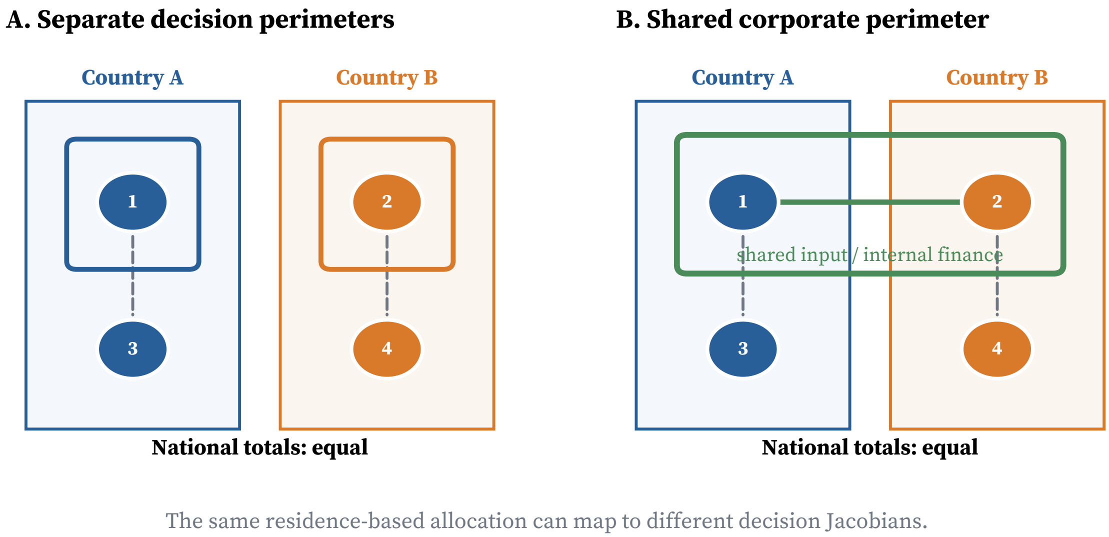

Each panel holds the residence-based allocation fixed. The green perimeter in panel B joins activities 1 and 2 inside one decision problem and introduces a shared-input or internal-finance link.

3. A Theory of Policy-Valid Aggregation

3.1. Regularity and hidden control perturbations

Fix a baseline (a_0, s_0, p_0) satisfying $F(a_0; s_0, p_0) = 0$. The analysis uses three benchmark assumptions.

ASSUMPTION 1 (Smooth regular equilibrium). *$F : \mathbb{R}^m \times \mathcal{S} \times \mathbb{R}^q \to \mathbb{R}^m$ is C^2 near the baseline and $J_0 = D_a F(a_0; s_0, p_0)$ is nonsingular. The local equilibrium supplied by the implicit-function theorem is denoted $a(s, p)$, and $J(s, p) := D_a F(a(s, p); s, p)$.*

ASSUMPTION 2 (Institutionally fixed aggregation). *q_R is C^1 and has constant rank near s_0. The accounting map M in equation (2) is fixed on that neighborhood.*

ASSUMPTION 3 (Pure control perturbations). *There is a linear space $\mathcal{V}_C(s_0) \subseteq \ker Dq_R(s_0)$ such that every $h \in \mathcal{V}_C(s_0)$ satisfies $D_h M = 0$, $D_h F_p = 0$, and $D_h H^L = 0$. The perturbation may change H^C, H^G, and Γ.*

The state-direction derivatives of J, M, and F_p below are total derivatives evaluated along $a(s, p)$. Assumption 3 isolates propagation through the organization of the firm. If accounting weights or direct policy incidence also vary with control, the response derivative contains the corresponding additional terms. A direct target component $G(S, p)$ contributes $D_h G_p$.

The implicit-function theorem gives

$$D_p a = -J^{-1}F_p, \qquad R_p = D_p Y = -MJ^{-1}F_p. \tag{9}$$

THEOREM 1 (Control transmission). *Under assumptions 1–3, every hidden control perturbation $h \in \mathcal{V}_C(s_0)$ changes the policy response by*

$$\begin{aligned} D_h R_p &= MJ^{-1}(D_h J)J^{-1}F_p \\ &= MJ^{-1}D_h H^C J^{-1}F_p + MJ^{-1}D_h H^G J^{-1}F_p - MJ^{-1}D_h \Gamma J^{-1}F_p. \end{aligned} \tag{10}$$

For a finite menu of policy directions $v_1, \dots, v_K \in \mathbb{R}^q$, define the baseline incidence matrix

$$P = [F_p v_1, \dots, F_p v_K].$$

The entire family is locally insensitive to h if and only if $MJ^{-1}(D_h J)J^{-1}P = 0$.

The three factors in equation (10) trace the path from a change in organization to a change in the national response. The rightmost factor selects the decision margins moved by policy. The middle factor changes the equilibrium operator, and the leftmost factor records what reaches the aggregate target. A hidden control change matters for the target when it survives each of these maps.

PROPOSITION 2 (Sufficient conditions for aggregation). *If $D_h H^C = D_h H^G = D_h \Gamma = 0$ for every $h \in \mathcal{V}_C(s_0)$, then $D_h R_p = 0$ for every policy in the declared family.*

3.2. Primitive response dimension

To relate this derivative to a finite set of measurable coordinates, suppose that organizational effects are generated by a fixed number of primitives.

