Chasing Demand

Arnav Deshpande
NYU

Martin Rotemberg
NYU

Sharon Traiberman
NYU

August 2026

Abstract

We quantify the contribution of foreign demand to country-level growth in a multi-sector Eaton–Kortum framework. Log-linearizing the equilibrium around observed trade shares yields reduced-form equations for wages and sectoral revenues that depend on demand shocks, country-level innovations, and trade elasticities. We estimate model primitives using the network granular instrumental variables approach of Chodorow-Reich, Gabaix and Viviano (2026), exploiting how idiosyncratic shocks propagate through trading partners. Applied to world trade tables over 1990–2017, foreign demand reallocation accounts for 0.6% of the variance of decadal sectoral revenue growth. This contribution is concentrated among countries whose production reallocates toward sectors facing rising foreign demand—a margin we call demand tropism. Moving from the least- to the most-tropic countries, the pass-through of foreign demand into revenue growth rises roughly fivefold.

I Introduction

In the Prebisch–Singer tradition, countries that remain specialized in primary commodities struggle to grow, as the price of primary commodities is broadly falling relative to other tradeables (Harvey et al., 2010). We argue that their core insight is more general: countries grow, in part, by producing goods and services whose global demand is expanding (Traiberman and Rotemberg, 2023; Reed, 2024).

We quantify this demand-based rationale for growth using a multi-sector quantitative trade framework. Building on Ricardian models of trade and technology (Eaton and Kortum, 2002; Caliendo and Parro, 2015), we develop a structural decomposition of country growth into supply and demand components. Our goal is to take demand shifts seriously when undertaking a quantitative accounting of growth, using widely-used frameworks (e.g.,

We are grateful to participants at the 2026 Power, Markets, and Strategy in a Changing Global Order conference hosted by the IMF, PIER, and the Bank of Thailand, and in particular to Juthathip Jongwanich for a thoughtful discussion and to Laura Alfaro, Yan Bai, Emine Boz, Jesus Fernandez-Villaverde, Pierre-Olivier Gourinchas, and Nat Tharnpanich for their suggestions and guidance. For any questions or comments, please reach out to us at ad6790@nyu.edu, mr3019@nyu.edu, and st1012@nyu.edu. 19 West 4th Street, 6th Floor. New York, NY 10012.

Levchenko and Zhang, 2016). We find that foreign demand shocks are relevant for domestic growth.

We show not only that demand shocks matter directly, but also that shifting supply toward demand matters. We call this demand tropism, slightly abusing a biological term that means "directional growth in response to external environmental stimulus." In the same way plants benefit from turning their stems towards the sun and their roots towards moisture, we show that countries benefit when the structure of production tilts towards sectors with increasing demand.

We focus on heterogeneous exposure because countries differ markedly in their initial specialization patterns, a long-standing observation in the trade and development literature (Imbs and Wacziarg, 2003; Hidalgo and Hausmann, 2009; Costinot, Donaldson and Komunjer, 2012). Sectoral growth in global demand is also dispersed.

A natural first pass for using demand changes for identification is a shift-share style design: interact variation in what countries make with changes in demand across sectors. Quantifying how this variation maps into revenue growth, however, requires an estimate of the trade elasticities, which govern how demand shifts in one market propagates into producers' revenue elsewhere. That said, the observed comovements of supply and demand load on common global factors such as aggregate productivity cycles, financial cycles, and the diffusion of particular technologies that are unlikely to be orthogonal to a country's own fundamentals (Borusyak, Hull and Jaravel, 2022; Goldsmith-Pinkham, Sorkin and Swift, 2020). Relatedly, when trade linkages are dense, a country and its partners may grow together for other reasons (Manski, 1993).

To identify the trade elasticities in the face of these confounds, we apply the network granular instrumental variables (*network GIV*) approach of Chodorow-Reich, Gabaix and Viviano (2026), which extends the granular instrumental variables idea of Gabaix and Koijen (2024) to settings with general network propagation. The methodology exploits the granularity of the market: while the distribution of shocks is independent of size, a particular shock matters more if it affects a particularly large agent (Gabaix, 2011). The key variation we exploit is that countries do not export to an aggregated global market, but concentrate on particular destinations. We use the structure of the trade model to separate the observed comovement of countries into low-dimensional common components, which are not used for identification, and a residual of idiosyncratic country-level shocks, which are the primitive instruments. Identification rests on the assumption that, after the common components are removed, the remaining shocks are uncorrelated across countries. Given the model, we jointly recover these idiosyncratic shocks and the trade elasticities that govern how they transmit through the global economy. Because the demand shocks are measured directly, the growth consequences of demand follow from the model once the elasticities are in hand.

We implement the network GIV inside a multi-sector Eaton–Kortum economy. Log-linearizing the equilibrium conditions (accounting for changes in demand) around an initial trade equilibrium yields two reduced-form equations. The first is a wage equation in which each country's wage change is a network-weighted combination of all other countries' wage changes, demand shocks in every destination-sector market, and a country-level residual that absorbs productivity and trade-cost shocks. The second is a sectoral revenue equation in which each country-sector's revenue growth depends on the same network of wages and demand shocks, plus a country-sector residual.

The model has a common propagation structure that is a known function of observable trade shares and of the sectoral trade elasticities. We can invert the wage equation at any candidate elasticities to recover country-level innovations. We then purge these innovations of common factors and use each country's leave-one-out innovation to instrument for the propagated wage and demand effects in the revenue equation of its trading partners. Intuitively, we ask whether country i's sectoral growth is predicted, in the cross-section of sectors and over time, by the idiosyncratic wage innovations of *other* countries, with the magnitude of that prediction pinning down the trade elasticities: since the trade elasticities govern the elasticity of bilateral trade shares to relative costs, a foreign wage innovation propagates into a domestic revenue response whose size scales with the trade elasticities, and matching the empirical strength of this mapping to its model counterpart identifies the trade elasticities uniquely. Note that the demand shocks themselves are directly measured.

We estimate a trade elasticity of around 4. We also extend the approach to allow the elasticity to differ between goods and services and find a higher value for goods (around 4) than for services (closer to 2.5), consistent with services trade being less elastic.

The log-linearized revenue equation expresses each country-sector's growth as the sum of a demand component, a wage-innovation component propagated through the trade network, and a residual. We focus on the part of demand that is most plausibly exogenous to the producer—the direct effect of foreign destinations reallocating their expenditure across sectors. Foreign demand shocks are a function of observed sourcing shares and demand shifts alone. We find that foreign demand reallocation accounts for around 0.6% of the variance of sectoral revenue growth over our sample.

The growth effects from changes in foreign demand depend on a country's demand tropism. We measure tropism for each country-decade as the within-country correlation, across sectors, between the change in a sector's revenue share and its foreign-demand component. Foreign demand translates into growth far more strongly for higher tropism countries: moving from the least- to the most-tropic quartile of country-windows, the pass-through of foreign-demand reallocation into revenue growth rises roughly fivefold, from 0.4 to 2.

We emphasize that our results do not speak to the efficiency of tropism or directly to policy proposals. Low tropism may reflect producers unable or unwilling to reallocate resources across production tasks, or it may reflect optimal choices given other features of the economy (Asker, Collard-Wexler and De Loecker, 2014). Suggestively, we do find that policy may play a role in promoting responsiveness to foreign demand, as countries with the highest tropism are those with the fewest restrictions.

Policymakers understand the importance of "making products the world wants" (Albanese, 2024), and in particular the importance of making goods that are in demand in "tomorrow's economy" (von der Leyen, 2023). They emphasize looking toward the "future" when considering industrial policy (Lim, 2023). Policymakers have in particular emphasized this with respect to clean energy, where the U.S. Department of Energy has cast the global transition as a "market opportunity that will have reached at least $23 trillion" by the end of the decade (Granholm, 2022). The same conviction leads resource-rich economies to reposition their production toward this anticipated demand: for instance Indonesia restricted raw nickel exports in order to capture the battery and electric-vehicle value chain, with President Joko Widodo predicting that electric vehicles would soon "replace maybe more than 50% of the existing market demand" (Lim, 2023).

Our paper sits at the intersection of three literatures. First, we contribute to a tradition that uses quantitative trade models to decompose welfare and growth into structural channels (Eaton and Kortum, 2002; Costinot, Donaldson and Komunjer, 2012; Caliendo and Parro, 2015; Levchenko and Zhang, 2016). Relative to this work, our advance is taking demand shifts seriously as a primitive source of variation, and using the network GIV to separately identify supply and demand contributions without resorting to timing assumptions or external instruments.

Second, shift-shares have a long tradition in empirical international trade, as a large body of research has exploited the interaction of (endogenous) sectoral composition with (plausibly exogenous) shifts in demand: at the firm level (di Giovanni, Levchenko and Méjean, 2014; Hummels et al., 2014; Berman, Berthou and Héricourt, 2015; Kramarz, Martin and Méjean, 2020; Esposito, 2022; di Giovanni, Levchenko and Méjean, 2024; Garin and Silvério, 2024; Dhyne et al., 2025) and at subnational units (Góes et al., 2024; Boehm et al., 2026). We build on this work in studying how demand affects national growth.[1]

Third, we contribute to a literature that emphasizes the importance of sectoral compo-

[1] A related literature uses aggregate destination-level shocks to study growth (Arora and Vamvakidis, 2005; Blattman, Hwang and Williamson, 2007; Schmitt-Grohé and Uribe, 2018; Gruss, Nabar and Poplawski-Ribeiro, 2020; Bastos, 2020; Fajgelbaum et al., 2024).

sition for growth and reallocation. Our finding that idiosyncratic demand shifts are quantitatively relevant for growth complements work on demand at the producer level (Hottman, Redding and Weinstein, 2016) and on structural change (Comin, Lashkari and Mestieri, 2021). Our work also complements research on how supply-side forces shape trade flows (Lashkaripour and Lugovskyy, 2023; Bartelme et al., 2025). A long tradition in development shows that the composition of production matters for growth beyond its contribution to current output, whether through the sophistication of the export basket (Hausmann, Hwang and Rodrik, 2007), the reallocation of labor across sectors of differing productivity (McMillan and Rodrik, 2011; Rodrik, 2016), or the difficulties that resource-abundant economies face in diversifying (van der Ploeg, 2011). That work asks which sectors a country should be in. Tropism is a complementary concept, taking demand rather than supply as the moving part: given where global expenditure is heading, which countries move with it?

II Theory

This section develops the framework we take to the data. First, we describe a multi-sector Eaton–Kortum economy, log-linearizing its equilibrium conditions around an initial trade equilibrium. This yields two reduced-form equations—one for country wages, one for sectoral revenues—that are jointly driven by demand shocks, country-level innovations to productivity and trade costs, and sectoral idiosyncratic shocks. Finally, we show how we identify the trade elasticities θ using the cross-country propagation of idiosyncratic wage innovations.

II.A Setup and equilibrium

The world consists of N countries and J sectors. Each country is inhabited by a representative household with labor endowment L_n, paid wage $w_{n,t}$. Preferences over sectoral consumption are Cobb–Douglas with time-varying expenditure shares $\alpha^j_{n,t}$, so that sectoral expenditure is

$$X^j_{n,t} = \alpha^j_{n,t}\, w_{n,t}\, L_n. \tag{1}$$

Within a sector, a continuum of varieties is aggregated with a constant elasticity of substitution. Productivities in sector j for country i are drawn from a Fréchet distribution with location $\lambda^j_{i,t}$ and a sector-specific dispersion parameter θ^j, which is also the sector-j trade elasticity. Exporting variety j from i to n requires an iceberg cost of $d^j_{ni,t} \geq 1$. We collect the elasticities in the vector $\theta \equiv (\theta^1, \ldots, \theta^J)$; the common-elasticity case $\theta^j \equiv \theta$ is nested as a special case. In the empirics we take two parameterizations to the data: a single common elasticity, and a two-group split in which θ^j equals one value for goods sectors and another for services.

Standard arguments (Eaton and Kortum, 2002; Caliendo and Parro, 2015) imply that the sectoral price index in country n is

$$p^j_{n,t} = A^j \left[\sum_i \lambda^j_{i,t} \big(w_{i,t}\, d^j_{ni,t} \big)^{-\theta^j} \right]^{-1/\theta^j}, \tag{2}$$

and that the share of country n's sector-j expenditure allocated to goods from origin i is

$$\pi^j_{ni,t} = \frac{\lambda^j_{i,t} \big(w_{i,t}\, d^j_{ni,t} \big)^{-\theta^j}}{\sum_h \lambda^j_{h,t} \big(w_{h,t}\, d^j_{nh,t} \big)^{-\theta^j}}. \tag{3}$$

Wages are determined by balanced trade: each country's total expenditure equals its total sales,

$$\sum_j \sum_i X^j_{n,t} \pi^j_{ni,t} = \sum_j \sum_i X^j_{i,t} \pi^j_{in,t}. \tag{4}$$

Equations (1)–(4) characterize the equilibrium in each period t given primitives $\{\lambda^j_{i,t}, d^j_{ni,t}, \alpha^j_{n,t}, L_n\}$ and the sector elasticities $\{\theta^j\}$.

We log-differentiate (2)–(4) around an initial equilibrium, treating the productivity shocks $\hat{\lambda}^j_{i,t}$ and trade-cost shocks $\hat{d}^j_{ni,t}$ as unobserved primitives that will be absorbed into a residual at the country and country-sector level. This yields a log-linear system in the equilibrium changes $(\hat{p}^j_n, \hat{\pi}^j_{ni}, \hat{X}^j_n, \hat{w}_n)$:

$$\hat{X}^j_n = \hat{\alpha}^j_n + \hat{w}_n, \tag{5}$$

$$\hat{p}^j_n = \sum_i \pi^j_{ni} \big(\hat{w}_i + \hat{d}^j_{ni} \big) - \tfrac{1}{\theta^j} \sum_i \pi^j_{ni} \hat{\lambda}^j_i, \tag{6}$$

$$\hat{\pi}^j_{ni} = -\theta^j \big(\hat{w}_i + \hat{d}^j_{ni} - \hat{p}^j_n \big) + \hat{\lambda}^j_i \tag{7}$$

The expenditure equation is an accounting identity: under Cobb–Douglas preferences, sectoral expenditure moves one-for-one with the income of the buyer and with the share of income the buyer wishes to allocate to sector j. The change in country n's sector-j price index is a trade-share-weighted average of changes in delivered costs from each origin (wages plus iceberg costs), less an offset for productivity improvements. Country n's import share from origin i responds to relative cost movements: i gains share when its delivered cost falls relative to the sectoral price index, with the strength of the response scaled by $\{\theta^j\}$.

The log-linearized trade-balance condition, after substituting (5) and (7), gives a reduced-form equation for wage changes. Define baseline revenue weights

$$\omega_{in}^{j} \equiv \frac{X_i^j \pi_{in}^j}{\sum_r \sum_h X_h^r \pi_{hn}^r}. \tag{8}$$

The weights ω_{in}^j record the share of country n's total revenue earned by selling sector-j goods to country i. Let $\bar{\omega}_n^j \equiv \sum_i \omega_{in}^j$ denote country n's revenue share in sector j, and $\bar{\theta}_n \equiv \sum_j \theta^j \bar{\omega}_n^j$ its revenue-weighted average trade elasticity. Then

$$\hat{w}_n = \frac{1}{1+\bar{\theta}_n}\left[\sum_j \sum_i \omega_{in}^j \hat{\alpha}_i^j + \sum_j \sum_i \omega_{in}^j \hat{w}_i + \sum_j \sum_i \omega_{in}^j \theta^j \hat{p}_i^j - \sum_j \sum_i \omega_{in}^j \theta^j \hat{d}_{in}^j + \nu_n\right], \tag{9}$$

where ν_n collects the productivity terms $(\hat{\lambda})$ entering through $\hat{\pi}_{in}^j$. A country's wage rises when its export markets shift their expenditure toward sectors in which it sells (the first term), when its trading partners' wages rise (the second term), when sectoral prices rise in its destination markets (the third term), and when it faces lower iceberg trade costs as a seller (the fourth term). The factor $1/(1+\bar{\theta}_n)$ dampens these effects: higher trade elasticities make foreign demand more responsive to prices, attenuating how much of a shock is absorbed in wages. The damping is country-specific, governed by $\bar{\theta}_n$, the average elasticity across the sectors in which n earns its revenue, and collapses to $1/(1+\theta)$ when elasticities are common.

