Review of International Economics

WILEY

ORIGINAL ARTICLE OPEN ACCESS

Markup Centrality and International Incidence in Global Production Networks

Constantin Colonescu

Department of Economics, MacEwan University, Edmonton, Alberta, Canada

Correspondence: Constantin Colonescu (ColonescuC@macewan.ca)

Received: 21 February 2026 | **Revised:** 15 April 2026 | **Accepted:** 22 April 2026

Keywords: compound markups | global value chains | input-output price model | international incidence | market power | pass-through | production networks

ABSTRACT

This paper develops a framework to measure how markups amplify prices through global production networks and to attribute final-demand price wedges to upstream country-industry sources. The approach defines a compound markup as the ratio of observed prices to counterfactual pure-cost prices that would prevail if all markups in the network were one. Using an input-output price model under a Cobb-Douglas benchmark, the method yields an implementable mapping from direct markups to network-propagated wedges and an exact decomposition into upstream contributions. The framework is implemented with the World Input-Output Database 2016 release. The decomposition identifies upstream nodes with high markup centrality and quantifies the international incidence of the aggregate wedge, including destination-specific attributions that separate domestic from foreign contributions and an exact ordered ranking of foreign-incidence exposure across WIOD destinations. Compound markups provide a rigorous accounting of how market power propagates through global value chains, with most of the aggregate wedge explained by a small number of propagation rounds and concentrated upstream contributors.

JEL Classification: D43, L11, L12, F14, C67

1 | Introduction

Market power is typically measured with a *direct* markup: the wedge between a producer's price and marginal cost evaluated at observed input prices (Hall 1988; De Loecker and Warzynski 2012). In global value chains, however, observed intermediate-input prices already embed upstream pricing wedges. As a result, a small direct markup at a downstream node can coexist with a large price cost wedge relative to a benchmark that strips out upstream markups from intermediate prices. This paper develops an implementable way to construct that benchmark and to attribute the resulting final-demand wedge to upstream country-industry sources.

I define a *pure-cost* price system as the counterfactual in which all output markups in the network are set to one, holding technologies, primary-input prices, and the input-output coefficient matrix fixed. The resulting *compound markup* κ_j is the ratio of the observed price of node j to its pure-cost benchmark price and therefore measures the cumulative markup wedge embedded along all upstream production paths. Under a Cobb-Douglas unit-cost benchmark, log compound markups satisfy a simple linear system, $\ln \kappa = (I - B^{\top})^{-1} \ln \mu$, which maps a vector of direct markups into network-propagated wedges. The same structure delivers an exact upstream attribution of the aggregate wedge, $\ln \kappa^{agg} = \sum_i C_i \ln \mu_i$, where $C = (I - B)^{-1} s$ are *markup-centrality* weights.

Implementing the framework with WIOD 2016 (baseline year 2014), the final-demand-weighted aggregate direct markup is $\mu^{\text{agg}} = 1.076$ while the aggregate compound markup is $\kappa^{\text{agg}} = 1.196$, implying that embedded upstream wedges raise the aggregate final-demand price cost wedge by about 11% relative to the direct-markup benchmark. The node-level distribution is also shifted upward: the median stabilized direct markup is 1.002, the median compound markup is 1.143, and the median amplification factor κ_i/μ_i is 1.103. Upstream incidence is concentrated: the top 50 upstream nodes account for 54.0% of $\ln \kappa^{\text{agg}}$. A destination-specific attribution shows substantial cross-border incidence: foreign upstream markups account for 20.6% of the U.S. destination wedge, 44.0% for Germany, and 51.2% for the United Kingdom. Across WIOD destinations, foreign-incidence shares range from 3.4% in India to 93.4% in Luxembourg, with Germany slightly above the sample median and the United Kingdom in the upper third of the distribution. Because WIOD prices are recorded in the WIOT valuation convention, these measured "markups" should be read as price cost wedges within that accounting system rather than as pure firm-level monopoly margins; taxes, subsidies, trade and transport margins, balancing adjustments, and imputed rents are part of the measured wedge and matter especially for sectors such as real estate.

The remainder of the paper proceeds as follows. Section 3 develops the mapping from direct to compound markups and the upstream decomposition. Section 4 and Section 5 describe the WIOD data and the construction and stabilization of direct markups. Section 6 summarizes the computational implementation, and Section 7 reports baseline results and the international-incidence extension.

2 | Related Literature

This paper bridges two literatures. First, it builds on production-based approaches to markup measurement (Hall 1988; De Loecker and Warzynski 2012; De Loecker et al. 2020) but emphasizes that in a production network, the relevant benchmark for incidence is not marginal cost at observed input prices, but a counterfactual that removes upstream markup wedges from intermediate-input prices. Second, it draws on input-output and production-network frameworks (Leontief 1941; Miller and Blair 2009; Acemoglu et al. 2012; Carvalho 2014; Baqaee and Farhi 2020; Bigio and La'O 2020) to propagate wedges and to obtain exact network-consistent attributions.

Conceptually, the compound markup object generalizes classic successive markup (double marginalization) logic (Spengler 1950; Tirole 1988) to high-dimensional global networks with many paths and feedback loops.

It is also related to global value chain measures of distance to final demand (e.g., upstreamness) (Antràs et al. 2012) and to prior WIOD-based evidence on markups and upstreamness (Colonescu 2021).

3 | Theory: Direct and Compound Markups in a Production Network

This section states the theoretical objects needed for the empirical analysis:

i. a *pure-cost* price benchmark that removes markup wedges from intermediate-input prices while holding technology and primary-input prices fixed;
ii. an implementable mapping from a measured vector of *direct* markups to *compound* markups in a production network; and
iii. an exact upstream attribution of the aggregate wedge. Detailed derivations and broader generalizations are provided in Colonescu (2026). The goal here is to present an empirically oriented summary and fix notation.

3.1 | Pure-Cost Prices and the Compound Markup

The empirical question is how large the price cost wedge becomes once upstream markup wedges embedded in intermediate-input prices are taken into account. Consider an economy with N country-industry nodes indexed by $j \in \{1, \ldots, N\}$. Node j produces one good, uses intermediate inputs sourced from other nodes, and uses a composite primary input (labor and capital) priced at an exogenous factor price. Let p_j be the observed basic price and let μ_j be the *direct* markup, defined as the ratio of price to marginal cost evaluated at observed input prices.

To obtain a transparent and IO-implementable mapping, I use a Cobb-Douglas unit-cost representation as a benchmark:

$$c_j(p, w) = \frac{1}{\phi_j}\left(\prod_{i=1}^{N} p_i^{\beta_{ij}}\right) w_j^{\alpha_j}, \quad \alpha_j > 0, \quad \beta_{ij} \geq 0, \quad \sum_i \beta_{ij} = 1 - \alpha_j, \tag{1}$$

where w_j is a composite primary-input price, ϕ_j is a productivity shifter, α_j is the primary-input share in *resource* costs, and β_{ij} is the share of input i in node j's resource costs. Observed prices satisfy the standard markup-pricing relationship,

$$p_j = \mu_j \, c_j(p, w). \tag{2}$$

Define the *pure-cost benchmark* as the counterfactual economy that keeps technology (ϕ, α, β) and primary-input prices w fixed but sets all output markups to one. Denote by p^0 the corresponding pure-cost price vector. The *compound markup* of node j is then

$$\kappa_j \equiv \frac{p_j}{p_j^0}.$$

By construction, κ_j isolates the part of the observed price that is attributable to markup wedges embedded in intermediate-input prices (and the node's own markup), abstracting from differences in technology or primary-factor prices.