ASSUMPTION 4 (Finite control-factor structure). *There are fixed matrices $B_1, \dots, B_R \in \mathbb{R}^{m \times m}$ and predetermined linear functionals $\ell_r : \mathcal{V}_C(s_0) \to \mathbb{R}$ such that*

$$D_h J = \sum_{r=1}^{R} \ell_r(h) B_r. \tag{11}$$

Let $L_C(h) = (\ell_1(h), \dots, \ell_R(h))'$ and define $A_{\mathcal{P}}$ by stacking $\mathrm{vec}(MJ^{-1}B_r J^{-1}P)$ as columns. Thus $A_{\mathcal{P}} : \mathbb{R}^R \to \mathbb{R}^{d_Y K}$, $L_C : \mathcal{V}_C(s_0) \to \mathbb{R}^R$, and the vectorized hidden response $\mathcal{H}_{\mathcal{P}} : \mathcal{V}_C(s_0) \to \mathbb{R}^{d_Y K}$ have fixed domains and codomains.

THEOREM 2 (Primitive hidden-response dimension). *The hidden response factors as*

$$\mathcal{H}_{\mathcal{P}} = A_{\mathcal{P}} L_C, \tag{12}$$

and

$$\operatorname{rank}(\mathcal{H}_{\mathcal{P}}) \le \min\{\operatorname{rank}(A_{\mathcal{P}}), \operatorname{rank}(L_C)\} \le R. \tag{13}$$

If $A_{\mathcal{P}}$ has full column rank R and L_C is onto $\mathbb{R}^R$, then $\operatorname{rank}(\mathcal{H}_{\mathcal{P}}) = R$.

Assumption 4 is a maintained restriction on the economic model. The resulting upper bound applies to the control subspace specified there. Extending that bound to the full national kernel requires the spanning and exclusion conditions in corollary 1. An empirical matrix with two columns has rank at most two by construction, so its column count describes the dimension of the measured projection. An economic claim about the full kernel requires the additional conditions.

3.3. Closure and minimum repair

Let $T_{\mathcal{P}} : \mathcal{S} \to \mathbb{R}^{d_T}$ collect the declared target vector and define the national vertical space

$$\mathcal{V}_R(s_0) = \ker Dq_R(s_0). \tag{14}$$

DEFINITION 1 (First-order counterfactual closure). *The national representation is first-order closed for $T_{\mathcal{P}}$ at s_0 if there exists a linear map $A : T_{x_0}\mathcal{X} \to \mathbb{R}^{d_T}$ such that $DT_{\mathcal{P}}(s_0) = ADq_R(s_0)$.*

THEOREM 3 (Policy-invariant aggregation characterization). *Suppose q_R and $T_{\mathcal{P}}$ are C^1 and q_R has constant rank near s_0. First-order closure holds at s_0 if and only if*

$$\ker Dq_R(s_0) \subseteq \ker DT_{\mathcal{P}}(s_0). \tag{15}$$

If the inclusion holds throughout a sufficiently small constant-rank neighborhood U with connected local vertical sections, there is a C^1 map $\overline{T}_{\mathcal{P}} : q_R(U) \to \mathbb{R}^{d_T}$, where $q_R(U)$ carries its constant-rank image-manifold structure, such that $T_{\mathcal{P}}|_U = \overline{T}_{\mathcal{P}} \circ q_R|_U$.

The derivative condition is pointwise. The constant-rank neighborhood in the second part specifies where it can be integrated into a local factorization. Finite repair asks how much error remains under finite state and policy displacements when the representation is approximately repaired.

Define the hidden-response operator and its rank by

$$\mathcal{L}_{\mathcal{P}}(s_0) = DT_{\mathcal{P}}(s_0)|_{\mathcal{V}_R(s_0)}, \qquad r_{\mathcal{P}}(s_0) = \operatorname{rank} \mathcal{L}_{\mathcal{P}}(s_0). \tag{16}$$

An augmentation $b : \mathcal{S} \to \mathbb{R}^d$ is a first-order repair if (q_R, b) satisfies the kernel criterion.

THEOREM 4 (Minimum state repair). *Every d-dimensional C^1 first-order repair satisfies*

$$d \geq r_{\mathcal{P}}(s_0). \tag{17}$$

Conversely, there is a C^1 augmentation b^* defined on a neighborhood of s_0 and valued in $\mathbb{R}^{r_{\mathcal{P}}(s_0)}$ whose differential at s_0 separates exactly the quotient $\mathcal{V}_R(s_0)/\ker \mathcal{L}_{\mathcal{P}}(s_0)$ and makes (q_R, b^*) first-order closed at s_0.