II.B Wages and Revenue

Let $\hat{w}_t \in \mathbb{R}^N$, $\hat{p}_t \in \mathbb{R}^{NJ}$, and $a_t \in \mathbb{R}^{NJ}$ collect the log changes $\hat{w}_{n,t}$, $\hat{p}_{n,t}^j$, and $\hat{\alpha}_{n,t}^j$. Let Π denote the $NJ \times N$ matrix of trade shares, with row (n, j) equal to $(\pi_{n1}^j, \ldots, \pi_{nN}^j)$. Define matrices $A(\theta) \in \mathbb{R}^{N \times NJ}$, $B_w(\theta) \in \mathbb{R}^{N \times N}$, $C_p(\theta) \in \mathbb{R}^{N \times NJ}$ by

$$A_{n,(i,j)} = \frac{\omega_{in}^j}{1+\bar{\theta}_n}, \qquad (B_w)_{ni} = \frac{\sum_j \omega_{in}^j}{1+\bar{\theta}_n}, \qquad (C_p)_{n,(i,j)} = \frac{\omega_{in}^j \theta^j}{1+\bar{\theta}_n}.$$

Stacking equation (9) across countries, substituting $\hat{p}_t = \Pi \hat{w}_t + \delta_t$ from (6), and defining $\Gamma(\theta) \equiv B_w(\theta) + C_p(\theta)\Pi$ gives

$$\big(I - \Gamma(\theta)\big)\hat{w}_t = A(\theta) a_t + \varepsilon_t. \tag{10}$$

The composite country-level innovation $\varepsilon_t \in \mathbb{R}^N$ captures all unobserved primitives at the country level, in particular productivity shocks $\hat{\lambda}_{i,t}^j$ and trade-cost shocks $\hat{d}_{ni,t}^j$. Each country's wage change is a weighted combination of observable demand shocks a_t, other countries' wage changes (through Γ), and a country-specific innovation ε.

Nominal homogeneity implies that the wage system (10) determines wages only up to a common scalar: the rows of $\Gamma(\theta)$ sum to one, so scaling all nominal wages leaves trade shares and all real outcomes unchanged. Rather than singling out a numeraire country, we impose the symmetric normalization that log wage changes are mean zero across countries in every period, and solve the wage block on that $(N-1)$-dimensional subspace. Concretely, with $P \equiv I_N - \frac{1}{N}\mathbf{1}\mathbf{1}'$ the cross-country demeaning projection and Q an orthonormal basis of $\mathbf{1}^\perp$, the wage innovation at candidate θ is

$$\hat{\varepsilon}_t(\theta) \;=\; P\Big[\big(I_N - \Gamma(\theta)\big)\hat{w}_t \;-\; A(\theta)\, a_t\Big], \tag{11}$$

and the wage-propagation matrix is the projected inverse $\Phi_w(\theta) \equiv Q\big(Q'(I_N - \Gamma(\theta))Q\big)^{-1}Q'$.

Sectoral revenue of country i in sector j is $R_i^j = \sum_n X_n^j \pi_{ni}^j$, and its log change $g_i^j \equiv \hat{R}_i^j$ satisfies

$$g_i^j = \sum_n \phi_{ni}^j\big(\hat{X}_n^j + \hat{\pi}_{ni}^j\big), \qquad \phi_{ni}^j \equiv \frac{X_n^j \pi_{ni}^j}{R_i^j}, \tag{12}$$

where the sourcing weights ϕ_{ni}^j record how much of country i's sector-j sales are made to destination n. Substituting the expenditure equation (5) and trade-share equation (7), collecting wage terms, and absorbing productivity and trade-cost terms into a residual η_i^j, equation (12) becomes

$$g_i^j = \underbrace{\sum_n \phi_{ni}^j \hat{\alpha}_n^j}_{\text{demand shifts}} + \underbrace{\sum_n \phi_{ni}^j \hat{w}_n - \theta^j \hat{w}_i + \theta^j \sum_n \phi_{ni}^j \sum_h \pi_{nh}^j \hat{w}_h}_{\text{wage propagation}} + \eta_i^j. \tag{13}$$

Define, for each sector j, the $N \times N$ matrix

$$\tilde{H}_j^R(\theta) \equiv \Phi_j - \theta^j I_N + \theta^j \Phi_j \Pi_j, \tag{14}$$

where $\Phi_j = [\phi_{ni}^j]_{i,n}$ and $\Pi_j = [\pi_{nh}^j]_{n,h}$ is the sector-j block of Π. Stacking across sectors into an $NJ \times N$ matrix $\tilde{H}^R(\theta)$, and letting D be the $NJ \times NJ$ block-diagonal matrix whose j-th block is Φ_j, equation (13) in matrix form is

$$g_t = \tilde{H}^R(\theta)\hat{w}_t + D\, a_t + \eta_t. \tag{15}$$

Substituting $\hat{w}_t$ from the wage reduced form yields

$$g_t = \underbrace{\Big[\tilde{H}^R(\theta)\Phi_w(\theta)A(\theta) + D\Big]}_{M(\theta)}\; a_t + \underbrace{\tilde{H}^R(\theta)\Phi_w(\theta)}_{J_w(\theta)}\; \varepsilon_t + \eta_t, \tag{16}$$

where $\Phi_w(\theta) \equiv (I - \Gamma(\theta))^{-1}$ is the wage propagation matrix.

Equation (16) shows that sectoral revenue growth is the sum of three pieces: a direct response to observed demand shocks propagated through the network ($M(\theta)a_t$), the response to country-level wage innovations propagated through the network ($J_w(\theta)\varepsilon_t$), and a sectoral residual (η_t). We can calculate all of the components of the equation with θ and from additional moments (Π, Φ) that are directly measurable from bilateral trade flows.

The matrix $\tilde{H}^R(\theta)$ satisfies $\tilde{H}^R(\theta)\mathbf{1}_N = \mathbf{1}_{NJ}$: a uniform change in all wages moves every country-sector's predicted revenue growth one for one. Because the level of wages is undetermined, predicted revenue growth is therefore determined only up to a period-specific scalar, which we absorb with a time fixed effect.

II.C Identification

The estimating equations (10) and (16) depend on the trade elasticities. We now show how the trade elasticities are identified from the joint structure of the two equations, given observable $\{g_t, \hat{w}_t, a_t, \Pi_t\}$.

The identification problem. Consider the revenue equation (16) at a candidate θ. Both the coefficient $M(\theta)$ on demand shocks and the coefficient $J_w(\theta)$ on wage innovations are functions of θ through the network propagation. Naive OLS of g_t on a_t fails for two reasons. First, demand shocks a_t are correlated with the sectoral residual η_t through common global factors—business-cycle innovations, technology waves, commodity-price movements—that shift demand and productivity in tandem across countries (Borusyak, Hull and Jaravel, 2022; Goldsmith-Pinkham, Sorkin and Swift, 2020). Second, even after those factors are removed, the wage-propagation term $J_w(\theta)\varepsilon_t$ is unobserved and correlated across countries through the trade network, creating a classical reflection problem (Manski, 1993): a country's revenue comoves with its partners' wages for reasons that the model attributes to θ, but which OLS cannot distinguish from unobserved cross-country spillovers in η.

We address both problems by recovering ε_t from the wage equation and using each country's innovation to instrument for the propagated wage effects in the revenue equation of its trading partners, following the approach developed by Chodorow-Reich, Gabaix and Viviano (2026).

At any candidate θ, equation (11) gives the implied wage innovation $\hat{\varepsilon}_t(\theta)$ as a direct function of observable data. At the true parameter, $\hat{\varepsilon}_t(\theta_0) = \varepsilon_t$ exactly. The innovation ε_t is a composite of structural productivity and trade-cost shocks, and it is in general contaminated by common global factors: a common productivity cycle, for example, will load on every country's innovation. Following Chodorow-Reich, Gabaix and Viviano (2026), we purge these common components by removing principal components from the panel $\{\hat{\varepsilon}_t(\theta)\}_{t=1}^T$, yielding a factor-purged innovation $\tilde{\varepsilon}_t(\theta)$. At θ_0, $\tilde{\varepsilon}_t(\theta_0)$ is the factor-purged true innovation.

Identification ultimately rests on the size of demand: large sector-by-destination combinations generate idiosyncratic innovations that are detectable in aggregate trade flows even after factor removal. A shock to a large trading partner has a larger aggregate effect than an otherwise similar shock to a small trading partner. The leave-one-out instrument we construct below exploits exactly this: it isolates the part of country c's sectoral revenue growth that the model attributes to the propagated wage innovations of the *other* countries in the network. The granular-residual logic of Gabaix and Koijen (2024) is that idiosyncratic shocks to large units survive aggregation. In our setting this is reinforced by trade-share heterogeneity: countries are not symmetric in the network, so even mid-sized economies generate identifying variation when their trading partners have concentrated exposure to them.

Equation (16) says that country i's sectoral revenue growth in sector j depends, through $J_w(\theta)$, on the innovations of *all* countries, including country i's own innovation. To construct a valid instrument for country i, we use only innovations to countries other than i. For each country c and period t, define the leave-one-out innovation vector

$$\tilde{\varepsilon}_t^{(-c)}(\theta) \equiv \tilde{\varepsilon}_t(\theta) - \tilde{\varepsilon}_{c,t}(\theta)\, e_c, \tag{17}$$

where e_c is the standard basis vector zeroing out country c. The instrument for any country-sector $k = (i, j)$ with $i = c$ is the model-implied contribution of non-c innovations to c's revenue growth:

$$z_{k,t}(\theta) = J_w(\theta)(k, :)\, \tilde{\varepsilon}_t^{(-c)}(\theta). \tag{18}$$

The instrument is valid if, after factor removal, innovations and sectoral residuals are independent across countries.

Assumption 1 (Cross-country independence)**.** *Let $\tilde{\varepsilon}_t$ denote the factor-purged wage innovations and $\tilde{\eta}_t$ the factor-purged sectoral residuals. For all countries $c \neq c'$ and all sectors:*

(i) $\mathrm{Cov}(\tilde{\varepsilon}_{c,t}, \tilde{\varepsilon}_{c',t}) = 0$.

(ii) $\mathrm{Cov}(\tilde{\eta}_{k,t}, \tilde{\eta}_{k',t}) = 0$ *whenever country-sector pairs k, k' are in different countries.*

(iii) $\mathrm{Cov}\left(\tilde{\eta}_{k,t}, \tilde{\varepsilon}_{c',t}\right) = 0$ *for every country-sector k located in country c.*

Parts (i) and (ii) require that, after controlling for a low-dimensional set of common global factors, what remains in country-level wage innovations and in country-sector residuals is idiosyncratic. This is a version of the granular-residual assumption familiar from the GIV literature (Gabaix and Koijen, 2024). Part (iii) is the key exclusion restriction. Within the model, a foreign innovation $\varepsilon_{c'}$ affects country c's sectoral growth *only* through the trade network—a channel fully captured by the $J_w(\theta)\varepsilon_t$ term in equation (16). Whatever remains in η after that propagation is by construction orthogonal to foreign innovations. The restriction is substantive: it would fail if foreign country-level shocks propagate to domestic sectoral outcomes through channels outside the model, such as FDI, technology diffusion, migration, or financial linkages.

Under Assumption 1, the instrument satisfies the exclusion restriction at θ_0. Define the residual at candidate θ as

$$r_t(\theta) \equiv g_t - M(\theta)a_t - J_w(\theta)\hat{\varepsilon}_t(\theta), \tag{19}$$

and its factor-purged, time-demeaned counterpart $\tilde{r}_t(\theta)$.

Proposition 1 (Identifying moment)**.** *Under Assumption 1, at $\theta = \theta_0$,*

$$\mathbb{E}\big[z_{k,t}(\theta_0)\,\tilde{r}_{k,t}(\theta_0)\big] = 0 \qquad \textit{for every country-sector } k.$$

Proof. At θ_0, equation (11) gives $\hat{\varepsilon}_t(\theta_0) = \varepsilon_t$ exactly. Substituting into (19) and using (16) yields $r_t(\theta_0) = \eta_t$, so $\tilde{r}_t(\theta_0) = \tilde{\eta}_t$. For $k = (c, j)$, the instrument is a linear combination of $\{\tilde{\varepsilon}_{c',t}\}_{c' \neq c}$. By Assumption 1(iii), each of these is orthogonal to $\tilde{\eta}_{k,t}$. Hence

$$\mathbb{E}[z_{k,t}(\theta_0)\,\tilde{r}_{k,t}(\theta_0)] = \sum_{c' \neq c} J_w(\theta_0)(k, c')\,\mathbb{E}[\tilde{\varepsilon}_{c',t}\,\tilde{\eta}_{k,t}] = 0. \qquad \square$$

Proposition 1 establishes that the moment is zero at θ_0. Identification further requires that it be nonzero away from θ_0. At a candidate $\theta \neq \theta_0$, the recovered innovation $\hat{\varepsilon}_t(\theta)$ is a linear combination of the true innovation and the demand shocks, with coefficients that depend on $\Gamma(\theta_0)$ and $\Gamma(\theta)$. This induces two sources of misfit in the residual $r_t(\theta)$. First, $M(\theta)a_t$ is the wrong linear combination of demand shocks. Second, $J_w(\theta)\hat{\varepsilon}_t(\theta)$ mixes innovations across countries in a way that depends on θ, so the residual at the wrong θ loads on the innovations of countries other than c. The instrument $z_{k,t}(\theta)$ is designed to load precisely on those non-c innovations. The two move together through the trade network in a way that is nonzero whenever $\theta \neq \theta_0$, driving the moment away from zero. Note that we require both the revenue and wage equations in order to identify θ.

II.D Returns to Scale

The baseline model holds sectoral productivity fixed as demand reallocates. A natural concern is that this understates the importance of demand, because with scale economies, demand tropism is additionally associated with a productivity payoff. There are many potential mechanisms that drive increasing returns. Learning by doing accumulates know-how with output (Arrow, 1962); exporting itself may raise productivity through exposure to foreign buyers and technology (Clerides, Lach and Tybout, 1998; De Loecker, 2013; Atkin, Khandelwal and Osman, 2017); and expanding sectors select toward their more productive firms and reallocate factors toward them (Melitz, 2003; Pavcnik, 2002). In lieu of modeling these forces directly, we take a sufficient-statistics type of approach, allowing sectoral productivity to change with sectoral scale, following Kucheryavyy, Lyn and Rodríguez-Clare (2023), which we describe in Appendix B.

III Data

Our main data source is the EORA Multi-Region Input–Output database. Relative to similar sources that also report world trade tables, the advantage of EORA is its wide coverage (with over 150 countries), though the costs are that information on production is relatively aggregated, with 25 consistent production sectors, and some of the data are imputed (Antrás and Chor, 2022; Caliendo and Parro, 2022; Caliendo et al., 2023). In order to have information on wages, we merge EORA to the Penn World Tables, using data on reported labor shares. The estimation sample is $N = 136$ countries (including an RoW aggregate).[2]

In the Penn World Tables, we impute each country's wages as

$$w_{n,t} = \frac{\text{labsh}_{n,t} \cdot R_{n,t}}{\text{emp}_{n,t}},$$

where $\text{labsh}_{n,t}$ is the labor share of GDP from PWT, $R_{n,t} = \sum_j R^j_{n,t}$ is within-sample total revenue, and $\text{emp}_{n,t}$ is total employment in millions of persons engaged.

We construct log changes over rolling ten-year windows: 1990–1999, 1991–2000, and so on through 2008–2017, for a total of nineteen overlapping windows. Working with ten-year horizons rather than year-to-year fluctuations lets us identify from larger movements in demand. We cluster standard errors at the country level to account for the resulting overlap. Our baseline analysis drops seven sectors that are essentially non-tradable: construction, public administration, private households, retail trade, maintenance and repair, electricity, gas and water, and a residual "other" category.

