Shift-Share Designs with Network Propagation

Stan de Ruijter*

July 2026

Abstract

Shift-share (Bartik) designs measure a region's exposure to industry shocks by weighting national shocks by its own-industry employment shares, implicitly assuming each shock affects only its own industry's workers. This assumption does not hold when industries are linked through a production network: a shock to one industry propagates to its suppliers and customers, so regional outcomes respond to both the direct shock and its network-propagated consequences, and the standard Bartik loads an omitted network-exposure term. We develop a closed-form diagnostic for this contamination, $\phi = \text{tr}(\mathbf{QP})/\text{tr}(\mathbf{Q}) \geq 1$ (for the un-residualized co-location matrix), which depends only on the share and propagation matrices. This diagnostic is a design baseline rather than a bound: under instrumental variables the realized contamination ϕ^{IV} also reflects the shock covariance, and in the application below it is roughly twice as large. The correction is straightforward—premultiply the shocks by the Leontief inverse before weighting by shares—and it nests the standard Bartik and requires only a public input–output matrix. We apply the framework to the Autor et al. (2013) China Shock, where the standard 2SLS coefficient is 2.40 times the corrected one; the two nonetheless deliver nearly identical population-weighted aggregate predicted impacts. This cancellation holds under an exact, falsifiable condition: the regional network multiplier must be uncorrelated with instrument exposure, a property of the $(\mathbf{Q}, \mathbf{P}, \mathbf{w}, \mathbf{g})$ design that the China Shock nearly satisfies but fails in its supply-chain-dense subregions. The cancellation concerns the aggregation of partial-equilibrium predicted values rather than general-equilibrium impact. What the correction changes is geography. Of the 100 commuting zones most exposed under the standard measure, 34% are not among the corrected top 100, and on held-out geographies the corrected exposure raises out-of-sample predictive fit from 0.080 to 0.262. We interpret this gain as a check on predictive stability rather than as causal validation of the correction.

JEL Codes: C26, C36, F14, L14, R23

Keywords: Shift-share instruments, Bartik estimator, production networks, input-output, instrumental variables, network propagation

*Department of Economics, University College London, Centre for Research and Analysis of Migration (CReAM), and Rockwool Foundation Berlin. email: uctpsld@ucl.ac.uk.

1 Introduction

Shift-share (Bartik) instruments have become a workhorse design for measuring how regional outcomes respond to industry-level shocks, and they underpin a generation of work on trade exposure, immigration, employment growth, and the diffusion of technology across local labor markets. The standard estimator assigns each region a weighted average of industry-level shocks, with weights equal to the region's lagged industry composition. Built into that exposure measure is an assumption that is easy to miss: each industry's shock affects only its own workers. The assumption is convenient, and for many questions it does no harm. But industries are connected through a production network, and once we take that seriously a shock to industry l propagates through input-output linkages to the suppliers and customers of l, so the regional outcome responds to both the direct shock and its network-propagated consequences. The standard exposure measure then omits the propagated channel, and the standard Bartik instrument loads that omitted term onto its coefficient. We should be careful about what this does and does not imply. A researcher whose object is the total reduced-form incidence of the shock, own-industry plus propagated, is not thereby misspecified; the gap bites when the target is the own-industry effect or the separation of the direct and propagated channels, and its main consequence, as we show, falls on the geography of measured impact rather than on its aggregate. The concern applies most cleanly to settings with input-output linkages, the canonical Autor et al. (2013) (henceforth ADH) China Shock, but it is not specific to production networks: any shift-share design whose units are connected through an observable network (migration networks, social networks, commuting flows) is potentially subject to the same misspecification.

This paper develops the econometrics of shift-share designs under network propagation. We establish three results: a closed-form characterization of the contamination the standard Bartik inherits from omitting the propagated channel; a corrected exposure measure that nests the standard Bartik and is consistent under the same identification regime; and an exact, falsifiable condition (our central theoretical contribution) under which the standard and corrected estimators deliver identical aggregate predicted impacts, which turns a numerical near-coincidence in any one application into a testable property of the design. The correction's economic payoff is geographic: the contamination can leave coefficient-level diagnostics and population-weighted aggregate impacts close to undisturbed while it relocates measured exposure across regions, and that relocation is the object with policy stakes. Let $B_r = \mathbf{s}_r'\mathbf{g}$ denote the standard direct-exposure Bartik and $N_r = \mathbf{s}_r'(\mathbf{P} - \mathbf{I})\mathbf{g}$ the corresponding network-only exposure that captures propagated shocks net of the own-industry component, where $\mathbf{P}$ is a propagation matrix (canonically the Leontief inverse $\mathbf{L} = (\mathbf{I} - \mathbf{A})^{-1}$). Under a two-channel data-generating process, the standard OLS estimator converges to a known mixture $\beta_D + \beta_N(\phi - 1)$ of the direct and network channels, where $\phi - 1$ is exactly the OLS-projection coefficient of N_r on B_r, a textbook omitted-variable-bias story applied to network-propagated exposure. Equivalently, the standard Bartik's per-region prediction error decomposes as $\Delta_r^{\text{std}} = \beta_N B_r(\phi - m_r)$, where $m_r = \tilde{B}_r / B_r$ is the region-level effective multiplier. This single equation organizes the three empirical objects the paper studies: ϕ is the projection component that drives the slope bias; its population-weighted gap $\phi - \psi^{\text{agg}}$ governs aggregate cancellation; and the cross-region

variance of Δ_r^{std} governs geographic misallocation. Three quantities arise on the same $(\mathbf{S}, \mathbf{P})$ design and the paper tracks them throughout. The *structural* contamination multiplier $\phi^{\text{iid}} = \text{tr}(\mathbf{QP})/\text{tr}(\mathbf{Q}) \geq 1$ is an ex-ante diagnostic, computable from $(\mathbf{S}, \mathbf{P})$ alone before any regression is run. The *realized* contamination factor ϕ^{IV} incorporates the cross-covariance of the instrument and the endogenous shock and gives the actual bias of the Bartik 2SLS in any given application; it can differ substantially from ϕ^{iid} when shocks are correlated. The *aggregate amplification* ψ^{agg} measures the population-weighted ratio of corrected to direct exposure and governs whether the coefficient bias passes through to aggregate predicted impacts. The corrected estimator pre-multiplies the shock vector by $\mathbf{P}$ before weighting by shares; consistency follows under the same exogenous-shocks identification regime of Borusyak et al. (2022) (henceforth BHJ) that supports the standard Bartik. We address GPS-style (Goldsmith-Pinkham et al., 2020) share exogeneity in Section 2.5 (where it supports the single-regressor corrected estimator but leaves the channel splits only weakly identified without shift-exogenous variation) and develop the framework correspondences to AKM exposure-robust inference (Adão et al., 2019) and to BH formula-instrument optimality (Borusyak and Hull, 2023, 2025) alongside.

We apply the framework to ADH. The standard 2SLS coefficient on Chinese import exposure is -0.612 (s.e. 0.066, first-stage $F = 61.7$); replacing the direct-exposure Bartik with the network-inclusive Bartik and instrumenting with its other-country analog yields a coefficient of -0.255 (s.e. 0.023, first-stage $F = 192.9$). The two coefficients differ by a factor of 2.40. The headline adopts ADH's timing, in which the endogenous exposure uses start-of-period employment shares and only the instrument's shares are decade-lagged; the fully predetermined variant with lagged shares on both sides gives a smaller ratio of roughly 1.6, and Section 4.1 reports both. The two-regressor decomposition that separates direct and network-only exposure identifies $\hat{\beta}_D = -0.305$ and $\hat{\beta}_N = -0.233$: roughly half of the standard coefficient's magnitude loads onto exposure that propagates through the supply chain rather than onto own-industry displacement. The structural diagnostic in this design is $\hat{\phi}^{\text{iid}} = 1.110$; the realized coefficient-contamination factor (the unweighted bilinear form that absorbs the standard-versus-corrected coefficient gap) is $\hat{\phi}^{IV} = 2.316$; the aggregate amplification is $\hat{\psi}^{\text{agg}} = 2.766$.[1] The gap between $\phi^{\text{iid}} \approx 1.11$ and $\phi^{IV} \approx 2.32$ reflects shock covariance: Chinese import shocks are clustered at the three-digit NAICS level (apparel industries together, electronics industries together) on industry groupings that are themselves IO-linked, and the realized factor absorbs the gap between the standard and corrected coefficients almost exactly. A Neumann-order decomposition shows that first- and second-order supplier-customer linkages account for 87.1% of the amplification, so the correction does not depend on exotic high-order network terms.

One might read the coefficient ratio of 2.40 as evidence that the China Shock literature overstated the aggregate labor-market cost of Chinese import competition. The data do not bear this out. When we translate the standard and corrected coefficients into predicted impacts at the commuting-zone level and aggregate, the two estimators nearly coincide:

[1] The aggregate-cancellation comparison of Section 5 uses the population-weighted, WLS-faithful realized factor $\phi^{IV} = 2.67$ against $\psi^{\text{agg}} = 2.77$ (a design-consistent gap of 3.5%); it differs from the coefficient-contamination factor 2.316 quoted here only by the WLS weighting, and the two should not be differenced against ψ^{agg} interchangeably. Proposition 7 states the cancellation identity on the WLS-faithful objects, and Table 2 catalogs every amplification-factor variant with its weighting, vintage, and value.

the interquartile-range effect is -1.063 percentage points under the standard estimator and -1.196 percentage points under the corrected one, a gap of 11.2% despite the per-unit ratio of 2.40. The reason is a property of the design, which the framework characterizes exactly. The aggregate prediction error decomposes as $\sum_r w_r \Delta_r^{\text{std}} = \beta_N \bar{B}_w (\phi - \psi^{\text{agg}})$: aggregate predictions cancel exactly when $\phi = \psi^{\text{agg}}$, and the gap between the two is governed by a covariance identity:

$$\psi^{\text{agg}} \;=\; \bar{m}^{\text{pop}} \;+\; \frac{\text{Cov}_{\text{pop}}(m_r,\, B_r)}{\bar{B}_w},$$

where $m_r = \tilde{B}_r / B_r$ is the region-level effective multiplier and $\bar{B}_w$ the population-weighted mean direct exposure. The aggregate gap between coefficient bias and exposure understatement is therefore governed entirely by the cross-region covariance of multiplier and direct exposure. In ADH this covariance is meaningfully negative: upstream-supplier regions, those specialized in process industries (petroleum refining, primary metals, basic chemicals, paper) whose output feeds many downstream manufacturers, have high multipliers but low direct exposure to the consumer-goods industries that dominated Chinese import growth. Coefficient inflation and exposure understatement both arise from ignoring the propagation matrix, and they move in lockstep through the same network geometry, so at the aggregate they cancel. This cancellation is conditional rather than automatic: it holds to the extent that $\text{Cov}_{\text{pop}}(m_r, B_r)$ is small, a falsifiable property of the $(\mathbf{Q}, \mathbf{P}, \mathbf{w}, \mathbf{g})$ design that the pooled China Shock nearly satisfies but that fails in its most supply-chain-dense subregions (Section 5.3).

The correction's economic content lies in the spatial distribution of measured impact. The MSE of $\Delta_r^{\text{std}} = \beta_N B_r (\phi - m_r)$ across commuting zones is dominated by its cross-region variance rather than its population-weighted mean. Decomposing the CZ-level mean-squared prediction error between the standard and corrected impacts into a level component and a CZ-specific reallocation component yields a level component of less than 1%; 99.8% is cross-CZ reallocation, with a state-block bootstrap confidence interval for that share lying entirely above 99%. Of the 100 commuting zones most exposed under the standard measure, 34% are not in the corrected top 100; among the top 25, 68% are not. The corrected ranking correlates more strongly with realized manufacturing employment changes than the standard's, producing an out-of-sample R^2 of 0.262 against the standard's 0.080, more than three times larger, with the gain robust to leave-one-Census-division and leave-one-state cross-validation. The zones the correction upweights are the industrial Midwest, Great Plains, and Gulf Coast regions specialized in upstream process industries whose exposure is carried by supply-chain linkages rather than by own-industry overlap with Chinese imports; the zones it downweights are the apparel, footwear, and furniture belts in the Carolinas, Appalachia, and the South, directly targeted by Chinese imports but with simple supply chains. A place-based policy targeted on the standard Bartik therefore misdirects resources away from supply-chain-exposed commuting zones: an equal-budget Trade Adjustment Assistance allocation across the ten standard-measure most-exposed CZs leaves 57.7% of need-weighted welfare on the table at $K = 10$, falling to 14.4% at $K = 100$ but not disappearing. Summed across communities the standard and corrected estimators produce a similar total; what differs is which communities carry that total.

The paper relates to four strands of the shift-share literature. The identification-content

papers of Goldsmith-Pinkham et al. (2020) and Borusyak et al. (2022) establish what the Bartik recovers under share and shift exogeneity respectively; we show that the standard Bartik is *misspecified*, its coefficient contaminated by the omitted network-propagated channel, under a correctly measured propagation matrix even when their shock- and share-exogeneity assumptions hold, because the fault lies in the exposure map rather than in the exogeneity of the shocks. We develop the framework correspondences to AKM exposure-robust inference (Adão et al., 2019) and to BH-style formula instruments (Borusyak and Hull, 2023, 2025) in Section 2.5. Autor et al. (2025) review the trade-shock shift-share literature and explicitly note that the standard analysis estimates "relative effects" that abstract from spillovers between commuting zones; our network-OVB framework formalizes one mechanism for those spillovers and quantifies how they redistribute measured exposure across regions. Goldsmith-Pinkham et al. (2024) document a parallel "contamination bias" in shift-share designs arising from heterogeneous treatment effects interacting with shared regressors; the mechanism there is heterogeneity in β, whereas ours is omission of the network-propagated exposure under the single-Bartik specification, and the two coexist in applied settings. On the applied side, Acemoglu et al. (2016) (henceforth AADHP) augment the ADH analysis with upstream and downstream exposure regressors constructed from the Leontief and Ghosh inverses; our two-regressor OVB formalizes and extends their approach by deriving the closed-form bias of the standard estimator and reparameterizing onto the identity-subtracted basis (B_r, N_r) with $N_r = \mathbf{s}_r'(\mathbf{P} - \mathbf{I})\mathbf{g}$, which subtracts the identity component of the Leontief regressor and yields stable coefficients with strong first stages (Section 4). Pierce et al. (2024) document a downstream cost-reduction channel through which cheaper Chinese inputs benefit U.S. customer producers; in our three-channel decomposition (Appendix B) the Pierce–Schott benefit operates only through non-manufacturing customers and is dominated by a local-supply-disruption channel, producing a net negative downstream effect in the ADH setting. A parallel literature (Adão et al., 2025; Caliendo et al., 2019; Galle et al., 2023; Baqaee and Farhi, 2019) embeds IO linkages in full general-equilibrium models; these structural approaches capture the complete GE response at the cost of stronger modeling assumptions, while our correction uses only the publicly available IO matrix and a single matrix multiplication. Dustmann and Schönberg (2025, §3.6.1) review shift-share identification in migration applications and argue that BHJ-style shock-independence is unlikely to hold when a few dominant origin countries drive the overall inflow (e.g. Mexicans in the U.S.); the network-augmented approach we develop is the natural remedy, and we apply the framework to immigration in the online appendix. Gulek (2026) provides closely related evidence using firm-to-firm VAT data in Turkey, showing that immigration shocks propagate through input-output linkages with magnitudes dictated by network centrality, and that aggregating spillovers can revise the labor-market incidence of immigration substantially when the host region is network-central.

The rest of the paper is organized as follows. Section 2 develops the framework: the two-channel data-generating process, the bias formula and its connection to the realized IV contamination factor, the corrected estimator and its inference, the framework correspondences to BHJ, GPS, AKM, and BH, and the covariance identity linking CZ-level overstatement to aggregate cancellation. Section 3 describes the ADH application and the construction of the propagation matrix from BEA input-output tables. Section 4 reports

the main 2SLS results and the two-regressor OVB decomposition. Section 5 translates the coefficient-level contamination into predicted impacts at the commuting-zone and aggregate levels and documents the geographic reallocation. Section 6 reports validation and robustness checks. Section 7 ports the framework to a second canonical design, the Bartik (1991) employment-growth Bartik—the same IO propagation matrix, a very different shock vector—and closes with a practitioner's checklist. Section 8 concludes. Appendix A contains proofs; Appendix B reports the upstream/downstream three-channel extension; Appendix C documents full-economy ($K = 483$) extensions; Appendix D reports Monte Carlo validation. An online appendix illustrates the framework's portability across three further shift-share settings: modern-era Card (2001) immigration with a native-migration network, the COVID labor-demand shock with the Facebook Social Connectedness Index, and the Acemoglu and Restrepo (2020) robotics shock with the LODES commuting matrix.

2 Framework

2.1 Setup and the Two-Channel DGP

Consider an economy with R regions indexed by r and K industries indexed by k. Let $s_{rk} \geq 0$ denote the share of industry k in region r, with $\sum_k s_{rk} = 1$. Collect the shares in the $R \times K$ matrix $\mathbf{S}$. Each industry receives an exogenous shock g_k, collected in the vector $\mathbf{g}$. The industries are connected through a production network whose propagation structure is summarized by a $K \times K$ matrix $\mathbf{P}$ with non-negative entries and diagonal entries $P_{kk} \geq 1$. The leading case is the Leontief inverse $\mathbf{P} = \mathbf{L} \equiv (\mathbf{I} - \mathbf{A})^{-1}$, where $\mathbf{A}$ is the matrix of direct input requirements: an industry-l shock propagates to industry k with weight L_{kl} through demand-side input-output linkages.[2] Define the $K \times K$ *co-location matrix*

$$\mathbf{Q} \;\equiv\; \mathbf{S}'\mathbf{S}, \qquad Q_{kl} \;=\; \sum_r s_{rk}\, s_{rl}. \tag{1}$$

$\mathbf{Q}$ is symmetric and positive semi-definite. Its diagonal $Q_{kk} = \sum_r s_{rk}^2$ is a Herfindahl-type index of industry k's geographic concentration; its off-diagonal Q_{kl} measures the extent to which industries k and l share the same geographic footprint.

Regional outcomes respond to both direct exposure and network-propagated exposure. We write the data-generating process in two-channel form:

$$y_r \;=\; \beta_D\, \mathbf{s}_r'\mathbf{g} \;+\; \beta_N\, \mathbf{s}_r'(\mathbf{P} - \mathbf{I})\mathbf{g} \;+\; \varepsilon_r, \tag{2}$$

with $B_r \equiv \mathbf{s}_r'\mathbf{g}$ the direct-exposure Bartik and $N_r \equiv \mathbf{s}_r'(\mathbf{P} - \mathbf{I})\mathbf{g}$ the network-only exposure that captures propagated shocks net of the own-industry component already in B_r. The pair (β_D, β_N) is the primary policy-relevant estimand: direct exposure displaces own-industry workers, while network-propagated exposure operates on upstream suppliers and downstream customers through changes in input demand and intermediate supply,

[2] Downstream (supply-side) propagation operates through the Ghosh inverse $\mathbf{G} = (\mathbf{I} - \mathbf{B})^{-1}$. Appendix B extends the two-channel framework of this section into a three-channel form that separately identifies upstream and downstream propagation, and reports the corresponding empirical decomposition.

and the two channels imply different policy responses (own-industry adjustment assistance versus place-based programs targeted on supply-chain-exposed regions, as we develop in Section 5.8). The *single-index special case* imposes $\beta_D = \beta_N \equiv \beta$, collapsing the DGP to

$$y_r \;=\; \beta\, \mathbf{s}_r' \mathbf{P} \mathbf{g} \;+\; \varepsilon_r. \tag{3}$$

This restriction asserts that the per-unit causal effect is the same whether a shock arrives directly through own-industry employment or through a propagated supplier/customer linkage. It is a strong assumption in settings where direct displacement and network-mediated reallocation operate through different mechanisms (different labor-supply margins, different speeds of adjustment) and a natural baseline when the channel of action is factor reallocation across industries within a local labor market; we report results for both the single-index single-regressor specification and the unrestricted two-regressor form, and the single-regressor estimand under heterogeneous channels is characterized in Proposition 3 below. The two-channel DGP nests the single index as a restriction. We develop the theory in the two-channel form and report the more disaggregated three-channel extension (upstream **L** and downstream **G** separately) in Appendix B.

The measure is based on the relationship between the cross-region covariance of $B_r = \mathbf{s}_r' \mathbf{g}$ and the omitted network-only exposure $N_r = \mathbf{s}_r'(\mathbf{P} - \mathbf{I})\mathbf{g}$: this is why the co-location matrix $\mathbf{Q}$ emerges. Expanding the cross-product,

$$\sum_r B_r\, N_r \;=\; \sum_{k,\,k',\,l} Q_{k,k'}\, g_k\, (\mathbf{P} - \mathbf{I})_{k',l}\, g_l, \tag{4}$$

where k is the direct-shock industry carrying B_r through $s_{r,k}$, k' is the propagation-receiver industry carrying N_r through $s_{r,k'}$, and l is the shock origin whose propagated effect on k' enters via $(\mathbf{P} - \mathbf{I})_{k',l}$. The bias-relevant co-location is $Q_{k,k'}$, between the direct and receiver industries, not between the receiver and the shock origin. The shock origin l does not need to share regions with anything: g_l is a national industry-level shock, it propagates to industry k' everywhere via the IO network, and it enters region r's outcome through $s_{r,k'}$. Cross-region propagation is therefore captured fully by the corrected exposure $\tilde{B}_r = \mathbf{s}_r' \mathbf{P} \mathbf{g}$, regardless of the geographic location of l. The OVB picks out only the part that varies across regions in a way correlated with B_r, which is the part driven by co-location of the direct industry k and the receiver industry k'. When direct and receiver industries are in disjoint regions ($Q_{k,k'} = 0$) the propagated effect still affects outcomes: the standard Bartik assigns zero impact to regions that lack direct exposure but participate in the network, while the corrected estimator correctly loads propagated effects onto those regions (Section 5). The missing piece, however, does not correlate with B_r across regions and so does not bias the coefficient. The bias is a genuine omitted-variable-bias phenomenon driven by the joint geometry of **S** and **P**, not a rescaling artifact.

Applied users do not, ultimately, care about the slope coefficient in isolation: they care about which places the shocks hit, and by how much. Define the region-level true impact under the DGP (2) as

$$T_r \;\equiv\; \beta_D\, B_r \;+\; \beta_N\, N_r, \tag{5}$$

and the standard Bartik prediction $\hat{T}_r^{\text{std}} \equiv \hat{\gamma}_{SS}\, B_r$. Anticipating the slope-bias result of

Proposition 1 below and writing $m_r \equiv \tilde{B}_r / B_r$ for the region-level effective multiplier developed in Section 2.6, the asymptotic per-region prediction error of the standard estimator decomposes as

$$\Delta_r^{\text{std}} \;\equiv\; \text{plim}_{K\to\infty}\big(\hat{\gamma}_{SS}\big)\, B_r \;-\; T_r \;=\; \beta_N\, B_r\,(\phi - m_r), \tag{6}$$

where ϕ is the projection coefficient identified in Proposition 1. Equation (6) organizes the remainder of this section as follows: Section 2.2 identifies the projection component ϕ (the slope-bias factor), Sections 2.3–2.4 characterize realized contamination and the estimator hierarchy, and Section 2.6 treats the cross-region distribution of Δ_r^{std} itself—whose population-weighted mean governs aggregate cancellation and whose cross-region variance governs geographic misallocation. Under the two-regressor estimator $\Delta_r^{\text{2reg}} = o_p(1)$ uniformly in r; the corrected single-regressor delivers a different error formula characterized in Corollary 2 below.

The propagation matrix $\mathbf{P} = \mathbf{L}$ is not merely an accounting matrix of dollar input requirements; to first order it is the kernel mapping industry shocks into regional employment responses. Under constant-returns technologies with input–output coefficients $\mathbf{A}$, gross output solves $\mathbf{X} = (\mathbf{I} - \mathbf{A})^{-1}\mathbf{f} = \mathbf{L}\mathbf{f}$, so a final-demand shock $\mathbf{g}$ moves output by $\Delta\mathbf{X} = \mathbf{L}\mathbf{g}$; with employment proportional to gross output at the first order ($\ell_k = \eta_k X_k$), industry k's employment responds by $\eta_k[\mathbf{L}\mathbf{g}]_k$ and region r's by $\mathbf{s}_r'\mathbf{L}\mathbf{g}$, up to a labor–output scalar absorbed into the coefficient. The cross-industry kernel is therefore $\mathbf{L}$, derived rather than assumed: a labor-share-weighted Leontief and the plain Leontief share the same off-diagonal propagation structure. Demand shocks propagate upstream through $\mathbf{L}$ while cost/supply shocks propagate downstream through the Ghosh inverse $\mathbf{G}$; this grounds the three-channel decomposition of Appendix B. The argument is first-order—it abstracts from input substitution, endogenous factor prices, and labor mobility—which is the relevant approximation precisely because the first-order term is empirically dominant in input–output networks (Baqaee and Rubbo, 2023); the misspecification result below is then stated relative to a derived employment-response kernel rather than an imposed one.

2.2 Structural Bias Under OLS

We adapt the standard shift-share regularity conditions of Borusyak et al. (2022) (henceforth BHJ), to the network-propagated setting.[3]

[3]Assumption 1 is BHJ shift-exogeneity restated in our notation. Assumption 2 is the BHJ denominator-concentration regularity that ensures no single industry dominates the co-location matrix $\mathbf{Q}$ or its propagated counterpart $\mathbf{QP}$; we call it "many industries" rather than "many shocks" since the substantive content is about the design matrix, not the distribution of the shocks. Assumption 3 is the standard i.i.d. condition; it is relaxed to correlated shocks in Section 2.3 and fully in Corollary 5 of Appendix A. It also posits mean-zero shocks while the realized ADH import shocks are strictly positive; the regression constant and period fixed effects absorb the common level, the standard recentering device in this literature (see the BH recentering correspondence in Appendix A). The many-industries regularity is an asymptotic benchmark; in the ADH application the realized shock vector is concentrated at the three-digit NAICS level with roughly 21 effective clusters. The paper's substantive inference argument is therefore finite-shock rather than asymptotic: Section 6.2 documents leave-one-industry-out and leave-one-cluster-out jackknife stability, Section 6.4 reports the Adão et al. (2019) exposure-robust variance computed on network-adjusted shares, which comes out

Assumption 1 (Shock Exogeneity)**.** $\mathbb{E}[g_k \mid \mathbf{S}, \varepsilon] = 0$ for every industry k.

Assumption 2 (Many Industries)**.** Along a sequence in which the design size $\mathrm{tr}(\mathbf{Q}) = \sum_r \sum_k s_{rk}^2$ diverges (under the within-region normalization $\sum_k s_{rk} = 1$, $\mathrm{tr}(\mathbf{Q}) = \sum_r \mathrm{HHI}_r \in [R/K,\ R]$; when regional concentration is bounded away from zero, $\mathrm{HHI}_r \geq \underline{h} > 0$, this gives $\mathrm{tr}(\mathbf{Q}) = \Theta(R)$ and the divergence is driven by the number of regions $R \to \infty$, with industries K growing jointly): (i) $\max_k Q_{kk}/\mathrm{tr}(\mathbf{Q}) \to 0$; (ii) $\max_k [\mathbf{QP}]_{kk}/\mathrm{tr}(\mathbf{QP}) \to 0$; (iii) $\sum_{k \neq l} Q_{kl}^2 = o\big([\mathrm{tr}(\mathbf{Q})]^2\big)$; and (iv) $\sum_{k \neq l} [\mathbf{QP}]_{kl}^2 = o\big([\mathrm{tr}(\mathbf{QP})]^2\big)$. Condition (iv) is a primitive restriction on the propagated design, stated directly rather than derived: a norm bound on $\mathbf{P}$ alone does not transfer the off-diagonal mass bound in (iii) from $\mathbf{Q}$ to $\mathbf{QP}$. We write $\mathrm{plim}_{K \to \infty}$ throughout as shorthand for limits along this sequence.

Assumption 3 (i.i.d. Shocks)**.** $\{g_k\}_{k=1}^K$ are i.i.d. with $\mathbb{E}[g_k] = 0$ and $\mathrm{Var}(g_k) = \sigma^2 > 0$.

With Assumption 3 the standard shift-share estimator admits a clean closed-form probability limit.

Proposition 1 (Two-Channel Bias Decomposition)**.** *Under the two-channel DGP* (2) *and Assumptions 1–3, the standard shift-share estimator from $y_r = \alpha + \gamma B_r + u_r$ satisfies*

$$\mathrm{plim}_{K \to \infty} \hat{\gamma}_{SS} \ = \ \beta_D \ + \ \beta_N \cdot (\phi - 1), \tag{7}$$

where the contamination multiplier *is*

$$\phi \ \equiv \ \frac{\mathrm{tr}(\mathbf{QP})}{\mathrm{tr}(\mathbf{Q})} \ = \ \frac{\sum_k \sum_l Q_{k,l} P_{k,l}}{\sum_k Q_{k,k}} \ \geq \ 1, \tag{8}$$

with equality iff (i) $P_{kk} = 1$ for every k with $Q_{kk} > 0$ and (ii) $P_{kl} = 0$ for every off-diagonal pair (k, l) with $Q_{kl} > 0$. Propagation that occurs strictly outside the co-location support of $\mathbf{Q}$ leaves $\phi = 1$ even when $\mathbf{P} \neq \mathbf{I}$: in that case the slope coefficient is unbiased but the regional prediction errors Δ_r^{std} of equation (6) *are still generically nonzero through $m_r \neq 1$. The weight $\phi - 1$ depends only on the share matrix $\mathbf{S}$ and the propagation matrix $\mathbf{P}$; it is invariant to the shock distribution.*

Proof in Appendix A. In the single-index special case $\beta_D = \beta_N \equiv \beta$, the probability limit collapses to $\beta \cdot \phi$: the standard estimator converges to the causal effect scaled by the contamination multiplier. In the general two-channel case, $\phi - 1$ is the bias weight that the standard Bartik places on the network-only channel, and the standard coefficient is a structural mixture of the direct and network-propagated causal effects.

Lemma 1 ($\phi - 1$ as an OLS Projection)**.** *The sample OLS projection of N_r on B_r has coefficient*

$$\hat{\gamma}_{N|B} \ = \ \frac{\sum_r B_r N_r}{\sum_r B_r^2} \ = \ \frac{\mathbf{g}'\mathbf{Q}(\mathbf{P} - \mathbf{I})\mathbf{g}}{\mathbf{g}'\mathbf{Q}\mathbf{g}}, \tag{9}$$

with $\mathrm{plim}_{K \to \infty} \hat{\gamma}_{N|B} = \phi - 1$ under Assumptions 2 and 3. The finite-sample OVB identity $\hat{\gamma}_{SS} = \hat{\beta}_D^* + \hat{\beta}_N^* \cdot \hat{\gamma}_{N|B}$ holds exactly, where $(\hat{\beta}_D^*, \hat{\beta}_N^*)$ are the OLS coefficients from the two-regressor regression of y_r on (B_r, N_r).

modestly more conservative than the state-clustered main-table standard errors.

Proof in Appendix A. The lemma gives $\phi - 1$ a sample-level interpretation: it is the coefficient of a "first-stage" regression of the omitted network-only exposure on the Bartik. This is the precise sense in which the Bartik and the network-only exposure are mechanically linked, and the OVB formula is not an asymptotic approximation but an exact finite-sample decomposition—in the demeaned, unweighted, control-free case. With an intercept the identity holds after centering B_r and N_r; with population weights w_r and additional controls $\mathbf{X}_r$ the identity holds replacing $\mathbf{Q}$ by the residualized, weighted co-location matrix $\tilde{\mathbf{Q}}$ defined in Corollary 4 (Appendix A).

When the regression includes population weights w_r and controls $\mathbf{X}_r$, the bias formula generalizes with $\mathbf{Q}$ replaced by the residualized, weighted co-location matrix $\tilde{\mathbf{Q}} = (\mathbf{M}_X\mathbf{S})'\mathbf{W}(\mathbf{M}_X\mathbf{S})$, where $\mathbf{M}_X$ is the weighted residual-maker from the controls (Corollary 4). In the ADH setting the residualized contamination multiplier is $\hat{\tilde{\phi}} \approx 1.08$, close to the raw $\hat{\phi} = 1.110$, because the ADH controls (Census division fixed effects, lagged demographics) are nearly orthogonal to the IO-weighted share structure that drives the contamination multiplier. The residualized $\tilde{\mathbf{Q}}$ is symmetric and positive semi-definite by construction but, unlike the raw $\mathbf{Q}$, can have negative off-diagonal entries; the non-negativity, monotonicity, and equality-support claims of Proposition 9 (Appendix A) extend only to the un-residualized $\mathbf{Q}$.

2.3 From Structural ϕ to Realized Bias

The contamination multiplier $\phi = \mathrm{tr}(\mathbf{QP})/\mathrm{tr}(\mathbf{Q})$ of Proposition 1 is an *ex-ante design diagnostic*: under i.i.d. shocks it depends only on $(\mathbf{S}, \mathbf{P})$ and can be computed without running a regression. The i.i.d. assumption is a simplifying device that delivers a closed-form, design-only diagnostic; the BHJ identification result that the Bartik works under exogenous shocks and endogenous shares does not require i.i.d. and survives weak dependence under a uniform row-sum bound on $\mathbf{\Sigma}$ (Corollary 5 of Appendix A). In an application whose shock vector is not i.i.d. the realized bias can differ substantially from this structural baseline—in ADH the realized factor is the IV analog $\phi^{IV} = 2.316$ rather than the i.i.d. $\phi^{\text{iid}} = 1.110$. Equation (6) accordingly carries a generic ϕ; wherever the paper attaches a number to the symbol we subscript the variant explicitly (ϕ^{iid}, ϕ^{corr}, ϕ^{IV}), and in the ADH application the value that enters Δ_r^{std} is the realized ϕ^{IV}. Two generalizations of the i.i.d. baseline are relevant, in sequence: first, correlated shocks under OLS; second, the IV analog.

Correlated shocks. Replace Assumption 3 with mean-zero shocks with arbitrary covariance $\mathrm{Cov}(\mathbf{g}) = \mathbf{\Sigma}$ (maintaining Assumptions 1 and 2 together with the uniform row-sum bound of Corollary 5). Under the single-index DGP the probability limit of the standard Bartik estimator generalizes to

$$\phi^{\text{corr}}(\mathbf{\Sigma}) \;=\; \frac{\mathrm{tr}(\mathbf{QP\Sigma})}{\mathrm{tr}(\mathbf{Q\Sigma})} \;=\; \frac{\sum_{k,l}\sigma_{kl}[\mathbf{QP}]_{kl}}{\sum_{k,l}\sigma_{kl}Q_{kl}}, \tag{10}$$

which collapses to ϕ when $\mathbf{\Sigma} = \sigma^2\mathbf{I}$ and is $\geq \phi$ under additional alignment conditions (Corollary 6, Appendix A). The cross-covariance kernel σ_{kl} re-weights which (k,l) pairs enter the ratio: when shocks are clustered at a coarser industry-group level (so that within-group pairs (k,l) have systematically larger σ_{kl} than between-group pairs), and if those within-group pairs are also IO-linked, the numerator's $[\mathbf{QP}]_{kl}$ is systematically large relative to the denominator's Q_{kl} on the relevant support, pushing ϕ^{corr} above the i.i.d. baseline. This mechanism is empirically active in any application whose shock vector is clustered at a level coarser than the propagation-matrix granularity but where within-cluster pairs share IO linkages—a property that holds in the ADH application (Section 4.4).

The IV extension. Under 2SLS with an external instrument, the bias depends on the cross-covariance of the instrument and the endogenous shocks rather than on the variance of either alone.

Corollary 1 (IV Bias Under Network Propagation)**.** *Suppose the true DGP is the single-index form* (3) *with primitive shock* $\mathbf{g}^{US}$*, the endogenous variable is* $X_r = \mathbf{s}_r'\mathbf{g}^{US}$*, and the instrument is a shift-exogenous analog* $Z_r = \mathbf{s}_r'\mathbf{g}^{OC}$ *with* $\mathbb{E}[g_k^{OC} \mid \mathbf{S}, \varepsilon] = 0$*. Then*

$$\hat{\beta}_{2SLS} \;-\; \beta \cdot \phi^{IV}(\mathbf{g}^{OC}, \mathbf{g}^{US}) \;=\; o_p(1), \qquad \phi^{IV}(\mathbf{g}^{OC}, \mathbf{g}^{US}) \;\equiv\; \frac{(\mathbf{g}^{OC})'\mathbf{QP}\mathbf{g}^{US}}{(\mathbf{g}^{OC})'\mathbf{Q}\mathbf{g}^{US}}. \tag{11}$$

Under mean-zero shocks with cross-covariance $\sigma_{kl}^{OC,US} \equiv \text{Cov}(g_k^{OC}, g_l^{US})$ *and the denominator-concentration regularity (Assumption 2: no single industry dominates the trace of* $\mathbf{Q}$ *or* $\mathbf{QP}$*),*

$$\text{plim}_{K\to\infty}\, \phi^{IV}(\mathbf{g}^{OC}, \mathbf{g}^{US}) \;=\; \frac{\sum_{k,l}\, \sigma_{kl}^{OC,US}\, [\mathbf{QP}]_{kl}}{\sum_{k,l}\, \sigma_{kl}^{OC,US}\, Q_{kl}}. \tag{12}$$

Proof in Appendix A. The population limit (12) collapses to the structural $\phi = \text{tr}(\mathbf{QP})/\text{tr}(\mathbf{Q})$ when the cross-covariance kernel has Kronecker structure $\sigma_{kl}^{OC,US} = \tau^2\delta_{kl}$, i.e. when the instrument and endogenous shocks are independent across industries and share a common cross-variance at each industry: in that case only diagonal terms $[\mathbf{QP}]_{kk}$ and Q_{kk} survive in both numerator and denominator, and the constant τ^2 cancels. When instead the two shock vectors share a common off-diagonal correlation structure (a generic feature of designs where instrument and endogenous shocks are clustered at the same level, for example both clustered at a coarser industry-group level shared between the two), ϕ^{IV} is close to the OLS correlated-shocks object ϕ^{corr} of (10): the IV itself does not introduce additional bias beyond what shock correlation in the endogenous vector already delivers. The empirical decomposition of ϕ^{IV} for the ADH application appears in Section 4.4.

2.4 The Corrected Estimator and Its Inference

The correction to the standard Bartik admits two natural estimators. The single-regressor corrected Bartik $\tilde{B}_r = \mathbf{s}_r'\mathbf{Pg}$ is the natural single-index/projection target and is consistent

for the causal coefficient when direct and network channels share a common per-unit effect; the two-regressor specification on (B_r, N_r) is the methodological gold standard that separately identifies (β_D, β_N) whenever BHJ-style shift-exogenous instruments and the relevant rank conditions are available. We develop both, characterize when the single-index restriction binds, and then state the per-region prediction errors of the three estimators (standard, corrected single-regressor, two-regressor) in a single corollary that maps back to equation (6).

Definition 1 (Network-Inclusive Bartik Instrument)**.** The *network-inclusive Bartik instrument* pre-multiplies the shock vector by the propagation matrix before weighting by shares:

$$\tilde{B}_r \;\equiv\; \mathbf{s}_r'\mathbf{P}\mathbf{g} \;=\; B_r + N_r. \tag{13}$$

Proposition 2 (Consistency of the Single-Regressor Corrected Estimator)**.** *Under the single-index DGP* (3) *and Assumptions 1–3, the network-adjusted shift-share estimator*

$$\hat{\gamma}_{NSS} \;=\; \frac{\sum_r \tilde{B}_r\, y_r}{\sum_r \tilde{B}_r^2}$$

satisfies $\operatorname{plim}_{K\to\infty} \hat{\gamma}_{NSS} = \beta$.

Proof in Appendix A. Consistency extends to weakly dependent shocks under a row-sum bound on the covariance matrix $\mathbf{\Sigma} = \operatorname{Cov}(\mathbf{g})$ (Corollary 5, Appendix A).

Under the general two-channel DGP (2) with $\beta_D \neq \beta_N$, the single-regressor corrected estimator converges to a projection of the two causal effects.

Proposition 3 (Projection Under Heterogeneous Direct and Network Effects)**.** *Under the two-channel DGP* (2) *and Assumptions 1–3,*

$$\operatorname{plim}_{K\to\infty} \hat{\gamma}_{NSS} \;=\; \lambda\,\beta_D + (1-\lambda)\,\beta_N, \tag{14}$$

where, under i.i.d. shocks, the projection weight $\lambda \equiv \operatorname{tr}(\mathbf{QP})/\operatorname{tr}(\mathbf{P}'\mathbf{QP}) \in [0, 1/\phi]$ *is pinned down by* $(\mathbf{S}, \mathbf{P})$ *alone. Under a general shock covariance the design-only trace ratio is replaced by the realized shock cross-products, yielding the IV-corrected analog* λ^{IV} *of equation* (18) *below.*

Proof in Appendix A. The single-regressor corrected estimator therefore recovers a convex combination of β_D and β_N whose weight λ is pinned down by $(\mathbf{S}, \mathbf{P})$ alone under i.i.d. shocks. The two-regressor corrected estimator that instruments (B_r, N_r) jointly recovers (β_D, β_N) separately under the same regularity conditions.

The three estimators (standard, corrected single-regressor, two-regressor) generate three distinct per-region prediction errors, which the following corollary collects.

Corollary 2 (Prediction Errors of the Three Estimators)**.** *Under the two-channel DGP* (2) *and Assumptions 1–3, the asymptotic per-region prediction errors of the standard Bartik, the corrected*

single-regressor, and the two-regressor specification satisfy

$$\Delta_r^{\text{std}} = \beta_N B_r (\phi - m_r), \tag{15}$$

$$\Delta_r^{\text{corr}} = (\beta_D - \beta_N)\big[\lambda N_r - (1-\lambda) B_r\big], \tag{16}$$

$$\Delta_r^{\text{2reg}} = o_p(1), \quad \textit{uniformly in } r, \tag{17}$$

where $m_r = \tilde{B}_r / B_r$, ϕ is the realized contamination factor of Section 2.3, and λ is the projection weight of Proposition 3.

Proof in Appendix A. The corollary makes the estimator hierarchy operational: under exogeneity and rank, the two-regressor specification has zero asymptotic prediction error at every region; the corrected single-regressor has zero error only when channels coincide ($\beta_D = \beta_N$, so the bracket vanishes) or in the trivial special case $\lambda N_r = (1-\lambda)B_r$; and the standard Bartik has zero error only when the realized ϕ equals the region-specific multiplier m_r—a condition that fails generically across regions even when it holds on the population-weighted aggregate. The cross-region variance of Δ_r^{std} is the geographic-misallocation object reported empirically in Section 5.

