The Micro Effects of Aggregate Shocks in Endogenous Trade Networks *

Carlos Bolivar

University of Minnesota

May 31, 2026

Link to Latest Draft

Abstract

This paper studies how trade liberalization affects aggregate risk. I develop a multi-country model where firms choose suppliers trading off cost against risk. Calibrating to Long-WIOD data, I find that openness and risk are non-monotone: moderate liberalization reduces GDP variance through diversification, but when trade barriers are sufficiently low, production concentrate sourcing in productive central suppliers, raising efficiency but amplifying shock propagation and aggregate risk. In addition, under incomplete markets, firms price the risk against domestic rather than world conditions, generating inefficiencies in the production network.

Keywords: Aggregate shocks, production networks, international trade, uncertainty, global financial crisis.

JEL classification: F44, F14, E32, D85, C68.

*I thank Manuel Amador, Timothy Kehoe, Javier Bianchi, and Marco Basseto for all their continuous support, guidance, and helpful comments on various stages of the project. I also thank Zhen Huo, Ilenin Kondo, Doireann Fitzgerald and Joseph Mullins for helpful conversations, as well as seminar participants at the University of Chile, CUNEF, the Richmond Fed, Vigo 2025 conference, Wisconsin-Minnesota workshop, and the trade workshop at the University of Minnesota.

1 Introduction

Trade increases productivity by connecting producers to the most efficient suppliers in the world. As trade barriers fall, production concentrates in a smaller set of highly productive hubs that serve as key input suppliers to many countries, such as Taiwan in semiconductor manufacturing. This concentration is at the core of the gains from trade: the most productive producers expand their reach to different locations, and buyers benefit from lower costs. But concentration comes at a price. As shown by Acemoglu, Carvalho, Ozdaglar, and Tahbaz-Salehi (2012), when a small number of industries account for a disproportionate share of input-output linkages, granular shocks generate stronger aggregate fluctuations. The force that makes trade beneficial, the emergence of productive central suppliers, is the same force that makes the global economy more volatile. In this paper, I argue that the tension between the productivity gains from concentration and the aggregate risk it creates has first-order consequences for aggregate volatility.

To study this tension, I first document how that network concentration and aggregate risk changed over fifty years of global input-output data. I then build a multi-country model in which the network and aggregate risk are jointly determined, and show that trade liberalization has a non-monotone effect on GDP variance: moderate liberalization reduces it through diversification, but deep liberalization reverses the gains as sourcing reconcentrates in central suppliers.

I extend Kopytov, Mishra, Nimark, and Taschereau-Dumouchel (2024) to an open economy with heterogeneous trade costs, endogenous labor supply, and incomplete markets. The key equilibrium object is a *risk-adjusted price*: the sum of the supplier's expected production cost, a trade cost, and an insurance premium that rises with the covariance between the supplier's price and the buyer's domestic consumption. Firms face uncertainty about global, country-specific, industry-specific, and country-industry pair (granular) productivity shocks. Before these shocks are realized, they choose suppliers to minimize risk-adjusted costs.

The main mechanism of the model works as follows. Trade distortions fragment the global production network into domestic clusters. Reducing these distortions has two effects. First, it allows firms to diversify their supply in more countries, as in Caselli, Koren, Lisicky, and Tenreyro (2020). Second, it facilitates the emergence of globally central suppliers: highly productive industries that become deeply interconnected because lower trade barriers make them accessible to many trading partners.

Diversification lowers exposure to granular and country-specific shocks, thus reducing aggregate risk. However, as the network reorganizes around central suppliers, shocks to those producers are strongly propagated through the Leontief inverse to their trading partners, increasing aggregate risk.

The net effect of lower trade on aggregate risk depends on how firms adjust their sourcing decisions. In the model, the optimal response of the firms changes over the business cycle. Central suppliers offer lower expected prices, but their insurance premium, the covariance between the supplier's price and the buyer's domestic consumption, rises with centrality, so when global uncertainty rises, above-average centrality suppliers become relatively more expensive and lose market share. In episodes like 2008, this reallocation partially offsets the amplification mechanism. However, under incomplete markets, firms discount profits using the domestic pricing kernel. As a result, the competitive equilibrium is inefficient: each firm prices risk using the covariance with its own domestic price index, while a planner would account for the effect on all countries' price indices.

The model was estimated using WIOD data for 24 countries and 23 industries from 1967 to 2014. I recover bilateral trade costs as structural residuals from observed network shares and risk-adjusted prices. This extends the inversion approach of Bonadio, Huo, Levchenko, and Pandalai-Nayar (2025) to a setting where the time variation in the expenditure shares reflects changes in trade costs, productivity, and aggregate risk. Next, I validate a mechanism of the model in the data: in a panel regression, above-average-centrality industries lose statistically significant market share when global variance rises, and the model reproduces this pattern with similar magnitude.

Finally, I present two counterfactuals. First, progressively removing trade costs from their 2014 levels toward free trade reveals a non-monotone relationship between openness and aggregate risk: GDP variance falls by up to 13% when roughly two-thirds of remaining trade barriers are removed, then partially rebounds to 2% below baseline as the economy approaches full free trade. Second, evaluating the observed 1967-2014 trade liberalization, I find that it reduced GDP variance for 15 of 24 economies, including the United States, Japan, and the United Kingdom, and increased variance for the remaining nine, notably Mexico, India, China, and Brazil.

Related Literature The main precursors are Kopytov et al. (2024) and Huo, Levchenko, and Pandalai-Nayar (2024). I build on Huo et al. (2024) by endogenizing the global production network. Relative to Kopytov et al. (2024), add trade cost, incomplete markets, and endogenous labor supply. Also, I replace the quadratic adjustment cost with an entropy (KL divergence) cost, which yields a softmax closed-form solution for optimal shares and eliminates corner solutions, making the computation of the competitive equilibrium tractable.

The main quantitative results are related most closely to Caselli et al. (2020), who decompose trade's effect on GDP volatility into a diversification and a specialization channel and find that diversification dominates for most countries. I introduce a third channel, network concentration, through which trade liberalization creates globally central suppliers, amplifying shock propagation. I show that for large hub economies this channel can reverse the sign of the net effect of trade on aggregate risk.

Two closely related papers are Fan and Luo (2025) and Kleinman, Liu, and Redding (2025). Fan and Luo (2025) study how the uncertainty of the tariff *policy* reshapes sourcing; I study *productivity* shocks that propagate through the production network, so riskiness is determined by a supplier's network position rather than its bilateral tariff exposure. Kleinman et al. (2025) study trade of final goods under uncertainty; while I study both final and intermediate goods, where the input-output links amplify shocks through the Leontief inverse.

The paper also connects to the growing literature on how incomplete markets modify the gains from trade. Fitzgerald (2025) shows that for volatile developing economies, balanced-trade gains are only one-third of complete market gains because agents cannot smooth consumption across states ex post. Allen and Atkin (2022) document that incomplete insurance halves the gains from trade integration for Indian farmers by distorting ex-ante crop portfolio choices. Adamopoulos and Leibovici (2024) find a pecuniary externality in food trade under incomplete markets that justifies large import tariffs for food-importing nations. My paper contributes a distinct mechanism: incomplete markets prevent firms from pricing the aggregate risk consequences of their sourcing decisions, leading to a misallocation of risk across the network rather than insufficient ex-post consumption smoothing.

The empirical motivation connects to work on time-varying volatility and production networks. Bloom, Floetotto, Jaimovich, Saporta-Eksten, and Terry (2018) document that plant-level TFP dispersion roughly doubles during recessions; I find an analogous pattern in an international panel of 24 countries and 23 industries. The time variation in network concentration is related most closely to Carvalho and Gabaix (2013), who show that changing Domar weights account for the Great Moderation and its partial reversal. I document an analogous variation in the global production network at business cycle frequency. On the comovement side, Kose and Yi (2001) show that trade in intermediates raises bilateral GDP comovement beyond what final-goods trade implies; Johnson (2014) confirm this using value-added trade flows. vom Lehn and Winberry (2022) provide complementary evidence that the U.S. investment network has changed over time with consequences for business cycle comovement.

The paper also relates to several broader literatures: on endogenous production networks (Arkolakis, Huneeus, and Miyauchi, 2025; Acemoglu and Azar, 2020; Elliott, Golub, and Leduc, 2022; Taschereau-Dumouchel, 2025; Acemoglu and Tahbaz-Salehi, 2025), on aggregation and distortions in networks (Acemoglu et al., 2012; Jones, 2011; Baqaee and Farhi, 2019; Bigio and La'O, 2020; Baqaee and Farhi, 2024), on trade and aggregate volatility (Caselli et al., 2020; di Giovanni, Levchenko, and Méjean, 2018), on tail risk in production networks (Dew-Becker, 2023; Dew-Becker, Tahbaz-Salehi, and Vedolin, 2021), on the estimation of trade cost (Eaton and Kortum (2002); Waugh (2010)), and on the granular origins of aggregate fluctuations (Gabaix, 2011; di Giovanni, Levchenko, and Méjean, 2014).

Outline. Section 2 presents the empirical motivation. Section 3 presents the model. Section 4 studies the optimal network choice and Section 5 the effect of aggregate shocks. Section 6 describes the calibration. Section 7 presents the quantitative results. Section 9 concludes.

2 Empirical Motivation

This section documents two sets of empirical patterns using the Long-run WIOD (Woltjer, Gouma, and Timmer, 2021; Timmer, Dietzenbacher, Los, Stehrer, and De Vries, 2015), which covers 24 countries, 23 industries, and 48 years from 1967 to 2014.[1] First, aggregate risk is time-varying and driven by a common global component. Second, the global production network changed substantially over the same period, both in its degree of concentration and in the geographic distribution of central nodes.

[1]These economies account for over 85% of world GDP throughout the sample.

2.1 Time-Varying Aggregate Risk

I measure aggregate risk by decomposing value-added growth into a global component $\tilde{f}_t$ common to all country-industry pairs, a country component $\tilde{f}_{c,t}$, an industry component $\tilde{f}_{i,t}$, and a granular residual $\tilde{u}_{ci,t}$. I estimate these components via PCA following Caselli et al. (2020) and obtain time-chaining variances with GARCH(1,1) for each series. Figure 1 shows the resulting decomposition for 1967–2014.

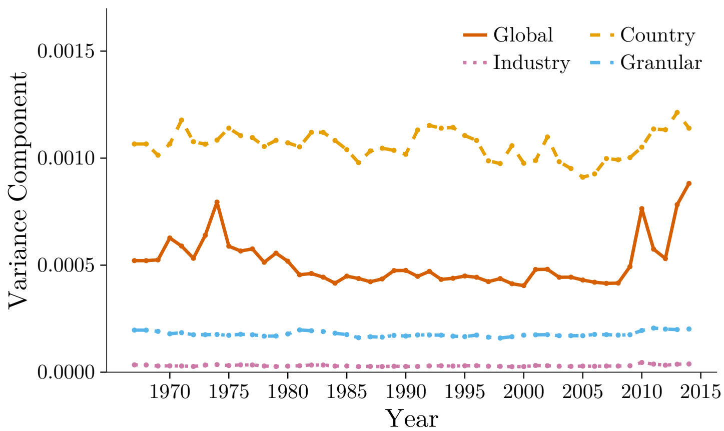

Figure 1: Time-Varying Aggregate Risk

Notes: The plot shows the time-varying variance of each component in the reduced-form decomposition $\tilde{f}_{ci,t} = \tilde{f}_t + \tilde{f}_{c,t} + \tilde{f}_{i,t} + \tilde{u}_{ci,t}$ of observed value-added growth, estimated by PCA following Caselli et al. (2020) with time-varying variances from a panel GARCH(1,1). WIOD data: 24 countries, 23 industries, 1967–2014.

The global component drives the time variation in aggregate risk: its variance shows the largest swings over the sample, while the country, industry, and granular components remain relatively stable. The country component is larger in level than the global component throughout most of the sample. In the sample, we find episodes of high aggregate volatility are global events in which the variance of the common factor affecting all country-industry pairs increases, rather than coincident national or granular shocks.

The decomposition of value-added reveals two model-relevant patterns. First, the country component remains quantitatively important, consistent with Caselli et al. (2020), who show that country-specific productivity shocks are more important for aggregate volatility compared to industry-specific shocks. Second, a global component emerges clearly that is spiking at each crisis. This is consistent with Huo et al. (2024), who show that production-network transmission alone cannot match the observed degree of international comovement, so correlated shocks are required to match the observed comovement in the data. The global component in value-added growth thus motivates the study of the effect of changes in the variance of a global shock.

2.2 A Changing Global Production Network

As shown by Acemoglu et al. (2012) concentration is a key measure for how strongly the network propagate productivity shocks. I measure network concentration using the Herfindahl-Hirschman Index of sales shares (Domar weights), $H_t = \sum_{c,i} \omega_{ci,t}^2$, where $\omega_{ci,t}$ denotes the ratio of country-industry (c, i) gross output to world GDP. Figure 2 plots H_t over 1967–2014.

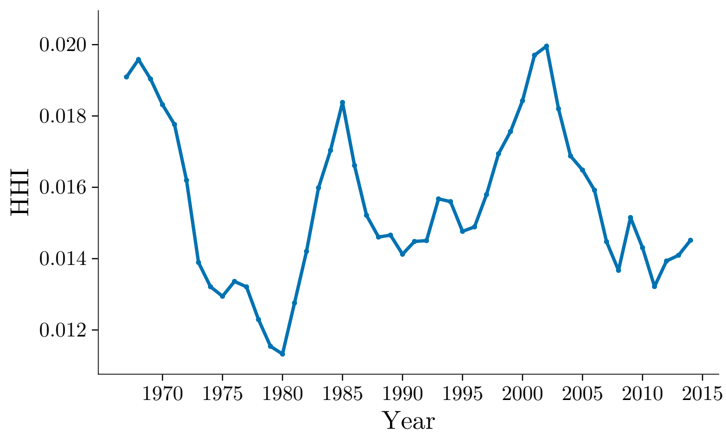

Figure 2: Herfindahl Index of Global Sales Shares, 1967–2014

Notes: The figure plots $H_t = \sum_{c,i} \omega_{ci,t}^2$ where $\omega_{ci,t}$ is the ratio of country-industry (c, i) gross output to world GDP. Dashed vertical lines mark 1982, 1991, and 2008. WIOD data, 24 countries, 23 industries, 1967–2014.

The index oscillates between roughly 0.012 and 0.020 over fifty years, with identifiable reversals at three episodes around 1982, 1991, and 2008. The network is therefore an actively adjusting object that moves over the business cycle.

In addition, the geographic distribution of central nodes shifted over the period. Figure 3 shows the global trade network at two points in time, 1967 and 2014.

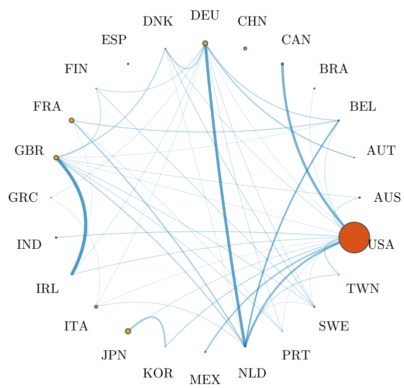

(a) Trade Network, 1967

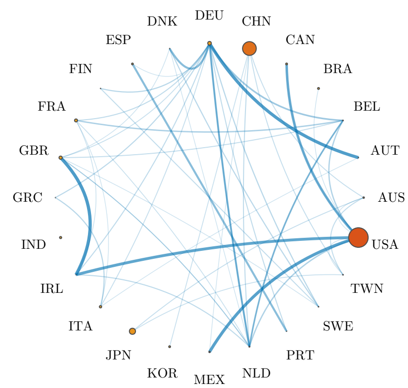

(b) Trade Network, 2014

Figure 3: Global Trade Network at Two Points in Time

Notes: Node size proportional to country Domar weight $\omega_{c,t} = \sum_i \omega_{ci,t}$. Edge thickness proportional to bilateral intermediate trade flows. WIOD data, 24 countries, 1967 and 2014.

The node size is proportional to each country's sale share (aggregate Domar weight) and edge thickness is proportional to bilateral trade flows. The United States is a dominant hub in both years, but the surrounding structure changed: Japan rose as a second major hub through the 1980s before partially retreating, China grew from a peripheral node in 1967 to one of the largest hubs by 2014. Also the overall network became more interconnected as edge became more thickened showing an increase in trade.

2.3 Discussion

The two patterns documented in this section are the joint object this paper seeks to explain. Aggregate risk shapes the network through firms' sourcing decisions. However, the network in turn determines how risk propagates. The model is designed to disentangle this feedback.

In particular, the model decomposes changes in the network into two sources: changes in trade costs and changes in uncertainty about productivity shocks. I then use the model's equilibrium conditions to identify both objects from the data, and to quantify the effects of trade liberalization.

3 Environment

There is a set of C countries and J distinct global industries. Each country has one representative producer in each industry, yielding $N = C \times J$ representative firms that produce differentiated intermediate goods. Each country has a representative firm that combines intermediate goods from all N suppliers into a final good that is sold to a representative household. Finally, each household supplies labor L_c to the J industries within its country c, consumes, and owns the firms.

The timing is as follows. At the beginning of each period t, firms observe the current conditional distribution of productivity shocks, characterized by the mean θ_t and the covariance $\Sigma_{\epsilon,t}$, and choose the production technique to minimize risk-adjusted costs under this distribution. Next, after shocks are realized, markets clear. All firms and households choose the demand and supply of all intermediate and final goods. There are no dynamic choices linking equilibria across periods: each period is a static game in which the network is chosen given current beliefs about shock dispersion.

3.1 Households

The representative household in the country c derives its utility from consumption and leisure. Preferences have the following form:

$$W_c = \sum_t \frac{\left(Z_{c,t} - \sum_i^J \frac{\left(L_{ci,t}\right)^{1+\frac{1}{\psi}}}{1+\frac{1}{\psi}}\right)^{1-\gamma}}{1-\gamma} \tag{1}$$

I follow Greenwood, Hercowitz, and Huffman (1988) in defining quasi-linear preferences, which eliminates any income effect on labor supply. Households choose consumption after uncertainty is realized. Thus, in each period and for each state of the world, the budget constraint is:

$$P_{c,t} Z_{c,t} = \sum_i W_{ci,t} L_{ci,t} + \delta_{c,t} \tag{2}$$

where $P_{c,t}$ is the price index of the final good, $W_{ci,t}$ is the nominal wage in industry (c, i). Also, I allow for trade imbalances by introducing exogenous, time-varying trade deficits $\delta_{c,t}$, following the approach of Dekle, Eaton, and Kortum (2007) and Caliendo and Parro (2015), among others. Trade deficits are predetermined at the time households make their consumption decision and are financed externally, so that markets are incomplete in the sense of Heathcote and Perri (2002): there is no state-contingent asset trade, and households cannot insure against aggregate risk through financial markets.[2] The optimality condition for labor in each industry is given by the following expression:

$$L_{ci,t} = \left(\frac{W_{ci,t}}{P_{c,t}}\right)^{\psi} \tag{3}$$

Equation (3) gives the standard optimal condition that relates labor supply to the real wage. Finally, I define the household's price kernel as follows:

$$\Lambda_{c,t} = \frac{\left(Z_{c,t} - \sum_{i=1}^{J}\left(L_{ci,t}\right)^{1+\frac{1}{\psi}}\right)^{-\gamma}}{P_{c,t}}$$

which measures the marginal utility of consumption for households in each state of the world.

3.2 Intermediate Goods Producers

First Stage Before productivity shocks are realized, the representative firm in country c, industry i chooses a production technique $\eta_{ci,t}$. The firm operates a Cobb-Douglas technology:

$$F_{ci,t}(\eta_{ci,t}) = e^{A_{ci,t}(\eta_{ci,t})+\epsilon_{ci,t}}\; L_{ci,t}^{\mu_{ci,t}}\left(\prod_{\hat{c},k} X_{ci,\hat{c}k,t}^{\eta_{ci,\hat{c}k,t}}\right)^{1-\mu_{ci,t}} \tag{4}$$

where $X_{ci,\hat{c}k,t}$ is the intermediate input from industry k in country $\hat{c}$, $L_{ci,t}$ is labor, and $\epsilon_{ci,t}$ is an exogenous productivity shock. The parameter $\mu_{ci,t} \in (0,1)$ is the Cobb-Douglas labor share, which varies over time. The vector $\eta_{ci,t} = \{\eta_{ci,\hat{c}k,t}\}$ specifies the production technique, constrained to the unit simplex $\sum_{\hat{c},k}\eta_{ci,\hat{c}k,t} = 1$. The firm maximize expected profits discounted by the household pricing kernel $\Lambda_{c,t}$:

$$\eta^{\star}_{ci,t} \in \arg\max_{\eta_{ci,t}\in\mathcal{A}} \mathbf{E}\left[\Lambda_{c,t}\left(P_{ci,t}Y_{ci,t} - W_{ci,t}\,L_{ci,t} - \sum_{\hat{c},k} P_{c,\hat{c}k,t}\,X_{ci,\hat{c}k,t}\right)\right] \tag{5}$$

[2] Heathcote and Perri (2002) show that models with financial autarky perform well in accounting for business-cycle comovement. I allow fixed trade deficits to match the data while maintaining incomplete markets.

Deviating from the ideal technique η^0_{ci} reduces productivity:

$$A_{ci,t}(\eta_{ci,t}) = -\kappa^I_i \sum_{\hat{c},k} \eta_{ci,\hat{c}k,t} \log \frac{\eta_{ci,\hat{c}k,t}}{\eta^0_{ci,\hat{c}k}} \tag{6}$$

The adjustment cost is an industry-specific parameter, and $\kappa^I_i > 0$ governs the cost of deviating from the ideal sourcing pattern. The entropy functional keeps all shares strictly positive for any finite κ, avoiding corner solutions. In Appendix C I show an example show how the optimal network under this cost compare to the quadratic cost in Acemoglu and Azar (2020) and describe how this cost help to make the solution of the competitive equilibrium more tractable.

Intermediate goods are subject to iceberg trade costs. The price of good $(\hat{c}, k)$ paid by buyers in country c is $P_{c,\hat{c}k,t} = P_{\hat{c},\hat{c}k,t}\, \tau^I_{c,\hat{c}k,t}$.

Second Stage In the second stage, given the production function, firms choose their demand for each input after observing prices. The first-order conditions yield the standard unit cost function:

$$K_{ci,t}(\eta_{ci}, P) = e^{-A_{ci,t}(\eta_{ci,t})-\epsilon_{ci,t}}\; W^{\mu_{ci,t}}_{ci,t} \left(\prod_{\hat{c},k} P^{\eta_{ci,\hat{c}k,t}}_{c,\hat{c}k,t} \right)^{1-\mu_{ci,t}} \tag{7}$$

and the first-stage problem becomes:

$$\eta^\star_{ci,t} \in \arg \min_{\eta_{ci,t} \in \mathcal{A}} \mathbf{E}\big[\Lambda_{c,t}\, Y_{ci,t}\, K_{ci,t}(\eta_{ci}, P)\big] \tag{8}$$

3.3 Final Goods Producers

Each country has a representative final goods producer that combines intermediate goods into a consumption bundle. The production function is:

$$F_{c,t}(\alpha_{c,t}) = e^{A_{c,t}(\alpha_{c,t})} \prod_{\hat{c},k} X^{\alpha_{c,\hat{c}k,t}}_{c,\hat{c}k,t} \tag{9}$$

where $\alpha_{c,t}$ lies in the unit simplex and deviating from the ideal technique α^0_c reduces productivity as in equation (6), with adjustment cost κ^F. The technique choice problem is analogous to that of intermediate goods producers: final goods firms choose $\alpha_{c,t}$ before shocks are realized to minimize expected costs discounted by $\Lambda_{c,t}$. Trade costs in final goods are $\tau^F_{c,\hat{c},t}$.

3.4 Stochastic Structure

Uncertainty comes from TFP shocks. The productivity of each industry decomposes into global, country, industry, and granular components:

$$\epsilon_{ci,t} = g_t + \chi_{c,t} + \zeta_{i,t} + u_{ci,t} \tag{10}$$

Each component follows an independent AR(1) process with persistence ρ_k and time-varying innovation variance $\sigma^2_{k,t}$ for $k \in \{g, \chi, \zeta, u\}$.

3.5 Equilibrium

We now define a competitive equilibrium. Since all firms are competitive, equilibrium prices equal costs. The equilibrium of the economy with the endogenous network is defined as follows:

Definition 1. (Competitive Equilibrium) An equilibrium is a sequence of allocations for intermediate firms $\{Y_{ci,t}, X_{ci,\hat{c}k,t}, \eta_{ci,t}\}_{i,c}$; a sequence of allocations for final good producers $\{Y_{c,t}, X_{c,\hat{c}k,t}, \alpha_{c,t}\}$; a sequence of allocations for households $\{Z_{c,t}, L_{ci,t}\}_c$; and a sequence of prices $\{P_{ci,t}, W_{ci,t}, P_{c,t}\}_{ci}$ such that all agents in the economy solve their problem.