For each year t, destination n, sector j, and origin i, bilateral expenditure $\exp(n, j, i, t)$

[2] We retain countries that have nonmissing labor-share and employment data in PWT for the full sample period and that EORA reports throughout. Thirty EORA countries have consistent employment data in PWT, but not labor shares, and we absorb them into a Rest-of-World (RoW) aggregate using each country's PWT employment and the cross-country median labor share.

is n's total absorption of sector-j goods from i, combining intermediate use by n's industries with final use across four final-demand categories (household, non-profit, government, and gross fixed capital formation). Trade shares are $\pi^j_{ni,t} = \exp(n, j, i, t)/\sum_{i'} \exp(n, j, i', t)$ and total sectoral expenditure is $X^j_{n,t} = \sum_i \exp(n, j, i, t)$. Sectoral revenue is $R^j_{i,t} = \sum_n X^j_{n,t}\pi^j_{ni,t}$, and nominal sectoral revenue growth is $g^j_{i,t} = \log R^j_{i,t_{\text{end}}} - \log R^j_{i,t_{\text{start}}}$.

The demand shock is the log change in country n's final-demand expenditure share on sector j,

$$a^j_{n,t} = \log \alpha^j_{n,t_{\text{end}}} - \log \alpha^j_{n,t_{\text{start}}}, \qquad \alpha^j_{n,t} = \frac{F^j_{n,t}}{\sum_{j'} F^{j'}_{n,t}},$$

with $F^j_{n,t} = \sum_{i \in \text{sample}} \sum_{\text{cat} \in \mathcal{C}} FD_t[(i, j), (n, \text{cat})]$ the total final demand by n for sector-j goods.

III.A Dispersion in Production

We start by showing that countries face heterogeneous exposure to global demand reallocation. Consider a Bartik-style measure that combines initial sectoral revenue shares with the leave-one-out shift in global final demand at the sector level,

$$\text{Bartik}_{n,t} = \sum_j s^j_{n,t_{\text{start}}} \cdot \log\left(\frac{F^j_{-n,t_{\text{end}}}}{F^j_{-n,t_{\text{start}}}}\right),$$

where $s^j_{n,t_{\text{start}}}$ is the share of sector j in country n's total revenue at the start of period t, and $F^j_{-n,\cdot}$ is total global final demand for sector-j goods excluding country n's own purchases.

Figure 2 Panel A shows that there is substantial heterogeneity in global sectoral growth. The interquartile range of sectoral demand growth is around 30 log points. The heterogeneity is both within and between broad sectors, for instance the interquartile range of sectoral demand growth within heavy manufacturing is also around 30 log points. Countries also differ in their specialization. Figure 2 Panel B plots the cosine similarity between each country's production-share vector and the global average: most countries are near 0.9, but a tail of specialized economies falls toward 0.7, with a standard deviation of about 0.1.

Specialization is highly persistent over the sample. Table 1 reports the Pearson correlation between a country-sector's revenue share at year t and its share at year $t + h$, both for raw shares and after subtracting the cross-country average within each sector-year. The raw correlation at a fifteen-year horizon is 0.96, and even after stripping out global sector-year trends it remains 0.87: countries that are idiosyncratically concentrated in a broad sector tend to stay idiosyncratically concentrated.

Because of the persistent dispersion in specialization, there is large dispersion in exposure to the sectoral growth rates. Figure 3 shows many countries specialize in sectors where global demand has grown much faster than average, and, correspondingly, countries whose sectors have been shrinking. The standard deviation of this exposure per decade is roughly 3 log points.

IV Estimation

The GMM estimator combines the wage equation (Equation 10) and the revenue equation (Equation 16). At any candidate θ, the wage equation pins down the structural innovation $\hat{\varepsilon}_t(\theta) = (I - \Gamma(\theta))\hat{w}_t - A(\theta)a_t$ as a deterministic function of observable wages, demand shocks, and trade shares. The revenue equation gives the residual $r_t(\theta) = g_t - M(\theta)a_t - J_w(\theta)\hat{\varepsilon}_t(\theta)$. At $\theta = \theta_0$, $\hat{\varepsilon}_t(\theta_0) = \varepsilon_t$ and $r_t(\theta_0) = \eta_t$. We consider two parameterizations: a single common elasticity ($\theta^j \equiv \theta$), and a goods/services split in which θ^j takes one value across goods sectors and another across services. Each is estimated in two steps: an identity-weighted first step, and a second step that reweights each moment by the inverse of its estimated variance, down-weighting the noisier country-sector moments to improve precision.[3] These appear as the unweighted and weighted columns of the tables below.

For our procedure, we remove one factor from the wage innovations and three from the revenue residuals. The factors stand in for the global forces that move many countries at once, such as productivity and financial cycles and commodity-price swings. We take out more from the revenue panel because it is exposed to an additional layer of comovement: global sectoral waves, which have no counterpart at the country level. We show our resusults are robust to alternative choices.

IV.A First stage

The instrument $z_{k,t}$ is the model-implied contribution of non-c wage innovations to country-sector k's revenue growth. The relevant first stage projects the full network-propagated wage component $J_{w,t}\tilde{\varepsilon}_t(k,:)$ onto $z_{k,t}$, evaluated at θ. Table 2 reports OLS slopes by period and pooled.

To isolate the within-period variation that identifies θ, we subtract period-specific means before pooling. Pooling all windows, the period-demeaned slope of $J_w\tilde{\varepsilon}$ on the leave-one-out instrument is 0.22 (s.e. 0.09, clustered by country). Within macro-periods the slope is 0.21 for windows beginning in the 1990s and 0.40 for the 2000s; the estimate for windows beginning in 2007–2008 is larger (1.95) but uses less data and is correspondingly more imprecise. Figure

[3] Efficient GMM weights the moments by the inverse of their covariance matrix S (Hansen, 1982). With NJ moment conditions and only T periods, the sample analog of S has rank at most $T - 1$ and cannot be inverted, so the fully efficient weighting is infeasible. The diagonal weighting we use is similar in spirit.

1 plots the demeaned binscatter underlying the pooled-demeaned slope. With the exception of a few outlier bins, the relationship is fairly linear.

IV.B Estimate of θ

Our baseline common elasticity is $\hat{\theta} = 4.17$ (Table 3), in the standard range of the Eaton–Kortum literature (e.g. Caliendo and Parro, 2015), and relatively stable across weighting and across sample restrictions on countries and sectors. Allowing the elasticity to differ between goods and services yields a higher value for goods (around 4) than for services (around 2.5), consistent with services being less substitutable across borders.

V Decomposition of Revenue Growth

The revenue equation (16) decomposes sectoral revenue growth into a demand component, a wage-innovation component, and a residual. We focus on the part of demand that is most plausibly exogenous to the producer: changes in expenditure shares of foreign destinations. Split the impact matrix D into D^O, which keeps the own-country entries ϕ^j_{ii} of each block, and $D^F \equiv D - D^O$, which collects the foreign entries ϕ^j_{ni} for $n \neq i$. Revenue growth is then

$$g_t = \underbrace{D^F a_t}_{\text{foreign demand}} + \underbrace{(M - D)\, a_t}_{\text{GE wage channel of demand}} + \underbrace{D^O a_t + J_w\, \varepsilon_t + \eta_t}_{\text{own demand and supply}}. \tag{20}$$

The first term is the direct impact of foreign expenditure reallocation, $D^F a$. It is parameter-free, depending only on observed sourcing weights and demand shifts. The second is the general-equilibrium channel through which all demand shocks, foreign and domestic, move revenue by moving wages. The third collects everything else: own-country expenditure-share shifts, wage innovations, and the sectoral residual.

Before quantifying foreign demand's contribution to revenue growth, we first describe the variation in demand shocks themselves.

V.A The Patterns of Demand Reallocation

Figure 4 plots the distribution of the foreign-demand component $D^F a$. Its standard deviation is around 5% over a decade, with most country-sectors close to zero and a small number facing much larger exposures, both positive and negative. The shocks are broadly centered on zero, since demand tilting toward one sector requires tilting away from another.

Table 5 shows that total direct demand Da accounts for about 10% of the variance of revenue growth, dominated by own-country movements. The foreign component $D^F a$ is much smaller, around 0.6% of the variance, though not negligible. Importantly, the two demand shocks covary, which further contributes to the variance of growth.

V.B Heterogeneity by Sector

The aggregate shares mask substantial variation across sectors. Table 6 computes the within-sector variance shares under our baseline demand shocks, measured from final-demand expenditure shares. Mining and quarrying stands out, with 3.5 percent of its revenue-growth variance attributable to foreign demand movements alone and 18 percent to total demand; wood and paper and fishing follow at 1.3 percent, then textiles (0.9), transport (0.7), and recycling (0.6). At the bottom sit services with little cross-border sourcing—financial intermediation, education and health, post and telecommunications. Table A1 repeats the exercise measuring demand by total-absorption expenditure shares, so that a destination's demand for a sector includes its industries' intermediate purchases. The ranking is similar, but the values are generally smaller as total-absorption is more persistent than final-use composition.

Table 4 aggregates the country-sector cells to the country level, in decadal log points. Qatar's +0.040 means that, averaged across the sample, foreign expenditure shares shifted toward Qatar's sectoral mix by roughly four log points per decade; Belarus's −0.021 is the largest movement in the opposite direction. The top of the distribution is dominated by hydrocarbon and mineral exporters—Qatar, Oman, Iraq, Kuwait, and Venezuela on the energy side, Mongolia and Niger on the mineral side—all with positive values. That demand has on balance shifted toward producers of primary inputs pushes against the classic Prebisch–Singer pessimism that these countries would face a secular decline in relative demand. That said, the countries benefiting the most from demand shocks are those that produce minerals and hydrocarbons, not in agricultural staples.

V.C Foreign Demand and Growth

Table 7 shows how demand shocks matter for realized growth, regressing g on $D^F a$ pooled across country-sector-period observations. The pooled slope is 1.11. A slope above one is consistent with the model: $D^F a$ is positively correlated with the larger own-demand component $D^O a$, so the OLS coefficient absorbs both the direct effect of $D^F a$ and part of the correlated own-demand variation. Figure 5 plots the corresponding binscatter; the relationship is fairly linear throughout the distribution. The pass-through is markedly stronger in richer countries—1.43 for high-income economies, 1.02 for upper-middle income, and 0.76 for lower-middle and low income.

Table 7 shows the direct effect of demand shocks. We can also use the model to study the general-equilibrium effects of demand shocks on revenue. Equation (16) decomposes revenue growth as $g_t = M(\theta)\, a_t + J_w(\theta)\, \varepsilon_t + \eta_t$. The first term, $M(\hat{\theta})\, a$, is the model-implied general equilibrium response to the realized demand shocks, combining each country-sector's direct exposure to demand shocks through the trade network with the indirect adjustment that operates through equilibrium wages.

Table 8 regresses observed revenue growth g on $M(\hat{\theta})\,a$, standardized by its sample-wide standard deviation. Both within and across countries, a one standard deviation increase in $M(\hat{\theta})\,a$ is associated with a 0.19 increase in ten-year log revenue growth for our preferred specification (column 2). The coefficient is similar when looking within continents.

V.D Robustness

Before turning to tropism, we study how robust the results in Table 8 are to alternative modeling and data choices.

Our baseline measure of demand shocks come from final-demand expenditure shares, for which we think foreign changes are plausibly exogenous to domestic productivity innovations. Table A2 instead constructs the shocks from total-absorption expenditure shares, including intermediate purchases. The slope is somewhat larger (0.24 against 0.19) than those in Table 8, but broadly similar.[4]

Appendix B extends the model to sector-specific external economies of scale. We estimate modest scale effects, where doubling a sector's employment raises its productivity by about three percent. Table A3 recomputes the model-implied response $M(\hat{\theta}, \hat{\kappa})a$ under the estimated scale economies and find estimates that are very close to those in Table 8.

Table A4 re-runs the Table 8 for a variety of alternative choices for cleaning the data including changing the sample and removing fewer factors, finding broadly similar effects.

VI Demand Tropism

We call the alignment between a country's production shifts and the foreign demand it faces *demand tropism*. For each country n and rolling ten-year window, tropism is the within-country, cross-sector correlation between the change in a sector's revenue share Δs_n^j and its foreign-demand component $D^F a_n^j$. A value near $+1$ means production reallocated toward the sectors whose foreign demand was rising; a value near -1 means it reallocated toward those whose foreign demand was falling.

Figure 6 plots the distribution of tropism across country-windows: it is centered near zero but spreads widely, from roughly -0.6 to $+0.6$, so countries differ substantially in how closely their production tracks demand. Table 9 ranks countries by mean tropism. Italy (0.41), Trinidad and Tobago, Slovenia, Germany, and Japan sit at the top; Iraq (-0.72), Qatar (-0.68), Angola, and Nigeria sit at the bottom. A negative value is the perverse

[4]The estimates are also more precise but this is partly mechanical, as expenditure inclusive of intermediate purchases and revenue inclusive of intermediate sales comes from the same input-output flows.

case—production shares moving away from the sectors gaining foreign demand and toward those losing it. Several of the countries with the largest aggregate demand gains in Table 4—Qatar, Iraq, and Angola—appear here at the bottom: their trading partners demanded more of their goods, but their production response moved the wrong way.

Tropism is an important mediator of how foreign demand maps into growth. Table 10 splits the regression of g on $D^F a$ by tropism quartile of the country-window cell. The slope rises monotonically from 0.39 in the least-tropic quartile to 1.96 in the most-tropic: a country in the top quartile converts a given foreign-demand shift into roughly five times the revenue growth of one in the bottom quartile. The second panel instead shows an interaction. The interaction of $D^F a$ with tropism has a coefficient of 1.69, larger than the coefficient on the direct effect. Note that tropism is not per se a measure of country success, as it has a small and insignificant direct effect on growth.

Figure 7 shows the four quartile binscatters on common axes. In the least-tropic quartile the relationship between $D^F a$ and growth is nearly flat; in the most-tropic it is steep and roughly linear across the full range of demand shifts. Note that high tropism is not unambiguously good for growth, as it raises sensitivity to demand in both directions. High-tropism countries grow faster when foreign demand moves toward them, and slower when it moves away.

Table 11 reports the share of the variance of g attributable to this amplification channel, $\hat{\beta}_3 \times \text{tropism} \times D^F a$. Across all country-sectors it accounts for 0.24% of the variance of growth, slightly less than half of the direct foreign-demand effect in Table 5 ($\text{var}(D^F a) = 0.6\%$). The contribution is concentrated in the tails of the tropism distribution: 0.73% in the least-tropic quartile and 0.89% in the most-tropic, against essentially zero in the middle two quartiles where production responds neither positively nor negatively to demand.

A natural question is the source of demand tropism. While a full-fledged investigation is beyond the scope of this paper, one natural candidate would be the reallocation of capital after demand changes. To provide preliminary evidence on this channel, we use measures from Fernández et al. (2016) on openness to foreign capital. Using both measures of overall inflow restrictions and specifically equity inflow restrictions, we find in Table A5 that more tropic countries have less restrictive policies.

VII Conclusion

We have used a multi-sector Eaton–Kortum framework, combined with the network granular instrumental variables approach of Chodorow-Reich, Gabaix and Viviano (2026), to quantify the contribution of shifts in global demand to country-level revenue growth. Four patterns emerge from the EORA trade tables. First, countries differ markedly in what they produce, even at the level of broad sectors: cross-country dispersion in sectoral revenue shares is substantial and persistent. Second, destinations differ markedly in how their final-demand expenditure evolves over time, with dispersion in sectoral demand growth that is large both within and across broad sectors. Third, the interaction of these two facts matters quantitatively for sectoral revenue growth at medium horizons. Fourth, how strongly demand translates into growth depends on *demand tropism*—how closely a country's production moves toward the sectors whose foreign demand is rising; the most-tropic countries convert a given foreign-demand shift into several times the revenue growth of the least-tropic, and the amplification runs in both directions, raising exposure to adverse demand shifts as much as to favorable ones.

It is worth emphasizing that these findings are not directly policy relevant. We do not try to document a market failure or a wedge between private and social returns to producing in any particular sector, and our framework is silent on the question of whether countries currently allocate production efficiently across the sectors we study. What our results do suggest is that the sectoral composition of what a country produces interacts with the sectoral composition of what its trading partners want, and that the demand side of this interaction has been a driver of sectoral revenue growth. Policymakers care about the future evolution of the global economy, and our results show that they are correct that this matters for domestic growth.