3.2 | The Network Mapping Implemented in WIOD

Taking logs of Equations (1) and (2) yields a linear system in log prices. Let $\ell \equiv \ln p$ and $m \equiv \ln \mu$. Collect all terms held fixed in the pure-cost benchmark into φ (technology and primary-input prices). The log price system can be written as

$$\ell = B^{\top}\ell + m + \varphi,$$

where B is the matrix of intermediate-input resource-cost shares with entries $B_{ij} \equiv \beta_{ij}$. Under the standard productivity condition $\rho(B^{\top}) < 1$, the system has the unique solution

$$\ell = (I - B^{\top})^{-1}(m + \varphi) \equiv L(m + \varphi), \tag{3}$$

where $L \equiv (I - B^{\top})^{-1}$ is the Leontief inverse in the price system. Setting $m = 0$ gives pure-cost log prices,

$$\ell^{0} = L\varphi. \tag{4}$$

Subtracting Equation (4) from Equation (3) delivers the empirical mapping from direct to compound markups:

$$\ln \kappa = L \ln \mu = (I - B^{\top})^{-1} \ln \mu. \tag{5}$$

Equation (5) is the core bridge between measured *direct* markups and *compound* markups: it is a production-network multiplier applied to log direct markups, holding technology and primary-input prices fixed.

3.3 | From WIOD Value Coefficients to the Propagation Matrix

World input-output tables (Timmer et al. 2015; WIOD 2021) report intermediate expenditures at observed prices. Let z_{ij} be the expenditure of node j on inputs from node i and let x_j be gross output. The *value coefficient* matrix is

$$A^{val}_{ij} \equiv \frac{z_{ij}}{x_j}. \tag{6}$$

These are *revenue-based* input shares. The mapping Equation (5), however, is governed by the *resource-cost* shares in marginal cost, collected in B.

Under markup pricing, revenue shares and marginal-cost shares are linked by the node's direct markup. This yields the implementable identity used in the empirical construction:

$$B = A^{val}\,\mathrm{diag}(\mu). \tag{7}$$

Equation (7) explains why a direct-markup vector is needed even when the goal is an accounting-style attribution: it converts observed value coefficients into the cost-share matrix that governs how markup wedges propagate through the network. It also motivates the stabilization steps in the empirical section, which ensure nonnegative implied value-added shares and a well-defined inverse in Equation (5).

3.3.1 | Walk Interpretation and Feedback

The Leontief inverse admits the Neumann-series expansion

$$L = (I - B^{\top})^{-1} = \sum_{t=0}^{\infty} (B^{\top})^{t}, \tag{8}$$

which implies

$$\ln \kappa = \sum_{t=0}^{\infty} (B^{\top})^{t} \ln \mu. \tag{9}$$

The terms with $t = 0$ and $t = 1$ correspond to the node's own wedge and the first upstream round of embedded wedges; higher-order terms capture longer upstream chains. When the production network contains cycles, there are infinitely many upstream walks; their cumulative effect is summarized by the infinite series, which converges under $\rho(B^{\top}) < 1$.

3.3.2 | One-Round Benchmark

For comparison with vertical-chain intuition (successive markup stacking along immediate supply links; see Colonescu 2023), define a one-round compound markup that truncates after the first upstream round:

$$\ln \kappa^{(1)} \equiv (I + B^{\top}) \ln \mu. \tag{10}$$

The difference between total and one-round compounding collects all higher-order propagation effects (including feedback loops):

$$\ln \kappa - \ln \kappa^{(1)} = \sum_{t=2}^{\infty} (B^{\top})^{t} \ln \mu. \tag{11}$$

3.4 | Aggregate Wedges and an Exact Upstream Attribution

Let $s \in \mathbb{R}^{N}_{+}$ be a weight vector that sums to one. In the baseline application, s is the vector of global final-demand expenditure shares across country-industry nodes. Combining Equation (5) with the definition of L gives an exact decomposition of the aggregate log wedge:

$$s^{\top} \ln \kappa = s^{\top}(I - B^{\top})^{-1} \ln \mu = \sum_{i} C_i \ln \mu_i, \qquad C \equiv (I - B)^{-1} s. \tag{12}$$

The element C_i is the *markup centrality* weight of node i: it is large when node i is upstream-relevant for downstream uses that matter for final demand. The contribution $C_i \ln \mu_i$ therefore combines a node's direct wedge with its network position. Grouping and summing these contributions by upstream country or sector yields the decompositions reported in the empirical sections. When reporting aggregate markups in levels, we use the corresponding final-demand-weighted geometric mean, $\exp\left(s^{\top} \ln x\right)$, which is the natural aggregator in the log-linear Cobb-Douglas system.

4 | Data and Input-Output Objects

4.1 | WIOD Input-Output Tables

We use the World Input–Output Database (WIOD) 2016 release (Timmer et al. 2015; WIOD 2021). The WIOD world input-output tables (WIOTs) provide, for each year, a balanced matrix of intermediate transactions across country-industry nodes and a block of final demand. The baseline year is 2014, the last year in the WIOD 2016 time span.

The WIOT classification contains 56 industries and 44 country aggregates (43 countries plus a Rest-of-World aggregate), yielding $N = 56 \times 44 = 2464$ nodes. Each node is indexed by a country c and an industry code k; in theory, a node corresponds to an index j.

4.1.1 | Valuation and Interpretation

WIOD transactions are reported in current prices and are balanced across countries and industries. Like other input-output systems, WIOD incorporates taxes, subsidies, trade and transport margins, balancing adjustments, and imputed-rent conventions. The analysis takes the WIOT valuation as given. Therefore, the measured "markups" should be interpreted as wedges between observed prices and benchmark costs within the valuation convention of the WIOT. This caveat is quantitatively relevant for the interpretation of the large real-estate contribution reported below: part of that measured wedge can reflect valuation conventions and fiscal wedges, not only pure market power.

4.2 | Intermediate Coefficient Matrix and Final-Demand Weights

Let Z be the $N \times N$ matrix of intermediate expenditures, where z_{ij} is the expenditure of node j on intermediate inputs from node i. Let x be the $N \times 1$ vector of gross outputs. Define the value coefficient matrix A^{val} as in Equation (6):

$$A^{val}_{ij} \equiv \frac{z_{ij}}{x_j}.$$

In this normalization, each column of A^{val} indicates the intermediate input composition of a unit of gross output, valued at observed prices.

Let Y denote the final demand block. Let f be the $N \times 1$ vector of total final demand across categories and destinations (row sums of Y). The final-demand weight of node j is

$$s_j \equiv \frac{f_j}{\sum_i f_i}. \tag{13}$$

In practice, small negative entries can arise from balancing and rounding in world input-output tables. The implementation rescales weights to sum to one and verifies that any negative weights are numerically negligible.

4.2.1 | Derived Dataset

All results are produced from a derived dataset for 2014 that contains the key objects needed for compound-markup measurement and attribution. In particular, it reports direct markup measures (both raw and stabilized, together with stabilization diagnostics), compound markups computed from the network mapping in Equation (5), final-demand weights from Equation (13), markup centrality from Equation (12), and each node's upstream contribution $C_i \ln \mu_i$ to the aggregate log wedge. These objects are sufficient to reproduce the mapping and the exact decomposition without reconstructing intermediate steps from the raw WIOD files.

Small negative entries in the WIOD final-demand block can arise from balancing and rounding. In the baseline year, the total negative mass of final-demand weights is 4.9×10^{-8}; we retain these negligible entries and renormalize s so that weights sum to one.