COROLLARY 1 (Primitive rank-dimensional repair). *Suppose $\mathcal{H}_{\mathcal{P}}$ is a coordinate block of $DT_{\mathcal{P}}(s_0)|_{\mathcal{V}_C(s_0)}$ and has rank R. Every repair then has dimension at least R. This bound is attained by a predetermined C^1 control map $b_0 : \mathcal{S} \to \mathbb{R}^R$ when there is a direct-sum decomposition $\mathcal{V}_R(s_0) = \mathcal{V}_C(s_0) \oplus \mathcal{V}_O(s_0)$ satisfying $Db_0|_{\mathcal{V}_C} = L_C$, $Db_0|_{\mathcal{V}_O} = 0$, $DT_{\mathcal{P}}|_{\mathcal{V}_O} = 0$, and $\ker L_C \subseteq \ker(DT_{\mathcal{P}}|_{\mathcal{V}_C})$.*

The proofs use kernel inclusion and rank–nullity, but the objects entering those arguments are economic. The map q_R is fixed by the statistical institution, $T_{\mathcal{P}}$ is fixed by the policy exercise, and the corporate primitives determine how hidden control directions change J^{-1}. A repair built directly from the response coordinates uses the target derivative itself.[6] Interpreting measured control coordinates as that repair requires the spanning conditions in corollary 1.

3.4. Approximate repair and finite policy changes

Suppose the policymaker can add d coordinates to the reported state. The remaining question is how much counterfactual accuracy those coordinates can preserve. Fix a constant-rank chart $s(z)$ of a national fiber, with $q_R(s(z)) = x_0$ and $z \in \overline{B}_\rho \subset \mathbb{R}^k$. This chart fixes the units in which hidden-state variation is measured. Let $Y(z, \tau)$ collect the declared outcomes and horizons under a scalar policy displacement τ along a fixed direction, and define the policy increment $G(z, \tau) = Y(z, \tau) - Y(z, 0)$. Write $R(z) = \partial_\tau Y(z, 0)$ and $L = D_z R(0)$. On this fiber, L is the hidden-response operator for the target $T_{\mathcal{P}} = R$. Equip the state and target spaces with specified Euclidean norms, and order the singular values of L as $\sigma_1(L) \geq \cdots \geq \sigma_k(L) \geq 0$, padding with zeros where needed.

ASSUMPTION 5 (Uniform finite-experiment regularity). *The fiber chart and Y extend to an open neighborhood of $\overline{B}_\rho \times [-\bar{\tau}, \bar{\tau}]$. The response R is twice continuously differentiable, Y is twice continuously differentiable in τ, and throughout that set $\|D_z^2 R(z)\| \leq K_s$ and $\|\partial_\tau^2 Y(z, \tau)\| \leq K_p$. The first bound is the bilinear operator norm.*

[6] Such a response-coordinate repair proves attainability within the smooth representation class. Treating an observed corporate-control index as the attaining record is an additional measurement claim, governed here by the primitive spanning conditions.

For $0 \le d < k$, let a continuous state encoder $b : \overline{B}_\rho \to \mathbb{R}^d$ supply the additional coordinates, and let g be any decoder. Define the best attainable error at policy displacement τ by

$$(18) \qquad e_d(\rho, \tau) = \inf_{b \text{ continuous}, g} \sup_{\|z\| \le \rho} \|G(z, \tau) - g(b(z), \tau)\|, \qquad \epsilon(\rho, \tau) = \frac{|\tau| K_s \rho^2 + K_p \tau^2}{2}.$$

THEOREM 5 (Policy-response compression frontier). *Under assumption 5, for $\rho > 0$ and $|\tau| \le \bar{\tau}$,*

$$(19) \qquad \max\{0, |\tau|\rho\sigma_{d+1}(L) - \epsilon(\rho, \tau)\} \le e_d(\rho, \tau) \le |\tau|\rho\sigma_{d+1}(L) + \epsilon(\rho, \tau).$$

The upper bound is attained as a bound by the leading right-singular coordinates $b_d(z) = (v_1' z, \ldots, v_d' z)$ and decoder $g_d(b_d(z), \tau) = \tau R(0) + \tau \sum_{j=1}^d \sigma_j u_j v_j' z$. For the first-order experiment $G(z, \tau) = \tau[R(0) + Lz]$, the equality $e_d = |\tau|\rho\sigma_{d+1}(L)$ holds even over continuous nonlinear encoders and arbitrary decoders. For $d \ge k$, an identity encoder and the exact response decoder give $e_d = 0$.

The first-order equality is the finite-dimensional Hilbert-space case of nonlinear and manifold widths (DeVore et al. 1989; Dahlke et al. 2006). The Bernstein-width lower bound uses antipodal coincidence (Borsuk 1933), while singular-value truncation attains the upper bound (Eckart and Young 1936). Standard manifold widths require both the encoder and decoder to be continuous. Here the lower bound uses continuity of the encoder, and the linear decoder that attains the bound is itself continuous, so the two decoder classes have the same value on this ellipsoid.

In the economic application, this width measures response information discarded by q_R. The radius ρ, policy displacement τ, state and target metrics, and curvature budget determine its units and the range over which the finite bound is informative. For $\tau \ne 0$, the lower bound is positive when $\sigma_{d+1} > K_s\rho/2 + K_p|\tau|/(2\rho)$. Baseline allocations enter separately through $Y(z, 0)$.