3.6 Definitions and Notation

I close the description fo the model by introduce some key objects.

Input-Output and Final Goods Matrices. The IO matrix $\boldsymbol{\eta}_t$ is the $N \times N$ matrix whose $(ci, \hat{c}k)$th element is the Cobb-Douglas coefficient controlling the intensity at which industry ci uses inputs from industry $\hat{c}k$. The Final Goods matrix $\boldsymbol{\alpha}_t$ stores the Cobb-Douglas coefficients in the final goods production function; rows for buyer industries within the same country share the same values, since there is one representative final-good producer per country. Let $\boldsymbol{\mu}_t$ be the diagonal $N \times N$ matrix of labor shares $\mu_{ci,t}$.

Domar Weights. The Domar weight of industry ci is its share of world GDP measured in sales:

$$\omega_{ci,t} = \frac{P_{ci,t} Y_{ci,t}}{GDP} \tag{11}$$

I also define the country Domar weight $\Omega_c = \sum_i \omega_{ci,t}$ and the industry Domar weight $\hat{\Omega}_i = \sum_c \omega_{ci,t}$.

4 Equilibrium Network Formation

I analyze the model backwards. First, I characterize prices given a network. Next, I use these results to characterize the optimal choice of network.

4.1 Equilibrium Given a Network

Given a the aggregate state in the second-stage are $s = \{\alpha^\star, \eta^\star, \Sigma, \theta, \tau\}$ and the stochastic state ϵ. Throughout, I denote $\boldsymbol{x}$ as the vector containing all values of log of $X_{ci,t}$.

Domar Weights. The relative size of each industry in the global economy is summarized by its Domar weight as follows:

Lemma 1. *Define $\boldsymbol{D}_\delta$ the deficit as a share of income. The Domar weight vector $\boldsymbol{\omega}(s)$ is the unit eigenvector with eigenvalue 1 of:*

$$\boldsymbol{\omega}(s) = \left[\boldsymbol{\alpha}^{\star\top}(I + \boldsymbol{D}_\delta)\,\boldsymbol{\mu} + \boldsymbol{\eta}^{\star\top}(I - \boldsymbol{\mu})\right]\boldsymbol{\omega}(s)$$

Proof. See Appendix B.1. □

The term $\boldsymbol{\alpha}^{\star\top}(I + \boldsymbol{D}_\delta)\boldsymbol{\mu}$ accounts for revenue from sales of final goods (including the trade deficit adjustment), while $\boldsymbol{\eta}^{\star\top}(I - \boldsymbol{\mu})$ accounts for revenue from sales of intermediate goods. Market shares increase either because other intermediate-good producers increase their use of the good, or because final-good producers increase their expenditure shares. Given the Cobb-Douglas structure, Domar weights do not change with the stochastic shocks ϵ_t.

The relative size of final versus intermediate goods trade depends on the labor share vector $\boldsymbol{\mu}$. Higher labor shares direct more production income to households, increasing the role of the final goods network $\boldsymbol{\alpha}^\star$ relative to the intermediate input network $\boldsymbol{\eta}^\star$. The deficit matrix $\boldsymbol{D}_\delta$ corrects for trade imbalances in the data, ensuring that countries whose expenditures exceed their income are properly represented in the model's expenditure flows.

Prices. Equilibrium prices are determined by the network and the realization of productivity shocks.

Lemma 2. *Let $\boldsymbol{p}$ be the log price vector and ϵ the vector of TFP shocks. Then:*

$$\boldsymbol{p}(s,\epsilon) = \mathcal{L}(s)\,(B(s) - \epsilon),$$

where the extended Leontief inverse is:

$$\mathcal{L}(s) = \left(I - \boldsymbol{\mu}\,\frac{\psi}{1+\psi}\,\boldsymbol{\alpha}^\star - (1-\boldsymbol{\mu})\,\boldsymbol{\eta}^\star\right)^{-1},$$

and the deterministic price component is:

$$B(s) = -A(\eta^\star_{ci}) - \frac{\psi}{1+\psi}\,\mu\,A(\alpha^\star_c) + \text{diag}\!\left(\tfrac{\psi}{1+\psi}\,\mu\,\alpha^\star + (I-\mu)\,\eta^\star\right)\log\tau + \frac{1}{1+\psi}\,(\log\mu + \log\omega(s)).$$

Proof. See Appendix B.2. □

Whereas the matrices $\boldsymbol{\alpha}^\star$ and $\boldsymbol{\eta}^\star$ record *direct* links between producers, the Leontief inverse records *direct and indirect* exposures through the entire network.This Leontief inverse generalizes the classical one by adding the term $\boldsymbol{\mu}\,\frac{\psi}{1+\psi}\,\boldsymbol{\alpha}^\star$. This additional term comes from the household labor supply condition and captures how a TFP shock in any industry affects wages and, through wages, production costs in all other industries.[3] The elasticity parameter ψ determines the relative strength of the two transmission channels. When labor supply is inelastic, wages respond only weakly, so changes in intermediate input prices become the main transmission channel. By contrast, when labor supply is elastic, final-good prices have a larger impact on wages, making final-good trade the key

[3] Huo et al. (2024) derive this extended Leontief inverse in a model with CES production.

transmission channel.

The vector $B(s)$ summarizes how the network tilts prices in the absence of shocks. The terms $A(\alpha_c^\star)$ and $A(\eta_{ci}^\star)$ capture how costly supplier reallocation lowers effective productivity and raises prices in industries that deviate from their ideal technology. The trade cost component makes prices higher in industries that source intensively from foreign inputs. The terms $\log \boldsymbol{\mu}$ and $\log \boldsymbol{\omega}(s)$ capture the fact that industries that use labor more intensively, and industries with larger market shares, must pay higher wages to compensate households, who supply labor with and elasticity ψ.

The final-good price index aggregates supplier prices using the optimal network weights $\alpha_c^\star$, adjusted for trade costs in final goods.

Corollary 1 (Final-good price index)**.** *The log price index of the final good in country c is:*

$$p_c(s,\epsilon) = -A(\alpha_c^\star) + \boldsymbol{\alpha}_c^\star \left(\boldsymbol{p}(s,\epsilon) + \log \tau^F\right).$$

Pricing Kernel. The pricing kernel connects price realizations to the household's marginal utility of consumption and is the key object for evaluating the insurance value of each supplier.

Lemma 3. *The log pricing kernel of the representative household in country c is:*

$$\log \Lambda_c(s,\epsilon) = (\gamma - 1)\, p_c(s,\epsilon) + C_c(s),$$

where $C_c(s) = -\gamma\left(1 + \frac{1}{\psi}\right)\log\left(\sum_i (\mu_{ci}\,\omega_{ci}(s))^{\frac{\psi}{1+\psi}}\right) + \gamma \log\left(1 + \frac{1}{\psi}\right) + \gamma \log \psi$ *depends on the network but is independent of shocks.*

Proof. See Appendix B.6. □

The separation in Lemma 3 is what keeps the first-stage problem tractable. Both $C_c(s)$ and $\mathrm{Cov}(p_c, p_{\hat{c}k})$ are determined by the network, but the insurance value of each supplier depends only on the covariance term, not on $C_c(s)$. Endogenous labor supply therefore does not add a new channel to the first-stage problem compared to the baseline model Kopytov et al. (2024).

Expectations. The final ingredient for the optimal network problem is the distribution of prices,

Lemma 4. *Let* $\epsilon_t \sim \mathcal{N}(\theta_t, \Sigma_{\epsilon,t})$. *Then* $\boldsymbol{p}(s,\epsilon)$ *is normally distributed with:*

$$\mathbf{E}\big[\boldsymbol{p}(s,\epsilon)\big] = \mathcal{L}(s)\,(B(s)-\theta),$$
$$\mathbf{V}\big[\boldsymbol{p}(s,\epsilon)\big] = \mathcal{L}(s)\,\Sigma_\epsilon\,\mathcal{L}(s)^\top.$$

Proof. See Appendix B.4. □

The Lemma 4 determines both expected costs and the covariance term entering the risk-adjusted price. Log-normality follows from the Cobb-Douglas production function combined with normally distributed shocks, which together ensure that prices are linear in ϵ through the Leontief inverse. The variance formula $\mathcal{L}(s)\Sigma_\epsilon\mathcal{L}(s)^\top$ shows how the network amplifies and transmits the primitive shock variances.

4.2 Optimal Production Networks

Firms in both final and intermediate industries choose a supplier vector to minimize expected discounted production costs, using the domestic pricing kernel derived in Lemma 3 the problem of firms in the first stage can be characterized as follows.

Proposition 1 (Risk-Adjusted Prices)**.** *Let* $m = \alpha_c$ *for final-good producers and* $m = \eta_{ci}$ *for intermediate-good producers. In equilibrium, their choice satisfies:*

$$m^\star(s) \in \arg\max_{m\in\Delta}\; -\kappa_b^m \sum_{\hat{c},k} m_{\hat{c}k}(\log m_{\hat{c}k} - \log m^0_{\hat{c}k}) \;-\; \sum_{\hat{c},k} m_{\hat{c}k}\, R_{\hat{c}k}(s),$$

where the risk-adjusted price of supplier $(\hat{c},k)$ *for buyer* b *is:*

$$R_{b,\hat{c}k}(s) = E[p_{\hat{c}k}(s,\epsilon)] + \log \tau^m_{b,\hat{c}k} + (\gamma-1)\,\mathrm{Cov}(p_c(s,\epsilon),\, p_{\hat{c}k}(s,\epsilon))$$

and p_c *is the domestic final-good price index of the buyer's country* c.

Proof. See Appendix B.6. □

The risk-adjusted price $R_{\hat{c}k}(s)$ combines three components. The first, $E[p_{\hat{c}k}]$, captures the expected productivity of the supplier: firms with lower expected prices reflect either higher own productivity or access to productive upstream inputs. The second, $\log \tau_{\hat{c}k}$, introduces a standard gravity channel: holding productivity fixed, firms prefer geographically close or domestically sourced suppliers. The third, $(\gamma - 1)\mathrm{Cov}(p_c, p_{\hat{c}k})$, is the insurance value of the supplier: as shown in the proof, household consumption in country c is a decreasing function of p_c, so firms prefer suppliers whose prices are low precisely when the domestic price index is high—suppliers that are cheap in bad times provide valuable insurance and have a lower risk-adjusted price.

Risk-adjusted prices in Proposition 1 nests those in Kopytov et al. (2024) as a special case: with a single country and no trade costs ($\tau_{\hat{c}k} = 1$).[4] The open-economy extension adds two forces. First, trade costs introduce heterogeneous barriers that break the symmetry across potential suppliers, generating the gravity structure needed to match observed bilateral trade patterns. Second, the covariance term $\mathrm{Cov}(p_c, p_{\hat{c}k})$ now reflects cross-country heterogeneity: in a multi-country network with incomplete markets, it depends on each country's equilibrium final-good network $\alpha_c^\star$, so the insurance value of the same supplier differs across buyers . The solution of the optimal network is:

Lemma 5 (Optimal Network)**.** *Let $m \in \{I, F\}$. The optimal network satisfies the softmax:*

$$m_{\hat{c}k}^\star(s) = \frac{m_{\hat{c}k}^0 \, \exp\!\left(-R_{\hat{c}k}(s)/\kappa_b^m\right)}{\sum_{(\hat{c}',k')} m_{\hat{c}'k'}^0 \, \exp\!\left(-R_{\hat{c}'k'}(s)/\kappa_b^m\right)}, \tag{12}$$

where κ_b^m is the adjustment cost (κ_i^I for intermediates, κ^F for final goods).

Proof. See Appendix B.7. □

The equilibrium network is a softmax function of risk-adjusted prices, with κ_b^m acting as the temperature: large κ_b^m keeps shares close to the ideal m^0, while small κ_b^m concentrates sourcing on the cheapest supplier. The equilibrium solves a fixed point: risk-adjusted prices $R_{b,\hat{c}k}$ depend on the optimal networks, and optimal networks depend on risk-adjusted prices. Under the entropy cost, expenditure shares are positive as long as $m^0 > 0$. Appendix C compares the softmax to CES and quadratic adjustment costs and describe how it helps in the numerical solution of the model.

[4]Under $\tau = 1$, Lemma 4 and the single-country restriction imply $R_{\hat{c}k}(s) = E[p_{\hat{c}k}(s, \epsilon)] + (\gamma - 1)\, e_{\hat{c}k}^\top \mathcal{L}(s)\Sigma_\epsilon \mathcal{L}(s)^\top m^\star$, where $e_{\hat{c}k}$ is the standard basis vector.

4.3 Existence and Uniqueness

The competitive equilibrium is a fixed point that links the aggregate network to firms' best responses: each firm's sourcing decision alters the covariance structure of prices, which then feeds back into the optimal choices of all other firms. This general-equilibrium feedback can, in principle, generate multiple equilibria. Because the problem does not admit a planner's representation under incomplete markets. Consequently, the approach of Kopytov et al. (2024), who establish uniqueness via a planning problem, is not available here. Instead, I write the competitive equilibrium as a static game to proof existence and derive sufficient conditions for uniqueness.

Proposition 2 (Existence and Uniqueness)**.** *Let $\kappa^m > 0$ for each margin $m \in \{\alpha, \eta\}$. Then the equilibrium $s^*, p(s^*, \epsilon)$ exists. Moreover, s^* is unique if:*

$$\bar{\kappa} \;>\; \frac{1}{2}\,\frac{(1+\psi)^2}{\underline{\mu}^3}\left[(1-\underline{\mu}) + \frac{\psi\,\bar{\mu}}{1+\psi}\right]\left[2(\gamma-1)(1+\psi)\,\bar{\sigma}^2 \;+\; (\bar{\theta}+\bar{B})\,\underline{\mu}\right]$$

where $\bar{A} = \max_{s\in\mathcal{S}} \|A(m^\star)\|_\infty$ bounds the entropy adjustment cost, $\bar{B} := \bar{A} + \frac{\psi}{1+\psi}\bar{\mu}\,\bar{A} + \overline{\log\tau} + \frac{1}{1+\psi}(\overline{\log\mu} + \overline{\log\omega})$, $\bar{\kappa} = \min_m\{\kappa^m\}$, $\bar{\sigma}^2$ is the maximum eigenvalue of Σ_ϵ, and $\underline{\mu}, \bar{\mu}$ are the minimum and maximum labor shares.

Proof. See Appendix D. □

Existence follows from Brouwer's theorem: the strategy space is a product of simplices, prices are uniquely pinned down by the Leontief inverse for any given network, and the entropy adjustment cost makes each firm's best response a smooth function of risk-adjusted prices. So the fixed-point mapping from the aggregate network to best responses is continuous on a compact convex set and must have a fixed point.

Uniqueness requires controlling the general equilibrium feedbacks. The right-hand side of the condition measures the maximum sensitivity of risk-adjusted prices to changes in the aggregate network: it is increasing in risk aversion γ and in shock's variance $\bar{\sigma}^2$. The condition says that if adjustment costs $\bar{\kappa}$ exceed this bound, firms respond weakly enough to changes in the covariance structure that the best-response map is a contraction and the equilibrium is unique. Intuitively, large adjustment costs anchor firms near their ideal technology m^0, preventing multiple equilibria.

4.4 Welfare and Production

Next, I describe the effect of changes in the network $\{\alpha^\star, \eta^\star\}$ in macroeconomic outcomes.

Real Value Added. I characterize the log of real value added in each country-industry pair as a function of the network and the realization of productivity shocks.

Lemma 6. *The log of real value added a_{ci} is:*

$$a_{ci}(s, \epsilon) = \epsilon_{ci} + A(\eta^\star_{ci}) + \mu_{ci}\, l_{ci}(s, \epsilon)$$

where log labor supply $l_{ci}(s, \epsilon)$ is:

$$l_{ci}(s, \epsilon) = \frac{\psi}{1+\psi}\Big[-p_c(s, \epsilon) + \log \mu_{ci} + \log \omega_{ci}(s)\Big]$$

Proof. See Appendix B.3. □

Lemma 6 decomposes value added into three components. The first, ϵ_{ci}, is the TFP shock. The second, $A(\eta^\star_{ci})$, is the endogenous productivity determined by the choice of network. The third is the endogenous labor supply response: a TFP shock anywhere in the network propagates into the domestic price index p_c, changing the real wage and shifting labor supply and value added across all country-industry pairs. This channel was first emphasized by Backus, Kehoe, and Kydland (1992) as the source of international comovement. Lemma 6 is closely related to Huo et al. (2024), but departs from their result because they characterize deviations from the steady state in a CES model.

Welfare. I express welfare for the representative household in country c in terms of the equilibrium network, the expected price index, and the variance of prices.

Lemma 7. *At each period t, welfare in country c is:*

$$W_c(s) = b + \left(1 + \tfrac{1}{\psi}\right) \log\left(\sum_i \left(\omega_{ci}(s)\, \mu_{ci}\right)^{\frac{\psi}{1+\psi}}\right) - \mathbb{E}\big[p_c(s, \epsilon)\big] + \tfrac{1}{2}(1-\gamma)\, \mathrm{Var}\big(p_c(s, \epsilon)\big)$$

where b is a constant.

Proof. See Appendix B.8. □

Welfare aggregates three forces that operate through the network. The first term captures the relative size of the domestic economy: households in larger economies face higher labor demand, work more, and consume more. The second term, $\mathbb{E}[p_c]$, measures the cost of the consumption basket: lower expected prices and raises purchasing power and increases consumption. The third term captures country-specific risk: under incomplete markets, $\text{Var}(p_c)$ measures the volatility of consumption and reduces welfare when $\gamma > 1$.

Sourcing decisions by any producer in the economy affect welfare by reshuffling production across countries, changing the efficiency in the production of the consumption basket through $\mathbb{E}[p_c]$, and changing country-specific risk through $\text{Var}(p_c)$.

5 Network and Aggregate Risk

In this section, I characterize how the network adjust in response to changes in risk and trade cost.

5.1 Centrality

I begin by providing the definition of centrality in this economy which will be useful in the subsequent analysis.

Definition 2 (Centrality)**.** The centrality of country-industry pair (c, i) is the row sum of the Leontief inverse:

$$v_{ci} = \sum_{\hat{c}k} \mathcal{L}_{ci,\hat{c}k}$$

The exposure of country c to central industries is:

$$v_c = \sum_{\hat{c}k} \alpha^{\star}_{c,\hat{c}k}\, v_{\hat{c}k}$$

The row sum v_{ci} measures how sensitive industry (c, i)'s price a productivity shock. So prices of central sector comove more strongly with productivity shocks and propagate them more strongly.

In contrast to Acemoglu et al. (2012), with endogenous labor supply, centrality is not characterized by the Domar weight. Centrality in this model differs from the Domar weight because the Leontief inverse weights final and intermediate sales asymmetrically by ψ, while Domar weights treat them symmetrically.[5]

5.2 Centrality-Risk Reallocation

Now, I show how firms reallocate sourcing away from central suppliers when uncertainty in productivity shocks rises. The focus is on the response to an increase in σ_g^2.

Proposition 3. *The first-order effect of a change in aggregate risk σ_g^2 on expenditure shares is:*

$$\frac{\partial m^\star_{ci,\hat{c}k}}{\partial \sigma_g^2} = -\frac{(\gamma - 1)}{\kappa^m}\, m^\star_{ci,\hat{c}k} \left(\nu_{\hat{c}k} - \bar{\nu}_{ci}\right) \nu_c \tag{13}$$

where $\bar{\nu}_{ci} = \sum_{\hat{c}k' \in N^m} m^\star_{ci,\hat{c}k'}\, \nu_{\hat{c}k'}$ is the average centrality weighed by the share among suppliers.

Proof. See Appendix E. □

The sign of $\partial m^\star / \partial \sigma_g^2$ depends on whether the supplier's centrality $\nu_{\hat{c}k}$ exceeds the buyer's average $\bar{\nu}_{ci}$: suppliers with above-average centrality lose market share when global risk rises, and suppliers with below-average centrality gain. The adjustment is amplified by ν_c, the country's aggregate exposure to central industries through its final-goods consumption basket. Because all buyers in the same country share the same domestic pricing kernel, ν_c is common to all intermediate producers in the country c.

This result extends to Kopytov et al. (2024), who show that industries with higher granular variance lose market share. Here, the channel is different: an industry can have low own variability $\sigma^2_{u_{ci}}$ and still be a risky input if it is highly central, because its price loads on the global factor through the Leontief inverse. The risk of each supplier is therefore driven by the position of the network. Another implication of this result is that if central suppliers lose market shares when the variance of the global shock increases the concentration of the economy should decrease.

In this section, I focus on the effect of uncertainty in global shocks. However, the result applies to increases in uncertainty of country-specific shocks and industry-specific

[5] The result is related to Baqaee and Farhi (2024) who show that revenue-based and cost-based transmission measures also differ in economies with wedges.

shocks. When the variance of a country-specified shock increases, central sectors within the country become more vulnerable.

5.3 Trade and Aggregate Risk

To build intuition, I consider a two-country economy — Home (H) and Foreign (F) — each with a single intermediate good producer. Labor shares differ: $\mu_H > \mu_F$, so the foreign use intermediates more intensively. Foreign is the star: $\theta_F > \theta_H$. To isolate the intermediate-trade channel, I shut down the final-goods trade: each country consumes only its own final good ($\alpha_{H,H} = \alpha_{F,F} = 1$). The only adjustment margin is intermediate supply $\eta_c = \{\eta_{c,H}, \eta_{c,F}\}$, subject to the iceberg trade cost τ^I. Let $\epsilon_c = g + u_c$.

Since the Home final good uses only the Home intermediate, $p_H^F = c + p_H$. The variance is decomposed as follows.

$$\text{Var}(p_H) \;=\; \underbrace{\sigma_g^2\, \nu_H^2}_{\text{Global}} \;+\; \underbrace{\sigma_u^2 \left(\mathcal{L}_{HH}^2 + \mathcal{L}_{HF}^2\right)}_{\text{Granular}} \tag{14}$$

where $\nu_H = \mathcal{L}_{HH} + \mathcal{L}_{HF}$ is the centrality of the Home sector (Definition 2). The global shock enters through ν_H because it hits both sectors uniformly and the Leontief loadings add up. Granular shocks enter through individual elements $\mathcal{L}_{HH}$ and $\mathcal{L}_{HF}$ because each is sector-specific.

Under autarky ($\tau^I > \tau^\star$), the Leontief inverse is diagonal: $\mathcal{L}_{HH} = (1+\psi)/\mu_H$, $\mathcal{L}_{HF} = 0$. All risk comes from the Home sector's own shocks. As intermediate trade opens ($\tau^I < \tau^\star$), $\eta_{H,F}$ rises and $\mathcal{L}_{HF} > 0$: Home's price now loads on Foreign shocks through the input-output network.

This has two opposing effects on variance. The sum of squares $\mathcal{L}_{HH}^2 + \mathcal{L}_{HF}^2$ decreases as loadings spread across two suppliers, reducing the exposure to granular shocks. This is the Caselli et al. (2020) channel. Also, as the foreign sector buys more from the more productive domestic goods, the domestic producer becomes central and ν_H increases, increasing the volatility generated by the aggregate shock.

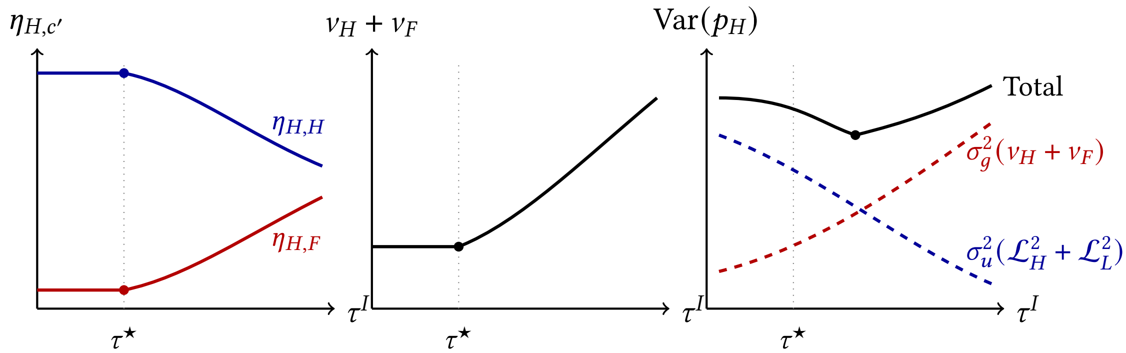

(a) Intermediate sourcing (b) Centrality $\nu_H + \nu_F$ (c) Variance decomposition

Figure 4: Intermediate Trade Liberalization: Sourcing, Centrality, and Volatility

Notes: Two-country economy with intermediate trade only ($\alpha_{H,H} = \alpha_{F,F} = 1$). Foreign is the star: $\theta_F > \theta_H$, $\mu_H > \mu_F$. $\sigma_{u_c}^2 = \sigma_u^2$ for both countries. $\tau^\star$ is the threshold below which Home sources intermediates from Foreign.