4.2.2 | Feasibility Diagnostics

In a small number of country-industry cells, WIOD balancing can imply intermediate-input column sums in A^{val} that are very close to (or slightly above) one. Before constructing $\boldsymbol{B}$, the implementation applies a minimal column-rescaling step (defined in Equation (23)) that preserves within-column input composition while restoring a strictly positive value-added share floor $\alpha_{\min}$. In 2014, only three columns require rescaling, and their combined final-demand weight is 2.0×10^{-5} (0.002% of global final demand); reported aggregates are unchanged at the stated precision if the rescaling is omitted and the feasibility cap is allowed to bind.

5 | Estimating Direct Markups From WIOD Socio-Economic Accounts

This section describes how the vector of direct markups μ is constructed. The approach follows the logic of De Loecker and Warzynski (2012) but is implemented at the country-industry level using WIOD Socio-Economic Accounts (SEA). The objective is not to propose a new markup estimator but to produce a transparent and reproducible markup vector that can be mapped into compound markups.

5.1 | Direct Markup Construction

Let v denote a variable input used in production. Under standard assumptions, the producer's first-order condition for v implies that the markup equals the ratio of the output elasticity of v to the expenditure share of v in revenue:

$$\mu = \frac{\theta^v}{s^v}. \tag{14}$$

Here θ^v is the output elasticity of the variable input and s^v is the share of that input's expenditure in revenue. The identity is algebraic: it follows from cost minimization and requires only that there exists a wedge between price and marginal cost (a markup).

In annual SEA, the distinction between "variable" and "fixed" inputs is necessarily approximate. Following the logic of De Loecker and Warzynski (2012) but adapting to highly aggregated WIOD data, I define the variable input as a composite of intermediate inputs and labor compensation. WIOD SEA reports intermediate expenditures II, compensation of employees $COMP$, and gross output GO in current prices. Define the corresponding expenditure shares:

$$s^{II}_{c,k,t} \equiv \frac{II_{c,k,t}}{GO_{c,k,t}}, \tag{15}$$

$$s^{L}_{c,k,t} \equiv \frac{COMP_{c,k,t}}{GO_{c,k,t}}, \tag{16}$$

$$s^{var}_{c,k,t} \equiv \frac{II_{c,k,t} + COMP_{c,k,t}}{GO_{c,k,t}} = s^{II}_{c,k,t} + s^{L}_{c,k,t}. \tag{17}$$

Under a Cobb-Douglas benchmark, if intermediates and labor are treated as variables at the annual frequency, the elasticity of the variable composite is the sum of the intermediates and labor elasticities:

$$\theta^{var}_{k} \equiv \theta^{II}_{k} + \theta^{L}_{k}. \tag{18}$$

The raw direct markup estimate is then

$$\mu^{raw}_{c,k,t} \equiv \frac{\hat{\theta}^{var}_{k}}{s^{var}_{c,k,t}}, \qquad \hat{\theta}^{var}_{k} \equiv \hat{\theta}^{II}_{k} + \hat{\theta}^{L}_{k}. \tag{19}$$

This composite-input choice is pragmatic at the level of country-industry annual accounts. It reduces sensitivity to any single expenditure item and limits the frequency with which mechanically small denominators generate extreme markups. It also aligns the markup vector with the interpretation of $\ln \mu$ as a nonnegative wedge propagated through the production network; additional stabilizations described in Section 5.2 ensure feasibility of the network mapping.

5.1.1 | Estimating Elasticities of Intermediates and Labor

Elasticities are estimated using the WIOD SEA panel (countries × industries × years). WIOD SEA provides a gross-output quantity index GO^{QI}, an intermediate-input quantity index II^{QI}, labor input LAB, and capital input K, together with current-price values for gross output and intermediate inputs. Quantity indices are used so that production elasticities are estimated in real terms, reducing mechanical correlations between prices and quantities.

For each industry k, the baseline specification is a Cobb-Douglas production function in logs:

$$\ln GO^{QI}_{c,k,t} = \theta^{II}_{k} \ln II^{QI}_{c,k,t} + \theta^{K}_{k} \ln K_{c,k,t} + \theta^{L}_{k} \ln LAB_{c,k,t} + \eta_c + \tau_t + \varepsilon_{c,k,t}, \tag{20}$$

where η_c and τ_t are country and year fixed effects. Estimation is performed separately by industry k to allow elasticities to differ across industries, reflecting differences in technology and input intensity. This separate-by-industry approach is also consistent with the use of WIOD's industry classification, which is relatively coarse and may mask within-industry heterogeneity; allowing elasticities to differ across industries is therefore a minimal flexibility.

5.1.2 | Identification and Limitations

At the country-industry level, Equation (20) should be interpreted as a reduced-form mapping between input quantities and output quantities, not as a fully identified structural production function. Country and year fixed effects absorb time-invariant cross-country differences and global shocks. Remaining endogeneity concerns (e.g., input choices correlated with productivity shocks) are not fully addressed at this level of aggregation. The goal is not causal identification of elasticities but a plausible and transparent mapping from observables to markups suitable for network propagation.

5.1.3 | Plausibility and Benchmarking

The resulting markup vector is intended as a conservative input for the network mapping rather than as a definitive estimate of firm-level market power. In 2014, after stabilization, the median direct markup is 1.002 and the 99th percentile is 2.319; the final-demand-weighted aggregate direct markup is $\mu^{agg} = 1.076$ (Table 3). These magnitudes are below typical firm-level averages (e.g., De Loecker et al. (2020)), as expected under aggregation and national-accounting measurement.

5.1.4 | Industry Coverage and Exclusions

Some WIOD industries correspond to public services or activities where standard production-function interpretation is less meaningful. When μ^{raw} cannot be computed reliably (because elasticities or required inputs are missing or nonpositive), the implementation sets the direct markup to one.

Estimated elasticities are economically plausible; intermediate-input elasticities tend to be higher in manufacturing than in services.

5.2 | Stabilization and Feasibility

Two issues arise when implementing Equation (19) in world input-output data:

i. small variable-input expenditure shares can mechanically generate extreme markups, and

ii. the mapping from direct to compound markups requires the network price system to be productive (invertible).

5.2.1 | Flooring and Missing Values

In principle, Equation (14) permits $\mu < 1$ if the estimated elasticity is smaller than the expenditure share. In micro data, $\mu < 1$ can reflect measurement error, nonconstant returns, or misspecification of the variable input. In the present context, where markups are measured at a high level of aggregation, $\mu < 1$ is best interpreted as noise rather than as evidence of systematic "negative market power."

The baseline implementation, therefore, floors raw markups at one:

$$\mu_j^{floor} = \max\{\mu_j^{raw}, 1\}.$$

If μ_j^{raw} is missing, μ_j^{floor} is set to one. This conservative choice avoids generating negative log wedges and imposes the economically natural restriction that markups are weakly above one before network propagation.

5.2.2 | Shrinkage for Very Small Variable-Input Shares

When s^{var} is small, the ratio in Equation (19) can mechanically explode even if the elasticity estimate is moderate. Because variable-input shares can be small in some service sectors, a small number of extreme markups can dominate network propagation and aggregate decompositions. To reduce sensitivity to small denominators, the implementation uses *shrinkage* toward $\mu = 1$ for observations with very low variable-input shares.

Let s_0 denote a shrinkage threshold (baseline $s_0 = 0.10$). Define a reliability weight

$$r_{c,k,t} \equiv \min\left\{1, \frac{s_{c,k,t}^{var}}{s_0}\right\}. \tag{21}$$

Then define a shrunk markup by shrinking *log* markups toward zero:

$$\ln \mu^{shrunk} \equiv r \cdot \ln \mu^{floor}. \tag{22}$$

This operation has two desirable properties. First, it is scale-free: it shrinks proportional deviations from one, rather than deviations in levels. Second, it is conservative: it does not alter markups when variable-input shares are moderate ($r = 1$), but it limits the influence of observations with extremely small variable-input shares ($r < 1$).