References

Albanese, Anthony. 2024. "Australian Minerals Industry Parliamentary Dinner." *Speech, Prime Minister of Australia, Canberra.*

Antrás, Pol, and Davin Chor. 2022. "Global value chains." *Handbook of International Economics*, 5: 297–376.

Arora, Vivek, and Athanasios Vamvakidis. 2005. "How Much Do Trading Partners Matter for Economic Growth?" *IMF Staff Papers*, 52(1): 24–40.

Arrow, Kenneth J. 1962. "The Economic Implications of Learning by Doing." *The Review of Economic Studies*, 29(3): 155–173.

Asker, John, Allan Collard-Wexler, and Jan De Loecker. 2014. "Dynamic Inputs and Resource (Mis)Allocation." *Journal of Political Economy*, 122(5): 1013–1063.

Atkin, David, Amit K. Khandelwal, and Adam Osman. 2017. "Exporting and Firm Performance: Evidence from a Randomized Experiment." *The Quarterly Journal of Economics*, 132(2): 551–615.

Bartelme, Dominick, Arnaud Costinot, Dave Donaldson, and Andrés Rodríguez-Clare. 2025. "The Textbook Case for Industrial Policy: Theory Meets Data." *Journal of Political Economy*, 133(5): 1527–1573.

Bastos, Paulo. 2020. "Exposure of Belt and Road Economies to China Trade Shocks." *Journal of Development Economics*, 145: 102474.

Berman, Nicolas, Antoine Berthou, and Jérôme Héricourt. 2015. "Export Dynamics and Sales at Home." *Journal of International Economics*, 96(2): 298–310.

Blattman, Christopher, Jason Hwang, and Jeffrey G. Williamson. 2007. "Winners and Losers in the Commodity Lottery: The Impact of Terms of Trade Growth and Volatility in the Periphery 1870–1939." *Journal of Development Economics*, 82(1): 156–179.

Boehm, Christoph E., Aaron Flaaen, Nitya Pandalai-Nayar, and Jan Schlupp. 2026. "The Local-Area Incidence of Exporting." *Journal of International Economics*, 161: 104257.

Borusyak, Kirill, Peter Hull, and Xavier Jaravel. 2022. "Quasi-Experimental Shift-Share Research Designs." *Review of Economic Studies*, 89(1): 181–213.

Caliendo, Lorenzo, and Fernando Parro. 2015. "Estimates of the Trade and Welfare Effects of NAFTA." *The Review of Economic Studies*, 82(1): 1–44.

Caliendo, Lorenzo, and Fernando Parro. 2022. "Trade Policy." *Handbook of International Economics*, 5: 219–295.

Caliendo, Lorenzo, Robert C. Feenstra, John Romalis, and Alan M. Taylor. 2023. "A Second-Best Argument for Low Optimal Tariffs on Intermediate Inputs." *Journal of International Economics*, 145: 103824.

Chodorow-Reich, Gabriel, Xavier Gabaix, and Davide Viviano. 2026. "Propagation of Shocks in Networks: Identification and Applications." Working Paper.

Clerides, Sofronis K., Saul Lach, and James R. Tybout. 1998. "Is Learning by Exporting Important? Micro-Dynamic Evidence from Colombia, Mexico, and Morocco." *The Quarterly Journal of Economics*, 113(3): 903–947.

Comin, Diego, Danial Lashkari, and Martí Mestieri. 2021. "Structural Change with Long-Run Income and Price Effects." *Econometrica*, 89(1): 311–374.

Costinot, Arnaud, Dave Donaldson, and Ivana Komunjer. 2012. "What Goods Do Countries Trade? A Quantitative Exploration of Ricardo's Ideas." *Review of Economic Studies*, 79(2): 581–608.

De Loecker, Jan. 2013. "Detecting Learning by Exporting." *American Economic Journal: Microeconomics*, 5(3): 1–21.

Dhyne, Emmanuel, Ayumu Ken Kikkawa, Toshiaki Komatsu, Magne Mogstad, and Felix Tintelnot. 2025. "Firm Responses and Wage Effects of Foreign Demand Shocks with Fixed Labor Costs and Monopsony." *American Economic Review*, 115(12): 4328–4368.

di Giovanni, Julian, Andrei A. Levchenko, and Isabelle Méjean. 2014. "Firms, Destinations, and Aggregate Fluctuations." *Econometrica*, 82(4): 1303–1340.

di Giovanni, Julian, Andrei A. Levchenko, and Isabelle Méjean. 2024. "Foreign Shocks as Granular Fluctuations." *Journal of Political Economy*, 132(2): 391–433.

Eaton, Jonathan, and Samuel Kortum. 2002. "Technology, Geography, and Trade." *Econometrica*, 70(5): 1741–1779.

Eora. 2026. "Frequently Asked Questions."

Esposito, Federico. 2022. "Demand Risk and Diversification through International Trade." *Journal of International Economics*, 135: 103562.

Fajgelbaum, Pablo, Pinelopi Goldberg, Patrick Kennedy, Amit Khandelwal, and Daria Taglioni. 2024. "The U.S.–China Trade War and Global Reallocations." *American Economic Review: Insights*, 6(2): 295–312.

Fernández, Andrés, Michael W. Klein, Alessandro Rebucci, Martin Schindler, and Martín Uribe. 2016. "Capital Control Measures: A New Dataset." *IMF Economic Review*, 64(3): 548–574.

Gabaix, Xavier. 2011. "The Granular Origins of Aggregate Fluctuations." *Econometrica*, 79(3): 733–772.

Gabaix, Xavier, and Ralph S. J. Koijen. 2024. "Granular Instrumental Variables." *Journal of Political Economy*, 132(7): 2274–2303.

Garin, Andrew, and Filipe Silvério. 2024. "How Responsive Are Wages to Firm-Specific Changes in Labour Demand? Evidence from Idiosyncratic Export Demand Shocks." *Review of Economic Studies*, 91(3): 1671–1710.

Góes, Carlos, Juan Segnana, Raymond Robertson, and Gladys Lopez-Acevedo. 2024. "Local Labor Market Dynamics and Export Shocks: Theory and Evidence from Indonesia." Working Paper.

Goldsmith-Pinkham, Paul, Isaac Sorkin, and Henry Swift. 2020. "Bartik Instruments: What, When, Why, and How." *American Economic Review*, 110(8): 2586–2624.

Granholm, Jennifer M. 2022. "Remarks Opening the GCEAF: Accelerating the Clean Energy Transition." *Speech, Global Clean Energy Action Forum, U.S. Department of Energy, Pittsburgh.*

Gruss, Bertrand, Malhar Nabar, and Marcos Poplawski-Ribeiro. 2020. "Growth Accelerations and Reversals in Emerging Market and Developing Economies: External Conditions and Domestic Amplifiers." *Open Economies Review*, 31: 753–786.

Hansen, Lars Peter. 1982. "Large Sample Properties of Generalized Method of Moments

Estimators." *Econometrica*, 50(4): 1029–1054.

Harvey, David I., Neil M. Kellard, Jakob B. Madsen, and Mark E. Wohar. 2010. "The Prebisch-Singer Hypothesis: Four Centuries of Evidence." *The Review of Economics and Statistics*, 92(2): 367–377.

Hausmann, Ricardo, Jason Hwang, and Dani Rodrik. 2007. "What You Export Matters." *Journal of Economic Growth*, 12(1): 1–25.

Hidalgo, César A., and Ricardo Hausmann. 2009. "The Building Blocks of Economic Complexity." *Proceedings of the National Academy of Sciences*, 106(26): 10570–10575.

Hottman, Colin J., Stephen J. Redding, and David E. Weinstein. 2016. "Quantifying the Sources of Firm Heterogeneity." *The Quarterly Journal of Economics*, 131(3): 1291–1364.

Hummels, David, Rasmus Jørgensen, Jakob R. Munch, and Chong Xiang. 2014. "The Wage Effects of Offshoring: Evidence from Danish Matched Worker-Firm Data." *American Economic Review*, 104(6): 1597–1629.

Imbs, Jean, and Romain Wacziarg. 2003. "Stages of Diversification." *American Economic Review*, 93(1): 63–86.

Kramarz, Francis, Julien Martin, and Isabelle Méjean. 2020. "Volatility in the Small and in the Large: The Lack of Diversification in International Trade." *Journal of International Economics*, 122: 103276.

Kucheryavyy, Konstantin, Gary Lyn, and Andrés Rodríguez-Clare. 2023. "Grounded by Gravity: A Well-Behaved Trade Model with Industry-Level Economies of Scale." *American Economic Journal: Macroeconomics*, 15(2): 372–412.

Lashkaripour, Ahmad, and Volodymyr Lugovskyy. 2023. "Profits, Scale Economies, and the Gains from Trade and Industrial Policy." *American Economic Review*, 113(10): 2759–2808.

Levchenko, Andrei A, and Jing Zhang. 2016. "The evolution of comparative advantage: Measurement and welfare implications." *Journal of Monetary Economics*, 78: 96–111.

Lim, Lionel. 2023. "Indonesia's Decision to Ban Nickel Ore Exports Was a Surprising Success for President Jokowi. Will It Work for His Successors?" Fortune.

Manski, Charles F. 1993. "Identification of Endogenous Social Effects: The Reflection Problem." *Review of Economic Studies*, 60(3): 531–542.

McMillan, Margaret S., and Dani Rodrik. 2011. "Globalization, Structural Change and Productivity Growth." National Bureau of Economic Research NBER Working Paper 17143.

Melitz, Marc J. 2003. "The Impact of Trade on Intra-Industry Reallocations and Aggregate Industry Productivity." *Econometrica*, 71(6): 1695–1725.

Pavcnik, Nina. 2002. "Trade Liberalization, Exit, and Productivity Improvements: Evidence from Chilean Plants." *The Review of Economic Studies*, 69(1): 245–276.

Reed, Tristan. 2024. "Export-Led Industrial Policy for Developing Countries: Is There a Way to Pick Winners?" *Journal of Economic Perspectives*, 38(4): 3–26.

Rodrik, Dani. 2016. "Premature Deindustrialization." *Journal of Economic Growth*, 21(1): 1–33.

Schmitt-Grohé, Stephanie, and Martín Uribe. 2018. "How Important Are Terms-of-

Trade Shocks?" *International Economic Review*, 59(1): 85–111.

Traiberman, Sharon, and Martin Rotemberg. 2023. "Precautionary Protectionism." *Journal of International Economics*, 145: 103836.

van der Ploeg, Frederick. 2011. "Natural Resources: Curse or Blessing?" *Journal of Economic Literature*, 49(2): 366–420.

von der Leyen, Ursula. 2023. "Special Address by President von der Leyen at the World Economic Forum." *Speech, World Economic Forum Annual Meeting, Davos.*

Figure 1: First Stage: Foreign Shocks And Domestic Revenue

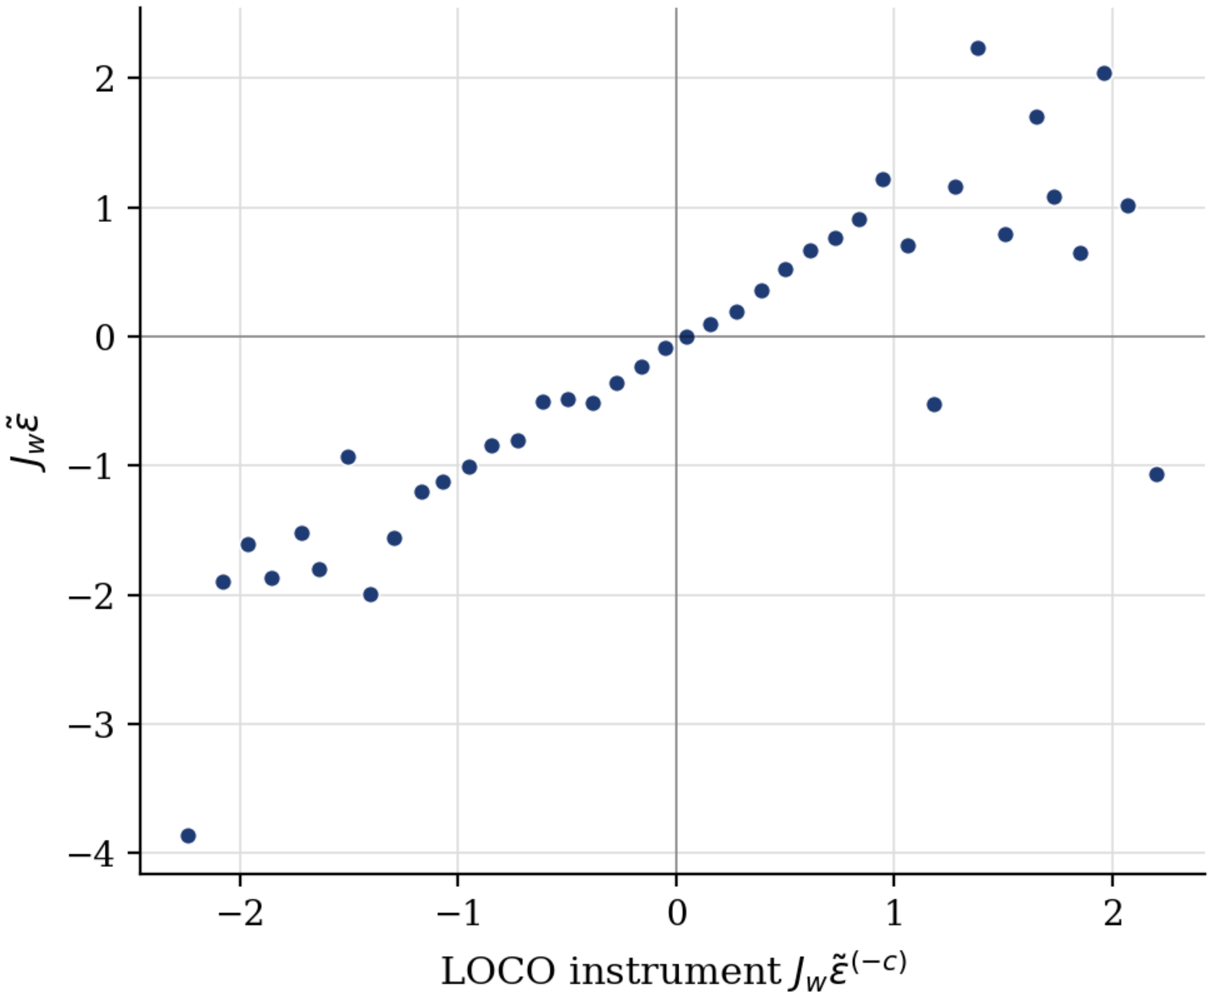

Notes: Pooled OLS regression of $J_w\,\tilde{\varepsilon}_t$ on the leave-one-out instrument $J_w\,\tilde{\varepsilon}_t^{(-c)}$, after subtracting period-specific means from both variables. $\tilde{\varepsilon}_t$ is the country-level wage shock recovered from the wage equation and purged of common time-series factors; J_w is the matrix mapping wage shocks to country-sector revenue growth at the estimated θ, so $J_w\,\tilde{\varepsilon}_t$ is the model-implied revenue effect of the realized wage shocks. The leave-one-out instrument $J_w\,\tilde{\varepsilon}_t^{(-c)}$ zeroes out country c's own wage shock before applying J_w, so the value at observation (c, j, t) depends only on other countries' wage shocks. Each observation is a country-sector-period.