Figure 4 illustrates. Panel (a) shows intermediate sourcing shares as τ^I falls: Home shifts toward the more productive Foreign supplier. Panel (b) shows that Home's centrality ν_H rises as trade opens, because sourcing from the more intermediate-intensive Foreign sector pulls Home into a more interconnected network. Panel (c) shows the resulting U-shape in $\text{Var}(p_H)$: close to autarky, the granular component $\sigma_u^2(\mathcal{L}_{HH}^2 + \mathcal{L}_{HF}^2)$ falls faster than the global component $\sigma_g^2 \nu_H^2$ rises, so diversification dominates. But as $\eta_{H,F}$ grows further, ν_H continues to rise and the global component eventually dominates, reversing the initial variance reduction.

The example takes σ_g^2 as fixed. But in equilibrium, when aggregate risk rises, firms respond by shifting sourcing away from central suppliers as shown in the previous subsection. This endogenous response mitigates the effect of centrality in risk in periods of high uncertainty.

6 Quantitative Analysis

I calibrate the model using WIOD data for 24 countries and 23 industries from 1967 to 2014; labor shares $\mu_{ci,t}$ vary over time and are taken directly from the data period by period and the remaining parameters are estimated as described below.

6.1 Trade Elasticities and Trade Costs

I identify the ideal technologies η^0, the time-varying bilateral trade costs $\tau_{b,k,t}$, and the adjustment cost parameters κ_b^m by combining the model's optimality conditions with observed input-output and value-added data.

Ideal technologies. Following Kopytov et al. (2024), I identify η^0 from the time-series average of the shares as follows:

$$m^0_{\hat{c}k,j,t} = \frac{1}{T}\sum m^{\text{data}}_{\hat{c}k,(c',j),t}$$

Trade Costs. Trade costs are identified from observed expenditure shares of producers in the same industry but in different locations. I normalize $\log \tau^m_{c,\hat{c}k,t} = 0$ for $c = \hat{c}$, so domestic purchases are frictionless.

Proposition 4. *Normalize* $\log \tau^m_{c,\hat{c}k,t} = 0$ *for* $c = \hat{c}$. *Let* $m \in \{\eta, \alpha\}$. *The trade costs solve*

$$\Big(I - \mathbf{D}\,\mathcal{L}\left[(I-\mu)\eta^\top \quad \tfrac{\psi}{1+\psi}\mu\alpha^\top\right]\Big)\boldsymbol{\tau} = \left[\mathbf{d} + \mathbf{D}\,\mathcal{L}\left(-A(\eta^\star) - \tfrac{\psi}{1+\psi}\mu A(\alpha^\star) + \tfrac{1}{1+\psi}(\log\mu + \log\omega) - \hat{\theta}\right)\right]$$

where $\mathbf{d}$ *is the vector of data-only constants*

$$d^m_{(\hat{c},c,k),t} = \kappa^m \log\frac{m^{\text{data}}_{ck,ck,t}}{m^{\text{data}}_{\hat{c}k,ck,t}} - \kappa^m \log\frac{m^0_{ck,ck,t}}{m^0_{\hat{c}k,ck,t}} + (\gamma-1)\big(\text{Cov}(p_{c,t}, p_{ck,t}) - \text{Cov}(p_{c,t}, p_{\hat{c}k,t})\big).$$

Also, $\mathbf{D}$ *be the difference operator* $[\mathbf{Dx}]_{(\hat{c},c,k)} = x_{ck} - x_{\hat{c}k}$.

Proof. See Appendix H. □

d^m can be estimated from the data and has three terms: the cross-buyer difference of observed shares, the cross-buyer difference of expected productivity, and a risk correction. The risk correction reflects that, under incomplete markets, the same supplier carries a different insurance value for buyers in different countries because their consumption baskets differ.

This contrasts with the gravity inversion of Bonadio et al. (2025) by adding a risk correction.[6] At $\gamma = 1$ (risk neutrality), the covariance term vanishes and the formula reduces to a gravity equation in log-share and log-price differences. With $\gamma > 1$, some cross-country variation in sourcing reflects insurance motives rather than trade frictions, and the risk term filters this out. The correction matters quantitatively in crisis periods, when the covariance differences spike and firms reallocate for insurance reasons unrelated to any change in bilateral frictions.

Identification of η^0, α^0 and τ. As in standard CES type models (see, for example, Bonadio et al. (2025)), the estimation of the level of τ cannot be separated from the technological parameters (η^0 and α^0). However, since technology is assumed to be constant over time, changes in τ that we will use in the counterfactual exercises do not depend on the technological parameters.

6.2 Stochastic Structure

Productivity Shocks I start by inverting the lemma 6, which characterizes VA in terms of productivity shocks given the empirical network, to estimate productivity shocks in every period.

Common Factors The next step is to perform the decomposition of the shocks. In this step I follow Caselli et al., 2020. I estimate the following factor structure.

$$\epsilon^D_{ci,t} = g^D_t + \xi^D_{c,t} + \omega^D_{i,t} + \iota^D_{ci} \tag{15}$$

Estimating time-varying uncertainty. I measure changes in aggregate uncertainty by first estimating a time-varying covariance matrix $\hat{\Sigma}^a_t$ for each via panel GARCH(1,1). Then we recover the covariance of primitive TFP shocks by inverting the structural relationship from Lemma 6:

$$\hat{\Sigma}^\epsilon_t = \left(I + \frac{\psi}{1+\psi}\,\boldsymbol{\mu}_t\,\alpha^\star_t\,\mathcal{L}_t\right)^{-1}\,\hat{\Sigma}^a_t\,\left(\left(I + \frac{\psi}{1+\psi}\,\boldsymbol{\mu}_t\,\alpha^\star_t\,\mathcal{L}_t\right)^{-1}\right)^\top, \tag{16}$$

[6] See also Eaton and Kortum (2002) and Arkolakis, Costinot, and Rodríguez-Clare (2012) for classical references on estimation of trade cost using gravity equations.

Details are provided in Appendix H.

7 Quantitative Results

This section evaluates the estimated model along three dimensions. First, I report the calibrated parameters and estimated time-varying uncertainty. Second, we provide some statistics on the model fit. Third, I test the centrality-risk mechanism of Proposition 3 in the data and the model: when global risk increases, buyers shift away from high-centrality suppliers. The model generates the same pattern with a smaller magnitude.

7.1 Parameters.

The model requires six parameters (Table 1). Labor shares μ_j range from 0.20 to 0.68 between industries and the Frisch elasticity $\psi = 0.72$ is taken directly from the estimation Huo et al. (2024).

Table 1: Parameters

External		
Parameter	Value	Related to
μ_j	$[0.2, 0.68]$	Labor shares from WIOD
ψ	0.72	Frisch elasticity (Huo et al. (2024))
Calibrated		
γ	2.0	Risk aversion
κ	5.0	Adjustment cost (uniform)
Estimated		
ρ	0.888	Avg. AR(1) persistence
ϱ	0.13	GARCH – innovation weight
P	0.49	GARCH – persistence

Notes: The range for μ_j reports the 10th and 90th percentiles across industries. κ is calibrated by matching model CE shares to data. GARCH parameters are estimated from VA-based factor innovations using a panel specification.

The risk aversion is $\gamma = 2.0$ and the average adjustment cost $\kappa = 5.0$, calibrated by the matching of the implied network model with the data. The GARCH parameters governing time-varying volatility are estimated separately for each factor from value-added innovations: the reported values $\varrho = 0.13$ and $P = 0.49$ are averages across the four factors, implying moderate volatility clustering with a half-life of roughly 1.5 years.

7.2 Estimated Uncertainty

The estimation recovers four components of the variance that vary over time (Figure 5). The global variance $\hat{\sigma}^2_{g,t}$ fluctuates between 0.0052 and 0.0067 (Panel a). The series troughs in the early 1970s, rises through the 1980s, and peaks around 1992 and 1995. Because the GARCH conditional variance is backward-looking, the 2008 crisis appears as a spike in the late 2000s, with the series reaching its sample maximum by 2014.

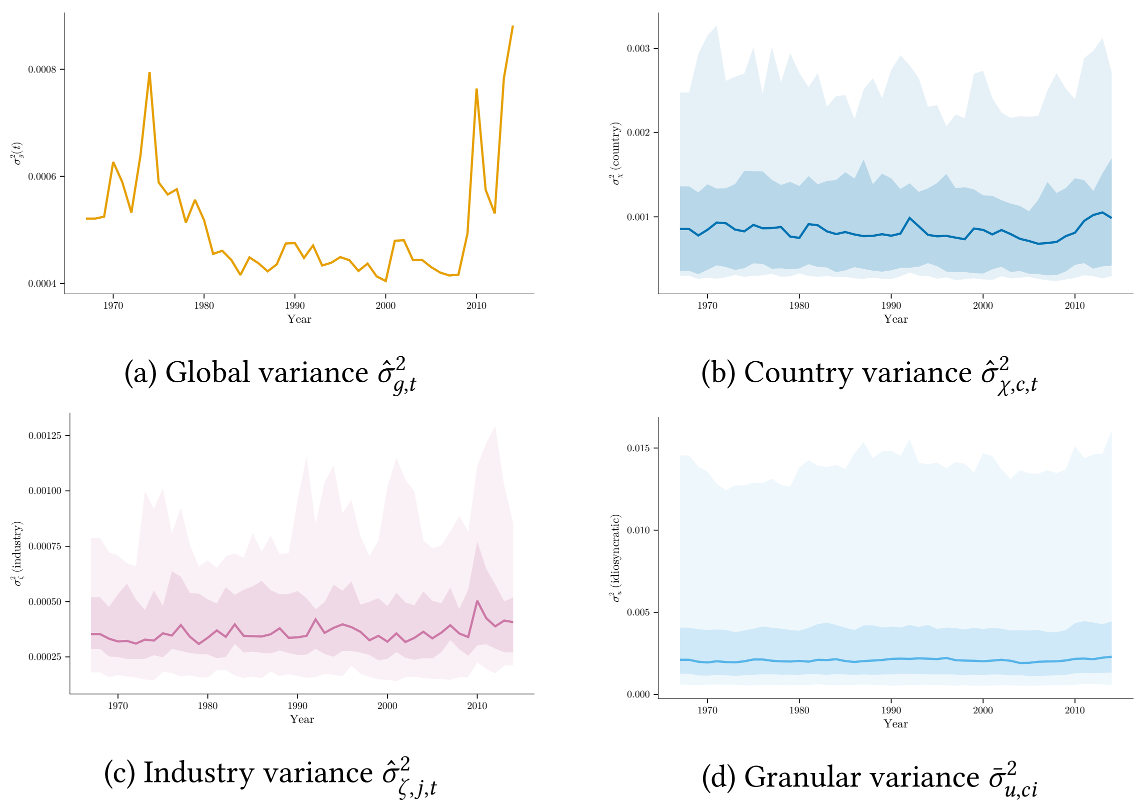

(a) Global variance $\hat{\sigma}^2_{g,t}$

(b) Country variance $\hat{\sigma}^2_{\chi,c,t}$

(c) Industry variance $\hat{\sigma}^2_{\zeta,j,t}$

(d) Granular variance $\bar{\sigma}^2_{u,ci}$

Figure 5: Changes in Uncertainty

Country-specific variances $\hat{\sigma}^2_{\chi,c,t}$ show a wide cross-sectional distribution (Panel b).

The median is around 0.001–0.002, while the upper tail reaches 0.008 around 1990. The distribution peaks in the early 1990s and collapses after 2010. Industry variances are smaller than country variances on average (Panel c). The bulk of the distribution lies below 0.002, with the upper tail reaching 0.003 around 1980. Granular variances $\hat{\sigma}^2_{u,ci}$ are similar in magnitude to the industry component, with the upper tail around 0.003 – 0.004 and a modest time variation spiking near 2010 (Panel d). The global component is the most persistent and displays the clearest trend.

7.3 Model Fit

Table 2 evaluates the fit of the model in two dimensions. The Sector-level Domar weights in the competitive equilibrium correlate 0.74 with their data counterparts. The model reproduces the cross-section and time variation of the GDP variance with a correlation of 0.78 across all country-year pairs. Scatter plots for both metrics appear in the Appendix H.6.

Table 2: Model Fit

Metric	Correlation
Domar weights ($N = 552$)	0.74
$\text{Var}(\text{GDP}_c)$ (all c, t pairs)	0.78

Figure 6 plots the time series of network concentration in data and model.

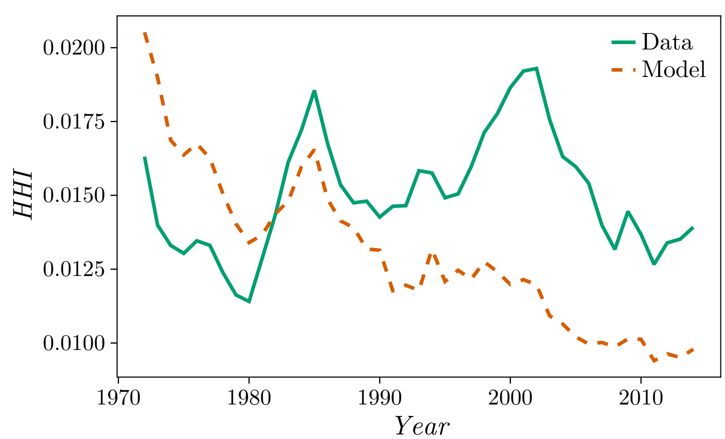

Figure 6: Network Concentration: Data vs. Model

Notes: HHI of Domar weights over time. Solid: data. Dashed: model competitive equilibrium. Note the different y-axis scales.

Both series decline through the 1970s, reach a peak in the mid-1980s, and then trend downward. The model predicts somewhat lower concentration than the data after 1985, with the gap most visible in the late 1990s when data concentration rebounds while the model continues to drift down.

7.4 The Centrality-Risk Mechanism in the Data

The Proposition 3 predicts that when the aggregate risk increases, each buyer substitutes away from suppliers whose centrality exceeds the buyer's average. Because this reallocation is common among buyers, it has two testable aggregate implications. First, network concentration should fall when global risk rises: as all buyers simultaneously reduce their exposure to central suppliers, the Herfindahl index of Domar weights declines. Second, at the micro level, high-centrality country-industry pairs should lose market share when global risk is high. I test both predictions using the estimated global variance $\hat{\sigma}^2_{g,t}$:

$$H_t \;=\; \beta_0 + \beta_1\, \hat{\sigma}^2_{g,t} + \varepsilon_t, \qquad\qquad \Delta\omega_{ci,t} \;=\; \beta\, v_{ci,t-1} \times \hat{\sigma}^2_{g,t-1} \;+\; \delta_{ci} \;+\; \varepsilon_{ci,t} \tag{17}$$

where H_t is the Herfindahl index of the Domar weights, $v_{ci,t-1}$ is the lagged centrality (Definition 2), and δ_{ci} are fixed effects of the country-industry.

	Data		Model	
	H_t	$\Delta\omega_{ci,t}$	H_t	$\Delta\omega_{ci,t}$
$\sigma^2_{g,t}$	-7.6289^{***}		-5.8948^{*}	
	(2.9048)		(3.0330)	
$v_{ci,t-1} \times \sigma^2_{g,t-1}$		-2.8869^{***}		-0.3568
		(0.3297)		(0.5487)
Pair FE	No	Yes	No	Yes
R^2	0.1608	0.0037	0.0950	0.0002
N	38	20,424	38	20,424

Notes: Columns 1 and 3: dependent variable is $H_t = \sum_{c,i} \omega^2_{ci,t}$; time-series regression on estimated global variance. Columns 2 and 4: dependent variable is $\Delta\omega_{ci,t}$; $v_{ci,t-1}$ is lagged centrality; country-industry fixed effects included. Standard errors in parentheses. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

Table 3: The Centrality-Risk Mechanism

Table 3 reports the results. In the aggregate regression (columns 1 and 3), the concentration of the network decreases when the global variance increases, both in the data (-7.63) and in the model (-5.89). In the micro regression (columns 2 and 4), above-average-centrality industries lose market share when global variance is high: the data coefficient is -2.89 and the model generates -0.35, same direction but smaller in magnitude.

As shown by Acemoglu et al. (2012) the concentration of the economy is an important statistic for aggregate risk, as it shapes how strongly productivity shocks propagate through the network and generate fluctuations in GDP. This result illustrates how in the data the network response mitigates the increased uncertainty by reducing the concentration in central suppliers, and the results also show how the mechanism of the model can explain part of this pattern.

8 Counterfactual Analysis

In this section I study the relation of trade cost with aggregate risk measured by the variance of GDP. The analysis proceeds in two steps. First, we trace the full relationship between openness and aggregate outcomes by progressively removing trade costs from their estimated 2014 levels toward free trade. This exercise isolates the non-monotonic interaction between diversification and concentration. Second, we evaluate the actual trade liberalization that occurred between 1967 and 2014, asking whether the observed reduction in trade costs raised or lowered aggregate volatility.

8.1 Trade Liberalization: From Estimated Costs to Free Trade

We reduce all bilateral trade costs proportionally from their estimated 2014 values to free trade, parameterized by $\varphi \in [0, 1]$, where $\varphi = 0$ is the 2014 baseline and $\varphi = 1$ is free trade. At each φ we solve for the competitive equilibrium network and compute the concentration and GDP variance. All other primitives (Σ, α^0, η^0) are fixed to their 2014 values.

Figure 7 reports the results. The network concentration increases monotonically in φ (Panel a). HHI increases from 0.0094 at the 2014 baseline to roughly 0.0205 at free trade, more than doubling. As barriers fall, the most productive hubs become accessible to all buyers, and the network refocuses around them, so the concentration force dominates throughout the liberalization path.

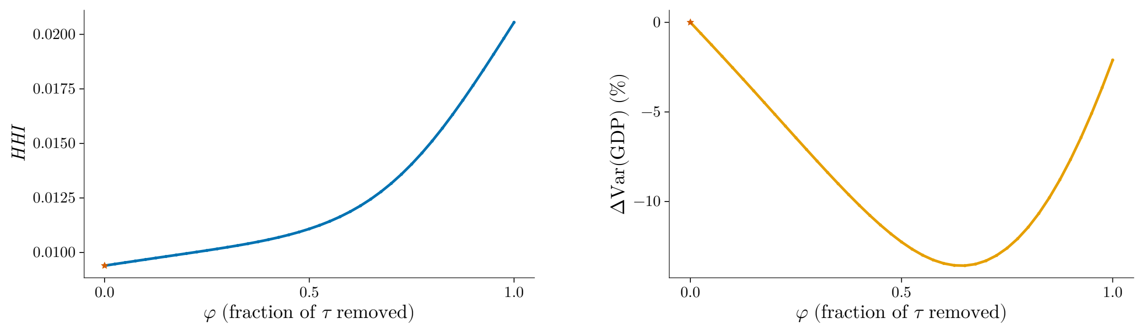

(a) Network concentration

(b) GDP variance

Figure 7: Trade Liberalization from 2014 Baseline to Free Trade

Notes: φ is the fraction of estimated 2014 bilateral trade costs removed. $\varphi = 0$: 2014 baseline. $\varphi = 1$: free trade. All other primitives held fixed at 2014 values.

The GDP variance is U-shaped in φ (Panel b). The Variance decreases by up to 13% around $\varphi = 0.65$, then partially recovers, ending at about -2 per cent in free trade. The initial decline reflects diversification: as buyers spread sourcing across more countries, granular and country shocks are diversified. The partial reversal reflects concentration: when trade is nearly free, the network collapses onto a small set of central suppliers whose shocks propagate globally and erode most of the diversification gains.

Two results stand out. First, the relationship between openness and aggregate risk is non-monotone: moderate liberalization reduces volatility, but deep liberalization reverses most of those gains. Second, this non-monotonicity arises from the endogenous network response.

8.2 The Effect of Trade Liberalization: 1967–2014

The second experiment evaluates the actual liberalization of trade over the sample period. We compare the competitive equilibrium at estimated 1967 trade costs with the equilibrium at 2014 trade costs, keeping risk primitives fixed, to isolate the effect of changing τ on aggregate outcomes.

Figure 8 summarizes the results. The Panel (a) plots the change in the median bilateral $\log \tau$ relative to 1967. The final costs of good trade τ^F steadily decline throughout the sample, falling by roughly five log points by 2014. The intermediate cost of good trade τ^I shows a smaller decline of about three log points, most of which occurred after the 1990s.

The bilateral pattern of these changes is heterogeneous: intra-European pairs experience the largest reductions in trade costs, while many intercontinental corridors show only modest declines (Appendix Figures 15a–15b). In addition, Appendix H.6 compares our baseline estimates ($\gamma = 2$) with a case where risk does not play a role ($\gamma = 1$). In the case of no risk role, the estimated reduction in trade cost is lower for both intermediate and final goods.

Panel (b) breaks down how the volatility of GDP changes, by final and intermediate goods for each country. Each bar shows the difference between the variance of GDP under the trade costs that prevail in 1967 and the variance of GDP under the trade costs that prevail in 2014. A positive bar means that higher trade barriers produce more variance, so trade liberalization reduced volatility for that country. The intermediate cost channel isolates the effect of changing τ^I only (holding τ^F at 2014 levels), the final cost channel isolates τ^F only, and the bars are ordered by the combined total effect.

For most countries, lower trade costs of intermediate inputs led to lower variance in GDP, with the largest reductions in Ireland, Japan, and Australia. Also, for most countries, lower trade costs for final goods led to a higher variance in GDP, with the highest increases in China (−32%) and Denmark (−22%). This split reflects the fact that lower intermediate trade costs primarily deliver diversification benefits (spreading sourcing across more suppliers), while lower final trade costs reorganize consumption baskets toward central foreign producers, raising exposure to global shocks.

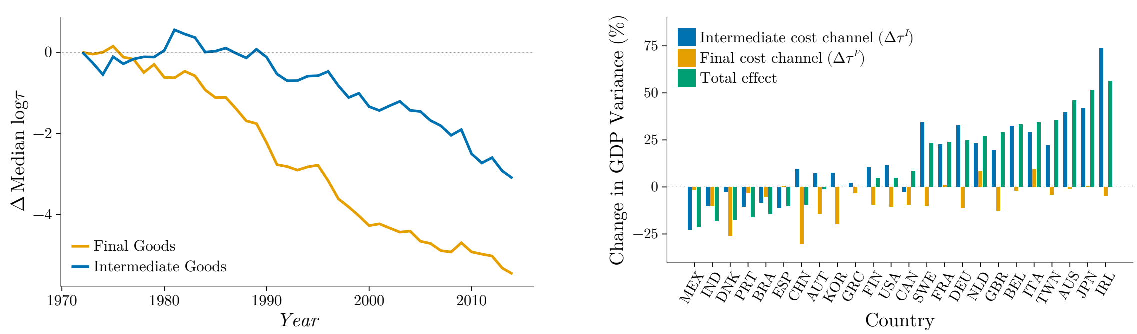

(a) Change in Trade Costs

(b) Change in GDP Variance by Country

Figure 8: Trade Liberalization and Aggregate Risk: 1967–2014

Notes: Panel (a): change in median bilateral log τ relative to 1967, separately for final goods τ^F and intermediate goods τ^I. Panel (b): change in $\mathrm{Var}(\mathrm{GDP}_c)$ between the CE at 1967 and 2014 trade costs, decomposed into the intermediate cost channel ($\Delta\tau^I$ only) and the final cost channel ($\Delta\tau^F$ only), together with the total effect.

The net effects are heterogeneous. For European hub economies (DEU, ITA, BEL, NLD, FRA, GBR, IRL) and several Asia-Pacific economies (JPN, AUS, TWN), lower trade costs led to lower GDP variance: the diversification benefit from lower intermediate trade costs dominates. For Mexico, lower trade costs led to a higher variance in GDP, with the largest net effect (−20%), driven primarily by the reduction in the trade cost of intermediate inputs: as Mexico became more deeply embedded in regional production networks, its exposure to North American supply shocks increased. For several mid-sized economies (USA, CAN, FIN, KOR, AUT), the reductions in intermediate and final trade costs nearly cancel, and the net effect is small.

These results contrast with Caselli et al. (2020). There are two important differences between their framework and mine that explain the different results presented in Panel (b). First, they include a specialization channel: countries concentrate production in comparative-advantage sectors, which can be more volatile, partially offset by diversification benefits. My model does not feature this specialization channel. Second, the elasticity of substitution between suppliers differs between the two models. In Caselli et al. (2020), the substitution elasticity between producers in the same industry but in different countries is linear, while between producers in different industries is CES. In contrast, my model has two relevant elasticities of substitution. In the first stage, the substitution elasticity is constant and governed by κ, which determines how producers reorganize the network in response to reductions in trade costs. In the second stage, the elasticity is Cobb-Douglas, which controls the benefits of diversification in terms of risk.