5.2.3 | Winsorization of the Upper Tail

Even with shrinkage, a small number of nodes can exhibit very large markups due to outliers in elasticities or shares. Following a standard robust-statistics device for limiting the influence of extreme observations (Tukey and McLaughlin 1963; Wilcox 2016), the baseline implementation therefore winsorizes μ^{shrunk} at percentile $p = 0.995$:

$$\mu_j^{wins} = \min\{\mu_j^{shrunk}, Q_p(\mu^{shrunk})\},$$

where $Q_p(\cdot)$ denotes the empirical p-quantile. The goal is not to discard high markups as implausible, but to prevent a handful of extreme values from dominating network propagation and to ensure numerical stability.

5.2.4 | Feasibility Cap and Value-Added Share Floor

The mapping from direct to compound markups requires the matrix $\boldsymbol{I} - \boldsymbol{B}^{\top}$ to be invertible, equivalently $\rho(\boldsymbol{B}^{\top}) < 1$. In a Cobb-Douglas benchmark, this holds when each node has a strictly positive value-added share in *resource* costs:

$$\alpha_j = 1 - \sum_i B_{ij} > 0.$$

Using Equation (7) and letting A^{bal} denote the (column-rescaled) value-coefficient matrix used in the computations, the column sum of $\boldsymbol{B}$ is

$$\sum_i B_{ij} = \mu_j \sum_i A_{ij}^{bal}.$$

In practice, a few WIOT columns can have intermediate-input column sums close to (or slightly above) one because balancing residuals are absorbed by the value-added block. Before constructing $\boldsymbol{B}$, we apply a minimal column rescaling: for any column j with $\sum_i A_{ij}^{val} > 1 - \alpha_{\min}$, rescale that column by γ_j:

$$\gamma_j \equiv \min\left\{1, \frac{1 - \alpha_{\min} - \varepsilon}{\sum_i A_{ij}^{val}}\right\}, \tag{23}$$

and the coefficient matrix used in the computations is $A^{bal} \equiv A^{val}\,\mathrm{diag}(\gamma)$. This preserves within-column input composition and reallocates only the residual intensity to value added. In 2014, the adjustment applies to only three WIOT columns.

To ensure feasibility and numerical stability, the implementation imposes a lower bound $\alpha_{\min}$ on the value-added share:

$$\alpha_j \geq \alpha_{\min} > 0. \tag{24}$$

Operationally, this is enforced by a cap on μ_j:

$$\mu_j \leq \mu_j^{cap} \equiv \frac{1 - \alpha_{\min}}{\sum_i A_{ij}^{bal}}. \tag{25}$$

The final direct markup used in the compound-markup computation is

$$\mu_j^{used} \equiv \min\{\mu_j^{wins}, \mu_j^{cap}\}. \tag{26}$$

In the baseline implementation, the cap binds rarely; when it binds, it reflects a column of A^{val} with a very high intermediate share (sometimes above one), which is best interpreted as a measurement or balancing artifact in the WIOT.

5.2.5 | Diagnostic Summary

Table 1 summarizes how the stabilization pipeline reshapes the direct-markup distribution. In 2014, 986 SEA observations (44.2%) have $\mu^{raw} < 1$ and are floored at one; 231 WIOT nodes (9.4%) have missing markups imputed to one; and 1217 nodes (49.4%) therefore enter the network mapping at exactly one after the full stabilization pipeline. At the same time, shrinkage and winsorization compress the upper tail, reducing the maximum direct markup from 6.202 in raw form to 3.026 in the final vector

TABLE 1 | Stabilization diagnostics for the direct markups μ (2014).

Stage/diagnostic	Value
SEA observations with defined raw markups	2233
Raw $\mu^{raw} < 1$ (SEA)	986 (44.2%)
WIOT nodes with missing μ imputed to 1	231 (9.4%)
Final stabilized $\mu = 1$ (WIOT nodes)	1217 (49.4%)
Maximum raw direct markup	6.202
Maximum stabilized direct markup	3.026
WIOT columns rescaled for feasibility (A balancing)	3 (0.122%)
Feasibility cap binds (nodes)	1 node ($< 0.1\%$)

Note: The final stabilized markup vector incorporates flooring at one, reliability shrinkage for low variable-input shares, winsorization of the upper tail, imputation of missing values to one, and the feasibility cap used in the network mapping.

used in the network mapping. Only three WIOT columns require the minimal A-balancing rescaling, and the feasibility cap binds for one node (less than 0.1% of the sample), so the main quantitative effect of the pipeline is to discipline the lower tail and a small number of upper-tail outliers rather than to alter the bulk of moderate markups.

6 | Computing Compound Markups and Upstream Attribution

This section describes how the stabilized direct markups are mapped into compound markups and how aggregate wedges are decomposed into upstream contributions.

6.1 | Network Mapping and Compound Markups

Given the (column-balanced) value coefficient matrix A^{bal} from the WIOT and the stabilized markup vector μ^{used}, construct

$$B = A^{bal} \operatorname{diag}(\mu^{used}).$$

The feasibility cap in Equation (25) ensures that the column sums of B are bounded above by $1 - \alpha_{\min}$. This implies a uniform lower bound on value-added shares:

$$\alpha_j = 1 - \sum_i B_{ij} \geq \alpha_{\min}.$$

A strictly positive value-added share floor rules out nearly singular price systems and yields practical numerical benefits. In particular, it implies bounds on the Neumann series in Equation (8) and therefore bounds on the size of amplification effects, as discussed in Colonescu (2026).

The implied value-added share in resource costs is

$$\alpha_j \equiv 1 - \sum_i B_{ij}. \tag{27}$$

6.1.1 | Solving for Compound Markups

The compound markups satisfy the linear system Equation (5):

$$(I - B^{\top}) \ln \kappa = \ln \mu^{used}. \tag{28}$$

Rather than inverting the matrix, the implementation solves Equation (28) using standard linear solvers. This yields $\ln \kappa$ and hence κ. Numerical accuracy is verified by checking that the residual $(I - B^{\top}) \ln \kappa - \ln \mu^{used}$ is close to zero.

For each node, the ratio κ/μ measures the extent to which upstream markups embedded in intermediate input prices amplify the wedge beyond the node's own direct markup. In the two-stage chain, κ/μ reduces to μ_i^{β}; in the network, it aggregates many upstream markups with weights determined by L.

6.2 | Upstream Attribution and Markup Centrality

Let s be final-demand weights defined in Equation (13). Markup centrality is computed from the dual system

$$(I - B)C = s, \tag{29}$$

which corresponds to $C = (I - B)^{-1} s$ in Equation (12). The contribution of node i to the aggregate log wedge is

$$\text{contrib}_i \equiv C_i \ln \mu_i^{used}. \tag{30}$$

Summing across nodes yields $\ln \kappa^{\text{agg}} = \sum_i \text{contrib}_i$ exactly. This is the empirical counterpart of the theoretical decomposition: it assigns the aggregate final-demand wedge to upstream sources in a way that is consistent with network propagation.

Two additional diagnostic objects help interpret the cross-section. The row sum of the Leontief inverse in the price system, $\sum_i L_{ji}$, summarizes how exposed node j is, on average, to upstream markup wedges. The upstream contribution term $C_j \ln \mu_j$ (defined in Equation (30)) measures node j's contribution to the aggregate wedge; by construction, $\sum_{j=1}^{N} C_j \ln \mu_j = \ln \kappa^{\text{agg}}$.