Figure 2: Sectoral Demand Growth and Country Specialization

Panel A: Global sectoral demand growth

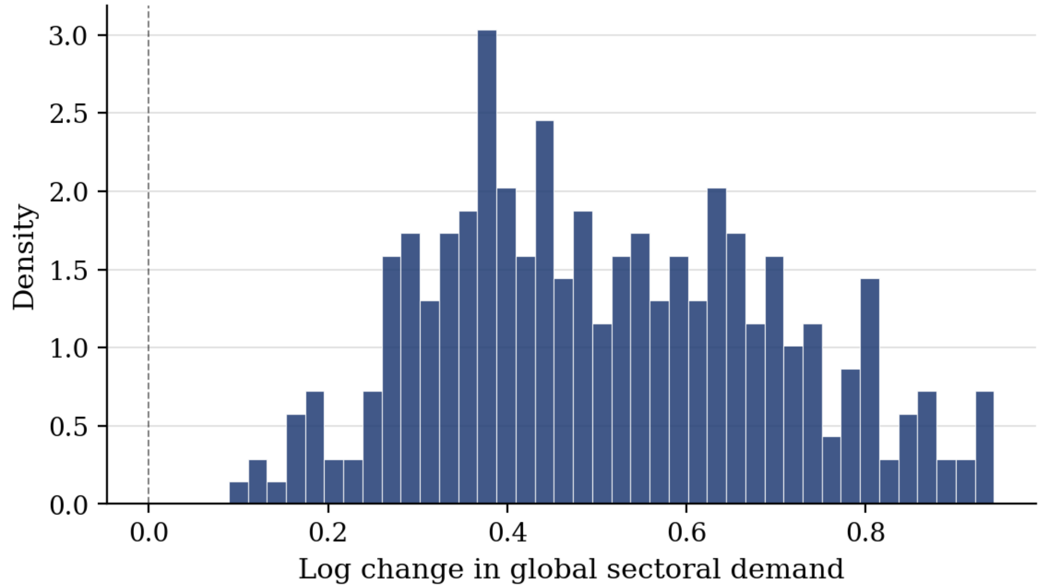

Panel B: Cosine similarity to global production

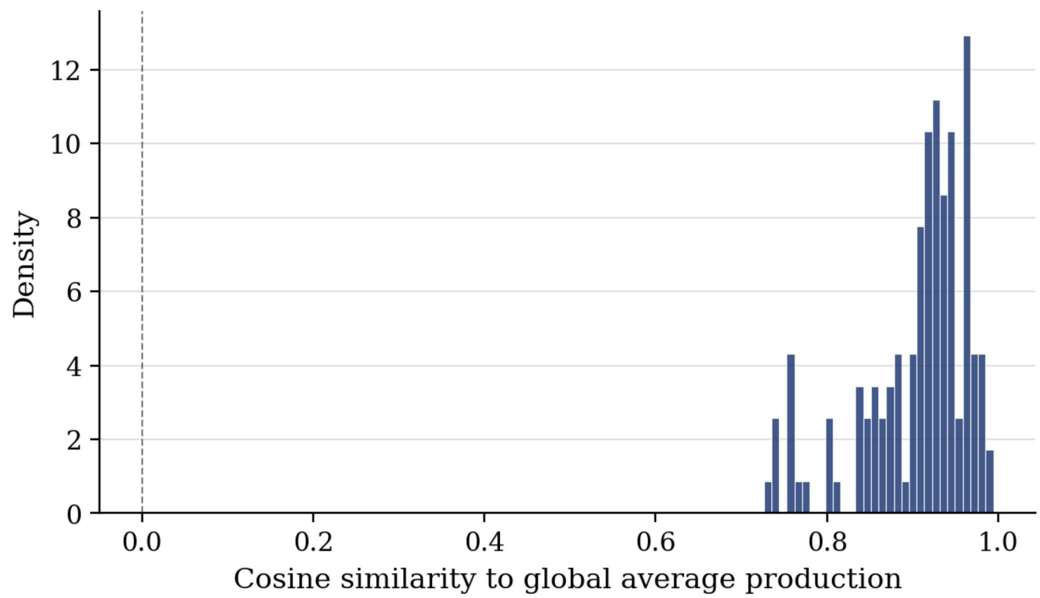

Notes: Panel A shows the distribution of log changes in global final demand for each sector across all rolling ten-year windows. Panel B shows the cosine similarity between each country's sectoral production-share vector and the global average, measured at the start of the sample.

Figure 3: Distribution of Bartik Exposure to Global Demand

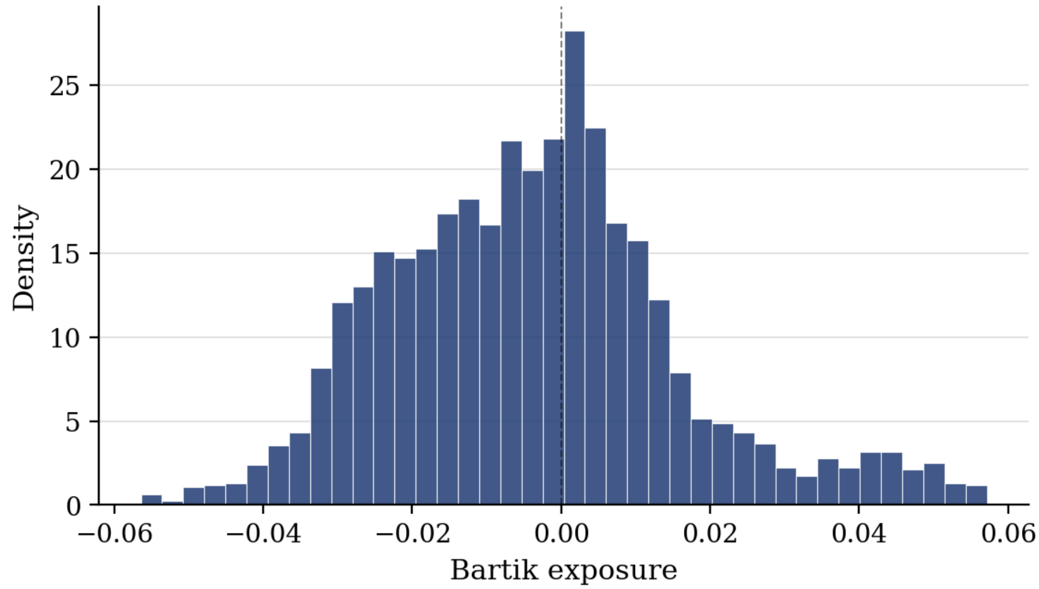

Notes: Each country's exposure is computed period by period using initial sectoral revenue shares as weights on the leave-one-out log change in global final demand for each sector. Initial revenue shares are constructed from the EORA multi-region input-output tables at the start of each rolling ten-year window. The figure pools the period-demeaned values across three macro-periods: 1990–1999, 2000–2009, and 2010–2017.

Figure 4: Distribution of Foreign Demand $D^F a$ Across Country-Sectors

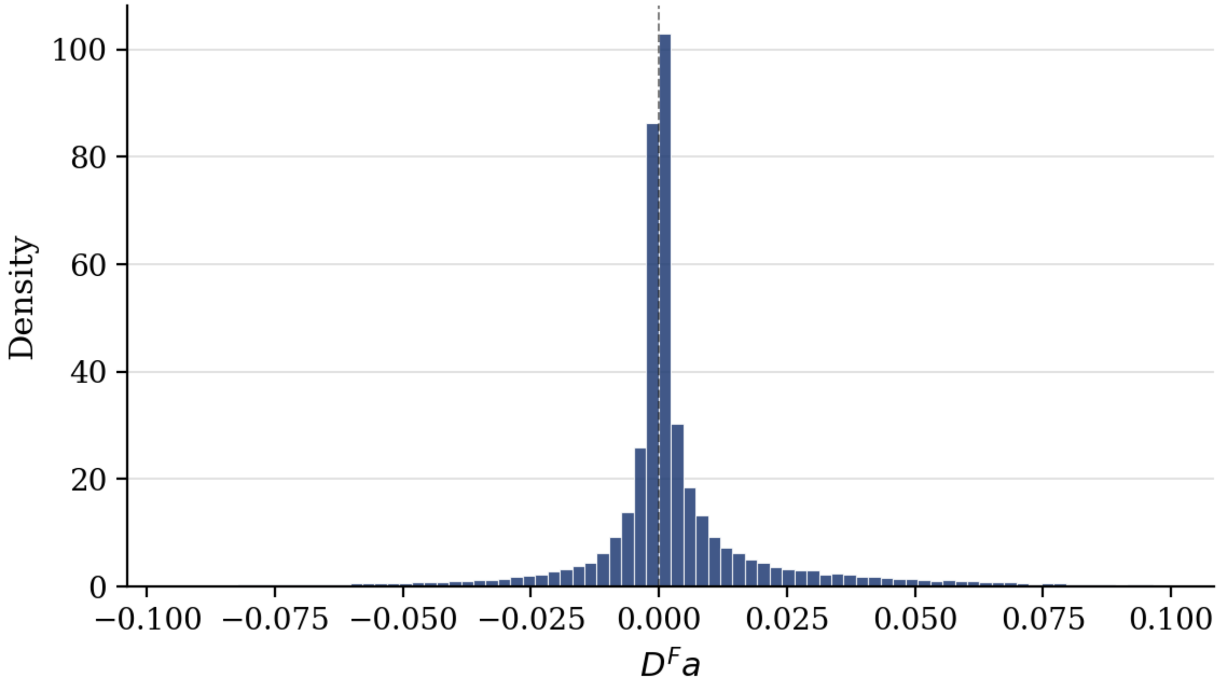

Notes: Cross-sectional histogram of $D^F a = \sum_{n \neq i} \phi_{ni}^j \hat{\alpha}_n^j$ across country-sectors, pooled across rolling ten-year windows.

Figure 5: Foreign Demand and Revenue Growth

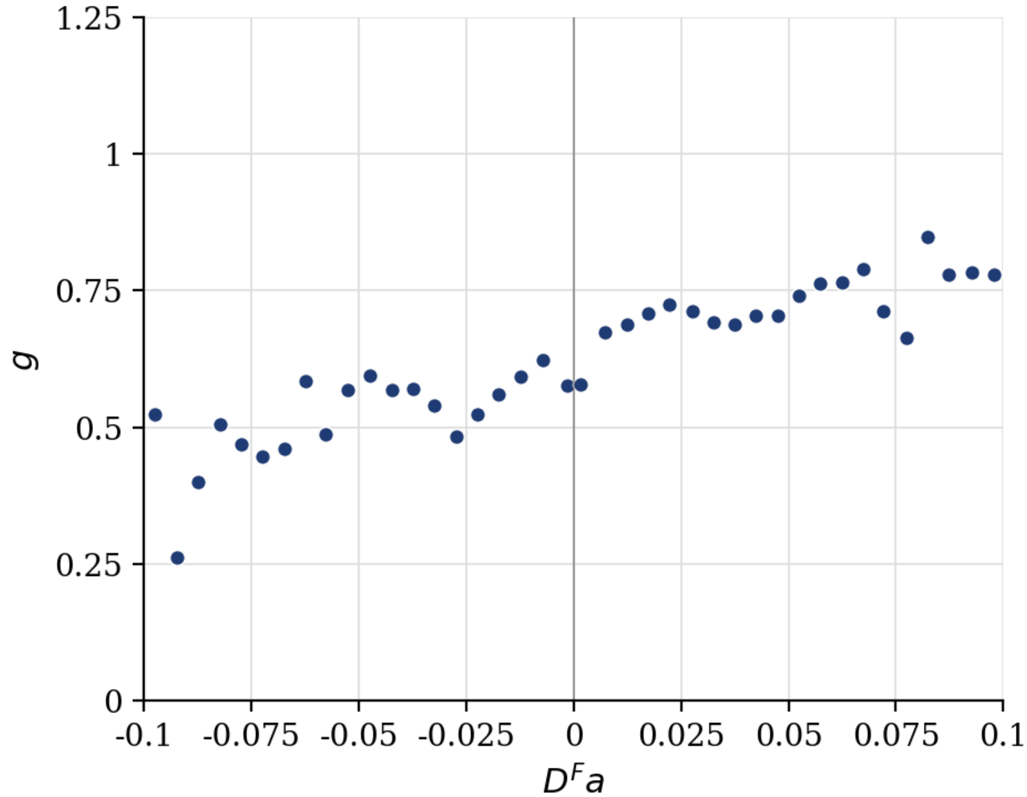

Notes: Binscatter of revenue growth g on the foreign-demand component $D^F\, a = \sum_{n \neq i} \phi^j_{ni} \hat{\alpha}^j_n$, pooled across all country-sector-window observations.

Figure 6: Distribution of Production Tropism Across Country-Windows

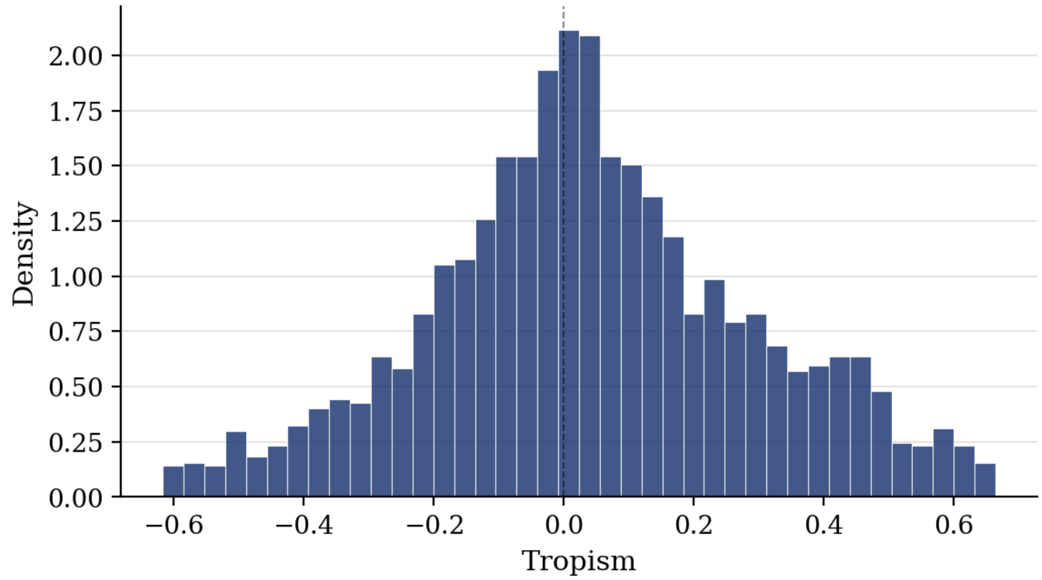

Notes: For each country and rolling ten-year window, tropism is the within-country, cross-sector correlation between the change in revenue share Δs_n^j and the foreign-demand component $D^F a_n^j$. Pooled across windows, 1990–1999 through 2008–2017.