The split between final and intermediate goods helps to clarify the role of trade in shaping aggregate risk. Intermediate input trade is largely stabilizing: it allows producers to diversify input sources across foreign suppliers, reducing exposure to any single country's shocks. The Final-goods trade reorganizes consumption baskets toward globally central foreign producers, whose prices are more sensitive to global shocks through the input-output network.

8.3 Sensitivity

Quantitative results depend on three parameters: Frisch elasticity ψ, risk aversion γ, and adjustment cost κ. Appendix H.6 reports a sensitivity analysis in which each parameter varies while the others are held at baseline and the model is re-estimated. The main result is robust for all parameter values. Trade liberalization continues to increase the variance in GDP for large economies in all specifications.

The magnitude of the change in the average variance is more sensitive to κ (Figure 21): low adjustment costs allow for the reallocation of a large network and amplify the variance effect, while high κ anchors the network and dampens it. The sensitivity to ψ is also substantial, with a higher elasticity of the labor supply reducing the variance effect. The sensitivity to γ is modest: the mean variance change varies by less than 0.1 percentage points in $\gamma \in [2, 10]$.

9 Conclusions

This paper studies how aggregate shocks reshape production networks through endogenous choice of suppliers, generating two-way feedback between network formation and shock propagation. I develop a quantitative model in which firms choose suppliers before productivity shocks are realized, trading off price efficiency against risk exposure.

Three insights emerge. First, increasing global uncertainty shifts trade from central to peripheral suppliers, as well-connected nodes offer the worst insurance against aggregate risk. Second, openness and risk are non-monotone: moderate liberalization cuts GDP variance by up to 13 percent through diversification, but deep liberalization erodes most of these gains as networks reconcentrate. Third, the 1967-2014 trade liberalization increased the variance for small open economies but decreased it for large hubs like the United States, contrasting with the universal gains found by Caselli et al. (2020) under exogenous networks.

References

Acemoglu, D. and P. D. Azar (2020): "Endogenous production networks," *Econometrica*, 88, 33–82.

Acemoglu, D., V. M. Carvalho, A. Ozdaglar, and A. Tahbaz-Salehi (2012): "The network origins of aggregate fluctuations," *Econometrica*, 80, 1977–2016.

Acemoglu, D., A. Ozdaglar, and A. Tahbaz-Salehi (2017): "Microeconomic origins of macroeconomic tail risks," *American Economic Review*, 107, 54–108.

Acemoglu, D. and A. Tahbaz-Salehi (2025): "The Macroeconomics of Supply Chain Disruptions," *Review of Economic Studies*, 92, 656–695.

Adamopoulos, T. and F. Leibovici (2024): "Trade Risk and Food Security," *Federal Reserve Bank of St. Louis Working Paper*.

Allen, T. and D. Atkin (2022): "Volatility and the Gains from Trade," *Econometrica*, 90, 2053–2092.

Arkolakis, C., A. Costinot, and A. Rodríguez-Clare (2012): "New Trade Models, Same Old Gains?" *American Economic Review*, 102, 94–130.

Arkolakis, C., F. Huneeus, and Y. Miyauchi (2025): "Production Network Formation, Trade, and Welfare," Tech. rep., CESifo Working Paper.

Backus, D. K., P. J. Kehoe, and F. E. Kydland (1992): "International real business cycles," *Journal of political Economy*, 100, 745–775.

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

——— (2024): "Networks, barriers, and trade," *Econometrica*, 92, 505–541.

Bigio, S. and J. La'O (2020): "Distortions in Production Networks," *Quarterly Journal of Economics*, 135, 2187–2253.

Bloom, N., M. Floetotto, N. Jaimovich, I. Saporta-Eksten, and S. J. Terry (2018): "Really Uncertain Business Cycles," *Econometrica*, 86, 1031–1065.

Bonadio, B., Z. Huo, A. A. Levchenko, and N. Pandalai-Nayar (2025): "Globalization, structural change and international comovement," *Journal of Monetary Economics*, 151, 103745.

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

Carvalho, V. M. and X. Gabaix (2013): "The Great Diversification and Its Undoing," *American Economic Review*, 103, 1697–1727.

Caselli, F., M. Koren, M. Lisicky, and S. Tenreyro (2020): "Diversification through trade," *The Quarterly Journal of Economics*, 135, 449–502.

Dekle, R., J. Eaton, and S. Kortum (2007): "Unbalanced Trade," *American Economic Review*, 97, 351–355.

Dew-Becker, I. (2023): "Tail risk in production networks," *Econometrica*, 91, 2089–2123.

Dew-Becker, I., A. Tahbaz-Salehi, and A. Vedolin (2021): "Skewness and time-varying second moments in a nonlinear production network: theory and evidence," Tech. rep., National Bureau of Economic Research.

di Giovanni, J., A. A. Levchenko, and I. Méjean (2014): "Firms, Destinations, and Aggregate Fluctuations," *Econometrica*, 82, 1303–1340.

——— (2018): "The Micro Origins of International Business-Cycle Comovement," *American Economic Review*, 108, 82–108.

Eaton, J. and S. Kortum (2002): "Technology, geography, and trade," *Econometrica*, 70, 1741–1779.

Elliott, M., B. Golub, and M. V. Leduc (2022): "Supply network formation and fragility," *American Economic Review*, 112, 2701–2747.

Fan, J. and W. Luo (2025): "Global Production Networks with Global Uncertainty," *Available at SSRN 5282371*.

Fitzgerald, D. (2025): "3-D Gains from Trade," *Federal Reserve Bank of Minneapolis Working Paper*.

Gabaix, X. (2011): "The granular origins of aggregate fluctuations," *Econometrica*, 79, 733–772.

Greenwood, J., Z. Hercowitz, and G. W. Huffman (1988): "Investment, Capacity Utilization and the Real Business Cycle," *American Economic Review*, 78, 402–417.

Heathcote, J. and F. Perri (2002): "Financial autarky and international business cycles," *Journal of monetary Economics*, 49, 601–627.

Huo, Z., A. Levchenko, and N. Pandalai-Nayar (2024): "International comovement in the global production network," *Review of Economic Studies*, rdae033.

Johnson, R. C. (2014): "Trade in intermediate inputs and business cycle comovement," *American Economic Journal: Macroeconomics*, 6, 39–83.

Jones, C. I. (2011): "Intermediate Goods and Weak Links in the Theory of Economic Development," *American Economic Journal: Macroeconomics*, 3, 1–28.

Kleinman, B., E. Liu, and S. J. Redding (2025): "International Trade in an Uncertain World," Manuscript, Stanford University and Princeton University.

Kopytov, A., B. Mishra, K. Nimark, and M. Taschereau-Dumouchel (2024): "Endogenous production networks under supply chain uncertainty," *Econometrica*, 92, 1621–1659.

Kose, M. A. and K.-M. Yi (2001): "International Trade and Business Cycles: Is Vertical Specialization the Missing Link?" *American Economic Review*, 91, 371–375.

Taschereau-Dumouchel, M. (2025): "Cascades and Fluctuations in an Economy with an Endogenous Production Network," *Review of Economic Studies.*

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, 575–605.

vom Lehn, C. and T. Winberry (2022): "Investment Network, Sectoral Comovement, and the Changing U.S. Business Cycle," *Quarterly Journal of Economics*, 137, 387–433.

Waugh, M. E. (2010): "International Trade and Income Differences," *American Economic Review*, 100, 2093–2124.

Woltjer, P., R. Gouma, and M. P. Timmer (2021): "The Long-Run WIOD: A Global Input–Output Database Covering the Years 1965–2000," GGDC Research Memorandum 186, Groningen Growth and Development Centre.

A Additional Empirical Evidence

A.1 Data Construction

The empirical analysis uses the Long-WIOD dataset, a balanced panel of world input-output tables spanning 1967–2014. The dataset covers 24 major economies and 23 sectors at the ISIC Rev. 3 level, yielding 552 country-sector pairs observed over 48 years. Tables 5 and 4 list the countries and sectors included. These economies account for over 85% of world GDP throughout the sample period.

Sources and chain-linking. The Long-WIOD combines two underlying datasets. The historical tables for 1967–2000 are drawn from Woltjer et al. (2021), who construct a retrospective world input-output database by harmonizing national accounts and bilateral trade data for the pre-WIOD period. The tables for 2001–2014 come from the WIOD 2016 release (?). The two series are chain-linked at the 2000–2001 overlap year by rescaling the historical tables so that the level of each bilateral flow matches the WIOD 2016 value in 2001, ensuring consistency of the intermediate input shares across the splice point.

Key variables. From the input-output tables I construct three main objects used in the empirical analysis and in the structural estimation. *Domar weights* $\omega_{ci,t}$ are the ratio of gross output of country-sector (c, i) to world GDP in year t, computed directly from the supply-use accounts. *Expenditure shares* $\eta_{ci,c'i',t}$ and $\alpha_{c,c'i',t}$ are intermediate and final demand shares, respectively, obtained by dividing each bilateral flow by the total purchases of the buying sector. *Value added* $a_{ci,t}$ is log real value added, deflated using the gross output price index from the socioeconomic accounts following Huo et al. (2024).

Trade deficit series. Country-level trade deficits $\delta_{c,t}$ are computed as the difference between total expenditure (absorption) and total income (value added plus net taxes) in each year, expressed as a share of world GDP. The series is taken directly from the supply-use accounts and requires no additional assumptions.

Sample coverage. The 24 economies in the sample include all G7 members, the four largest emerging markets represented in the WIOD (Brazil, China, India, Mexico), and the major European economies. The 23 sectors span agriculture, mining, manufacturing (14 subsectors), and services (7 subsectors), covering all tradeable and non-tradeable industries at a level of disaggregation sufficient to identify heterogeneous trade costs across country-sector pairs. The balanced panel structure—all 552 pairs observed in all 48 years—is essential for the factor decomposition in Section 2.1, which requires a complete cross-section in each period to identify the global and country-specific components of VA growth.

Code	Description
AtB	Agriculture, Hunting, Forestry & Fishing
C	Mining & Quarrying
D15t16	Food, Beverages & Tobacco
D17t19	Textiles, Leather & Footwear
D21t22	Pulp, Paper, Printing & Publishing
D23	Coke, Refined Petroleum & Nuclear Fuel
D24	Chemicals & Chemical Products
D25	Rubber & Plastics
D26	Other Non-Metallic Mineral Products
D27t28	Basic Metals & Fabricated Metal Products
D29	Machinery, nec
D30t33	Electrical & Optical Equipment
D34t35	Transport Equipment
Dnec	Manufacturing, nec; Recycling
E	Electricity, Gas & Water Supply
F	Construction
G	Wholesale & Retail Trade
H	Hotels & Restaurants
I60t63	Transport & Storage
I64	Post & Telecommunications
J	Financial Intermediation
K	Real Estate, Renting & Business Activities
LtQ	Public Administration & Other Services

Table 4: Sectors in the Long-WIOD Panel (ISIC Rev. 3)

AUS – Australia	AUT – Austria	BEL – Belgium	BRA – Brazil
CAN – Canada	CHN – China	DEU – Germany	DNK – Denmark
ESP – Spain	FIN – Finland	FRA – France	GBR – United Kingdom
GRC – Greece	IND – India	IRL – Ireland	ITA – Italy
JPN – Japan	KOR – South Korea	MEX – Mexico	NLD – Netherlands
PRT – Portugal	SWE – Sweden	TWN – Taiwan	USA – United States

Table 5: Countries in the Long-WIOD Panel

A.2 QQ Plots and the Nature of Aggregate Risk

Figure 9 compares the empirical quantiles of growth rates against the normal distribution. Panel 9a shows that GDP growth exhibits heavy tails, particularly in the left tail where extreme negative outcomes occur far more frequently than a constant-variance normal distribution predicts. Panel 9b shows that sectoral value-added growth exhibits even heavier tails in both directions, consistent with large common shocks that hit individual sectors simultaneously. This discrepancy points to episodes of abnormal synchronization: in most years the cross-sectional standard deviation of sectoral VA growth is approximately 12% and roughly 25% of sectors experience negative growth, but in 2009 dispersion nearly doubled to 18% and the fraction of declining sectors rose to over 80%. The joint increase in dispersion and synchronization is the signature of a large common factor, not of large idiosyncratic shocks (see also Acemoglu, Ozdaglar, and Tahbaz-Salehi, 2017).

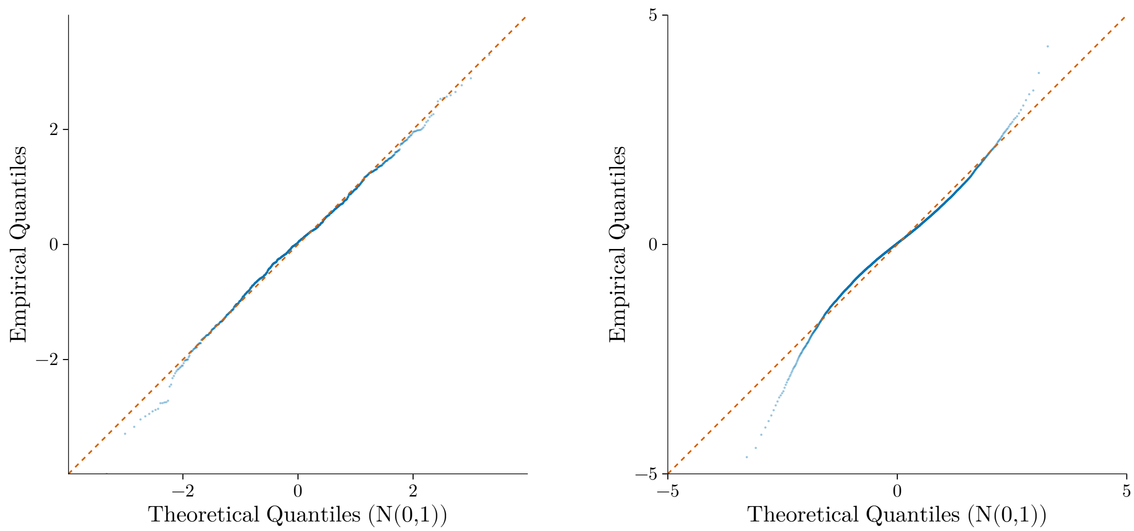

(a) Country GDP growth

(b) Sectoral VA growth

Figure 9: Q-Q Plots of Growth Rates Against the Normal Distribution

Notes: Each panel plots empirical quantiles of annual growth rates against theoretical normal quantiles with the same mean and variance. WIOD data, 23 countries, 23 industries, 1967–2014.

A.3 Country-Level Variance Decomposition

Figure 10 reports the share of each country's GDP variance accounted for by each factor component, averaged over the full sample. The global factor dominates for most countries but with substantial heterogeneity: it accounts for over 75% of GDP variance in France, the United Kingdom, and Germany—open economies deeply integrated into the global production network—but under 35% in China and Brazil, whose large domestic markets make them less exposed to global fluctuations. This heterogeneity foreshadows the model's prediction that the flight-to-safety mechanism has heterogeneous effects depending on each country's position in the global network.

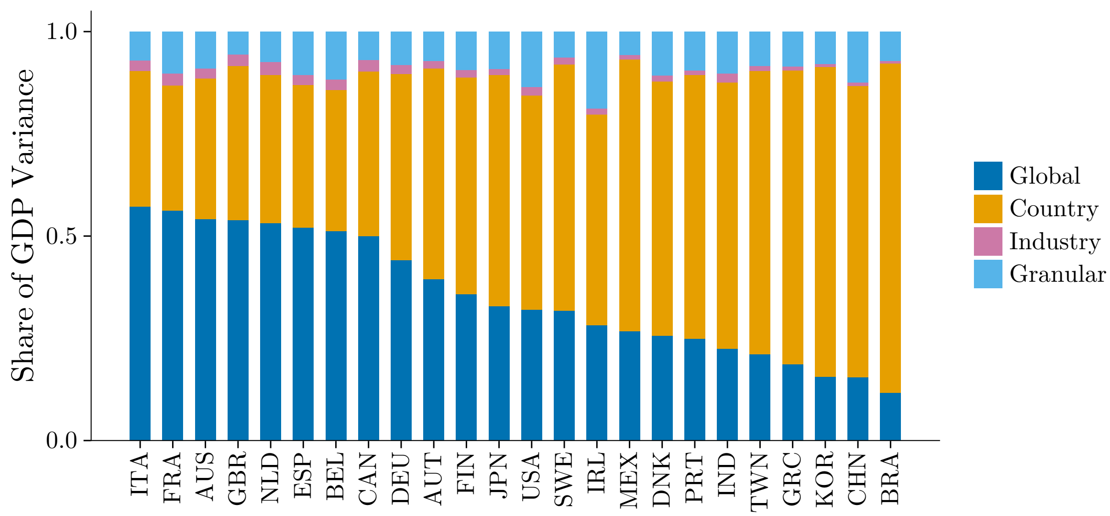

Figure 10: Share of GDP Variance by Component, by Country

Notes: Each bar shows the share of a country's GDP variance accounted for by the global (σ_g^2), country (σ_χ^2), industry (σ_ζ^2), and idiosyncratic (σ_u^2) components of the factor decomposition, averaged over 1967–2014. Countries sorted by global component share. WIOD data.

A second complementary statistic is the average pairwise correlation of VA growth rates across all country-industry pairs in each year:

$$\bar{\rho}_t = \frac{2}{N(N-1)} \sum_{(ci)<(c'i')} \rho_{ci,c'i',t} \tag{18}$$

where $\rho_{ci,c'i',t}$ is the correlation between the VA growth rates of pairs (ci) and $(c'i')$ computed over a rolling window centered at t. When $\bar{\rho}_t$ is high, country-industry units move together more than usual. This is a measure of synchronization.

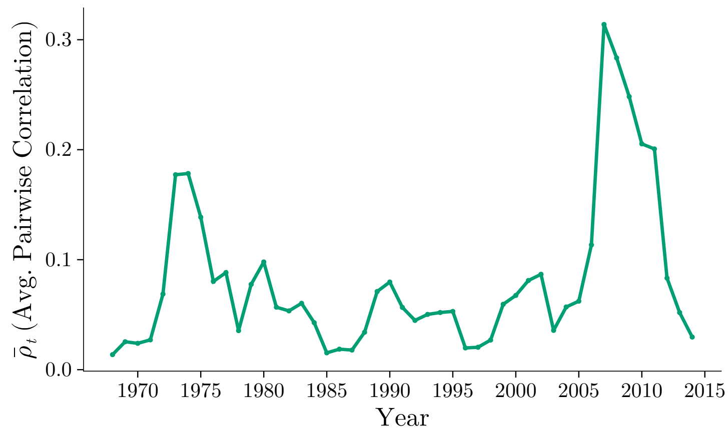

Figure 11: Average pairwise correlation $\bar{\rho}_t$

Notes: Shows the average pairwise correlation $\bar{\rho}_t$ of VA growth across all $N(N-1)/2$ country-industry pairs, computed with a five-year rolling window. WIOD data: 23 countries, 23 industries, 1967-2014.

Figure 11 shows that $\bar{\rho}_t$ is time-varying, with pronounced spikes around 1974, 1980, and 2008: in normal years the average pair of country-industry units moves with a correlation below 0.08, rising to 0.18 at the peak of the 1974 episode and reaching 0.31 in 2008.

A.4 Intermediate Network Snapshots

Figure 12 shows the global trade network in 1984 and 2000. The 1984 network shows Japan's emergence as a major hub: its node has grown substantially relative to 1967, reflecting the expansion of Japanese manufacturing during the 1970s and early 1980s. By 2000, Japan remains large but China begins to appear as a non-trivial node, while the US retains its dominant position. These intermediate snapshots confirm that the structural transformation of the network was gradual, driven primarily by Japan's rise and China's emergence rather than by any discrete event.

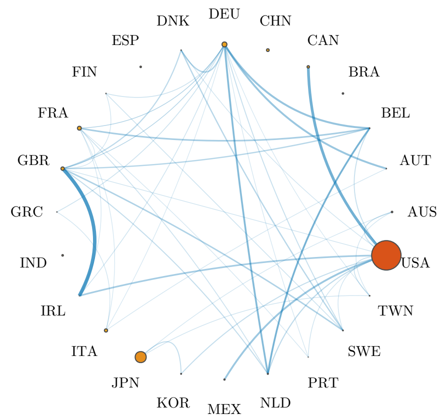

(a) Trade Network, 1984

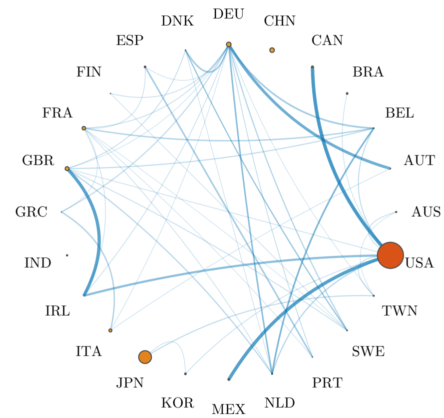

(b) Trade Network, 2000

Figure 12: Global Trade Network: Intermediate Snapshots

Notes: Node size proportional to country Domar weight $\omega_{c,t}$. Edge thickness proportional to bilateral intermediate trade flows above the 75th percentile. WIOD data, 23 countries, 1984 and 2000.

B Proofs of Theoretical Sections

B.1 Lemma 1

Proof. We begin with the revenue of sector ci, which equals its sales to final-good producers plus its sales to intermediate-good producers. The optimality conditions of buyers give the expenditure of final-good producer $\hat{c}$ on supplier ci and the expenditure of intermediate producer $\hat{c}k$ on supplier ci:

$$\tau_{\hat{c},ci} P_{ci} Z_{\hat{c},ci} = P_{\hat{c}} Z_{\hat{c}} \alpha_{\hat{c},ci} \tag{19}$$

$$\tau_{\hat{c}k,ci} P_{ci} Z_{\hat{c}k,ci} = (1 - \mu_{\hat{c}k}) P_{\hat{c}k} Y_{\hat{c}k} \eta_{\hat{c}k,ci} \tag{20}$$

where $\alpha_{\hat{c},ci}$ is the share of buyer $\hat{c}$'s final demand from supplier ci, and $\eta_{\hat{c}k,ci}$ is the share of buyer $\hat{c}k$'s intermediate inputs from supplier ci. Summing over all buyers:

$$P_{ci} Y_{ci} = \sum_{\hat{c}} P_{\hat{c}} Z_{\hat{c}} \alpha_{\hat{c},ci} + \sum_{\hat{c}} \sum_{k} (1 - \mu_{\hat{c}k}) P_{\hat{c}k} Y_{\hat{c}k} \eta_{\hat{c}k,ci} \tag{21}$$

Next, we use the budget constraint of the households, $P_{\hat{c}} Z_{\hat{c}} = \sum_k \mu_{\hat{c}k} P_{\hat{c}k} Y_{\hat{c}k} + \delta_{\hat{c}}$, where $\delta_{\hat{c}}$ is the trade deficit of country $\hat{c}$, to get:

$$P_{ci} Y_{ci} = \sum_{\hat{c}} \alpha_{\hat{c},ci} \left(\sum_{k} \mu_{\hat{c}k} P_{\hat{c}k} Y_{\hat{c}k} + \delta_{\hat{c}} \right) + \sum_{\hat{c}} \sum_{k} (1 - \mu_{\hat{c}k}) P_{\hat{c}k} Y_{\hat{c}k} \eta_{\hat{c}k,ci} \tag{22}$$

In vector notation, with $\boldsymbol{\alpha}$ and $\boldsymbol{\eta}$ defined so that row = buyer and column = supplier:

$$\boldsymbol{PY} = \left(\boldsymbol{\alpha}^\top (I + \boldsymbol{D}_\delta) \boldsymbol{\mu} + \boldsymbol{\eta}^\top (I - \boldsymbol{\mu}) \right) \boldsymbol{PY} \tag{23}$$

where $\boldsymbol{D}_\delta$ is the diagonal matrix of deficit-to-income ratios as defined in Lemma 1.

The column sums of the matrix $\boldsymbol{\alpha}^\top (I + \boldsymbol{D}_\delta) \boldsymbol{\mu} + \boldsymbol{\eta}^\top (I - \boldsymbol{\mu})$ equal one by construction (since for each supplier ci, $\sum_{\hat{c}} \alpha_{\hat{c},ci} (1 + D_{\delta,\hat{c}}) \mu_{\hat{c}} + \sum_{\hat{c}k} \eta_{\hat{c}k,ci} (1 - \mu_{\hat{c}k}) = 1$), so $\lambda = 1$ is an eigenvalue and a solution exists. I impose normalization $\sum \omega_{ci} = 1$ to identify a unique solution, choosing world GDP as the numeraire.