The curvature constants can also be read through the corporate decision perimeter. On a smooth fiber where $R(z) = -MJ(z)^{-1}P$, suppose M and P are fixed and $\|J^{-1}\| \le \kappa$, $\|D_z J\| \le J_1$, and $\|D_z^2 J\| \le J_2$. Then one may take

$$(20) \qquad K_s = \|M\|\|P\|(2\kappa^3 J_1^2 + \kappa^2 J_2).$$

The inverse equilibrium Jacobian amplifies the hidden first-order response as well as the finite-state remainder. Policy-dependent strategic feedback also enters K_p, as shown in the numerical illustration. The frontier is measured in employment units. Corollary A1 carries the bound from a measured-index map to an economic fiber by accounting for metric distortion and response-projection discrepancy.

3.5. Policy menus and target scope

Adding policies or horizons stacks rows in the hidden-response system, so the rank can only increase. Each rank statement therefore refers to a particular policy, set of outcomes, horizons, sample window, and primitive upper bound. A state that is sufficient for one monetary experiment may require more coordinates for a family that also includes fiscal, regulatory, or cross-border policies. Firm-level response images that point in different directions can raise the dimension of the aggregate repair. If all firms share an r-dimensional response image, the aggregate rank remains bounded by r.

If q_R assigns the same reported state to s and s' while $\|T_{\mathcal{P}}(s) - T_{\mathcal{P}}(s')\| = \Delta > 0$, any deterministic predictor based on that report has worst-case error of at least $\Delta/2$. The quantitative collision exercise expresses this separation in economic units.

4. Measuring Corporate Control

The application studies euro-area manufacturing under ECB monetary innovations. Eurostat inward FATS assigns the residence of control to the ultimate controlling institutional unit (Eurostat 2026). Control is defined through authority over an affiliate's general policy, including the appointment of directors. This measure describes control geography; internal liquidity sharing and operational integration require separate evidence. For each manufacturing margin m, let W_{mct} denote world-classified local activity, F_{mct} foreign-controlled activity, and E_{mct} extra-EU-controlled activity. The primitive shares are

$$s^F_{mct} = F_{mct}/W_{mct}, \qquad s^E_{mct} = E_{mct}/W_{mct}. \tag{21}$$

Employment, tangible investment, turnover, and value added each provide one margin. The two indices average the corresponding within-vintage standardized shares and require at least three margins in the closed unit interval. A policy observation in year t is paired with the index from $t-1$. The primary window, 2009Q2–2020Q4, uses the legacy production system throughout. The data record includes the EBS extension, source flags, and transformation dictionary.

The state file contains 222 country-years across the two production vintages, and the two indices have a correlation of 0.780.[7] Among the 220 observations with all four margins in both coordinates, the three split-half correlations range from 0.954 to 0.968. First-principal-component scores correlate above 0.9999 with the equal-weight composites.

[7]The primary response window ends in 2020Q4 so that every included state coordinate comes from the legacy FATS production system. The later production vintage is used in the state-file comparison.

These calculations describe internal consistency. They leave room for source errors shared across the four margins. Margin-exclusion estimates use the alternative shock benchmark. The F/W and E/W indices measure activity-weighted control geography. The conditional share $g^{E|F} = E/F$, where $E/F = (E/W)/(F/W)$, is used only as a separate illustration moment.

The monetary data contain 315 EA-MPD event dates from January 1999 through October 2025 (Altavilla et al. 2019). EA-MPD reports a policy-release window, a press-conference window, and the union of the two. This union is the main treatment.[8] A short-rate factor formed from one-, three-, six-, and twelve-month OIS changes explains 82.70 percent of complete-event variance. Contractionary monetary news raises rates and lowers equities, while information news raises both (Jarociński and Karadi 2020). Orthogonal sign-restricted components follow the construction in Arias et al. (2018). The reference component uses the midpoint angle −0.691244 in the admissible interval $(-1.382489, 0)$ determined by event covariance and the sign restrictions. The shocks have unit event variance, and quarterly shocks sum the event contributions. The analysis evaluates the full admissible set and labels the alternative standardized prior-mean benchmark used in robustness specifications.

5. Do National Aggregates Close Monetary Policy Responses?

5.1. National benchmark state

The national benchmark contains eight quarter-end variables: lagged manufacturing-production growth, headline inflation, energy inflation, released GDP growth, NFC credit growth, the NFC lending rate, trade openness, and the manufacturing value-added share. Each enters in levels and in interactions with the monetary and information innovations. Country and calendar-quarter effects complete the benchmark. Write $q_N(S_{ct})$ for this observed representation. The regression asks whether the two corporate coordinates add information to the conditional response projection generated by q_N.

In the theory, q_R denotes the national state available to the statistical institution. The empirical projection based on q_N transfers to that broader object when changes in the indices admit a lift into $\ker Dq_R$ and the response variation associated with national coordinates omitted from q_N is bounded by the projection discrepancy in corollary A1.[9] The candidate corporate state consists of the one-year-lagged foreign-control and extra-EU-control indices. The primary window ends in 2020Q4, which keeps every state coordinate within the legacy FATS production system.

[8] Separate-window specifications retain the release and conference components and are reported with their common samples.

[9] The empirical q_N contains the eight national controls and their stated interactions. The broader q_R is the institutional state in the theory.