The projection weight λ is informative in three ways. First, it gives the ex-ante weighted-average interpretation of the single-regressor corrected estimator under heterogeneous channels: $\text{plim}\hat{\gamma}_{NSS} = \lambda\beta_D + (1-\lambda)\beta_N$, with λ pinned down by the design alone. Second, it is a diagnostic for whether the single-regressor specification is more informative about the direct or the network channel: large λ (sparse propagation, $\mathbf{P}$ close to $\mathbf{I}$) means the single-regressor estimator is essentially β_D, while small λ means it is essentially β_N, and the practitioner should choose specification accordingly. Third, comparing the structural $\hat{\lambda}^{\text{struct}} = \text{tr}(\mathbf{QL})/\text{tr}(\mathbf{L}'\mathbf{QL})$ to the implied $\hat{\lambda}^{\text{impl}} = (\hat{\gamma}_{NSS} - \hat{\beta}_N)/(\hat{\beta}_D - \hat{\beta}_N)$ from the joint two-regressor estimates is a useful diagnostic for two-channel DGP misspecification, but it is not a clean over-identification test in the IV setting: the structural λ^{struct} is the OLS/i.i.d. projection weight, while the appropriate analog for the IV-corrected single-regressor uses the realized cross-products of instrument and endogenous shock, namely

$$\lambda^{IV} \equiv \frac{(\mathbf{g}^{OC})'\mathbf{QPg}^{US}}{(\mathbf{g}^{OC})'\mathbf{P}'\mathbf{QPg}^{US}}, \tag{18}$$

in the parallel of Corollary 1. Under correlated shocks the structural and realized λ are distinct objects, and the appropriate over-identification-style comparison is between λ^{IV} and λ^{impl} rather than between λ^{struct} and λ^{impl}. In ADH the three values are $\hat{\lambda}^{\text{struct}} = 0.76$, $\hat{\lambda}^{IV} = 0.45$, and $\hat{\lambda}^{\text{impl}} = 0.30$: the realized IV weight sits between the structural i.i.d. value and the implied two-regressor value, much closer to the latter. The apparent gap between the structural λ^{struct} and the implied λ^{impl} in the IV setting is therefore largely a property of the i.i.d. approximation being wrong for the actual ADH shock vector rather than evidence of two-channel DGP misspecification: once shock correlation is accounted for in the realized projection weight, the gap to the implied weight is roughly one third of the i.i.d. gap. On its own, λ is not informative about the relative magnitudes of β_D and β_N: it is a projection weight rather than a ratio of effects. What it cannot do is separate the two channels tightly; that still requires a two-regressor specification with sufficient first-stage

strength to identify (β_D, β_N) separately—which the BHJ regime delivers in ADH and a GPS-only design delivers only weakly (Section 2.5).

Inference with network-adjusted shares. The network-inclusive Bartik can be rewritten as a shift-share with *network-adjusted shares*:

$$\tilde{B}_r \;=\; \sum_l \tilde{s}_{rl}\, g_l, \qquad \tilde{s}_{rl} \;\equiv\; \sum_k s_{rk}\, P_{kl}. \tag{19}$$

Each primitive shock g_l now enters through a share $\tilde{s}_{rl}$ that aggregates direct ownership (s_{rl}) with the shares of all industries that supply l's production chain through the network (weighted by P_{kl}). The Adão et al. (2019) exposure-robust variance formula applied to $(\tilde{\mathbf{S}}, \mathbf{g})$ delivers valid inference; using the original shares $\mathbf{S}$ in the variance formula produces an inconsistent estimate.

Proposition 4 (Network-Adjusted Exposure-Robust Inference)**.** *Let $\hat{\gamma}_{NSS}$ solve the regression $y_r = \alpha + \gamma\, \tilde{B}_r + \varepsilon_r$ with residuals $\hat{\varepsilon}_r$ and let $\tilde{s}_{rl}$ be as in* (19)*. Suppose the denominator-concentration regularity of Assumption 2 applies to the network-adjusted share matrix $\tilde{\mathbf{S}}$ (in place of $\mathbf{S}$): no single primitive shock l accounts for more than a vanishing fraction of $\sum_r w_r \tilde{s}_{rl}^2$, and the off-diagonal sums $\sum_{l \neq l'} (\sum_r w_r \tilde{s}_{rl} \tilde{s}_{rl'})^2$ are of smaller order than $[\sum_l \sum_r w_r \tilde{s}_{rl}^2]^2$. The exposure-robust variance estimator for $\hat{\gamma}_{NSS}$ is*

$$\hat{V}_{NSS} \;=\; \frac{1}{\hat{D}_{NSS}^2} \sum_{l=1}^{K} \hat{v}_l^2, \qquad \hat{v}_l \;=\; \tilde{g}_l \sum_r w_r\, \tilde{s}_{rl}\, \hat{\varepsilon}_r, \tag{20}$$

with $\hat{D}_{NSS} = \sum_r w_r (\tilde{B}_r - \bar{\tilde{B}})^2$ and $\tilde{g}_l$ the shock residualized at the shock level (at minimum demeaned with the importance weights $n_l = \sum_r w_r \tilde{s}_{rl}$): the score multiplies the share-weighted residual sum by the shock, since the sampling variation being estimated is that of $\sum_l g_l \big(\sum_r w_r \tilde{s}_{rl} \varepsilon_r\big)$. At $\mathbf{P} = \mathbf{I}$ the formula reduces to the Adão et al. (2019) estimator on the original shares. Substituting s_{rl} for $\tilde{s}_{rl}$ produces an inconsistent variance estimate: the naive formula indexes residuals by original industry k, missing the cross-k correlation introduced by each primitive shock g_l acting on every industry that supplies l through the network.

The network-adjusted concentration condition is a substantive restriction: $\tilde{\mathbf{S}}$ can be more concentrated than $\mathbf{S}$ when the propagation matrix $\mathbf{P}$ has columns with large column sums (a few highly central upstream industries propagating broadly), in which case the AKM asymptotics under network adjustment require a finite-sample diagnostic. As a check, we report the Rotemberg-weight effective N for both $\mathbf{S}$ and $\tilde{\mathbf{S}}$ in Table 7.
Proof in Appendix A. Section 6.4 quantifies the gap empirically in the ADH setting.

Robustness to a mismeasured propagation matrix. The correction replaces $\mathbf{g}$ with $\mathbf{Pg}$ through a measured propagation matrix, so a natural concern is that using a national input–output table in place of region-specific linkages injects measurement error into $\mathbf{P}$. At first glance, it may seem that this delivers the classical attenuation toward zero; the induced bias is not of that form. Appendix A.4 derives it exactly (Proposition 10): the

bias runs entirely through the covariance of the instrument with the propagation-error term, so its sign is not pinned down a priori, and it is governed by an estimable, signed statistic bounded by a Cauchy–Schwarz ceiling (Corollary 7). Estimating that statistic against state-specific Leontief tables in Section 6.7 gives $\hat{\kappa} = -10.36\%$ (against a ceiling of 17.69%), an over-correction rather than the conservative lower bound that the attenuation intuition would assert.

2.5 Identification Regimes

The bias formulation of Proposition 1 and the corrected estimator of Definition 1 interact with the four leading frameworks for shift-share identification. The structural bias weight $\phi - 1$ depends only on $(\mathbf{S}, \mathbf{P})$ and is invariant to which framework supplies identification, but the mechanism through which it manifests differs across regimes and each perspective offers a distinct diagnostic. Most empirical work in the shift-share literature operates in either the BHJ shift-exogeneity regime or the GPS share-exogeneity regime; the two impose materially different constraints on what the corrected estimator can identify.

BHJ shift exogeneity. Borusyak et al. (2022) (BHJ) identify the Bartik through quasi-random assignment of shocks. Their shock-level equivalence representation, $\hat{\gamma}_{SS} = \sum_k g_k \hat{y}_k / \sum_k g_k \hat{B}_k$ with $\hat{y}_k = \sum_r s_{rk} y_r$ and $\hat{B}_k = [\mathbf{Qg}]_k$, makes plain that misspecifying the exposure formula—using $\mathbf{g}$ where the true exposure operator is $\mathbf{Pg}$—contaminates the share-weighted outcome $\hat{y}_k$, which captures effects on industry k from all shocks $\{g_l\}_l$ that propagate to k through the network and are picked up via geographic co-location ($Q_{kl} > 0$). Under BHJ shift exogeneity our two-regressor and three-regressor (Appendix B) corrected estimators are jointly identified using shock-based instruments. ADH is a case in point: the other-country shock $\mathbf{g}^{OC}$ supplies the independent shift-exogenous variation that the multi-regressor decomposition requires (Section 4).

GPS share exogeneity. Goldsmith-Pinkham et al. (2020) (GPS) instead identify off the shares: each industry share s_{rk} is a valid instrument under $\mathbb{E}[s_{rk}\, \varepsilon_r] = 0$, and the Bartik is a Rotemberg-weighted GMM estimator using all K shares as instruments. Network propagation generically violates the GPS exclusion restriction for the standard exposure, since each share s_{rk} predicts both direct and propagated exposure: the just-identified IV estimate using share s_{rk} alone satisfies $\hat{\beta}_k = \beta_D + \beta_N \cdot \mathrm{Cov}(s_{rk}, N_r)/\mathrm{Cov}(s_{rk}, B_r)$, with the second term generically nonzero whenever industry k is co-located with IO-linked neighbors. The Rotemberg decomposition (Section 4.4, Table 7) makes explicit which industries drive the contamination. A practical limit concerns the multi-regressor channel splits, and it is worth stating precisely, once. With a single realized shock vector the exposure regressors B_r, N_r^L, N_r^G are fixed linear transformations of the same shares matrix, with coefficient vectors $\mathbf{g}$, $(\mathbf{L} - \mathbf{I})\mathbf{g}$, $(\mathbf{G} - \mathbf{I})\mathbf{g}$. The rank condition for a multi-regressor split then requires only linear independence of these transformed exposures, which holds generically; the exact failure occurs only on a measure-zero set of designs. The binding problem is weak identification: with no independent shock-level variation the split rests entirely on the fixed geometry of those vectors, which are typically nearly collinear, so channel coefficients are poorly

determined in practice. GPS-only designs therefore support the single-regressor corrected estimator (interpreted through Proposition 3 as a projection of the channel effects) but deliver only weakly identified channel splits without an external shift-exogenous instrument; wherever the paper refers back to this limit, this is the statement meant. Section 7 works this out for the Bartik (1991) employment-growth Bartik. Dustmann and Schönberg (2025, §3.6.1) note that the BHJ-style shock-independence assumption is implausible in migration applications when a small number of origin countries dominate the inflow, since the within-origin shocks are correlated and the BHJ shock-level LLN fails, which they argue is the empirically relevant case for U.S. immigration. The network-augmented approach we develop is one natural remedy when the design admits a network structure: applying the framework to modern Card-style immigration with a migration-network propagation matrix is the subject of the online appendix. Autor et al. (2025) adopt the BHJ framework explicitly for ADH-style trade-shock applications, choosing shift exogeneity over share exogeneity on the basis of substantial industry-level variation in Chinese import competition; our network-OVB correction operates within their identifying regime and quantifies the contamination from omitted propagated exposure that BHJ alone does not address.

AKM exposure-robust inference. Adão et al. (2019) (AKM) document that residuals in shift-share regressions are correlated across regions sharing industry composition, requiring an exposure-robust variance correction. Network propagation creates an additional source of cross-regional residual correlation through shared exposure to network-propagated shocks; Proposition 4 above formalizes the network-adjusted exposure-robust variance, replacing original shares s_{rk} with network-adjusted shares $\tilde{s}_{rl} = \sum_k s_{rk} P_{kl}$. Substituting the original shares into the AKM formula produces an inconsistent variance estimate; the implied standard error is about 36% smaller than under network-adjusted shares in the ADH application (Section 6.4).

BH formula instruments. Borusyak and Hull (2023) cast the Bartik as a "formula instrument"—a deterministic combination of exogenous shocks and endogenous exposure—and Borusyak and Hull (2025) characterize optimality and recentering in this class. Our corrected Bartik $\tilde{B}_r = \mathbf{s}_r' \mathbf{P} \mathbf{g}$ is the formula implied by the correctly specified network exposure map under the linear two-channel DGP (2): the contamination multiplier $\phi - 1$ is the closed-form cost of using the standard formula $\mathbf{s}_r' \mathbf{g}$ when the correct exposure operator is $\mathbf{s}_r' \mathbf{P}$. Formal BH-style optimality requires the additional recentering and efficiency conditions of Borusyak and Hull (2025), which we do not prove here; restricted transformations of $\mathbf{P}$ (truncations to first-order neighbors only, or to within-sector links) discard propagation signal under a known and correct $\mathbf{P}$ with i.i.d. shocks but may improve robustness when $\mathbf{P}$ is noisy or overfit. The exposure-map framing of our correction is therefore complementary to the formula-instrument efficiency program: our claim concerns misspecification of the exposure map rather than optimality within the formula class.

Contamination bias. Goldsmith-Pinkham et al. (2024) document a parallel "contamination bias" in shift-share designs in which OLS estimates of the effect of industry k's shock are contaminated by effects of other shocks under treatment-effect heterogeneity. Their mechanism (heterogeneous β_k interacting with shared regressors) is parallel to but distinct from ours (homogeneous β_D and β_N interacting with an omitted network-propagated regressor); the two coexist in applied designs and both can be read off the per-industry Rotemberg-weight decomposition. The framework correspondences, compact closed-form connections to BHJ, GPS, AKM, and BH, are formalized in Appendix A.

2.6 Prediction Errors and Incidence

The per-region prediction error introduced in equation (6), $\Delta_r^{\text{std}} = \beta_N B_r(\phi - m_r)$, organizes three empirical objects within a single algebraic identity: the population-weighted mean of Δ_r^{std} governs *aggregate cancellation* (when standard and corrected aggregates near-coincide despite a large slope-bias factor); the projection component ϕ alone governs the *slope bias* (the coefficient-level contamination of Proposition 1); and the cross-region variance/MSE of Δ_r^{std} governs *geographic misallocation* (which CZs are wrongly identified as exposed). Proposition 1 establishes that the standard Bartik coefficient is a structural mixture of direct and network-propagated causal effects. The task of this subsection is to translate that coefficient-level contamination into statements about predicted impacts at the region and aggregate levels. The conditions under which the standard and corrected aggregates coincide can then be stated as algebraic equalities involving these quantities; whether they are satisfied numerically in any given application is a property of the design, addressed empirically in Section 5.

Define the region-level effective multiplier

$$m_r \;\equiv\; \frac{\mathbf{s}_r'\mathbf{P}\mathbf{g}}{B_r} \;=\; \frac{\tilde{B}_r}{B_r}, \tag{21}$$

so that $N_r/B_r = m_r - 1$, and the *channel ratio* $\theta \equiv \beta_N/\beta_D$ measuring the relative strength of the network channel.

Proposition 5 (Region-Level Overstatement Ratio)**.** *Under the two-channel DGP* (2)*, the standard and corrected predicted impacts at region r satisfy*

$$\frac{\hat{y}_r^{\text{std}}}{\hat{y}_r^{\text{corr}}} \;=\; \frac{1 + \theta\,(\phi - 1)}{1 + \theta\,(m_r - 1)}, \tag{22}$$

where $\hat{y}_r^{\text{std}} = \hat{\gamma}_{SS}\, B_r$ and $\hat{y}_r^{\text{corr}} = \hat{\beta}_D\, B_r + \hat{\beta}_N\, N_r$. The numerator is constant across regions (it uses only the structural contamination multiplier ϕ); the denominator varies with region composition through m_r. In the single-index case $\theta = 1$ the ratio simplifies to ϕ/m_r.

Proof in Appendix A. Averaging (22) over population weights w_r yields the aggregate overstatement ratio and the condition under which it equals one.

Corollary 3 (Aggregate Amplification and the Aggregate Overstatement Ratio)**.** *Let $\bar{B}_w \equiv \sum_r w_r B_r$ and define the population-weighted aggregate amplification*

$$\psi^{\text{agg}} \;\equiv\; \frac{\sum_r w_r\, \mathbf{s}_r' \mathbf{P} \mathbf{g}}{\sum_r w_r\, B_r} \;=\; 1 + \frac{\sum_r w_r N_r}{\bar{B}_w}. \tag{23}$$

The aggregate overstatement ratio of the standard estimator's predicted impact relative to the corrected estimator's is

$$\frac{T^{\text{std}}}{T^{\text{corr}}} \;=\; \frac{1 + \theta\,(\phi - 1)}{1 + \theta\,(\psi^{\text{agg}} - 1)}, \tag{24}$$

which equals one exactly when $\phi = \psi^{\text{agg}}$. Both ϕ and ψ^{agg} are functions of $(\mathbf{S}, \mathbf{P}, \mathbf{w}, \mathbf{g})$ alone; neither depends on the causal coefficients (β_D, β_N).

The corollary makes precise that the aggregate disagreement between standard and corrected predictions is governed by the gap between ϕ and ψ^{agg}, scaled by the channel ratio θ. Equivalently, summing the per-region prediction error (6) weighted by w_r yields

$$\sum_r w_r\, \Delta_r^{\text{std}} \;=\; \beta_N\, \bar{B}_w\, \big(\phi \;-\; \psi^{\text{agg}}\big), \tag{25}$$

so the aggregate-cancellation condition $\phi = \psi^{\text{agg}}$ is exactly the condition $\sum_r w_r \Delta_r^{\text{std}} = 0$. The two amplification factors arise from the same act of omitting the propagation matrix but operate through different statistical objects: ϕ is the cross-region projection coefficient that drives the slope bias, while ψ^{agg} is the population-weighted ratio of corrected to direct aggregate exposure. Whether the two coincide is a property of the design $(\mathbf{S}, \mathbf{P}, \mathbf{w}, \mathbf{g})$. The next proposition decomposes ψ^{agg} into components, and so makes the dependence on the design explicit.

Proposition 6 (Covariance Identity for the Aggregate Gap)**.** *With the population-weighted unconditional multiplier $\bar{m}^{\text{pop}} \equiv \sum_r w_r m_r$, the population-weighted aggregate multiplier decomposes as*

$$\psi^{\text{agg}} \;=\; \bar{m}^{\text{pop}} \;+\; \frac{\text{Cov}_{\text{pop}}(m_r,\, B_r)}{\bar{B}_w}. \tag{26}$$

Because $m_r = 1 + N_r / B_r$, the governing covariance can be written equivalently in terms of the multiplier m_r or the exposure ratio N_r / B_r:

$$\begin{aligned} \psi^{\text{agg}} - \bar{m}^{\text{pop}} \;&=\; \frac{\text{Cov}_{\text{pop}}(m_r, B_r)}{\bar{B}_w}, \\ \text{sign}\big(\psi^{\text{agg}} - \bar{m}^{\text{pop}}\big) \;&=\; \text{sign}\big(\text{Cov}_{\text{pop}}(m_r,\, B_r)\big) \;=\; \text{sign}\big(\text{Cov}_{\text{pop}}(N_r / B_r,\, B_r)\big). \end{aligned} \tag{27}$$

The sign of the aggregate gap is governed by the covariance of the multiplier m_r (equivalently the exposure ratio N_r / B_r) with direct exposure B_r, not by $\text{Cov}_{\text{pop}}(N_r, B_r)$: a region's contribution to aggregate amplification is its multiplier m_r, weighted by its exposure mass $w_r B_r$, so high-multiplier regions pull the aggregate up only insofar as they also carry direct exposure.

Proof in Appendix A. The identity (26) is sample-level and requires no asymptotic argument. Whether ψ^{agg} lies near the realized contamination factor ϕ^{realized}—and there-

fore whether the aggregate predictions of standard and corrected estimators numerically coincide—depends on the joint distribution of $(\mathbf{S}, \mathbf{P}, \mathbf{w}, \mathbf{g})$ in the application at hand. Section 5.1 reports the realized values for ADH and the within-sample subsamples in which the alignment fails.

The next proposition characterizes exactly when the aggregate cancels, and isolates a portable structural class under which it cancels for any shock distribution and any weights.

Proposition 7 (Exact Cancellation Condition and a Sufficient Class)**.** *Let $Z_r \equiv \mathbf{s}_r' \mathbf{g}^{OC}$ denote the instrument's direct exposure, take the realized IV contamination factor in the population-weighted (WLS-faithful) form $\phi^{IV} = \sum_r w_r Z_r \tilde{B}_r / \sum_r w_r Z_r B_r$ that matches a population-weighted 2SLS,[4] and let $p_r \propto w_r B_r$ be the exposure measure on the support $B_r > 0$. Then:*

(i) (Exact condition.) *The realized IV factor and the aggregate amplification obey the sample-level identity*

$$\phi^{IV} - \psi^{\text{agg}} \;=\; \frac{\text{Cov}_p(m_r,\, Z_r)}{\mathbb{E}_p[Z_r]}, \tag{28}$$

so aggregate cancellation $\phi^{IV} = \psi^{\text{agg}}$—equivalently $\sum_r w_r \Delta_r^{\text{std}} = 0$—holds if and only if $\text{Cov}_p(m_r, Z_r) = 0$: under the exposure measure p, the regional multiplier m_r is uncorrelated with instrument exposure. An equivalent slope-balance form is $\text{Cov}_w(Z_r, N_r) = (\psi^{\text{agg}} - 1)\,\text{Cov}_w(Z_r, B_r)$.

(ii) (Distribution-free sufficient class.) *If $\mathbf{S}(\mathbf{P} - c\mathbf{I})\mathbf{g} = \mathbf{0}$ for some scalar c—equivalently the regional multiplier is uniform, $m_r \equiv c$ for all r; in particular whenever $\mathbf{g}$ is a right-eigenvector of $\mathbf{P}$—then $\phi^{\text{corr}} = \phi^{IV} = \psi^{\text{agg}} = c$ for any weights $\mathbf{w}$ and any shocks $(\mathbf{g}, \mathbf{g}^{OC})$, and cancellation holds region by region ($\Delta_r^{\text{std}} = 0$ for every r), not merely in aggregate. (The shock-free i.i.d. diagnostic $\phi^{\text{iid}} = \text{tr}(\mathbf{QP})/\text{tr}(\mathbf{Q})$ does not depend on $\mathbf{g}$, so it need not equal c under this condition; it does so only under the stronger restriction that $\mathbf{P} = c\mathbf{I}$ on the support of $\mathbf{Q}$.)*

(iii) (Falsifiable bound.) *The relative aggregate gap is bounded by the cross-region dispersion of the multiplier:*

$$\left| \frac{\phi^{IV}}{\psi^{\text{agg}}} - 1 \right| \;\leq\; \frac{\text{CV}_p(Z_r)\,\text{sd}_p(m_r)}{\psi^{\text{agg}}}, \tag{29}$$

with $\text{CV}_p(Z_r) = \text{sd}_p(Z_r)/\mathbb{E}_p[Z_r]$. The master scale of the gap is $\text{sd}_p(m_r)$: as the multiplier becomes uniform across regions the gap vanishes at first order.

Proof in Appendix A. Proposition 7 turns aggregate cancellation from an application-specific coincidence into a portable, falsifiable statement, checkable ex ante from $(\mathbf{S}, \mathbf{P}, \mathbf{g}, \mathbf{g}^{OC}, \mathbf{w})$ without estimating a regression: cancellation is exact when the network multiplier m_r is orthogonal to instrument exposure under the exposure measure, and its

[4]The factor ϕ^{IV} of (11) uses the unweighted co-location matrix $\mathbf{Q} = \mathbf{S}'\mathbf{S}$; a population-weighted 2SLS instead realizes the weighted co-location $\mathbf{Q}_w = \mathbf{S}'\mathbf{W}\mathbf{S}$, i.e. the weighted form stated here, and the two coincide only under uniform weights. We use the weighted form throughout the ADH cancellation analysis because the ADH regressions are population-weighted; the unweighted-$\mathbf{Q}$ analog satisfies the parallel identity (28) with $p_r \propto B_r$ and Z_r replaced by Z_r/w_r.

failure is bounded by how dispersed m_r is across regions. In ADH (population-weighted, $\mathbf{P} = \mathbf{L}$), the realized factor is $\phi^{IV} = 2.67$ against $\psi^{\text{agg}} = 2.77$, a gap of -0.095 (an implied -3.5% of ψ^{agg}) reproduced exactly by $\text{Cov}_p(m_r, Z_r)/\mathbb{E}_p[Z_r]$ with $\text{Cov}_p(m_r, Z_r) = -0.283$ (correlation -0.10). The design does not satisfy the exact sufficient class—the eigenvector residual $\|(\mathbf{P} - c\mathbf{I})\mathbf{g}\|/\|\mathbf{P}\mathbf{g}\| = 0.39$ and the multiplier dispersion $\text{sd}_p(m_r) = 1.04$ are both well away from zero, so the ADH near-cancellation is a measurable near-miss of the condition rather than an identity, and Section 5.3 exhibits the condition failing outright in the high-IO-density tercile.

The framework now distinguishes three amplification factors on the same $(\mathbf{S}, \mathbf{P})$ design:

(i) the *structural* contamination multiplier $\phi = \text{tr}(\mathbf{QP})/\text{tr}(\mathbf{Q})$ of Proposition 1—the design's sensitivity to network propagation under the i.i.d. null, an ex-ante diagnostic;

(ii) the *realized* contamination factor—either the correlated-shocks object ϕ^{corr} of (10) under OLS or its IV analog ϕ^{IV} of Corollary 1—the actual bias of the estimator under the application's shock distribution;

(iii) the *aggregate amplification* ψ^{agg} of Corollary 3—the ratio of corrected to direct exposure at the population-weighted mean.

All three reduce to functions of $(\mathbf{S}, \mathbf{P})$ alone when shocks are i.i.d. and weights are uniform. They separate once the realized shock distribution and population weights are introduced, and the distance between the realized contamination factor and ψ^{agg} governs the gap between the coefficient ratio and the aggregate-impact ratio. The values taken by these objects in any particular application are properties of the realized design and are reported empirically; Section 4.4 gives the ADH values. Table 2 is a scorecard of every variant of these factors used in the paper, with its definition, weighting, share vintage, and ADH value.

The five diagnostic objects that the framework above produces are collected in Table 1, which also gives the ADH values that the empirical sections of the paper recover. The first three rows are the scalar amplification factors developed in this section; the fourth row summarizes the cross-region dispersion of the regional multiplier m_r that drives geographic incidence; and the fifth row reports the share of the standard estimator's MSE that comes from the cross-region variance of Δ_r^{std} rather than its population-weighted mean.

3 Data and Construction

3.1 The ADH Setting

We apply our methodology to the canonical Autor et al. (2013) China Shock analysis (hereafter ADH), which constituted a major shock to U.S. local labor markets. The unit of observation is a commuting zone (CZ) r in period t, with two stacked periods (1990–2000, 2000–2007) yielding $R = 722$ commuting zones observed twice ($N = 1,444$). The

main outcome is the change in manufacturing employment as a share of the working-age population ($\Delta\mathcal{L}_{rt}^{\text{mfg}}$, in percentage points[5]). The endogenous variable is U.S. import exposure from China per worker:

$$\Delta IPW_{rt}^{US} \;=\; \sum_{k} \frac{\mathcal{L}_{rk,t-1}}{\mathcal{L}_{r,t-1}} \cdot \frac{\Delta M_{kt}^{US}}{\mathcal{L}_{k,t-1}}, \tag{30}$$

a Bartik instrument with shares $s_{rk} = \mathcal{L}_{rk,t-1}/\mathcal{L}_{r,t-1}$ and shocks $g_k = \Delta M_{kt}^{US}/\mathcal{L}_{k,t-1}$.[6] Following Autor et al. (2013) we instrument with the analogous measure constructed from other high-income countries' imports from China,

$$\Delta IPW_{rt}^{OC} \;=\; \sum_{k} \frac{\mathcal{L}_{rk,t-1}}{\mathcal{L}_{r,t-1}} \cdot \frac{\Delta M_{kt}^{OC}}{\mathcal{L}_{k,t-1}}, \tag{31}$$

which isolates China's export-supply shocks from U.S. demand conditions. Two share vintages recur throughout the paper, and we fix the vocabulary here. The endogenous exposure uses *start-of-period* shares, built from County Business Patterns employment in 1991 for the first period and 1999 for the second; this matches ADH's own timing, in which the endogenous exposure is constructed from beginning-of-period employment. The instrument uses *decade-lagged* shares, from 1980 employment for the first period (imputed following Dorn (2009) at SIC87dd) and 1991 for the second; the decade lag keeps contemporaneous labor-market adjustments out of the instrument's shares, again following ADH. Trade changes are measured over 1991–1999 for the first period and 1999–2007 for the second (ADH's first-period window runs to 2000). Constructions that instead lag the shares on both sides—the fully predetermined variant used in several robustness exercises—are labeled as such wherever they appear, and Table 4 catalogs every construction. Controls include Census division fixed effects and the lagged CZ characteristics standard in Autor et al. (2013): manufacturing employment share, college-educated share, foreign-born share, female employment share, routine-occupation share, and an offshorability index. Regressions are weighted by start-of-period CZ population and standard errors are clustered at the state level. Three estimation samples recur, and we define them once. The estimation panel is all 722 commuting zones in both periods ($N = 1,444$). Objects built on the regional multiplier $m_r = \tilde{B}_r/B_r$ are defined on the positive-exposure support $B_r > 0$, which covers 1,433 of the 1,444 CZ-periods (714 CZs in the first period, 719 in the second). The period-2 ranking, policy-targeting, and out-of-sample exercises of Sections 5–6 use the 714 period-2 CZs with non-negligible direct exposure ($|B_r| > 0.01$); the bootstrap confidence-interval table reports its point estimates

[5] We use $\mathcal{L}$ for employment levels to distinguish from the Leontief inverse $\mathbf{L}$, which appears throughout.

[6] The U.S.-from-China flows ΔM_{kt}^{US} are taken directly from the Autor et al. (2013) replication files, which apply the Feenstra et al. (2005) cleaning that reassigns Chinese-origin goods transshipped through Hong Kong back to China. The reassignment raises ADH's 2007 aggregate China-import figure by roughly 16% relative to the raw U.S. Census country-of-origin total. The reassignment is empirically active in the data but does not affect the cross-industry distribution of shocks materially: cross-source validation against UN Comtrade, CEPII BACI, and WIOD-2013 yields Pearson correlations on changes of 0.999, 0.997, and 0.975 respectively at the SIC87dd industry level (Appendix C.4).

on the full 722-CZ period-2 sample, as its notes state.

3.2 The Propagation Matrix

Our main data source for the propagation matrix $\mathbf{P} = \mathbf{L}$ is the Bureau of Economic Analysis (BEA) 1997 detailed benchmark input-output tables (Bureau of Economic Analysis, 2002). These provide Make and Use matrices at the detailed industry level, from which we compute the industry-by-industry direct requirements matrix $\mathbf{A} = \mathbf{DB}$ (with $\mathbf{D}$ the market shares matrix and $\mathbf{B}$ the commodity direct requirements matrix) and the Leontief inverse $\mathbf{L} = (\mathbf{I} - \mathbf{A})^{-1}$, following standard BEA methods. For specifications with manufacturing employment as the outcome, we use the $K = 342$ manufacturing submatrix of $\mathbf{L}$; this captures both direct manufacturing-to-manufacturing linkages and indirect linkages mediated through non-manufacturing sectors (for example, steel reaching auto manufacturing indirectly through transportation services). We match BEA industries to ADH's SIC87dd codes via a concordance chain, matching 392 of 397 ADH industries; the five unmatched industries are dropped from both the share and the shock construction. Four industry granularities recur in the paper, and we fix the labels here: the BEA detailed manufacturing submatrix used for the headline ($K = 342$); the ADH SIC87dd industries to which it is concorded (392 of 397 matched); the BEA detailed full economy used for the three-regressor and full-economy checks ($K = 483$); and the BEA summary level used for the aggregation-robustness checks ($K = 53$).

An issue with our propagation matrix is the potential endogeneity of the IO coefficients: if Chinese competition causes U.S. producers to restructure supply chains, the matrix $\mathbf{A}$ is endogenous to the treatment (Acemoglu and Azar, 2020). Four considerations mitigate the concern. First, we use the 1997 benchmark IO table, which predates the bulk of the China Shock: the shock intensified over 2000–2007, so the IO structure is pre-determined for the second period of our panel. For the first period (1990–2000) the 1997 matrix is technically post-treatment, and pre-determination holds only in the weaker sense that the treatment's impact on the IO structure through 1997 was quantitatively small. As a direct check, re-estimating the main specifications on the clean 2000–2007 period alone, for which the 1997 table is unambiguously pre-determined, yields a standard coefficient of -0.542 and a corrected coefficient of -0.199 (standard/corrected ratio 2.73 on 722 commuting zones), preserving the standard-exceeds-corrected pattern of the pooled estimates. The paper's central geographic-incidence results are likewise robust to this clean period: re-running the reallocation suite and out-of-sample fit on 2000–2007 alone (training and evaluating on the single period) yields a geographic share of mean-squared prediction error of 99.9% (state-block bootstrap CI $[99.7, 100.0]$) and a corrected-versus-standard out-of-sample ranking ratio of 3.82, both in line with, if anything stronger than, the pooled values of Section 5. Due to the lack of a period-matched pre-1990 IO benchmark at the BEA detailed level, a backcast is not available: the available pre-1997 benchmarks are far coarser (roughly 65 sectors), and the aggregation sensitivity documented in Section 6.3 would confound such a 65-sector-versus-342-industry comparison with vintage drift, so we do not construct one. Second, IO structure is highly stable over time: the correlation between 1997 and 2002 Leontief-inverse entries is 0.994, and the vintage diagnostic $\hat{\phi}$ is nearly identical across the two benchmarks (manufacturing $\hat{\phi} \approx 1.10$ at both; Table 14 and Section 6.3 report the

full panel over 1997–2017). Third, vintage drift that does appear is not concentrated in China-shocked sectors: Section 6.3 reports that a cluster-level regression of per-cluster Leontief-drift on cluster-level China-shock exposure returns a negative coefficient that is economically negligible (though statistically distinguishable from zero), evidence against the reading in which the instrument propagates through supply chains severed by the shock itself. Fourth, any endogenous adjustment of the IO coefficients is itself a form of mismeasurement of the propagation matrix, whose consequence for the slope is governed by the same exact decomposition we use for the regional-IO concern (the covariance of the induced exposure error with the instrument; Appendix A.4) and is bounded directly in Section 6.7.

A separate concern, distinct from temporal endogeneity, is that the BEA national IO table may not accurately represent region-specific input-output structures. We construct $\mathbf{A}$ at the national level from the BEA Make-Use tables; if regions deviate, idiosyncratically or systematically, from the national IO structure (e.g., a primary-metals-heavy region in the Great Lakes uses a different supplier mix than a primary-metals-heavy region in the Pacific Northwest), the correction's network-adjusted exposure $\tilde{s}_{rl} = \sum_k s_{rk} L_{kl}^{\text{nat}}$ is constructed with measurement error in L_{kl} relative to the true $L_{kl,r}$. There are two responses. First, the consequence for the IV slope is governed by an exact object rather than a sign intuition: the network-adjusted coefficient is biased only to the extent that the *induced exposure error* $\widehat{X}_r - X_r^*$ (the gap between national-IO and true regional-IO exposure) covaries with the instrument, $\hat{\beta}/\beta - 1 = -\text{Cov}(\widehat{Z}, \widehat{X} - X^*)/\text{Cov}(\widehat{Z}, \widehat{X})$ (Appendix A.4). The direction is not signed in general, because $\tilde{s}_{rl}$ and its instrument are nonlinear in $\mathbf{A}$ through $(\mathbf{I} - \mathbf{A})^{-1}$, so mean-zero regional deviations in $\mathbf{A}$ need not attenuate the coefficient toward the standard, and we make no lower-bound claim. But the same decomposition yields a Cauchy–Schwarz sensitivity bound and a directly estimable bias diagnostic, both of which we evaluate in Section 6.7: benchmarked against state-specific intra-state Leontief inverses, the implied proportional slope bias is 10.36% with a worst-case bound of 17.69%. Second, the BEA's Regional Input-Output Modeling System (RIMS II) and IMPLAN regional multiplier files are state-level summaries that aggregate to roughly the national pattern, which suggests systematic regional deviations are quantitatively small at the resolution of the BEA-detail propagation matrix. Our approach is related to recent work using province-specific IO tables in non-U.S. settings (e.g. Gulek, 2026), which provides additional evidence on how regional IO heterogeneity affects shift-share propagation analyses; we treat the regional-specificity question as a robustness exercise—Section 6.7 bounds the consequences of national-versus-regional IO mismatch directly and finds them small—and retain the national-aggregate IO as the baseline. Section 6.7 reports two direct robustness exercises: an $\mathbf{A}$-perturbation simulation and a state-specific intra-state Leontief substitution using the EPA's state-level bi-regional Use tables, both of which leave the corrected coefficient close to its national-IO baseline.

3.3 Practical Interpretation of the Corrected Exposure

The network-inclusive Bartik instrument can be written as a shift-share with network-adjusted shares. To see this, start from $\tilde{B}_r = \mathbf{s}_r' \mathbf{L}\mathbf{g} = \sum_k s_{rk} [\mathbf{L}\mathbf{g}]_k$ and rearrange the summa-

tions:

$$\tilde{B}_r \;=\; \sum_{l=1}^{K} \tilde{s}_{rl}\, g_l, \qquad \tilde{s}_{rl} \;=\; \sum_{k=1}^{K} s_{rk}\, L_{kl}. \tag{32}$$

Each primitive shock g_l now enters CZ r with a weight $\tilde{s}_{rl}$ that aggregates direct ownership of industry l (the $k = l$ term, contributing $s_{rl}L_{ll}$) with the share of every industry k that supplies l's production chain, directly or indirectly (weighted by the Leontief coefficient L_{kl}). The correction therefore reweights each CZ's exposure to each primitive shock, upweighting CZs whose industries are dense in upstream or downstream linkages to the shocked industries. To better understand this result, consider a CZ that manufactures aluminum: it gains extra exposure to aircraft shocks because aircraft manufacturing is a major aluminum buyer; a CZ that manufactures semiconductors gains extra exposure to electronic-computer shocks for the same reason. The practical consequence is that two CZs with identical direct exposure B_r can have substantially different network-inclusive exposure $\tilde{B}_r$, with the gap reflecting local supply-chain density rather than direct final-goods overlap with Chinese imports.

3.4 Contamination Multiplier Diagnostics

Before estimating any regression we compute the ex-ante contamination multiplier from the share matrix and the Leontief inverse alone. With the decade-lagged share construction,

$$\hat{\phi} \;=\; \frac{\mathrm{tr}(\hat{\mathbf{Q}}\hat{\mathbf{L}})}{\mathrm{tr}(\hat{\mathbf{Q}})} \;=\; 1.110, \tag{33}$$

where $\hat{\mathbf{Q}} = \mathbf{S}'\mathbf{S}$ is the empirical co-location matrix. This estimate is stable across periods and vintages: $\hat{\phi} = 1.120$ using the 1980 lagged shares (period 1) and $\hat{\phi} = 1.110$ using the 1991 lagged shares (period 2), for a pooled decade-lagged value of 1.115; the pooled start-of-period shares give 1.113. These results suggest that ϕ reflects a structural feature of U.S. manufacturing geography and IO linkages rather than a period-specific artifact. The residualized multiplier $\hat{\tilde{\phi}}$ that accounts for the ADH controls and population weights (Corollary 4, Appendix A) is ≈ 1.08, modestly smaller than the raw diagnostic because the ADH controls are nearly orthogonal to the IO-weighted share structure. Both diagnostics indicate structural contamination weights of 8–11%: modest in isolation, but enough—when combined with application-specific causal coefficients and with the realized shock covariance documented in Section 4.4—to produce large coefficient differences in practice.

Table 14 in Section 6.3 reports $\hat{\phi}$ across every BEA detailed benchmark from 1997 to 2017 for both the manufacturing-only and full-economy IO systems. This relationship remains robust even after accounting for vintage: manufacturing $\hat{\phi}$ ranges from 1.10 to 1.14 across vintages, full-economy from 1.10 to 1.15.

4 Main ADH Results

4.1 Standard vs Corrected 2SLS

The coefficient estimates for both specifications are reported in Table 3: the standard Bartik (Panel A) and the single-regressor network-corrected Bartik (Panel B). The standard 2SLS coefficient is -0.612 (s.e. 0.066, first-stage $F = 61.7$), closely replicating the Autor et al. (2013) estimate of -0.596 (s.e. 0.100, $F = 47.6$). The single-regressor corrected coefficient (replacing the direct-exposure Bartik B_r with the Leontief-propagated $\tilde{B}_r = \mathbf{s}_r'\mathbf{L}\mathbf{g}$ and instrumenting with its other-country analog $\tilde{Z}_r = \mathbf{s}_r'\mathbf{L}\mathbf{g}^{OC}$) is -0.255 (s.e. 0.023, first-stage $F = 192.9$).[7] The coefficient ratio is 2.40. The corrected first-stage F is roughly three times the standard's (and about four times the $F = 47.6$ of the ADH replication). This reflects the construction rather than stronger identification: pre-multiplying both the endogenous exposure and its other-country instrument by the same Leontief operator raises their correlation, so the larger F is a mechanical consequence of the shared transformation. Identification continues to rest on the exogeneity of $\mathbf{g}^{OC}$ in the BHJ regime, and the same caveat applies to the two-regressor first stages in Section 4.2.

The headline builds the endogenous Bartik from start-of-period (1991/1999) employment shares and instruments it with the decade-lagged (1980/1991) other-country Bartik—ADH's own timing, in which the endogenous exposure uses beginning-of-period employment and only the instrument is lagged (Section 3.1). Because the instrument is predetermined and shift-exogenous, identification does not require the endogenous regressor to be lagged: the start-of-period shares measure the industrial composition through which the shock actually propagates, and instrumenting them with a lagged other-country analog is the standard remedy for composition endogeneity. The choice does affect the level of the ratio. The fully predetermined variant, with decade-lagged shares on both sides at BEA detail, gives a standard coefficient of -0.266 and a corrected coefficient of -0.165, a ratio of roughly 1.6, against the ADH-timing headline ratio of 2.40. The qualitative result is the same in both: the correction shrinks the coefficient and the realized contamination factor exceeds one. The difference between the two ratios is itself an instance of the $\phi^{\text{iid}} \to \phi^{IV}$ mechanism, in that the realized factor depends on the covariance of the instrument and endogenous shares with the shock, which the share vintage shifts, while the i.i.d. diagnostic ϕ^{iid} is vintage-robust (Section 4.4). Table 4 lays out every construction (ADH replication, BEA detail and summary, full economy) with its diagnostic and its standard and corrected coefficients.

The coefficient progression across columns (OLS, OLS + lagged CZ controls, OLS + Census division fixed effects, 2SLS, 2SLS + state fixed effects; Table 3) shows that the corrected estimator is substantially more stable under the standard menu of control variables than the standard estimator. Comparing coefficient estimates across the menu, we observe large movements in the standard coefficient – from -0.406 (simple OLS) to -0.596 (2SLS) to -0.582 (2SLS with state FE). The corrected coefficient moves from -0.222 to -0.255 to

[7] The corrected point estimate is sensitive to the input–output vintage used to build $\mathbf{L}$: it is -0.255 on the 1997 benchmark and -0.374 on the 2002 benchmark (a movement of roughly 47%; Section 6.3, Table 14). The contamination diagnostic $\hat{\phi}$ and the sign and direction of the correction are stable across vintages; it is the level that moves, and our claims rest on the direction and the diagnostic.