Alternative representation with the standard Leontief inverse. Separate the intermediate and final-demand components:

$$\boldsymbol{\omega} = \boldsymbol{\eta}^\top (I - \boldsymbol{\mu})\, \boldsymbol{\omega} + \boldsymbol{\alpha}^\top \boldsymbol{\mu}\, \boldsymbol{\omega}.$$

The second term is final demand. Define the country-level absorption $D_c = \sum_k \mu_{ck}\, \omega_{ck}$ and the standard IO Leontief inverse $\tilde{\mathcal{L}} = \left(I - \boldsymbol{\eta}^\top (I - \boldsymbol{\mu})\right)^{-1}$. Rearranging:

$$\boldsymbol{\omega} = \tilde{\mathcal{L}}\, \boldsymbol{\alpha}^\top \boldsymbol{\mu}\, \boldsymbol{\omega}.$$

The vector $\boldsymbol{\alpha}^\top \boldsymbol{\mu}\, \boldsymbol{\omega}$ has element (ci) equal to $\sum_{c'} \alpha_{c',ci}\, D_{c'}$, i.e. the total final demand directed to sector (ci) from all countries. In scalar notation:

$$\omega_{ci} = \sum_{\hat{c},k} \tilde{\mathcal{L}}_{ci,\hat{c}k} \sum_{c'} \alpha_{c',\hat{c}k}\, D_{c'}. \tag{24}$$

This is the representation used in the computational implementation: the standard Leontief inverse $\tilde{\mathcal{L}}$ propagates final demand through the intermediate network to obtain Domar weights.

□

B.2 Lemma 2

Proof. **B.2.1 Nominal GDP**

The starting point of the proof is the Domar weights that we characterize in Lemma **??**.

$$P_{ci} Y_{ci} = \omega_{ci} \tag{25}$$

Taking logs we have:

$$\log P_{ci} + \log Y_{ci} = \log(\omega_{ci}) \tag{26}$$

B.2.2 Labor Demand

The FOC of firms for labor are as follows:

$$W_{ci} L_{ci} = \mu_{ci} P_{ci} Y_{ci} \tag{27}$$

$$\log P_{ci} + \log Y_{ci} + \log \mu_{ci} = \log W_{ci} + \log L_{ci} \tag{28}$$

Then:

$$\log L_{ci} = -\log W_{ci} + \log \mu_{ci} + \log(\omega_{ci}) \tag{29}$$

B.2.3 Labor supply

Also, FOC of households implies:

$$P_c(L_{ci})^{\frac{1}{\psi}} = W_{ci} \tag{30}$$

$$\log L_{ci} = [\log W_{ci} - \log P_c]\psi \tag{31}$$

Combining both:

$$\log W_{ci} = \frac{\psi}{1+\psi} \log P_c + \frac{1}{1+\psi}[\log \mu_{ci} + \log(\omega_{ci})] \tag{32}$$

B.2.4 Demand of Inputs

From FOC of intermediate inputs of firms:

$$\log P_{ci} = -\epsilon_i - A(\eta_i^\star) + \mu_{ci} \log W_{ci} + (1-\mu_{ci}) \sum_{\hat{c}} \sum_j \eta_{ci,\hat{c}j}[\log P_{\hat{c}j} + \log(\tau_{c,\hat{c}j})] \tag{33}$$

$$\begin{aligned} \log P_{ci} = -\epsilon_i - A(\eta_i^\star) + \mu_{ci} \left[\frac{\psi}{1+\psi} \log P_c + \frac{1}{1+\psi} (\log \mu_{ci} + \log(\omega_{ci})) \right] + \\ (1-\mu_{ci}) \sum_{\hat{c}} \sum_j \eta_{ci,\hat{c}j}[\log P_{\hat{c}j} + \log(\tau_{c,\hat{c}j})] \end{aligned} \tag{34}$$

$$\begin{aligned} \log P_{ci} = -\epsilon_i - A(\eta_i^\star) + \mu_{ci} \left[\frac{\psi}{1+\psi} \left[-A(\alpha_c^\star) + \sum_{\hat{c}} \sum_j \alpha^\star_{c,\hat{c}j}[\log P_{\hat{c}j} + \log(\tau_{c,\hat{c}j})] \right] + \right. \\ \left. \frac{1}{1+\psi} (\log \mu_{ci} + \log(\omega_{ci})) \right] + (1-\mu_{ci}) \sum_{\hat{c}} \sum_j \eta_{ci,\hat{c}j}[\log P_{\hat{c}j} + \log(\tau_{c,\hat{c}j})] \end{aligned} \tag{35}$$

Turning in matrix notation:

$$\log P[I - \frac{\psi}{1+\psi}\mu\alpha^\star - (I-\mu)\eta] =$$
$$-\epsilon - A(\eta_i^\star) - \frac{\psi}{1+\psi}\mu A(\alpha_c^\star) + \text{diag}\left[\left(\frac{\psi}{1+\psi}\mu\alpha^\star + (I-\mu)\eta\right)\log(\tau)\right] + \frac{1}{1+\psi}\left(\log\mu + \log(\omega)\right) \tag{36}$$

Using the definition of the Leontief Inverse:

$$\log P = \tag{37}$$
$$\mathcal{L}(\alpha^\star)\left(-\epsilon - A(\eta_i^\star) - \frac{\psi}{1+\psi}\mu A(\alpha_c^\star) + \text{diag}\left(\frac{\psi}{1+\psi}\mu\alpha^\star + (I-\mu)\eta\right)\log(\tau) + \frac{1}{1+\psi}\left(\log\mu + \log(\omega)\right)\right)$$

So, applying the definition in the lemma:

$$\boldsymbol{p} = -\mathcal{L}\,\epsilon + \overline{p}(s) \tag{38}$$

□

B.3 Proof of Lemma 6

Proof. The proof starts from the definition of VA:

$$VA = e^{\epsilon A(\eta_{ci})} L_{ci}^{\mu_{ci}} \tag{39}$$

Taking logs lead to the first part of the Lemma. The labor supply expression follows from the proof of Lemma 2:

$$l_{ci} = \frac{\psi}{1+\psi}\left[-p_c + \log\mu_{ci} + \log(\omega_{ci})\right] \tag{40}$$

□

B.4 Proof of Corollary 4

Proof. The result follows directly from taking derivatives of:

$$\mathbf{E}[p_{ci}] = \mathcal{L}\,(B(s) - \theta) \tag{41}$$

$$\mathbf{V}\,[p_{ci}] = \mathcal{L}\,\Sigma\,(\,\mathcal{L})^{\mathrm{T}} \tag{42}$$

□

B.5 Proof of Lemma 9

Proof. By Hulten's theorem, real GDP growth in country c is the local-Domar-weighted sum of real value-added growth across sectors:

$$g_{c,t} = \sum_{i \in N_c} \hat{\omega}_{ci,t-1}\, g_{ci,t}$$

where $\hat{\omega}_{ci} = P_{ci} Y_{ci} / GDP_c$ is the local Domar weight and $g_{ci,t} = a_{ci,t} - a_{ci,t-1}$. By Lemma 6, $a_{ci} = \epsilon_{ci} + A(\eta^{\star}_{ci}) + \mu_{ci}\, l_{ci}$, so:

$$g_{ci,t} = \Delta\epsilon_{ci,t} + \Delta A(\eta^{\star}_{ci,t}) + \mu_{ci}\, \Delta l_{ci,t}.$$

Substituting yields the result. □

B.6 Proof of Proposition 1

Proof. The proof has three steps:

Step 1: Firms Objective Starting from equation (6). For any input-bundle shares, the problem is the following.

$$\max \quad E\left[\Lambda_c\, Y\left(P_{\mathrm{rev}} - K_{\mathrm{cost}}\right)\right],$$

with $(P_{\mathrm{rev}}, K_{\mathrm{cost}}) = (P_c, K_c(\alpha, P))$ for the finals and $(P_{\mathrm{rev}}, K_{\mathrm{cost}}) = (P_{ci}, K_{ci}(\eta, P))$ for the intermediates. Since P_{rev} does not depend on the shares, choosing the shares is equivalent to *minimizing* the discounted expected cost:

$$\min \quad E[\Lambda_c \, Y \, K_{\text{cost}}] \, .$$

Log-moment step. Write $k_{\text{cost}} = \ln K_{\text{cost}}$ and $\lambda_c = \ln \Lambda_c$, $y = \ln Y$. By log-normal identity,

$$\ln E\left[e^{\lambda_c + y + k_{\text{cost}}}\right] \;=\; E[\lambda_c + y + k_{\text{cost}}] \;+\; \tfrac{1}{2}\,\text{Var}(\lambda_c + y + k_{\text{cost}}).$$

Only the part that varies with the shares matters. Using $\frac{1}{2}\,\partial\,\text{Var}(X) = \text{Cov}(X, \partial X)$ and $\partial k_{\text{cost}}/\partial(\text{share on }(c', k)) = p_{c'k}$, the share-wise index that multiplies each supplier (c', k) is

$$\underbrace{E[p_{c'k}] + \ln \tau_{\bullet,c'k}}_{\text{mean delivered price}} \quad + \quad \underbrace{\text{Cov}\big(\lambda_c + y + k_{\text{cost}},\; p_{c'k}\big)}_{\text{risk term}},$$

where $\tau_{\bullet,c'k}$ denotes the relevant iceberg factor (final or intermediate).

Final-good problem (to get (23)). Here $y = y_c$ and $k_{\text{cost}} = k_c$. Under equilibrium conditions for finals,

$$y_c = -\,p_c, \qquad k_c = p_c.$$

Therefore

$$\text{Cov}(\lambda_c + y_c + k_c,\; p_{c'k}) = \text{Cov}(\lambda_c,\; p_{c'k}) + \underbrace{\text{Cov}(-p_c,\; p_{c'k}) + \text{Cov}(p_c,\; p_{c'k})}_{=0} \, .$$

Hence the per-supplier index for finals is

$$E[p_{c'k}] + \ln \tau_{c,c'k} \;+\; \text{Cov}(\lambda_c,\; p_{c'k}) \, .$$

Intermediate-good problem. For intermediates we must use sectoral (own-good) demand Y_{ci}; then $y = \ln Y_{ci} = y_{ci}$ and $k_{\text{cost}} = k_{ci}$. Under assumptions for intermediates,

$$y_{ci} = -\,p_{ci}, \qquad k_{ci} = p_{ci}.$$

Thus

$$\mathrm{Cov}(\lambda_c + y_{ci} + k_{ci},\ p_{c'k}) = \mathrm{Cov}(\lambda_c,\ p_{c'k}) + \underbrace{\mathrm{Cov}(-p_{ci},\ p_{c'k}) + \mathrm{Cov}(p_{ci},\ p_{c'k})}_{=0}.$$

Hence the per-supplier index for intermediates is

$$E[p_{c'k}] + \ln \tau_{ci,c'k} \ + \ \mathrm{Cov}(\lambda_c,\ p_{c'k})\ .$$

Substituting the same pricing-kernel mapping for λ_c (exactly as stated in the text immediately before (24)) yields Eq. (24).

Stacking. Summing these per-supplier indexes with weights $\alpha_{c,c'k}$ (finals) and $\eta_{ci,c'k}$ (intermediates) gives the objective indexes used in (23) and (24), respectively.

Step 2, Pricing kernel

Household pricing kernel in per-sector form. From the budget constraint $P_c z_c = \sum_i W_{ci} L_{ci}$ and the per-sector labor FOC $W_{ci}/P_c = L_{ci}^{1/\psi}$, we obtain

$$z_c = \sum_i \frac{W_{ci}}{P_c} L_{ci} = \sum_i L_{ci}^{1+\frac{1}{\psi}}.$$

Since the disutility of labor is $\sum_i L_{ci}^{1+1/\psi}/(1+1/\psi)$, the GHH net consumption simplifies to

$$z_c - \sum_i \frac{L_{ci}^{1+\frac{1}{\psi}}}{1+\frac{1}{\psi}} = \left(1 - \frac{1}{1+1/\psi}\right)\sum_i L_{ci}^{1+\frac{1}{\psi}} = \frac{1}{\psi+1}\sum_i L_{ci}^{1+\frac{1}{\psi}}.$$

The stochastic discount factor is

$$\Lambda_c = \frac{u'\left(z_c - \sum_i \frac{L_{ci}^{1+1/\psi}}{1+1/\psi}\right)}{P_c} = \frac{\left(\frac{1}{\psi+1}\sum_i L_{ci}^{1+\frac{1}{\psi}}\right)^{-\gamma}}{P_c} = (\psi+1)^{\gamma}\left(\sum_i L_{ci}^{1+\frac{1}{\psi}}\right)^{-\gamma} P_c^{-1},$$

so that

$$\log \Lambda_c = -\gamma\ \log \sum_i L_{ci}^{1+\frac{1}{\psi}} \ + \ \gamma \log(\psi+1) \ - \ p_c(s,\epsilon).$$

From the per-sector equilibrium relation derived in Lemma 2 (combining labor demand $W_{ci}L_{ci} = \mu_{ci}\omega_{ci}$ with labor supply $L_{ci} = (W_{ci}/P_c)^{\psi}$):

$$L_{ci} = \left(\frac{\mu_{ci}\omega_{ci}}{P_c}\right)^{\frac{\psi}{1+\psi}}, \qquad L_{ci}^{1+\frac{1}{\psi}} = \frac{\mu_{ci}\omega_{ci}}{P_c},$$

where the second equality uses $\frac{\psi}{1+\psi} \cdot \left(1 + \frac{1}{\psi}\right) = 1$. Summing across sectors,

$$\sum_i L_{ci}^{1+\frac{1}{\psi}} = \frac{1}{P_c} \sum_i \mu_{ci}\omega_{ci}.$$

Substituting,

$$\begin{aligned}\log \Lambda_c &= -\gamma \left[-p_c + \log \sum_i \mu_{ci}\omega_{ci} \right] + \gamma \log(\psi + 1) - p_c \\ &= (\gamma - 1)\, p_c \; - \; \gamma \log \sum_i \mu_{ci}\omega_{ci} \; + \; \gamma \log(\psi + 1).\end{aligned}$$

Thus

$$\log \Lambda_c(s, \epsilon) = (\gamma - 1)\, p_c(s, \epsilon) + C_c(s),$$

where

$$C_c(s) = -\gamma \, \log \sum_i \mu_{ci}(s)\, \omega_{ci}(s) \; + \; \gamma \log(\psi + 1)$$

depends only on the network s (through the labor-share-weighted Domar weights $\{\mu_{ci}\omega_{ci}\}$) and parameters, but not on shocks.

Step 3: Combining both For final-good producers in country c, the (log) cost index is

$$k_c(\alpha_c; \alpha^\star) = -A_c(\alpha_c) + \sum_{\hat{c},k} \alpha_{c,\hat{c}k} \big(p_{\hat{c}k} + \log \tau_{c,\hat{c}k}\big),$$

while the equilibrium final-good price index satisfies

$$p_c(\alpha^\star) = -A_c(\alpha_c^\star) + \sum_{\hat{c},k} \alpha^\star_{c,\hat{c}k} \big(p_{\hat{c}k} + \log \tau_{c,\hat{c}k}\big).$$

Since $C_c(s)$, $A_c(\cdot)$ and $\log \tau_{c,\hat{c}k}$ are deterministic, they drop out of covariances, and we obtain

$$\begin{aligned}
\operatorname{Cov}\big(\lambda_c, k_c(\alpha_c;\alpha^\star)\big) &= (\gamma-1)\operatorname{Cov}\Big(p_c(\alpha^\star), k_c(\alpha_c;\alpha^\star)\Big) \\
&= (\gamma-1)\operatorname{Cov}\Big(\sum_{\hat{c}',j}\alpha^\star_{c,\hat{c}'j}p_{\hat{c}'j}, \sum_{\hat{c},k}\alpha_{c,\hat{c}k}p_{\hat{c}k}\Big) \\
&= (\gamma-1)\sum_{\hat{c},k}\sum_{\hat{c}',j}\alpha_{c,\hat{c}k}\,\alpha^\star_{c,\hat{c}'j}\operatorname{Cov}\big(p_{\hat{c}k}, p_{\hat{c}'j}\big).
\end{aligned}$$

The final-good producer in country c chooses α_c to minimize expected (log) cost adjusted for risk:

$$\min_{\alpha_c\in\mathcal{A}} \mathbb{E}\big[k_c(\alpha_c;\alpha^\star)\big] \;+\; \operatorname{Cov}\big(\lambda_c, k_c(\alpha_c;\alpha^\star)\big).$$

Substituting the expressions above and rearranging,

$$\begin{aligned}
&\mathbb{E}\big[k_c(\alpha_c;s)\big] + \operatorname{Cov}\big(\lambda_c, k_c(\alpha_c;s)\big) \\
&= -A_c(\alpha_c) + \sum_{\hat{c},k}\alpha_{c,\hat{c}k}\Big(\mathbb{E}[p_{\hat{c}k}] + \log\tau_{c,\hat{c}k} + (\gamma-1)\sum_{\hat{c}',j}\alpha^\star_{c,\hat{c}'j}\operatorname{Cov}\big(p_{\hat{c}k}, p_{\hat{c}'j}\big)\Big).
\end{aligned}$$

Defining the *risk-adjusted delivered price* as

$$\mathcal{R}_{c,\hat{c}k}(s) := \mathbb{E}[p_{\hat{c}k}] + \log\tau_{c,\hat{c}k} + (\gamma-1)\sum_{\hat{c}',j}\alpha^\star_{c,\hat{c}'j}\operatorname{Cov}\big(p_{\hat{c}k}, p_{\hat{c}'j}\big),$$

and using the entropy adjustment cost $A_c(\alpha_c) = \kappa^F\sum_{\hat{c},k}\alpha_{c,\hat{c}k}\log(\alpha_{c,\hat{c}k}/\alpha^0_{\hat{c}k})$, the problem of the final-good producer is

$$\alpha^\star_c \in \arg\min_{\alpha_c\in\mathcal{A}} \kappa^F\sum_{\hat{c},k}\alpha_{c,\hat{c}k}\log\frac{\alpha_{c,\hat{c}k}}{\alpha^0_{\hat{c}k}} + \sum_{\hat{c},k}\alpha_{c,\hat{c}k}\,\mathcal{R}^F_{c,\hat{c}k}(s).$$

For an intermediate producer in sector (c,i), the (log) cost index is

$$k_{ci}(\eta_{ci};s) = -A_{ci}(\eta_{ci}) + \sum_{\hat{c},k}\eta_{ci,\hat{c}k}\big(p_{\hat{c}k} + \log\tau_{ci,\hat{c}k}\big),$$

while the household SDF for country c is the same Λ_c as above. Using again that constants drop from covariances, we obtain

$$\operatorname{Cov}\big(\lambda_c, k_{ci}(\eta_{ci}; s)\big) = (\gamma - 1)\operatorname{Cov}\Big(p_c(\alpha^\star), \sum_{\hat{c},k} \eta_{ci,\hat{c}k} p_{\hat{c}k}\Big)$$
$$= (\gamma - 1)\sum_{\hat{c},k}\sum_{\hat{c}',j} \eta_{ci,\hat{c}k}\, \alpha^\star_{c,\hat{c}'j} \operatorname{Cov}\big(p_{\hat{c}k}, p_{\hat{c}'j}\big).$$

Thus the intermediate producer solves

$$\min_{\eta_{ci}\in\mathcal{E}} \mathbb{E}\big[k_{ci}(\eta_{ci}; s)\big] \;+\; \operatorname{Cov}\big(\lambda_c, k_{ci}(\eta_{ci}; \eta^\star)\big),$$

that is,

$$-A_{ci}(\eta_{ci}) + \sum_{\hat{c},k} \eta_{ci,\hat{c}k}\Big(\mathbb{E}[p_{\hat{c}k}] + \log \tau_{ci,\hat{c}k} + (\gamma - 1)\sum_{\hat{c}',j} \alpha^\star_{c,\hat{c}'j} \operatorname{Cov}\big(p_{\hat{c}k}, p_{\hat{c}'j}\big)\Big).$$

Replacing the risk-adjusted prices and writing the entropy adjustment cost $A_{ci}(\eta_{ci}) = \kappa^I \sum_{\hat{c},k} \eta_{ci,\hat{c}k} \log(\eta_{ci,\hat{c}k}/\eta^0_{ci,\hat{c}k})$, the problem of the intermediate producer is

$$\eta^\star_{ci} \in \arg\min_{\eta_{ci}\in\mathcal{E}} \kappa^I \sum_{\hat{c},k} \eta_{ci,\hat{c}k} \log \frac{\eta_{ci,\hat{c}k}}{\eta^0_{ci,\hat{c}k}} + \sum_{\hat{c},k} \eta_{ci,\hat{c}k}\, \mathcal{R}_{ci,\hat{c}k}(s).$$

Which closes the proof

□

B.7 Proof of Lemma 5

We consider the problem for a generic buyer with shares $s \in \mathcal{A}$ (the unit simplex), ideal technology s^0, adjustment cost $\kappa > 0$, and risk-adjusted prices $\mathcal{R}$. The entropy adjustment cost problem is:

$$s^\star \in \arg\min_{s\in\mathcal{A}} \left\{ \kappa \sum_k s_k \log \frac{s_k}{s^0_k} + \sum_k s_k \mathcal{R}_k \right\},$$

subject to $\sum_k s_k = 1$ and $s_k \geq 0$ for all k. We assume $s^0_k > 0$ and $\sum_k s^0_k = 1$.

Step 1: Forming the Lagrangian. We introduce a Lagrange multiplier λ for the equality constraint. Since $\partial[s_k \log(s_k/s^0_k)]/\partial s_k = \log(s_k/s^0_k) + 1 \to -\infty$ as $s_k \to 0^+$ whenever $s^0_k > 0$, the marginal cost of reducing any share diverges at the boundary, so the minimizer lies in the interior of the simplex and the non-negativity constraints are never active. The Lagrangian is:

$$\mathcal{L}(s,\lambda) = \kappa \sum_k s_k \log \frac{s_k}{s_k^0} + \sum_k s_k \mathcal{R}_k + \lambda \left(1 - \sum_k s_k\right).$$

Step 2: First-Order Conditions. Differentiate $\mathcal{L}$ with respect to each s_k:

$$\frac{\partial \mathcal{L}}{\partial s_k} = \kappa \left(\log \frac{s_k}{s_k^0} + 1\right) + \mathcal{R}_k - \lambda = 0.$$

Solving for s_k:

$$\log \frac{s_k}{s_k^0} = \frac{\lambda - \kappa - \mathcal{R}_k}{\kappa},$$

$$s_k = s_k^0 \exp\left(\frac{\lambda - \kappa - \mathcal{R}_k}{\kappa}\right) = s_k^0 \exp\left(-\frac{\mathcal{R}_k}{\kappa}\right) \cdot \exp\left(\frac{\lambda - \kappa}{\kappa}\right).$$

Step 3: Enforcing the Constraint. Summing over k and imposing $\sum_k s_k = 1$:

$$\exp\left(\frac{\lambda - \kappa}{\kappa}\right) \sum_k s_k^0 \exp\left(-\frac{\mathcal{R}_k}{\kappa}\right) = 1.$$

Defining the partition function $Z = \sum_k s_k^0 \exp(-\mathcal{R}_k/\kappa)$, we have $\exp((\lambda-\kappa)/\kappa) = 1/Z$. Substituting back:

$$s_k^\star = \frac{s_k^0 \exp(-\mathcal{R}_k/\kappa)}{\sum_{k'} s_{k'}^0 \exp(-\mathcal{R}_{k'}/\kappa)}.$$

This is the softmax of $-\mathcal{R}_k/\kappa$ with base measure s^0. All shares are strictly positive whenever $s_k^0 > 0$. The result holds for both intermediate producers (with $s = \eta_{ci}$, $\kappa = \kappa^I$) and final-good producers (with $s = \alpha_c$, $\kappa = \kappa^F$), using the corresponding risk-adjusted prices $\mathcal{R}^I$ and $\mathcal{R}^F$.