5.2. Conditional response closure

For country c, quarter t, and horizon h, estimate

$$
\begin{aligned}
\Delta_h n_{ct} = {} & \alpha_{c,h} + \tau_{t,h} + \sum_{k=1}^{8} \left(\theta_{kh} X_{kct} + \omega_{kh} X_{kct} \varepsilon_t^p + \zeta_{kh} X_{kct} \varepsilon_t^I \right) + g_h(B_{c,t-1}) \\
& + \delta_{Fh} B_{c,t-1}^F \varepsilon_t^p + \delta_{Eh} B_{c,t-1}^E \varepsilon_t^p \\
& + \rho_{Fh} B_{c,t-1}^F \varepsilon_t^I + \rho_{Eh} B_{c,t-1}^E \varepsilon_t^I + \eta_{c,t+h}.
\end{aligned}
\tag{22}
$$

Here $\Delta_h n_{ct}$ is cumulative manufacturing-employment log growth from t through $t + h$, ε_t^p is the contractionary monetary component, and ε_t^I is the information component. Calendar-quarter effects absorb the common shock level. The state-dependent monetary response is $(\delta_{Fh}, \delta_{Eh})$, and proposition 1 predicts $\delta_{Fh} < 0$ and $\delta_{Eh} > 0$. The horizon-specific restriction is

$$
H_0: \quad \delta_{Fh} = \delta_{Eh} = 0. \tag{23}
$$

Rejecting equation (23) means that the state directions summarized by the two corporate coordinates shift the declared policy response conditional on the national benchmark. The intraday policy window provides temporal assignment. Predetermined state coordinates, national-state shock interactions, and information-shock interactions form the conditioning structure. Under the event-timing exclusion, sign-decomposition, and common-shock assumptions, ε_t^p is the assigned treatment and $(\delta_{Fh}, \delta_{Eh})$ traces its conditional response surface across corporate states.

At four quarters, a one-unit increase in the foreign-control index lowers the response to a unit monetary innovation by 0.466 log percentage points. The extra-EU loading is 0.631. Their pointwise intervals are $[-0.939, -0.005]$ and $[0.257, 1.002]$, with signs matching equation (8). Alternative shock and national-state specifications accompany the full admissible-set analysis, where point estimates and uniform sign inference are kept separate.

Two checks examine what the intraday innovation contains. Conditioning on six public-information variables available before the announcement preserves the directions of the four-quarter loadings, although the foreign-control interval includes zero. Separating the policy-release and press-conference windows places the stronger joint evidence in the release window. The employment frontier remains similar when the public-information set is added. These comparisons address predictable news and the composition of the announcement; common samples, loading estimates, intervals, and joint tests are reported for each window and maturity basis. The event-window exclusion continues to carry the causal interpretation.

FIGURE 2. Corporate Control and the Employment Response

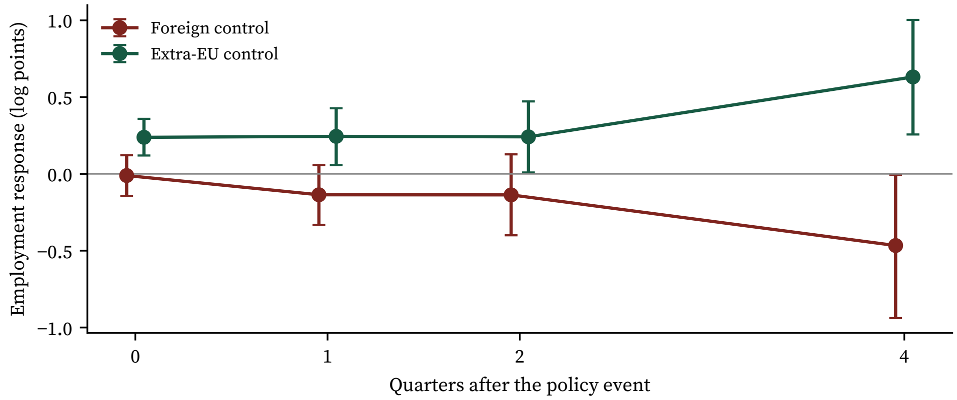

Points use the central admissible monetary decomposition on the common 47-quarter support of the twelve-row system. Bars are pointwise 95 percent basic intervals from 9,999 aligned wild calendar-quarter draws. Every equation includes the eight national variables in levels and interacted with both shocks. Employment is cumulative $100\Delta \log N$ from quarter t through horizon h. These intervals condition on the measured state and event covariance.

5.3. The identified response contrast

The assigned treatment is the intraday monetary innovation. Its causal interpretation rests on the event-window exclusion, the stated rate–equity restrictions, and orthogonality to country-specific outcome innovations conditional on the benchmark. The coefficients describe how the response to this treatment varies across lagged FATS states. This is a test of policy-relevant state representation with corporate control used as a measured candidate coordinate. Persistent differences in national transmission can also contribute to the response surface. Estimating the causal effect of assigning corporate control would require a separate estimand and design.

An industry bridge compares 24 manufacturing divisions exposed to the same national policy environment and weights each cell by preceding-year employment.[10] At impact, the foreign-control and extra-EU loading pair is $(-0.133, 0.104)$. The simultaneous intervals over the eight horizon-specific loadings exclude zero, while the joint-zero and joint one-sided sign values are 0.0001 and 0.0002. The supporting specifications include the full horizon path, the unweighted design, country-history reassignment, national within-country estimates, and affiliate-data requirements. Control assignment remains distinct from accompanying financing and accounting changes.