Figure 7: Foreign Demand and Revenue Growth, by Tropism Quartile

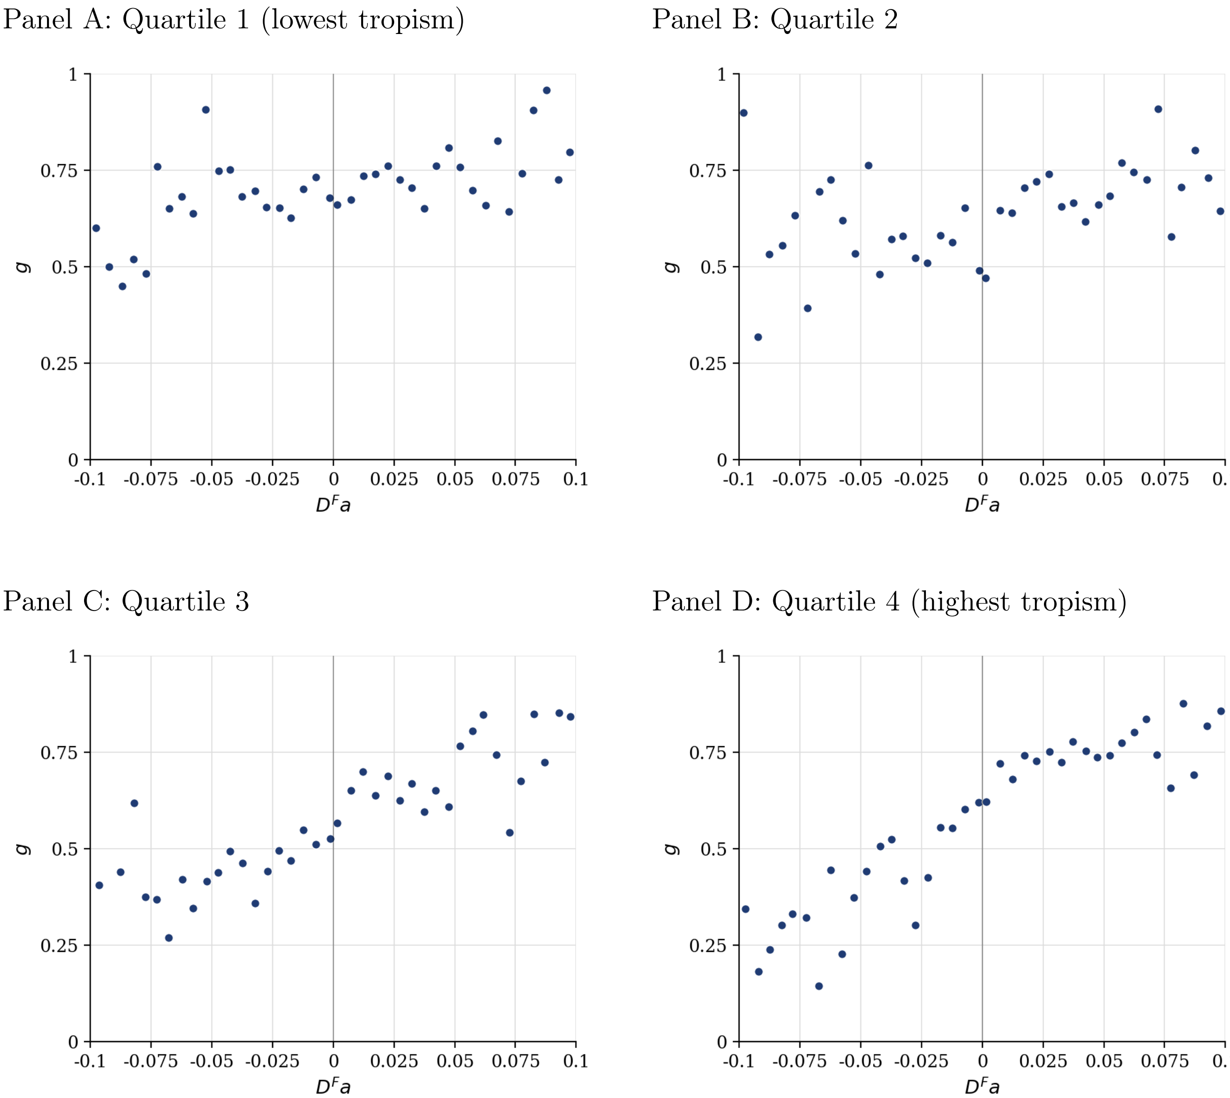

Notes: Binscatter of revenue growth g on the foreign-demand component $D^F a$, splitting the sample by tropism quartile of the country-window cell. Axes are common across panels.

Table 1: Persistence of Country-Sector Revenue Shares

Horizon	Raw shares	Sector-year demeaned
1 year	0.997	0.989
5 years	0.986	0.956
10 years	0.973	0.915
15 years	0.958	0.870

Notes: Pearson correlation of country-sector revenue shares (sector j as a share of country n's total revenue) at year t with shares at year $t+h$, pooled across all country-sectors and all valid $(t, t+h)$ pairs in the 1990–2008 window. Column 1 reports raw correlations. Column 2 first subtracts the cross-country average within each (sector, year) cell.

Table 2: First Stage: Wage-Shock Revenue Effects Regressed on the Leave-One-Out Instrument

Period	$\hat{\beta}$	SE	N
Pooled (demeaned)	0.222	(0.090)	46,512
1990s	0.213	(0.086)	24,480
2000s	0.403	(0.349)	17,136
2010s	1.947	(0.624)	4,896

Notes: OLS regression of $J_w\,\tilde{\varepsilon}_t$ on the LOCO instrument $J_w\,\tilde{\varepsilon}_t^{(-c)}$, where country c's wage innovation is zeroed out before J_w is applied. Each observation is a country-sector-period. The pooled regression stacks all periods; subtracts period-specific means before pooling. The 1990s row averages bins starting in the 1990s, the 2000s row bins starting in the 2000s, and the 2010s row covers bins starting in 2007.

Table 3: Estimates of θ

Sector group	$\hat{\theta}$ (unweighted)	$\hat{\theta}$ (weighted)
Combined	5.72	4.17
Goods	4.509	3.919
Services	2.467	2.449

Notes: GMM estimates of the trade elasticity θ from rolling ten-year windows, 1990–1999 through 2008–2017. In the first row we estimate one θ for the entire economy. The unweighted column reports identity-weighted GMM estimates; the weighted column applies a diagonal weighting matrix from the unweighted estimates. For the second and third row, we allow θ to differ between goods-producing and services-producing sectors. Goods sectors are agriculture, mining, food and beverage, and manufacturing; services sectors are the remaining tradeable sectors.

Table 4: Countries with Largest Aggregate Foreign-Demand Reallocation

Country	Revenue-weighted mean $D^F a$
Qatar	0.0396
Mongolia	0.0337
Oman	0.0245
Iraq	0.0194
Niger	0.0172
Kuwait	0.0153
Trinidad and Tobago	0.0136
Venezuela	0.0092
Macao	0.0075
Angola	0.0075
Largest negative:	
Belarus	-0.0208
Canada	-0.0090
Malaysia	-0.0086
Singapore	-0.0080
Mexico	-0.0073
Fiji	-0.0073
Bolivia	-0.0064
New Zealand	-0.0061
Taiwan	-0.0056
Chile	-0.0055

Notes: For each country i and ten-year window, $\bar{D}^F a_{i,t} = \sum_j s^j_{i,t} D^F a^j_{i,t}$ is the revenue-weighted average across sectors of the foreign-demand reallocation effect, where $s^j_{i,t}$ is country i's sectoral revenue share at the start of the window. Reported values are averages across rolling ten-year windows from 1990–1999 through 2008–2017, in decadal log points. The table reports the ten countries with the largest positive values and the ten with the largest negative values.

Table 5: Demand and the Variance of Growth

	Share of var(g)
var($D^O a$)	9.6%
var($D^F a$)	0.6%
$2\,\mathrm{cov}(D^O a, D^F a)$	0.4%
var(Da) = var($D^O a + D^F a$)	10.6%

Notes: Each row reports the corresponding statistic as a share of var(g), the variance of revenue growth across all country-sector-period observations. $D^O\,a = \phi^j_{ii}\,\hat{\alpha}^j_i$ is the direct effect of own-country expenditure-share reallocation on producer revenue. $D^F\,a = \sum_{n \neq i} \phi^j_{ni}\,\hat{\alpha}^j_n$ is the direct effect of foreign expenditure-share reallocation. Their sum $D\,a$ is total direct demand, and the relationship $\mathrm{var}(D\,a) = \mathrm{var}(D^O\,a) + \mathrm{var}(D^F\,a) + 2\,\mathrm{cov}(D^O\,a, D^F\,a)$ holds exactly across rows. Pooled across rolling ten-year windows, 1990–1999 through 2008–2017.

Table 6: Demand-Driven Variance of Revenue Growth, by Sector

Sector	$\text{var}(D^F a)/\text{var}(g)$	$\text{var}(Da)/\text{var}(g)$
Mining and Quarrying	3.5%	17.9%
Wood and Paper	1.3%	13.5%
Fishing	1.3%	22.0%
Textiles and Wearing Apparel	0.9%	7.6%
Transport	0.7%	12.4%
Recycling	0.6%	6.7%
Petroleum, Chemical and Non-Metallic Mineral Products	0.4%	6.2%
Other Manufacturing	0.3%	8.3%
Transport Equipment	0.2%	5.5%
Metal Products	0.2%	7.3%
Wholesale Trade	0.2%	16.9%
Electrical and Machinery	0.1%	5.8%
Agriculture	0.1%	12.8%
Food & Beverages	0.05%	3.6%
Hotels and Restaurants	0.04%	4.9%
Post and Telecommunications	0.02%	9.6%
Financial Intermediation and Business Activities	0.01%	14.2%
Education, Health and Other Services	0.01%	8.3%

Notes: Variance shares of revenue growth g attributable to total demand $D\,a$ and to its foreign-only component $D^F\,a$, computed within each sector. Sectors are sorted in descending order by foreign-demand variance share. Pooled across rolling ten-year windows from 1990–1999 through 2008–2017.

Table 7: Foreign Demand Shocks and Growth

Group	$\hat{\beta}$	SE	N
All countries	1.111	(0.113)	46,512
High income	1.425	(0.202)	18,468
Upper-middle income	1.022	(0.254)	12,996
Lower-middle and low income	0.758	(0.139)	15,048

Notes: OLS regression of revenue growth g on the foreign-demand component $D^F a$, pooled across rolling ten-year windows from 1990–1999 through 2008–2017. Each observation is a country-sector-window. The first row pools all countries; the next three rows split the sample by World Bank FY2025 income classification. Standard errors in parentheses are clustered at the country level.

Table 8: Revenue Growth and the Model-Implied GE Response to Demand

	(1)	(2)	(3)	(4)
Panel A: Common parameters				
Ma (standardized)	0.2059	0.1945	0.2016	0.1864
	(0.0319)	(0.0347)	(0.0273)	(0.0288)
Panel B: Separate parameters for Goods and Services				
Ma (standardized)	0.2046	0.1929	0.2006	0.1851
	(0.0303)	(0.0330)	(0.0256)	(0.0269)
Observations	46,512	46,512	46,512	46,512
Countries	136	136	136	136
Country FE	No	Yes	No	Yes
Continent × year FE	No	No	Yes	Yes

Notes: OLS regressions of revenue growth g on the model-implied GE response $M(\hat{\theta})\,a$ from Equation (16), evaluated at the estimated trade elasticity. Panel A uses a common $\hat{\theta}$ estimated across all sectors; Panel B uses separate $\hat{\theta}$ estimates for goods-producing and services-producing sectors. In each panel, $M(\hat{\theta})\,a$ is scaled by its sample-wide standard deviation, so the reported coefficient is the change in g associated with a one-standard-deviation increase in the model-implied response. Each observation is a country-sector-window. Standard errors in parentheses are clustered at the country level. Pooled across rolling ten-year windows from 1990–1999 through 2008–2017.

Table 9: Countries by Tropism

Country	Mean tropism	T
Italy	0.409	19
Trinidad and Tobago	0.400	19
Slovenia	0.399	19
Germany	0.379	19
Japan	0.374	19
Belarus	0.333	19
Tajikistan	0.324	19
Slovakia	0.311	19
Belgium	0.298	19
Fiji	0.295	19
Bottom:		
Iraq	-0.722	19
Qatar	-0.682	19
Angola	-0.549	19
Nigeria	-0.400	19
Bolivia	-0.289	19
North Macedonia	-0.264	19
Sri Lanka	-0.253	19
Dominican Republic	-0.242	19
El Salvador	-0.227	19
Laos	-0.227	19

Notes: For each country, mean tropism is averaged across rolling ten-year windows, 1990–1999 through 2008–2017. The table reports the ten countries with the highest mean tropism and the ten with the lowest.

Table 10: Tropism and Demand Sensitivity

	$\hat{\beta}$	SE	N
Q1 (least tropic)	0.392	(0.164)	11,628
Q2	0.507	(0.154)	11,628
Q3	1.213	(0.196)	11,628
Q4 (most tropic)	1.957	(0.158)	11,628
Pooled interaction:			
Constant	0.606	(0.023)	46,512
$D^F a$	1.046	(0.103)	
Tropism	-0.089	(0.071)	
$D^F a \times$ Tropism	1.688	(0.272)	

Notes: Panel A reports OLS regressions of revenue growth g on the foreign-demand component $D^F a$ within each tropism quartile of the country-window cell. Panel B reports a pooled regression of g on $D^F a$, tropism, and their interaction; the reported N applies to the entire pooled regression. Standard errors in parentheses are clustered at the country level. Pooled across rolling ten-year windows from 1990–1999 through 2008–2017.

Table 11: Variance Share of the Tropism Channel

Group	N	$\text{var}(\hat{\beta}_3 \cdot \text{Trop} \cdot D^F a)/\text{var}(g)$
All observations	46,512	0.24%
Q1 (least tropic)	11,628	0.73%
Q2	11,628	0.00%
Q3	11,628	0.02%
Q4 (most tropic)	11,628	0.89%

Notes: Variance share of revenue growth g attributable to the tropism channel, $\hat{\beta}_3 \times \text{tropism}_{n,t} \times D^F a^j_{n,t}$, where $\hat{\beta}_3$ is the interaction coefficient from the pooled regression in Table 10. Reported overall and by tropism quartile. Pooled across rolling ten-year windows from 1990–1999 through 2008–2017.

Appendix

A Estimation and Computation

This appendix discusses our implementation choices.

A.A Normalization of the wage block

The model is homogeneous in nominal wages: the rows of $\Gamma(\theta)$ sum to one, so $(I-\Gamma)\mathbf{1}_N = 0$ and the wage system (10) determines $\hat{w}_t$ only up to a common scalar. Some normalization is required before the system can be solved. We impose that log wage changes are mean zero across countries in every period. Operationally, with $P \equiv I_N - \frac{1}{N}\mathbf{1}\mathbf{1}'$ the cross-country demeaning projection and Q any $N \times (N-1)$ orthonormal basis of $\mathbf{1}^{\perp}$, the wage innovation is the demeaned residual

$$\hat{\varepsilon}_t(\theta) \;=\; P\big[(I-\Gamma(\theta))\,\hat{w}_t - A(\theta)\,a_t\big],$$

and every inverse of $I-\Gamma$ (and, in Appendix B, of $I-\Gamma^s$) is the projected inverse

$$\Phi_w(\theta) \;\equiv\; Q\big(Q'(I-\Gamma(\theta))\,Q\big)^{-1}Q',$$

which is invariant to the choice of basis Q. Under this normalization no country plays a special role: the estimates are exactly equivariant to relabeling of countries (verified numerically on the estimation data), and the leave-one-out instrument treats every country symmetrically—country c's own innovation is zeroed when constructing c's instruments, for every c. Note that, due to the features of the leave-one-out instrument, had we instead normalized countries to a particular base country, then our results *would* be sensitive to the choice of numeraire.

A.B Residuals, factor removal, and weighting

At each candidate parameter the residuals are computed from the structural equations using the normalization above. The revenue residual is demeaned within each period, absorbing the undetermined level. Global comovement is then removed from both panels by principal components: we subtract the projection onto the top $K_\varepsilon = 1$ left singular vector of the wage-innovation panel and the top $K_r = 3$ left singular vectors of the revenue-residual panel. The purged innovations enter the instrument (18) through the model loading J_w, with the own country's entry zeroed.

The moment vector stacks one condition per country-sector, $\mathbb{E}[z_{k,t}\,\tilde{\eta}_{k,t}] = 0$, giving $NJ =$ 2,448 moments from $T = 19$ windows. A full efficient weighting matrix is infeasible at these dimensions—the sample covariance of the moments has rank at most $T-1$—so the second step uses diagonal weights, the reciprocal of each moment's time-series variance evaluated at the first-step estimate, normalized to mean one. The first step uses identity weights.

A.C Parameterization and optimization

The θ-only objectives are one-dimensional and smooth; each step is minimized by golden-section search on $[1, 15]$. The multi-parameter objectives are minimized by Nelder–Mead in transformed coordinates that keep the search unconstrained in particular $\log\theta^j$ for the trade elasticities. The optimizer is started from a grid of points spanning the interior and the near-boundary region. Candidates at which the projected wage solve is ill conditioned, or at which a parameter leaves its admissible region, receive a large penalty, so the optimizer cannot exploit near-singular configurations.

A.D Data quality and robustness

EORA's principal advantage is coverage: harmonized input–output tables for over 150 countries across nearly three decades. That coverage is achieved by reconciling heterogeneous national sources through an optimization protocol, and the reconciliation leaves inconsistencies that bear on our estimation. This appendix documents those we have encountered and reports how our results respond to them.

EORA contains a 'visible seam' in the tables in 2015–16 (Eora, 2026), with large variablity in cells. To the extent that the seam moves many country-sectors in a common direction, the factor removal of Section IV absorbs part of it by construction. We nonetheless report estimates excluding windows ending after 2015 in Table A4, column (5).

Eora (2026) also note that the data contain small values where one might expect zero values, and suggest resolving the model after zeroing these cells. We undertake this approach in Table A4, columns (8) and (9). Column (8) zeroes small origin-destination-sector cells (less than \$10,000), and column (9) zeros small origin-destination cells (less than \$100,000).