B.8 Proof of Lemma 7

From the budget constraint $P_c Z_c = \sum_i W_{ci} L_{ci}$ and the per-sector labor FOC $W_{ci}/P_c = L_{ci}^{1/\psi}$, we have

$$Z_c = \sum_i \frac{W_{ci}}{P_c} L_{ci} = \sum_i L_{ci}^{1+\frac{1}{\psi}}.$$

The per-sector disutility of labor in the household objective is then a fixed fraction of Z_c:

$$\sum_i \frac{L_{ci}^{1+\frac{1}{\psi}}}{1+\frac{1}{\psi}} = \frac{1}{1+1/\psi} Z_c = \frac{\psi}{1+\psi} Z_c,$$

so net consumption simplifies to

$$Z_c - \sum_i \frac{L_{ci}^{1+\frac{1}{\psi}}}{1+\frac{1}{\psi}} = Z_c - \frac{\psi}{1+\psi} Z_c = \frac{1}{1+\psi} Z_c.$$

With CRRA utility,

$$U\left(Z_c - \sum_i \frac{L_{ci}^{1+1/\psi}}{1+1/\psi}\right) = \frac{1}{1-\gamma}\left(\frac{Z_c}{1+\psi}\right)^{1-\gamma}.$$

Using the per-sector equilibrium relation from Lemma 2, $L_{ci}^{1+1/\psi} = \mu_{ci}\omega_{ci}/P_c$, we obtain

$$Z_c = \sum_i L_{ci}^{1+\frac{1}{\psi}} = \frac{1}{P_c}\sum_i \mu_{ci}\omega_{ci},$$

and therefore

$$\log Z_c(s,\epsilon) = -p_c(s,\epsilon) + \log \sum_i \mu_{ci}\omega_{ci}.$$

Since $Z_c - \sum_i \frac{L_{ci}^{1+1/\psi}}{1+1/\psi} = Z_c/(1+\psi)$, we have

$$\log\left(Z_c - \sum_i \frac{L_{ci}^{1+1/\psi}}{1+1/\psi}\right) = -p_c(s,\epsilon) + \log \sum_i \mu_{ci}\omega_{ci} - \log(1+\psi).$$

Assume $\log\big(Z_c - \sum_i L_{ci}^{1+1/\psi}/(1+1/\psi)\big)$ is (approximately) normal. By the lognormal identity,

$$\log \mathbb{E}\big[U(\cdot)\big] = \log\left(\tfrac{1}{1-\gamma}\right) + (1-\gamma)\,\mathbb{E}\left[\log\Big(Z_c - \sum_i \tfrac{L_{ci}^{1+1/\psi}}{1+1/\psi}\Big)\right] + \tfrac{1}{2}(1-\gamma)^2 \operatorname{Var}\left(\log\Big(Z_c - \sum_i \tfrac{L_{ci}^{1+1/\psi}}{1+1/\psi}\Big)\right).$$

From the expression above,

$$\mathbb{E}\left[\log\!\left(Z_c - \sum_i \frac{L_{ci}^{1+1/\psi}}{1+1/\psi}\right)\right] = -\,\mathbb{E}[p_c(s,\epsilon)] + \log\sum_i \mu_{ci}\omega_{ci} - \log(1+\psi),$$

and, since the remaining terms are deterministic,

$$\operatorname{Var}\left(\log\!\left(Z_c - \sum_i \frac{L_{ci}^{1+1/\psi}}{1+1/\psi}\right)\right) = \operatorname{Var}\big(p_c(s,\epsilon)\big).$$

Putting these together,

$$\begin{aligned}\log\mathbb{E}[U(\cdot)] &= \underbrace{\log\!\left(\tfrac{1}{1-\gamma}\right) - (1-\gamma)\log(1+\psi)}_{=:\,\tilde{c}} + (1-\gamma)\ \log\sum_i \mu_{ci}\omega_{ci} \\ &\quad - (1-\gamma)\,\mathbb{E}[p_c(s,\epsilon)] + \tfrac{1}{2}(1-\gamma)^2\ \operatorname{Var}\big(p_c(s,\epsilon)\big).\end{aligned}$$

Define welfare as the log certainty equivalent $W_c(s) = \log \mathrm{CE}_c$. Since $\mathrm{CE}_c^{1-\gamma} = \mathbb{E}[(\text{net consumption})^{1-\gamma}]$, we have $\log \mathrm{CE}_c = \frac{1}{1-\gamma}\log\mathbb{E}[(\text{net consumption})^{1-\gamma}]$. Dividing by $(1-\gamma)$ and absorbing the deterministic constants into b:

$$W_c(s) = b \ + \ \log\sum_i \mu_{ci}\omega_{ci} \ - \ \mathbb{E}[p_c(s,\epsilon)] \ + \ \tfrac{1}{2}(1-\gamma)\ \operatorname{Var}\big(p_c(s,\epsilon)\big),$$

where b collects all terms that depend on parameters but not on the network or shocks.

C Example Economy

This appendix compares three formulations of the network choice problem — CES, quadratic adjustment costs, and entropy (softmax) — using the two-country economy of Section 4.

Setup

There are two countries, Home (H) and Foreign (F), each with a single intermediate good producer. Labor shares differ: $\mu_H > \mu_F$, so Foreign uses intermediates more intensively and has higher centrality ($\nu_F > \nu_H$). Foreign is the star: $\theta_F > \theta_H$. The Home final-good producer allocates shares $\alpha_H = \{\alpha_{H,H},\ \alpha_{H,F}\}$ over the two suppliers subject to $\alpha_{H,H} + \alpha_{H,F} = 1$. Let $\epsilon_c = g + u_c$ with $g \sim N(0, \sigma_g^2)$ and $u_c \sim N(0, \sigma_u^2)$ independent.

Risk-Adjusted Prices

By Proposition 1, the risk-adjusted price of supplier c' for the Home final-good producer is

$$R_{H,c'} = E[p_{c'}] + \log \tau_{H,c'} + (\gamma - 1)\,\mathrm{Cov}(p_H,\ p_{c'}).$$

The relative risk-adjusted price is

$$\Delta R_H \equiv R_{H,F} - R_{H,H} = \underbrace{(\theta_H - \theta_F)}_{\text{productivity gap<0}} + \underbrace{\log \tau_{H,F}}_{\text{trade cost}} + \underbrace{(\gamma - 1)\nu_H(\nu_F - \nu_H)\sigma_g^2}_{\text{risk premium}}. \tag{43}$$

Three Formulations of Optimal Shares

1. CES. With a CES aggregator of elasticity σ and delivered prices $q_{c'} = \tau_{H,c'}\, e^{p_{c'}}$, expenditure shares satisfy

$$\log\left(\frac{\alpha_{H,F}^{\text{CES}}}{\alpha_{H,H}^{\text{CES}}}\right) = \text{const} + (1 - \sigma)\,\Delta R_H,$$

so the Hicks–Allen elasticity is $\sigma^{\text{CES}} = \sigma$ for all ΔR_H. The elasticity is a global constant, independent of the current allocation, aggregate risk, or trade costs. Shares can reach zero: if ΔR_H is sufficiently large, $\alpha_{H,F} = 0$ and the supplier is dropped entirely.

2. Quadratic adjustment costs. Under a quadratic penalty $\kappa \sum_{c'} (\alpha_{H,c'} - \alpha^0_{H,c'})^2$, the optimal shares are linear in risk-adjusted prices:

$$\alpha^{\text{quad}}_{H,F} = \max\left\{\alpha^0_{H,F} - \frac{1}{4\kappa}\,\Delta R_H,\ 0\right\}, \qquad \alpha^{\text{quad}}_{H,H} = \min\left\{\alpha^0_{H,H} + \frac{1}{4\kappa}\,\Delta R_H,\ 1\right\}.$$

The effective elasticity is state-dependent:

$$\sigma^{\text{quad}}(\Delta R_H) = \frac{4\kappa}{(4\kappa\,\alpha^0_{H,F} - \Delta R_H)(4\kappa\,\alpha^0_{H,H} + \Delta R_H)},$$

which varies with ΔR_H and diverges as shares approach the boundary of the simplex. Corner solutions arise when $|\Delta R_H|$ exceeds $4\kappa\alpha^0$: the non-negativity constraint binds and a supplier is dropped from the active set. Handling corners requires tracking the active set and checking KKT conditions at each iteration, which complicates computation in large-scale models. Moreover, the Lipschitz constant of the best-response map is $1/\kappa$, which makes the contraction condition for uniqueness more restrictive.

3. Entropy (softmax) adjustment costs. Under the KL divergence penalty used in this paper, the optimal shares are a softmax (Lemma 5):

$$\alpha^\star_{H,c'} = \frac{\alpha^0_{H,c'}\ \exp(-R_{H,c'}/\kappa)}{\sum_{c''} \alpha^0_{H,c''}\ \exp(-R_{H,c''}/\kappa)}.$$

The log ratio of shares is linear in relative risk-adjusted prices:

$$\log\left(\frac{\alpha^\star_{H,F}}{\alpha^\star_{H,H}}\right) = \log\left(\frac{\alpha^0_{H,F}}{\alpha^0_{H,H}}\right) - \frac{\Delta R_H}{\kappa}, \tag{44}$$

so the Hicks–Allen elasticity is

$$\sigma^{\text{soft}} = \frac{1}{\kappa} \quad \forall\ \Delta R_H. \tag{45}$$

Comparison

The entropy formulation combines two advantages relative to the alternatives. First, shares are strictly positive for any finite κ: because the KL divergence diverges as any share approaches zero, the softmax automatically keeps all suppliers active without requiring a max operator or active-set tracking. This simplifies computation in the multi-country model with $N = C \times J$ suppliers, where managing corners at each iteration would be costly. Second, the Lipschitz constant of the softmax best response is $1/(2\kappa)$ rather than $1/\kappa$ under quadratic costs. This relaxes the contraction condition for uniqueness by a factor of two, making it easier to establish a unique equilibrium for a given κ.

The softmax shares the constant-elasticity property with CES: the log ratio of shares is linear in relative risk-adjusted prices, with slope $1/\kappa$. However, unlike CES, the softmax anchors shares to the base measure α^0, ensuring that even uncompetitive suppliers retain a positive (though potentially small) share. This anchoring is economically motivated by the entropy interpretation: deviating from the ideal technology α^0 is costly, so firms maintain diversified supplier portfolios even when risk-adjusted prices favor concentration.

Comparative Statics in the Example

The comparative statics from Section 4 follow directly from (44) and (43):

$$\frac{\partial}{\partial \sigma_g^2} \log\left(\frac{\alpha^{\star}_{H,F}}{\alpha^{\star}_{H,H}}\right) = -\frac{(\gamma - 1)\nu_H(\nu_F - \nu_H)}{\kappa} < 0.$$

A rise in aggregate risk shifts sourcing away from the more central Foreign supplier, consistent with Proposition 3. Similarly,

$$\frac{\partial}{\partial \log \tau_{H,F}} \log\left(\frac{\alpha^{\star}_{H,F}}{\alpha^{\star}_{H,H}}\right) = -\frac{1}{\kappa} < 0 :$$

higher trade costs reduce the share of the Foreign supplier.

Although the elasticity is constant, the *level* of shares is nonlinear in ΔR_H through the softmax. As $\alpha_{H,F}$ grows large under trade liberalization, further reductions in τ produce smaller absolute changes in shares (since the softmax saturates), even though the log ratio continues to shift at rate $1/\kappa$. This generates the concentration dynamics documented in Figure 4: moderate liberalization diversifies sourcing, but deep liberalization reconcentrates the network in the productive Foreign supplier, raising aggregate volatility through the global component of (14).

D Existence and Uniqueness

This appendix provides detailed derivations for existence and uniqueness of the equilibrium. The strategy is to cast the equilibrium as a static game to analyze existence and uniqueness.

Assumptions Let us begin by stating the assumptions we require in the proof

Assumption 1 (Assumption)**.**

1. *The strategy set $\mathcal{A}$ is the unit simplex, and the productivity function $A_{ci,t}$ is bounded below.*
2. *Input shares satisfy $1 > \mu_{ci,t} > 0$ for all country-industries ci and periods t.*
3. ***Curvature:** $\kappa_i^I, \kappa^F > 0$ for all industries i.*
4. ***Finite frictions:** $\tau_{b,k,t} \geq 1$ for all bilateral pairs (b,k) and periods t.*
5. *There exists $\bar{A} > -\infty$ such that, for all feasible α_c, $A(\alpha_c) \geq \bar{A}$*

It is important to highlight that no new assumptions are introduced here. This statement serves merely as a summary of previously established assumptions needed for this proof.

Notation for uniform bounds. For scalar parameters, I use overline and underline for the maximum and minimum across sectors:

$$\underline{\mu} := \min_k \mu_k, \qquad \overline{\mu} := \max_k \mu_k.$$

For vector-valued objects, I use overline to denote the ℓ_∞ norm:

$$\overline{\log \tau} := \|\log \tau\|_\infty, \qquad \overline{\log \mu} := \|\log \mu\|_\infty, \qquad \overline{\log \omega} := \|\log \omega\|_\infty.$$

D.1 Game Formulation

Players, Strategies I formulate the equilibrium as a game. Using the notation of my model, the players are intermediate goods producers $k = 1, \ldots, N$ and final goods producers $c = 1, \ldots, C$. Since there is a representative firm for each sector, each k and each c can be understood as a type of producer. The actions of each producer are the share vectors

$\eta_k \in \mathcal{A}$ (intermediate goods) and $\alpha_c \in \mathcal{A}$ (final goods), respectively. The strategy space is

$$\mathcal{M} := \left(\prod_{k=1}^{N} \mathcal{A}\right) \times \left(\prod_{c=1}^{C} \mathcal{A}\right),$$

Given that $\mathcal{A}$ is the simplex, $\mathcal{M}$ is nonempty, compact, and convex. The state is the matrix of risk-adjusted prices $R(s)$.

Aggregate State For any $s = (\eta, \alpha) \in \mathcal{M}$, the spectral radius of $H(s) = (I - \mu)\eta + \frac{\psi}{1+\psi}\mu\alpha$ is strictly less than one (since row sums are bounded by $1 - \underline{\mu}/(1 + \psi) < 1$), so the Leontief inverse $\mathcal{L}(s) = (I - H(s))^{-1}$ exists by the Neumann series. By Lemma 2, $\mathbf{p}(s, \epsilon) = \mathcal{L}(s)(B(s) - \epsilon)$ is the unique price vector. $\mathbb{E}[\mathbf{p}(s, \epsilon)]$ and $\mathrm{Cov}[\mathbf{p}(s, \epsilon), \mathbf{p}(s, \epsilon)]$ are therefore unique functions of s. Thus, the aggregate state $R(s)$ from Proposition 1 is a unique function of the aggregate network.

Payoff As shown in Proposition 1, the payoff (negative cost) of each type of producer is:

$$J(c, \alpha_c, s) = -\kappa^F \sum_j \alpha_{c,j} \log \frac{\alpha_{c,j}}{\alpha^0_{c,j}} - \sum_j \alpha_{c,j} R_{cj}(s),$$

$$J(k, \eta_k, s) = -\kappa^I_i \sum_j \eta_{k,j} \log \frac{\eta_{k,j}}{\eta^0_{k,j}} - \sum_j \eta_{k,j} R_{kj}(s),$$

where the first term in each expression is the entropy (KL divergence) adjustment cost penalizing deviations from the ideal network.

Best Response Given s, let $\eta^*(s)$ and $\alpha^*(s)$ be the solution of Lemma 5. Then, best-response map is:

$$\Phi : \mathcal{M} \to \mathcal{M}, \qquad \Phi(s) := (\eta^*(s), \alpha^*(s)).$$

D.2 Existence via Fixed Point

We are ready to prove the first claim of the proposition, existence:

Claim (Continuity of Φ). For every $s \in \mathcal{M}$: (i) $R(s)$ is well defined and continuous in s; (ii) Lemma 5 has a unique solution in $R(s)$; (iii) the map $s \mapsto (\eta^*(s), \alpha^*(s))$ is continuous. Hence Φ is a continuous self-map of $\mathcal{M}$.

Proof. (i) By Lemma 1 the normalized solution $p(s, \epsilon)$ exists, is unique, and continuously depends on s. Then $R(s)$ is also unique and continuous. (ii) The KL divergence $\sum_j s_j \log(s_j/s_j^0)$ is strictly convex on the relative interior of the simplex, so each objective has a unique minimizer. Moreover, since $s_j \log(s_j/s_j^0) \to +\infty$ as $s_j \to 0^+$ whenever $s_j^0 > 0$, the minimizer lies in the interior of the simplex (all shares are strictly positive). (iii) By Lemma 5, the best response is the softmax $s_j^* = s_j^0 \exp(-R_j/\kappa^m)/Z$, which is a C^∞ function of $R(s)$. Since $R(s)$ is continuous in s by (i), the composition $s \mapsto \Phi(s)$ is continuous. Therefore, Φ is a continuous self-map of $\mathcal{M}$. □

Claim (Existence via fixed point). There exists a competitive equilibrium $(s^\star, R(s^\star))$ with $s^\star = (\eta^*, \alpha^*) \in \mathcal{M}$.

Proof. By Claim 1, $\Phi : \mathcal{M} \to \mathcal{M}$ is a continuous self-map of a compact convex nonempty set. By Brouwer's fixed point theorem, there exists $s^* = (\eta^*, \alpha^*) \in \mathcal{M}$ with $\Phi(s^*) = s^*$. Given s^*, solve Lemma 2 to obtain the unique normalized price vector $p(s^\star, \epsilon)$. By construction, each η_k^* and α_c^* optimize their respective problems at $R(s^*)$ and prices satisfy unit costs; hence $(s^*, p(s^*, \epsilon))$ is a competitive equilibrium. □

D.3 GE Effects

To prove uniqueness, we show that the best response map Φ is a contraction in the space of networks. Here, prices, covariances and the risk adjusted costs depend on the whole network. Instead of relying on monotonicity, I bound how much mean prices and covariances can change when the network moves. The maximum change in expected prices and covariances are defined as:

$$L_{pp} := \sup_{s_1 \neq s_2} \frac{\| \operatorname{Cov}(p(s_1), p(s_1)) - \operatorname{Cov}(p(s_2), p(s_2)) \|_\infty}{\|s_1 - s_2\|_\infty} \tag{46}$$

$$L_p := \sup_{s_1 \neq s_2} \frac{\left\| \mathbb{E}[p(s_1)] - \mathbb{E}[p(s_2)] \right\|_\infty}{\|s_1 - s_2\|_\infty}, \qquad s = (\alpha, \eta). \tag{47}$$

Bound on the Change of Price covariance. Recall from Corollary 1.

$$\operatorname{Cov}(p(s_1), p(s_1)) = \mathcal{L} \Sigma_\epsilon \mathcal{L}^\top \tag{48}$$

Claim. There exists a finite constant L_p such that (46) holds. Moreover, one can choose

$$L_{pp} \;\le\; \frac{2\,\bar{\sigma}^2(1+\psi)^3}{\underline{\mu}^3}\,\left[(1-\underline{\mu}) + \tfrac{\psi}{1+\psi}\,\overline{\mu}\right] \tag{49}$$

Proof. Let $s_1 = (\alpha_1, \eta_1)$ and $s_2 = (\alpha_2, \eta_2)$ be two feasible networks, and denote

$$\Delta s := s_1 - s_2, \quad \Delta\mathcal{L} := \mathcal{L}(s_1) - \mathcal{L}(s_2), \quad \Delta B := B(s_1) - B(s_2).$$

Then,

$$\mathrm{Cov}(p(s_1), p(s_1)) - \mathrm{Cov}(p(s_2), p(s_2)) = \mathcal{L}_1\Sigma_\epsilon(\Delta\mathcal{L})^\top + (\Delta\mathcal{L})\Sigma_\epsilon\mathcal{L}_2^\top, \tag{50}$$

Step 1: Bound on $\mathcal{L}(s)$ Recall

$$\mathcal{L}(s) = \left(I - (I-\mu)\eta - \tfrac{\psi}{1+\psi}\,\mu\alpha\right)^{-1}, \qquad s = (\eta, \alpha),$$

and define

$$H(s) := (I-\mu)\eta + \tfrac{\psi}{1+\psi}\,\mu\alpha, \qquad \mathcal{L}(s) = (I - H(s))^{-1}.$$

The row sum of row k of $H(s)$ is

$$\sum_j H_{kj}(s) = (1-\mu_k) + \tfrac{\psi}{1+\psi}\mu_k = 1 - \tfrac{1}{1+\psi}\mu_k.$$

Hence, using $\mu_k \ge \underline{\mu}$,

$$\|H(s)\|_\infty = \max_k \sum_j |H_{kj}(s)| = \max_k \sum_j H_{kj}(s) \le 1 - \tfrac{1}{1+\psi}\underline{\mu} =: x < 1.$$

By the Neumann series,

$$\mathcal{L}(s) = (I - H(s))^{-1} = \sum_{n=0}^{\infty} H(s)^n,$$

and thus

$$\|\mathcal{L}(s)\|_\infty \le \sum_{n=0}^{\infty} \|H(s)\|_\infty^n \le \sum_{n=0}^{\infty} x^n = \frac{1}{1-x} = \frac{1}{\frac{1}{1+\psi}\underline{\mu}} = \frac{1+\psi}{\underline{\mu}},$$

Step 2: Bounds on $\Delta\mathcal{L}$. Write $H_\ell := H(s_\ell)$ and $\mathcal{L}_\ell := \mathcal{L}(s_\ell)$ for $\ell = 1, 2$. Using the matrix identity

$$A^{-1} - B^{-1} = A^{-1}(B - A)B^{-1},$$

with $A = I - H_1$ and $B = I - H_2$, we obtain

$$\mathcal{L}_1 - \mathcal{L}_2 = (I - H_1)^{-1} - (I - H_2)^{-1} = \mathcal{L}_1 \,(H_1 - H_2)\, \mathcal{L}_2.$$

Taking $\|\cdot\|_\infty$ norms and applying the submultiplicativity of the operator norm,

$$\|\mathcal{L}_1 - \mathcal{L}_2\|_\infty \le \|\mathcal{L}_1\|_\infty \,\|H_1 - H_2\|_\infty \,\|\mathcal{L}_2\|_\infty.$$

By step 1,

$$\|\mathcal{L}_1\|_\infty,\ \|\mathcal{L}_2\|_\infty \le \Big(\frac{1+\psi}{\underline{\mu}}\Big)^2,$$

so

$$\|\mathcal{L}_1 - \mathcal{L}_2\|_\infty \le \Big(\frac{1+\psi}{\underline{\mu}}\Big)^2 \|H_1 - H_2\|_\infty. \tag{51}$$

We now bound $\|H_1 - H_2\|_\infty$ in terms of Δs. By definition,

$$H_1 - H_2 = (I - \mu)(\eta_1 - \eta_2) + \tfrac{\psi}{1+\psi}\,\mu(\alpha_1 - \alpha_2) = (I - \mu)\Delta\eta + \tfrac{\psi}{1+\psi}\,\mu\Delta\alpha.$$

For any row k,

$$\sum_j |(H_1 - H_2)_{kj}| \le (1 - \mu_k) \sum_j |\Delta\eta_{kj}| + \tfrac{\psi}{1+\psi}\mu_k \sum_j |\Delta\alpha_{kj}|.$$

Thus,

$$\|H_1 - H_2\|_\infty \le (1 - \underline{\mu})\, \|\Delta\eta\|_\infty + \tfrac{\psi}{1+\psi}\, \overline{\mu}\, \|\Delta\alpha\|_\infty.$$

If we define

$$\|\Delta s\|_\infty := \max\{\|\Delta\eta\|_\infty,\ \|\Delta\alpha\|_\infty\},$$

then

$$\|H_1 - H_2\|_\infty \le \Big[(1 - \underline{\mu}) + \tfrac{\psi}{1+\psi}\, \overline{\mu}\Big]\, \|\Delta s\|_\infty = K_H\, \|\Delta s\|_\infty.$$

Substituting this into (51) yields

$$\|\mathcal{L}(s_1) - \mathcal{L}(s_2)\|_\infty \le \left(\frac{1+\psi}{\underline{\mu}}\right)^2 K_H \,\|\Delta s\|_\infty,$$

Taking $\|\cdot\|_\infty$ norms of (50) and using submultiplicativity:

$$\|\Delta\mathrm{Cov}\|_\infty \le \|\mathcal{L}_1\|_\infty \,\|\Sigma_\epsilon\|_\infty \,\|(\Delta\mathcal{L})^\top\|_\infty + \|\Delta\mathcal{L}\|_\infty \,\|\Sigma_\epsilon\|_\infty \,\|\mathcal{L}_2^\top\|_\infty.$$

Note that $\|A^\top\|_\infty = \|A\|_1$ (maximum column sum). For the Leontief inverse with non-negative entries, $\|\mathcal{L}^\top\|_\infty = \max_j \sum_i \mathcal{L}_{ij}$ (maximum column sum). Since all entries are non-negative and each column sum is bounded by $N \cdot \max_{ij} \mathcal{L}_{ij} \le N(1+\psi)/\underline{\mu}$, and similarly for $\Delta\mathcal{L}$, we obtain:

$$\|\Delta\mathrm{Cov}\|_\infty \le 2\,\bar{\sigma}^2\,\frac{(1+\psi)^2}{\underline{\mu}^2}\cdot\frac{(1+\psi)}{\underline{\mu}}\,K_H\,\|\Delta s\|_\infty,$$

where the factor $\bar{\sigma}^2 = \|\Sigma_\epsilon\|_\infty$ bounds the shock covariance. This yields the stated bound on L_{pp}. □

Bound on the Change of Mean Prices Recall from Corollary 1. that expected prices can be written as

$$\mathbb{E}[p(s)] = -\,\mathcal{L}(s)\,\theta + \mathcal{L}(s)\,B(s). \tag{52}$$

where

$$B(s) = -A(\eta^\star) - \frac{\psi}{1+\psi}\,\mu\,A(\alpha^\star) + \mathrm{diag}\Big(\tfrac{\psi}{1+\psi}\,\mu\,\alpha^\star + (I-\mu)\,\eta^\star\Big)\log\tau + \frac{1}{1+\psi}(\log\mu + \log\omega) \tag{53}$$

Claim (Bounded sensitivity of mean prices). There exists a finite constant L_p such that (47) holds. Moreover, one can choose

$$L_p \;\le\; \left(\overline{\theta} + \bar{B}\right)\left(\frac{1+\psi}{\underline{\mu}}\right)^2\left[(1-\underline{\mu}) + \tfrac{\psi\,\overline{\mu}}{1+\psi}\right], \tag{54}$$

where

$$\bar{B} := \bar{A} + \frac{\psi}{1+\psi}\,\overline{\mu}\,\bar{A} + \overline{\log\tau} + \frac{1}{1+\psi}\big(\overline{\log\mu} + \overline{\log\omega}\big). \tag{55}$$

Proof. Let $s_1 = (\alpha_1, \eta_1)$ and $s_2 = (\alpha_2, \eta_2)$ be two feasible networks, and denote

$$\Delta s := s_1 - s_2, \quad \Delta\mathcal{L} := \mathcal{L}(s_1) - \mathcal{L}(s_2), \quad \Delta B := B(s_1) - B(s_2).$$

From (52),

$$\mathbb{E}[p(s_1)] - \mathbb{E}[p(s_2)] = -\Delta\mathcal{L}\,\theta + \Delta\mathcal{L}\,B(s_2) + \mathcal{L}(s_1)\,\Delta B. \tag{56}$$

Taking $\|\cdot\|_\infty$ norms and using the triangle inequality,

$$\big\|\mathbb{E}[p(s_1)] - \mathbb{E}[p(s_2)]\big\|_\infty \le \|\Delta\mathcal{L}\|_\infty \left(\|\theta\|_\infty + \|B(s_2)\|_\infty\right) + \|\mathcal{L}(s_1)\|_\infty\,\|\Delta B\|_\infty. \tag{57}$$

Step 1 and 2: Bounds on $\mathcal{L}(s)$ and $\Delta\mathcal{L}$. Take steps 1 and 2 from the proof of claim D.3

Step 3: uniform bound on $B(s)$. From (53)

$$\|\mu A(\alpha^\star)\|_\infty \le \overline{\mu}\,\|A(\alpha^\star)\|_\infty \le \overline{\mu}\,\bar{A}.$$

The trade cost term is bounded by

$$\left\|\operatorname{diag}\!\left(\tfrac{\psi}{1+\psi}\,\mu\,\alpha^\star + (I-\mu)\,\eta^\star\right)\log\tau\right\|_\infty \le \|\log\tau\|_\infty,$$

since for each row ci, the diagonal entry $\frac{\psi}{1+\psi}\mu_{ci}\sum_{\hat{c}k}\alpha^\star_{ci,\hat{c}k} + (1-\mu_{ci})\sum_{\hat{c}k}\eta^\star_{ci,\hat{c}k} = \frac{\psi}{1+\psi}\mu_{ci} + (1-\mu_{ci}) \le 1$. Finally,

$$\left\|\frac{1}{1+\psi}(\log\mu + \log\omega)\right\|_\infty \le \frac{1}{1+\psi}\left(\|\log\mu\|_\infty + \|\log\omega\|_\infty\right).$$

Combining these, we obtain the uniform bound

$$\|B(s)\|_\infty \le \bar{B} \quad \text{for all } s, \tag{58}$$

with $\bar{B}$ given by (55).