[10]The weighting fixes the industry composition before the contemporaneous monetary event; the robustness exercise uses equal division weights.

5.4. The cost of compressing the response state

Let $\ell_h = 100\log(N_{t+h}/N_{t-1})$ for $h \in \{0,1,2,4\}$ and measure error by employment RMS loss, $\|\ell\|_{\mathrm{RMS}} = (\sum_h \ell_h^2/4)^{1/2}$. One log percentage point equals 0.01 in log employment. Stack the four unstandardized corporate-state loading rows in D_N and set $L_N = D_N/2$. The hidden state is $z = (\Delta B^F, \Delta B^E)'$ in fixed index units, while τ is an event-standard-deviation innovation at the central admissible rotation. The fitted increment is $\widehat{G}_N(z,\tau) = \tau[\widehat{R}_N(0) + L_N z]$. These choices fix the metric and scale of the approximation exercise.

The singular values of L_N are 0.4502 and 0.0661. The corresponding coordinates are $b_1 = -0.5470\Delta B^F + 0.8371\Delta B^E$ and $b_2 = 0.8371\Delta B^F + 0.5470\Delta B^E$. The first preserves 97.89 percent of the fitted squared employment-response norm. With d retained coordinates, the fitted error is $|\tau|\rho\sigma_{d+1}$; the error is zero when both coordinates are retained. At $(\rho,\tau) = (0.25, 1)$, the first coordinate lowers employment RMS error from 0.1125 to 0.0165 log percentage points. This is a reduction of 0.0960, or 85.3 percent.

The reference radius is close to the 90th percentile, 0.2436, of 178 observed same-vintage annual index changes. Because these movements also change national conditions, their scale serves only to calibrate a conditional exercise in index space. Radii 0.5 and 1 provide stress scenarios. Under the covariance ellipse $z = \Sigma_B^{1/2}u$, $\|u\| \le 0.25$, the frontier is $(0.0603, 0.0170, 0)$. The reported frontier therefore depends on the state metric as well as the outcome weights.

Across 315 complete monetary-event windows in EA-MPD, a central monetary innovation of one event standard deviation moves the average factor formed from the 1-month, 3-month, 6-month, and 1-year OIS rates by 2.789 basis points. The maturity-specific impacts range from 2.568 to 3.029 basis points. At a benchmark of one million manufacturing workers, the local log-to-level derivative converts the zero-coordinate point error to 1,125 employment equivalents, with interval [312, 1,939]. The one-coordinate error corresponds to 165 workers, with interval [0, 979], and the point reduction corresponds to 960, with aligned-score interval [286, 1,598]. These conversions give external units to the fitted four-horizon RMS error. A realized employment calculation would also require a state displacement, a signed response path over horizons, and the employment base at each horizon.

For a finite response map with derivative L_N at the reference state, the same radius and policy size give

$$\max\{0, \widehat{e}_1^{\mathrm{lin}} - 0.03125K_s - 0.5K_p\} \le e_1 \le \widehat{e}_1^{\mathrm{lin}} + 0.03125K_s + 0.5K_p, \tag{24}$$

where $\widehat{e}_1^{\mathrm{lin}} = 0.25\widehat{\sigma}_{N2} \simeq 0.016523$. The units of K_s are employment RMS log points per squared index unit per policy unit; those of K_p are employment RMS log points per squared policy unit. The scenario $K_s = 0.09003$, $K_p = 0$ gives $[0.0137, 0.0193]$. Metric distortion and

FIGURE 3. Employment-Unit Repair Frontiers

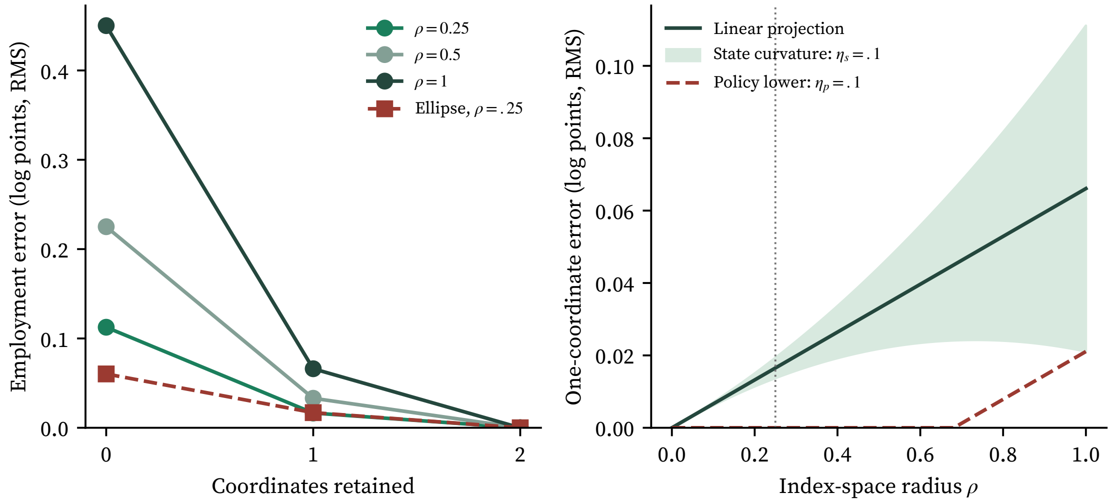

The central admissible shock has unit event variance. Left: fitted frontiers for the reference radius 0.25, stress radii 0.5 and 1, and the covariance ellipse at radius 0.25. Right: one-coordinate linear error, a maintained state-curvature interval with $K_s = 0.2\widehat{\sigma}_{N1}$ and $K_p = 0$, and a policy-curvature lower bound with $K_s = 0$ and $K_p = 0.2\widehat{\sigma}_{N1}$. The vertical line marks the reference radius. These curvature scenarios are separate from sampling intervals.

discrepancy between the measured projection and an economic fiber enter as additional maintained quantities in the finite-experiment interpretation.