7 | Results

This section expands the baseline empirical results along four dimensions: the node-level distributions of raw, stabilized, and compound markups; final-demand aggregates and propagation rounds; a broad manufacturing services comparison; and the upstream-node and international-incidence decompositions.

7.1 | Distributions and Cross-Sectional Amplification

Figure 1 juxtaposes the two distributional comparisons that matter most for interpretation. Panel (a) compares the raw and stabilized direct-markup distributions; panel (b) compares the stabilized direct-markup distribution to the compound-markup

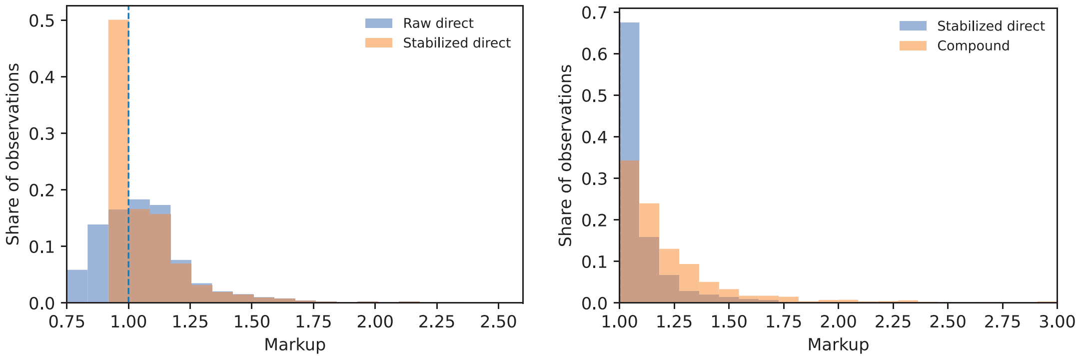

(a) Raw and stabilized direct markups. (b) Stabilized direct and compound markups.

FIGURE 1 | Markup distributions. The dashed vertical line marks $\mu = 1$. The stabilized direct-markup series and the compound-markup series are bounded below by one in the baseline implementation. For readability, the x-axis is truncated at 2.6 in panel (a) and at 3.0 in panel (b); Table 2 reports upper-tail percentiles and maxima. (a) Raw and stabilized direct markups, (b) Stabilized direct and compound markups.

TABLE 2 | Distributional summaries for direct and compound markups (2014).

Statistic	Raw direct μ^{raw}	Stabilized direct μ	Compound κ
Mean	1.058	1.108	1.270
Standard deviation	0.356	0.239	0.400
25th percentile	0.904	1.000	1.072
Median	1.027	1.002	1.143
75th percentile	1.143	1.130	1.304
Interquartile range	0.239	0.130	0.232
95th percentile	1.487	1.457	1.985
99th percentile	2.419	2.319	3.177
Maximum	6.202	3.026	5.120
Skewness	5.457	4.830	4.066

Note: Raw direct markups are reported for the 2233 SEA observations where μ^{raw} is defined. Stabilized direct markups and compound markups are reported for all 2464 WIOT nodes after flooring, shrinkage, winsorization, imputation of missing values to one, and the feasibility cap.

distribution generated by network propagation. Because the stabilized direct-markup vector and the resulting compound markups are bounded below by one in the baseline implementation, the figure uses overlaid histograms rather than kernel densities so that the plotted support matches the underlying data. To keep the bulk of the distributions legible in a side-by-side layout, panel (a) truncates the x-axis at 2.6 and panel (b) at 3.0; Table 2 reports the upper-tail percentiles and maxima that lie beyond those cutoffs.

The stabilization pipeline changes the direct-markup distribution in economically transparent ways. Flooring and imputation materially reshape the lower tail: 44.2% of defined raw markups fall below one, 9.4% of WIOT nodes are missing and are imputed to one, and 49.4% of nodes therefore enter the network mapping at exactly one. At the same time, shrinkage and winsorization compress the right tail, reducing the maximum direct markup from 6.202 in raw form to 3.026 in the final vector. The raw

TABLE 3 | Aggregate wedges and loop amplification (baseline, 2014).

	μ^{agg}	$\kappa^{(1),\text{agg}}$	κ^{agg}	$\kappa^{\text{agg}}/\kappa^{(1),\text{agg}}$	Δ^{HO}
Baseline	1.076	1.124	1.196	1.064	0.062

Note: $\Delta^{HO} \equiv \ln \kappa^{\text{agg}} - \ln \kappa^{(1),\text{agg}}$ is the higher-order (two or more rounds) contribution to the aggregate log wedge.

observations below one should therefore be read as estimation noise in the elasticity-share ratio at this level of aggregation, whereas the stabilized markups that enter the network mapping are conservatively bounded at one. These conservative adjustments are visible in panel (a) of Figure 1 and are summarized in Tables 1 and 2.

Even after those conservative adjustments, the network-propagated distribution is shifted sharply upward. The compound-markup mean is 1.270 versus 1.108 for the stabilized direct markup, the median rises from 1.002 to 1.143, and the 99th percentile rises from 2.319 to 3.177 (Table 2). The median amplification factor κ_i/μ_i is 1.103, indicating that embedded upstream wedges are quantitatively important well beyond a few extreme nodes.

7.2 | Aggregate Wedge and Propagation Rounds

The final-demand-weighted aggregate wedges are reported in Table 3. The aggregate compound markup is $\kappa^{\text{agg}} = 1.196$ versus an aggregate direct markup of $\mu^{\text{agg}} = 1.076$. The one-round benchmark $\kappa^{(1)}$ isolates the direct wedge plus the first upstream embedding; the ratio $\kappa^{\text{agg}}/\kappa^{(1),\text{agg}} = 1.064$ shows that higher-order propagation and feedback loops add a meaningful but smaller increment.

A Neumann-series decomposition of the walk expansion in Equation (9) shows rapid decay with path length: the direct term ($t = 0$) accounts for 40.9% of $\ln \kappa^{\text{agg}}$; adding one-round propagation ($t = 1$) raises the cumulative share to 65.5%; terms $t = 2$ and $t = 3$ add 13.9% and 8.1%; and all paths of length $t \geq 6$ jointly contribute 4.9%. Thus, most aggregate compounding is captured by the first few propagation rounds.

7.3 | Manufacturing Versus Services

To provide a broad sectoral contrast, I classify manufacturing as WIOD industries C10–C33 and services as G45–U; the remaining sectors (primary activities, utilities, and construction) are reported for completeness. For each downstream bundle g, Table 4 reports the final-demand share of that bundle, the group-specific aggregate wedges obtained by renormalizing final-demand weights within the bundle, the median node-level direct and compound markups within the bundle, the median amplification factor, and the share of the group-specific aggregate wedge accounted for by the top 50 upstream contributors. These group-specific aggregates are the direct broad-sector counterparts to the economy-wide aggregates reported above (Table 3).

Manufacturing displays the larger within-group wedge. Its group-specific direct aggregate is 1.102 and its compound aggregate is 1.323, implying amplification of about 20.0% relative to the direct benchmark. Services have lower direct and compound medians (1.000 and 1.087, respectively) and lower aggregate amplification (1.077), consistent with weaker intermediate intensity on average. At the same time, services account for 62.0% of final demand and, when grouped on the upstream side, 55.8% of the global aggregate wedge. Manufacturing contributes 31.6% of the global wedge. Thus, the global wedge is service-heavy in origin even though manufacturing bundles exhibit stronger within-group compounding. The service bundle is also more concentrated: the top 50 upstream contributors account for 65.1% of its aggregate log wedge, versus 57.0% for manufacturing.