In our own data quality checks, we find that a small number of country-sectors report year-on-year changes in the intermediates matrix too large to be obviously plausible, with domestic flows differing by factors in the hundreds between adjacent years. These are concentrated among small, largely closed economies, predominantly in the early 1990s, and the pattern is not consistent with a change in currency of account. Because our wage measure is $w_{n,t} = \text{labsh}_{n,t}\, R_{n,t}/\text{emp}_{n,t}$, these movements enter as extreme decadal log wage changes.

Their consequences differ between the two parts of the paper. For the effects of demand on growth, they are minor. The data jumps are confined to the domestic intermediate-use block of the T matrix; bilateral trade flows with other countries are unaffected. Because the affected countries are small and nearly autarkic in the early 1990s, their trade shares in other countries' expenditure are negligible. The inflated domestic cells push each country's own import shares toward zero, further isolating the distortion. Through the lens of the model, the contestability operator S_j vanishes under autarky, so the contaminated revenue observations transmit no measurable effect to any other country's outcomes.

For the estimation of θ, they matter more. Identification comes from cross-country propagation of wage innovations. Observations with extreme wage changes can exert substantial influence on the curvature of the GMM objective, disproportionate to the identifying information they carry: propagation from a near-autarkic country is weak, by the same logic that makes these observations harmless downstream. We therefore winsorize decadal log wage differences at 2.5, which bounds their influence on the objective. Table A4, column (3) reports the corresponding estimates. To be fully transparent, if we were to instead drop the countries that ever have decadal log wage differences above 2.5, we would not estimate a θ above one.[5]

[5] In the conference version of the draft, we fixed a numeraire country (Luxembourg) rather than normalizing the mean to zero. We switched our approach after realizing that our estimates changed depending on the numeraire. However, in this original approach we did able to estimate a θ close to the baseline when

B Marshallian Externalities

The baseline model holds sectoral productivity fixed as demand reallocates. In this section, we discuss how we adapt Kucheryavyy, Lyn and Rodríguez-Clare (2023) to account for Marshallian returns to scale

B.A Environment

The only change to the environment of Section II.A is that productivity in a country-sector increases with the scale of activity there. Efficiency of a variety in country i, sector j, is $z_i^j(\omega)\,(L_{i,t}^j)^{\varphi^j}$, where $z_i^j(\omega)$ is drawn from the Fréchet distribution with location $\lambda_{i,t}^j$ and shape θ^j as before, $L_{i,t}^j$ is *aggregate* sector-j employment in i, and $\varphi^j \geq 0$ is the sector-j scale elasticity. The externality is external to the firm and internal to the country-sector: individual producers take $L_{i,t}^j$ as given, so perfect competition and zero profits are undisturbed.

Standard Eaton–Kortum algebra then delivers the equilibrium of Section II.A with one substitution: the Fréchet location $\lambda_{i,t}^j$ is replaced everywhere by $\lambda_{i,t}^j\,(L_{i,t}^j)^{\kappa^j}$, where

$$\kappa^j \;\equiv\; \theta^j \varphi^j \tag{21}$$

is the composite scale parameter.[6] In particular the trade share (3) becomes

$$\pi_{ni,t}^j \;=\; \frac{\lambda_{i,t}^j\left(L_{i,t}^j\right)^{\kappa^j}\left(w_{i,t}\,d_{ni,t}^j\right)^{-\theta^j}}{\sum_h \lambda_{h,t}^j\left(L_{h,t}^j\right)^{\kappa^j}\left(w_{h,t}\,d_{nh,t}^j\right)^{-\theta^j}}, \tag{22}$$

and the price index (2) is modified identically inside the bracketed sum. At $\varphi^j = 0$ the model of the main text obtains.

With increasing returns the equilibrium *in levels* need not be unique; Kucheryavyy, Lyn and Rodríguez-Clare (2023) characterize the admissible region, and we return to it in Section B.F. Our empirical objects, however, are local: we log-linearize around the *observed* equilibrium, so what the analysis requires is that the observed allocation be interior and that the log-linear system be invertible there.

B.B Log-linearization and sectoral labor

Because the level system is the baseline system under the replacement $\lambda \mapsto \lambda L^{\kappa}$, the log-linearized system (5)–(7) is the baseline system under the substitution

$$\hat{\lambda}_{i,t}^j \;\longmapsto\; \hat{\lambda}_{i,t}^j + \kappa^j\,\hat{L}_{i,t}^j. \tag{23}$$

The new terms involve sectoral employment growth, which is itself an equilibrium object. With labor the only factor and the externality external to firms, sectoral zero profits equate the sector-j wage bill to sector-j revenue. Our wage measure in Section III is $w_{i,t} = \text{labsh}_{i,t} R_{i,t}/\text{emp}_{i,t}$, which treats the labor share as uniform across sectors within a

---

dropping countries with large wage swings, probably because of their limited trade with Luxembourg.

[6] κ^j plays the role of α_k in Kucheryavyy, Lyn and Rodríguez-Clare (2023). Everything below depends on (θ^j, φ^j) only through (θ^j, κ^j), so we state results, and estimate, in (θ, κ) and back out $\hat{\varphi}^j = \hat{\kappa}^j/\hat{\theta}^j$. We reserve φ^j for the scale elasticity; the sourcing weights ϕ_{ni}^j of equation (12) are unrelated.

country while letting it move over time; the sectoral wage bill is then $w_{i,t}L^j_{i,t} = \text{labsh}_{i,t}\, R^j_{i,t}$, so that

$$\hat{L}^j_{i,t} \;=\; g^j_{i,t} - \hat{w}_{i,t} + \hat{\ell}_{i,t}, \qquad \hat{\ell}_{i,t} \;\equiv\; \Delta \log \text{labsh}_{i,t}, \tag{24}$$

B.C The revenue block

For each sector define the $N \times N$ *contestability operator*

$$S_j \;\equiv\; I_N - \Phi_j \Pi_j, \qquad (S_j)_{ih} \;=\; \mathbb{1}\{i = h\} - \sum_n \phi^j_{ni}\, \pi^j_{nh}. \tag{25}$$

The (i, h) entry measures the destination overlap between sources i and h: how much of i's sector-j revenue is earned in destinations where h holds market share. Its diagonal, $1 - \sum_n \phi^j_{ni}\pi^j_{ni}$, is large when i's own share in its destination markets is small—when i's markets are contestable. Note that $\tilde{H}^R_j(\theta^j)$ of equation (14) can be written compactly as $\tilde{H}^R_j(\theta^j) = \Phi_j - \theta^j S_j$.

Lemma 1 (Spectral structure). *$\Phi_j\Pi_j$ is row-stochastic, so $S_j\mathbf{1}_N = 0$. All eigenvalues of $\Phi_j\Pi_j$ are real and lie in $[0, 1]$; hence all eigenvalues of S_j are real and lie in $[0, 1]$, and for any $\kappa^j \in [0, 1)$ the matrix $I_N - \kappa^j S_j$ is invertible, with convergent Neumann series $\sum_{m\geq 0}(\kappa^j S_j)^m$. At $\kappa^j = 1$, $I_N - S_j = \Phi_j\Pi_j$ is invertible if and only if Π_j is nonsingular. Under autarky, $S_j = 0$.*

Proof. Row-stochasticity follows because Φ_j and Π_j are each row-stochastic. Writing $\Phi_j\Pi_j = D^{-1}_{R^j}\Pi'_j D_{X^j}\Pi_j$, with D_{R^j} and D_{X^j} the diagonal matrices of sectoral revenues and expenditures, the product is similar (via $D^{1/2}_{R^j}$) to the symmetric positive semidefinite matrix $D^{-1/2}_{R^j}\Pi'_j D_{X^j}\Pi_j D^{-1/2}_{R^j}$, so its spectrum is real and nonnegative; interiority ($R^j_i > 0$) makes the similarity well defined. Being nonnegative and row-stochastic, its spectral radius is one. Hence $\text{spec}(S_j) = 1 - \text{spec}(\Phi_j\Pi_j) \subset [0, 1]$, and the eigenvalues of $I_N - \kappa^j S_j$ lie in $[1 - \kappa^j, 1]$. At $\kappa^j = 1$, singularity of $D^{-1}_{R^j}\Pi'_j D_{X^j}\Pi_j$ is equivalent, positive diagonal scalings aside, to singularity of Π_j. Under autarky $\Phi_j = \Pi_j = I_N$. □

In the baseline system the productivity shocks enter sectoral revenue growth with loading S_j—that is, the productivity component of $\eta_{j,t}$ in (15) is $S_j\hat{\lambda}_{j,t}$ plus trade-cost terms—so the substitution (23) adds the observable term $\kappa^j S_j \hat{L}_{j,t}$ to each sector block:

$$g_{j,t} \;=\; \Phi_j\, a_{j,t} + \tilde{H}^R_j(\theta^j)\, \hat{w}_t + \kappa^j S_j\, \hat{L}_{j,t} + \eta_{j,t}, \tag{26}$$

where $g_{j,t}, a_{j,t}, \hat{L}_{j,t}, \eta_{j,t} \in \mathbb{R}^N$ are the sector-j blocks. That $S_j\mathbf{1} = 0$ is economically transparent: a uniform expansion of every country's sector-j employment leaves all relative competitiveness, and hence all trade shares, unchanged—only *relative* scale matters in a share system. Eliminating $\hat{L}_{j,t} = g_{j,t} - \hat{w}_t + \hat{\ell}_t$ with (24) and using $\tilde{H}^R_j(\theta^j) - \kappa^j S_j = \Phi_j - (\theta^j + \kappa^j)S_j = \tilde{H}^R_j(\theta^j + \kappa^j)$,

$$\big(I_N - \kappa^j S_j\big)\, g_{j,t} \;=\; \Phi_j\, a_{j,t} + \tilde{H}^R_j\big(\theta^j + \kappa^j\big)\, \hat{w}_t + \kappa^j S_j\, \hat{\ell}_t + \eta_{j,t}, \tag{27}$$

where the labor-share term is observable and carried through estimation. Stacking across sectors, with $K \equiv \operatorname{blockdiag}\big(I_N - \kappa^j S_j\big)_{j=1}^{J}$ and $\tilde H^R(\theta{+}\kappa)$ the $NJ \times N$ matrix stacking the $\tilde H^R_j(\theta^j + \kappa^j)$,

$$K\, g_t \;=\; D\, a_t + \tilde H^R(\theta{+}\kappa)\, \hat w_t + (I_{NJ} - K)\, E_w \hat\ell_t + \eta_t, \tag{28}$$

using $\operatorname{blockdiag}(\kappa^j S_j) = I_{NJ} - K$ and E_w as defined below.

Relative to (15), the externality does exactly two things. First, wage propagation operates at the *effective elasticity* $\theta^j + \kappa^j = \theta^j(1 + \varphi^j)$: a fall in a country's relative wage improves its competitiveness directly (elasticity θ^j) and, by expanding its sectoral employment, raises its productivity and improves competitiveness again (the additional $\theta^j\varphi^j$). Second, shocks are passed through the within-sector *scale multiplier* $K_j^{-1} = \sum_{m\geq 0}(\kappa^j S_j)^m$, which amplifies their nonuniform contestability eigenmodes and leaves the uniform mode untouched ($S_j\mathbf{1} = 0$); individual countries can lose through the negative business-stealing cross-loadings. Round m of the series is the m-th round of feedback: a revenue impulse raises employment one-for-one by (24), raises productivity by φ^j, and feeds back into revenue with positive own-loading (agglomeration, strongest where markets are contestable) and negative cross-loadings proportional to destination overlap (business stealing). Under autarky $S_j = 0$ and the channel is dead: scale economies show up in revenue dynamics only where market shares can move.

B.D The wage block

The wage equation changes too, and it is instructive to derive it by aggregation. Balanced trade with a fixed labor endowment implies $\hat w_{n,t} = \hat R_{n,t} = \sum_j \bar\omega^j_n\, g^j_{n,t}$, since $\bar\omega^j_n = \sum_i X^j_i \pi^j_{in}/R_n = R^j_n/R_n$ is precisely sector j's share of country n's revenue. Multiplying each sector equation (26) by $\bar\omega^j_n$, summing over j, moving the $-\theta^j \hat w_{n,t}$ terms to the left, and dividing by $1 + \bar\theta_n$ reproduces the baseline objects $\Gamma(\theta)$ and $A(\theta)$ of (10) exactly, plus one new observable term:

$$\hat w_t \;=\; \Gamma(\theta)\, \hat w_t + A(\theta)\, a_t + C_L\, \hat L_t + \varepsilon_t, \qquad \big[C_L\big]_{n,(i,j)} \;=\; \frac{\kappa^j\, \bar\omega^j_n}{1 + \bar\theta_n}\, \big(S_j\big)_{ni}, \tag{29}$$

where $\hat L_t = g_t - E_w \hat w_t + E_w \hat\ell_t$ stacks (24) and E_w is the $NJ \times N$ matrix of J stacked copies of I_N, so that $E_w \hat w_t$ repeats the wage vector across sectors. The sign is intuitive: $[S_j \hat L_{j,t}]_n > 0$ when n expands sector-j employment relative to its rivals' destination-weighted expansion, which raises n's competitiveness, its export revenue, and hence its wage. At $\kappa = 0$, $C_L = 0$ and (29) is (10).

The aggregation derivation also shows in what sense the wage block is unchanged: as a level condition it is the same balanced-trade identity as in the baseline, and the innovation it delivers is the revenue-share-weighted aggregate of the sectoral residuals, $\varepsilon_{n,t} = (1 + \bar\theta_n)^{-1} \sum_j \bar\omega^j_n\, \eta^j_{n,t}$. Two implications are worth recording. First, own-country correlation between ε and η is generic—it is built into the model, in the baseline as much as here—which is precisely why the instrument of Section II.C is constructed leave-one-out: part (iii) of Assumption 1 is required, and used, only across countries.[7] Second, what changes with $\kappa > 0$

[7] In the data the wage equation retains independent content because $\hat w$ is measured from PWT labor shares and employment rather than constructed from EORA wages; the wedge between the two measurements is

is only the *estimating* equation: dropping the observable $C_L\hat{L}_t$ term is not a simplification but a misspecification, whose consequences for the instrument we spell out in Section B.G.

B.E Reduced form

Solving (28) for g_t (Lemma 1 guarantees K^{-1} exists for $\kappa^j < 1$), substituting into (29), and collecting wage terms yields the scale-adjusted wage system

$$\big(I_N - \Gamma^s\big)\hat{w}_t \;=\; A^s a_t + C_L K^{-1}\eta_t + \varepsilon_t, \qquad \begin{aligned} \Gamma^s &\equiv \Gamma(\theta) - C_L E_w + C_L K^{-1}\tilde{H}^R(\theta{+}\kappa), \\ A^s &\equiv A(\theta) + C_L K^{-1} D. \end{aligned} \tag{30}$$

Nominal homogeneity is preserved: $S_j\mathbf{1}_N = 0$ implies $C_L\mathbf{1}_{NJ} = 0$ and $K^{-1}\mathbf{1}_{NJ} = \mathbf{1}_{NJ}$, and $\tilde{H}^R(\theta{+}\kappa)\mathbf{1}_N = \mathbf{1}_{NJ}$ holds at the shifted elasticity, so the rows of Γ^s sum to one exactly as those of Γ do. Wages are therefore again determined only up to a common scalar, and we impose the same normalization as in the main text: wage changes are mean zero across countries in every period, and the wage block is solved on the mean-zero subspace (Appendix A). Let Q be an orthonormal basis of $\mathbf{1}^\perp$ and define the projected inverse

$$\Phi^s_w \;\equiv\; Q\big(Q'(I_N - \Gamma^s)Q\big)^{-1}Q',$$

which does not depend on the choice of basis. Then $\hat{w}_t = \Phi^s_w\big[A^s a_t + C_L K^{-1}\,\eta_t + \varepsilon_t\big]$, and substituting back into (28),

$$g_t \;=\; M(\theta,\kappa)\,a_t \;+\; J_w(\theta,\kappa)\,\varepsilon_t \;+\; J_\eta(\theta,\kappa)\,\eta_t, \tag{31}$$

with

$$\begin{aligned} &M(\theta,\kappa) = K^{-1}\Big[\tilde{H}^R(\theta{+}\kappa)\,\Phi^s_w\,A^s + D\Big], \qquad J_w(\theta,\kappa) = K^{-1}\,\tilde{H}^R(\theta{+}\kappa)\,\Phi^s_w, \\ &J_\eta(\theta,\kappa) = K^{-1}\Big[\tilde{H}^R(\theta{+}\kappa)\,\Phi^s_w\,C_L K^{-1} + I_{NJ}\Big]. \end{aligned} \tag{32}$$

The undetermined wage level translates, through $\tilde{H}^R(\theta{+}\kappa)\mathbf{1}_N = \mathbf{1}_{NJ}$, into a period-specific scalar in predicted revenue growth with no structural content, absorbed by demeaning the residual within each period exactly as at $\kappa = 0$. At $\varphi = 0$: $K = I$, $C_L = 0$, $\Gamma^s = \Gamma$, $A^s = A$, so M and J_w collapse to their expressions in (16) and $J_\eta = I_{NJ}$. The structural shocks (ε_t, η_t) are the same objects as in the baseline, so (31) nests (16) exactly. The one genuinely new feature is $J_\eta \neq I$: with scale economies, sectoral shocks anywhere move employment, hence productivity, hence revenues everywhere, so the reduced-form residual of g_t is a dense mixture of η's across countries. This matters for estimation and we return to it below.