The aligned score draws give a 95 percent operator-norm radius of 0.3254. Weyl's bound then yields simultaneous intervals of $[0.0312, 0.1939]$ for zero-coordinate loss and $[0, 0.0979]$ for one-coordinate loss. The first interval places a positive lower bound on response loss within the measured two-index projection. Evaluating the singular-value gap directly on the same draws gives a basic interval of $[0.0286, 0.1598]$ for the point reduction of 0.0960. This functional interval is reported alongside the operator-norm envelope. Combining sampling error with the continuous shock set gives $[0, 0.1406]$ for one-coordinate loss. Curvature budgets, sampling bands, shock-set variation, and transfer from the measured projection to an economic fiber address separate parts of the exercise.

Corollary A1 states the transfer condition. With sampling radius q, the full-fiber zero-coordinate lower bound is positive when

$$m\widehat{\sigma}_1 - Mq - \delta > \frac{\epsilon}{|\tau|\rho}. \tag{25}$$

For the identity measured projection, $m = M = 1$, $\delta = 0$, and $\epsilon = 0$. The remaining margin is $\widehat{\sigma}_1 - q = 0.1248$ response units, which equals the 0.0312 lower endpoint after multiplication by the radius 0.25. An economy-wide statement must specify the lift constants, projection discrepancy, and curvature remainder in equation (25). The fitted projection provides the empirical component of this inequality.

Adding credit and investment at the same horizons with fixed outcome-standard-deviation weights changes the singular values to 0.5011 and 0.2053. The second direction now accounts for 14.37 percent of squared response norm, compared with 2.11 percent for employment alone. This remains a two-column projection with its corresponding rank ceiling. The covariance-weighted KP procedure gives $p = 0.0132$. The primary Chen–Fang procedure uses an unweighted residual distance, includes the rank-zero null stratum, and retains rank one with an adjusted p-bound of 0.3848. Both calculations fit a rank-one space, so the weighting helps explain their difference. The object of interest remains target-specific approximation loss together with its uncertainty.

5.5. Shock-set uncertainty and the sign target

The rotation varies over the continuous admissible shock set while the event covariance, FATS transformation, sample, and national benchmark remain fixed. The four-quarter employment point estimates range from −0.607 to −0.112 for foreign control and from 0.457 to 0.633 for extra-EU control. A one-sided intersection–union test asks whether the first loading is negative and the second positive throughout the set. It gives $p = 0.3679$ from the component tests and full angle sweep.

Across the retained rotations, the second singular value of the twelve-row system ranges from 0.099 to 0.334. The union band over sampling and rotation is $[0, 0.898]$. The employment-only interval uses a different target metric. Coefficient bands, shock scales, rank distances, and direct coefficient checks use the same admissible rotation set.

5.6. Financing evidence

Related evidence describes several mechanisms through which joint control can affect transmission. Biermann and Huber (2024) trace a German parent-credit shock through internal funding to the sales and employment of foreign affiliates. Cetorelli and Goldberg (2012) document internal funding adjustments in global banks, and Zhang et al. (2025) connect subsidiaries' investment responses during local monetary contractions to their access to internal capital. In my legacy-shock auxiliary analysis, credit quantities and lending rates give a centered wild-bootstrap block statistic of 65.32 and $p = 0.0005$ at the 1,999-draw resolution. The five-block Holm value is 0.0025. Other channels use their own data frequencies and inference rules.

6. A Numerical Illustration of State Repair

A two-action economy illustrates what can happen when a reported state pools different decision perimeters. Its equilibrium Jacobian contains common-control curvature, a shared-input margin, strategic interaction, and an extra-EU financing buffer. Published FATS shares determine three calibration moments. Policy-response coefficients are reserved for the comparison with the model. The equations, parameter mapping, and solution checks preserve that separation. The exercise asks two questions: whether these organizational channels can generate response differences at the same national state, and whether the resulting response pattern agrees with the empirical contrast. The first is a property of the specified economy; the second compares its implications with coefficients estimated from the data.

A matching rule selects 140 pairs using proximity in the national state and distance in the control state. Their model responses are then compared. The resulting divergence illustrates the half-separation bound for a pooled prediction. Model-internal errors and the model–data comparison use the legacy shock. State and policy curvature bounds include control-dependent incidence and a joint-remainder check. These constants remain in model units. The empirical frontier uses separately specified sensitivity budgets in employment units.

7. Conclusion

A national state can preserve the current allocation while pooling economies with different policy-response systems. Corporate control is one source of this variation because it joins decisions across locations. The rank of the hidden response gives the number of local coordinates required to separate those systems. Its singular values measure the loss from retaining fewer coordinates, and curvature terms extend that calculation to finite experiments.