Step 3: bound on ΔB. We bound the Lipschitz constant of each term in $B(s)$ from (53).

Entropy terms. Under the softmax (12), all equilibrium shares satisfy $m_{\hat{c}k} \ge$

$m^0_{\min}\exp(-2\bar{R}/\kappa_{\min})$, where $\bar{R} := \bar{\theta} + \overline{\log\tau} + (\gamma-1)\bar{\sigma}^2(1+\psi)/\underline{\mu}$ bounds risk-adjusted prices and $\kappa_{\min} := \min_m \kappa^m$. Denote this lower bound by $\underline{m} > 0$. The entropy function $A(m) = -\kappa\sum_{\hat{c}k} m_{\hat{c}k}\log(m_{\hat{c}k}/m^0_{\hat{c}k})$ has gradient $\partial A/\partial m_{\hat{c}k} = -\kappa(\log(m_{\hat{c}k}/m^0_{\hat{c}k})+1)$, bounded in absolute value by $\kappa(\log(1/\underline{m}) + \log(1/m^0_{\min}) + 1)$. Since $\|\Delta m\|_\infty \le \|\Delta s\|_\infty$, the entropy terms contribute at most

$$C_A := \bar{\kappa}\big(|\log\underline{m}| + |\log m^0_{\min}| + 1\big)\big(1 + \tfrac{\psi}{1+\psi}\overline{\mu}\big)$$

to the Lipschitz constant, where $\bar{\kappa} := \max_m \kappa^m$.

Trade cost term. The diagonal term $\operatorname{diag}(\frac{\psi}{1+\psi}[\mu\alpha^\star + (I-\mu)\eta^\star])\log\tau$ is linear in (α,η), so its Lipschitz constant is $\frac{\psi}{1+\psi}\|\log\tau\|_\infty$.

Domar weight term. The Domar weight $\omega(s)$ solves the eigenvector equation $\omega = [\alpha^{\star\top}(I+D_\delta)\mu + \eta^{\star\top}(I-\mu)]\,\omega$. By implicit differentiation, $\|\Delta\log\omega\|_\infty \le C_\omega\|\Delta s\|_\infty$ where $C_\omega := \|\log\omega\|_\infty/(1-\rho_2)$ and $\rho_2 < 1$ is the second-largest eigenvalue of the input-output matrix, bounded away from 1 by the strictly positive labor shares $\mu > 0$.

Combining these three components:

$$\|\Delta B\|_\infty \le C_B\,\|\Delta s\|_\infty, \tag{59}$$

where

$$C_B := C_A + \frac{\psi}{1+\psi}\,\|\log\tau\|_\infty + \frac{C_\omega}{1+\psi}. \tag{60}$$

Step 4: collecting bounds. Substituting (58), and (59) into (57), we obtain

$$\big\|\mathbb{E}[p(s_1)] - \mathbb{E}[p(s_2)]\big\|_\infty \le \left(\frac{1+\psi}{\underline{\mu}}\right)^2 K_H\big(\|\theta\|_\infty + \bar{B}\big)\|\Delta s\|_\infty + \frac{1+\psi}{\underline{\mu}}\,C_B\,\|\Delta s\|_\infty.$$

Absorbing the second term into the first by enlarging $\bar{B}$ if necessary, we obtain the simpler bound

$$\big\|\mathbb{E}[p(s_1)] - \mathbb{E}[p(s_2)]\big\|_\infty \le \left(\overline{\theta} + \bar{B}\right)\left(\frac{1+\psi}{\underline{\mu}}\right)^2 K_H\,\|\Delta s\|_\infty,$$

where $K_H = (1-\underline{\mu}) + \frac{\psi\overline{\mu}}{1+\psi}$, which implies (54) by the definition (47) of L_p. □

D.4 Contraction and Uniqueness

The dependence of risk adjusted costs on the network s is summarized by

$$\|R(s_1) - R(s_2)\|_\infty \;\le\; \Lambda\,\|s_1 - s_2\|_\infty, \qquad \Lambda := |1-\gamma|\,L_{pp} + L_p,$$

for all $s_1, s_2 \in \mathcal{M}$, where $\|\cdot\|_\infty$ denotes the sup-norm and $\|s\|_\infty := \max\{\|\eta\|_\infty, \|\alpha\|_\infty\}$ for $s = (\eta, \alpha)$.

Claim (Uniqueness). If

$$\min\{\kappa^F, \kappa^I\} > \Lambda/2,$$

then the best response operator Φ is a contraction on $\mathcal{M}$. Hence, $s^* = (\eta^*, \alpha^*)$ and the corresponding normalized price vector p^* are unique.

Proof. Step 1 (Inner problems via Lemma 5). By Lemma 5, for each $m \in \{I, F\}$ and each pair (c, k) the optimal share $s^m_{ck}(s)$ satisfies the softmax

$$s^m_{ck}(s) = \frac{s^{0,m}_{ck}\,\exp\!\big(-R_{ck}(s)/\kappa^m\big)}{\sum_{(c',k')} s^{0,m}_{c'k'}\,\exp\!\big(-R_{c'k'}(s)/\kappa^m\big)}.$$

The Jacobian of the softmax with respect to R is

$$\frac{\partial s^m_{ck}}{\partial R_{c'k'}} = \frac{1}{\kappa^m}\big(s^m_{ck}\,s^m_{c'k'} - s^m_{ck}\,\mathbf{1}_{(c,k)=(c',k')}\big).$$

For any row (c, k), the absolute row sum of this Jacobian is

$$\sum_{(c',k')}\left|\frac{\partial s^m_{ck}}{\partial R_{c'k'}}\right| = \frac{s^m_{ck}}{\kappa^m}\Big[(1 - s^m_{ck}) + \sum_{(c',k')\neq(c,k)} s^m_{c'k'}\Big] = \frac{s^m_{ck}}{\kappa^m}\cdot 2(1 - s^m_{ck}) \le \frac{1}{2\kappa^m},$$

where the last inequality uses $s(1-s) \le 1/4$ for $s \in [0, 1]$, so $2s(1-s) \le 1/2$. By the mean value theorem, for any R_1, R_2,

$$|s^m_{ck}(R_1) - s^m_{ck}(R_2)| \le \frac{1}{2\kappa^m}\,\|R_1 - R_2\|_\infty.$$

Taking the sup over (c, k) gives

$$\|s^m(s_1) - s^m(s_2)\|_\infty \;\leq\; \frac{1}{2\kappa^m}\,\|R(s_1) - R(s_2)\|_\infty, \qquad m \in \{I, F\}. \tag{61}$$

Identifying $s^F(s) \equiv \alpha^*(s)$ and $s^I(s) \equiv \eta^*(s)$, (61) becomes

$$\|\alpha^*(s_1) - \alpha^*(s_2)\|_\infty \leq \frac{1}{2\kappa^F}\,\|R(s_1) - R(s_2)\|_\infty, \qquad \|\eta^*(s_1) - \eta^*(s_2)\|_\infty \leq \frac{1}{2\kappa^I}\,\|R(s_1) - R(s_2)\|_\infty.$$

Step 2 (Sensitivity of risk-adjusted costs). By the bounds on mean prices and covariances,

$$\|R(s_1) - R(s_2)\|_\infty \;\leq\; \Lambda\,\|s_1 - s_2\|_\infty \qquad \text{for all } s_1, s_2 \in \mathcal{M}.$$

Step 3 (Contraction of Φ). The aggregate best response operator is

$$\Phi(s) := (\eta^*(s), \alpha^*(s)).$$

Using the product sup-norm on $\mathcal{M}$ and the bounds above,

$$\begin{aligned}
\|\Phi(s_1) - \Phi(s_2)\|_\infty &= \max\left\{\|\eta^*(s_1) - \eta^*(s_2)\|_\infty,\ \|\alpha^*(s_1) - \alpha^*(s_2)\|_\infty\right\} \\
&\leq \max\left\{\frac{1}{2\kappa^I}, \frac{1}{2\kappa^F}\right\}\|R(s_1) - R(s_2)\|_\infty \\
&\leq \frac{\Lambda}{2\min\{\kappa^F, \kappa^I\}}\,\|s_1 - s_2\|_\infty.
\end{aligned}$$

If $2\min\{\kappa^F, \kappa^I\} > \Lambda$, equivalently $\min\{\kappa^F, \kappa^I\} > \Lambda/2$,

$$\frac{\Lambda}{2\min\{\kappa^F, \kappa^I\}} < 1,$$

so Φ is a contraction on the complete metric space $(\mathcal{M}, \|\cdot\|_\infty)$. By Banach's fixed-point theorem, Φ admits a unique fixed point s^*. The corresponding normalized price vector $p^* = p(s^*, \epsilon)$ is unique by Lemma 2, which yields uniqueness of the competitive equilibrium. □

Remark 1 (Empirical verification)**.** *The uniqueness condition can be verified directly from data without knowledge of τ or ε. The only place trade costs enter the bound is through $\overline{\log\tau}$ in $\bar{B}$. By Proposition 4,* $\log\tau$ *is an exact linear function of κ and observables:* $\mathbf{M}\log\tau = -\kappa\,\mathbf{M}(\Delta\log\eta^{\text{data}}) - (\gamma - 1)\mathbf{M}(\Delta Q)$*, where* $\mathbf{M}$ *is the cross-buyer differencing operator from*

Proposition 4 (taking the difference between the domestic and foreign buyer's log shares for each supplier) and Q is the covariance correction computed from $(\eta^{\text{data}}, \mu, \hat{\Sigma}_\varepsilon)$ via the Leontief inverse. Substituting into the uniqueness condition yields $\bar{\kappa} > (a + b\,\bar{\kappa})/2$ for data-computable constants a, b, which reduces to the closed-form threshold $\bar{\kappa} > a/(2 - b)$ whenever $b < 2$. The condition is therefore verifiable at the estimated parameters from value-added growth and observed network shares alone.

E Effect of Aggregate TFP Shocks

E.1 Proof of Proposition 3

Step 1: Risk-Adjusted Prices. By Proposition 1, the risk-adjusted price of supplier $(\hat{c}, k)$ for a buyer in country c is:

$$R_{c,\hat{c}k} = E[p_{\hat{c}k}] + \log \tau_{c,\hat{c}k} + (\gamma - 1)\,\mathrm{Cov}(p_c,\ p_{\hat{c}k}). \tag{62}$$

The global component of the covariance is $\mathrm{Cov}(p_c, p_{\hat{c}k})\big|_g = v_{\hat{c}k}\, v_c\, \sigma_g^2$, where $v_{\hat{c}k} = \sum_j \mathcal{L}_{\hat{c}k,j}$ is the centrality of supplier $(\hat{c}, k)$ and $v_c = \sum_{c'j} \alpha^{\star}_{c,c'j}\, v_{c'j}$ is the country exposure (Definition 2). The labor supply feedback $\psi/(1+\psi)$ is already embedded in v through the Leontief inverse $\mathcal{L}$.

Step 2: Derivative of Risk-Adjusted Price.

$$\frac{\partial R_{c,\hat{c}k}}{\partial \sigma_g^2} = (\gamma - 1)\, v_{\hat{c}k}\, v_c. \tag{63}$$

Step 3: Softmax Derivative. By Lemma 5, the optimal share is a softmax with temperature κ^m. The standard softmax derivative gives:

$$\frac{\partial m^{\star}_{ci,\hat{c}k}}{\partial \sigma_g^2} = -\frac{1}{\kappa^m}\, m^{\star}_{ci,\hat{c}k} \left[\frac{\partial R_{c,\hat{c}k}}{\partial \sigma_g^2} - \sum_{\hat{c}'k'} m^{\star}_{ci,\hat{c}'k'} \frac{\partial R_{c,\hat{c}'k'}}{\partial \sigma_g^2} \right]. \tag{64}$$

Substituting Step 2:

$$\frac{\partial m^{\star}_{ci,\hat{c}k}}{\partial \sigma_g^2} = -\frac{(\gamma - 1)}{\kappa^m}\, m^{\star}_{ci,\hat{c}k} \left(v_{\hat{c}k} - \bar{v}_{ci} \right) v_c \tag{65}$$

where $\bar{v}_{ci} = \sum_{\hat{c}'k'} m^{\star}_{ci,\hat{c}'k'}\, v_{\hat{c}'k'}$ is the share-weighted average centrality. The same derivation applies to both $\eta^{\star}$ (with κ^I) and $\alpha^{\star}$ (with κ^F). □

F GDP and Value Added Characterization

F.1 Output

The nominal output or Gross Domestic Product (GDP) for country c is the sum of the value of the goods produced in the country minus the value of intermediate inputs. This coincides with the total value added earned by producers located in the country:

$$GDP_c = \sum_i \left(P_{ci} y_{ci} - \sum_k \sum_{\hat{c}} p_{ci,\hat{c}k} X_{ci,\hat{c}k} \right) = \sum_{i \in N_c} W_{ci} L_{ci} \tag{66}$$

Following the convention on national accounts, the real GDP of country c is:

$$G_c = \sum_i^J \left(\overline{P}_{ci} Y_{ci} - \sum_{\hat{c}} \sum_k \overline{P}_{\hat{c}k} Z_{ci,\hat{c}k} \right) \tag{67}$$

F.2 Real Value Added

Lemma 8. *The log of real VA, a_{ci}, is:*

$$a_{ci}(s,\epsilon) = \epsilon_{ci} + A(\eta^\star_{ci}) + \mu_{ci}\, l_{ci}(s,\epsilon)$$

where $l_{ci}(s,\epsilon)$ is the log of labor supply:

$$l_{ci}(s,\epsilon) = \frac{\psi}{1+\psi}\Big[-p_c(s,\epsilon) + \log \mu_{ci} + \log(\omega_{ci}(s)) \Big]$$

Proof. See Appendix B.3. □

Lemma 8 decomposes value added into productivity and labor supply. The stochastic component of labor supply is driven by the transmission of TFP shocks through wages: a productivity shock in any (c,i) changes the price index and the real wage, affecting labor supply and value added in all other pairs. This channel, first emphasized by Backus et al. (1992), generates international comovement.[7] The strength of the response depends on the

[7] See also Huo et al. (2024), who characterize deviations from steady-state VA using the Leontief inverse.

exposure of the final-good firm to each good, summarized by $\alpha_c^\star$.

The deterministic part depends on (i) sales and labor shares (Domar weights), since larger or more labor-intensive industries employ more labor in equilibrium; and (ii) the deterministic component of final-good prices, reflecting expenditures on trade costs. Changes in the production network affect the deterministic component of labor across sectors.

F.3 Influence Matrix

Corollary 2. *The influence matrix M, measuring the effect of a TFP shock in sector $\hat{c}k$ on real VA in sector ci, is:*

$$M_{ci,\hat{c}k}(s) = \mathbf{1}_{\{ci=\hat{c}k\}} + \mu_{ci} \frac{\psi}{1+\psi} \sum_{c'} \sum_{j} \alpha^\star_{c,c'j} \mathcal{L}_{c'j,\hat{c}k}(s)$$

The first term captures the direct effect of TFP shocks; the second captures the effect on endogenous labor supply. Define also:

$$\bar{v}_{ci} = \sum_{\hat{c}} \sum_{k} M_{ci,\hat{c}k} \tag{68}$$

which captures the total effect of a global shock on the VA of sector ci, both directly and through network connections. This differs from the centrality measure $v_{ci} = \sum_{\hat{c}k} \mathcal{L}_{ci,\hat{c}k}$ (Definition 2), which is the row sum of the Leontief inverse and measures the sensitivity of *prices*. The VA-based measure $\bar{v}_{ci}$ additionally includes the labor supply response through the $\mu_{ci} \frac{\psi}{1+\psi} \alpha^\star \mathcal{L}$ term.

F.4 Real GDP Growth

Lemma 9. *Let $\hat{\omega}_{ci}(s) = P_{ci} Y_{ci} / GDP_c$ be the local Domar weight and $g_{ci,t} = a_{ci,t} - a_{ci,t-1}$ the growth in VA. Then GDP growth is:*

$$g_{c,t}(s,\epsilon) = \sum_{i} \hat{\omega}_{ci,t-1}(s) \left(\Delta\epsilon_{ci,t} + \Delta A(\eta^\star_{ci,t}) + \mu_{ci,t}\, \Delta l_{ci,t} \right)$$

Proof. See Appendix B.5. □

Lemma 9 aggregates sectoral VA growth into GDP growth using base-year local Domar weights $\hat{\omega}_{ci} = P_{ci}Y_{ci}/GDP_c$. The first two terms inside the brackets provide a Hulten aggregation of changes in exogenous and endogenous TFP, while the last term captures how the equilibrium network generates growth through changes in labor supply. The formula is a first-order discrete-time approximation: when labor shares $\mu_{ci,t}$ vary over time, the exact first difference of $\mu_{ci}l_{ci}$ would include a $\Delta\mu_{ci}$ term, which is omitted here as second-order.

F.5 Growth Decomposition

Let $g_{ci,t} = a_{ci,t} - a_{ci,t-1}$ denote the growth in VA of country c, sector i from $t-1$ to t. Using Lemma 6 and the definition of M above, growth decomposes as:

$$g_{ci,t} = \underbrace{\sum_{\hat{c}k} M_{ci,\hat{c}k,t}\left(\Delta\hat{\epsilon}_{\hat{c}k,t} + \Delta\theta_{\hat{c}k,t}\right)}_{\text{Transmission of shocks}} + \underbrace{\sum_{\hat{c}k} \Delta M_{ci,\hat{c}k,t}\,\epsilon_{\hat{c}k,t-1} + \Delta A(\eta^{\star}_{ci,t})}_{\text{Network Adjustment}} \tag{69}$$

The first component is the standard Acemoglu et al. (2012) propagation of contemporaneous innovations ($\hat{\epsilon}_{\hat{c}k,t}$) and revisions in expected productivity ($\Delta\theta_{\hat{c}k,t}$) through the existing production structure, where $M_{ci,\hat{c}k,t}$ captures both the direct TFP effect and the labor supply response. The second component captures endogenous reallocation of the network: when uncertainty increases, $\Delta M_{ci,\hat{c}k,t}$ tilts exposure towards safer suppliers, potentially lowering expected growth while reducing risk. The term $\Delta A(\eta^{\star}_{ci,t})$ captures the change in endogenous productivity from the network adjustment.

G Variance and Covariance of GDP

This appendix characterizes the variance and covariance of GDP as functions of the production network and the distribution of TFP shocks. The results extend the decomposition in Huo et al. (2024) by isolating the interaction of aggregate shocks with the network.

G.1 Covariance Decomposition

Using the influence vectors v and M defined in Appendix F, the covariance of GDP between two countries decomposes as follows:

Proposition 5. *Let $\sigma_{c'j}$ be the standard deviation of sector-specific shocks. Then, given the network $\alpha^\star, \eta^\star$, the covariance in GDP between any two countries is:*

$$\text{Cov}(g_c, g_{\hat{c}}) = \sum_i \sum_k \hat{\omega}_{ci}\hat{\omega}_{\hat{c}k}\Bigg(\underbrace{\sigma_g^2}_{\textit{Global Shock}} + \underbrace{\sum_{c'} \sum_j \sigma_{c'j}^2 M_{ci,c'j}\, M_{\hat{c}k,c'j}}_{\textit{Sector-Specific Transmission}} + \underbrace{\sigma_g^2 \;[\bar{v}_{ci}\,\bar{v}_{\hat{c}k} - 1]}_{\textit{Global Transmission}} \Bigg)$$

Proof. See below. □

Proposition 5 extends the decomposition proposed by Huo et al. (2024) by separating the effect of global shocks on comovement into two components. The first, σ_g^2, is the direct mechanical effect: comovement increases because a common shock g_t enters each sector's productivity simultaneously, contributing σ_g^2 to the covariance of any pair of sectors regardless of the network structure. The second, $\sigma_g^2[\bar{v}_{ci}\bar{v}_{\hat{c}k} - 1]$, is the network amplification of the global shock. Here $\bar{v}_{ci} = \sum_{\hat{c}k} M_{ci,\hat{c}k}$ is the total effect of a uniform shock on sector ci's value added (Definition in Corollary 2). If the network had no amplification ($\bar{v}_{ci} = 1$ for all sectors), this term would be zero: the -1 subtracts the direct effect already counted in the first component. When $\bar{v}_{ci} > 1$, sector ci amplifies the global shock through its input-output connections, and the product $\bar{v}_{ci}\bar{v}_{\hat{c}k}$ captures the joint amplification across any pair of sectors.

Comovement between two countries increases when the largest sectors in both economies are exposed to the same sector-specific shocks, and when both countries are exposed to the aggregate shock through central sectors (high $\bar{v}_{ci}$ weighted by high ω_{ci}).

G.2 Variance of GDP

As a special case of the covariance decomposition, the variance of GDP for country c is (using local Domar weights $\hat{\omega}_{ci} = P_{ci}Y_{ci}/GDP_c$ from Lemma 9):

$$\text{Var}(g_c) = \sum_i \sum_k \hat{\omega}_{ci}\hat{\omega}_{ck}\Big(\sigma_g^2 + \sum_{c',j} \sigma_{c'j}^2 M_{ci,c'j} M_{ck,c'j} + \sigma_g^2[\bar{v}_{ci}\bar{v}_{ck} - 1]\Big)$$

This can be written more compactly as:

$$\mathrm{Var}(g_c) = \sigma_g^2 \left(\sum_i \hat{\omega}_{ci} \bar{v}_{ci} \right)^2 + \sum_{c',j} \sigma_{c'j}^2 \left(\sum_i \hat{\omega}_{ci} M_{ci,c'j} \right)^2$$

The first term shows that aggregate volatility depends on the local-Domar-weighted average of global influence $\bar{v}_{ci}$: countries whose large sectors are also highly exposed to the global factor have higher GDP volatility. The second term captures the contribution of sector-specific shocks, amplified by the network through the influence matrix M.