7.4 | Upstream Contributors and International Incidence

The exact attribution $\ln \kappa^{\text{agg}} = \sum_i C_i \ln \mu_i$ yields an additive ranking of upstream nodes. Table 5 reports the top 20 contributors. The ranking combines both components of upstream importance: the level of the direct markup and the node's network position relative to final demand. The list, therefore, contains both high-markup service nodes (administrative support, retail, finance, and real estate) and large manufacturing contributors (food, chemicals, and machinery).

The exact attribution implies strong concentration. In 2014, the top 50 upstream nodes account for 54.0% of the aggregate log wedge, the top 100 for 70.0%, and the top 250 for 88.4%. Aggregating contributions by upstream country, the leading contributors are the United States (17.8% of $\ln \kappa^{\text{agg}}$), China (14.6%), India (12.7%), Mexico (8.7%), and Japan (6.7%), with the remaining 20.1% distributed across smaller contributors. By upstream sector, real estate activities account for 13.5% of the aggregate wedge, followed by food manufacturing (8.4%), financial services (8.3%), wholesale trade (7.1%), and retail trade (5.1%). Because WIOD embeds taxes, subsidies, trade and transport margins, and imputed rents in the valuation convention, the prominence of real estate should be read as an accounting wedge within the table rather than as a pure measure of monopoly power.

International incidence is obtained by recomputing final-demand weights by destination, $s^{(d)}$, and applying the same decomposition with destination-specific markup centrality $C^{(d)} = (I - B)^{-1} s^{(d)}$. Table 6 reports destination-specific aggregate wedges and the domestic versus foreign shares of $\ln \kappa^{\text{agg},(d)}$ for three large destinations.

For the United States, 20.6% of the destination wedge is attributable to upstream nodes located abroad, with Mexico (7.1% of the destination wedge) and China (3.2%) as the largest foreign sources. For Germany, the foreign share is 44.0%, with China (4.4%) and the United States (4.1%) as the leading foreign contributors. For the United Kingdom, the foreign share is 51.2%, with Germany (5.7%) and the United States (5.3%) as the largest foreign sources. The high U.S. domestic share is consistent with the size of the U.S. economy and its relatively lower import dependence, whereas Germany and especially the United Kingdom rely more heavily on foreign intermediates and foreign upstream services embodied in absorption. More generally, the foreign-incidence share rises with the imported content of destination final demand and falls with destination size, so the cross-country pattern closely tracks the openness gradient in the WIOT.

Using the destination-specific final-demand block in the full WIOT, the exact foreign share of the destination-specific compound wedge can be computed for all WIOD countries. The full ordered ranking is reported in Appendix Figure A1. ROW is omitted because its direct markups are imputed to one in the baseline, which mechanically yields a foreign share of 100%. The exact ranking displays a clear openness gradient. India has the lowest foreign-incidence share at 3.4%, followed by Indonesia (7.1%) and Mexico (7.8%). At the other end, Luxembourg has by far the highest foreign share at 93.4%, followed by Estonia and Latvia at 72.6% each. Germany's foreign share of 44.0% sits slightly above the sample median of 37.9%, the United Kingdom's 51.2% lies in the upper third of the distribution, and the United States' 20.6% remains near the lower end. The ranking therefore reinforces the interpretation of Table 6: foreign incidence is primarily governed by the import content of absorption and by the reliance of final expenditure on foreign upstream services and intermediates.

8 | Discussion

The compound-markup perspective reframes a direct-markup estimate in a world with extensive intermediate-input trade. A direct markup compares price to marginal cost evaluated at observed input prices; when upstream sectors price above marginal cost, observed intermediate prices already embed upstream wedges. Compound markups make that embedding explicit by comparing observed prices to a pure-cost benchmark in which all output markups in the network are set to one, holding technology and primary-input prices fixed.

In the WIOD 2014 implementation, the gap between μ^{agg} and κ^{agg} is economically meaningful (Table 3) and the upstream attribution is concentrated: a relatively small set of upstream nodes accounts for a large share of the aggregate log wedge. The new distributional results show that this conclusion is not driven solely by the upper tail of raw markups. The stabilization pipeline is conservative and creates a large mass at one in the direct-markup

TABLE 4 | Broad-sector comparison: manufacturing, services, and other sectors (2014).

Broad sector	FD share (%)	med. μ	med. κ	med. κ/μ	$\mu^{\text{agg},(g)}$	$\kappa^{\text{agg},(g)}$	Top 50 share (%)
Manufacturing	20.3	1.083	1.247	1.141	1.102	1.323	57.0
Services	62.0	1.000	1.087	1.072	1.073	1.155	65.1
Other	17.7	1.000	1.123	1.100	1.056	1.201	62.8

Note: For each broad sector g, $\mu^{\text{agg},(g)}$ and $\kappa^{\text{agg},(g)}$ are computed by renormalizing final-demand weights within that sectoral bundle. "Top 50 share" is the share of the corresponding group-specific aggregate log wedge accounted for by the top 50 upstream contributors under the exact decomposition. The manufacturing/services split, therefore, compares both downstream bundles and the concentration of their upstream origins.

TABLE 5 | Top contributing upstream nodes to the aggregate log wedge $\ln \kappa^{\text{agg}} = \sum_i C_i \ln \mu_i$ (2014).

Country	Sector	μ	κ	C	$C \ln(\mu)$	Share (%)
CHN	Manufacture of food products, beverages and tobacco products	1.189	1.417	0.027	0.005	2.63
MEX	Administrative and support service activities	3.026	4.178	0.004	0.005	2.54
IND	Retail trade, except of motor vehicles and motorcycles	3.026	4.412	0.004	0.005	2.54
JPN	Real estate activities	1.619	1.683	0.009	0.004	2.39
CHN	Financial service activities, except insurance and pension funding	1.427	1.518	0.012	0.004	2.35
ITA	Real estate activities	2.617	2.929	0.004	0.004	2.34
BRA	Real estate activities	3.026	3.224	0.003	0.003	1.82
USA	Real estate activities	1.078	1.128	0.038	0.003	1.58
IND	Wholesale trade, except of motor vehicles and motorcycles	3.026	4.412	0.003	0.003	1.56
CHN	Wholesale trade, except of motor vehicles and motorcycles	1.144	1.239	0.021	0.003	1.55
USA	Manufacture of chemicals and chemical products	1.339	1.599	0.009	0.003	1.54
IND	Financial service activities, except insurance and pension funding	2.767	3.905	0.003	0.003	1.45
USA	Wholesale trade, except of motor vehicles and motorcycles	1.118	1.174	0.023	0.003	1.43
USA	Manufacture of food products, beverages and tobacco products	1.189	1.421	0.014	0.002	1.34
USA	Financial service activities, except insurance and pension funding	1.247	1.319	0.010	0.002	1.26
USA	Electricity, gas, steam and air conditioning supply	1.446	1.566	0.006	0.002	1.19
USA	Motion picture, video and television program production, sound recording and music publishing activities; programming and broadcasting activities	1.524	1.779	0.005	0.002	1.15
ESP	Real estate activities	2.124	2.222	0.002	0.002	1.00
CHN	Manufacture of machinery and equipment n.e.c.	1.104	1.325	0.017	0.002	0.96
USA	Construction	1.111	1.228	0.016	0.002	0.96

Note: Columns report the stabilized direct markup μ_i, the compound markup κ_i, the markup-centrality weight C_i, and each node's contribution $C_i \ln \mu_i$ to the aggregate wedge. "Share" denotes each node's share of $\ln \kappa^{\text{agg}}$.

vector, yet the compound-markup distribution remains substantially to the right of the direct-markup distribution (Figure 1).