At the central admissible shock, discarding both measured corporate coordinates leaves a positive lower bound on employment loss under the simultaneous sampling envelope. Retaining the first coordinate reduces the point loss by 85.3 percent. The value of a second coordinate depends on uncertainty in the remaining error and changes when the outcome family changes. Target weights, index-space metrics, curvature, and projection discrepancy determine how the local frontier maps into a finite economic experiment. Specifications that separate the policy-release and conference windows place the fitted response surface within the ECB announcement. The evidence concerns the representation of monetary responses across measured corporate states. The causal effect of assigning corporate control remains a separate object.

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.

Sungbae An, Yongsung Chang, and Sun-Bin Kim. Can a representative-agent model represent a heterogeneous-agent economy? *American Economic Journal: Macroeconomics*, 1(2):29–54, 2009.

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.

Adrien Auclert, Bence Bardóczy, Matthew Rognlie, and Ludwig Straub. Using the sequence-space jacobian to solve and estimate heterogeneous-agent models. *Econometrica*, 89(5):2375–2408, 2021. doi: 10.3982/ECTA17434.

David Rezza Baqaee and Emmanuel Farhi. The macroeconomic impact of microeconomic shocks: Beyond hulten's theorem. *Econometrica*, 87(4):1155–1203, 2019. doi: 10.3982/ECTA15202.

Régis Barnichon and Geert Mesters. A sufficient statistics approach for macro policy. *American Economic Review*, 113(11):2809–2845, 2023. doi: 10.1257/aer.20220581.

Marcus Biermann and Kilian Huber. Tracing the international transmission of a crisis through multinational firms. *Journal of Finance*, 79(3):1789–1829, 2024. doi: 10.1111/jofi.13338.

Karol Borsuk. Drei sätze über die n-dimensionale euklidische sphäre. *Fundamenta Mathematicae*, 20:177–190, 1933. doi: 10.4064/fm-20-1-177-190.

Tomás E. Caravello, Alisdair McKay, and Christian K. Wolf. Evaluating monetary policy counterfactuals: (when) do we need structural models? Working paper, March 4, 2026.

Nicola Cetorelli and Linda S. Goldberg. Banking globalization and monetary transmission. *Journal of Finance*, 67(5):1811–1843, 2012. doi: 10.1111/j.1540-6261.2012.01773.x.

Yongsung Chang, Sun-Bin Kim, and Frank Schorfheide. Labor-market heterogeneity, aggregation, and policy (in)variance of DSGE model parameters. *Journal of the European Economic Association*, 11(S1):193–220, 2013.

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

Antonio Coppola, Matteo Maggiori, Brent Neiman, and Jesse Schreger. Redrawing the map of global capital flows: The role of cross-border financing and tax havens. *Quarterly Journal of Economics*, 136(3):1499–1556, 2021.

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.

Wouter J. Den Haan and Pontus Rendahl. Solving the incomplete markets model with aggregate uncertainty using explicit aggregation. *Journal of Economic Dynamics and Control*, 34(1):69–78, 2010. doi: 10.1016/j.jedc.2008.12.008.

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

Carl Eckart and Gale Young. The approximation of one matrix by another of lower rank. *Psychometrika*, 1(3):211–218, 1936. doi: 10.1007/BF02288367.

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

Fatih Guvenen, Jr. Mataloni, Raymond J., Dylan G. Rassier, and Kim J. Ruhl. Offshore profit shifting and aggregate measurement: Balance of payments, foreign investment, productivity, and the labor share. *American Economic Review*, 112(6):1848–1884, 2022.

James J. Heckman. Building bridges between structural and program evaluation approaches to evaluating policy. *Journal of Economic Literature*, 48(2):356–398, 2010. doi: 10.1257/jel.48.2.356.

Leonid Hurwicz. On the structural form of interdependent systems. In Ernest Nagel, Patrick Suppes, and Alfred Tarski, editors, *Logic, Methodology and Philosophy of Science*, pages 232–239. Stanford University Press, 1962.

Marek Jarociński and Peter Karadi. Deconstructing monetary policy surprises—the role of information shocks. *American Economic Journal: Macroeconomics*, 12(2):1–43, 2020.

Yann Koby and Christian K. Wolf. Aggregation in heterogeneous-firm models: Theory and measurement. Working paper, July 15, 2020.

Per Krusell and Jr. Smith, Anthony A. Income and wealth heterogeneity in the macroeconomy. *Journal of Political Economy*, 106(5):867–896, 1998. doi: 10.1086/250034.

Jr. Lucas, Robert E. Econometric policy evaluation: A critique. *Carnegie-Rochester Conference Series on Public Policy*, 1:19–46, 1976. doi: 10.1016/S0167-2231(76)80003-6.

Jacob Marschak. Economic measurements for policy and prediction. In William C. Hood and Tjalling C. Koopmans, editors, *Studies in Econometric Method*, pages 1–26. Wiley, New York, 1953.

Jiarui Zhang, Yingying Shi, and Lei Li. Corporate response to monetary policies: Do foreign subsidiaries and local firms behave differently? *Journal of International Money and Finance*, 154:103302, 2025.