Decomposing $K^{-1} = I + (K^{-1} - I)$ separates the demand effect into three channels:

$$M(\theta,\kappa)\,a_t \;=\; \underbrace{D\,a_t}_{\substack{\text{impact}\\\text{(parameter-free)}}} \;+\; \underbrace{(K^{-1} - I)\,D\,a_t}_{\substack{\text{scale amplification}\\(\kappa\text{ only})}} \;+\; \underbrace{K^{-1}\tilde{H}^R(\theta{+}\kappa)\,\Phi^s_w\,A^s\,a_t}_{\text{GE wage channel}}. \tag{33}$$

The impact effect of a demand reallocation—including the foreign-demand object $D^F a$ of

part of what ε absorbs, exactly as in the baseline.

equation (20), the anchor of Sections V and VI—is still the parameter-free object of the main text. The externality adds, on top, a within-sector amplification $(K^{-1} - I)Da_t = \sum_{m\geq 1} \text{blockdiag}\big((\kappa^j S_j)^m\big) Da_t$ that redistributes the impulse toward agglomerating producers and away from their destination-overlapping rivals, and it modifies the general-equilibrium wage channel both through the effective elasticity and through A^s, which now includes the reallocation-to-scale-to-competitiveness route by which demand shocks move wages. For demand tropism the reading is direct: under $\kappa^j > 0$, reallocating production toward rising-demand sectors raises productivity in those sectors, so tropic countries receive the amplification term in addition to the impact effect measured in the main text.

B.F Existence, uniqueness, and the admissible region

Lemma 1 delivers the local counterpart of the equilibrium theory in Kucheryavyy, Lyn and Rodríguez-Clare (2023). For $\kappa^j < 1$ in every sector, the sectoral-allocation stage has a unique interior solution given wages (their Proposition 1), and the log-linear multiplier K^{-1} exists unconditionally with a convergent geometric-series interpretation. The knife-edge $\kappa^j = 1$ requires the nonsingularity of Π_j; this is the exact log-linear analog of their Assumption 1 (nonsingularity of the iceberg-cost kernel), since Π_j is that kernel up to positive row and column scalings, and it fails under free trade, where every row of Π_j is identical and the allocation is indeterminate. For $\kappa^j > 1$, multiplicity and complete specialization arise generically, the observed allocation may be a corner, and a log-linearization around an interior point is the wrong object; we therefore do not use the formulas above in that region and impose $\kappa^j \in [0, 1)$ in estimation. Two caveats delimit the scope of these statements. First, $\kappa^j < 1$ delivers uniqueness of the sectoral stage *given wages*; uniqueness of the full wage equilibrium with many countries and costly trade is an open question in Kucheryavyy, Lyn and Rodríguez-Clare (2023), who verify it numerically. Second, everything presumes an interior observed allocation, $R_i^j > 0$: under $\kappa^j > 0$ a zero cell is a corner at which the local expansion is invalid, so exact zero-revenue cells are excluded from estimation.

B.G Identification

At candidate (θ, κ), the structural equations (29) and (28) deliver the residuals directly from observables. With $P \equiv I_N - \frac{1}{N}\mathbf{11'}$ the cross-country demeaning projection of the main text,

$$\hat{\varepsilon}_t(\theta, \kappa) = P\Big[\big(I - \Gamma(\theta)\big)\,\hat{w}_t - A(\theta)\,a_t - C_L\,\hat{L}_t\Big], \tag{34}$$

$$\hat{\eta}_t(\theta, \kappa) = K\,g_t - D\,a_t - \tilde{H}^R(\theta{+}\kappa)\,\hat{w}_t - (I_{NJ} - K)\,E_w\hat{\ell}_t, \tag{35}$$

extending (11) and (19) by observable corrections only (the $\hat{L}$ terms in $C_L\hat{L}_t$ likewise carry the observable $\hat{\ell}_t$). Because the scale terms are observable, recovery at the truth is exact: $\hat{\varepsilon}_t(\theta_0, \kappa_0) = \varepsilon_t$ and $\hat{\eta}_t(\theta_0, \kappa_0) = \eta_t$. The residuals are then demeaned within period and factor-purged exactly as in Section IV, yielding $\tilde{\varepsilon}_t$ and $\tilde{\eta}_t$, and the instrument is the model-implied contribution of non-c innovations at the extended loading,

$$z_{k,t}(\theta, \kappa) \;=\; J_w(\theta, \kappa)(k, :)\;\tilde{\varepsilon}_t^{(-c)}(\theta, \kappa), \tag{36}$$

as in (17)–(18); under the mean-zero normalization every country's own entry is zeroed in its own instrument, with no asymmetry across countries.

One implementation point deserves emphasis. Under the extension the GMM residual must be the *structural* residual (35), not the reduced-form residual of the main text: identically in the data, $g_t - M(\theta,\kappa)a_t - J_w(\theta,\kappa)\hat{\varepsilon}_t = J_\eta(\theta,\kappa)\,\hat{\eta}_t$, and since J_η is dense, row k of the reduced-form residual mixes $\hat{\eta}$'s from all countries and would pair the instrument with own-country shocks of *other* countries. In the baseline the two residual concepts coincide ($J_\eta = I$), which is why the distinction never arises in Section IV.

Proposition 2 (Identifying moment with scale economies)**.** *Under Assumption 1, unchanged, and $\kappa_0^j < 1$ for all j: at $(\theta,\kappa) = (\theta_0,\kappa_0)$,*

$$\mathbb{E}\big[z_{k,t}(\theta_0,\kappa_0)\,\tilde{\eta}_{k,t}\big] = 0 \qquad \textit{for every country-sector } k.$$

Proof. At (θ_0,κ_0), exact recovery gives $\tilde{\eta}_{k,t}$ as the factor-purged true residual, and $z_{k,t}$ is a linear combination of $\{\tilde{\varepsilon}_{c',t}\}_{c'\neq c}$ with weights from row k of $J_w(\theta_0,\kappa_0)$. By Assumption 1(iii), each $\tilde{\varepsilon}_{c',t}$ with $c' \neq c$ is orthogonal to $\tilde{\eta}_{k,t}$, so every term in the combination has zero expectation, as in the proof of Proposition 1. □

The exclusion logic is thus unchanged; what the extension changes is the cost of ignoring it. Suppose the scale terms were instead absorbed into the wage residual—that is, the baseline recovery (11) were applied to data generated with $\kappa_0 > 0$. By (30) the recovered innovation would be $u_t = \varepsilon_t + C_L K^{-1}\,\eta_t$, which loads on *foreign* sectoral shocks through the dense matrix $C_L K^{-1}$. The leave-one-out instrument for country c would then be built from $\tilde{u}_{c'}$, $c' \neq c$, each containing $\eta_{c'}$, and would pair with a residual that also mixes $\eta_{c'}$; the product picks up own-country covariances $\mathrm{Cov}(\tilde{\eta}_{k'\in c'},\tilde{\varepsilon}_{c'})$, exactly the covariances that Assumption 1(iii) does not restrict—and, by the aggregation identity of Section B.D, cannot restrict, since they are mechanically nonzero. The exclusion restriction fails not because the economics changed but because the econometrician's ε is no longer ε. The corrections in (34)–(35) repair this at no cost in assumptions, which is why we estimate (θ,κ) jointly rather than treating scale as part of the innovation.

Identification of the additional parameter comes through observable curvature that θ alone cannot generate: κ^j moves the within-sector amplification $(K^{-1}-I)D$ of the parameter-free demand impulse, the effective elasticity $\theta^j + \kappa^j$ in the wage-propagation block, and the C_L correction, while θ^j separately enters the wage-side weights $\Gamma(\theta), A(\theta), \bar{\theta}_n$. A caveat follows from Lemma 1: in near-autarkic sectors $S_j \approx 0$ and the data carry little information about κ^j—where markets are not contestable, scale economies leave no footprint in trade flows. This is economics rather than a defect, but it motivates pooling: we mirror the parameterizations of the main text, a common (θ,κ) and a goods/services split $(\theta_g,\theta_s,\kappa_g,\kappa_s)$, rather than sector-by-sector scale elasticities.

B.H Estimation

Estimation proceeds exactly as in Section IV: two-step GMM on the moments $\mathbb{E}[z_{k,t}\,\tilde{\eta}_{k,t}] = 0$, with an identity-weighted first step, inverse-variance diagonal weighting in the second, the same factor counts, and the same rolling windows. The parameter space is (θ^j,κ^j) with $\kappa^j \in [0,1)$ enforced by a logistic transform; the objective is minimized from multiple starting points, including near-boundary values of κ, to guard against local minima (Appendix A); and the scale elasticity is recovered as $\hat{\varphi}^j = \hat{\kappa}^j/\hat{\theta}^j$. Note that at $\kappa = 0$ the objective coincides with the baseline objective to machine precision. Because the effective elasticity $\theta^j + \kappa^j$ governs wage propagation, the criterion is considerably more informative about the sum than about its split into substitution and scale; we therefore emphasize $\hat{\theta}^j + \hat{\kappa}^j$, which is stable across specifications and close to the constant-returns estimates, and read the division between $\hat{\theta}^j$ and $\hat{\kappa}^j$ with corresponding caution.

Table A1: Demand-Driven Variance of Revenue Growth, by Sector: Total-Absorption Demand Shocks

Sector	$\text{var}(D^F a)/\text{var}(g)$	$\text{var}(Da)/\text{var}(g)$
Mining and Quarrying	1.2%	14.4%
Textiles and Wearing Apparel	0.8%	6.1%
Fishing	0.5%	16.1%
Recycling	0.4%	10.5%
Metal Products	0.3%	5.0%
Transport Equipment	0.3%	7.3%
Other Manufacturing	0.1%	10.4%
Petroleum, Chemical and Non-Metallic Mineral Products	0.1%	3.4%
Transport	0.1%	4.2%
Wholesale Trade	0.1%	13.6%
Agriculture	0.1%	13.4%
Hotels and Restaurants	0.1%	10.4%
Electrical and Machinery	0.1%	4.2%
Wood and Paper	0.05%	3.3%
Food & Beverages	0.03%	3.0%
Post and Telecommunications	0.03%	11.1%
Education, Health and Other Services	0.01%	15.0%
Financial Intermediation and Business Activities	0.002%	12.1%

Notes: Variance shares of revenue growth g attributable to total demand $D\,a$ and to its foreign-only component $D^F\,a$, computed within each sector. Sectors are sorted in descending order by foreign-demand variance share. Relative to the main table, the demand shocks a are constructed from total-absorption expenditure shares (intermediate plus final use) rather than final-demand shares Pooled across rolling ten-year windows from 1990–1999 through 2008–2017.

Table A2: Revenue Growth and the Model-Implied GE Response to Demand: Total-Absorption Demand Shocks

	(1)	(2)	(3)	(4)
Panel A: Common parameters				
Ma (standardized)	0.2405	0.2375	0.2375	0.2291
	(0.0074)	(0.0085)	(0.0088)	(0.0104)
Panel B: Separate parameters for Goods and Services				
Ma (standardized)	0.2352	0.2314	0.2315	0.2223
	(0.0063)	(0.0070)	(0.0051)	(0.0061)
Observations	46,512	46,512	46,512	46,512
Countries	136	136	136	136
Country FE	No	Yes	No	Yes
Continent × year FE	No	No	Yes	Yes

Notes: OLS regressions of revenue growth g on the model-implied general-equilibrium response to demand, $M(\hat{\theta})a$, standardized by its sample-wide standard deviation. Relative to the main table, the demand shocks a are constructed from total-absorption expenditure shares (intermediate plus final use) rather than final-demand shares. Panel A uses a common $\hat{\theta}$; Panel B estimates separate parameters for goods and services. Each observation is a country-sector-window; the sample pools rolling ten-year windows from 1990–1999 through 2008–2017. Standard errors clustered at the country level in parentheses.

Table A3: Revenue Growth and the Model-Implied GE Response to Demand: Increasing Returns

	(1)	(2)	(3)	(4)
Panel A: Common parameters				
Ma (standardized)	0.2057	0.1943	0.2013	0.1860
	(0.0321)	(0.0349)	(0.0275)	(0.0290)
Panel B: Separate parameters for Goods and Services				
Ma (standardized)	0.2041	0.1923	0.2001	0.1846
	(0.0300)	(0.0326)	(0.0252)	(0.0266)
Observations	46,512	46,512	46,512	46,512
Countries	136	136	136	136
Country FE	No	Yes	No	Yes
Continent × year FE	No	No	Yes	Yes

Notes: OLS regressions of revenue growth g on the model-implied general-equilibrium response to demand under Marshallian external economies of scale. Each observation is a country-sector-window; the sample pools rolling ten-year windows from 1990–1999 through 2008–2017. Standard errors clustered at the country level in parentheses.

Table A4: Revenue Growth and the Model-Implied GE Response to Demand, Robustness to Data Choices

	(1) Baseline	(2) All sectors	(3) Winsorize	(4) Emp floor	(5) Pre-2015	(6) $K_r = 2$	(7) $K_r = 1$	(8) Zero cells <10k	(9) Zero blocks <100k
Ma (std.)	0.1945	0.1568	0.1913	0.2019	0.1963	0.1940	0.1944	0.1948	0.1943
	(0.0347)	(0.0487)	(0.0354)	(0.0372)	(0.0397)	(0.0348)	(0.0348)	(0.0331)	(0.0341)
N	46,512	64,600	46,512	42,750	41,616	46,512	46,512	46,512	46,512
Country FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes

Notes: Each column reports the slope from regressing revenue growth g on standardized $M(\hat{\theta})a$ with country fixed effects and country-clustered standard errors—the specification of column (2) of Table 8, with the model re-estimated for each column. (1) Baseline: final-demand demand shocks, $K_\varepsilon = 1$, $K_r = 3$. (2) All 25 sectors retained. (3) Wage changes winsorized at $|\hat{w}| \leq 2.5$ before estimation. (4) Countries with employment below 200,000 at the start of the sample folded into the rest-of-world aggregate. (5) Windows ending after 2015 dropped ($T = 17$) due to a "visible seam " in the data after this year (Eora, 2026). (6)–(7) Baseline data with two and one revenue factors removed ($K_r = 2$ and $K_r = 1$). (8) Transaction-level entries below \$10,000 set to zero ($K_r = 1$). (9) Aggregated bilateral flows below \$100,000 set to zero ($K_r = 1$).

Table A5: Foreign Capital Flows and Tropism

	Q1 (least tropic)	Q2	Q3	Q4 (most tropic)
Mean tropism	-0.378	-0.038	0.127	0.417
Mean *kai*	0.353	0.354	0.354	0.254
Median *kai*	0.240	0.286	0.244	0.135
Mean *eqi*	0.334	0.322	0.357	0.255
Median *eqi*	0.222	0.100	0.183	0.000
N	418	418	418	418

Notes: Tropism is the within-country cross-sector correlation between the change in production shares and foreign demand (DFa). *kai* is the Fernández et al. (2016) capital account openness index and *eqi* is the equity openness index. For both measures, 0 = open and 1 = restricted, averaged over each 10-year window. Unit of observation: country-window.