-0.252 across the same progression. The additional variation in the standard coefficient reflects the usual OLS-versus-IV gaps (attenuation from measurement error, demand-side endogeneity) amplified by the omitted network channel; the corrected coefficient, with its more complete exposure map, is better behaved.

Main-table standard errors cluster at the state level, matching Autor et al. (2013). State clustering captures regional dependence (across-CZ correlation induced by shared state-level labor-market shocks and common regional business cycles) that is orthogonal to the shift-share exposure structure and therefore not picked up by exposure-robust variance formulas alone. The Adão et al. (2019) exposure-robust variance for the corrected estimator is reported in Section 6.4 for comparison. The correctly implemented exposure-robust standard errors come out modestly larger than the state-clustered ones for both estimators, and all main-specification t-statistics remain highly significant under either method.

Figure 1 plots standard against network-adjusted exposure at the CZ level (each point a CZ-period); the points lie broadly along but meaningfully off a constant-multiple line. What stands out from the figure is not how much the two measures track each other but instead how much they diverge: a simple rescaling of the standard Bartik cannot reproduce the network-adjusted measure, because the ratio $m_r = \tilde{B}_r / B_r$ varies with CZ composition (Section 5). Figure 2 translates the exposure transformation into a reduced-form binscatter: the standard and network-adjusted other-country instruments are overlaid against the change in manufacturing employment share, with the corrected fit displaying tighter dispersion around its slope; Section 5.6 quantifies this predictive edge on held-out geographies.

The three-regressor 2SLS at full-economy granularity ($K = 483$), with direct, upstream-only, and downstream-only exposure regressors instrumented jointly by their other-country analogs, is reported in Table A.3. The Section 4.2 two-regressor decomposition that aggregates upstream and downstream into a single network channel is the specification we develop in detail in the body; the three-regressor refinement and the full-economy specifics are deferred to Appendix B and Appendix C.

4.2 Two-Regressor Direct + Network-Only Decomposition

The single-regressor corrected estimator aggregates the direct and network-only channels into a single coefficient (Proposition 3). The two-regressor corrected estimator separates them by adding the network-only exposure $N_r = \mathbf{s}_r'(\mathbf{L} - \mathbf{I})\mathbf{g}$ as a second endogenous regressor, instrumented by its other-country analog (Definition 1, two-channel form).

The two-regressor OLS (column 2) and 2SLS (column 3) estimates are reported in Table 5 alongside the standard specification (column 1). The 2SLS direct coefficient is $\hat{\beta}_D = -0.305$ (s.e. 0.107) and the network-only coefficient is $\hat{\beta}_N = -0.233$ (s.e. 0.050). Both Sanderson–Windmeijer (Sanderson and Windmeijer, 2016) first-stage F-statistics exceed conventional weak-IV thresholds: $F = 24.3$ on B_r and $F = 219.7$ on N_r. The large F on N_r in part reflects the near-deterministic mapping from the other-country instruments to the network-adjusted regressors (both are linear transforms of the same shock vectors) rather than an independent measure of identification strength; identification rests on the exogeneity of $\mathbf{g}^{OC}$ in the BHJ regime, not on the F magnitude (Appendix C discusses this mechanical near-determinism for the three-regressor specifications). The direct coefficient is about half

the standard Bartik's -0.612, and the network-only coefficient is of comparable magnitude to the direct: roughly half of the standard coefficient's magnitude loads onto exposure that propagates through the supply chain rather than onto own-industry displacement. A Wald test of the single-index restriction $\beta_D = \beta_N$ does not reject at conventional levels ($t = -0.48$, $p = 0.63$, with $\hat{\beta}_D - \hat{\beta}_N = -0.072$ and clustered s.e. 0.149): the data do not refuse the simpler structure that direct and network channels carry the same per-unit effect, although the two-regressor estimates with substantial weight on each channel are also consistent with the alternative. The two-channel OVB formula $\hat{\beta}_D + \hat{\beta}_N(\hat{\phi} - 1) \approx -0.305 + (-0.233)(0.11) \approx -0.33$ understates the realized standard coefficient of -0.612. This discrepancy could be attributed to the fact that the ADH shock vector is not i.i.d. but instead exhibits substantial within-sector correlation (Chinese import growth is clustered at the three-digit NAICS level), which amplifies the realized IV contamination factor above the structural i.i.d. diagnostic. Section 4.4 quantifies this amplification directly.

The identification regime for the two-regressor estimator is the Borusyak et al. (2022) (BHJ) shift-exogeneity regime: the other-country shocks $\mathbf{g}^{OC}$ provide an independent source of identifying variation on top of the shares, yielding two linearly independent columns in the instrument matrix. To better understand the role of this independent variation, consider the polar case where identification comes from Goldsmith-Pinkham et al. (2020) (GPS) share exogeneity alone, with a single realized shock vector: the two regressors are then fixed, nearly collinear transformations of the same shares matrix, so the split—though generically identified in the formal rank sense—is weakly identified in practice (Section 2.5). ADH's other-country instrument places the design firmly in the BHJ regime.

4.3 Interpreting the Channels

The two-regressor estimates partition the standard Bartik into a direct-displacement coefficient and a network-only coefficient that are precisely identified, individually significant, and economically distinct. The direct effect $\hat{\beta}_D = -0.305$ measures own-industry displacement among workers in CZs whose 1990 industry composition overlaps with the China-competing final-goods targets (apparel, footwear, electronic computers). The network-only effect $\hat{\beta}_N = -0.233$ captures employment loss in CZs that are exposed to the shock through their supply-chain neighbors rather than through their own industry mix: steel-mill and chemical regions that supply automotive and consumer-goods customers, foundries that supply electronics, and similar upstream-supplier patterns. Both coefficients are negative and of comparable magnitude. The standard Bartik's single coefficient, by combining the two, conflates direct and propagated exposure under a single weight that is structurally a mixture (Proposition 1) rather than a substantive average treatment effect.

There are two design features that make the (B_r, N_r) decomposition well-conditioned. First, the network-only regressor $N_r = \mathbf{s}_r'(\mathbf{L} - \mathbf{I})\mathbf{g}$ subtracts the identity from the Leontief inverse before pre-multiplying the shock vector, separating the within-industry roundabout multiplier from the strict cross-industry propagation that the direct Bartik does not already contain. Second, the BHJ shift-exogeneity regime supplies an independent column in the instrument matrix, the other-country analog $\mathbf{g}^{OC}$, so the rank condition for joint estimation is met. The payoff is visible in a direct comparison with the augmented-

regressor approach of Acemoglu et al. (2016) (Acemoglu, Autor, Dorn, Hanson, and Price; henceforth AADHP), who capture network propagation by adding a separately constructed upstream net-input IO exposure alongside the direct import-penetration measure rather than orthogonalizing the two channels. Because the upstream measure propagates the same import shock through the input–output network, it is nearly collinear with direct exposure (correlation 0.988 on the ADH data), and the augmented 2SLS is only weakly identified (the Sanderson–Windmeijer first-stage F-statistics are 5.1 and 5.1, both below conventional thresholds), returning the large, nearly canceling coefficients -2.051 and $+1.722$ reported in Table 5 columns 4–5. Subtracting the identity is precisely what removes the directly overlapping component from the network channel: the resulting N_r is well-separated from B_r (Sanderson–Windmeijer F of 24.3 and 219.7), so the two channels are each precisely identified. Our specification differs from AADHP, who model network propagation as a separately constructed upstream net-input measure: the advantage is one of construction and conditioning rather than of identifying information. Our network channel is a single Leontief-derived regressor N_r, differently constructed from AADHP's separately-built net-input measure, and within our own design (B_r, N_r) and $(B_r, \tilde{B}_r)$ are exact reparameterizations with identical fitted values. The same shift-exogeneity content, resting on $\mathbf{g}^{OC}$, is thus expressed in a better-conditioned, interpretable basis; we do not claim to extract more from the data than the augmented specification, only to recover stable, separately interpretable channel coefficients. This is the construction we use throughout.

Two finer disaggregations of the network-only channel are possible and clarify the economic mechanism. Splitting N_r into upstream (Leontief) and downstream (Ghosh) components recovers separate coefficients on supplier- and customer-mediated exposure under the same BHJ identification; further splitting each network channel by whether the originating industry is manufacturing or non-manufacturing isolates the Pierce et al. (2024) input-cost benefit channel from local-supply-disruption and local-multiplier mechanisms. We find that the dominant downstream effect is negative and runs through manufacturing customers losing their U.S. suppliers, with the Pierce–Schott input-cost benefit operating only through non-manufacturing customers and not large enough to overturn the net downstream sign. The full three-regressor and five-regressor decompositions are reported in Appendix B, alongside the AADHP confrontation and the supporting Monte Carlo and condition-number diagnostics. The body works with the two-regressor split throughout.

4.4 The Realized Contamination Factor

The i.i.d. prediction from the two-channel OVB formula (Proposition 1), evaluated at the estimated channel coefficients and the structural contamination multiplier $\hat{\phi} = 1.110$, is $\hat{\beta}_D + \hat{\beta}_N(\hat{\phi} - 1) \approx -0.305 + (-0.233)(0.11) \approx -0.33$. The realized standard 2SLS coefficient is -0.612. What accounts for the seeming discrepancy between the i.i.d. prediction and the realized coefficient? Corollary 1 locates its source in the realized IV contamination factor

$$\phi^{IV}(\mathbf{g}^{OC}, \mathbf{g}^{US}) \;=\; \frac{(\mathbf{g}^{OC})'\mathbf{Q}\mathbf{L}\mathbf{g}^{US}}{(\mathbf{g}^{OC})'\mathbf{Q}\mathbf{g}^{US}}, \tag{34}$$

whose raw sample counterpart, evaluated on the ADH shock vectors through the canonical exposure columns, is 2.39 unweighted and 2.67 under the population weights of the 2SLS.[8] Inside the estimated regressions the same object appears in two guises, both reported in Table 2. The controls-residualized, population-weighted 2SLS projection of N_r on B_r gives a realized factor of 2.316—the value that enters the finite-sample OVB identity of Lemma 1 and the number we quote as $\hat{\phi}^{IV}$ throughout. And the single-regressor coefficient ratio $\hat{\gamma}_{SS}/\hat{\gamma}_{NSS} = 2.403$ is, by construction, the realized factor of the single-regressor comparison. The realized IV factor is about 2.1× the structural i.i.d. diagnostic $\phi^{\text{iid}} = 1.110$. It is this gap that accounts for the coefficient ratio: the residual between the i.i.d. prediction -0.33 and the realized -0.60 is absorbed exactly by the $\phi^{\text{iid}} \to \phi^{IV}$ amplification.

In this section we define two channels through which the amplification operates. The OLS correlated-shocks factor $\phi^{\text{corr}} = (\mathbf{g}^{US})'\mathbf{Q}\mathbf{L}\mathbf{g}^{US}/(\mathbf{g}^{US})'\mathbf{Q}\mathbf{g}^{US} = 1.87$ is what the bias would be under OLS using the U.S. shock vector itself: this captures the effect of shock correlation in $\mathbf{g}^{US}$ alone, abstracting from the IV mechanism. The realized 2SLS factor $\phi^{IV} = 2.316$ is materially larger but only modestly so. The shock-correlation contribution $\phi^{\text{iid}} \to \phi^{\text{corr}}$ accounts for a multiplicative factor of $1.87/1.110 \approx 1.68$, while the IV-vs-OLS mechanism $\phi^{\text{corr}} \to \phi^{IV}$ accounts for $2.316/1.87 \approx 1.24$. Most of the gap between structural and realized contamination is therefore attributable to the within-NAICS-3 clustering of Chinese import shocks rather than to the IV-instrument-vs-endogenous-shock cross-covariance: an OLS shift-share regression on $\mathbf{g}^{US}$ would already exhibit roughly 1.7× the structural diagnostic, with the 2SLS IV contributing the residual amplification.

The realized ϕ^{IV} decomposes by propagation order. From the Neumann series $\mathbf{L} = \mathbf{I} + \mathbf{A} + \mathbf{A}^2 + \cdots$, the contamination multiplier admits the decomposition $\phi - 1 = \sum_{n=1}^{\infty} \phi^{(n)}$ with $\phi^{(n)} = \text{tr}(\mathbf{Q}\mathbf{A}^n)/\text{tr}(\mathbf{Q})$ (Proposition 9(ii) in Appendix A); we refer to $\phi^{(n)}/(\phi - 1)$ as the *Neumann-order share* of order n, the fraction of the total amplification carried by n-step propagation through the IO network. The Neumann-order shares under both the i.i.d. weighting and the ADH realized IV weighting are reported in Table 6; Figure 3 plots $|\phi^{(n)}|$ on a log scale against order n, superimposing the three series (i.i.d., ADH realized, BK realized) on the same axes. Under the population-weighted i.i.d. weighting, $\phi = 1.164$, with $\phi^{(1)} = 0.1237$ accounting for 75.5% of the $\phi - 1$ amplification and the first two orders accounting for 92.7%. Under the ADH realized weighting, $\phi^{\text{ADH}} = 2.262$, with $\phi^{(1),\text{ADH}} = 0.8034$ accounting for 63.6% and the first two orders accounting for 87.1%. The $\phi^{\text{iid}} \to \phi^{IV}$ amplification is present already in the first-order term, where $\phi^{(1),\text{ADH}}/\phi^{(1),\text{iid}} \approx 6.5$ is close to the overall realized/i.i.d. ratio: the amplification is not concentrated in exotic high-order network terms. The geometric decay—at order-to-order ratios below the spectral-radius bound $\rho(\hat{\mathbf{A}}) = 0.506$, which caps the limiting rate but not the transient one—is mechanical: by the $n = 5$ order, 99.7% (i.i.d.) or 99.3% (ADH) of the series is captured, and the higher-order tail is numerical finish rather than economic

[8] Although the factor is written with a single co-location matrix $\mathbf{Q}$, the headline builds the endogenous exposure from start-of-period shares and the instrument from decade-lagged shares. The weighted factor is exactly the cross-vintage matrix object $(\mathbf{g}^{OC})'(\mathbf{S}_Z'\mathbf{W}\mathbf{S}_X)\mathbf{L}\mathbf{g}^{US}/(\mathbf{g}^{OC})'(\mathbf{S}_Z'\mathbf{W}\mathbf{S}_X)\mathbf{g}^{US}$ in the decade-lagged instrument shares $\mathbf{S}_Z$, start-of-period endogenous shares $\mathbf{S}_X$, and population-weight matrix $\mathbf{W}$: the matrix form reproduces the region-level factor 2.67 to numerical precision (difference $4.4e-16$). The i.i.d. diagnostic is vintage-robust—$\phi^{\text{iid}} = 1.115$ on decade-lagged shares versus 1.180 on the cross-vintage weighted object.

content.

A complementary view of the same amplification is provided by the Rotemberg-weight decomposition. Under Goldsmith-Pinkham et al. (2020), the Bartik 2SLS equals a GMM estimator using industry shares as instruments, with each share s_{rk} delivering a just-identified IV estimate $\hat{\beta}_k$ and the overall Bartik a weighted average with Rotemberg weights α_k^{std}. The corrected Bartik admits an analogous decomposition with network-adjusted weights α_k^{corr} derived from the network-adjusted shares $\tilde{s}_{rl} = \sum_k s_{rk} L_{kl}$. The two weight distributions are compared in Table 7 and Figure 4. Top standard weights concentrate on direct China-competing final-goods targets (electronic computers, cut-and-sew apparel, dolls and toys, footwear); top corrected weights shift toward upstream suppliers of those targets (semiconductors, other electronic components, iron and steel mills). The effective number of identifying industries is similar between the two specifications (10.0 for the standard and 8.9 for the corrected under the inverse-Herfindahl count), but the Pearson correlation between the two weight distributions is only 0.596 and the Spearman rank correlation 0.691. The corrected estimator identifies off a substantially different set of industries than the standard, consistent with the qualitative interpretation that the correction is not a rescaling of the standard Bartik but a genuine reweighting toward the propagating industries that the standard exposure measure omits.

The ADH design therefore produces three amplification factors, each carrying a distinct economic interpretation. The structural diagnostic $\hat{\phi}^{\text{iid}} = 1.110$ depends only on $(\mathbf{S}, \mathbf{L})$ and gives the standard Bartik's network-channel overweight under a generic shock distribution—about 11%. The realized IV factor $\hat{\phi}^{IV} = 2.316$ translates the structural overweight into the observed coefficient ratio of 2.4 once shock covariance is incorporated, because the actual Chinese import shocks are correlated at the three-digit NAICS cluster level (apparel together, electronics together) and those clusters are both co-located and IO-linked. The aggregate amplification $\hat{\psi}^{\text{agg}} = 2.766$ is the population-weighted average of the propagated-to-direct exposure ratio: 2.8. Table 8 collects these and the other quantities most relevant to interpreting the correction. Per-unit coefficients differ by a factor of 2.40 (standard -0.612 vs corrected -0.255); exposure at the interquartile range differs by a factor of roughly $4.70/1.74 \approx 2.7$; the two factors nearly cancel when combined into predicted impact at the IQR (standard -1.063 vs corrected -1.196 percentage points, a ratio of 0.89, a gap of 11.2%). Top-K overlap between standard and corrected CZ rankings is imperfect (34% of the top 100 and 68% of the top 25 differ between the two measures), and the corrected ranking produces an out-of-sample R^2 of 0.262 against realized manufacturing employment changes, more than three times the standard's 0.080. Taken together, these results suggest that the correction reallocates measured exposure across places without changing the aggregate; the next two sections develop each half of that statement.

5 Aggregate Impacts and Geographic Misallocation

We have seen that the standard Bartik's single coefficient conflates the direct and network-propagated channels (Section 4), producing a 2.40-fold gap between standard and corrected coefficients. This section translates the coefficient-level contamination into predicted impacts at the commuting-zone and aggregate levels, documents the empirical aggregate

near-cancellation between standard and corrected predictions in ADH, accounts for it through the covariance-identity decomposition of Proposition 6, and then turns to the spatial distribution of measured impact, which is where the correction's economic content lies. Autor et al. (2025) review the trade-shock shift-share literature and explicitly frame the standard ADH coefficients as estimates of relative effects across CZs that abstract from cross-CZ spillovers; the network-OVB framework here makes one such spillover mechanism (propagation through the IO network) explicit and shows that, although the across-CZ spillover redistributes which regions absorb the shock, it leaves the population-weighted aggregate close to the standard's prediction in this application. Aggregate cancellation throughout this section compares the population-weighted sums of the two estimators' partial-equilibrium predicted values, $T^{\text{std}} = \sum_r w_r \hat{\gamma}_{SS} B_r$ and $T^{\text{corr}} = \sum_r w_r \hat{\gamma}_{NSS} \tilde{B}_r$. It does not claim that the standard Bartik recovers the true general-equilibrium macroeconomic response to the China Shock: that quantity would require GE closure across national prices, monetary policy, and aggregate labor reallocation, which a cross-sectional reduced-form estimator cannot deliver. The relevant statement is narrower. The paper's two estimators, both applied in partial equilibrium, produce nearly identical aggregated predicted values in this particular application because of the $(\mathbf{Q}, \mathbf{P}, \mathbf{w})$ alignment documented below.

5.1 CZ-Level and Aggregate Predicted Impacts

Applying Proposition 5 to ADH, the CZ-level predicted impact under the standard estimator is $\hat{y}_r^{\text{std}} = \hat{\gamma}_{SS}\, B_r$, and under the single-regressor corrected estimator is $\hat{y}_r^{\text{corr}} = \hat{\gamma}_{NSS}\, \tilde{B}_r$. By Corollary 2, the standard estimator's per-region prediction error decomposes as $\Delta_r^{\text{std}} = \beta_N B_r(\phi - m_r)$, and every aggregate and geographic statistic reported below is a moment of this object across commuting zones. Evaluated at the interquartile range of each exposure distribution, the standard estimator implies an effect of -1.063 percentage points (the product of $-0.612 \cdot 1.74$); the single-regressor corrected estimator implies -1.196 percentage points (the product of $-0.255 \cdot 4.70$). Despite a per-unit coefficient ratio of 2.40, the per-IQR predicted impacts differ by only 11.2%, with the standard slightly smaller in magnitude than the corrected (0.89 of the corrected impact). What this says is that coefficient-level overstatement and exposure-scale understatement nearly offset at the interquartile range.

Figure 5 plots the CZ-level predicted impacts: each point is a commuting zone with horizontal axis $\hat{\gamma}_{NSS}\, \tilde{B}_r$ and vertical axis $\hat{\gamma}_{SS}\, B_r$, with point size proportional to population. The points scatter around the 45-degree line rather than along a constant-multiple line: individual CZs differ markedly between standard and corrected predicted impacts even though the population-weighted aggregates nearly coincide. This is the pattern the CZ-level overstatement ratio (22) predicts, with a constant numerator $(1 + \theta(\phi - 1))$ and a denominator $(1 + \theta(m_r - 1))$ that varies across CZs through the effective multiplier m_r.

5.2 The Covariance Identity and Why Aggregates Cancel

An exact condition for the aggregate to cancel was established in Section 2.6 (Proposition 7): on the internally consistent, population-weighted design objects, $\phi^{IV} - \psi^{\text{agg}} =$

$\mathrm{Cov}_p(m_r, Z_r)/\mathbb{E}_p[Z_r]$, so the aggregate cancels exactly if and only if the regional multiplier m_r is uncorrelated with instrument exposure Z_r under the exposure measure $p_r \propto w_r B_r$. Applied to ADH with $\mathbf{P} = \mathbf{L}$, the identity holds to numerical precision and the design nearly satisfies the condition: $\phi^{IV} = 2.67$ against $\psi^{\text{agg}} = 2.77$, a gap of -0.095 (an implied -3.5% of ψ^{agg}) reproduced exactly by $\mathrm{Cov}_p(m_r, Z_r) = -0.283$ (correlation -0.10), well inside the dispersion bound (29). The covariance-identity decomposition of Proposition 6 gives the same picture from the ψ^{agg} side: $\psi_L^{\text{agg}} = 2.766 = 3.214 + -0.448$, a meaningfully negative covariance pulling ψ^{agg} down from the unconditional multiplier $\bar{m}_L^{\text{pop}}$ toward the realized contamination.

Two conventions for the gap should be kept distinct. The theorem-governed comparison places both factors on the same population-weighted design objects (the ϕ^{IV} above) and gives the -3.5% gap. When computed instead from the fitted coefficient ratio $\hat{\gamma}_{SS}/\hat{\gamma}_{NSS}$ (the controls-residualized realized factor) against ψ^{agg}, the signed gap is -13.1%, statistically indistinguishable from zero (state-cluster bootstrap $p = 0.184$). Both comparisons agree that the aggregate near-cancels; the design-consistent comparison is the tighter, and is the one the exact identity governs.[9] Equivalently, in terms of the aggregate prediction error of equation (25), $\sum_r w_r \Delta_r^{\text{std}} = \beta_N \bar{B}_w (\phi^{IV} - \psi^{\text{agg}})$ is near zero in ADH despite each amplification factor on the right-hand side individually exceeding two.

Figure 6 visualizes the underlying (B_r, m_r) relationship: each point is one CZ-period with the population-weighted regression overlaid. The population-weighted regression slope of m_r on B_r is unambiguously negative (this is $\mathrm{Cov}_{\text{pop}}(m_r, B_r)/\mathrm{Var}_{\text{pop}}(B_r)$, the object that signs the aggregate gap in the alternative form of (27)), so high-multiplier supplier regions carry systematically below-average direct exposure. We now turn to a more detailed analysis of this negative covariance in order to further study its geographic content, through the corresponding spatial reallocation pattern (Sections 5.4 and 5.5).

5.3 When the Alignment Fails

The structural-alignment condition does not hold universally. Splitting the ADH sample by a CZ-level IO-connectivity measure (the employment-share-weighted average Leontief column sum, the object that drives the regional multiplier m_r) into terciles yields realized coefficient ratios of 3.347 (low-IO), 2.989 (medium), and 1.113 (high), against nearly constant aggregate multipliers ψ_L^{agg} of 2.658, 2.787, and 2.786. In the low and medium terciles the two objects nearly coincide. By contrast, in the high-IO tercile the realized coefficient ratio ϕ^{realized} collapses while ψ^{agg} stays at its structural level, producing a substantial implied overstatement.

The failure is consistent with the covariance identity. In the high-IO subsample, upstream-supplier CZs dominate, compressing the cross-CZ variation in m_r and weakening the negative covariance that elsewhere pulls ψ^{agg} down toward ϕ. The same identity diagnoses the failure, in that computing $\mathrm{Cov}_{\text{pop}}(m_r, B_r)$ within any candidate subsample

[9] The residualized ψ^{agg} that would exactly pair with the controls-residualized coefficient ratio is ill-posed: after partialling out the controls the population-weighted denominator $\sum_r w_r B_r^{\text{res}}$ is numerically zero, so the coefficient-ratio comparison necessarily pairs a raw-aggregate ψ^{agg} with a residualized realized factor. Proposition 7 is stated on the internally consistent (both-raw, population-weighted) objects.

reveals whether the cross-region alignment that produces near-cancellation in the pooled sample is also present in the subsample. The empirical content of the high-IO tercile is that this alignment condition is non-generic: it holds in the pooled ADH design but fails within an internally meaningful sub-design. In the language of Proposition 7, the high-IO tercile is exactly where the exact condition $\mathrm{Cov}_p(m_r, Z_r) = 0$ fails: the multiplier–instrument correlation moves from 0.05 in the low-IO tercile to -0.22 in the high-IO tercile, and the corresponding design-consistent aggregate gap widens from near zero to -0.208. This connectivity split is deliberately different from the tercile split used for the coefficient-heterogeneity analysis of Section 6.5, which sorts CZs by their employment share in high-IO-density industries; each split matches the object it diagnoses (the multiplier here, the coefficient there), and the two tercile assignments do not coincide.

5.4 Mean-Squared Prediction Error and the Geographic Map

The standard and corrected estimators disagree sharply on per-unit coefficients (Section 4) but agree at the aggregate. Their remaining point of disagreement, which is large, is about which communities bear the cost of the China shock. Equivalently, by Corollary 2 the standard estimator's mean-squared prediction error is $\mathrm{MSE}(\Delta_r^{\mathrm{std}}) = \mathbb{E}_r[\beta_N^2 B_r^2 (\phi - m_r)^2]$, which decomposes into (i) the square of a constant mean overstatement (governed by $\phi - \psi^{\mathrm{agg}}$, which is near zero in ADH) and (ii) the cross-region variance of Δ_r^{std} driven by region-specific $(\phi - m_r)$. The latter is the geographic-misattribution component. The truth T_r is unobservable, so the measured object is the gap between the two estimators' predictions, $\mathrm{pe}_r = \hat{\gamma}_{SS} B_r - \hat{\gamma}_{NSS} \tilde{B}_r$; under the maintained consistency of the corrected estimator this proxies Δ_r^{std}, and the theoretical decomposition applies to it term by term. The decomposition yields the paper's central empirical finding: mean overstatement accounts for less than 1% of the MSE, and 99.8% is cross-CZ reallocation. This share is precisely estimated. A state-block subsampling bootstrap (1,000 replications, each dropping 4 of 48 states without replacement and refitting, with the pivotal rescaling of Politis and Romano, 1994) places its 95% confidence interval entirely above 99%, so the geographic-reallocation share is sharply distinguishable from any materially smaller value rather than a knife-edge point estimate. At $b/n \approx 0.92$ the resampling scheme sits closer to a delete-d jackknife than to small-b subsampling, so we read these intervals as approximate. The root mean-squared prediction error across commuting zones is about 1.05 percentage points, comparable in magnitude to the typical CZ-level predicted impact itself, and essentially all of this error is geographic reallocation rather than scale error.

Figure 7 maps the standard and network-adjusted exposure measures across CZs; Figure 8 maps the ratio of the two at full width, since the ratio carries the economic content. The standard map reproduces the familiar ADH pattern, with exposure concentrated in the textile, apparel, and furniture belt of the Carolinas, Appalachia, and the South, and additional hot spots in the industrial Midwest's electronic-computer and metal-goods clusters. The network-adjusted map shows a broadly similar but not identical geography: regions across the Great Plains, the Gulf Coast, and the industrial Midwest exhibit substantially higher exposure than the standard measure suggests. What the ratio map reveals is that the highest-ratio CZs specialize in upstream-supplier process industries

(petroleum refining, primary metals, basic chemicals, paper), industries whose output feeds many downstream manufacturers but whose own output is not a primary target of Chinese imports. The lowest-ratio zones are the traditional China Shock belt specializing in cut-and-sew apparel, footwear, hosiery, and furniture: directly targeted but with simple supply chains. The geographic pattern of network amplification depends not on the average Leontief multiplier (which is broadly similar across CZs in levels) but on how densely each CZ's direct-exposure composition is embedded in the supply-chain network, the same economic content that produces the negative covariance term documented above.

5.5 Rank Changes and Top-K Overlap

Of the 100 commuting zones most exposed under the standard measure, 34% are not in the corrected top 100; among the top 25, 68% are not in the corrected top 25. Figure 9 plots each CZ's standard-measure rank against its corrected-measure rank. The Spearman rank correlation is 0.87 (strong, but far from perfect), with a mean absolute percentile shift of 10.4 points (Table 10) and individual CZs moving by as many as 505 of 714 rank positions, about 71 percentile points (Table 9). Table 9 lists the top and bottom ten CZs by rank change, and the comparison shows that industrial Midwest and Great Plains CZs with heavy exposure to upstream-supplier industries (steel, chemicals, petroleum, paper) gain rank, while textile-belt CZs in the Carolinas lose rank. State-block subsampling bootstrap confidence intervals for these reallocation statistics (top-K targeting error, mean percentile shift, the Spearman correlation, and the geographic share of MSE) are reported in Table 10.

Regional RMSPE patterns reinforce the story. The textile-specialized East South Central region (Tennessee, Alabama, Mississippi, Kentucky) has the largest regional RMSPE at 2.95 percentage points, driven by high standard-measure ranks and much lower corrected-measure ranks. By contrast, the IO-dense East North Central region (the Great Lakes manufacturing belt) has a much smaller RMSPE of 0.66 percentage points, consistent with its industrial structure producing closer correspondence between the two measures.

5.6 Out-of-Sample Predictive Fit

If the geographic reallocation is real, the corrected measure's ordering of CZs should correlate more strongly with realized manufacturing employment changes than the standard's on held-out geographies. We now explore this possibility by examining the naive bivariate R^2 of each exposure measure against realized manufacturing employment changes, evaluated across held-out Census divisions and states. This is an out-of-sample predictive-stability test rather than a causal-validation test. These results must be interpreted with caution, because richer exposure maps can mechanically improve predictive fit by smoothing the shock distribution across industries, and we do not interpret the R^2 gain as proof that the correction's coefficient is the true causal effect; for that, the theoretical OVB argument of Section 2 and the finite-shock jackknife of Section 6.2 are the relevant diagnostics. What the predictive check does tell us is whether the corrected ordering remains a better ranking of CZs than the standard ordering when we hold out geographies entirely from the estimation step, and a mechanical-smoothing story that merely reweights exposure

should still degrade when tested on held-out regions. We find that the corrected measure achieves $R^2 = 0.262$, more than three times the standard measure's $R^2 = 0.080$. The gain is geographically consistent: a leave-one-Census-division-out cross-validation produces a minimum ratio of $3.07\times$ across the nine divisions, and a leave-one-state-out cross-validation produces a minimum ratio of $3.06\times$ across the 48 continental states, with the corrected estimator beating the standard estimator in every division and every state fold. Summed across CZs, the two estimators produce a similar total; the corrected measure is better at identifying which CZs carry that total on held-out samples.

This predictive gain is a finding about this design and outcome rather than a generic property of network correction. In the immigration application in the online appendix, the inflow-propagated corrected exposure produces a negative leave-one-state-out R^2 (-0.5393, against -0.1490 for the standard measure, both worse than the sample mean), so a richer exposure map does not always rank held-out outcomes better. Whether the correction improves out-of-sample ranking depends on whether the propagated channel carries real signal for the outcome in question; it does so for ADH manufacturing employment and does not for that immigration outcome. We read the ADH result accordingly, as evidence on this application rather than as a universal claim, and report R^2 levels (not merely the $3\times$ ratio) precisely so the reader can judge predictive content rather than a ratio of two small numbers.

To bolster the case that the corrected measure's edge is not mechanical smoothing, we present two additional results. First, partialling the controls and Census-division fixed effects out of both the outcome and the exposures (a Frisch–Waugh residualization) leaves the corrected measure explaining $3.81\times$ more of the residual outcome variation than the standard (0.112 versus 0.030, squared population-weighted correlations): the edge survives removing the variation the controls already explain, so it is not a relabeling of the control set. Second, a placebo that uses the same exposures to predict pre-period manufacturing-employment changes (1970–80 and 1980–90) finds near-zero predictive content for the corrected measure (0.040 and 0.005, against 0.383 in the China-shock period) and no systematic advantage over the standard, which indicates the correction is capturing China-mediated propagation rather than serving as a smoother proxy for generic manufacturing decline.

5.7 External Validation: Does the Reweighting Track Real Distress?

The reallocation and predictive-fit results establish that the correction shifts measured exposure toward supplier regions and that this ordering predicts realized manufacturing decline better out of sample. A sharper, and more skeptical, question is whether the up-weighted regions exhibit *independently measured* supplier-linked distress beyond what direct exposure alone would predict. To separate the network channel from direct exposure, we regress commuting-zone distress outcomes on the network-only exposure $N_r = \tilde{B}_r - B_r$, controlling for direct exposure B_r and the standard controls, with period fixed effects and state-clustered standard errors:

$$\text{distress}_{r,t} \;=\; \alpha \;+\; \delta\, N_r \;+\; \gamma\, B_r \;+\; \mathbf{x}_r'\boldsymbol{\beta} \;+\; \tau_t \;+\; \varepsilon_{r,t}. \tag{35}$$

A coefficient δ that is significant in the distress direction means network-up-weighted CZs suffer more distress than their direct exposure predicts, which is evidence that the reweighting tracks real supply-chain incidence rather than mechanically smoothing the exposure measure.

Our analysis uses four outcomes, each built from data independent of the paper's Census-based manufacturing-employment outcome (Table 19). Two come from County Business Patterns establishment employment: (O1) the change in total CZ manufacturing employment, and (O2) the change in employment in upstream-supplier industries, the 159 of 342 BEA industries whose output is pulled mainly by other industries' shocks rather than their own (those with $[(\mathbf{L}-\mathbf{I})\mathbf{g}]_l > g_l$), which are not themselves directly hit by China. Two come from the Census Bureau's Business Dynamics Statistics, aggregated from county $\times$ sector to commuting zones via the standard crosswalk: (O3) the cumulative manufacturing establishment-exit rate and (O4) the cumulative manufacturing job-destruction rate over each period.

We find that the network channel loads significantly in the distress direction for three of the four outcomes. Network-only exposure predicts greater manufacturing employment decline beyond direct exposure ($\hat{\delta} = -0.0128$, $p = 0.001$), and the effect is largest, and most distinctly network-specific, for upstream-supplier employment ($\hat{\delta} = -0.0264$, $p = 0.000$): the supplier industries the correction relies on decline more precisely in the CZs it up-weights. Gross establishment exits from the Business Dynamics Statistics tell the same story ($\hat{\delta} = 0.0044$, $p = 0.043$). The cumulative job-destruction rate is the one outcome that does not corroborate the pattern: the point estimate is small and, if anything, of the non-distress sign ($\hat{\delta} = -0.0037$, $p = 0.221$). On balance the externally measured distress outcomes corroborate that the corrected measure's spatial reweighting reflects real supplier-linked incidence of the China Shock rather than an artifact of the exposure construction.

5.8 Policy Stakes: Trade Adjustment Assistance Targeting

The geographic reallocation has direct implications for place-based policy. We conduct a scenario analysis: if the true incidence of the shock is better captured by network-corrected exposure than by direct exposure, then a place-based program whose allocation is determined by standard Bartik rankings misallocates resources relative to the corrected ranking. Consider a planner allocating a fixed Trade Adjustment Assistance (TAA) budget equally across the K most-exposed commuting zones, ranked by the standard exposure measure. Under a need-weighted welfare function in which each CZ's need is proportional to its working-age population times its network-corrected exposure, the standard ranking is misaligned with the welfare-maximizing allocation, and the misalignment is largest precisely at tight targeting thresholds. While beyond the scope of our work, a promising avenue for future research is to bring in external welfare targets, such as realized TAA certifications or survey-based hardship indexes, which would upgrade this from scenario analysis to direct validation; we interpret the numbers below as conditional on the need-proportional-to-corrected-exposure assumption.

At $K = 10$, only 3 of the ten standard-measure zones are in the corrected top 10, and the standard allocation leaves 57.7% of need-weighted welfare on the table (37.7% under

population-only weighting). The efficiency loss moderates at broader thresholds but does not disappear: 16.1% at $K = 50$, 14.4% at $K = 100$, and 11.8% at the top 20% of commuting zones. The losses are invariant to the absolute budget and arise solely from the re-ranking, in that a program targeting by standard exposure misdirects assistance away from upstream-supplier CZs whose supply-chain exposure the standard formula does not capture. Table 11 reports the full set of targeting thresholds, and Figure 10 plots the efficiency loss against targeting depth. The cost of the misallocation is largest at tight targeting thresholds, where place-based programs tend to concentrate. Summed across communities, the standard and corrected estimators produce nearly identical total predicted impact; what differs is which communities are carrying that total, and the channel composition of the effect.

6 Validation and Robustness

This section reports validation and robustness checks: pre-trends and placebo outcomes that address the canonical "correction reloads onto already-declining places" concern (Section 6.1); leave-one-top-industry-out and leave-one-cluster-out jackknives that probe finite-shock sensitivity (Section 6.2); stability across IO-table vintages and aggregation levels (Section 6.3); exposure-robust inference (Section 6.4); cross-subsample heterogeneity (Section 6.5); and a sensitivity simulation that perturbs the entries of $\mathbf{A}$ to bound the consequences of using a single national IO matrix in place of unobserved region-specific networks (Section 6.7). A first-order-versus-all-orders robustness check that confirms the correction does not depend on exotic high-order network terms is reported as part of the three-channel decomposition appendix (Section B.6).

6.1 Pre-Trends

Because the network correction shifts weight toward supplier-heavy commuting zones (Section 5.5), a natural concern is that the correction loads onto places whose manufacturing employment was already declining before the China Shock. One might argue that the geographic-misattribution finding is driven by the CZs the correction upweights being on steeper pre-existing decline paths than the ones it downweights, rather than by genuine China-mediated propagation. To address this concern, we run the canonical falsification test of Autor et al. (2013) Table 8 on both the standard and corrected measures: regress pre-1990 CZ-level manufacturing employment share changes on the forward-looking $\mathbf{g}^{US,P2}$ Bartik, using 1990 shares for both measures, and check whether the corrected coefficient is differentially predictive of pre-trend decline.

We estimate each CZ's pre-period manufacturing employment share change on a fixed forward-exposure Bartik,

$$\Delta \mathcal{L}^{\text{mfg}}_{r,[t-10,t]} \;=\; \alpha_d \;+\; \beta\, E_r \;+\; \eta_{rt}, \qquad E_r \in \{B^{P2}_r,\, \tilde{B}^{P2}_r\},$$

separately for the two pre-periods $t \in \{1980, 1990\}$, where the exposure $B^{P2}_r = \mathbf{s}'_{r,1990}\mathbf{g}^{US,P2}$

is the standard P2 Bartik and $\tilde{B}_r^{P2} = \mathbf{s}'_{r,1990}\mathbf{L}\mathbf{g}^{US,P2}$ is the corrected P2 Bartik. Both exposures use the same $\mathbf{g}^{US,P2}$ (the 2000–2007 U.S. China import shock) and the same 1990 lagged shares; the difference between the two is only the network amplification. We instrument each with its other-country analog and cluster standard errors at the state level, weighting by start-of-period CZ population and absorbing Census division fixed effects.

Table 12 reports the estimates for both pre-periods and for the two in-sample periods 1990–2000 and 2000–2007. The pre-period evidence is mixed across decades and should not be read as a blank check. For the immediate pre-period 1980–1990 the standard 2SLS coefficient is 0.021 (s.e. 0.039, $p = 0.585$) and the corrected coefficient is -0.014 (s.e. 0.017, $p = 0.407$); both are statistically indistinguishable from zero under state-clustered standard errors. This is the canonical falsification window, one decade before the China shock arrives, and the result is clean under both exposure measures. In contrast, the earlier pre-period 1970–1980 yields significant positive coefficients on both measures: standard 0.235 (s.e. 0.058, $p = 0.000$) and corrected 0.105 (s.e. 0.031, $p = 0.001$). The 1970s coefficients are statistically significant at conventional levels but go in the opposite direction to the China-Shock prediction: future-China-exposed commuting zones were growing their manufacturing employment share during the 1970s rather than declining. The pattern is one of boom-then-bust: the regions that specialize in the industries China later competed with were absorbing manufacturing employment relative to the national trend in the 1970s, peaking around 1980, and contracting thereafter, so the sign of the 1970s placebo is structurally inconsistent with a "correction reloads onto already-declining places" story. The pre-trend concern that would invalidate the main result is a significant negative coefficient on either measure in 1980–1990 (the same direction as the main effect and the decade immediately adjacent to the sample), and neither measure produces one. In both pre-periods the corrected coefficient is smaller in absolute magnitude than the standard (in 1970–1980 less than half the standard, in 1980–1990 near zero for both measures), so the network correction is, if anything, less susceptible to pre-trend contamination than the standard Bartik. The upstream-supplier Midwest and Gulf Coast regions the correction upweights are not distinctively declining pre-1990; the correction's spatial reallocation reflects supply-chain propagation of the China Shock rather than selection onto pre-existing decline.

In the in-sample periods the expected pattern recovers: the 1990–2000 standard and corrected 2SLS coefficients are -0.162 (s.e. 0.038) and -0.071 (s.e. 0.022), and the 2000–2007 counterparts are -0.393 (s.e. 0.053) and -0.206 (s.e. 0.031). Both measures identify a significant negative effect in each main period, with the standard/corrected ratio close to the pooled 2.40 documented in Section 4. The transition from null adjacent-decade pre-trends (1980–1990) to negative main-period effects (1990–2000, 2000–2007) is synchronous with the onset of Chinese import competition, and is preceded by the sign-opposite manufacturing-share growth of the 1970s boom decade. $N = 722$ commuting zones per period. A broader placebo-outcome test with public-sector and agricultural employment share changes requires Census sector aggregates that are not part of the main analysis pipeline; we leave that extension to future work.