The broad-sector split is also informative. Manufacturing bundles exhibit stronger within-group amplification than service bundles (Table 4), which is consistent with heavier intermediate-input use and longer production chains. At the same time, service sectors account for a majority of the aggregate wedge on the upstream side because they are used pervasively throughout the network. In that sense, the aggregate compound wedge combines manufacturing-style propagation mechanics with service-sector centrality. The destination results point in the same direction.

TABLE 6 | Destination-specific aggregate compound wedges and domestic versus foreign incidence (2014).

Destination	$\kappa^{\text{agg},(d)}$	Domestic share (%)	Foreign share (%)
USA	1.153	79.4	20.6
DEU	1.132	56.0	44.0
GBR	1.099	48.8	51.2

Note: For each destination d, the table reports the final-demand-weighted aggregate compound wedge $\kappa^{\text{agg},(d)}$ and the shares of $\ln \kappa^{\text{agg},(d)}$ attributable to upstream nodes located domestically versus abroad, using the exact decomposition $\ln \kappa^{\text{agg},(d)} = \sum_i C_i^{(d)} \ln \mu_i$.

The exact benchmark destinations already show that foreign incidence can be substantial even for large absorbing economies, and the all-destination ranking in Appendix Figure A1 places the highest foreign-incidence shares in small open economies while leaving large, relatively self-contained destinations near the lower end of the distribution.

Several limitations matter for interpretation. Direct markups are measured at the country-industry level and inherit the limitations of production-function approaches at high levels of aggregation; they should be viewed as transparent inputs for the network mapping rather than definitive measures of firm-level pricing power. The mapping itself is comparative static and holds the input-output structure fixed, abstracting from substitution or reallocation in response to changes in markups. Because WIOD incorporates taxes, subsidies, trade and transport margins, and imputed rents in its valuation convention, the empirical objects are wedges within the WIOT accounting framework rather than pure monopoly margins. This is especially relevant for sectors such as real estate, which account for a large share of the aggregate wedge.

9 | Robustness and Extensions

The baseline implementation uses three conservative stabilizations: a minimum value-added share $\alpha_{\min} = 0.01$ to guarantee feasibility of the price system, a shrinkage threshold $s_0 = 0.10$ to attenuate mechanically extreme markups when variable-input shares are very small, and winsorization of the upper tail at $p = 0.995$. Sensitivity checks varying $(\alpha_{\min}, s_0, p)$ yield aggregate compound markups in a narrow range (approximately 1.175–1.203) and similar higher-order amplification factors (about 1.06), indicating that the headline results are not an artifact of a single tuning choice. Using an intermediates-only markup base produces substantially larger network wedges (aggregate $\kappa^{\text{agg}} \approx 2.08$), illustrating that how "variable inputs" are defined matters quantitatively, even though the network-propagation and attribution logic is unchanged.

A separate concern is bias in the estimated production elasticities that feed into the direct-markup vector. To gauge the magnitude of that channel, Table 7 rescales all estimated variable-input elasticities uniformly by ±5% and ±10%, then reapplies the full stabilization pipeline and the network mapping. A +5% elasticity shift raises the aggregate direct markup from 1.076 to 1.100, but it raises the aggregate compound markup more strongly,

TABLE 7 | Sensitivity of aggregate wedges to uniform shifts in estimated variable-input elasticities.

Elasticity shift	μ^{agg}	$\kappa^{(1),\text{agg}}$	κ^{agg}	$\kappa^{\text{agg}} / \mu^{\text{agg}}$
−10%	1.038	1.058	1.084	1.044
−5%	1.055	1.087	1.130	1.071
Baseline	1.076	1.124	1.196	1.111
+5%	1.100	1.168	1.281	1.165
+10%	1.126	1.217	1.388	1.232

Note: Each row rescales the estimated variable-input elasticities uniformly, reconstructs the direct-markup vector, reapplies flooring, shrinkage, winsorization, imputation, and the feasibility cap, and then recomputes the network mapping. The exercise is a back-of-the-envelope sensitivity check rather than a structural identification test.

from 1.196 to 1.281. A +10% shift raises the aggregate compound markup to 1.388. Under downward shifts, the aggregate compound markup falls to 1.130 and 1.084 for the −5% and −10% cases, respectively. This is the practical sense in which network propagation can amplify elasticity-level measurement bias: the direct-markup input moves first, and the Leontief propagation map then magnifies part of that movement through repeated upstream embedding.

More broadly, the framework is modular: any vector of direct markups at the country-industry level can be fed into the mapping in Section 3 to obtain compound markups and upstream attribution. This makes it straightforward to

i. plug in alternative markup estimates (including aggregates of firm-level markups when micro data are available), and

ii. extend the analysis over time using WIOD's panel structure to decompose changes in κ^{agg} into changes in direct markups versus changes in network structure.

10 | Conclusion

This paper proposes a theory-based, implementable measure of *compound markups* in global production networks and an exact upstream attribution of final-demand wedges. By comparing observed prices to a pure-cost counterfactual that removes all upstream markups, the framework separates local pricing wedges from the wedges embedded in intermediate-input prices and provides a network-consistent incidence accounting via markup centrality. In WIOD 2014, network embedding raises the aggregate final-demand wedge from $\mu^{\text{agg}} = 1.076$ to $\kappa^{\text{agg}} = 1.196$; manufacturing bundles exhibit the strongest within-group amplification, while service sectors account for most of the aggregate wedge because they are pervasive upstream suppliers. Exact destination-specific decompositions show that foreign incidence ranges from 3.4% in India to 93.4% in Luxembourg, with the United States at 20.6%, Germany at 44.0%, and the United Kingdom at 51.2%. Destination-specific weights reveal substantial international incidence for large absorbing economies, and the broader cross-destination ranking places foreign incidence at the low end for large, relatively self-contained destinations and at the high end for small open economies. The approach offers a transparent bridge between direct-markup measurement and production-network analysis and can be combined with alternative markup inputs and applied over time as input-output structures evolve.

Funding

The author has nothing to report.

Ethics Statement

Constantin Colonescu is the sole author of this article and is responsible for all aspects of the research.

Conflicts of Interest

The author declares no conflicts of interest.

Data Availability Statement

The data that support the findings of this study are available in the World Input-Output Database 2016 Release, 2000–2014 at https://dataverse.nl/dataset.xhtml?persistentId=doi:10.34894/PJ2M1C. These data were derived from the following resources available in the public domain:—WIODS in R, https://dataverse.nl/file.xhtml?fileId=199101&version=2.1—Socio-Economic Accounts, https://dataverse.nl/file.xhtml?fileId=199095&version=2.1—SUT International, https://dataverse.nl/file.xhtml?fileId=199100&version=2.1.

References

Acemoglu, D., V. M. Carvalho, A. Ozdaglar, and A. Tahbaz-Salehi. 2012. "The Network Origins of Aggregate Fluctuations." *Econometrica* 80, no. 5: 1977–2016.

Antràs, P., D. Chor, T. Fally, and R. Hillberry. 2012. "Measuring the Upstreamness of Production and Trade Flows." *American Economic Review* 102, no. 3: 412–416.

Baqaee, D. R., and E. Farhi. 2020. "Productivity and Misallocation in General Equilibrium." *Quarterly Journal of Economics* 135, no. 1: 105–163.

Bigio, S., and J. La'O. 2020. "Distortions in Production Networks." *Quarterly Journal of Economics* 135, no. 4: 2187–2253.

Carvalho, V. M. 2014. "From Micro to Macro via Production Networks." *Journal of Economic Perspectives* 28, no. 4: 23–48.