6.2 Finite-Shock Jackknife

The i.i.d. asymptotics of Proposition 1 are an idealized benchmark. In practice the realized ADH shock vector is concentrated at the NAICS-3 level with roughly 21 effective clusters; the consistency argument of Corollary 5 accommodates weak dependence but still requires a large effective number of primitive shocks. The substantive validity of the estimator in ADH therefore rests on finite-shock arguments—shock-level sensitivity analyses that probe whether the coefficient is stable to dropping any single industry or cluster—rather than on the asymptotic regime itself. The following subsection reports those diagnostics, and together with the exposure-robust inference of Section 6.4 constitutes the paper's primary inference argument; the i.i.d. asymptotics are retained as a conceptual benchmark against which the finite-shock behavior is compared. We probe this sensitivity with two jackknife exercises at the shock level for the single-coefficient main specification: leaving one BEA detail industry out at a time ($K = 342$ replicates) and leaving one three-digit NAICS cluster out at a time (21 replicates). For each replicate we zero the shock of the excluded industry (or cluster), rebuild the Bartik and its other-country IV analog with the same shares and Leontief inverse as the main specification, and rerun the pooled 2SLS.

The baseline single-regressor estimates on the $K = 342$ BEA-detail panel are -0.266 (s.e. 0.032) standard and -0.165 (s.e. 0.016) corrected; these are smaller in level than the headlines of Section 4 because the jackknife rebuilds the Bartik from decade-lagged employment shares on both sides (the fully predetermined variant), whereas the headline endogenous Bartik uses start-of-period shares.[10] Table 13 reports the jackknife distributions.

We find that under the leave-one-industry-out jackknife the standard coefficient moves by at most 0.0512 (jackknife SD 0.0746) and the corrected coefficient by at most 0.0589 (jackknife SD 0.0663). Both maximum deviations come from dropping a single electronic-computer industry (BEA code 334111): its removal moves the standard estimate by about 0.05 and the corrected by about 0.06, small relative to the baseline point estimates. Under the stricter leave-one-NAICS-3-cluster-out jackknife, the analogous maxima are 0.0978 (standard, SD 0.1001) and 0.1263 (corrected, SD 0.1213); the most influential cluster is NAICS 334 (computers and electronics, 22 BEA detail codes) for both specifications, which is the single cluster most exposed to the China Shock and therefore unsurprisingly load-bearing. Excluding the NAICS 334 cluster the remaining 20 clusters deliver standard and corrected coefficients within 0.02 of the respective baselines. The corrected coefficient's jackknife SD is comparable to the standard's in absolute terms: neither specification relies on a tail of exotic industry-level variation, and the network correction does not amplify the sensitivity to the top cluster beyond what the standard specification already exhibits. This result is consistent with the Neumann analysis of Section 4.4: most of the realized contamination factor is loaded onto first- and second-order propagation from a small number of heavily-shocked industries, and the NAICS 334 cluster sits at the top of that list.

[10] The headline and the jackknife panels use the same industry classification (BEA detail, $K = 342$). They differ in the share vintage of the endogenous regressor: the main-text headlines -0.612 and -0.255 instrument a start-of-period (1991/1999) endogenous Bartik with the decade-lagged (1980/1991) other-country Bartik, while the jackknife uses decade-lagged shares on both sides (the natural construction for direct industry-level shock perturbation). This share-vintage change, not any difference in industry granularity, accounts for essentially all of the level gap to the jackknife baselines $-0.266/-0.165$, which preserve the standard-exceeds-corrected pattern. Table 4 catalogues the constructions.

A placebo ϕ test complements the jackknife by probing whether the structural diagnostic itself is an artifact of generic matrix structure. Figure 11 reports the distribution of $\hat{\phi}_{\text{rand}} = \text{tr}(\mathbf{Q}\mathbf{L}_{\text{rand}})/\text{tr}(\mathbf{Q})$ under 1,000 random draws of $\mathbf{L}_{\text{rand}}$ from two shuffle schemes. Scheme A independently shuffles the off-diagonal entries of each row of $\mathbf{L}$, preserving row sums and row-level sparsity while scrambling the identity of the supplying columns; Scheme B further constrains each row-shuffle to remain within each of the 21 three-digit NAICS column blocks, a strictly tighter null. Under Scheme A no draw exceeds the empirical $\hat{\phi}$ ($p < 0.001$); under the stricter Scheme B the empirical value sits at the 97th percentile ($p = 0.034$). That result further rules out the possibility that the contamination multiplier merely reflects a generic matrix-smoothing artifact: it instead reflects specific element-wise alignment of $\mathbf{Q}$ and $\mathbf{L}$, the economic content of the co-location matrix emphasized in Section 2.1.

6.3 Aggregation and IO-Table Vintage

Table 14 reports the diagnostic $\hat{\phi}$ computed from every BEA detailed benchmark from 1997 to 2017. Manufacturing $\hat{\phi}^{\text{mfg}}$ shows a modest upward trend from 1.100 in 1997 to 1.138 in 2017, consistent with increasing supply-chain integration; full-economy $\hat{\phi}^{\text{full}}$ is stable at 1.10–1.15 throughout. The diagnostic $\hat{\phi}$ is highly stable across vintages, but the corrected 2SLS point coefficient is more vintage-sensitive: it is -0.255 using the 1997 IO table and -0.374 using the 2002 table (Table 14), a movement of roughly 47%. Both remain of the same sign and order of magnitude and far from the standard estimate; the robustness we document is of the structural diagnostic $\hat{\phi}$ and the qualitative direction of the correction, not point-coefficient invariance across IO vintages. A sharper concern is that the $+0.038$ drift in $\hat{\phi}^{\text{mfg}}$ might be concentrated in industry clusters that co-offshored with China during the 2000s, which would weaken the 1997 pre-determination argument. A cluster decomposition of the drift (Figure 12 and Table 15) helps to rule out this concern: 85% of the total drift comes from three clusters (Food, Transp. equip., Electrical equip.), none of which are China-shocked sectors, and a cluster-level regression of per-cluster drift on China-shock exposure returns a coefficient that is negative and economically negligible ($-1.85e-04$, $p = 0.011$); with only 17 cluster observations we put the weight on the magnitude rather than the p-value.

Industry aggregation also matters. At the summary level ($K = 53$), the diagnostic contamination multiplier is substantially larger ($\hat{\phi} = 1.26$) than at the detail level ($\hat{\phi} = 1.11$), because aggregation inflates the apparent IO linkages: within-summary-industry flows between detailed industries are absorbed into the diagonal of the aggregated matrix, mechanically increasing the off-diagonal linkages that drive the bias formula. Comparing coefficient estimates across aggregation levels, we observe that moving from the summary ($K = 53$) to the detail ($K = 342$) system, the standard 2SLS coefficient moves from -0.383 to -0.266 and the corrected from -0.287 to -0.165 (Table 16, which compares like-for-like BEA constructions). The qualitative pattern is preserved at both levels (the standard coefficient exceeds the corrected in magnitude and both remain negative and of the same order), but the point levels are not invariant to aggregation. These lagged-share BEA coefficients also differ from the main-text headlines (-0.612 standard, -0.255 corrected), which build the endogenous Bartik from start-of-period shares; Table 4 catalogues the

constructions. The detail-level results are preferred throughout the main analysis because they more accurately represent the true IO structure.

6.4 Exposure-Robust Inference

For the single-instrument corrected estimator, Proposition 4 establishes that the Adão et al. (2019) exposure-robust variance must be computed with the network-adjusted shares $\tilde{s}_{rl} = \sum_k s_{rk} L_{kl}$; using the original shares s_{rk} produces an inconsistent estimate. Table 17 reports the comparison. We find that the AKM-naive standard error (computed with original shares) is about 36% smaller than the AKM-correct standard error (computed with network-adjusted shares): a practical warning that using the original shares in the variance formula produces confidence intervals that are too narrow. The correctly computed exposure-robust standard errors come out modestly larger than the state-clustered ones used in the main tables (0.0263 versus 0.0228 for the corrected estimator, 0.0845 versus 0.0659 for the standard), and every main-specification estimate remains highly significant under either method, with t-statistics between -7.2 and -11.2. Two implementation caveats attach to the AKM variant: the regional controls have no shift-share structure, so the shocks in the score can be residualized only on the period constants, and the shock-level influence is concentrated in relatively few three-digit NAICS cells. We therefore keep the state-clustered intervals as the headline convention (matching ADH) and complement them with the few-cluster-robust NAICS-3 wild bootstrap reported below. Because the single-instrument corrected and standard 2SLS specifications are just-identified, we also report Anderson–Rubin (Anderson and Rubin, 1949) weak-instrument-robust 95% confidence sets (which coincide with the conditional-likelihood-ratio set in the just-identified case): $[-0.301, -0.213]$ for the corrected estimator and $[-0.763, -0.498]$ for the standard. Both are bounded, exclude zero, and lie close to the corresponding Wald intervals (as expected given the strong first stages, $F = 192.9$ corrected and 61.7 standard), so weak-instrument-robust inference does not overturn either estimate. Finally, to address the finite-shock concern directly (the realized shock vector has roughly 21 effective three-digit NAICS clusters), we run a Webb (2023) six-point cluster wild bootstrap ($9,999$ replications) that clusters the exposure-robust score at the NAICS-3 shock level (complementing the leave-one-cluster-out jackknife of Section 6.2). The corrected coefficient's 95% bootstrap confidence interval is $[-0.256, -0.253]$ and the standard's is $[-0.617, -0.607]$, both with bootstrap $p < 0.0001$: significance survives few-effective-cluster inference. Consistent with the pattern above, these shock-side intervals are tighter than the state-clustered intervals reported in the main tables, which remain our preferred convention.

6.5 Heterogeneity by IO Density

A correctly specified estimator should produce coefficients that vary across subsamples in theoretically interpretable ways rather than arbitrary ones. We next consider heterogeneity according to local IO density. Splitting the ADH sample at the median of CZ-level employment share in high-IO-connectivity industries, the corrected Bartik coefficient moves from -0.270 in the high-IO half to -0.199 in the low-IO half (small deviations from the pooled

-0.255), while the standard Bartik moves from -0.338 to -0.596. Extending to terciles, Figure 13 and Table 18 show that the corrected estimator's coefficients are monotone in IO density (-0.197, -0.250, -0.298, moving in a consistent direction from low to high tercile), consistent with the network channel mattering more in supply-chain-dense regions, while the standard Bartik's coefficients are non-monotone and move in the opposite direction (-0.645, -0.537, -0.293). The tercile Wald test rejects coefficient equality at the 1% level for both estimators ($p = 0.002$ corrected, $p = 0.003$ standard), so the subsample differences are statistically detectable in this well-powered sample; what distinguishes the corrected from the standard is not that either is strictly stable across subsamples but that the corrected's heterogeneity is monotone and theoretically interpretable (increasing magnitude with IO density), while the standard's is not. Combined with the failure of structural alignment in the high-IO tercile documented in Section 5.3, the heterogeneity analysis provides an economically interpretable picture of how the coefficient contamination varies with local supply-chain density.

6.6 Contamination Bias from Heterogeneous Direct Effects

Goldsmith-Pinkham et al. (2024) show that a regression with several treatments can attribute to one treatment the heterogeneous effects of another. The concern here is that the network coefficient $\hat{\beta}_N$ could absorb cross-industry heterogeneity in the direct effect $\beta_{D,k}$ rather than measure genuine propagation. We provide two pieces of evidence that it does not. First, this form of contamination operates through negative Rotemberg weights, and in this design the negative-weight mass is small: it sums to 0.037 for the standard Bartik and 0.012 for the corrected one (Table 7), so each estimator is close to a convex average of industry-level effects and the scope for contamination is limited. Second, we test the channel directly by saturating the direct channel. We replace the single direct exposure B_r with sector-specific direct Bartiks, one per three-digit NAICS manufacturing sector (the level at which the Chinese import shocks cluster), each instrumented by its other-country analog, while keeping the single network-only regressor N_r. This lets the direct per-unit effect differ freely across the 21 sectors, so any direct-effect heterogeneity that N_r might otherwise proxy is absorbed by the saturated direct channel. The network coefficient is essentially unchanged: $\hat{\beta}_N = -0.328$ (s.e. 0.072) under the 21-sector saturation and -0.306 (s.e. 0.057) under a coarser two-digit (3-sector) saturation, against -0.233 (s.e. 0.050) in the baseline two-regressor specification, all significant at the 1% level. If anything the network coefficient grows slightly in magnitude once direct heterogeneity is absorbed, the opposite of what upward contamination would produce, so the network channel is not a re-badging of heterogeneous direct effects.

6.7 Sensitivity to Regional IO Heterogeneity

The corrected estimator uses a single national IO matrix $\mathbf{A}$ (BEA 1997 detail) for every commuting zone. If true regional input–output coefficients $\mathbf{A}_r$ deviate from the national matrix—because, say, a Detroit auto plant draws on a different upstream input mix than a Texas auto plant—then $\tilde{B}_r = \mathbf{s}_r'(\mathbf{I} - \mathbf{A})^{-1}\mathbf{g}$ measures the network-amplified exposure

with error, and the corrected 2SLS coefficient is biased toward the standard's. Gulek (2026) documents substantial province-level heterogeneity in IO structure for Turkey; in the U.S., commercial benchmarks (BEA's RIMS II, IMPLAN) provide regional multipliers but are not freely incorporable into a single-pipeline replication, and the EPA's open-source state input–output tables (Li et al., 2024) provide bi-regional Make and Use tables for all fifty states plus the District of Columbia at the BEA Summary level. We provide three pieces of evidence in favor of insensitivity to this concern: an exact bias decomposition, a sensitivity simulation (Path A below), and a direct state-level robustness check (Path B).

Exact bias decomposition. The consequence of using a national matrix for the IV *slope* is governed exactly: the corrected coefficient is biased only to the extent that the induced exposure error $\widehat{X}_r - X_r^*$ (the gap between national-IO and true regional-IO exposure) covaries with the instrument, $\hat{\beta}/\beta - 1 = -\mathrm{Cov}(\widehat{Z}, \widehat{X} - X^*)/\mathrm{Cov}(\widehat{Z}, \widehat{X})$ (Proposition 10). This replaces a classical errors-in-variables intuition—that mean-zero regional deviations of $\mathbf{A}_r$ would attenuate the corrected coefficient toward the standard and make the 2.40× ratio a lower bound—which is not a theorem here: $\tilde{B}_r$ and its instrument are nonlinear in $\mathbf{A}$ through $(\mathbf{I} - \mathbf{A})^{-1}$, so the attenuation direction is not guaranteed and we claim no formal lower bound. The decomposition is the more useful object because it is both signed by an estimable covariance and bounded: by Cauchy–Schwarz, $|\hat{\beta}/\beta - 1| \leq [\mathrm{sd}(\widehat{X} - X^*)/\,\mathrm{sd}(\widehat{X})]\,/\,|\mathrm{Corr}(\widehat{Z}, \widehat{X})|$ (Corollary 7), so a strong first stage and a small induced exposure error cap the slope bias whatever the sign of the matrix error. Paths A and B evaluate both the magnitude and the *a priori* ambiguous direction directly.

Path A: A-perturbation sensitivity simulation. As a robustness check, we simulate a counterfactual scenario in which every nonzero entry of $\mathbf{A}$ is perturbed multiplicatively, $A_{kl}^{\text{noisy}} = A_{kl}\,(1 + \varepsilon_{kl})$ with $\varepsilon_{kl} \sim \mathcal{N}(0, \sigma^2)$ truncated below at zero, recompute the noisy Leontief inverse $\mathbf{L}^{\text{noisy}} = (\mathbf{I} - \mathbf{A}^{\text{noisy}})^{-1}$, and rerun the full 2SLS with the noisy $\tilde{B}_r$ on both sides of the IV. We do this for 200 draws at each of $\sigma \in \{0.05, 0.10, 0.20, 0.50\}$. The $\sigma = 0.05$ benchmark corresponds to 5% noise on each input coefficient—roughly the level of plausible regional-vintage drift documented in Section 6.3—and $\sigma = 0.50$ is an upper bound far exceeding any plausible regional dispersion. Figure 14 reports the resulting distribution of the corrected coefficient. At $\sigma = 0.05$ the simulated coefficients have mean -0.262 and standard deviation 0.002, indistinguishable from the noiseless baseline of -0.262; at $\sigma = 0.10$ the mean moves to -0.262 (sd 0.004); at $\sigma = 0.20$, -0.261 (sd 0.007); and at the implausibly large $\sigma = 0.50$, -0.258 (sd 0.021). Across all four perturbation levels the simulated mean stays within -1.68% of the noiseless baseline, and even the $\sigma = 0.50$ band excludes the standard 2SLS coefficient -0.612 by a wide margin.

Path B: state-specific Leontief substitution. These tables provide bi-regional Use tables, grounded in Freight Analysis Framework trade flows, that split each state's economy into the state itself and the rest of the United States at the BEA Summary level, for 51 states. Li et al. (2024) note that the tables impose the national input mix on every state's production technology, with state heterogeneity entering only through inter-state commodity flows;

the table accordingly separates flows within the state, from the state to the rest of the country, from the rest of the country into the state, and within the rest of the country, and these blocks differ across states because states differ in how much they source in-state versus out-of-state. We use the within-state block to construct, for each state s, an intra-state direct-requirements matrix $\mathbf{A}_s^{\text{intra}}$ (within-state intermediate purchases divided by state gross output) and the corresponding $\mathbf{L}_s^{\text{intra}} = (\mathbf{I} - \mathbf{A}_s^{\text{intra}})^{-1}$, restricted to the 19 BEA Summary manufacturing codes after aggregating shares and shocks from BEA Detail to BEA Summary. We then assign each commuting zone the $\mathbf{L}_s^{\text{intra}}$ of its state and rerun the full 2SLS with these state-specific Leontiefs on both sides of the IV.

The state-restriction is intentionally severe: $\mathbf{L}_s^{\text{intra}}$ ignores all cross-state propagation, so the average column sum of state-specific Leontiefs is 1.087 versus 1.353 for the national-summary baseline—propagation through the supply chain is roughly 20% smaller when restricted to in-state linkages. Despite this, the corrected coefficient barely moves: the summary-level baseline (using a national $\mathbf{L}$ at BEA Summary, with detail shares and shocks aggregated to summary) is -0.690 (s.e. 0.075), and the state-intra version is -0.712 (s.e. 0.091)—a shift of only 3.05% relative to the summary baseline.[11] Applying the exact decomposition of Proposition 10 to this benchmark makes the bias precise. Treating the state-specific intra-state Leontief as the regional "truth," the induced exposure error $\widehat{d}_r = \mathbf{s}_r'(\mathbf{L}^{\text{intra}}_{\text{state}(r)} - \mathbf{L}^{\text{nat}})\mathbf{g}$ has a standard deviation equal to 15.34% of the exposure's own variation, and with a first-stage correlation of 0.867 the estimable proportional slope bias is $\hat{\kappa} = -10.36\%$ (residualized and population-weighted to match the 2SLS; -0.155 unweighted), against a Cauchy–Schwarz ceiling of 17.69% (Corollary 7). The sign is the over-correction direction: the in-state restriction strips cross-state linkages, so the national matrix over-propagates exposure in high-instrument regions and the national-IO coefficient is, if anything, slightly too small in magnitude. This is the opposite of the conservative "lower bound" that the naive attenuation intuition would assert, and it is exactly the ambiguity that Proposition 10 formalizes. The full re-estimation with state-specific Leontiefs on both sides moves the coefficient in the same (more-negative, closer-to-standard) direction by only 3.05%. Taken together, both the realized bias and its ceiling are small under this deliberately severe in-state benchmark.

Two extensions remain open: a further direct test using a CFS-based interregional reweighting (the natural Path C), and a re-estimation in which the state tables' assumption of homogeneous production technology is itself relaxed using state-by-industry value-added shares. We leave both, since Paths A and B together already show that the corrected coefficient is insensitive to plausible departures from the single-national-IO assumption.

[11] The summary-level baseline -0.690 differs in level from the detail-level headline because aggregating to BEA Summary collapses much within-cluster network propagation into the diagonal of $\mathbf{A}$ (Section 6.3); the Path B comparison is internal to the summary level, not an alternative headline. Vintage gap: these tables ship the 2017 benchmark; ADH spans 1991–2007. We treat them as the closest publicly available state-level IO benchmark and accept the roughly 10-to-26-year vintage gap as a robustness-check caveat, consistent with the small drift documented at the vintage level in Section 6.3.

7 Portability: The Blanchard–Katz Employment-Growth Bartik

This section applies the two-channel framework to a second canonical shift-share design: the Bartik (1991) employment-growth Bartik in the Blanchard and Katz (1992) form used by Goldsmith-Pinkham et al. (2020), on the same share matrices, Leontief inverse, and commuting zones as the main ADH analysis. The section does two things: it shows portability (the framework applies to any shift-share design with a well-defined propagation matrix), and it distinguishes two mechanisms that can produce aggregate cancellation.

7.1 The BK Setup

The BK shift-share (Blanchard–Katz form, following the Goldsmith-Pinkham-Sorkin-Swift formulation) uses national industry employment growth rates as shocks, $g_k^{BK} = \Delta \ln \mathcal{L}_{k,t}^{\text{nat}}$, and regional industry employment shares as weights. The reduced-form regression is

$$\Delta \mathcal{L}_{rt}^{\text{mfg}} \;=\; \alpha_t + \beta^{BK} \cdot B_r^{BK} + \varepsilon_{rt}, \qquad B_r^{BK} = \mathbf{s}_r' \mathbf{g}^{BK},$$

with standard OLS (there is no equivalent of the ADH other-country instrument in this setting). Applying the two-channel correction, the network-corrected BK Bartik is $\tilde{B}_r^{BK} = \mathbf{s}_r' \mathbf{L} \mathbf{g}^{BK}$ and the analogous corrected specification estimates $\hat{\beta}^{BK,\text{corr}}$. The standard and corrected BK estimates are reported in Table 20. The coefficient ratio of 2.25 is comparable to the ADH ratio of 2.40, despite a very different shock vector. The two-channel framework of Section 2 therefore produces a similar pattern of standard-vs-corrected divergence in both applications. The picture for the underlying economic mechanism, however, is more complicated: it differs in one important respect.

7.2 Two Distinct Mechanisms for Aggregate Cancellation

Our empirical work above shows the form of the link between aggregate cancellation and the covariance identity of Proposition 6 for ADH: the channel ratio $\theta = \beta_N / \beta_D$ is non-negligible, but the gap $\phi - \psi^{\text{agg}}$ is small because $\text{Cov}_{\text{pop}}(m_r, B_r)$ is meaningfully negative. Call this the structural cancellation mechanism.

In contrast to ADH, a different mechanism operates in BK. The coefficient estimates from the BK two- and three-regressor OLS decompositions are reported in Table 21. At the three-regressor level,

$$\hat{\beta}_D^{BK} = 1.110, \qquad \hat{\beta}_L^{BK} = 0.004, \qquad \hat{\beta}_G^{BK} = 0.035;$$

at the two-regressor level the network-channel coefficient is $\hat{\beta}_N^{BK} = 0.038$ (s.e. 0.313), essentially zero, against a direct coefficient of $\hat{\beta}_D^{BK} = 1.133$ (s.e. 0.639). The channel ratio $\hat{\theta}^{BK} = 0.03$ is therefore near zero. When $\theta \approx 0$, the two-channel aggregate ratio (24) simplifies to $1 + 0 \cdot (\phi - 1) \approx 1$ regardless of ϕ: aggregate cancellation is mechanical

(a small network-channel causal coefficient) rather than structural (a covariance-driven alignment between ϕ and ψ^{agg}).

This case carries a sharp practical lesson. The standard-versus-corrected comparison alone is misleading, because it ignores whether the network channel is active. When $\beta_N \approx 0$ there is no omitted network channel, so the standard Bartik is itself consistent ($\text{plim}\hat{\gamma}_{SS} = \beta_D$ by Proposition 1), while the single-regressor corrected estimator is attenuated toward zero (Proposition 3: its probability limit is $\lambda\beta_D$ with $\lambda \leq 1/\phi < 1$). The BK standard/corrected ratio of 2.25 therefore reflects the corrected estimator being biased toward zero rather than the standard being contaminated, the opposite of the ADH reading. The general lesson is that a large standard-versus-corrected coefficient ratio is not, by itself, evidence of network contamination: it is diagnostic of contamination only when the network channel is active ($\beta_N \neq 0$), which is exactly what the two-regressor split tests. The implementation advice is direct: estimate the two-regressor (or three-regressor) decomposition first; adopt the single-regressor correction only when β_N is detectably nonzero; and where it is not, keep the standard Bartik, which is the consistent estimator there—the correction would be the biased choice.

7.3 Cross-Application Summary

The diagnostic across applications is summarized in Table 22. In both ADH and BK the structural diagnostic $\hat{\phi}^{\text{iid}} \approx 1.11$ is identical (they share the same $(\mathbf{Q}, \mathbf{L})$), but the realized coefficient ratio differs slightly and the mechanism of aggregate near-cancellation differs qualitatively. The distinction matters for practice. To assess a new shift-share design, a researcher should test both the structural cancellation condition ($\phi \approx \psi^{\text{agg}}$, computable from the share matrix and propagation matrix alone) and the mechanical cancellation condition ($\beta_N \approx 0$, requiring the two-regressor decomposition to be estimable). Aggregate agreement driven by a small network channel in a given outcome does not imply that aggregate agreement will hold for a different outcome or in a different application; aggregate agreement driven by structural alignment of co-location and population weights is a property of the design that transfers across outcomes.

One caution attaches to this reading: the mechanical-cancellation interpretation for BK rests on the network-channel estimates being empirically close to zero in this particular sample and outcome, and the BK point estimates are imprecise ($\hat{\beta}_D^{BK} = 1.133$ with s.e. 0.639 is not significant at the 5% level), so the mechanism call is made on noisy coefficients. In reality, a BK-style application whose three-regressor specification revealed a detectably nonzero network-channel coefficient would render the two-regressor form misspecified, and the researcher would then need to use the three-regressor decomposition or the single-regressor corrected estimator (interpreted via Proposition 3 as a projection of the two causal effects) rather than relying on a zero-β_N simplification. The two-regressor-vs-three-regressor choice in BK is therefore an empirical finding about this outcome, not a modeling commitment that would carry over to other BK-style panels.

7.4 Using the Framework in a New Application

The workflow the paper implies can be collected as a checklist.

1. Compute the ex-ante diagnostic $\hat{\phi}^{\text{iid}} = \text{tr}(\mathbf{QP})/\text{tr}(\mathbf{Q})$ from the share and propagation matrices, before running any regression (Proposition 1); with controls and weights, use the residualized $\tilde{\phi}$ of Corollary 4.

2. Where shift-exogenous instruments exist, estimate the two-regressor decomposition (B_r, N_r) and test $\beta_N = 0$ (Section 4.2).

3. If the network channel is live, apply the correction: report the two-regressor split, or the single-regressor corrected estimator read through the projection weight λ (Proposition 3). If $\beta_N \approx 0$, keep the standard Bartik—it is the consistent estimator there.

4. Check aggregate safety ex ante: cancellation holds exactly when $\text{Cov}_p(m_r, Z_r) = 0$ under the exposure measure, and its failure is bounded by the dispersion of the regional multiplier m_r (Proposition 7).

5. Inspect the geography regardless of the aggregate: rank changes, top-K overlap, and the cross-region variance of Δ_r^{std} (Section 5).

6. Compute exposure-robust standard errors on the network-adjusted shares $\tilde{s}_{rl}$, with the residualized shock in the score (Proposition 4).

8 Conclusion

This paper shows how a production network may affect the shift-share design. When the design is applied to industries connected through such a network, the standard exposure map **Sg** is misspecified relative to the network-propagated map **SPg**, and the standard Bartik coefficient is a structural mixture of direct and network-propagated causal effects: a misspecification of the exposure map rather than a failure of shock or share exogeneity. The mixture weight is a closed-form object, $\phi - 1 = \text{tr}(\mathbf{Q}(\mathbf{P} - \mathbf{I}))/\text{tr}(\mathbf{Q})$, that depends only on the share matrix and the propagation matrix. A single-regressor correction that pre-multiplies the shock vector by the propagation matrix is consistent under the same shift-exogeneity regime as the standard Bartik (given a predetermined, correctly measured propagation matrix and the network-adjusted concentration conditions of Proposition 2), and a two-regressor form separates the direct and network-only channels under BHJ-style shift exogeneity. Both corrections admit clean exposure-robust inference through a network-adjusted-shares variant of the Adão et al. (2019) formula. A covariance identity yields, in turn, an exact, falsifiable condition (the paper's central theoretical result) under which the standard and corrected estimators deliver identical aggregate predicted impacts even when their coefficients differ sharply.

We apply the framework to the Autor et al. (2013) China Shock. Here we find that the standard Bartik 2SLS coefficient of -0.612 is a factor 2.40 larger in magnitude than the corrected coefficient of -0.255. Translated into predicted impacts at the interquartile range of exposure, the two estimators nearly coincide (-1.063 vs -1.196 percentage points, a gap of 11.2%). The coincidence is conditional rather than automatic: by the covariance identity

above, aggregate predictions coincide only to the extent that the covariance between CZ-level effective multipliers and direct exposure is small. In ADH it is meaningfully negative: upstream-supplier regions (petroleum, primary metals, basic chemicals, paper) have high multipliers but low direct exposure, so the pooled design nearly satisfies the condition while its supply-chain-dense subregions violate it. Coefficient inflation and exposure understatement both arise from ignoring the propagation matrix and move through the same network geometry; whether they cancel in aggregate is the testable question the condition answers.

The correction's economic content lies instead in the spatial distribution of measured impact. The mean-squared prediction error between the standard and corrected CZ-level impacts is 99.8% cross-CZ reallocation; 34% of the top-100 standard-measure zones are not in the corrected top 100, and 68% of the top-25 are not in the corrected top 25. The corrected measure's ordering of CZs correlates more strongly with realized manufacturing employment changes, producing an out-of-sample R^2 more than three times the standard's, a gain robust to leave-one-division-out and leave-one-state-out cross-validation. While the two estimators sum across communities to a similar total, the corrected measure is better at identifying which communities carry that total. A place-based policy built on the standard Bartik therefore misdirects resources away from supply-chain-exposed commuting zones.

The framework extends naturally to any shift-share design whose units are connected through an observable network. Section 7 ports it to the Bartik (1991) employment-growth Bartik; an online appendix applies it, in turn, to modern-era Card (2001) immigration with a native-migration matrix, the COVID labor-demand shock with the Bailey et al. (2018) Social Connectedness Index, and the Acemoglu and Restrepo (2020) robotics shock with the LODES commuting matrix. Each application's structural diagnostic, aggregate amplification, and geographic reallocation pattern can be read off the same $(\mathbf{Q}, \mathbf{P}, \mathbf{w})$ apparatus developed here, though the magnitude of the geographic-incidence payoff is not uniform across settings: it is most pronounced in the China Shock, where direct exposure and network multipliers are strongly misaligned, and in each extension depends on how far the propagation matrix departs from proportionality to direct exposure. While our focus has been on the national-network case, parallel questions about spatially disaggregated linkages remain unanswered; the natural next refinement is to relax the national-network assumption using such data on the relevant linkages.

References

ACEMOGLU, D., D. AUTOR, D. DORN, G. H. HANSON, AND B. PRICE (2016): "Import Competition and the Great US Employment Sag of the 2000s," *Journal of Labor Economics*, 34, S141–S198.

ACEMOGLU, D. AND P. D. AZAR (2020): "Endogenous Production Networks," *Econometrica*, 88, 33–82.

ACEMOGLU, D. AND P. RESTREPO (2020): "Robots and Jobs: Evidence from US Labor Markets," *Journal of Political Economy*, 128, 2188–2244.

ADÃO, R., C. ARKOLAKIS, AND F. ESPOSITO (2025): "General Equilibrium Effects in Space: Theory and Measurement," *American Economic Journal: Macroeconomics*, conditionally accepted. Previously NBER Working Paper 25544.

ADÃO, R., M. KOLESÁR, AND E. MORALES (2019): "Shift-Share Designs: Theory and Inference," *The Quarterly Journal of Economics*, 134, 1949–2010.

ANDERSON, T. W. AND H. RUBIN (1949): "Estimation of the Parameters of a Single Equation in a Complete System of Stochastic Equations," *Annals of Mathematical Statistics*, 20, 46–63.

ANDREWS, D. W. K. (1988): "Laws of Large Numbers for Dependent Non-Identically Distributed Random Variables," *Econometric Theory*, 4, 458–467.

AUTOR, D., D. DORN, G. HANSON, M. R. JONES, AND B. SETZLER (2025): "Places versus people: the ins and outs of labor market adjustment to globalization," in *Handbook of Labor Economics*, Elsevier, vol. 6, chap. 7, 549–657.

AUTOR, D. H., D. DORN, AND G. H. HANSON (2013): "The China Syndrome: Local Labor Market Effects of Import Competition in the United States," *American Economic Review*, 103, 2121–2168.

BAILEY, M., R. CAO, T. KUCHLER, AND J. STROEBEL (2018): "The Economic Effects of Social Networks: Evidence from the Housing Market," *Journal of Political Economy*, 126, 2224–2276.

BAQAEE, D. R. AND E. FARHI (2019): "The Macroeconomic Impact of Microeconomic Shocks: Beyond Hulten's Theorem," *Econometrica*, 87, 1155–1203.

BAQAEE, D. R. AND E. RUBBO (2023): "Micro Propagation and Macro Aggregation," *Annual Review of Economics*, 15, 647–676.

BARTIK, T. J. (1991): *Who Benefits from State and Local Economic Development Policies?*, Kalamazoo, MI: W.E. Upjohn Institute for Employment Research.

BLANCHARD, O. J. AND L. F. KATZ (1992): "Regional Evolutions," *Brookings Papers on Economic Activity*, 1992, 1–75.

BORUSYAK, K. AND P. HULL (2023): "Nonrandom Exposure to Exogenous Shocks," *Econometrica*, 91, 2155–2185.

——— (2025): "Optimal Formula Instruments," Working Paper 33594, National Bureau of Economic Research.

BORUSYAK, K., P. HULL, AND X. JARAVEL (2022): "Quasi-Experimental Shift-Share Research Designs," *The Review of Economic Studies*, 89, 181–213.

BUREAU OF ECONOMIC ANALYSIS (2002): "Benchmark Input-Output Accounts of the United States, 1997," Tech. rep., U.S. Department of Commerce, detailed Make and Use Tables. Available at `https://www.bea.gov/industry/input-output-accounts-data`.

CALIENDO, L., M. DVORKIN, AND F. PARRO (2019): "Trade and Labor Market Dynamics: General Equilibrium Analysis of the China Trade Shock," *Econometrica*, 87, 741–835.

CARD, D. (2001): "Immigrant Inflows, Native Outflows, and the Local Labor Market Impacts of Higher Immigration," *Journal of Labor Economics*, 19, 22–64.

DORN, D. (2009): "Essays on Inequality, Spatial Interaction, and the Demand for Skills," Dissertation, University of St. Gallen. Industry crosswalks available at `https://www.ddorn.net/data.htm`.

DUSTMANN, C. AND U. SCHÖNBERG (2025): "Linking empirical evidence to theory: a framework for understanding immigration's labor market effects," in *Handbook of Labor Economics*, Elsevier, vol. 6, chap. 6, 447–548.

ECKERT, F., T. C. FORT, P. K. SCHOTT, AND N. J. YANG (2024): "Imputing Missing Values in the US Census Bureau's County Business Patterns," *NBER Working Paper*.

FEENSTRA, R. C., R. E. LIPSEY, H. DENG, A. C. MA, AND H. MO (2005): "World Trade Flows: 1962-2000," NBER Working Paper 11040, National Bureau of Economic Research, documents the construction of cleaned bilateral trade flows including reassignment of Hong Kong re-exports of Chinese-origin goods back to China.

GALLE, S., A. RODRÍGUEZ-CLARE, AND M. YI (2023): "Slicing the Pie: Quantifying the Aggregate and Distributional Effects of Trade," *The Review of Economic Studies*, 90, 215–251.

GOLDSMITH-PINKHAM, P., P. HULL, AND M. KOLESÁR (2024): "Contamination Bias in Linear Regressions," *American Economic Review*, 114, 4015–4051.

GOLDSMITH-PINKHAM, P., I. SORKIN, AND H. SWIFT (2020): "Bartik Instruments: What, When, Why, and How," *American Economic Review*, 110, 2586–2624.

GULEK, A. (2026): "Effects of Immigrants on Non-host Regions: Evidence from the Syrian Refugees in Turkey," Job Market Paper. Available at `https://ahmetgulek.github.io/Gulek_jmp.pdf`.

LI, M., W. W. INGWERSEN, B. YOUNG, J. VENDRIES, C. BIRNEY, AND S. CASHMAN (2024): "Methods and data needed to construct a state economic input–output model that ties to national accounts: An illustration using U.S. states," *International Regional Science Review*, 47, 325–360, r package `stateior` version 0.4.0 available at `https://github.com/USEPA/stateior`.

PIERCE, J. R., P. K. SCHOTT, AND C. TELLO-TRILLO (2024): "To Find Relative Earnings Gains After the China Shock, Look Outside Manufacturing and Upstream," Working Paper 32438, National Bureau of Economic Research.

POLITIS, D. N. AND J. P. ROMANO (1994): "Large Sample Confidence Regions Based on Subsamples under Minimal Assumptions," *Annals of Statistics*, 22, 2031–2050.

SANDERSON, E. AND F. WINDMEIJER (2016): "A Weak Instrument *F*-test in Linear IV Models with Multiple Endogenous Variables," *Journal of Econometrics*, 190, 212–221.

WEBB, M. D. (2023): "Reworking Wild Bootstrap-Based Inference for Clustered Errors," *Canadian Journal of Economics*, 56, 839–858.

Table 1. Network-OVB Diagnostic Objects (ADH)

Object	Formula	Diagnoses	ADH value
ϕ^{iid}	$\text{tr}(\mathbf{QP})/\text{tr}(\mathbf{Q})$	Ex-ante co-location / projection risk	1.110
ϕ^{IV}	$(\mathbf{g}^{OC})'\mathbf{QP}\mathbf{g}^{US}/(\mathbf{g}^{OC})'\mathbf{Q}\mathbf{g}^{US}$	Realized slope contamination	2.316
ψ^{agg}	$\sum_r w_r \tilde{B}_r / \sum_r w_r B_r$	Aggregate exposure amplification	2.766
m_r dispersion	IQR of $\tilde{B}_r / B_r$ across regions	Geographic incidence risk	1.267
$\text{MSE}(\Delta_r^{\text{std}})$	$\mathbb{E}_r\left[\beta_N^2 B_r^2 (\phi - m_r)^2\right]$	Geographic share of standard's MSE	99.8% geo.

Notes: Five diagnostic objects produced by the framework of Section 2 and their ADH values. The first three rows are scalar amplification factors; the fourth summarizes cross-region dispersion of the regional multiplier $m_r = \tilde{B}_r / B_r$; the fifth reports the share of the standard estimator's MSE attributable to the cross-region variance of Δ_r^{std} rather than to its population-weighted mean. ADH values are reported in Sections 4.4 (slope and aggregate amplification) and 5 (region-level dispersion and MSE share). The structural ϕ^{iid} here uses the unweighted co-location matrix $\hat{\mathbf{Q}} = \mathbf{S}'\mathbf{S}$; the population-weighted convention $\mathbf{Q}^w = \mathbf{S}' \text{diag}(\mathbf{w})\mathbf{S}$ used in the Neumann decomposition of Table 6 yields $\phi^{\text{iid}} = 1.164$, the small difference reflecting population weighting of the co-location structure rather than a distinct estimand.

Table 2. Amplification-Factor Scorecard (ADH)

Factor	Variant	Weighting / share vintage	ADH value
Structural (i.i.d., design-only): $\phi^{\text{iid}} = \text{tr}(\mathbf{QP})/\text{tr}(\mathbf{Q})$			
ϕ^{iid}	unweighted $\mathbf{Q} = \mathbf{S}'\mathbf{S}$	unweighted, decade-lagged 1991 shares	1.110
ϕ^{iid}_{w}	weighted $\mathbf{Q}_w = \mathbf{S}'\mathbf{W}\mathbf{S}$	population-weighted, decade-lagged	1.164
$\phi^{\text{iid}}_{\text{pool}}$	pooled across periods	unweighted, decade-lagged	1.115
$\phi^{\text{iid}}_{\text{cv}}$	cross-vintage co-location	weighted, start-of-period $\mathbf{S}_X$ / decade-lagged $\mathbf{S}_Z$	1.180
Realized under OLS: $\phi^{\text{corr}} = (\mathbf{g}^{US})'\mathbf{QP}\mathbf{g}^{US}/(\mathbf{g}^{US})'\mathbf{Q}\mathbf{g}^{US}$			
ϕ^{corr}	US shock on both sides	unweighted, decade-lagged	1.87
Realized under IV: $\phi^{IV} = (\mathbf{g}^{OC})'\mathbf{QP}\mathbf{g}^{US}/(\mathbf{g}^{OC})'\mathbf{Q}\mathbf{g}^{US}$			
ϕ^{IV}_{raw}	raw bilinear, canonical exposures	unweighted, start-of-period / decade-lagged	2.39
ϕ^{IV}_{w}	WLS-faithful, weighted $\mathbf{Q}_w$	population-weighted, start-of-period / decade-lagged	2.67
ϕ^{IV}_{Neum}	bilinear on lagged-share $\mathbf{Q}_w$ (Table 6)	population-weighted, decade-lagged both sides	2.262
ϕ^{IV}	$1 + 2$SLS projection of N_r on B_r	controls-residualized, weighted, start-of-period / decade-lagged	2.316
ϕ^{IV}_{1reg}	coefficient ratio $\hat{\gamma}_{SS}/\hat{\gamma}_{NSS}$	controls-residualized, weighted, start-of-period / decade-lagged	2.403
Aggregate amplification: $\psi^{\text{agg}} = \sum_r w_r \tilde{B}_r / \sum_r w_r B_r$			
ψ^{agg}	unweighted-$\mathbf{Q}$ analog	population sums	2.766
ψ^{agg}_{w}	WLS-faithful	population sums, weighted $\mathbf{Q}_w$	2.77

Notes: Every amplification-factor variant used in the paper, with its definition, weighting and share vintage, and ADH value. "Decade-lagged" shares are the instrument's 1980/1991 shares; "start-of-period" shares are the endogenous 1991/1999 shares of the headline (Section 3.1); $\mathbf{S}_X$ and $\mathbf{S}_Z$ denote the endogenous and instrument share matrices. The single canonical pair used for the aggregate-cancellation statement (Proposition 7) is the WLS-faithful $\phi^{IV}_{w} = 2.67$ against $\psi^{\text{agg}}_{w} = 2.77$, a design-consistent gap of -3.5%; the controls-residualized projection factor $\phi^{IV} = 2.316$ (the value quoted as $\hat{\phi}^{IV}$ in the text, computed as one plus the population-weighted 2SLS projection of N_r on B_r with the full control set) and the unweighted $\psi^{\text{agg}} = 2.766$ should not be differenced across weighting conventions. The structural ϕ^{iid} is vintage- and weighting-robust (every variant near 1.1–1.2); the realized factors are larger because the ADH shocks are correlated within three-digit NAICS clusters that are co-located and IO-linked (Section 4.4).

Table 3. Standard vs. Network-Corrected Shift-Share Estimates: Manufacturing Employment

	(1)	(2)	(3)	(4)	(5)
Panel A: Standard Bartik					
Import Exposure	−0.406	−0.171	−0.280	−0.596	−0.582
	(0.048)	(0.028)	(0.059)	(0.100)	(0.105)
R^2	0.166	0.447	0.379	—	—
First-Stage F				47.6	46.6
Panel B: Network-Corrected Bartik ($\tilde{B}_r = \mathbf{s}_r'\mathbf{L}\mathbf{g}$)					
Network-Adjusted Exposure	−0.222	−0.191	−0.245	−0.255	−0.252
	(0.034)	(0.019)	(0.021)	(0.023)	(0.022)
R^2	0.228	0.494	0.462	—	—
First-Stage F				192.9	190.8
Memo: headline like-for-like comparison (both the paper's own BEA-detail Bartik with start-of-period endogenous shares, 2SLS with controls): standard $\tilde{B}_r$, -0.612 (0.066), $F = 61.7$; network-corrected $\tilde{B}_r$, -0.255 (0.023), $F = 192.9$; ratio 2.40. This is the headline standard/corrected pair, built identically except for the network adjustment; Panel A column (4) instead reports the Autor et al. (2013) replication on ADH's published exposure variable.					
Observations	1,444	1,444	1,444	1,444	1,444
Dep. Var. Mean	−1.623	−1.623	−1.623	−1.623	−1.623
Controls	No	Yes	No	Yes	Yes
State FE	No	No	Yes	No	Yes
Estimator	OLS	OLS	OLS	2SLS	2SLS

Notes: Dependent variable is the change in manufacturing employment / working-age population (percentage points). Panel A uses the standard ADH import-exposure measure and the other-country instrument; column (4) reproduces the Autor et al. (2013) 2SLS estimate. Panel B replaces both with network-adjusted measures, applying the BEA 1997 manufacturing Leontief submatrix ($K = 342$ industries) to BEA-level employment shares (County Business Patterns employment allocated to BEA codes via the concordance of Section 3.2); the endogenous Bartik uses start-of-period shares and its instrument decade-lagged shares. The memo line beneath Panel B reports the headline like-for-like comparison: the paper's own BEA-detail standard Bartik (-0.612) and the network-corrected estimate (-0.255) are built identically (BEA detail, start-of-period endogenous shares) and differ only in the network adjustment, so the headline ratio 2.40 is recoverable from a single construction. The BEA-detail standard Bartik (-0.612) differs from the Panel A column-(4) ADH replication (-0.596) only in using the paper's BEA-level shares rather than ADH's published exposure variable; the two are close. Table 4 catalogues how this headline construction relates to the lagged-share and coarser-aggregation estimates; the three-regressor full-economy OVB at $K = 483$ is in Table A.3 (Appendix C). Controls are Census division dummies, lagged manufacturing share, college share, foreign-born share, female employment share, routine occupation share, and offshorability index. Standard errors clustered at the state level in parentheses. Observations are weighted by start-of-period CZ population. $N = 722$ CZs $\times$ 2 periods (1990–2000, 2000–2007).

Table 4. Specification Map: Standard and Corrected Coefficients Across Industry Systems

Specification	Bartik construction	$\hat{\phi}$	Standard 2SLS	Corrected 2SLS
Headline (main text)	BEA detail (342), start-of-period	1.110	-0.612 (0.066)	-0.255 (0.023)
ADH replication	ADH original exposure variable	—	-0.596 (0.100)	—
BEA detail	BEA detail (342), lagged both sides	1.11	-0.266 (0.032)	-0.165 (0.016)
BEA summary	BEA summary (53), lagged both sides	1.26	-0.383 (0.064)	-0.287 (0.039)
Full economy	BEA full (483, 3-reg.)	—	—	-0.274 (direct)

Notes: Each row is a distinct shift-share construction; the coefficients differ because of the industry aggregation and the share vintage used to build the Bartik, not because of an inconsistency. All of the paper's own constructions (rows 1 and 3–5) are built at the BEA industry level: ADH's SIC87dd codes are used only to allocate County Business Patterns employment into BEA codes, so the resulting shares are indexed by BEA codes. The *headline* main-text specification (reported in Tables 1 and 8 and Section 4) builds the endogenous Bartik from start-of-period employment shares (1991/1999) and instruments it with the decade-lagged (1980/1991) other-country Bartik, matching ADH's timing (Section 3.1). The *BEA detail* and *BEA summary* rows are the fully predetermined variant with decade-lagged shares on both sides (the construction used for the jackknife and aggregation robustness checks, Tables 13 and 16); moving the endogenous Bartik from start-of-period to decade-lagged shares (with the industry classification held fixed at BEA detail) accounts for essentially all of the headline-versus-detail level difference (corrected: -0.255 to -0.165). Row 2 is the Autor et al. (2013) replication using their original published exposure variable, which the headline standard coefficient closely matches. The *full-economy* row is the three-regressor $K = 483$ OVB decomposition (Table A.3; Appendix C), for which the single standard/corrected split does not apply. First-stage F statistics are 61.7 (headline standard), 47.6 (ADH replication), and 192.9 (headline corrected). Standard errors (state-clustered, population-weighted) in parentheses; samples as in the cited tables.

Table 5. Bias Decomposition: Direct vs. Network-Mediated Effects

	ADH	OVB Decomposition		AADHP (2016)	
	2SLS	OLS	2SLS	OLS	2SLS
	(1)	(2)	(3)	(4)	(5)
Direct (β_D)		-0.118	-0.305	-0.930	-2.051
		(0.057)	(0.107)	(0.319)	(0.422)
Network-only (β_N)		-0.232	-0.233		
		(0.036)	(0.050)		
Upstream (AADHP)				$+0.871$	$+1.722$
				(0.345)	(0.462)
Standard Bartik	-0.596				
	(0.100)				
Observations	1,444	1,444	1,444	1,444	1,444
First-Stage F	47.6		—		—
SW F, direct regressor			24.3		5.1
SW F, network regressor			219.7		5.1

Notes: Dependent variable: change in manufacturing employment / working-age population. Column 1: Autor et al. (2013) 2SLS replication. Columns 2–3: direct (B_r) and network-only ($N_r = \mathbf{s}_r'(\mathbf{L} - \mathbf{I})\mathbf{g}$) regressors, instrumented by other-country analogs. Columns 4–5: the Acemoglu et al. (2016) (AADHP) specification as those authors implement it, augmenting the ADH import-penetration exposure with a separately constructed upstream net-input IO exposure, each instrumented by its other-country analog. The two trade-exposure regressors are nearly collinear (correlation 0.988), so the augmented 2SLS is weakly identified and returns large, nearly canceling coefficients; this is the augmented-regressor instability the identity-subtracted (B_r, N_r) decomposition in columns 2–3 is designed to avoid. All include period and Census division fixed effects with lagged CZ controls. State-clustered standard errors. Columns 3 and 5 report Sanderson–Windmeijer first-stage F-statistics per endogenous variable; the AADHP F-statistics fall below conventional weak-instrument thresholds, reflecting the near-collinearity.

Table 6. Neumann-Series Decomposition of the Bias Factor ϕ

Order n	Structural i.i.d. $\mathrm{tr}(\mathbf{Q}\mathbf{A}^n)/\mathrm{tr}(\mathbf{Q})$	ADH realized (IV) bilinear, weighted	BK realized (OLS) quadratic, weighted
$n = 1$	0.1237	0.8034	0.5706
$n = 2$	0.0282	0.2967	0.2117
$n = 3$	0.0080	0.1030	0.0755
$n = 4$	0.0026	0.0371	0.0279
$n = 5$	0.0009	0.0137	0.0107
$n = 6$	0.0003	0.0052	0.0042
$n = 7$	0.0001	0.0020	0.0017
$n = 8$	0.0000	0.0008	0.0007
Cumulative $1 + \sum_{n=1}^{8} \phi^{(n)}$	1.1638	2.2618	1.9029
Residual ($n > 8$)	0.0000	0.0005	0.0005
Exact ϕ	1.1639	2.2624	1.9035
Share of $(\phi - 1)$ at $n \leq 1$	75.5%	63.6%	63.2%
Share of $(\phi - 1)$ at $n \leq 2$	92.7%	87.1%	86.6%
Share of $(\phi - 1)$ at $n \leq 5$	99.7%	99.3%	99.2%

Notes: Neumann decomposition of the bias factor $\phi = 1 + \sum_{n=1}^{\infty} \phi^{(n)}$, where each order n corresponds to supplier-chain depth: $n = 1$ captures direct supplier linkages, $n = 2$ suppliers of suppliers, and so on. All columns use the population-weighted co-location matrix $\mathbf{Q}^w = \mathbf{S}' \,\mathrm{diag}(w)\mathbf{S}$ and the BEA 1997 342-industry manufacturing $\mathbf{A}$ matrix. Column 1 reports the structural i.i.d. decomposition $\phi_{\mathrm{iid}}^{(n)} = \mathrm{tr}(\mathbf{Q}^w\mathbf{A}^n)/\mathrm{tr}(\mathbf{Q}^w)$, which depends only on $\mathbf{Q}^w$ and $\mathbf{A}$. Column 2 reports the ADH realized IV decomposition evaluated on the actual US and other-country shock vectors. Column 3 reports the Blanchard-Katz realized OLS decomposition. The structural i.i.d. column decays geometrically, at order-to-order ratios of 0.23–0.35 – below the asymptotic spectral-radius bound $\rho(\mathbf{A}) = 0.506$, which caps the limiting rate but not the transient one; the realized columns decay at 0.35–0.37, somewhat slower because correlated shocks preferentially activate higher-order supplier chains. The cumulative rows show how close the truncated sum comes to the exact ϕ.

Table 7. Rotemberg-Weight Decomposition of the Standard and Network-Corrected Bartik Estimators

Industry (BEA detail code)	α_k^{std}	α_k^{corr}	$\hat{\beta}_k^{\text{std}}$	$\hat{\beta}_k^{\text{corr}}$
Panel A: top 5 industries by $\lvert\alpha_k^{\text{std}}\rvert$				
Electronic computer mfg. (334111)	0.2489	0.1537	-0.531	-0.164
Cut and sew apparel mfg. (315200)	0.1005	0.0182	-1.045	-1.091
Doll, toy, and game mfg. (339930)	0.0911	0.0168	-0.162	-0.144
Footwear mfg. (316200)	0.0619	0.0095	-0.650	-0.694
Other computer peripheral equipment mfg. (334119)	0.0566	0.0348	-0.336	-0.181
Panel B: top 5 industries by $\lvert\alpha_k^{\text{corr}}\rvert$				
Semiconductors and related device mfg. (334413)	0.0458	0.2312	-0.753	-0.178
Electronic computer mfg. (334111)	0.2489	0.1537	-0.531	-0.164
All other electronic component mfg. (33441A)	0.0422	0.1473	-0.504	-0.150
Iron and steel mills (331111)	0.0090	0.0926	-2.012	-0.204
Other computer peripheral equipment mfg. (334119)	0.0566	0.0348	-0.336	-0.181
Panel C: top 5 industries by $\lvert\alpha_k^{\text{corr}} - \alpha_k^{\text{std}}\rvert$				
Semiconductors and related device mfg. (334413)	0.0458	0.2312	-0.753	-0.178
All other electronic component mfg. (33441A)	0.0422	0.1473	-0.504	-0.150
Electronic computer mfg. (334111)	0.2489	0.1537	-0.531	-0.164
Iron and steel mills (331111)	0.0090	0.0926	-2.012	-0.204
Cut and sew apparel mfg. (315200)	0.1005	0.0182	-1.045	-1.091

Notes: Rotemberg decomposition (Goldsmith-Pinkham et al., 2020) of the standard and network-corrected Bartik 2SLS estimators. Each 2SLS coefficient writes as $\hat{\beta} = \sum_{k,t} \alpha_{k,t} \hat{\beta}_{k,t}$, where $\hat{\beta}_{k,t}$ is the industry-period just-identified IV estimate on the residualized outcome and $\alpha_{k,t}$ is the Rotemberg weight. Weights are aggregated across the two periods; $\hat{\beta}_k$ is the weighted average of $\hat{\beta}_{k,t}$. Panel A reports the top five industries by $\lvert\alpha_k^{\text{std}}\rvert$, dominated by direct China Shock final-goods targets. Panel B reports the top five by $\lvert\alpha_k^{\text{corr}}\rvert$: the correction loads weight onto upstream industrial suppliers (semiconductors, other electronic components, iron and steel mills). Panel C reports the top five by the change $\lvert\alpha_k^{\text{corr}} - \alpha_k^{\text{std}}\rvert$. The Herfindahl effective number of industries moves from 10.0 (standard) to 8.9 (corrected); the Pearson correlation between the two weight vectors is 0.596.

Table 8. What Changes After Correction

Quantity	Standard	Corrected	Ratio or gap
Panel A: Per-unit coefficient			
2SLS coefficient	-0.612	-0.255	2.40
	(0.066)	(0.023)	
First-stage F	61.7	192.9	—
Panel B: Predicted impact at the IQR of exposure			
Exposure IQR (pp)	1.74	4.70	$4.70/1.74 \approx 2.7$
Implied effect at IQR (pp)	-1.063	-1.196	0.89
Panel C: Contamination and amplification factors			
Structural $\hat{\phi}^{\text{iid}} = \text{tr}(\mathbf{QL})/\text{tr}(\mathbf{Q})$		1.110	
OLS correlated-shocks $\hat{\phi}^{\text{corr}} = \mathbf{g}^{US\prime}\mathbf{QLg}^{US}/\mathbf{g}^{US\prime}\mathbf{Qg}^{US}$		1.87	
Realized 2SLS $\hat{\phi}^{IV}(\mathbf{g}^{OC}, \mathbf{g}^{US})$		2.316	
Aggregate $\hat{\psi}^{\text{agg}} = \sum_r w_r \tilde{B}_r / \sum_r w_r B_r$		2.766	
Single-regressor projection weight $\hat{\lambda}^{\text{struct}}$		0.76	
realized $\hat{\lambda}^{IV}$		0.45	
implied $\hat{\lambda}^{\text{impl}}$		0.30	
Wald test $H_0 : \beta_D = \beta_N$ (p-value)		0.63	
Neumann share $n \leq 2$ (realized)		87.1%	
Panel D: Geographic reallocation			
Top-25 miss rate		68%	
Top-100 miss rate		34%	
Cross-CZ share of MSE		99.8%	
Bivariate R^2 against $\Delta\mathcal{L}^{\text{mfg}}$	0.080	0.262	$\approx 3.3\times$

Notes: All entries are pooled ADH estimates (1990–2000 and 2000–2007, $N = 1,444$ CZ-periods), state-clustered standard errors, population-weighted. Panel A: the main-text headline standard 2SLS (-0.612) and single-regressor network-corrected 2SLS (-0.255), both built at the BEA detail level with a start-of-period-share endogenous Bartik; these headline coefficients, and how they relate to Table 3's ADH-replication column (-0.596), are catalogued in Table 4. Panel B: IQR-implied effects from $\hat{\gamma} \cdot \text{IQR}$ with the single-regressor corrected estimator replacing exposure B_r with $\tilde{B}_r$. Panel C: structural depends only on $(\mathbf{S}, \mathbf{L})$ (Proposition 1); realized IV is the controls-residualized realized factor of Section 4.4 (Corollary 1 gives the population object); aggregate amplification is from Corollary 3; Neumann share is the fraction of $\phi^{\text{ADH}} - 1$ contributed by the first two orders $n = 1, 2$ (Table 6). Panel D: geographic-misattribution diagnostics from Section 5.

Table 9. Commuting Zones with the Largest Exposure-Rank Changes Under the Network Correction

				Exposure rank			Lagged
#	CZ code	State	Top mfg. industries	Std.	Net.	Δ rank	mfg. sh.
Panel A: Commuting zones that become MORE exposed (input-intensive; understated by standard Bartik)							
1	31007	TX	Food, Petroleum/coal	607	102	-505	33.625
2	26004	MN	Paper, Wood	469	8	-461	31.586
3	31001	KS	Food, Fab. metals	637	196	-441	32.505
4	29003	KS	Food, Machinery	597	210	-387	35.167
5	31006	OK	Fab. metals, Printing	642	259	-383	27.884
6	31004	KS	Food, Machinery	644	268	-376	30.802
7	26104	MO	Food, Transp. equip.	437	124	-313	45.667
8	38401	ID	Wood, Paper	540	237	-303	20.669
9	10802	AL	Paper, Apparel	316	17	-299	35.008
10	26105	MO	Food, Printing	483	185	-298	55.242
Panel B: Commuting zones that become LESS exposed (input-simple; overstated by standard Bartik)							
1	26410	ND	Furniture, Machinery	123	432	+309	10.987
2	28304	NE	Misc. mfg, Nonmet. min.†	205	511	+306	2.708
3	26404	ND	Misc. mfg, Machinery	155	449	+294	9.587
4	30602	NM	Elec. equip., Wood	200	494	+294	4.813
5	33601	OK	Plastics/rubber, Machinery	113	376	+263	15.527
6	29104	KS	Food, Elec. equip.	131	391	+260	22.818
7	32306	TX	Apparel, Fab. metals	58	317	+259	14.452
8	32304	TX	Apparel, Food	224	478	+254	10.748
9	20002	ME	Paper, Food	148	397	+249	18.537
10	03201	LA	Wood, Elec. equip.	68	304	+236	22.044

Notes: Ten commuting zones with the largest negative (Panel A) and positive (Panel B) exposure-rank changes between the standard and network-corrected Bartik measures for the 2000–2007 China Shock window ($N = 714$ CZs, Period 2). Ranks are assigned with 1 as most exposed under each measure; Δ rank = Net. rank − Std. rank (negative values move up, more exposed). Lagged manufacturing share is the 1990 CZ manufacturing employment share (percent). The "Top mfg. industries" column lists the two largest three-digit NAICS manufacturing categories by CZ employment share, from the 1991 lagged share matrix. A dagger (†) indicates a CZ whose 1991 shares are suppressed and the 1999 shares are substituted. Over the full sample the Spearman rank correlation is 0.87; confidence intervals for the reallocation statistics are in Table 10.

Table 10. Bootstrap Confidence Intervals on the Reallocation Statistics

Statistic	Point estimate	95% CI lower	95% CI upper
Top-25 targeting error (%)	68.0	64.2	79.5
Top-100 targeting error (%)	34.0	31.1	35.0
Mean |percentile shift|	10.4	10.0	10.7
Spearman ρ (standard vs. corrected)	0.874	0.865	0.883
Geographic share of MSE (%)	99.8	99.4	100.0

Notes: Point estimates on the full period-2 sample (722 commuting zones, 2000–2007). Confidence intervals are Politis-Romano pivotal intervals from a state-block subsampling bootstrap: each of 1000 replications drops 4 of 48 states without replacement, refits the three IV specifications on the surviving panel, and recomputes the statistic on period 2. We subsample without replacement because rank-threshold statistics are mechanically biased upward under with-replacement bootstraps (repeated draws force top-K overlap below 100%). The kept subsample is large relative to the panel ($b = 44$ of $n = 48$ states, $b/n \approx 0.92$, closer to a delete-d jackknife than to small-b subsampling); the intervals apply the $\sqrt{b/n}$ pivotal rescaling of Politis and Romano (1994) and should be read as approximate. The top-K targeting error is the percentage of CZs in the standard top K that are not in the corrected top K. The geographic share of MSE is $\mathrm{var}(\mathrm{pe}_r - \overline{\mathrm{pe}})/\mathrm{mean}(\mathrm{pe}_r^2)$ for the CZ-level prediction-error difference $\mathrm{pe}_r = \hat{\gamma}_{\mathrm{std}} B_r - \hat{\beta}\tilde{B}_r$. The CI upper bound on the MSE share is physically truncated at 100%.

Table 11. Welfare Loss from Using the Standard Bartik Ranking to Allocate a Fixed Trade Adjustment Assistance Budget

		Set overlap		Need-wgt. welfare		Efficiency loss	
K	Top-K % of N	$\lvert A \cap B \rvert$ (CZs)	Share (%)	Std. W^{std}	Corr. W^{corr}	Need-wgt. $1 - W^{\text{std}}/W^{\text{corr}}$	Pop.-wgt. $1 - W^{\text{std}}/W^{\text{corr}}$
10	1.4	3	30	0.020	0.048	57.7%	37.7%
25	3.5	8	32	0.072	0.086	16.2%	8.4%
36	5.0	18	50	0.089	0.102	13.0%	5.2%
50	7.0	29	58	0.111	0.133	16.1%	11.5%
72	10.1	44	61	0.150	0.168	11.0%	2.5%
100	14.0	66	66	0.179	0.209	14.4%	8.8%
143	20.0	94	66	0.250	0.283	11.8%	6.0%

Notes: Welfare-loss exercise for a fixed TAA-style budget allocated equally across the K most-exposed commuting zones under the standard (A) vs. network-corrected (B) Bartik ranking. Sample is the $N = 714$ valid commuting zones in Period 2 (2000–2007). Observed pre-Shock TAA outlays were approximately $321 million, computed from the Autor et al. (2013) replication files' per-capita TAA transfers summed over CZs in the base year; the welfare-loss ratios are invariant to the absolute budget. The set-overlap column $\lvert A \cap B \rvert$ counts CZs in both top-K sets. Need-weighted welfare is $W^{\text{spec}} = \sum_r n_r \cdot \mathbb{1}\{r \in \text{top-}K^{\text{spec}}\} / \sum_r n_r$ with n_r proportional to working-age population times positive corrected shock incidence per worker. Efficiency loss columns report $1 - W^{\text{std}}/W^{\text{corr}}$ under the need-weighted and population-only weightings. Both columns would read 0% under a planner with the corrected ranking, so their values measure the welfare cost of misattributing exposure.

Table 12. Pre-Trend Falsification: Standard and Corrected P2 Bartik vs Pre-1990 CZ Manufacturing Employment Change

	Pre-periods		In-sample periods	
	1970–1980	1980–1990	1990–2000	2000–2007
	(1)	(2)	(3)	(4)
Standard Bartik B_r^{P2}	0.235	0.021	−0.162	−0.393
	(0.058)	(0.039)	(0.038)	(0.053)
	[$p = 0.000$]	[$p = 0.585$]		
Corrected Bartik $\tilde{B}_r^{P2}$	0.105	−0.014	−0.071	−0.206
	(0.031)	(0.017)	(0.022)	(0.031)
	[$p = 0.001$]	[$p = 0.407$]		
Observations	722	722	722	722
Estimator	2SLS	2SLS	2SLS	2SLS

Notes: Dependent variable is the CZ-level change in manufacturing employment share during the period in each column header (percentage points). Columns (1)–(2) are pre-period falsification tests following Autor et al. (2013) Table 8: the Bartik regressor B_r^{P2} (standard) or $\tilde{B}_r^{P2}$ (corrected) is constructed from the forward $\mathbf{g}^{US,P2}$ shock vector (the 2000–2007 China import shock at BEA detail level) and 1990 lagged shares, identical across columns; the instrument is the corresponding other-country analog $\mathbf{g}^{OC,P2}$. The pre-period outcomes are drawn from the ADH replication data at year $t \in \{1970, 1980\}$. Columns (3)–(4) report in-sample analogs. All specifications include Census division fixed effects, are weighted by the start-of-period CZ population, and cluster standard errors at the state level.

Table 13. Single-Regressor Finite-Shock Jackknife: Leave-One-Industry-Out and Leave-One-NAICS3-Cluster-Out

	Standard Bartik	Corrected Bartik
Panel A: Baseline (BEA-detail, pooled 2SLS)		
Coefficient	-0.266	-0.165
State-clustered SE	(0.032)	(0.016)
Panel B: Leave-one-industry-out (342 replicates)		
Jackknife SD	0.0746	0.0663
Max $\lvert\hat{\beta}_j - \hat{\beta}_0\rvert$	0.0512	0.0589
Most influential industry	BEA 334111 (electronic computers)	
Panel C: Leave-one-NAICS3-cluster-out (21 replicates)		
Jackknife SD	0.1001	0.1213
Max $\lvert\hat{\beta}_c - \hat{\beta}_0\rvert$	0.0978	0.1263
Most influential cluster	NAICS 334 (computer and electronic products, 22 detail codes)	

Notes: Pooled 2SLS on the BEA-detail panel ($N = 1,444$ CZ-periods, $K = 342$ manufacturing industries). Dependent variable: change in CZ manufacturing employment share (pp). Endogenous variable: standard Bartik $B_r = \mathbf{s}_r'\mathbf{g}$ in column 1 and corrected Bartik $\tilde{B}_r = \mathbf{s}_r'\mathbf{L}\mathbf{g}$ in column 2; each instrumented by its other-country analog. Controls: period dummy, Census division FE, lagged CZ characteristics (ADH main-specification control set), weighted by start-of-period CZ population, state-clustered standard errors. Panel B jackknife: for each $k \in \{1, \ldots, 342\}$ zero g_k^{US} and g_k^{OC} in both periods, rebuild B_r, $\tilde{B}_r$, and their IV analogs, rerun 2SLS. Panel C jackknife: same, with BEA detail codes grouped by modal three-digit NAICS. Jackknife SD is $\sqrt{\frac{n-1}{n}\sum_j(\hat{\beta}_j - \bar{\hat{\beta}})^2}$; Max is $\max_j \lvert\hat{\beta}_j - \hat{\beta}_0\rvert$ against the baseline.

Table 14. Robustness to IO Table Vintage (1997–2017)

	1997	2002	2007	2012	2017
Manufacturing only (K^{mfg})					
$\hat{\phi}^{mfg}$	1.100	1.112	1.135	1.137	1.138
K^{mfg} mapped	342	342	334	334	334
Full economy (K^{full})					
$\hat{\phi}^{full}$	1.148	1.123	1.110	1.110	1.120
K^{full} mapped	483	478	446	446	446
Corrected 2SLS	-0.255	-0.374	—	—	—
	(0.023)	(0.046)			

Notes: $\hat{\phi} = \text{tr}(\mathbf{QL})/\text{tr}(\mathbf{Q})$ computed for each BEA detailed benchmark IO table. Industry codes mapped across vintages using a hierarchical concordance. Coverage is 98–100% of 1997 manufacturing codes and 92–99% of full-economy codes. The corrected 2SLS is re-estimated only for the 1997 and 2002 benchmarks, the vintages that predate or overlap the 1991–2007 sample window; the 2007–2017 benchmarks postdate the treatment period (and shift to NAICS-basis classifications with partial concordance to the estimation panel), so they enter the table only through the $\hat{\phi}$ diagnostic.

Table 15. Cluster Decomposition of the 1997–2017 Drift in $\hat{\phi}^{\text{mfg}}$

NAICS-3 cluster	K_C	Drift contrib. $\Delta\hat{\phi}_C$	Share of total $(\Delta\hat{\phi}_C/\Delta\hat{\phi})$	China shock $\bar{g}_C^{\text{USCH}}$	1997 IO density $\overline{(\mathbf{L}-\mathbf{I})_{\text{rs},C}}$
Food (311)	15	+0.0195	+50%	2.78	0.076
Transp. equip. (336)	16	+0.0097	+25%	2.01	0.097
Electrical equip. (335)	11	+0.0040	+10%	17.93	0.270
Machinery (333)	15	+0.0031	+8%	10.36	0.144
Plastics/rubber (326)	7	+0.0022	+6%	9.84	0.443
Fabricated metals (332)	13	+0.0019	+5%	6.65	0.720
Chemicals (325)	11	+0.0016	+4%	8.04	1.382
Misc. mfg (339)	9	+0.0011	+3%	39.05	0.039
Petroleum/coal (324)	3	+0.0006	+2%	1.20	2.218
Paper (322)	4	+0.0005	+1%	2.46	1.206
Nonmetallic minerals (327)	8	+0.0005	+1%	5.88	0.113
Bev./tobacco (312)	4	+0.0004	+1%	0.19	0.031
Furniture (337)	5	+0.0000	+0%	23.82	0.017
Primary metals (331)	1	−0.0000	-0%	1.53	0.892
Textile prods. (314)	2	−0.0014	-3%	33.83	0.040
Textile mills (313)	1	−0.0018	-5%	0.68	0.505
Computers/electronics (334)	17	−0.0031	-8%	52.42	1.288
Total drift	142	+0.0389	100%	—	—

Notes: Cluster decomposition of the 1997–2017 drift in $\hat{\phi}^{\text{mfg}} = \text{tr}(\mathbf{QL})/\text{tr}(\mathbf{Q})$, using the trace identity $\text{tr}(\mathbf{QL}) = \sum_k (\mathbf{QL})_{k,k}$ to split per-industry contributions, then aggregated to NAICS-3 manufacturing clusters. Drift shares sum to 100% by construction. Sample is the $K = 142$ BEA detail codes appearing verbatim in both the 1997 and 2017 vintage classifications; the restricted-set drift (+0.0389) matches the headline-panel drift (+0.0380) to within 0.001. $\mathbf{Q}$ is held fixed at 1991 lagged-shares, isolating drift in $\mathbf{L}$ from drift in shares. The China-shock column is the employment-weighted average USCH import shock per worker for 2000–2007; the IO-density column is the employment-weighted row-sum of $(\mathbf{L}_{1997} - \mathbf{I})$, an industry-level measure of upstream linkage strength.

Table 16. Sensitivity to Industry Aggregation Level

	Summary (53)	Detail (342)
$\hat{\phi}$	1.26	1.11
Standard 2SLS	-0.383	-0.266
	(0.064)	(0.032)
Corrected 2SLS	-0.287	-0.165
	(0.039)	(0.016)

Notes: Summary = 53 BEA summary manufacturing industries; Detail = 342 BEA detailed manufacturing industries. Both columns build the entire Bartik—shares, shock, and Leontief—on BEA industries (the detail column is the BEA-detail baseline of Section 6.2), and both use BEA 1997 benchmark IO tables. The main-text headline coefficients (-0.612 standard, -0.255 corrected) use start-of-period (rather than decade-lagged) shares for the endogenous Bartik, not a different industry classification, and are catalogued in Table 4. 2SLS with full controls; state-clustered standard errors.

Table 17. Inference Comparison: State-Clustered vs. Exposure-Robust Standard Errors

	SE Type	Coef	SE	t-stat
Standard 2SLS	State-clustered	-0.612	0.0659	-9.28
	AKM exposure-robust	-0.612	0.0845	-7.24
Corrected 2SLS	State-clustered	-0.255	0.0228	-11.19
	AKM naive (original $\mathbf{S}$)	-0.255	0.0170	-15.00
	AKM correct ($\tilde{\mathbf{S}} = \mathbf{SP}$)	-0.255	0.0263	-9.68

Notes: All specifications include period dummies, Census division dummies, and lagged CZ controls, weighted by start-of-period CZ population. State-clustered: Eicker-Huber-White SEs clustered at the state level. AKM: exposure-robust SEs following Adão et al. (2019), implemented per Proposition 4: the score aggregates residuals over the instrument's decade-lagged shares at the primitive-shock level and multiplies by the residualized other-country shock ($\hat{v}_l = \tilde{g}_l^{OC} \sum_r w_r \tilde{s}_{rl}^Z \hat{\varepsilon}_r$, shocks demeaned within period with importance weights), with the denominator built from control-residualized instrument and endogenous exposures. For the corrected estimator, "naive" AKM aggregates over the original lagged shares $\mathbf{S}$; "correct" AKM uses the network-adjusted shares $\tilde{\mathbf{S}} = \mathbf{SP}$. $N = 1{,}444$.

Table 18. Heterogeneity by IO Network Density Tercile

	Low	Medium	High
Standard 2SLS	−0.645	−0.537	−0.293
	(0.137)	(0.091)	(0.128)
Corrected 2SLS	−0.197	−0.250	−0.298
	(0.014)	(0.031)	(0.044)
Std/Corr ratio	3.28	2.15	0.98
Standard: max−min spread		0.353	
Joint Wald p (df = 2)		0.003	
Corrected: max−min spread		0.101	
Joint Wald p (df = 2)		0.002	

Notes: Commuting zones sorted into terciles by employment share in high-IO-density industries. 2SLS with full controls; state-clustered standard errors.

Table 19. External Validation: Network Exposure and Independently Measured Distress

	Distress outcome (CZ × period)	Source	$\hat{\delta}$ on N_r (s.e.)	p
O1	Manufacturing employment, $\Delta \log$ ($\downarrow$)	CBP	-0.0128 (0.0039)	0.001
O2	Upstream-supplier employment, $\Delta \log$ ($\downarrow$)	CBP	-0.0264 (0.0047)	0.000
O3	Establishment-exit rate, cumulative ($\uparrow$)	BDS	0.0044 (0.0022)	0.043
O4	Job-destruction rate, cumulative ($\uparrow$)	BDS	-0.0037 (0.0030)	0.221

Notes: Each row is a separate regression of equation (35): the distress outcome on network-only exposure $N_r = \tilde{B}_r - B_r$, direct exposure B_r, the standard CZ controls, and period fixed effects; population-weighted, state-clustered standard errors in parentheses; CZ × period observations ($N \approx 1{,}440$). The arrow marks the distress direction (employment outcomes fall, $\downarrow$; exit and destruction rates rise, $\uparrow$), so a coefficient of that sign with small p indicates that network-up-weighted commuting zones suffer more distress than their direct exposure B_r alone predicts. CBP = County Business Patterns establishment employment (manufacturing, SIC 2000–3999). BDS = Census Bureau Business Dynamics Statistics (manufacturing, NAICS 31–33), aggregated from county × sector to commuting zones via the standard county–CZ crosswalk (707 of 722 CZs covered; suppressed small cells dropped). Upstream-supplier industries (O2) are the 159 of 342 BEA industries with $[(\mathbf{L} - \mathbf{I})\mathbf{g}]_l > g_l$ (output pulled mainly by other industries' shocks). A parallel specification replacing N_r with the regional network multiplier m_r yields the same qualitative sign and significance pattern.

Table 20. Employment-Growth Bartik: Standard vs. Network-Adjusted

	Standard	Network-Adjusted
Employment growth Bartik	1.163	0.518
	(0.587)	(0.300)
Coefficient ratio		2.25
i.i.d. $\hat{\phi}$		1.110
Observations	1,444	1,444
Estimator	OLS	OLS

Notes: Reduced-form regressions of the change in manufacturing employment / working-age population on the BK Bartik instrument. The standard BK uses national industry employment growth rates as shocks; the network-adjusted BK pre-multiplies the same shocks by the BEA 1997 Leontief inverse. Industry shares at SIC87dd level from the Goldsmith-Pinkham et al. (2020) replication data, aggregated to BEA detail. Period dummies included; state-clustered standard errors; observations weighted by start-of-period CZ population. $N = 722$ CZs $\times$ 2 periods.

Table 21. BK Channel Decompositions: Two- and Three-Regressor OLS

	Manufacturing Employment Growth
Panel A: Two-regressor decomposition (B_r^{BK}, N_r^{BK})	
Direct B_r^{BK} (β_D)	1.133
	(0.639)
Network-only N_r^{BK} (β_N)	0.038
	(0.313)
Channel ratio $\hat{\theta}^{BK} = \hat{\beta}_N / \hat{\beta}_D$	0.03
Panel B: Three-regressor decomposition $(B_r^{BK}, N_r^{L,BK}, N_r^{G,BK})$	
Direct B_r^{BK} (β_D)	1.110
	(0.643)
Upstream $N_r^{L,BK}$ (β_L)	0.004
	(0.344)
Downstream $N_r^{G,BK}$ (β_G)	0.035
	(0.177)
Residualized condition number	2.47
VIF (downstream regressor)	1.96
Observations	1,444
Estimator	OLS

Notes: Two- and three-regressor OLS decompositions of the BK employment-growth Bartik at manufacturing-only granularity. Panel A splits exposure into the direct Bartik B_r^{BK} and the network-only exposure $N_r^{BK} = \mathbf{s}_r'(\mathbf{L} - \mathbf{I})\mathbf{g}^{BK}$; Panel B further separates upstream (Leontief) and downstream (Ghosh) channels. The condition-number and VIF diagnostics refer to the three-regressor design. Period dummies; state-clustered standard errors; population-weighted.

Table 22. Network Bias Across Applications

	China Shock (ADH)	Employment Growth (BK)
Structural diagnostic		
i.i.d. $\hat{\phi}$	1.110	1.110
Coefficient ratio		
Standard coefficient	-0.612	1.163
Corrected coefficient	-0.255	0.518
Coefficient ratio $\hat{\phi}^{\text{realized}}$	2.40	2.25
Estimator	2SLS	OLS
Magnitude cancellation		
Pop-weighted $\hat{\psi}^{\text{agg}}$	2.766	1.997
Aggregate gap $\hat{\psi}^{\text{agg}} - \hat{\phi}^{\text{realized}}$	$+0.36$	-0.25
Implied coefficient-vs-aggregate gap (signed)	-13.1%	$+12.5\%$
Bootstrap p-value (null: exact cancel)	0.20	0.69
Channel ratio		
$\hat{\theta} = \hat{\beta}_N / \hat{\beta}_D$	non-negligible	0.03 (near zero)
Cancellation mechanism	structural (Cov identity)	mechanical ($\beta_N \approx 0$)

Notes: Structural vs. mechanical cancellation. In ADH, θ is meaningful and the aggregate ratio approaches one because $\phi \approx \psi^{\text{agg}}$ via the negative Cov term of Proposition 6. In BK, $\theta \approx 0$: the per-unit network channel has essentially no causal effect on the outcome, so coefficient inflation mechanically cancels regardless of $\phi - \psi^{\text{agg}}$. Both produce standard-vs-corrected aggregate agreement at the IQR, but only ADH's is a structural property of $(\mathbf{Q}, \mathbf{L})$ and the population weights; BK's is a feature of the particular causal DGP at hand.

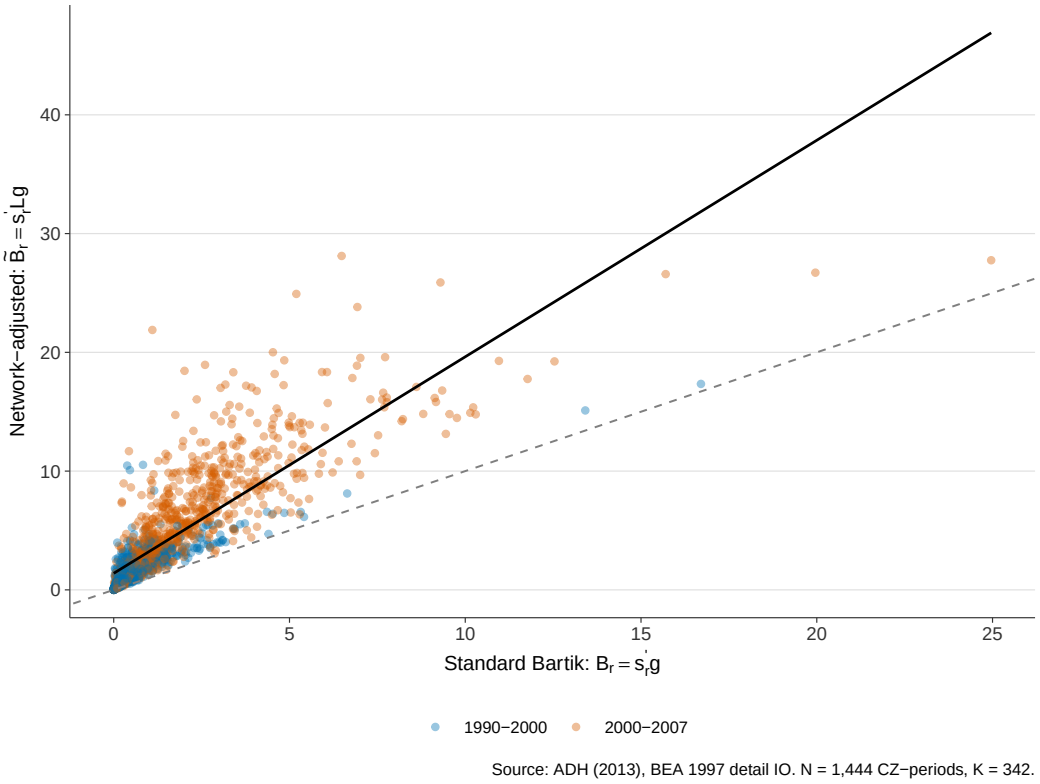


Figure 1. Standard vs. Network-Adjusted Bartik Exposure

Notes: Each point is a commuting zone. Horizontal axis: standard Bartik exposure $B_r = \mathbf{s}_r'\mathbf{g}$; vertical axis: network-adjusted exposure $\tilde{B}_r = \mathbf{s}_r'\mathbf{L}\mathbf{g}$. The 45-degree line is shown for reference. $N = 1,444$ (722 CZs $\times$ 2 periods).

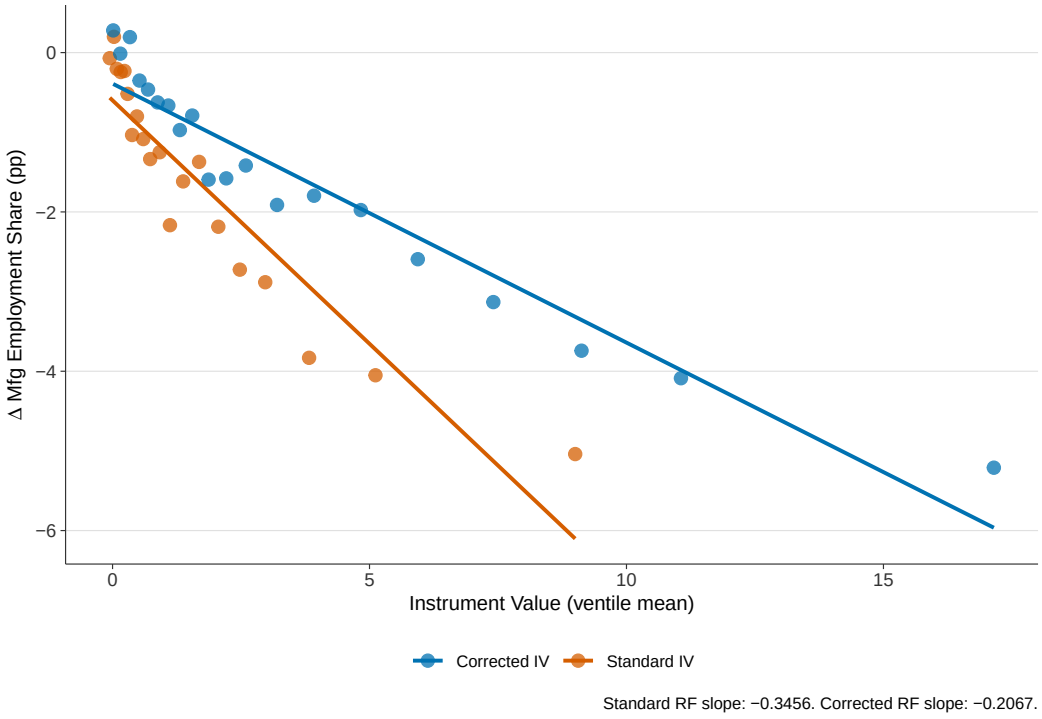


Figure 2. Reduced-Form Relationship: Standard vs. Network-Adjusted Instrument

Notes: Binscatter of the change in manufacturing employment share (vertical axis) against the other-country instrument (horizontal axis), with the standard and network-adjusted instruments overlaid on a single panel and distinguished by color. Linear fits are shown for each measure. $N = 1,444$.