Colonescu, C. 2021. "Price Markups and Upstreamness in World Input–Output Data." *Acta Universitatis Sapientiae, Economics and Business* 9, no. 1: 71–85. https://doi.org/10.2478/auseb-2021-0004.

Colonescu, C. 2023. "Compounded Markups in Complex Market Structures." *Athens Journal of Business & Economics* 9, no. 2: 131–144. https://doi.org/10.30958/ajbe.9-2-1.

Colonescu, C. 2026. "Compound Markups: A Theory of Price Amplification in Production Networks." Preprint. https://papers.ssrn.com/sol3/papers.cfm?abstract_id=6257378.

De Loecker, J., J. Eeckhout, and G. Unger. 2020. "The Rise of Market Power and the Macroeconomic Implications." *Quarterly Journal of Economics* 135, no. 2: 561–644.

De Loecker, J., and F. Warzynski. 2012. "Markups and Firm-Level Export Status." *American Economic Review* 102, no. 6: 2437–2471.

Hall, R. E. 1988. "The Relation Between Price and Marginal Cost in U.S. Industry." *Journal of Political Economy* 96, no. 5: 921–947.

Leontief, W. 1941. *The Structure of the American Economy, 1919–1929: An Empirical Application of Equilibrium Analysis*. Harvard University Press.

Miller, R. E., and P. D. Blair. 2009. *Input–Output Analysis: Foundations and Extensions*. 2nd ed. Cambridge University Press.

Spengler, J. J. 1950. "Vertical Integration and Antitrust Policy." *Journal of Political Economy* 58, no. 4: 347–352.

Timmer, M. P., E. Dietzenbacher, B. Los, R. Stehrer, and G. J. de Vries. 2015. "An Illustrated User Guide to the World Input–Output Database: The Case of Global Automotive Production." *Review of International Economics* 23, no. 3: 575–605.

Tirole, J. 1988. *The Theory of Industrial Organization*. MIT Press.

Tukey, J. W., and D. H. McLaughlin. 1963. "Less Vulnerable Confidence and Significance Procedures for Location Based on a Single Sample: Trimming/Winsorization 1." *Sankhyā: The Indian Journal of Statistics Series A* 25: 331–352.

Wilcox, R. R. 2016. *Introduction to Robust Estimation and Hypothesis Testing*. 4th ed. Academic Press (Elsevier).

WIOD. 2021. "World Input–Output Database 2016 Release, 2000–2014." DataverseNL. https://doi.org/10.34894/PJ2M1C.

Appendix A

All-Destination Foreign-Incidence Ranking

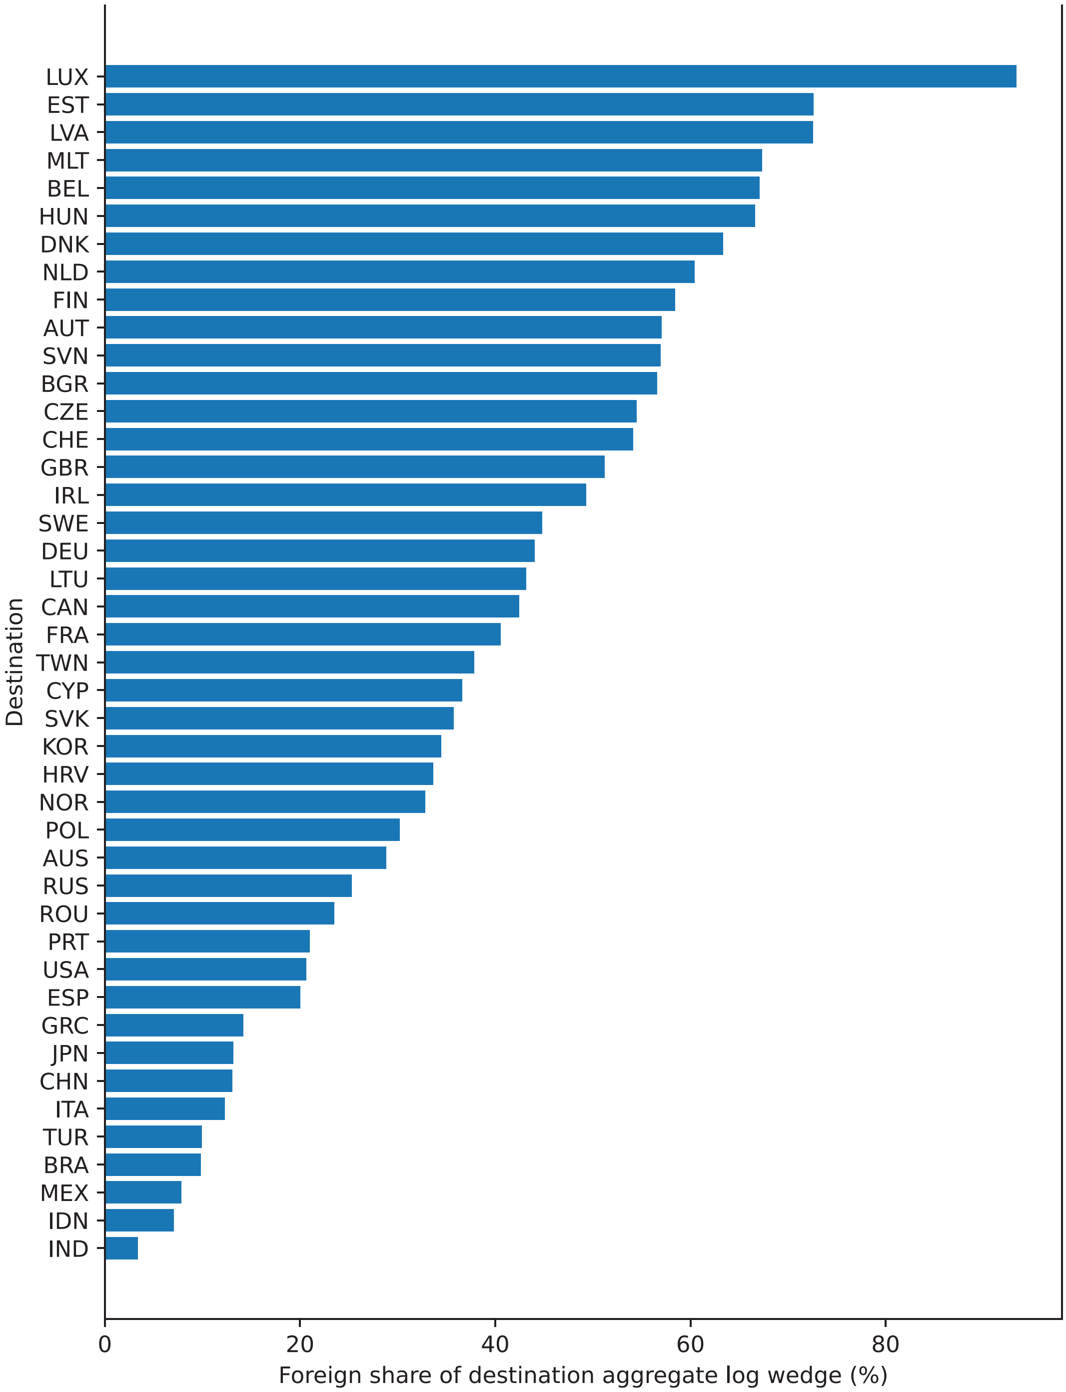

FIGURE A1 | Ordered exact foreign-incidence shares across WIOD destinations. The foreign share is the percentage of $\ln \kappa^{\text{agg},(d)}$ attributable to upstream nodes located outside destination d. ROW is omitted because its direct markups are imputed to one in the baseline.