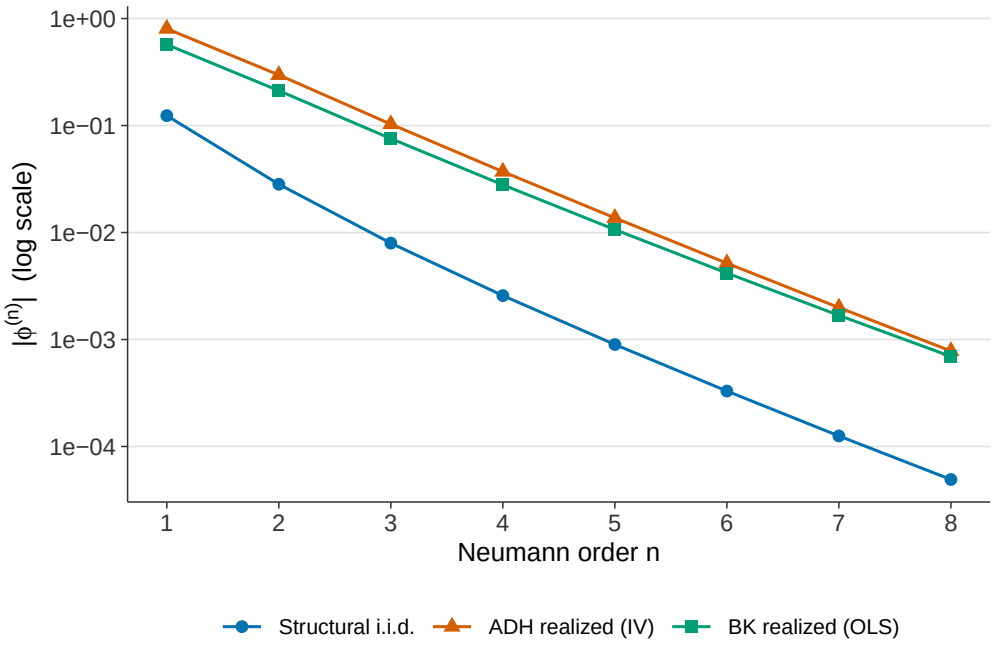


Figure 3. Neumann-order decay of $\phi^{(n)}$

Notes: $|\phi^{(n)}|$ on a log scale against Neumann order $n = 1, \ldots, 8$. Three series: structural i.i.d. ($\phi^{(n)} = \mathrm{tr}(\mathbf{Q}^w \mathbf{A}^n)/\mathrm{tr}(\mathbf{Q}^w)$ with population-weighted co-location matrix and BEA 1997 342-industry $\mathbf{A}$); ADH realized (bilinear form on ADH other-country and U.S. shock vectors); BK realized (quadratic form on BK shock vector). Spectral radius $\rho(\hat{\mathbf{A}}) = 0.506$; first two orders account for 92.7% of $\phi^{\mathrm{iid}} - 1$ and 87.1% of $\phi^{\mathrm{ADH}} - 1$.

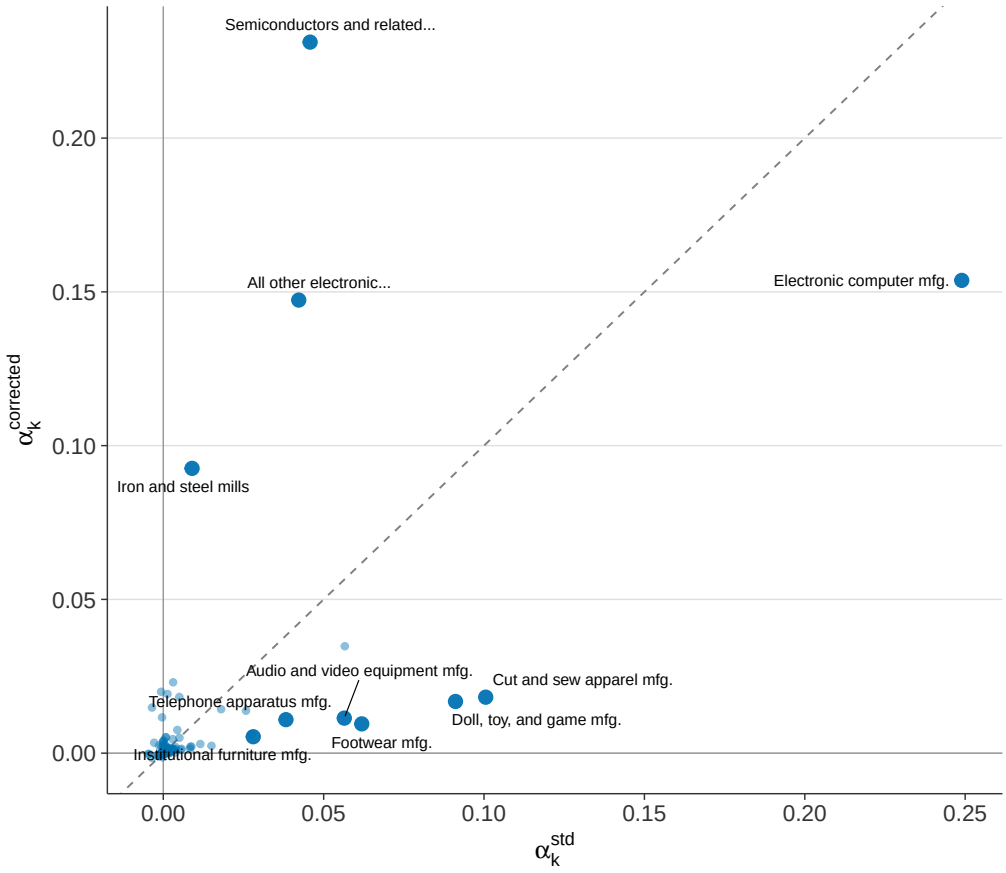


Figure 4. Standard versus Network-Corrected Rotemberg Weights

Notes: Each point is a BEA 1997 detail manufacturing industry; axes are the standard Bartik Rotemberg weight α_k^{std} and the network-corrected weight α_k^{corr}. The 45-degree line marks industries whose weight is unchanged. Pearson correlation 0.596; Spearman rank correlation 0.691. The ten industries with the largest absolute weight movement are labeled.

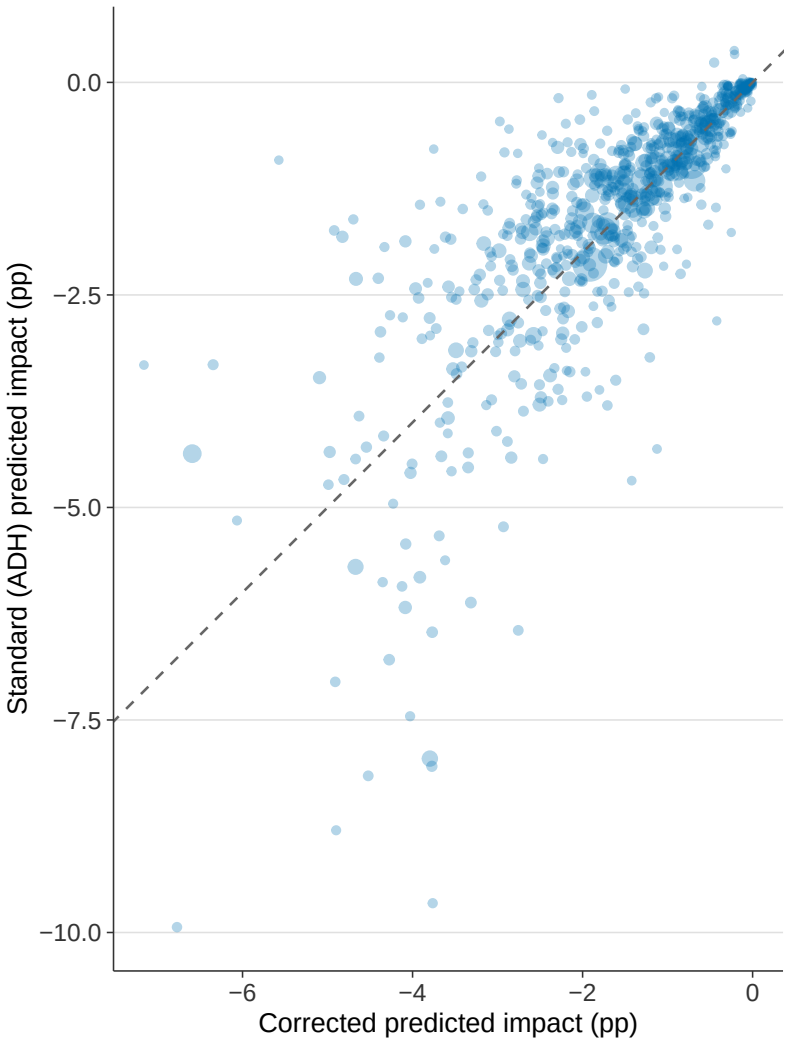


Figure 5. CZ-Level Predicted Impacts: Standard vs. Network-Corrected

Notes: Each point is a commuting zone. Horizontal axis: corrected predicted impact $\hat{\gamma}_{NSS} \cdot \tilde{B}_r$; vertical axis: standard predicted impact $\hat{\gamma}_{SS} \cdot B_r$. Dashed 45-degree line indicates identical predictions. Point size proportional to CZ population. $N = 722$ CZs, Period 2 (2000–2007). Two commuting zones with standard predicted impact below -10 pp are omitted from the plot for legibility.

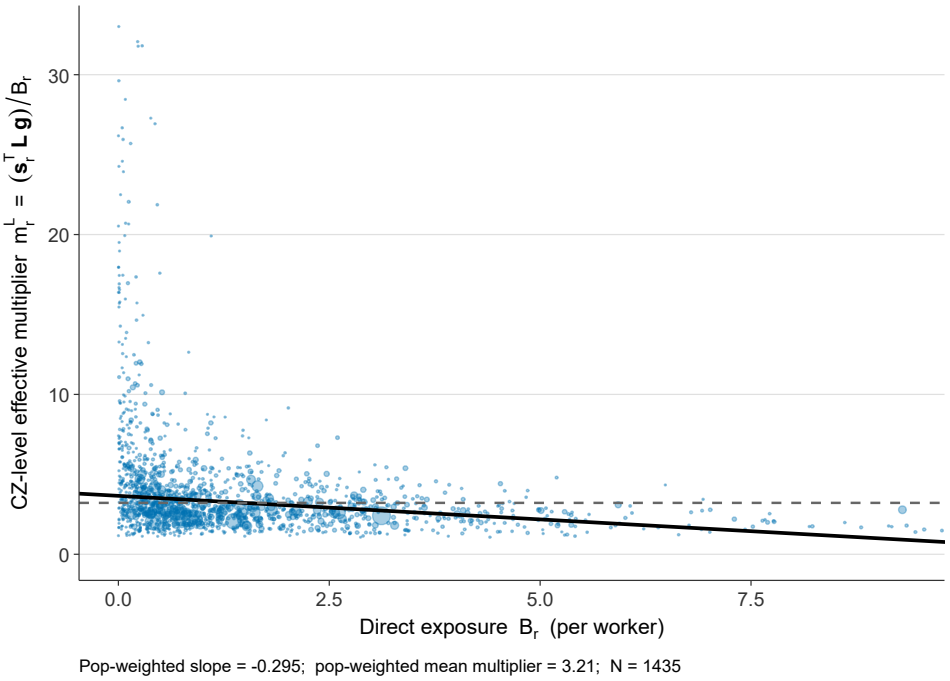
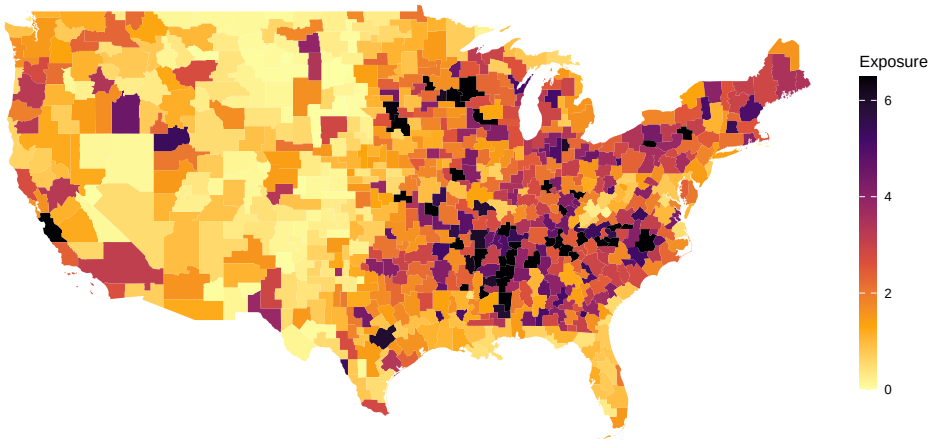
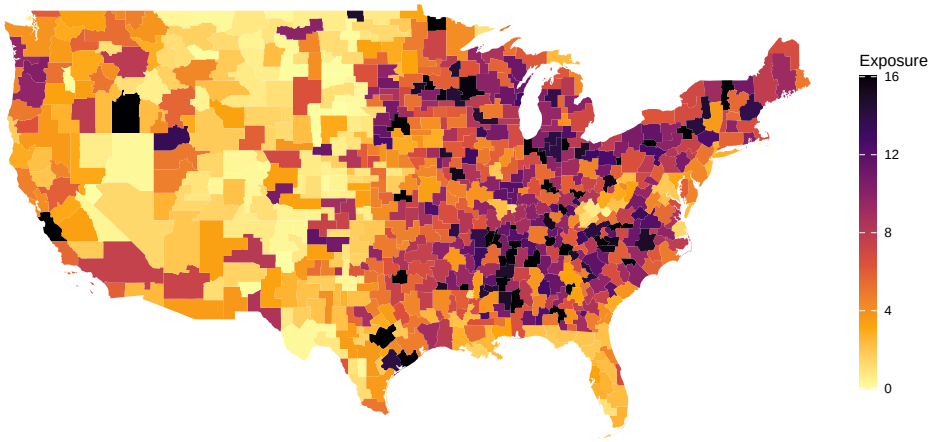


Figure 6. CZ-Level Leontief Multiplier m_r^L vs. Direct Exposure B_r

Notes: Each point is a CZ-period observation on the positive-exposure support $B_r > 0$ ($N = 1{,}433$; Section 3.1). Horizontal axis: $B_r = \mathbf{s}_r'\mathbf{g}$; vertical axis: $m_r^L = \mathbf{s}_r'\mathbf{L}\mathbf{g}/B_r$. Point size proportional to population weight. Solid line: population-weighted regression of m_r^L on B_r; slope is $\mathrm{Cov}_{\mathrm{pop}}(m_r, B_r)/\mathrm{Var}_{\mathrm{pop}}(B_r)$. Dashed grey line: population-weighted unconditional mean. Axes are truncated at the 99th percentile of each variable, which trims a thin tail of high-multiplier CZs.



(a) Standard Import Exposure $B_r = \mathbf{s}_r' \mathbf{g}^{US}$



(b) Network-Adjusted Import Exposure $\tilde{B}_r = \mathbf{s}_r' \mathbf{L} \mathbf{g}^{US}$

Figure 7. Geographic Distribution of Import Exposure Across Commuting Zones (Period 2: 2000–2007)

Notes: Panel (a): standard Bartik import exposure. Panel (b): network-adjusted exposure using the BEA 1997 Leontief inverse. Color scale top-coded at the 95th percentile. Alaska and Hawaii excluded. Figure 8 maps the ratio of the two measures.

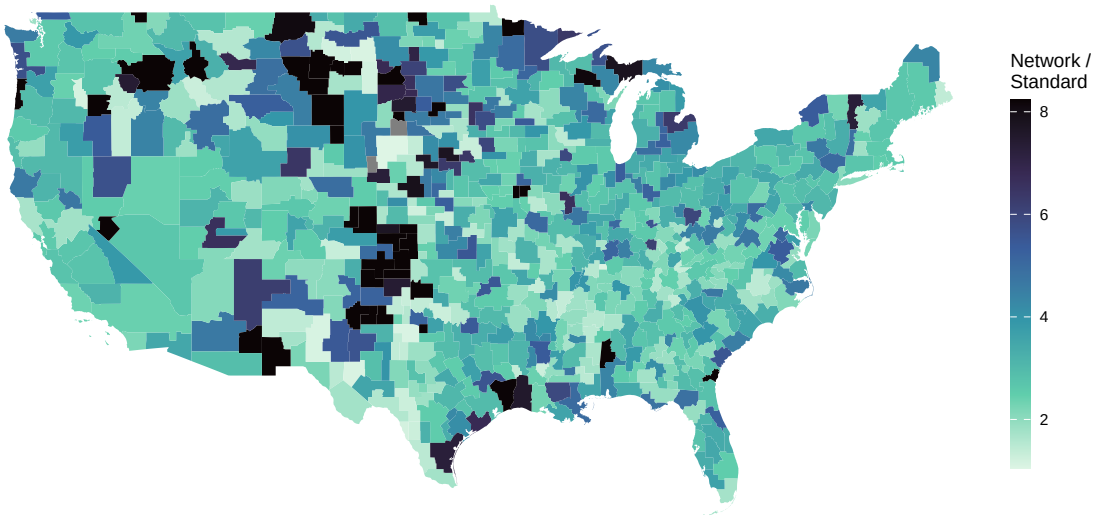


Figure 8. Network Amplification Ratio $\tilde{B}_r / B_r$ Across Commuting Zones (Period 2: 2000–2007)

Notes: Ratio of network-adjusted to standard import exposure. High-ratio commuting zones specialize in upstream process industries (petroleum refining, primary metals, basic chemicals, paper) whose exposure is carried by supply-chain linkages; low-ratio zones are the directly targeted apparel, footwear, and furniture belts with simple supply chains. Color scale top-coded at the 95th percentile. Alaska and Hawaii excluded.

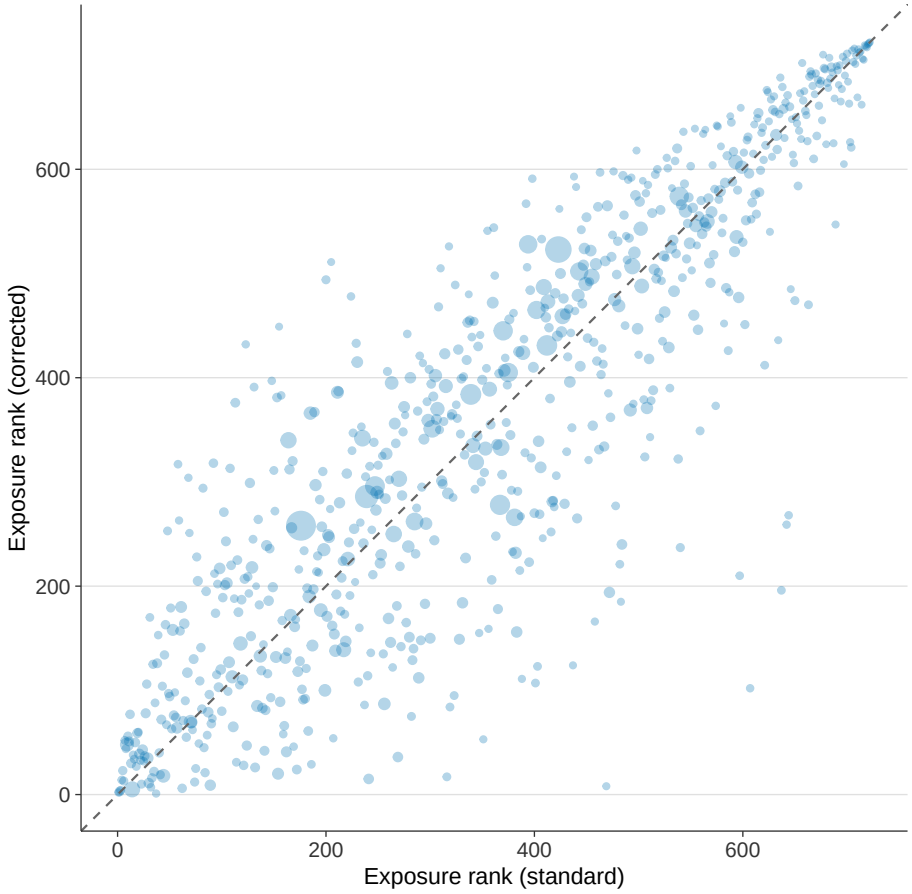


Figure 9. Geographic Misattribution: Exposure Rankings Under Standard vs. Network-Corrected Measure

Notes: Each point is a commuting zone in Period 2 (2000–2007). Horizontal axis: CZ's exposure rank under $B_r = \mathbf{s}_r'\mathbf{g}$; vertical axis: rank under $\tilde{B}_r = \mathbf{s}_r'\mathbf{L}\mathbf{g}$. Rank 1 = most exposed. Dashed 45-degree line marks identical rankings. Point size proportional to CZ population. Spearman rank correlation = 0.87.

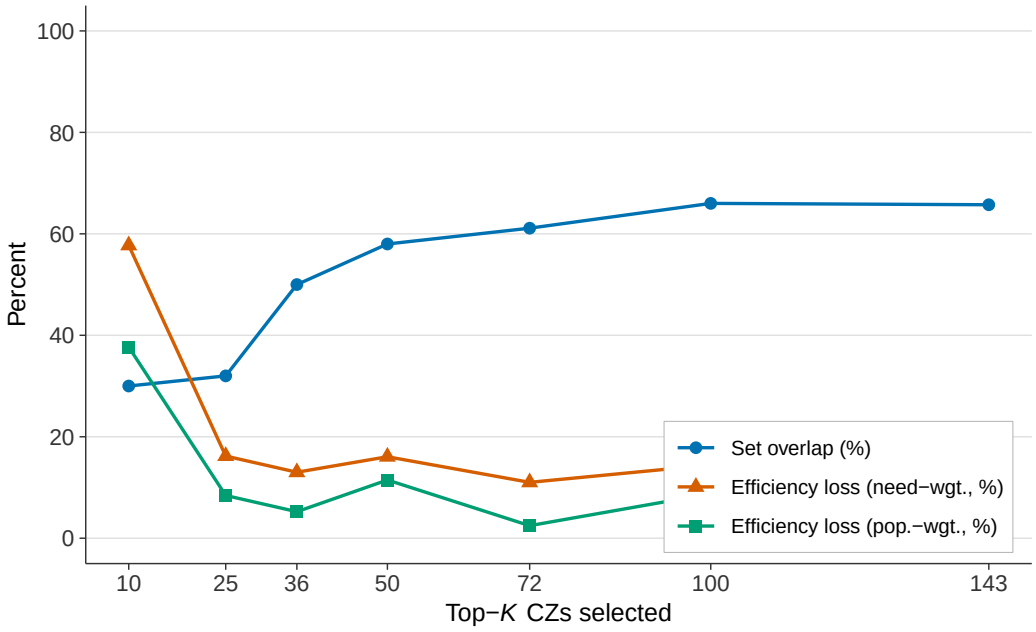


Figure 10. Welfare Loss from a Mis-Targeted TAA Budget Across Targeting Thresholds

Notes: Targeting accuracy of a fixed TAA budget allocated equally across the top-K most-exposed commuting zones under the standard vs. network-corrected Bartik ranking (Period 2, 2000–2007, $N = 714$ valid CZs). Set-overlap curve: $|\text{top-}K^{\text{std}} \cap \text{top-}K^{\text{corr}}|/K$. Efficiency-loss curves: $1 - W^{\text{std}}/W^{\text{corr}}$ under a need-weighted welfare function $n_r = \text{timepwt48}_r \cdot \max(\tilde{B}_r, 0)$ and a population-weighted welfare function $n_r = \text{timepwt48}_r$. The three series are distinguished by color and marker (see legend). Welfare-loss ratios are invariant to the absolute budget.

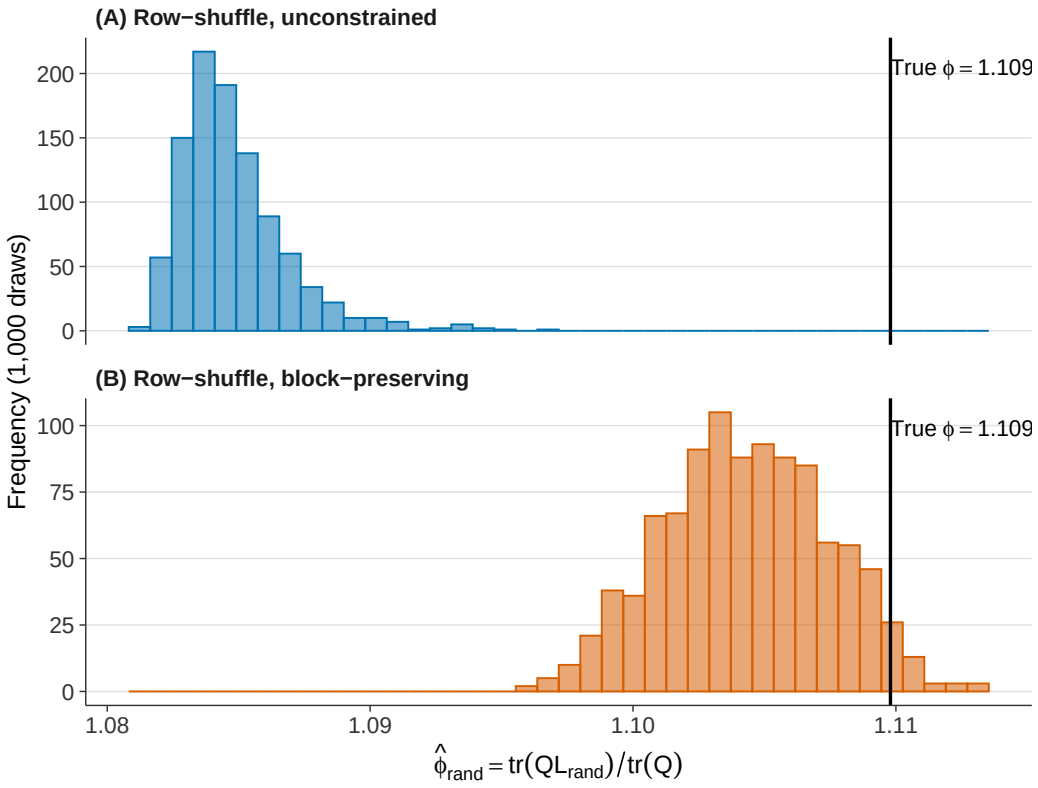


Figure 11. Placebo distribution of $\hat{\phi}$ under row-shuffled random propagation matrices

Notes: Distribution of $\hat{\phi}_{\text{rand}} = \text{tr}(\mathbf{QL}_{\text{rand}})/\text{tr}(\mathbf{Q})$ across 1,000 random draws of $\mathbf{L}_{\text{rand}}$ under two placebo schemes. Top panel: Scheme A (unconstrained row-shuffle preserving row sums and row-level sparsity). Bottom panel: Scheme B (block-preserving row-shuffle constrained to remain within 21 three-digit NAICS column blocks). Vertical black line marks the empirical value $\hat{\phi}$ from the true BEA 1997 detailed manufacturing Leontief inverse. Scheme A: $p < 0.001$; Scheme B: 97th percentile ($p = 0.034$).

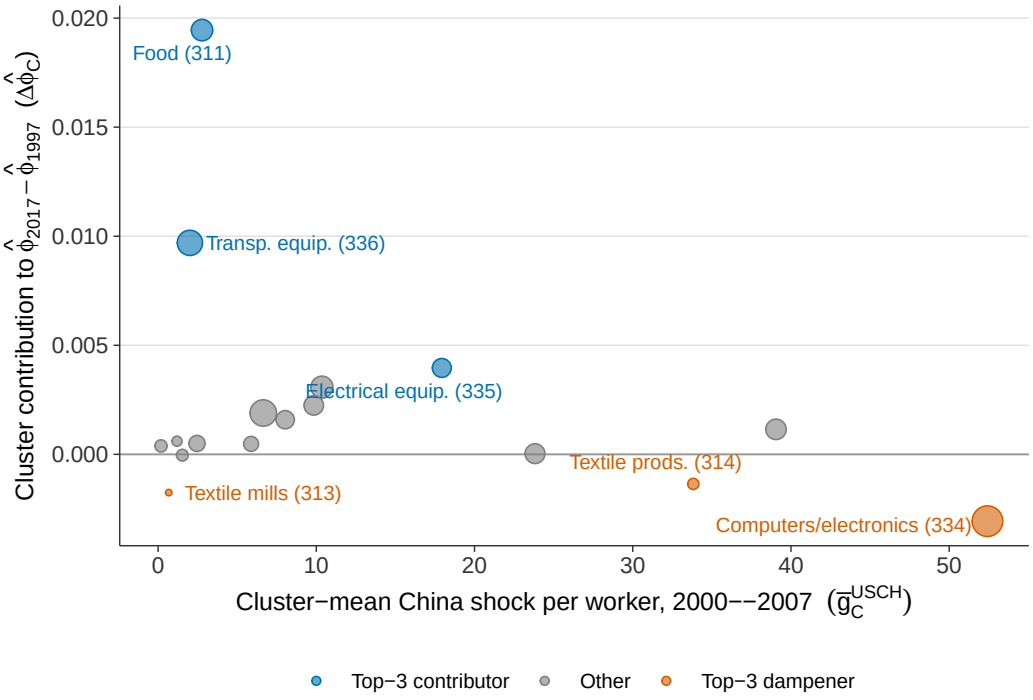


Figure 12. Cluster-Level Contributions to the 1997–2017 Drift in $\hat{\phi}^{mfg}$ vs. Cluster-Mean China Shock

Notes: Per-cluster contribution to the 1997–2017 drift in $\hat{\phi}^{mfg}$ ($\Delta\hat{\phi}_C$, vertical axis) vs. size-weighted mean 2000–2007 Chinese import shock per worker for each NAICS-3 cluster ($\bar{g}_C^{USCH}$, horizontal axis). 21 manufacturing clusters aggregated from the 142 BEA detail manufacturing codes common to both vintage classifications. Point size proportional to 1999 national employment.

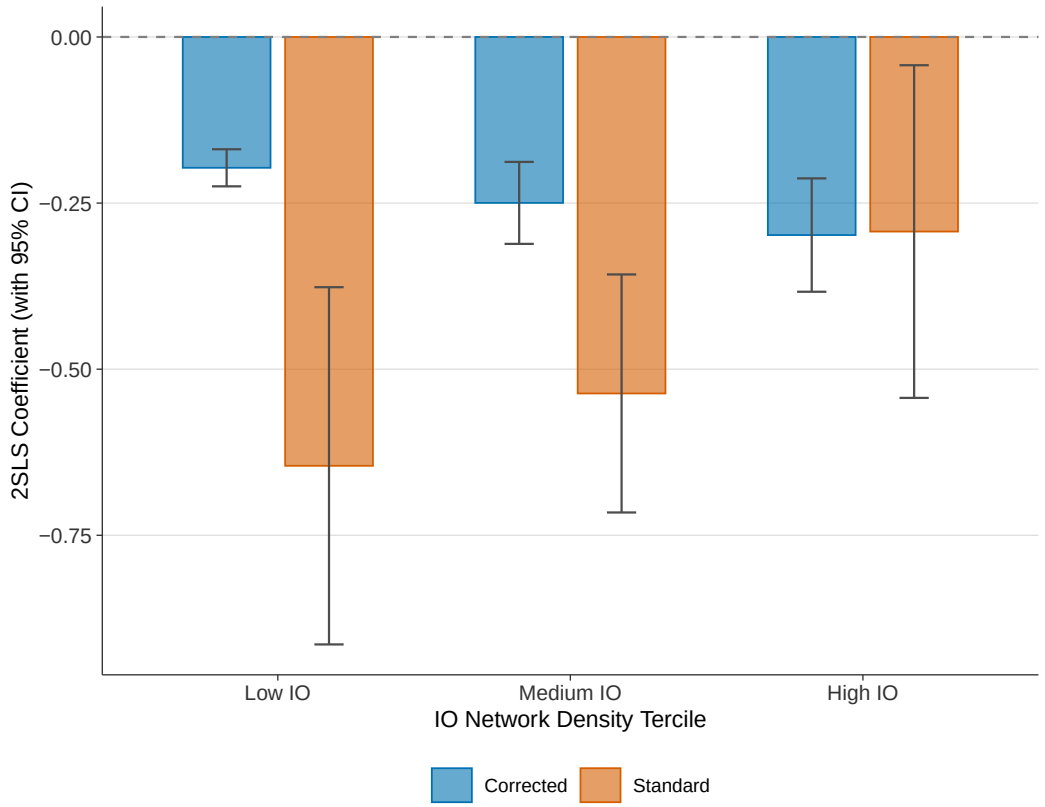


Figure 13. Heterogeneity by IO Network Density

Notes: Standard and corrected 2SLS coefficients on manufacturing employment by IO density tercile, distinguished by fill (see legend). Error bars are ±1 state-clustered standard error.

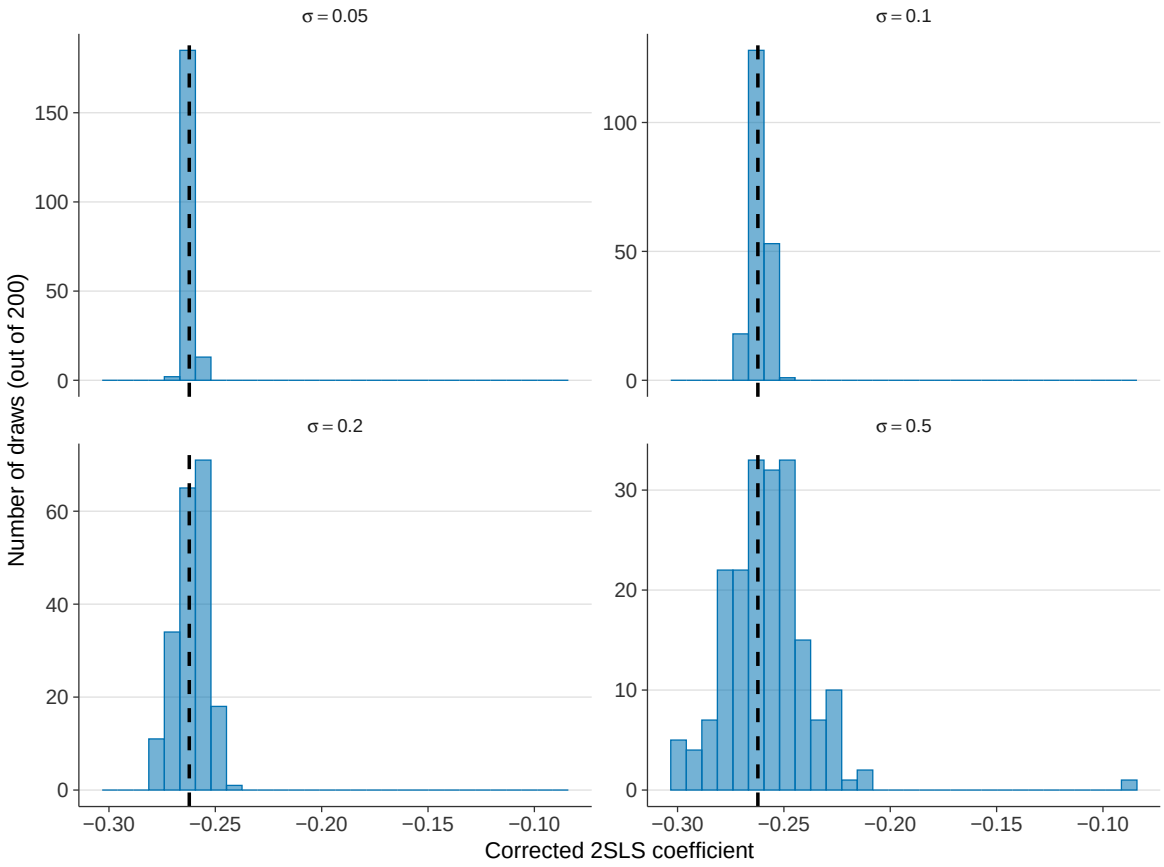


Figure 14. Sensitivity of the Corrected Coefficient to Noise in **A**

Notes: Distribution of the corrected single-regressor 2SLS coefficient over 200 draws per panel of multiplicative noise on the BEA 1997 detail **A** matrix: $A_{kl}^{\text{noisy}} = A_{kl}(1 + \varepsilon_{kl})$, $\varepsilon_{kl} \sim \mathcal{N}(0, \sigma^2)$, truncated below at zero. Each draw recomputes $\mathbf{L}^{\text{noisy}} = (\mathbf{I} - \mathbf{A}^{\text{noisy}})^{-1}$, the corrected $\tilde{B}_r$ on both sides of the IV, and the full pooled 2SLS specification. Black dashed line: noiseless baseline coefficient -0.262. The four panels span plausible-vintage ($\sigma = 0.05$) through implausibly-large ($\sigma = 0.50$) regional IO heterogeneity. Across all four levels the simulated mean stays within -1.68% of the noiseless baseline. 0.75% of $\sigma = 0.50$ draws produced a singular $\mathbf{I} - \mathbf{A}^{\text{noisy}}$ and were dropped.

A Proofs and Supporting Derivations

A.1 Main-Text Proofs

Proof of Proposition 1. The standard Bartik OLS estimator from $y_r = \alpha + \gamma B_r + u_r$ (after demeaning) is

$$\hat{\gamma}_{SS} \;=\; \frac{\sum_r B_r\, y_r}{\sum_r B_r^2} \;=\; \frac{\mathbf{g}'\mathbf{S}'\mathbf{y}}{\mathbf{g}'\mathbf{S}'\mathbf{S}\mathbf{g}}.$$

Substituting the two-channel DGP (2), $\mathbf{y} = \beta_D\,\mathbf{S}\mathbf{g} + \beta_N\,\mathbf{S}(\mathbf{P}-\mathbf{I})\mathbf{g} + \boldsymbol{\varepsilon}$, and using $\mathbf{Q} = \mathbf{S}'\mathbf{S}$:

$$\hat{\gamma}_{SS} \;=\; \beta_D \;+\; \beta_N \cdot \frac{\mathbf{g}'\mathbf{Q}(\mathbf{P}-\mathbf{I})\mathbf{g}}{\mathbf{g}'\mathbf{Q}\mathbf{g}} \;+\; \frac{\mathbf{g}'\mathbf{S}'\boldsymbol{\varepsilon}}{\mathbf{g}'\mathbf{Q}\mathbf{g}}. \tag{36}$$

Error term. Under Assumption 1, $\mathbf{g}'\mathbf{S}'\boldsymbol{\varepsilon} = \sum_k g_k\big(\sum_r s_{rk}\varepsilon_r\big)$, each summand with conditional mean zero. Summands are independent across k by Assumption 3; the numerator has variance $\sigma_\varepsilon^2\sigma_g^2\,\mathrm{tr}(\mathbf{Q})$, hence is $O_p(\mathrm{tr}(\mathbf{Q})^{1/2})$, while the denominator $\mathbf{g}'\mathbf{Q}\mathbf{g}$ concentrates at $\sigma_g^2\,\mathrm{tr}(\mathbf{Q}) = \Theta(\mathrm{tr}(\mathbf{Q}))$ by Assumption 2, so the ratio is $O_p(\mathrm{tr}(\mathbf{Q})^{-1/2}) \to 0$. Under the within-region normalization $\sum_k s_{rk} = 1$ with regional concentration bounded away from zero (Assumption 2), $\mathrm{tr}(\mathbf{Q}) = \Theta(R)$, so this is the joint $(R,K) \to \infty$ limit driven by the number of regions.

Trace-ratio limit. Under Assumption 3, only diagonal terms survive in the quadratic forms:

$$\mathbb{E}[\mathbf{g}'\mathbf{Q}\mathbf{g}] = \sigma^2\,\mathrm{tr}(\mathbf{Q}), \qquad \mathbb{E}[\mathbf{g}'\mathbf{Q}(\mathbf{P}-\mathbf{I})\mathbf{g}] = \sigma^2\big(\mathrm{tr}(\mathbf{Q}\mathbf{P}) - \mathrm{tr}(\mathbf{Q})\big).$$

Under Assumption 2, the LLN applies and Slutsky's theorem gives $\mathbf{g}'\mathbf{Q}(\mathbf{P}-\mathbf{I})\mathbf{g}/\mathbf{g}'\mathbf{Q}\mathbf{g} \xrightarrow{p} \phi - 1$, proving (7).

Non-negativity of $\phi - 1$. Using symmetry of $\mathbf{Q}$, decompose $\mathrm{tr}(\mathbf{Q}\mathbf{P}) = \sum_k Q_{k,k}P_{k,k} + \sum_{k\neq l} Q_{k,l}P_{k,l}$. Since $Q_{k,l} \geq 0$, $P_{k,l} \geq 0$, and $P_{k,k} \geq 1$, $\mathrm{tr}(\mathbf{Q}\mathbf{P}) \geq \mathrm{tr}(\mathbf{Q})$, so $\phi \geq 1$ with equality iff (i) $P_{kk} = 1$ for every k with $Q_{kk} > 0$ and (ii) $P_{kl} = 0$ for every off-diagonal pair (k,l) with $Q_{kl} > 0$. Propagation outside the co-location support of $\mathbf{Q}$ contributes zero to $\mathrm{tr}(\mathbf{Q}\mathbf{P}) - \mathrm{tr}(\mathbf{Q})$ and so does not bias the slope; however such propagation still enters the corrected exposure $\tilde{B}_r$ and so generically produces nonzero per-region prediction errors Δ_r^{std} through $m_r \neq 1$, even when $\phi = 1$. □

Proof of Lemma 1. Direct substitution yields $\sum_r B_r N_r \;=\; \mathbf{g}'\mathbf{S}'\mathbf{S}(\mathbf{P}-\mathbf{I})\mathbf{g} \;=\; \mathbf{g}'\mathbf{Q}(\mathbf{P}-\mathbf{I})\mathbf{g}$ and $\sum_r B_r^2 \;=\; \mathbf{g}'\mathbf{Q}\mathbf{g}$, establishing (9). Under Assumptions 3–2, the ratio converges to $\mathrm{tr}(\mathbf{Q}(\mathbf{P}-\mathbf{I}))/\mathrm{tr}(\mathbf{Q}) = \phi - 1$. The finite-sample OVB identity follows from the standard two-regressor OLS decomposition. □

Proof of Corollary 1. Substituting the DGP $y_r = \beta\,\mathbf{s}_r'\mathbf{P}\mathbf{g}^{US} + \varepsilon_r$ into the 2SLS formula with instrument $Z_r = \mathbf{s}_r'\mathbf{g}^{OC}$:

$$\hat{\beta}_{2SLS} \;=\; \frac{\sum_r Z_r y_r}{\sum_r Z_r X_r} \;=\; \beta \cdot \frac{(\mathbf{g}^{OC})'\mathbf{S}'\mathbf{S}\mathbf{P}\mathbf{g}^{US}}{(\mathbf{g}^{OC})'\mathbf{S}'\mathbf{S}\mathbf{g}^{US}} + \frac{(\mathbf{g}^{OC})'\mathbf{S}'\boldsymbol{\varepsilon}}{(\mathbf{g}^{OC})'\mathbf{S}'\mathbf{S}\mathbf{g}^{US}}.$$

The second term is $o_p(1)$ under exogeneity of $\mathbf{g}^{OC}$ (parallel to the Proposition 1 argument).

The first term equals $\beta \cdot \phi^{IV}(\mathbf{g}^{OC}, \mathbf{g}^{US})$. The population probability limit (12) follows from the quadratic-form identity $\mathbb{E}[\mathbf{g}'\mathbf{M}\mathbf{g}] = \mathrm{tr}(\mathbf{M}\boldsymbol{\Sigma}) = \sum_{k,l} M_{kl}\sigma_{kl}$ applied to the cross-covariance kernel. □

A.2 Exogeneity of Transformed Exposure

The corrected estimator and its two- and three-regressor variants use exposure measures of the form $\mathbf{s}_r'\mathbf{M}\mathbf{g}^{OC}$ with $\mathbf{M} \in \{\mathbf{L},\ \mathbf{L}-\mathbf{I},\ \mathbf{G},\ \mathbf{G}-\mathbf{I},\ \mathbf{L}-\mathbf{L}_{\mathrm{diag}}\}$. Assumption 1 is stated element-wise on the primitive shock $\mathbf{g}^{OC}$; the following lemma propagates it to any deterministic linear transformation of the shock vector, so transformed exposure measures inherit exogeneity without re-assumption.

Lemma 2 (Exogeneity Propagation Under Deterministic Transformation)**.** *Let $\mathbf{M} \in \mathbb{R}^{K\times K}$ be deterministic (a function of $\mathbf{P}$ and/or $\mathbf{S}$ only), and let $\tilde{g}_l \equiv (\mathbf{M}\mathbf{g}^{OC})_l = \sum_k M_{lk} g_k^{OC}$. Under Assumption 1,*

$$\mathbb{E}[\tilde{g}_l \mid \mathbf{S}, \boldsymbol{\varepsilon}] \;=\; \sum_k M_{lk}\, \mathbb{E}[g_k^{OC} \mid \mathbf{S}, \boldsymbol{\varepsilon}] \;=\; 0 \qquad \textit{for every } l.$$

Consequently, for any deterministic $R \times K$ weighting matrix $\mathbf{W}$ (including $\mathbf{S}$ itself and $\mathbf{S}\mathbf{M}$),

$$\mathbb{E}\Big[\sum_l W_{rl}\, \tilde{g}_l \cdot \varepsilon_r \,\Big|\, \mathbf{S}\Big] \;=\; \sum_l W_{rl}\, \mathbb{E}[\tilde{g}_l\, \varepsilon_r \mid \mathbf{S}] \;=\; 0 \qquad \textit{for every } r.$$

Instantiating with $\mathbf{M} = \mathbf{P} - \mathbf{I}$, $\mathbf{M} = \mathbf{L}$, or $\mathbf{M} = \mathbf{L} - \mathbf{L}_{\mathrm{diag}}$ and $W_{rl} = s_{rl}$ shows that $\mathbf{s}_r'\mathbf{M}\mathbf{g}^{OC}$ is a valid instrument whenever $\mathbf{g}^{OC}$ is valid for $\mathbf{s}_r'\mathbf{g}^{OC}$.

Proof. Linearity of conditional expectation. $\mathbb{E}[\tilde{g}_l \mid \mathbf{S}, \boldsymbol{\varepsilon}] = \sum_k M_{lk}\, \mathbb{E}[g_k^{OC} \mid \mathbf{S}, \boldsymbol{\varepsilon}]$ because M_{lk} is $(\mathbf{S}, \boldsymbol{\varepsilon})$-measurable (deterministic, in fact). Each summand is zero by Assumption 1. The unconditional orthogonality condition then follows by the law of iterated expectations. □

Remark 1 (Amplification of exclusion-restriction violations). If Assumption 1 fails by $\mathbb{E}[g_k^{OC} \mid \mathbf{S}, \boldsymbol{\varepsilon}] = \eta_k$ for some deviation vector η_k, the transformed deviation is $(\mathbf{M}\eta)_l = \sum_k M_{lk}\eta_k$, and the induced bias in the 2SLS estimator is bounded above by $\|\mathbf{M}\|_\infty \max_k |\eta_k|$ up to denominator-scaling. For the Leontief inverse at the BEA detailed level, $\|\mathbf{L}\|_\infty = \max_l \sum_k L_{lk}$ is the largest row sum of $\mathbf{L}$; in the 1997 BEA manufacturing matrix it is 12.9, attained at a highly central supplier row, against an average row sum of 1.62. Transformed exposure therefore does *not* introduce unbounded amplification of any original exclusion-restriction violation, but the worst-case rescaling of a violation concentrated on a central supplier is an order of magnitude, while the typical rescaling is the modest average multiplier; the bound is a genuine sensitivity multiplier and should be reported alongside the main specification.

Proof of Proposition 2. Under the single-index DGP (3),

$$\hat{\gamma}_{NSS} \;=\; \frac{\sum_r \tilde{B}_r y_r}{\sum_r \tilde{B}_r^2} \;=\; \beta + \frac{\mathbf{g}'\mathbf{P}'\mathbf{S}'\boldsymbol{\varepsilon}}{\mathbf{g}'\mathbf{P}'\mathbf{Q}\mathbf{P}\mathbf{g}}.$$

The numerator is $\sum_k g_k \sum_l P_{lk} \sum_r s_{rl}\varepsilon_r$, each summand with conditional mean zero under Assumption 1. Independence across k (Assumption 3) gives numerator variance of order $\mathrm{tr}(\mathbf{P}'\mathbf{Q}\mathbf{P})$, so the numerator is $O_p(\mathrm{tr}(\mathbf{P}'\mathbf{Q}\mathbf{P})^{1/2})$ while the denominator $\mathbf{g}'\mathbf{P}'\mathbf{Q}\mathbf{P}\mathbf{g}$ concentrates at $\Theta(\mathrm{tr}(\mathbf{P}'\mathbf{Q}\mathbf{P}))$ under Assumption 2. The ratio is $O_p(\mathrm{tr}(\mathbf{P}'\mathbf{Q}\mathbf{P})^{-1/2}) \to 0$ as the design grows, and $\mathrm{plim}\hat{\gamma}_{NSS} = \beta$. □

Proof of Proposition 3. Substitute (2) into $\hat{\gamma}_{NSS}$ and use $\tilde{B}_r = B_r + N_r$:

$$\hat{\gamma}_{NSS} \;=\; (\beta_D - \beta_N)\,\frac{\sum_r \tilde{B}_r B_r}{\sum_r \tilde{B}_r^2} + \beta_N + \frac{\sum_r \tilde{B}_r \varepsilon_r}{\sum_r \tilde{B}_r^2}.$$

The error term vanishes by the Proposition 2 argument, and the first ratio converges to

$$\frac{\mathbf{g}'\mathbf{P}'\mathbf{Q}\mathbf{g}}{\mathbf{g}'\mathbf{P}'\mathbf{Q}\mathbf{P}\mathbf{g}} \;\xrightarrow{p}\; \lambda, \qquad \lambda \;=\; \frac{\mathrm{tr}(\mathbf{Q}\mathbf{P})}{\mathrm{tr}(\mathbf{P}'\mathbf{Q}\mathbf{P})},$$

so $\mathrm{plim}\,\hat{\gamma}_{NSS} = \lambda\beta_D + (1-\lambda)\beta_N$. The upper bound $\lambda \le 1/\phi$ follows from Cauchy–Schwarz in the Frobenius inner product, $[\mathrm{tr}(\mathbf{Q}\mathbf{P})]^2 \le \mathrm{tr}(\mathbf{Q})\,\mathrm{tr}(\mathbf{P}'\mathbf{Q}\mathbf{P})$:

$$\lambda \;=\; \frac{\mathrm{tr}(\mathbf{Q}\mathbf{P})}{\mathrm{tr}(\mathbf{P}'\mathbf{Q}\mathbf{P})} \;\le\; \frac{\mathrm{tr}(\mathbf{Q})}{\mathrm{tr}(\mathbf{Q}\mathbf{P})} \;=\; \frac{1}{\phi}.$$

Non-negativity follows from non-negativity of both the numerator and denominator. □

Proof of Proposition 4. The shift-share estimator applied to $(\tilde{\mathbf{S}}, \mathbf{g})$ with $\tilde{s}_{rl} = \sum_k s_{rk} P_{kl}$ is consistent by Borusyak et al. (2022, Proposition 3) under Assumption 2 applied to $\tilde{\mathbf{S}}$; asymptotic normality and the exposure-robust variance estimator are those of Adão et al. (2019) applied to $(\tilde{\mathbf{S}}, \mathbf{g})$, which yields (20). Using the original shares s_{rk} indexes residuals by the wrong industry: since g_l affects corrected exposure through every industry k with $P_{kl} > 0$, the naive formula misses the cross-k residual correlation that the network structure introduces. The correct aggregation is by primitive shock index l. □

Proof of Proposition 5. Substitute $\mathrm{plim}\hat{\gamma}_{SS} \to \beta_D + \beta_N(\phi - 1)$ from Proposition 1 into $\hat{y}_r^{\mathrm{std}} = \hat{\gamma}_{SS} B_r$, and use $N_r = (m_r - 1)B_r$ in $\hat{y}_r^{\mathrm{corr}}$. Dividing numerator and denominator by $\beta_D B_r$ gives the result. The single-index collapse follows from $(1 + (\phi - 1))/(1 + (m_r - 1)) = \phi/m_r$. □

Proof of Corollary 3. Multiply numerator and denominator of (22) by B_r and sum over regions weighted by w_r:

$$T^{\mathrm{std}} \;=\; \beta_D \bar{B}_w \cdot [1 + \theta(\phi - 1)], \qquad T^{\mathrm{corr}} \;=\; \beta_D \bar{B}_w \cdot [1 + \theta(\psi^{\mathrm{agg}} - 1)],$$

and dividing gives (24). The exact-cancellation condition $\phi = \psi^{\mathrm{agg}}$ follows immediately. □

Proof of Proposition 6. Expand the numerator of ψ^{agg} as $\sum_r w_r\, \mathbf{s}_r'\mathbf{P}\mathbf{g} = \sum_r w_r\, m_r\, B_r$ (using $m_r = \tilde{B}_r / B_r$). Then $\psi^{\mathrm{agg}} = \sum_r w_r m_r B_r / \bar{B}_w$. Using the covariance identity $\sum_r w_r m_r B_r = \bar{m}^{\mathrm{pop}} \bar{B}_w + \mathrm{Cov}_{\mathrm{pop}}(m_r, B_r)$ with $\bar{m}^{\mathrm{pop}} = \sum_r w_r m_r$ gives (26). The identity is sample-level and

holds exactly. Multiplying by $\beta_N \bar{B}_w$ and using $\Delta_r^{\text{std}} = \beta_N B_r(\phi - m_r)$ from Corollary 2 gives the equivalent aggregate-error form $\sum_r w_r \Delta_r^{\text{std}} = \beta_N \bar{B}_w(\phi - \psi^{\text{agg}})$ reported in equation (25) of the main text. □

Proof of Proposition 7. Throughout, sums run over the support $B_r > 0$ on which $m_r = \tilde{B}_r / B_r$ is defined, and $p_r = w_r B_r / \sum_{r'} w_{r'} B_{r'}$.

(i) Since $\tilde{B}_r = m_r B_r$, the population-weighted aggregate amplification is a p-weighted mean of the multiplier,

$$\psi^{\text{agg}} \;=\; \frac{\sum_r w_r \tilde{B}_r}{\sum_r w_r B_r} \;=\; \sum_r p_r \, m_r \;=\; \mathbb{E}_p[m_r].$$

The WLS-faithful IV factor is a q-weighted mean of the same multiplier,

$$\phi^{IV} \;=\; \frac{\sum_r w_r Z_r \tilde{B}_r}{\sum_r w_r Z_r B_r} \;=\; \sum_r q_r \, m_r \;=\; \mathbb{E}_q[m_r], \qquad q_r \;\propto\; w_r Z_r B_r \;\propto\; p_r Z_r,$$

and normalizing gives $q_r = p_r Z_r / \mathbb{E}_p[Z_r]$. Hence

$$\phi^{IV} \;=\; \frac{\mathbb{E}_p[m_r Z_r]}{\mathbb{E}_p[Z_r]} \;=\; \frac{\mathbb{E}_p[m_r]\,\mathbb{E}_p[Z_r] + \text{Cov}_p(m_r, Z_r)}{\mathbb{E}_p[Z_r]} \;=\; \psi^{\text{agg}} + \frac{\text{Cov}_p(m_r, Z_r)}{\mathbb{E}_p[Z_r]},$$

which is (28); cancellation $\phi^{IV} = \psi^{\text{agg}}$ holds iff $\text{Cov}_p(m_r, Z_r) = 0$. Using $N_r = (m_r - 1)B_r$ and writing the condition in w-weighted moments of (Z_r, N_r, B_r), $\text{Cov}_p(m_r, Z_r) = 0$ is equivalent to $\text{Cov}_w(Z_r, N_r) = (\psi^{\text{agg}} - 1)\text{Cov}_w(Z_r, B_r)$, the slope-balance form. The equivalence to $\sum_r w_r \Delta_r^{\text{std}} = 0$ is equation (25).

(ii) If $\mathbf{S}(\mathbf{P} - c\mathbf{I})\mathbf{g} = \mathbf{0}$ then $\tilde{B}_r = \mathbf{s}_r'\mathbf{P}\mathbf{g} = c\,\mathbf{s}_r'\mathbf{g} = cB_r$ for every r, so $m_r \equiv c$. A weighted mean of a constant equals that constant, so $\psi^{\text{agg}} = \mathbb{E}_p[m_r] = c$ and $\phi^{IV} = \mathbb{E}_q[m_r] = c$. The correlated-shocks factor ϕ^{corr} also equals c: the condition is exactly $\mathbf{S}\mathbf{P}\mathbf{g} = c\,\mathbf{S}\mathbf{g}$, so $\mathbf{g}'\mathbf{Q}\mathbf{P}\mathbf{g} = (\mathbf{S}\mathbf{g})'(\mathbf{S}\mathbf{P}\mathbf{g}) = c\,(\mathbf{S}\mathbf{g})'(\mathbf{S}\mathbf{g}) = c\,\mathbf{g}'\mathbf{Q}\mathbf{g}$, giving $\phi^{\text{corr}} = \mathbf{g}'\mathbf{Q}\mathbf{P}\mathbf{g}/\mathbf{g}'\mathbf{Q}\mathbf{g} = c$ (and identically $\phi^{IV} = c$ with $\mathbf{g}^{OC}$ in the left factor). The shock-free i.i.d. diagnostic $\phi^{\text{iid}} = \text{tr}(\mathbf{Q}\mathbf{P})/\text{tr}(\mathbf{Q})$ does *not* depend on $\mathbf{g}$, so it need not equal c under this condition; it does so only if additionally $\mathbf{P} = c\mathbf{I}$ on the support of $\mathbf{Q}$. (Counterexample: $\mathbf{S} = \mathbf{Q} = \mathbf{I}$, $\mathbf{P} = \text{diag}(2, 3)$, $\mathbf{g} = (1, 0)'$, $c = 2$ satisfies the condition, with $\mathbf{g}$ even a right-eigenvector of $\mathbf{P}$, and $\phi^{\text{corr}} = \phi^{IV} = \psi^{\text{agg}} = 2$, yet $\phi^{\text{iid}} = \text{tr}(\mathbf{P})/2 = 2.5 \neq c$.) The per-region error is $\Delta_r^{\text{std}} = \beta_N B_r(\phi - m_r) = \beta_N B_r(c - c) = 0$ for every r. A right-eigenvector $\mathbf{P}\mathbf{g} = c\mathbf{g}$ gives $(\mathbf{P} - c\mathbf{I})\mathbf{g} = \mathbf{0}$, hence $\mathbf{S}(\mathbf{P} - c\mathbf{I})\mathbf{g} = \mathbf{0}$ a fortiori.

(iii) By the Cauchy–Schwarz inequality for the p-weighted inner product,

$$|\text{Cov}_p(m_r, Z_r)| \;\leq\; \text{sd}_p(m_r)\,\text{sd}_p(Z_r).$$

Dividing (28) by ψ^{agg} and using $\text{CV}_p(Z_r) = \text{sd}_p(Z_r)/\mathbb{E}_p[Z_r]$,

$$\left|\frac{\phi^{IV}}{\psi^{\text{agg}}} - 1\right| \;=\; \frac{|\text{Cov}_p(m_r, Z_r)|}{\mathbb{E}_p[Z_r]\,\psi^{\text{agg}}} \;\leq\; \frac{\text{sd}_p(m_r)\,\text{sd}_p(Z_r)}{\mathbb{E}_p[Z_r]\,\psi^{\text{agg}}} \;=\; \frac{\text{CV}_p(Z_r)\,\text{sd}_p(m_r)}{\psi^{\text{agg}}},$$

which is (29); as $\mathrm{sd}_p(m_r) \to 0$ the bound vanishes, recovering (ii) in the limit. □

Proof of Corollary 2. For (15), substitute Proposition 1's $\mathrm{plim}\hat{\gamma}_{SS} = \beta_D + \beta_N(\phi - 1)$ into $\hat{T}_r^{\mathrm{std}} = \mathrm{plim}(\hat{\gamma}_{SS})\, B_r$ and subtract $T_r = \beta_D B_r + \beta_N N_r$:

$$\Delta_r^{\mathrm{std}} \;=\; \beta_D B_r + \beta_N(\phi-1)B_r - \beta_D B_r - \beta_N N_r \;=\; \beta_N\big[(\phi-1)B_r - N_r\big] \;=\; \beta_N B_r(\phi - m_r),$$

using $N_r = (m_r - 1)B_r$. For (16), the corrected single-regressor prediction is $\hat{T}_r^{\mathrm{corr},1} = \mathrm{plim}(\hat{\gamma}_{NSS})\,\tilde{B}_r = [\lambda\beta_D + (1-\lambda)\beta_N](B_r + N_r)$ by Proposition 3. Subtracting T_r and collecting terms by causal coefficient,

$$\begin{aligned}\Delta_r^{\mathrm{corr}} &= [\lambda\beta_D + (1-\lambda)\beta_N](B_r + N_r) - \beta_D B_r - \beta_N N_r \\ &= \beta_D\big[\lambda N_r - (1-\lambda)B_r\big] \;-\; \beta_N\big[\lambda N_r - (1-\lambda)B_r\big] \\ &= (\beta_D - \beta_N)\big[\lambda N_r - (1-\lambda)B_r\big].\end{aligned}$$

For (17), under exogeneity (Assumption 1), rank, and Assumption 2, the two-regressor OLS/IV estimates of (β_D, β_N) are consistent, so $\hat{T}_r^{\mathrm{2reg}} - T_r = (\hat{\beta}_D - \beta_D)B_r + (\hat{\beta}_N - \beta_N)N_r = o_p(1)$ pointwise; uniformity follows from boundedness of (B_r, N_r) under bounded shares and shocks. □

Proposition 8 (Mean–Variance Decomposition of $\mathrm{MSE}(\Delta_r^{\mathrm{std}})$)**.** *Let $w_r > 0$ with $\sum_r w_r = 1$ denote the population weights used to define the aggregate amplification ψ^{agg}. Let $\bar{B}_w = \sum_r w_r B_r$ and $\bar{\tilde{B}}_w = \sum_r w_r \tilde{B}_r$, so that $\psi^{\mathrm{agg}} = \bar{\tilde{B}}_w / \bar{B}_w$. The population-weighted mean squared prediction error of the standard Bartik admits the exact decomposition*

$$\mathrm{MSE}_w\big(\Delta_r^{\mathrm{std}}\big) \;\equiv\; \sum_r w_r \big(\Delta_r^{\mathrm{std}}\big)^2 \;=\; \underbrace{\beta_N^2\, \bar{B}_w^2\, \big(\phi - \psi^{\mathrm{agg}}\big)^2}_{\textit{aggregate-level squared error}} \;+\; \underbrace{\beta_N^2\, \mathrm{Var}_w\big(\phi\, B_r - \tilde{B}_r\big)}_{\textit{cross-region (geographic) variance}} \quad . \tag{37}$$

The first term measures how far the population-weighted aggregate of the standard estimator's prediction departs from the true aggregate impact; the second term measures the cross-region dispersion of region-specific prediction errors around their population-weighted mean. The two components correspond, respectively, to a violation of aggregate cancellation ($\phi \neq \psi^{\mathrm{agg}}$) and to the geographic misattribution that the standard estimator generates conditional on the aggregate.

Proof. By Corollary 2, $\Delta_r^{\mathrm{std}} = \beta_N B_r(\phi - m_r) = \beta_N(\phi B_r - \tilde{B}_r)$, using $m_r B_r = \tilde{B}_r$. The standard mean-square identity for any real-valued sequence $\{X_r\}$ with probability weights w_r,

$$\sum_r w_r X_r^2 \;=\; \Big(\sum_r w_r X_r\Big)^2 + \mathrm{Var}_w(X_r),$$

applied with $X_r = \Delta_r^{\mathrm{std}}$ gives

$$\mathrm{MSE}_w\big(\Delta_r^{\mathrm{std}}\big) \;=\; \Big(\sum_r w_r \Delta_r^{\mathrm{std}}\Big)^2 + \mathrm{Var}_w\big(\Delta_r^{\mathrm{std}}\big).$$

By equation (25), $\sum_r w_r \Delta_r^{\mathrm{std}} = \beta_N \bar{B}_w(\phi - \psi^{\mathrm{agg}})$, so the squared aggregate-level term is

$\beta_N^2 \bar{B}_w^2 (\phi - \psi^{\text{agg}})^2$. The variance term factors as $\beta_N^2 \text{Var}_w(\phi B_r - \tilde{B}_r)$, which can be expanded as $\beta_N^2[\phi^2 \text{Var}_w(B_r) - 2\phi \text{Cov}_w(B_r, \tilde{B}_r) + \text{Var}_w(\tilde{B}_r)]$ if the individual second moments are useful for diagnostics. □

The decomposition is purely sample-level and does not depend on asymptotics. In the ADH application reported in Section 5.4, the aggregate-level squared error contributes less than 1% of the total MSE while the cross-region variance contributes 99.8%. The aggregate cancellation $\phi \approx \psi^{\text{agg}}$ documented in Section 5.2 is therefore precisely the statement that the first term in (37) is small relative to the second; the geographic misallocation reported in Sections 5.5–5.8 is the empirical content of the second term.

A.3 Additional Propositions

Corollary 4 (Bias Under Controls, Fixed Effects, and Weights)**.** *Suppose the researcher estimates $y_r = \alpha + \gamma B_r + \mathbf{X}_r'\boldsymbol{\delta} + u_r$ by weighted least squares with weights $w_r > 0$ and controls $\mathbf{X}_r$. Let $\mathbf{M}_X = \mathbf{I} - \mathbf{X}(\mathbf{X}'\mathbf{W}\mathbf{X})^{-1}\mathbf{X}'\mathbf{W}$ with $\mathbf{W} = \text{diag}(w_1, \ldots, w_R)$. Under Assumptions 1–2 with $\mathbb{E}[\mathbf{X}_r \varepsilon_r] = 0$,*

$$\text{plim}_{K\to\infty} \hat{\gamma}_{SS}^{\text{ctrl}} \;=\; \beta_D + \beta_N(\tilde{\phi} - 1), \qquad \tilde{\phi} = \frac{\text{tr}(\tilde{\mathbf{Q}}\mathbf{P})}{\text{tr}(\tilde{\mathbf{Q}})},$$

where $\tilde{\mathbf{Q}} = \mathbf{S}'\mathbf{W}\mathbf{M}_X\mathbf{S} = (\mathbf{M}_X\mathbf{S})'\mathbf{W}(\mathbf{M}_X\mathbf{S})$. The matrix $\tilde{\mathbf{Q}}$ is symmetric and positive semi-definite, but unlike the raw $\mathbf{Q}$ it can have negative off-diagonal entries; consequently the non-negativity, monotonicity, and equality-support results of Proposition 9 apply only to the unresidualized $\mathbf{Q}$ and do not extend to $\tilde{\phi}$ as a function of $\mathbf{P}$.

Proof. The WLS first-order conditions for $(\gamma, \boldsymbol{\delta})$ give $\hat{\gamma}_{SS}^{\text{ctrl}} = \mathbf{B}'\mathbf{W}\mathbf{M}_X\mathbf{y} / \mathbf{B}'\mathbf{W}\mathbf{M}_X\mathbf{B}$ after solving out $\boldsymbol{\delta}$. Substituting $\mathbf{B} = \mathbf{S}\mathbf{g}$ and $\mathbf{y} = \beta_D \mathbf{S}\mathbf{g} + \beta_N \mathbf{S}(\mathbf{P} - \mathbf{I})\mathbf{g} + \boldsymbol{\varepsilon}$, the ratio converges to $\beta_D + \beta_N \text{tr}(\tilde{\mathbf{Q}}(\mathbf{P} - \mathbf{I}))/\text{tr}(\tilde{\mathbf{Q}})$. □

Corollary 5 (Consistency Under Correlated Shocks)**.** *Replace Assumption 3 with: (i) $\boldsymbol{\Sigma} = \text{Cov}(\mathbf{g})$ satisfies the uniform row-sum bound $\max_k \sum_l |\sigma_{kl}| \le C_\Sigma < \infty$; (ii) $\{g_k\}$ is weakly dependent in the sense of Andrews (1988) so that an LLN applies to weighted sums $K^{-1}\sum_k g_k h_k$ with deterministic uniformly bounded weights; (iii) Assumption 2 is strengthened so that $\mathbf{g}'\mathbf{Q}^P\mathbf{g} / \mathbb{E}[\mathbf{g}'\mathbf{Q}^P\mathbf{g}] \xrightarrow{p} 1$ with $\mathbf{Q}^P = \mathbf{P}'\mathbf{Q}\mathbf{P}$. Under Assumption 1, $\text{plim}_{K\to\infty} \hat{\gamma}_{NSS} = \beta$.*

Proof. Write the numerator of $\hat{\gamma}_{NSS} - \beta$ as $\sum_k g_k h_k$ with $h_k = \sum_{l,r} P_{lk} s_{rl} \varepsilon_r$ $\mathbf{S}$-measurable. Condition (ii) applies the Andrews (1988) LLN: $K^{-1}\sum_k g_k h_k \xrightarrow{p} 0$. The denominator satisfies $\mathbb{E}[\mathbf{g}'\mathbf{Q}^P\mathbf{g}] = \text{tr}(\mathbf{Q}^P\boldsymbol{\Sigma})$; condition (i) bounds $|\text{tr}(\mathbf{Q}^P\boldsymbol{\Sigma})| \le C_\Sigma \text{tr}(\mathbf{Q}^P)$, which grows at rate $\text{tr}(\mathbf{Q}^P) = \Theta(R)$ under Assumption 2; condition (iii) delivers concentration. The ratio is $o_p(\text{tr}(\mathbf{Q}^P))/\Theta_p(\text{tr}(\mathbf{Q}^P)) \to 0$. □

Proposition 9 (Properties of ϕ)**.** *The contamination multiplier $\phi = \text{tr}(\mathbf{Q}\mathbf{P})/\text{tr}(\mathbf{Q})$ satisfies:*

(i) ***Non-negativity:*** *$\phi \ge 1$, with equality iff (i) $P_{kk} = 1$ for every k with $Q_{kk} > 0$ and (ii) $P_{kl} = 0$ for every off-diagonal pair (k, l) with $Q_{kl} > 0$. Propagation strictly outside the co-location support of $\mathbf{Q}$ leaves $\phi = 1$ but still affects regional prediction errors Δ_r^{std}.*

(ii) ***Neumann decomposition:*** *when* $\mathbf{P} = (\mathbf{I} - \mathbf{A})^{-1}$ *and* $\rho(\mathbf{A}) < 1$,

$$\phi \;=\; 1 + \sum_{n=1}^{\infty} \phi^{(n)}, \qquad \phi^{(n)} \;=\; \frac{\mathrm{tr}(\mathbf{Q}\mathbf{A}^n)}{\mathrm{tr}(\mathbf{Q})}.$$

(iii) ***Monotonicity:*** ϕ *is weakly increasing in each entry* $P_{k,l}$ *with* $Q_{k,l} > 0$ *(trivial from the sum-of-non-negative-terms form of* $\mathrm{tr}(\mathbf{QP}) = \sum_{k,l} Q_{k,l} P_{k,l}$*).*

(iv) ***Finiteness and bounds:*** *when* $\mathbf{P} = (\mathbf{I} - \mathbf{A})^{-1}$ *with* $\rho(\mathbf{A}) < 1$*, the Neumann series in (ii) converges, so* ϕ *is finite. The scalar bound* $\phi \le 1/(1 - \rho(\mathbf{A}))$ *holds when* $\mathbf{A}$ *is symmetric or normal (there the spectral radius controls each Neumann term,* $\phi^{(n)} \le \rho(\mathbf{A})^n$ *by the von Neumann trace inequality), but it can fail for a non-normal* $\mathbf{A}$*, whose transient powers* $\mathbf{A}^n$ *may load on the co-location support of* $\mathbf{Q}$ *by more than* $\rho(\mathbf{A})^n$*.*[12] *In the BEA 1997 application* $\rho(\mathbf{A}) = 0.506$ *and* $\phi = 1.110$ *sits comfortably below* $1/(1 - \rho(\mathbf{A}))$*.*

(v) ***Network centrality:*** *writing* $\phi = \sum_l c_l$ *with* $c_l = \sum_k Q_{k,l} P_{k,l} / \mathrm{tr}(\mathbf{Q})$*, shock-origin industries* l *with large column sums* $\sum_k P_{k,l}$ *(i.e., that propagate widely through the network) contribute more, weighted by their co-location with the receiving industries.*

(vi) ***Shock-distribution invariance:*** ϕ *depends only on* $(\mathbf{S}, \mathbf{P})$*; not on shock variance* σ^2*, realized shocks, or causal coefficients.*

Aggregation behavior is reported empirically in Table 16 rather than as a formal theorem: coarser BEA aggregations give higher $\hat{\phi}$ in the data, but the mechanical argument from inflation of off-diagonal co-location entries requires specific non-negativity and partition conditions on the aggregation matrix that are not universal.

Corollary 6 (Ordering Under Non-Negative Correlation and Variance–Centrality Alignment)**.** *With* $\mathbf{\Delta} \equiv \mathbf{Q}(\mathbf{P} - \mathbf{I})$ *and* $\langle \mathbf{A}, \mathbf{B} \rangle_F = \sum_{k,l} A_{kl} B_{kl}$*, the correlated-shocks factor satisfies*

$$\phi^{\mathrm{corr}} \;=\; 1 + \frac{\langle \mathbf{\Sigma}, \mathbf{\Delta} \rangle_F}{\langle \mathbf{\Sigma}, \mathbf{Q} \rangle_F}.$$

Under $\sigma_{kl} \ge 0$ *everywhere,* $\phi^{\mathrm{corr}} \ge 1$*. The ordering* $\phi^{\mathrm{corr}} \ge \phi^{\mathrm{hetero}} \ge \phi^{\mathrm{iid}} \ge 1$ *additionally requires positive alignment between shock variances and diagonal centrality, i.e.* $\sum_k \sigma_k^2 [\mathbf{QP}]_{kk} / \sum_k \sigma_k^2 Q_{kk} \ge \sum_k [\mathbf{QP}]_{kk} / \sum_k Q_{kk}$*.*

A.4 Measurement Error in the Propagation Matrix

The correction uses a single national propagation matrix $\mathbf{P} = (\mathbf{I} - \mathbf{A})^{-1}$ for every region. If the true region-specific matrix is $\mathbf{P}_r^*$ but the researcher uses $\widehat{\mathbf{P}}$, the network-adjusted exposure $\widehat{X}_r = \mathbf{s}_r' \widehat{\mathbf{P}} \mathbf{g}^{US}$ and instrument $\widehat{Z}_r = \mathbf{s}_r' \widehat{\mathbf{P}} \mathbf{g}^{OC}$ are constructed with error. A common intuition holds that such error attenuates the corrected coefficient toward the standard

[12] For instance, take $\mathbf{A}$ with sole nonzero entry $A_{21} = 0.95$: then $\rho(\mathbf{A}) = 0$, so the scalar bound would assert $\phi \le 1$, yet $\phi = 1.475$ at the all-ones co-location matrix $\mathbf{Q} = \mathbf{1}\mathbf{1}'$ (at $\mathbf{Q} = \mathbf{I}$ the off-diagonal load of $\mathbf{A}$ is invisible and $\phi = 1$). A finite bound that always holds is the operator norm $\phi \le \|\mathbf{P}\|_{\mathbf{Q}}$ in the $\mathbf{Q}$-seminorm.

Bartik, making the headline ratio a lower bound. That intuition is *not* a theorem: $\widehat{X}_r$ and $\widehat{Z}_r$ are nonlinear in $\mathbf{A}$ through $(\mathbf{I}-\mathbf{A})^{-1}$ (so even mean-zero error in $\mathbf{A}$ induces non-classical, non-mean-zero error in $\widehat{\mathbf{P}}$), and the relevant object is the alignment of the induced exposure error with the instrument, whose sign is not pinned down. What *is* available is an exact decomposition, a sign condition, and a sensitivity bound.

Proposition 10 (Exact IV bias decomposition under propagation-matrix error)**.** *Let the true model be $y_r = \beta X_r^* + \varepsilon_r$ with $X_r^* = \mathbf{s}_r'\mathbf{P}_r^*\mathbf{g}^{US}$, and let the researcher instrument $\widehat{X}_r = X_r^* + U_r$ with $\widehat{Z}_r$, where $U_r = \widehat{X}_r - X_r^*$ is the induced exposure error. If $\mathrm{Cov}(\widehat{Z},\varepsilon) = 0$ and $\mathrm{Cov}(\widehat{Z},\widehat{X}) \neq 0$, the IV estimand $\hat{\beta}_{\widehat{P}} = \mathrm{Cov}(\widehat{Z},y)/\mathrm{Cov}(\widehat{Z},\widehat{X})$ satisfies*

$$\frac{\hat{\beta}_{\widehat{P}}}{\beta} - 1 \;=\; -\frac{\mathrm{Cov}(\widehat{Z},U)}{\mathrm{Cov}(\widehat{Z},\widehat{X})}.$$

Propagation-matrix error biases the slope only through the covariance between the instrument and the induced exposure error: if $\mathrm{Cov}(\widehat{Z},U) = 0$ the slope is unbiased despite region-level exposure error; with a positively-normalized first stage, $\mathrm{Cov}(\widehat{Z},U) < 0$ inflates $|\hat{\beta}_{\widehat{P}}|$ (under-correction) and $\mathrm{Cov}(\widehat{Z},U) > 0$ deflates it (over-correction).

Proof. Since $y = \beta X^* + \varepsilon$ and $\mathrm{Cov}(\widehat{Z},\varepsilon) = 0$, $\hat{\beta}_{\widehat{P}} = \beta\,\mathrm{Cov}(\widehat{Z},X^*)/\mathrm{Cov}(\widehat{Z},\widehat{X})$. Substituting $X^* = \widehat{X} - U$ gives $\mathrm{Cov}(\widehat{Z},X^*) = \mathrm{Cov}(\widehat{Z},\widehat{X}) - \mathrm{Cov}(\widehat{Z},U)$; dividing and subtracting one yields the identity. □

Corollary 7 (Sensitivity bound)**.** *Under the conditions of Proposition 10, Cauchy–Schwarz gives*

$$\left|\frac{\hat{\beta}_{\widehat{P}}}{\beta} - 1\right| \;\leq\; \frac{\mathrm{sd}(\widehat{X}-X^*)}{\mathrm{sd}(\widehat{X})} \cdot \frac{1}{|\mathrm{Corr}(\widehat{Z},\widehat{X})|}.$$

A strong first stage and an induced exposure error that is small relative to observed exposure variation jointly cap the proportional slope bias, regardless of the sign or structure of the matrix error.

National-versus-regional mismatch. Specializing to $\widehat{\mathbf{P}} = \mathbf{P}^{\mathrm{nat}}$ and $\mathbf{P}_r^* = \mathbf{P}^{\mathrm{nat}} + \mathbf{\Delta}_r$, the induced error $U_r = -\mathbf{s}_r'\mathbf{\Delta}_r\mathbf{g}^{US}$ yields proportional slope bias $\mathrm{Cov}(\widehat{Z},\mathbf{s}'\mathbf{\Delta}\mathbf{g}^{US})/\mathrm{Cov}(\widehat{Z},\widehat{X})$. Given any regional IO proxy $\mathbf{P}_{\mathrm{state}(r)}$, this is directly estimable as

$$\hat{\kappa} \;=\; \frac{\mathrm{Cov}(\widehat{Z},\widehat{d})}{\mathrm{Cov}(\widehat{Z},\widehat{X})}, \qquad \widehat{d}_r = \mathbf{s}_r'\big(\mathbf{P}_{\mathrm{state}(r)} - \mathbf{P}^{\mathrm{nat}}\big)\mathbf{g}^{US},$$

with the matching bound $|\hat{\kappa}| \leq [\mathrm{sd}(\widehat{d})/\,\mathrm{sd}(\widehat{X})]\;/\;|\mathrm{Corr}(\widehat{Z},\widehat{X})|$. Section 6.7 reports $\hat{\kappa}$ using state-specific intra-state Leontief inverses.

Remark 2 (Two limiting cases). Setting $\widehat{\mathbf{P}} = \mathbf{I}$ (the standard Bartik) gives $U_r = -M_r$ with $M_r = \mathbf{s}_r'(\mathbf{P}-\mathbf{I})\mathbf{g}^{US}$, and Proposition 10 returns the omitted-network-bias formula $\hat{\beta}_{\mathrm{std}}/\beta - 1 = \mathrm{Cov}(Z_D,M)/\mathrm{Cov}(Z_D,D)$ of Proposition 1, the same contamination viewed

as the special case in which the entire network channel is the "error." The intermediate scalar-shrinkage path $\widehat{\mathbf{P}}_a = \mathbf{I} + a(\mathbf{P}^* - \mathbf{I})$ nests the standard Bartik ($a = 0$) and the correct correction ($a = 1$); under a positive first stage and a positive instrument–omitted-network covariance, an under-correction $a \in [0, 1)$ leaves the estimand between the standard and the true corrected coefficient, while $a > 1$ overshoots. This recovers the attenuation intuition as a benchmark for one special (scalar-contraction) error structure rather than as a general theorem.

A.5 Non-Co-Located Propagation: A Worked Example

The support-qualified equality condition for $\phi = 1$ in Proposition 1 and Proposition 9 can be made concrete with a small worked example. The example shows that propagation occurring strictly outside the co-location support of $\mathbf{Q}$ leaves the slope coefficient unbiased ($\phi = 1$) but still produces nonzero per-region prediction errors at every region. The economic content: a standard Bartik can pass every coefficient-level diagnostic the literature applies (no slope bias, well-calibrated standard errors) while systematically misallocating predicted impact across regions, and the only object that captures the misallocation is Δ_r^{std}.

Setup. Two regions ($R = 2$) and two industries ($K = 2$, indexed A and B). Region 1 fully specializes in industry A; region 2 fully specializes in industry B:

$$\mathbf{S} = \begin{pmatrix} 1 & 0 \\ 0 & 1 \end{pmatrix}, \qquad \mathbf{Q} = \mathbf{S}'\mathbf{S} = \begin{pmatrix} 1 & 0 \\ 0 & 1 \end{pmatrix}.$$

The co-location matrix is diagonal: the two industries share no geographic footprint. The propagation matrix has identity diagonals (no self-amplification) and symmetric off-diagonal propagation of strength $\delta > 0$ between the two industries:

$$\mathbf{P} = \begin{pmatrix} 1 & \delta \\ \delta & 1 \end{pmatrix}, \qquad \mathbf{QP} = \begin{pmatrix} 1 & \delta \\ \delta & 1 \end{pmatrix}.$$

Industry A's shock propagates to industry B with weight δ, and conversely. The off-diagonal of $\mathbf{P}$ falls on the off-diagonal of $\mathbf{Q}$, where $Q_{AB} = Q_{BA} = 0$.

Computing the contamination multiplier,

$$\phi = \frac{\operatorname{tr}(\mathbf{QP})}{\operatorname{tr}(\mathbf{Q})} = \frac{1+1}{1+1} = 1,$$

independent of δ. The support-qualified condition of Proposition 9 holds: $P_{AA} = P_{BB} = 1$ wherever $Q_{kk} > 0$, and $P_{AB}, P_{BA} > 0$ but $Q_{AB} = Q_{BA} = 0$, so the off-diagonal propagation contributes zero to $\operatorname{tr}(\mathbf{QP}) - \operatorname{tr}(\mathbf{Q})$. By Proposition 1, the population slope is uncontaminated: $\operatorname{plim}\hat{\gamma}_{SS} = \beta_D + \beta_N(\phi - 1) = \beta_D$, equal to the structural mixture $\beta_D + 0$ in the general two-channel case. In this $K = 2$ design there is no law of large numbers, so the statement is to be read as an expectation over i.i.d. shock draws—$\mathbb{E}[\mathbf{g}'\mathbf{Q}(\mathbf{P} - \mathbf{I})\mathbf{g}] = \sigma^2 \operatorname{tr}(\mathbf{Q}(\mathbf{P} - \mathbf{I})) = 0$ because $\mathbf{Q}$ has no off-diagonal support—not

as a large-K approximation. On any single realized draw the exact projection of Lemma 1 is $\hat{\gamma}_{N|B} = 2\delta g_A g_B/(g_A^2 + g_B^2)$, generically nonzero, so the example also illustrates in miniature the gap between the i.i.d. diagnostic and a realized contamination factor.

For a primitive shock vector $\mathbf{g} = (g_A, g_B)'$,

$$\begin{aligned} B_1 &= g_A, & N_1 &= \delta g_B, & \tilde{B}_1 &= g_A + \delta g_B, & m_1 &= 1 + \delta g_B/g_A, \\ B_2 &= g_B, & N_2 &= \delta g_A, & \tilde{B}_2 &= g_B + \delta g_A, & m_2 &= 1 + \delta g_A/g_B. \end{aligned}$$

The region-level effective multipliers m_r depend on the realized shock vector through the propagated component: region 1's effective exposure is its own industry-A shock plus a δ-weighted dose of industry-B's shock arriving through the network, and conversely. By Corollary 2,

$$\Delta_r^{\text{std}} \;=\; \beta_N B_r(\phi - m_r) \;=\; -\beta_N N_r,$$

which is nonzero at every region whenever $\delta > 0$ and $g_l \neq 0$ on the propagating industry: $\Delta_1^{\text{std}} = -\beta_N \delta g_B$ and $\Delta_2^{\text{std}} = -\beta_N \delta g_A$.

Numerical instance. Set $\delta = 0.5$, $g_A = 2$, $g_B = 1$, $\beta_D = \beta_N = 1$, and equal population weights $w_1 = w_2 = 1/2$. Then $B = (2, 1)$, $N = (0.5, 1)$, $\tilde{B} = (2.5, 2)$, true impacts $T_r = B_r + N_r$ give $T = (2.5, 2)$, and standard predictions $\hat{T}_r^{\text{std}} = B_r$ give $\hat{T}^{\text{std}} = (2, 1)$. (At this draw the realized projection is $\hat{\gamma}_{N|B} = 2\delta g_A g_B/(g_A^2 + g_B^2) = 0.4$, so a regression run on the single draw would return $\hat{\gamma}_{SS} = 1.4$; the predictions here are evaluated at the population slope $\beta_D = 1$.) The per-region prediction errors are

$$\Delta_1^{\text{std}} = 2 - 2.5 = -0.5, \qquad \Delta_2^{\text{std}} = 1 - 2 = -1.$$

The standard estimator under-predicts true impact at every region by a different amount, and the aggregate cancellation also fails: $\sum_r w_r \Delta_r^{\text{std}} = -0.75 \neq 0$. Checking against the aggregate identity (25), $\bar{B}_w = 1.5$, $\psi^{\text{agg}} = (0.5 \cdot 2.5 + 0.5 \cdot 2)/1.5 = 1.5$, and $\beta_N \bar{B}_w(\phi - \psi^{\text{agg}}) = 1.5 \cdot (1 - 1.5) = -0.75$, confirming the formula. The MSE decomposition of Proposition 8 gives an aggregate-level squared error of $(-0.75)^2 = 0.5625$ and a cross-region variance of $\text{Var}_w(\Delta_r^{\text{std}}) = 0.5 \cdot (-0.5 + 0.75)^2 + 0.5 \cdot (-1 + 0.75)^2 = 0.0625$, so the example is one in which the aggregate-level component dominates the MSE, the opposite of the ADH pattern, where the cross-region variance dominates by a factor of roughly 10^3.

The standard slope coefficient is correct ($\hat{\gamma}_{SS} \to \beta_D + 0$); $\phi = 1$ is a meaningful diagnostic that signals no slope contamination. But the standard estimator's region-level predictions are wrong at every region, and its population-weighted aggregate is also wrong. A practitioner reading only the slope and a confidence interval would see no problem; the prediction error Δ_r^{std} exposes the residual misallocation. In ADH the opposite alignment obtains ($\phi \neq 1$, slope contaminated, but $\phi \approx \psi^{\text{agg}}$ so aggregates cancel), and the geographic component of MSE is large. The two regimes share the same diagnostic apparatus: ϕ alone characterizes neither, while $\Delta_r^{\text{std}} = \beta_N B_r(\phi - m_r)$ characterizes both.

A.6 Framework Correspondences (BHJ, GPS, AKM, BH)

The contamination formula links to four leading shift-share identification frameworks.

BHJ equivalence. The standard shift-share estimator admits the shock-level representation $\hat{\gamma}_{SS} = \sum_k g_k \hat{y}_k / \sum_k g_k \hat{B}_k$ with $\hat{y}_k = \sum_r s_{rk} y_r$ and $\hat{B}_k = [\mathbf{Qg}]_k$. The network correction replaces $\hat{B}_k$ with $[\mathbf{QPg}]_k$ and s_{rk} with $\tilde{s}_{rk}$ in the numerator, preserving the BHJ shock-level structure.

GPS reduced-form identities. Under share exogeneity, the just-identified IV estimate using share s_{rk} as instrument satisfies $\hat{\beta}_k = \beta_D + \beta_N \text{Cov}(s_{rk}, N_r)/\text{Cov}(s_{rk}, B_r)$. The Rotemberg-weighted average of $\hat{\beta}_k$'s reconstructs the standard Bartik; replacing (s_{rk}, B_r) with $(\tilde{s}_{rl}, \tilde{B}_r)$ reconstructs the corrected Bartik, shifting the identifying industries toward upstream linkages (empirical Rotemberg shift reported in Section 4.4).

AKM residual structure under misspecification. The misspecified regression $y_r = \gamma B_r + u_r$ leaves residuals $u_r = \beta_N \, \mathbf{s}_r'(\mathbf{P} - \mathbf{I})\mathbf{g} + \varepsilon_r$, which are correlated across regions sharing industries linked to a common primitive shock. The correct variance formula (Proposition 4) aggregates by the primitive-shock index l using the network-adjusted shares $\tilde{s}_{rl}$.

BH recentering. Under i.i.d. mean-zero shocks, recentered Bartik $\tilde{B}_r^{BH} = B_r - \mathbb{E}_{\mathbf{g}}[B_r] = B_r$; the corrected Bartik $\tilde{B}_r = \mathbf{s}_r'\mathbf{Pg}$ is itself a linear formula instrument under the Borusyak and Hull (2023) classification, and the contamination multiplier $\phi - 1$ is the closed-form cost of the specific choice of formula ($\mathbf{s}_r'\mathbf{g}$ instead of $\mathbf{s}_r'\mathbf{Pg}$) in terms of the omitted exposure map.

B Three-Channel Extension: Separating Upstream and Downstream

The framework of Section 2 treats the network correction as a single-channel problem: the standard Bartik omits network-propagated exposure, and a single corrected exposure $\tilde{B}_r = \mathbf{s}_r'\mathbf{Pg}$ (with $\mathbf{P}$ a single propagation matrix, canonically the Leontief inverse $\mathbf{L}$) recovers the network-inclusive exposure. The two-regressor specification (B_r, N_r) developed in Section 2.4 is positioned as the methodological gold standard whenever shift-exogenous BHJ-style instruments and the relevant rank conditions are available, with the corrected single-regressor $\tilde{B}_r$ playing the role of a single-index/projection target under heterogeneous channels. The three-channel extension developed here, (B_r, N_r^L, N_r^G) with upstream (Leontief) and downstream (Ghosh) propagation identified separately, is a robustness and interpretation appendix rather than a competing main estimator. However, introducing the third channel in addition to the two-channel split matters for three reasons. First, the three-channel identification requires sufficient first-stage strength for the additional Ghosh-based regressor, which in the ADH application is weaker than the two-regressor

specification's first stages (Appendix C reports the relevant SW F-statistics). Second, the economics of the channels (direct displacement, upstream reduced input demand, downstream cheaper foreign inputs) requires a clean story to communicate, and the empirical signs in ADH (downstream channel net negative because local-supply-disruption dominates the Pierce–Schott cost-reduction benefit) require additional narration that would burden the main text. Third, the two-channel framework already captures the closed-form bias of the standard Bartik and the per-region prediction error $\Delta_r^{\text{std}} = \beta_N B_r(\phi - m_r)$ that organizes the paper's empirical findings; promoting three channels to the main framework would dilute that organizing principle without changing the headline results. We therefore develop the three-channel decomposition here as a sharpening of the two-channel result, useful for interpreting which network channels do the work, rather than as the recommended baseline specification. We view the three-channel form as a complement to the two-regressor specification, not as a replacement for it: applied users with strong first stages for both upstream and downstream regressors and a clear interpretive story may want to lead with it; otherwise the two-regressor specification is the recommended workflow.

In this appendix we develop the three-channel extension in which upstream (demand-side, Leontief) and downstream (supply-side, Ghosh) propagation are identified separately, following the applied motivation in Acemoglu et al. (2016) and Pierce et al. (2024).

B.1 Three-Channel DGP and Bias Decomposition

Let $\mathbf{G} = (\mathbf{I} - \mathbf{B})^{-1}$ denote the Ghosh inverse, where $\mathbf{B}$ is the matrix of direct output-allocation coefficients (B_{kl} is the share of industry k's output sold to industry l). The three-channel DGP allows the direct, upstream, and downstream exposure components to carry distinct causal effects:

$$y_r \;=\; \beta_D\, \mathbf{s}_r'\mathbf{g} \;+\; \beta_L\, \mathbf{s}_r'(\mathbf{L} - \mathbf{I})\mathbf{g} \;+\; \beta_G\, \mathbf{s}_r'(\mathbf{G} - \mathbf{I})\mathbf{g} \;+\; \varepsilon_r, \tag{38}$$

with regressors $B_r = \mathbf{s}_r'\mathbf{g}$, $N_r^L = \mathbf{s}_r'(\mathbf{L} - \mathbf{I})\mathbf{g}$, $N_r^G = \mathbf{s}_r'(\mathbf{G} - \mathbf{I})\mathbf{g}$.

Proposition 11 (Three-Channel Bias Decomposition)**.** *Under the DGP* (38) *and Assumptions 1–3, the standard estimator satisfies*

$$\operatorname{plim}_{K\to\infty} \hat{\gamma}_{SS} \;=\; \beta_D \;+\; \beta_L(\phi_L - 1) \;+\; \beta_G(\phi_G - 1), \tag{39}$$

where $\phi_L = \operatorname{tr}(\mathbf{QL})/\operatorname{tr}(\mathbf{Q})$ *and* $\phi_G = \operatorname{tr}(\mathbf{QG})/\operatorname{tr}(\mathbf{Q})$ *are the upstream and downstream contamination multipliers (both* ≥ 1*, with equality iff the corresponding propagation matrix is the identity).*

Proof. Apply the Proposition 1 argument separately to each channel. Under i.i.d. shocks the three trace ratios converge independently, and linearity of the OLS formula in the DGP coefficients delivers (39). □

Definition 2 (Three-Regressor Network-Corrected Estimator)**.** The three-regressor network-corrected estimator is the 2SLS estimator from

$$y_r \;=\; \alpha + \beta_D B_r + \beta_L N_r^L + \beta_G N_r^G + \mathbf{X}_r'\boldsymbol{\delta} + u_r, \tag{40}$$

with each of (B_r, N_r^L, N_r^G) instrumented by its analog $(B_r^{IV}, N_r^{L,IV}, N_r^{G,IV})$ constructed from an exogenous shock $\mathbf{g}^{IV}$.

Proposition 12 (Joint Consistency)**.** *Under Assumptions 1–3, the three-regressor estimator is jointly consistent for $(\beta_D, \beta_L, \beta_G)$ as $K \to \infty$, provided the instrument matrix $[\mathbf{B}^{IV}, \mathbf{N}^{L,IV}, \mathbf{N}^{G,IV}]$ spans three linearly independent directions in region-space.*

Proof. The argument parallels Proposition 1 applied jointly to the three endogenous regressors. The rank condition is not automatic: it requires an independent source of identifying variation beyond the shares matrix alone, typically a shift-exogenous instrument for the shocks. Under Borusyak et al. (2022) (BHJ) identification the other-country instrument supplies this; under Goldsmith-Pinkham et al. (2020) (GPS) identification alone, with a single realized shock vector, the three exposure regressors are fixed and typically nearly collinear transformations of the same shares matrix, so the rank condition holds only through the geometry of three nearly parallel vectors and identification is weak in practice (Section 2.5). □

B.2 Three-Channel Impacts and Cancellation

Define the channel-specific effective multipliers $m_r^c = \mathbf{s}_r' \mathbf{P}_c \mathbf{g} / B_r$ for $c \in \{L, G\}$, the channel ratios $\theta_c = \beta_c / \beta_D$, and the population-weighted aggregates $\psi_c^{\text{agg}} = (\sum_r w_r \, \mathbf{s}_r' \mathbf{P}_c \mathbf{g}) / \bar{B}_w$.

Proposition 13 (Three-Channel Overstatement)**.** *Under the DGP* (38)*,*

$$\frac{\hat{y}_r^{\text{std}}}{\hat{y}_r^{\text{corr}}} \;=\; \frac{1 + \theta_L(\phi_L - 1) + \theta_G(\phi_G - 1)}{1 + \theta_L(m_r^L - 1) + \theta_G(m_r^G - 1)},$$

and the aggregate ratio is

$$\frac{T^{\text{std}}}{T^{\text{corr}}} \;=\; \frac{1 + \theta_L(\phi_L - 1) + \theta_G(\phi_G - 1)}{1 + \theta_L(\psi_L^{\text{agg}} - 1) + \theta_G(\psi_G^{\text{agg}} - 1)},$$

equal to 1 *exactly when* $\theta_L(\phi_L - \psi_L^{\text{agg}}) + \theta_G(\phi_G - \psi_G^{\text{agg}}) = 0$*.*

The channel-level covariance identity $\psi_c^{\text{agg}} = \bar{m}_c^{\text{pop}} + \text{Cov}_{\text{pop}}(m_r^c, B_r) / \bar{B}_w$ (Proposition 6 applied channel-by-channel) decomposes the channel-specific gap between contamination and aggregate amplification into a mean-shift and a covariance term, exactly as in the two-channel case.

B.3 Empirical Three-Channel Decomposition in ADH

The body works with the two-regressor (B_r, N_r) split (Section 4.2); this subsection reports the three-regressor refinement that further separates the network-only channel into upstream (Leontief) and downstream (Ghosh) components. By decomposing the network-only channel into upstream and downstream pieces, we show which linkages carry the

network effect. At manufacturing granularity ($K = 342$), the three-regressor OLS estimates are

$$\hat{\beta}_D = -0.112, \qquad \hat{\beta}_L = -0.121, \qquad \hat{\beta}_G = -0.166.$$

The 2SLS equivalents on the full-economy share matrix ($K = 483$), reported in Table A.3, are

$$\hat{\beta}_D = -0.274, \qquad \hat{\beta}_L = +0.015, \qquad \hat{\beta}_G = -0.080,$$

with cluster-robust Sanderson–Windmeijer (Sanderson and Windmeijer, 2016) first-stage F-statistics of 548.4, 3387.8, and 273.6 on the three endogenous regressors, all far above conventional weak-instrument thresholds. (Table A.3's notes report per-endogenous cluster-robust Wald statistics for the same specification, a different statistic with visibly different values; both are labeled where they appear.) F-statistics in the thousands reflect the mechanical near-determinism of instrumenting three functions of the same underlying shock vector with three linear transformations of its other-country analog; the economically relevant identification argument is exogeneity of $\mathbf{g}^{OC}$ (BHJ; Lemma 2 propagates this to the transformed channels) rather than F-statistic magnitude. The residualized design-matrix condition number 1.50 (well below the conventional threshold of 30) confirms the three regressors span three linearly independent directions in region-space. Across the three channels we find that the direct channel is by far the largest in magnitude; the downstream channel is moderate and negative; the combined upstream channel is small and of opposite sign under the combined-upstream construction.

The direct effect $\hat{\beta}_D = -0.274$ captures own-industry displacement by direct China Shock exposure. It is smaller in magnitude than either the standard Bartik or the two-regressor OVB direct coefficient, because those specifications conflate own-industry displacement with network-propagated effects from co-located industries.

The downstream effect $\hat{\beta}_G = -0.080$ is moderate and negative. This sign is somewhat puzzling, since under the Pierce et al. (2024) input-cost-reduction mechanism the predicted sign is positive: U.S. customer industries should benefit from cheaper Chinese inputs displacing their former U.S. suppliers. The negative sign we estimate indicates that in ADH this Pierce–Schott benefit is dominated by two opposing forces: supply-chain disruption (customers lose their U.S. suppliers and cannot frictionlessly substitute) and local-multiplier contraction (when hit supplier industries shed employment, local spending contracts and pulls down neighboring customer industries). Our result is consistent with the existence of the Pierce–Schott mechanism while showing that it does not dominate the net downstream effect; Section B.4 below decomposes the downstream channel by whether the CZ's origin industries are manufacturing or non-manufacturing, with $\hat{\beta}_G^{\text{mfg}} = -0.158$ carrying the bulk of the effect and pointing to a local-supply-disruption interpretation.

The combined upstream effect $\hat{\beta}_L = +0.015$ is small in magnitude (an order smaller than the direct or downstream channel) and positive under the combined-upstream construction. Contrary to the naive supplier-loses-demand expectation, the sign comes out positive, not negative. The sign is not a headline finding: the combined coefficient's small magnitude means it carries little weight in predicted impacts, and the economic content of the upstream channel is revealed by the further mfg/nmfg split below. Through CZ manufacturing suppliers the coefficient is the naive negative ($\hat{\beta}_L^{\text{mfg}} = -0.110$), while

through CZ non-manufacturing suppliers (services, logistics, construction) it is small and positive. The combined positive is a net of two opposing components dominated by services exposure. We do not emphasize the combined sign as a structural finding and report the combined coefficient primarily to show that all three channels are precisely identified when instrumented jointly.

The three-regressor form $(B_r,\, N_r^L,\, N_r^G)$ with $N_r^c = \mathbf{s}_r'(\mathbf{P}_c - \mathbf{I})\mathbf{g}$ subtracts the identity from both Leontief and Ghosh inverses and is well-conditioned (diagnostics below). Our specification differs from Acemoglu et al. (2016), who instead capture network propagation by augmenting the ADH import-penetration regression with a separately constructed upstream net-input IO exposure, without orthogonalizing the two channels. Because that upstream measure propagates the same import shock through the input–output network, it is nearly collinear with direct import exposure (correlation 0.988 on the ADH data); Table 5 columns 4–5 reports the resulting two-regressor 2SLS, with coefficients $\hat{\beta}_D^{\text{AADHP}} = -2.051$ and $\hat{\beta}_{\text{Up}}^{\text{AADHP}} = +1.722$ that nearly cancel, large standard errors, and Sanderson–Windmeijer F-statistics (5.1 and 5.1) below conventional weak-instrument thresholds, reflecting the near-singularity of the augmented design. In contrast to that design, subtracting the identity from the Leontief and Ghosh inverses removes the directly overlapping component from each network channel, leaving regressors that span well-separated directions in region-space; the orthogonal basis is more informative because its coefficients have direct economic interpretation as per-unit direct, per-unit upstream-only, and per-unit downstream-only effects.

The three-regressor specification is well-conditioned in every standard diagnostic. First, the residualized design-matrix condition number is 1.50 at full economy and 1.66 at manufacturing-only granularity, far below the conventional threshold of 30, and all variance inflation factors are below 2.4. Second, a Monte Carlo exercise on a pure-direct DGP recovers the truth with mean biases of order 10^{-3} and no spurious $\hat{\beta}_G$ exceeding the empirical estimate (Appendix D). Third, a leave-one-three-digit-NAICS-cluster-out jackknife produces standard deviations well below the cluster-robust standard errors of Table 3.

B.4 Manufacturing/Non-Manufacturing Decomposition and Pierce–Schott Interpretation

Contrary to the naive supplier-loses-demand expectation, the combined upstream coefficient comes out positive, $\hat{\beta}_L > 0$. Under a five-regressor decomposition that separates upstream and downstream exposure by whether the CZ's origin industry is manufacturing or non-manufacturing,

$$\hat{\beta}_L^{\text{mfg}} = -0.110, \quad \hat{\beta}_L^{\text{nmfg}} = +0.013, \quad \hat{\beta}_G^{\text{mfg}} = -0.158, \quad \hat{\beta}_G^{\text{nmfg}} = +0.018,$$

the combined upstream coefficient is a net of two opposing components (Table A.1). Through a CZ's manufacturing suppliers the coefficient is the naive negative (suppliers lose demand when their customers contract), while through non-manufacturing suppliers (logistics, transport, business services) the coefficient is small and positive. The combined

positive $\hat{\beta}_L$ reflects the quantitative dominance of services exposure. The downstream decomposition is cleaner: the bulk of the effect is carried by CZ manufacturing customers of hit industries ($\hat{\beta}_G^{\text{mfg}} = -0.158$, larger in magnitude than the direct coefficient in the same five-regressor split; Table A.1), which points to a local-supply-disruption and local-multiplier mechanism rather than the Pierce et al. (2024) input-cost-reduction channel. The Pierce–Schott mechanism is consistent with the small positive $\hat{\beta}_G^{\text{nmfg}}$ (non-manufacturing customers may benefit from cheaper imported intermediate inputs), but it does not dominate the net downstream effect in the ADH setting. The single-regressor corrected Bartik of Section 2 aggregates these five components into one coefficient by construction; what the three-channel and five-regressor decompositions show is the economic content that the aggregation obscures.

B.5 Convention: $(\mathbf{L} - \mathbf{I})$ vs $(\mathbf{L} - \mathbf{L}_{\text{diag}})$

The three-channel form (38) defines the upstream regressor as $N_r^L = \mathbf{s}_r'(\mathbf{L} - \mathbf{I})\mathbf{g}$: this subtracts the identity from the Leontief inverse, leaving both off-diagonal (cross-industry) propagation *and* the within-industry round-about multiplier $(L_{kk} - 1) \geq 0$ that the Leontief diagonal places on each industry's own shock. In this subsection it is convenient to distinguish between this convention and an alternative that subtracts the full Leontief diagonal, assigning the within-industry multiplier to the direct channel:

$$\tilde{B}_r^{\text{alt}} = \mathbf{s}_r'\mathbf{L}_{\text{diag}}\mathbf{g} = \sum_k s_{rk} L_{kk} g_k, \qquad N_r^{L,\text{alt}} = \mathbf{s}_r'(\mathbf{L} - \mathbf{L}_{\text{diag}})\mathbf{g},$$

so that $\tilde{B}_r^{\text{alt}} + N_r^{L,\text{alt}} = \mathbf{s}_r'\mathbf{L}\mathbf{g}$ (the network-inclusive exposure is preserved). Under this convention β_D^{alt} carries the own-industry Leontief multiplier and β_L^{alt} captures strict cross-industry propagation.

On the same full-economy panel as the headline three-regressor specification, the alternative decomposition yields $\hat{\beta}_D^{\text{alt}} = -0.256$ (s.e. 0.039), $\hat{\beta}_L^{\text{alt}} = 0.015$ (s.e. 0.007), and $\hat{\beta}_G^{\text{alt}} = -0.079$ (s.e. 0.039); the original $(\mathbf{L} - \mathbf{I})$ decomposition on the same panel returns $(-0.274, 0.015, -0.080)$. Comparing the two conventions, we see that the direct coefficient moves modestly, reflecting the small mean own-industry multiplier in this Leontief matrix ($\bar{L}_{kk} = 1.039$, max 1.940), while the upstream and downstream coefficients are essentially unchanged. Because the Leontief diagonal carries little amplification away from unity at the BEA detailed level—input-output activity is dominantly cross-industry—the choice between the two conventions is quantitatively immaterial in ADH. For applications whose Leontief matrices have heavier diagonals (round-about-intensive industries, aggregated sectors) the conventions could diverge more; we report both as a standing diagnostic.

B.6 First-Order Versus All-Orders Propagation

The Leontief inverse $\mathbf{L} = \mathbf{I} + \mathbf{A} + \mathbf{A}^2 + \cdots$ folds in second- and higher-order propagation. One seemingly obvious objection to the correction is that its effect on the coefficient depends on exotic high-order network terms rather than on direct first-order supplier

and customer linkages. To address this concern, we replace $(\mathbf{L} - \mathbf{I})$ with the technical-coefficients matrix $\mathbf{A}$, which produces the Acemoglu et al. (2016)-style first-order construction, and replace $(\mathbf{G} - \mathbf{I})$ with the allocation matrix $\mathbf{B}$, which does the analogous thing on the downstream side.

The comparison is reported in Table A.2. The two-regressor corrected estimator with the first-order matrix yields coefficients close to the full-Leontief specification: the direct channel and the network-only channel are both negative, with comparable magnitudes. We find qualitatively very similar results from the three-regressor first-order decomposition, which reproduces the sign pattern of the full-Leontief/Ghosh decomposition reported above (Section B.3): small positive combined-upstream under the aggregate form, moderate negative downstream. The correction's coefficient shift is therefore present in direct first-order linkages; it is not an artifact of the higher-order tail of the Leontief expansion.

C Full-Economy Three-Regressor Extension

We chose the BEA manufacturing submatrix ($K = 342$ industries) because it is the natural aggregation for the China Shock, whose direct targets are all manufacturing. There are two reasons the larger granularity is nonetheless required. The three-channel extension of Appendix B and the multi-outcome analysis of this appendix require the full economy ($K = 483$ BEA detail industries, including non-manufacturing sectors), because downstream propagation from manufacturing to non-manufacturing industries (e.g., steel to automotive distribution and business services) is an explicit part of the economic mechanism, and outcome variables beyond manufacturing employment (total employment, unemployment, wages) can be affected through non-manufacturing channels. In this appendix we report the full-economy three-regressor 2SLS results and address concerns about concentration and residualization that arise at the larger granularity.

C.1 Full-Economy Three-Regressor 2SLS

The coefficient estimates for the $K = 483$ three-regressor 2SLS across outcomes are reported in Table A.3; Table A.4 reports the companion standard-versus-corrected comparison across the full ADH outcome menu. We find that the direct coefficient $\hat{\beta}_D = -0.274$ is the dominant channel for every outcome; the downstream coefficient is consistently moderate-to-large in magnitude and negative for employment outcomes; the upstream coefficient varies in sign across outcomes but is always small in magnitude. First-stage Sanderson–Windmeijer F-statistics are 548.4, 3387.8, and 273.6 for the three endogenous regressors, all far above conventional weak-IV thresholds, and the residualized design-matrix condition number is 1.50 (well below the conventional threshold of 30).

C.2 Concentration Caveat

Under the full-economy Leontief inverse the structural contamination multiplier is $\hat{\phi}_L^{\text{full}} = 1.149$ (about 1.15), modestly larger than the manufacturing-only $\hat{\phi}_L^{\text{mfg}} = 1.110$. This larger

granularity gives reason for caution on two counts. First, at $K = 483$ the denominator concentration diagnostic $\max_k [\mathbf{QL}]_{kk} / \text{tr}(\mathbf{QL})$ is larger than at $K = 342$: a small number of large non-manufacturing industries (real estate, business services, retail trade) carry disproportionate co-location mass, and the asymptotic approximation in Assumption 2 is correspondingly less tight. The concentration diagnostic is nonetheless well within the range for which Monte Carlo simulations recover the truth (Appendix D). Second, the full-economy analysis imputes non-manufacturing shares from the Eckert et al. (2024) harmonized County Business Patterns; the imputation is transparent and pre-determined relative to the treatment period, but non-manufacturing employment responses are subject to their own measurement error. Given these considerations, we treat the manufacturing-only specification as the main text headline and report full-economy results as the extension.

C.3 Manufacturing Submatrix: A Numerical Remark

Does the $\mathbf{L}^{\text{mfg}}$ used throughout the main text coincide with the manufacturing submatrix of the full-economy Leontief inverse, $\mathbf{L}^{\text{full}}[\text{mfg}, \text{mfg}]$? As a pure matrix-algebra statement the equality is not implied: block-inversion of $\mathbf{I} - \mathbf{A}$ under a manufacturing/non-manufacturing partition gives $\mathbf{L}^{\text{full}}[\text{mfg}, \text{mfg}] = (\mathbf{I} - \mathbf{A}_{\text{mm}} - \mathbf{A}_{\text{mn}}(\mathbf{I} - \mathbf{A}_{\text{nn}})^{-1}\mathbf{A}_{\text{nm}})^{-1}$, which reduces to the manufacturing-only inverse $(\mathbf{I} - \mathbf{A}_{\text{mm}})^{-1}$ only when $\mathbf{A}_{\text{mn}}(\mathbf{I} - \mathbf{A}_{\text{nn}})^{-1}\mathbf{A}_{\text{nm}} = \mathbf{0}$, i.e. when there are no through-non-manufacturing linkages between manufacturing industries. In our pipeline the two matrices are numerically identical in the BEA 1997 data (Frobenius norm of the difference is zero to machine precision) because the object used throughout the main text as $\mathbf{L}^{\text{mfg}}$ is constructed by extracting the manufacturing block of the full-economy Leontief inverse $\mathbf{L}^{\text{full}}$, not by independently inverting $(\mathbf{I} - \mathbf{A}_{\text{mm}})$ against the manufacturing-only block of the A-matrix. We made this construction choice because the extracted object correctly captures manufacturing-to-manufacturing propagation that routes through non-manufacturing intermediaries (for example, steel reaching automotive through transportation services), which a strict manufacturing-only $(\mathbf{I} - \mathbf{A}_{\text{mm}})^{-1}$ would omit. The notational shorthand $\mathbf{L}^{\text{mfg}} \equiv \mathbf{L}^{\text{full}}[\text{mfg}, \text{mfg}]$ is therefore accurate for our data, but it is a property of the specific pipeline construction rather than a general matrix-algebra invariance; readers transporting the framework to other contexts should verify which of the two objects is being used in their own pipelines.

C.4 Cross-Source Validation of the ADH China-Shock Vector

Our main data source for the ADH industry-level shock vector ΔM_{kt}^{US} is the Autor et al. (2013) replication files, which apply the Feenstra et al. (2005) cleaning that reassigns Chinese-origin goods transshipped through Hong Kong back to China. To assess how the headline coefficients change in response to this provenance, we reconstructed the same SIC87dd-level shock vector independently from three external sources: UN Comtrade HS6 trade flows (downloaded via the public API across Dorn's HS6-to-SIC87dd crosswalk), CEPII BACI bilateral trade flows (HS92 release, aggregated to SIC87dd), and WIOD 2013 (15 manufacturing industries, with the WIOD-to-SIC concordance applied). Cross-source

Pearson correlations against ADH at the SIC87dd industry level are 0.999 / 0.997 / 0.975 (Comtrade / BACI / WIOD) on the 1991–1999 changes, and 0.997 / 0.997 / 0.975 on the 1999–2007 changes after we drop three SIC industries (3651, 3663, 3944) for which Dorn's HS-1996/2002-based crosswalk fails to map post-2007 HS-revision codes; without that drop, the Comtrade 1999–2007 correlation falls to 0.748 (the BACI value is unaffected because CEPII maintains a live HS-vintage concordance). At the aggregate level, ADH's 2007 inflated total is about 16% higher than U.S. Census of-origin imports, attributable entirely to the Feenstra Hong Kong reassignment, which grows from roughly 5% of the total in 1991 to roughly 16% in 2007 as Hong Kong's share of Chinese-origin transshipment expanded. Taken together, the cross-source consistency at the cross-industry level, which is what the shift-share design loads on, supports the interpretation that the headline coefficients reflect the cross-sectional distribution of Chinese export-supply shocks rather than artifacts of the Feenstra cleaning or of Dorn's HS-to-SIC crosswalk.

D Monte Carlo Validation

We run two Monte Carlo exercises: one validating recovery under a null DGP; a second probing the role of shock correlation. The first shows that the three-regressor corrected estimator recovers the truth under a pure-direct DGP in which $\beta_L = \beta_G = 0$, ruling out the concern that the estimator might spuriously attribute variation to the network channels. The second shows that within-sector shock correlation amplifies the realized IV contamination factor above the structural i.i.d. diagnostic, directly supporting the bridge from ϕ^{iid} to ϕ^{IV} developed in Section 2.3.

D.1 Pure-Direct DGP Recovery

Under a pure-direct DGP $y_r = \beta_D \, \mathbf{s}_r' \mathbf{g} + \varepsilon_r$ with $\beta_D = -0.300$ and shocks drawn from the empirical ADH shock distribution, we generate 500 replicates on the actual ADH shares and Leontief inverse, and run the three-regressor 2SLS of Definition 2. Results:

- Mean bias on $\hat{\beta}_D$: -0.002 (true value -0.300). MC SD: 0.038.
- Mean bias on $\hat{\beta}_L$: 0.000 (true value 0.000). MC SD: 0.010.
- Mean bias on $\hat{\beta}_G$: 0.000 (true value 0.000). MC SD: 0.026.
- Fraction of replicates with spurious $|\hat{\beta}_G|$ exceeding the empirical point estimate $|-0.080|$: 0.000.

We find that all biases are of order 10^{-3} and that the null channels are never spuriously identified above the empirical magnitude. The check is not tautological: because the direct and network regressors are highly collinear by construction (the very feature that motivates the identity-subtracted basis of Section 4.2), this is why the key finite-sample question arising is whether estimation noise spuriously loads the null network channels, and the answer is that it does not. The three-regressor decomposition recovers the truth.

D.2 Correlated-Shocks Amplification

We replicated the same ADH three-regressor specification using shocks drawn from a correlated-shocks DGP in which shocks are clustered into 10 sectors of about 34 industries each (mirroring the three-digit NAICS structure of the ADH data). Within each sector the shocks have pairwise correlation ρ; across sectors they are independent. The true $\beta = -0.500$. Table A.5 reports 1,000-replicate results for $\rho \in \{0, 0.3, 0.5, 0.7, 0.9\}$.

We find that the corrected Bartik remains unbiased at every ρ (mean across replicates within 0.02 of -0.500). This exercise runs on a synthetic design (a random sparse $\mathbf{A}$ and its Leontief inverse, with random shares calibrated to the ADH dimensions $R = 722$, $K = 342$, and ten shock sectors), whose i.i.d. contamination factor is $\phi^{\text{iid}} = 3.01$ rather than the ADH value, so the magnitudes are not comparable to ADH's and the exercise speaks to the direction of the shock-correlation effect rather than to its size. At $\rho = 0$ the standard Bartik's ratio to the truth is 2.5, near this design's i.i.d. baseline; as within-sector correlation rises the ratio climbs monotonically to 7.4 at $\rho = 0.9$, well above the i.i.d. baseline. The main idea is that within-sector correlation preferentially activates cross-industry propagation pairs that are both co-located and IO-linked, producing the $\phi^{\text{iid}} \to \phi^{IV}$ amplification documented for ADH in Section 4.4 and pushing the realized bias above the i.i.d. trace diagnostic.

Appendix Tables

Table A.1. Upstream/Downstream Sector Heterogeneity: Manufacturing vs Non-Manufacturing Origin

	3-Regressor	5-Regressor (Split)
Direct β_D	-0.274	-0.098
	(0.042)	(0.030)
Upstream β_L (combined)	$+0.015$	
	(0.007)	
β_L^{mfg} (CZ mfg suppliers)		-0.110
		(0.023)
β_L^{nmfg} (CZ services suppliers)		$+0.013$
		(0.006)
Downstream β_G (combined)	-0.080	
	(0.039)	
β_G^{mfg} (CZ mfg customers)		-0.158
		(0.048)
β_G^{nmfg} (CZ services customers)		$+0.018$
		(0.028)
First-stage F (direct)	490	765
First-stage F (upstream mfg)	—	2333
First-stage F (upstream nmfg)	—	22269
First-stage F (downstream mfg)	—	160
First-stage F (downstream nmfg)	—	6749
Observations	1,444	1,444

Notes: Both columns run 2SLS on the 722-CZ $\times$ 2-period panel. Column 1 is the reference three-regressor spec. Column 2 splits each network-only regressor by the CZ's industry-of-origin: $N_r^{L,\text{mfg}} = \sum_{k \in \text{mfg}} s_{rk}[(\mathbf{L} - \mathbf{I})\mathbf{g}]_k$ restricts the outer sum to manufacturing rows of the share matrix (CZ's manufacturing industries that supply shock-hit industries), while $N_r^{L,\text{nmfg}}$ restricts to non-manufacturing (services, construction, utilities, transport). $N_r^{G,\text{mfg}}$ and $N_r^{G,\text{nmfg}}$ split the Ghosh regressor analogously. By construction $N_r^L = N_r^{L,\text{mfg}} + N_r^{L,\text{nmfg}}$ (and similarly for N_r^G). Each regressor is instrumented by its OC analog. Outcome: change in mfg employment / WAP.

Table A.2. Leontief vs. First-Order (AADHP-Style) Three-Regressor Decomposition

	Leontief/Ghosh $(\mathbf{L}-\mathbf{I}), (\mathbf{G}-\mathbf{I})$	First-order (AADHP) $\mathbf{A}, \mathbf{B}$
Direct β_D	-0.274	-0.268
	(0.042)	(0.036)
Upstream β_L/β_A	$+0.015$	$+0.017$
	(0.007)	(0.012)
Downstream β_G/β_B	-0.080	-0.076
	(0.039)	(0.059)
First-stage F (direct)	490	530
First-stage F (upstream)	16958	16047
First-stage F (downstream)	143	33
Observations	1,444	1,444

Notes: Both columns run the three-regressor 2SLS on the same 722-CZ $\times$ 2-period panel with the full ADH controls, timepwt48 weights, and state-clustered errors. Column 1 uses the full Leontief/Ghosh network-only measures $N_r^L = \mathbf{s}_r'(\mathbf{L}-\mathbf{I})\mathbf{g}$, $N_r^G = \mathbf{s}_r'(\mathbf{G}-\mathbf{I})\mathbf{g}$; column 2 uses the first-order AADHP-style analogs $M_r^A = \mathbf{s}_r'\mathbf{A}\mathbf{g}$, $M_r^B = \mathbf{s}_r'\mathbf{B}\mathbf{g}$. Because $(\mathbf{L}-\mathbf{I}) = \mathbf{A}+\mathbf{A}^2+\cdots$, the first-order measure captures only direct supplier/customer linkages whereas the Leontief measure captures all orders of propagation. Outcome: change in manufacturing employment / WAP.

Table A.3. Three-Regressor Full-Economy Decomposition: Direct, Upstream, and Downstream Channels

	Mfg Empl (483)	Non-Mfg Empl (483)	Total Empl (483)	Unempl. (483)	Log Wage (NM) (483)
Panel A: OVB three-regressor (full economy)					
Direct B_r	−0.274	−0.025	−0.299	+0.090	−0.359
	(0.042)	(0.057)	(0.059)	(0.030)	(0.127)
Upstream N_r^L	+0.015	+0.043	+0.057	+0.005	−0.050
	(0.007)	(0.016)	(0.019)	(0.006)	(0.026)
Downstream N_r^G	−0.080	−0.131	−0.211	+0.055	−0.169
	(0.039)	(0.076)	(0.103)	(0.043)	(0.135)
First-stage F direct B_r	489.9	489.9	489.9	489.9	489.9
First-stage F upstream N_r^L	16957.9	16957.9	16957.9	16957.9	16957.9
First-stage F downstream N_r^G	143.1	143.1	143.1	143.1	143.1
Panel B: Two-regressor OVB decomposition (B_r, N_r^L)					
Direct B_r	−0.259	−0.001	−0.259	+0.079	−0.327
	(0.034)	(0.064)	(0.061)	(0.028)	(0.141)
Upstream N_r^L	+0.014	+0.042	+0.056	+0.005	−0.051
	(0.007)	(0.017)	(0.020)	(0.006)	(0.026)
Panel C: Single corrected instrument $\tilde{B}_r = \mathbf{s}_r'\mathbf{Lg}$					
Full Leontief	+0.012	+0.041	+0.054	+0.005	−0.053
	(0.008)	(0.017)	(0.020)	(0.006)	(0.027)
Observations	1,444	1,444	1,444	1,444	1,444

Notes: Full-economy ($K = 483$) three-regressor, two-regressor, and single-instrument specifications. Panel A instruments (B_r, N_r^L, N_r^G) jointly with their other-country analogs constructed from eight high-income countries' imports from China (matching ADH). Panel B drops the downstream channel, and Panel C uses the single network-adjusted instrument $\tilde{B}_r = \mathbf{s}_r'\mathbf{Lg}$. Regressors: $B_r = \mathbf{s}_r'\mathbf{g}$ (direct), $N_r^L = \mathbf{s}_r'(\mathbf{L} - \mathbf{I})\mathbf{g}$ (upstream Leontief), $N_r^G = \mathbf{s}_r'(\mathbf{G} - \mathbf{I})\mathbf{g}$ (downstream Ghosh). All specifications use 1980 (P1) and 1991 (P2) lagged employment shares, built by combining Dorn-pipeline manufacturing shares with Eckert et al. (2024) non-manufacturing shares at the BEA detail level and row-renormalizing (combined row sums average 0.981 and 0.997 before renormalization). First-stage F-statistics are per-endogenous cluster-robust Wald tests of instrument relevance; they are not Sanderson–Windmeijer statistics, and so differ from the SW values 548.4, 3387.8, and 273.6 quoted for the same manufacturing-employment specification in Appendix B. Dependent variables are changes in each outcome as a share of working-age population, except Log Wage (NM) which is the change in log average non-manufacturing wage. Controls include period dummies, Census division dummies, and lagged CZ characteristics matching Table 3. $N = 1{,}444$; observations weighted by start-of-period CZ population; standard errors clustered at the state level.

Table A.4. Standard vs. Network-Corrected Estimates Across Outcomes (2SLS)

	Mfg Employment	Total Employment	Unemployment	NILF	Log Weekly Wage
Standard Bartik	−0.596	−0.774	0.221	0.553	−0.759
	(0.100)	(0.179)	(0.059)	(0.153)	(0.257)
Network-corrected $\tilde{B}_r$	−0.255	−0.306	0.087	0.219	−0.332
	(0.023)	(0.071)	(0.013)	(0.070)	(0.107)
Observations	1,444	1,444	1,444	1,444	1,444

Notes: Each cell is a separate 2SLS regression with the full control set; state-clustered standard errors in parentheses; population-weighted; $N = 1,444$ CZ-periods. The standard row replicates the ADH exposure and instrument. The corrected row uses the manufacturing-shares network-corrected Bartik of the main text: endogenous $\tilde{B}_r = \mathbf{s}_r' \mathbf{L}^{\text{mfg}} \mathbf{g}^{US}$ built from start-of-period shares, instrumented by the decade-lagged other-country analog $\mathbf{s}_r' \mathbf{L}^{\text{mfg}} \mathbf{g}^{OC}$ (the headline construction of Table 3). A full-economy single-Leontief corrected row is deliberately not reported: the full-economy co-location matrix is dominated by geographically ubiquitous service sectors (the concentration caveat of this appendix), and the reliable full-economy objects are the channel decompositions of Table A.3.

Table A.5. Monte Carlo Exercise D: Correlated Shocks

Within-sector	Standard Bartik		Corrected Bartik		Std/Corr
correlation ρ	Mean	SD	Mean	SD	Ratio
0.0	-1.227	1.469	-0.485	0.637	2.5
0.3	-2.903	6.149	-0.493	0.570	5.9
0.5	-3.241	6.755	-0.519	0.633	6.2
0.7	-3.531	7.227	-0.493	0.663	7.2
0.9	-3.678	7.095	-0.498	0.717	7.4
True β			-0.500	i.i.d. $\phi = 3.01$	

Notes: 1,000 Monte Carlo replications on a synthetic design (random $\mathbf{A}$, its Leontief inverse, and random shares) with $R = 722$ regions and $K = 342$ industries grouped into 10 sectors of ~ 34 industries. Within each sector, shocks have pairwise correlation ρ; across sectors, shocks are independent. The design's i.i.d. contamination factor is $\phi = 3.01$, which differs from the ADH value, so the exercise validates the direction of the shock-correlation effect rather than the ADH magnitudes. The corrected Bartik remains unbiased at every ρ. At $\rho = 0$ the standard/corrected ratio is 2.5, near this design's i.i.d. $\phi = 3.01$ (the small shortfall is finite-sample); as ρ increases the ratio climbs to 7.4 at $\rho = 0.9$, well above that baseline. The pattern is the Monte Carlo analog of the mechanism behind the $\phi^{\text{iid}} \to \phi^{IV}$ amplification in Section 4.4: within-sector correlation causes the realized IV bias to exceed the i.i.d. trace diagnostic.