G.3 Full Decomposition with Four Shock Components

In the full model with global (g_t), country ($\chi_{c,t}$), industry ($\zeta_{i,t}$), and idiosyncratic ($u_{ci,t}$) shocks, the covariance decomposes into:

$$\begin{aligned}
\mathrm{Cov}(g_c, g_{\hat{c}}) = & \underbrace{\sigma_g^2 \left(\sum_i \omega_{ci} \bar{v}_{ci} \right) \left(\sum_k \omega_{\hat{c}k} \bar{v}_{\hat{c}k} \right)}_{\text{Global}} + \underbrace{\sum_{c'} \sigma_{c'}^2 \left(\sum_i \omega_{ci} V^c_{ci,c'} \right) \left(\sum_k \omega_{\hat{c}k} V^c_{\hat{c}k,c'} \right)}_{\text{Country}} \\
& + \underbrace{\sum_j \sigma_j^2 \left(\sum_i \omega_{ci} V^i_{ci,j} \right) \left(\sum_k \omega_{\hat{c}k} V^i_{\hat{c}k,j} \right)}_{\text{Industry}} + \underbrace{\sum_{c',j} \sigma_{c'j}^2 \left(\sum_i \omega_{ci} M_{ci,c'j} \right) \left(\sum_k \omega_{\hat{c}k} M_{\hat{c}k,c'j} \right)}_{\text{Idiosyncratic}}
\end{aligned}$$

where $V^c_{ci,c'}$ and $V^i_{ci,j}$ are the influence vectors for country-specific and industry-specific shocks respectively.

G.4 Asymptotic Behavior

This section extends Acemoglu et al. (2012) by considering aggregate shocks. The following result characterizes the behavior of GDP comovement as the number of sectors grows.

As $N \to \infty$ with all sectors becoming arbitrarily small, the production network becomes irrelevant for GDP comovement. In this limiting case, comovement depends only on the aggregate (global) shock:

$$\lim_{N \to \infty} \mathrm{Cov}(g_c, g_{\hat{c}}) = \sigma_g^2$$

Similarly, the variance of GDP in the limit is explained by country-specific and global shocks alone. The network transmission terms vanish because with infinitely many small sectors, the diversification argument of Acemoglu et al. (2012) applies.

H Estimation Algorithm

This appendix describes the estimation procedure and counterfactual computation. The algorithm takes as given the balanced panel of $C = 24$ countries and $J = 23$ sectors over $T = 48$ years (1967–2014), constructed by chain-linking the Long-WIOD and WIOD-2016 databases at year 2000 (Step 0).

H.1 Inputs

The following objects are computed from data and do not depend on structural parameters:

- Observed networks: $\eta^{\text{data}}_{ci,\hat{c}k,t}$ (intermediate shares), $\alpha^{\text{data}}_{c,\hat{c}k,t}$ (final-demand shares).
- Intermediate cost shares $\mu_{ci,t}$ and Domar weights $\omega_{ci,t}$.
- Log value added $a_{ci,t}$.
- Ideal technology $\eta^0_{i,k}$: cross-country average of η^{data} under Assumption 1 (η^0 is an industry-pair blueprint).

The structural parameters are (ψ, γ, κ), where ψ is the Frisch elasticity, γ is risk aversion, and κ is the softmax temperature (uniform across buyers in the baseline).

H.2 Step 1: GARCH Estimation of $\hat{\Sigma}^a_t$ and Recovery of Σ^ε_t

The time-varying covariance matrix of TFP shocks is estimated from value-added data, independently of (γ, κ, τ).

1. **Factor decomposition.** Project $\tilde{a}_{ci,t}$ onto the factor structure:

$$a_{ci,t} = f^1_t + f^2_{c,t} + f^3_{j,t} + f^4_{ci,t}, \tag{70}$$

   estimated following Caselli et al. (2020) by running a panel reg with time effects, country and industry specific dummies.

2. **AR(1) persistence.** Estimate ρ_k for each factor $k \in \{g, \chi, \zeta, u\}$ by pooled OLS, clamping $|\rho_k| \leq 0.999$.

3. **Innovations.** Compute $v_{k,t} = k_t - \rho_k\, k_{t-1}$ for each factor. Winsorize at $3 \times$ MAD.

4. **Panel GARCH(1,1).** Estimate common (α, β) by maximizing the panel log-likelihood across all factor series simultaneously:

$$\hat{\sigma}^{2,a}_{k,t} = (1 - \alpha - \beta)\, \bar{\sigma}^2_k + \alpha\, v^2_{k,t-1} + \beta\, \hat{\sigma}^{2,a}_{k,t-1}, \tag{71}$$

   with $\hat{\sigma}^{2,a}_k(t_0) = \mathrm{Var}(v_k)$ and the constraint $\alpha + \beta \leq 0.80$.

5. **Assemble $\hat{\Sigma}^a_t$.** Construct the time-varying covariance matrix of value-added innovations. Its diagonal elements are $\hat{\sigma}^{2,a}_{ci,t} = \hat{\sigma}^{2,a}_{g,t} + \hat{\sigma}^{2,a}_{\chi_c,t} + \hat{\sigma}^{2,a}_{\zeta_j,t} + \hat{\sigma}^{2,a}_{u_{ci},t}$.

6. **Recover $\hat{\Sigma}^\varepsilon_t$.** Lemma 6 implies $\mathbf{a}_t = M_t\, \boldsymbol{\varepsilon}_t + \mathbf{D}_t$, where $\mathbf{D}_t$ is deterministic given the network state. It follows that

$$\hat{\Sigma}^\varepsilon_t = M_t^{-1}\, \hat{\Sigma}^a_t\, (M_t^{-1})^\top, \qquad M_t \equiv I + \tfrac{\psi}{1+\psi}\, \mathrm{diag}(\boldsymbol{\mu}_t)\, \alpha^\star_t\, L_t, \tag{72}$$

   evaluated at the estimated network $(\hat{\eta}^\star_t, \hat{\alpha}^\star_t)$.

H.3 Step 2: Trade Cost Identification

Bilateral trade costs are recovered exactly from the softmax FOCs at the observed economy each year, given the structural parameters (κ, γ, ψ) and the average TFP shock $\hat{\theta}$. No linearization around a steady state is used: every quantity below $(\eta^{\text{data}}, \alpha^{\text{data}}, \mu, \mathcal{L}, \mathrm{Cov}(\mathbf{p}))$ is computed at the data shares for the year being inverted.

The softmax FOC for buyer (c, k) purchasing from supplier $(\hat{c}, k')$ is:

$$\log \eta^\star_{ck,\hat{c}k'} = \log \eta^0_{ck,\hat{c}k'} - \kappa^{-1}\left[\mathbb{E}[\log p_{\hat{c}k'}] + \log \tau^\eta_{c,\hat{c}k'} + R^\eta_{ck,\hat{c}k'}\right] - \log Z_{ck}, \tag{73}$$

where $R^\eta_{ck,\hat{c}k'} = (\gamma - 1) \sum_m \mathcal{L}_{ck,m} \hat{\sigma}^2_m \mathcal{L}^\top_{\hat{c}k',m}$ is the risk premium and Z_{ck} is the partition function. The analogous FOC for α replaces $\eta \to \alpha$ and the buyer index by the country.

Asymmetric trade cost. Take the FOC at buyer (c, k) for the foreign supplier $(\hat{c}, k)$ and the domestic same-sector reference (c, k), which has $\log \tau^\eta_{c,ck} = 0$. Subtracting the two

FOCs (the partition function Z_{ck} cancels because the buyer is the same):

$$\log \tau^{\eta}_{c,\hat{c}k} = d^{\eta}_{(\hat{c},c,k)} + \big(\mathbb{E}[\log p_{ck}] - \mathbb{E}[\log p_{\hat{c}k}]\big), \tag{74}$$

where the data-only constant is

$$d^{\eta}_{(\hat{c},c,k)} = \kappa \, \log \frac{\eta^{\text{data}}_{ck,ck}}{\eta^{\text{data}}_{ck,\hat{c}k}} - \kappa \, \log \frac{\eta^{0}_{ck,ck}}{\eta^{0}_{ck,\hat{c}k}} + \big(R^{\eta}_{ck,ck} - R^{\eta}_{ck,\hat{c}k}\big). \tag{75}$$

The analogous formula for τ^{α} replaces $\eta \to \alpha$ at the buyer side and the destination by the country c. Unlike the symmetric (Head–Ries) closed form, this derivation does not average the FOC for $\hat{c} \to c$ with the FOC for $c \to \hat{c}$; the $\mathbb{E}[\log p]$ terms therefore survive and $\tau^{\eta}_{c,\hat{c}k} \neq \tau^{\eta}_{\hat{c},ck}$ in general.

Solving the system for $\boldsymbol{\tau}$. The closed form (74) expresses $\log \tau$ in terms of data and the price expectations $\mathbb{E}[\log p]$. But $\mathbb{E}[\log p]$ itself depends on τ through the price system, so (74) is implicit. We now derive the explicit closed-form solution.

Stack $\log \tau$ into a vector $\boldsymbol{\tau}$ indexed by directed pairs $(\hat{c}, c, k)$ and $\mathbb{E}[\log p]$ into a vector $\boldsymbol{p}$ indexed by sector cells (c, k). By Lemma 2, the price system is:

$$B = B_0 + (I - \mu)\, \eta^{\top} \tau^{\eta} + \tfrac{\psi}{1+\psi} \mu \, \alpha^{\top} \tau^{\alpha}, \qquad \boldsymbol{p} = \mathcal{L}\,(B - \hat{\theta}), \tag{76}$$

where B_0 collects the τ-free terms ($A(\eta^{\star})$, $A(\alpha^{\star})$, $\log \mu$, $\log \omega$). Both τ-induced terms in B are linear in $\boldsymbol{\tau}$, so we can write

$$B = B_0 + \mathbf{A}\,\boldsymbol{\tau}, \tag{77}$$

where $\mathbf{A}$ is the sparse matrix that maps a directed-pair entry $\tau^{m}_{c,\hat{c}k}$ to its contribution to $B_{(c,k)}$ via the $(I - \mu)\eta^{\top}$ block (for $m = \eta$) or the $\frac{\psi}{1+\psi} \mu \alpha^{\top}$ block (for $m = \alpha$). Substituting (77) into (76):

$$\boldsymbol{p} = \mathcal{L}\, B_0 + \mathcal{L}\,\mathbf{A}\,\boldsymbol{\tau} - \mathcal{L}\,\hat{\theta}. \tag{78}$$

Define the difference operator $\mathbf{D}$ that maps $\boldsymbol{p}$ to the cross-buyer differences appearing in (74):

$$[\mathbf{D}\,\boldsymbol{p}]_{(\hat{c},c,k)} = \mathbb{E}[\log p_{ck}] - \mathbb{E}[\log p_{\hat{c}k}]. \tag{79}$$

Each row of $\mathbf{D}$ has entries +1 in column (c, k) and −1 in column $(\hat{c}, k)$ and zeros elsewhere.

Stacking (74) across all $(\hat{c}, c, k)$:

$$\boldsymbol{\tau} = \mathbf{d} + \mathbf{D}\,\boldsymbol{p}, \tag{80}$$

where $\mathbf{d}$ collects the data-only constants $d^{m}_{(\hat{c},c,k)}$. Substituting (78) into (80):

$$\boldsymbol{\tau} = \mathbf{d} + \mathbf{D}\,\mathcal{L}\,B_0 + \mathbf{D}\,\mathcal{L}\,\mathbf{A}\,\boldsymbol{\tau} - \mathbf{D}\,\mathcal{L}\,\hat{\theta}. \tag{81}$$

Rearranging:

$$(I - \mathbf{D}\,\mathcal{L}\,\mathbf{A})\,\boldsymbol{\tau} = \mathbf{d} + \mathbf{D}\,\mathcal{L}\,(B_0 - \hat{\theta}). \tag{82}$$

The closed-form solution is:

$$\boldsymbol{\tau} = (I - \mathbf{D}\,\mathcal{L}\,\mathbf{A})^{-1}\left[\,\mathbf{d} + \mathbf{D}\,\mathcal{L}\,(B_0 - \hat{\theta})\,\right]. \tag{83}$$

All quantities $(\mathbf{D}, \mathcal{L}, \mathbf{A}, \mathbf{d}, B_0)$ are computed at the observed economy each year; $\hat{\theta}$ enters as an exogenous input from the outer loop. The inversion is one linear solve, with no iteration in $\boldsymbol{\tau}$.

Outer fixed-point loop on $\hat{\theta}$. The inversion above takes $\hat{\theta}$ as given. After recovering τ, productivity shocks are updated via $\hat{\varepsilon}_t = M_t^{+}(\mathrm{va}_t - D_t)$ where M_t, D_t are built from the data shares and the new τ, and the aggregate is $\hat{\theta} = \mathrm{mean}_t(\hat{\varepsilon}_t)$. The outer loop iterates the pair $(\tau, \hat{\theta})$ until $\|\Delta\hat{\theta}\|_\infty < 10^{-5}$ (typically 5–6 iterations). The system is strongly contractive.

Identification and active set. A direction $(\hat{c}, c, k)$ is identified when all four bilateral shares $\eta^{\text{data}}_{ck,ck}, \eta^{\text{data}}_{\hat{c}k,\hat{c}k}, \eta^{\text{data}}_{ck,\hat{c}k}, \eta^{\text{data}}_{\hat{c}k,ck}$ and the corresponding four prior cells exceed the floor 10^{-8}. Otherwise the cell is set to a prohibitive default $\bar{\tau} = 5$ (in logs), which acts as a soft block on those pairs in the CE.[8]

Comparison with the symmetric closed form. Imposing Head–Ries symmetry $\log \tau^{\eta}_{c,\hat{c}k} = \log \tau^{\eta}_{\hat{c},ck}$ and averaging the two reciprocal FOCs cancels the $\mathbb{E}[\log p]$ differences, yielding the per-cell scalar formula

$$\log \tau^{\eta,\text{sym}}_{(\hat{c},c,k)} = \tfrac{\kappa}{2}\log\frac{\eta^{\text{data}}_{\hat{c}\hat{c}}\,\eta^{\text{data}}_{cc}}{\eta^{\text{data}}_{\hat{c}c}\,\eta^{\text{data}}_{c\hat{c}}} - \tfrac{\kappa}{2}\log\frac{\eta^{0}_{\hat{c}\hat{c}}\,\eta^{0}_{cc}}{\eta^{0}_{\hat{c}c}\,\eta^{0}_{c\hat{c}}} + \tfrac{\gamma-1}{2}\Delta\mathrm{Cov}(\log p), \tag{84}$$

[8]The entropy formulation guarantees all equilibrium shares are strictly positive (Lemma 5). In practice, the softmax assigns exponentially small shares to inactive links, and the floor and prohibitive default ensure numerical stability without contradicting the no-corner-solutions property.

which forces $\tau_{AB} = \tau_{BA}$ and requires no linear-system solve. The asymmetric inversion of (82)–(74) retains the directionality that the averaging collapses, at the cost of solving a single linear system per year.

H.4 Step 3: Calibration of (κ, γ)

Given fixed ψ, the parameters (κ, γ) are calibrated by minimizing the distance between CE model shares and data shares:

$$(\hat{\kappa}, \hat{\gamma}) = \arg\min_{\kappa,\gamma} \frac{\mathrm{MSE}(\eta^{\mathrm{CE}} - \eta^{\mathrm{data}})}{\overline{\eta^{\mathrm{data}}}^2} + \frac{\mathrm{MSE}(\alpha^{\mathrm{CE}} - \alpha^{\mathrm{data}})}{\overline{\alpha^{\mathrm{data}}}^2}, \tag{85}$$

where η^{CE} is the converged CE equilibrium at the estimated $\hat{\tau}(\kappa, \gamma)$. The minimization uses grid search over (κ, γ) followed by Nelder-Mead refinement (Fminbox). The baseline calibration is $\psi = 0.72$, $\gamma = 2.0$, $\kappa = 5.0$.

H.5 Step 4: Competitive Equilibrium Solver

Given $(\hat{\tau}_t, \hat{\Sigma}^{\varepsilon}_t, \hat{\gamma}, \hat{\kappa}, \psi)$, the CE network $s^{\star}_t = (\eta^{\star}_t, \alpha^{\star}_t)$ is the fixed point of the following system.

Equilibrium objects. At candidate network $s^{(n)}_t$, using the (buyer, supplier) convention of the main text:[9]

$$\mathcal{L}^{(n)}_t = \left[I - \frac{\psi}{1+\psi}\, \mu\, \alpha^{\star(n)}_t - (I - \mu)\, \eta^{\star(n)}_t \right]^{-1}, \tag{86}$$

$$B^{(n)}_t = B^{(0),(n)}_t + d^{F,(n)}_t \log \hat{\tau}^F_t + d^{I,(n)}_t \log \hat{\tau}^I_t, \tag{87}$$

$$\mathbb{E}^{(n)}[\mathbf{p}_t] = \mathcal{L}^{(n)}_t\, (B^{(n)}_t - \hat{\theta}_t), \tag{88}$$

$$\mathrm{Cov}^{(n)}(\mathbf{p}_t) = \mathcal{L}^{(n)}_t\, \hat{\Sigma}^{\varepsilon}_t\, (\mathcal{L}^{(n)}_t)^{\top}, \tag{89}$$

where $\hat{\theta}_t = \hat{\rho}\, \hat{\varepsilon}_{t-1}$ is the conditional mean (set to zero in the baseline).

These are computed by the canonical function `compute_expected_prices` in `functions.jl`.

[9] The code stores IO matrices in transposed form (supplier × buyer), so the implementation transposes $\alpha^{\star}$ and $\eta^{\star}$ before computing $\mathcal{L}$.

Risk-adjusted prices. For intermediate buyer (c, i) and supplier $(\hat{c}, k)$:

$$R^{I,(n)}_{\hat{c}k}(ci) = \mathbb{E}^{(n)}[p_{\hat{c}k}] + \log \hat{\tau}^{I}_{c,\hat{c}k} + (\gamma - 1)\, \alpha_c^{\star(n)\top}\, \mathrm{Cov}^{(n)}(\mathbf{p})\, e_{\hat{c}k}. \tag{90}$$

Softmax update.

$$\eta^{(n+1)}_{ci,\hat{c}k} = \frac{\eta^{0}_{ci,\hat{c}k}\, \exp\Big(-R^{I,(n)}_{\hat{c}k}(ci)/\kappa\Big)}{\sum_{(\hat{c}',k')} \eta^{0}_{ci,\hat{c}'k'}\, \exp\Big(-R^{I,(n)}_{\hat{c}'k'}(ci)/\kappa\Big)}. \tag{91}$$

Under the entropy cost, shares are strictly positive and sum to one—no simplex projection is needed.

Solver. The fixed point is found using Newton-GMRES with temperature annealing (`solve_counterfactual` in `functions.jl`):

- Initialize at data shares $s^{(0)} = s^{\text{data}}$.
- JFNK: Jacobian-free Newton-Krylov with GMRES inner solver and backtracking line search.
- Temperature annealing: start at $\kappa_{\text{eff}} = \kappa \times T_{\text{mult}}$ (broad softmax), reduce T_{mult} toward 1 as residual decreases.
- Convergence: $\|s^{(n+1)} - s^{(n)}\|_\infty < 10^{-3}$.

Internal consistency. At the data equilibrium with estimated $\hat{\tau}_t$, the solver must converge in one iteration: $s^{(1)} = s^{(0)} = s^{\text{data}}$. This is verified by checking $\|F(s^{\text{data}})\|_\infty < 10^{-4}$.

H.6 Counterfactual Exercises

Trade cost homotopy. To compute equilibria at counterfactual $\tau^\star \neq \hat{\tau}$, the solver traces a path:

$$\tau(\varphi) = (1 - \varphi)\, \hat{\tau} + \varphi\, \tau^\star, \qquad \varphi \in [0, 1], \tag{92}$$

solving the CE at each step and using the previous solution as warm start. This is implemented in `run_counterfactuals.jl` and `run_theta_path.jl`.

Historical counterfactual ($\tau_{1977} \longrightarrow \tau_{2014}$). The welfare gain from trade liberalization is computed using hat algebra (`solve_hat_algebra.jl`), which works in changes $\hat{\eta} = \eta^\star/\bar{\eta}$ to avoid level effects that can overflow the consumption-equivalent formula. The homotopy proceeds in 20 steps with Anderson acceleration at each step.

Variance decomposition. The change in GDP variance from the counterfactual $\tau_{2014} \longrightarrow \tau_{1977}$ is decomposed into intermediate ($\Delta\tau^I$ only) and final ($\Delta\tau^F$ only) channels by solving the CE with each channel shut down independently.

Appendix: Additional Quantitative Results

This appendix reports additional detail on model fit, estimated trade costs, and sensitivity of the main results.

Model Fit

Figure 13 compares the Domar weights and the variance of GDP implied by the model with their data counterparts. Panel (a) is a scatter plot of country Domar weights for all 24 countries; Panel (b) shows the country-level variance of GDP in the data and in the model.

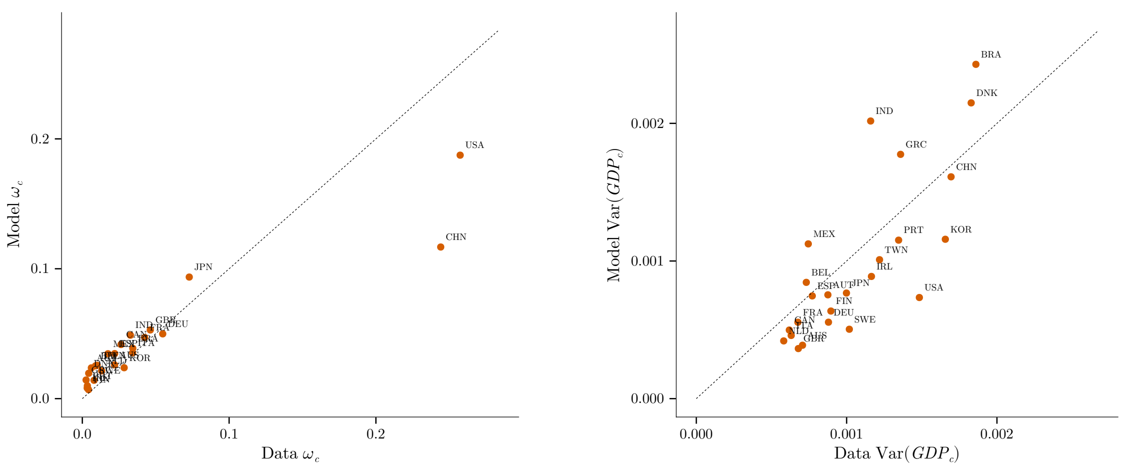

(a) Domar weights

(b) GDP Variance

Figure 13: Model Fit

Notes: Panel (a): each point is a country, time-averaged Domar weight ω_c. Panel (b): country-level GDP variance, data vs. model CE. The 45-degree line is shown for reference in both panels.

Trade Cost Estimation: Risk vs. No Risk

Figure 14 compares the estimated change in median bilateral trade costs under risk aversion ($\gamma = 2$) and under risk neutrality ($\gamma = 1$). Under $\gamma = 1$, the model attributes all variation in expenditure shares to trade costs alone, yielding a roughly three log-point decline in τ^F and essentially no change in τ^I by 2014. Under $\gamma = 2$, part of the reallocation in expenditure shares is rationalized as a response to evolving risk, so a steeper trade cost reduction is required to match the data: about five log points for τ^F and three log points for τ^I.

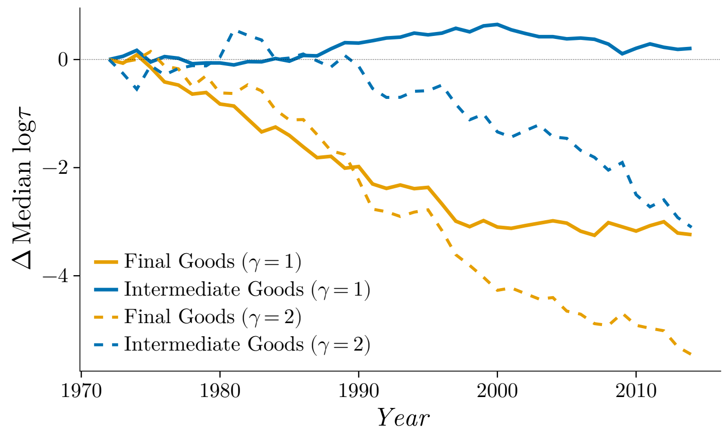

Figure 14: Estimated Trade Cost Changes: $\gamma = 1$ vs. $\gamma = 2$

Notes: Change in median bilateral log τ relative to 1967. Solid lines: $\gamma = 1$ (risk neutral). Dashed lines: $\gamma = 2$.

Figures 15a and 15b show the bilateral heterogeneity in trade cost changes from 1967 to 2014 for final and intermediate goods respectively.

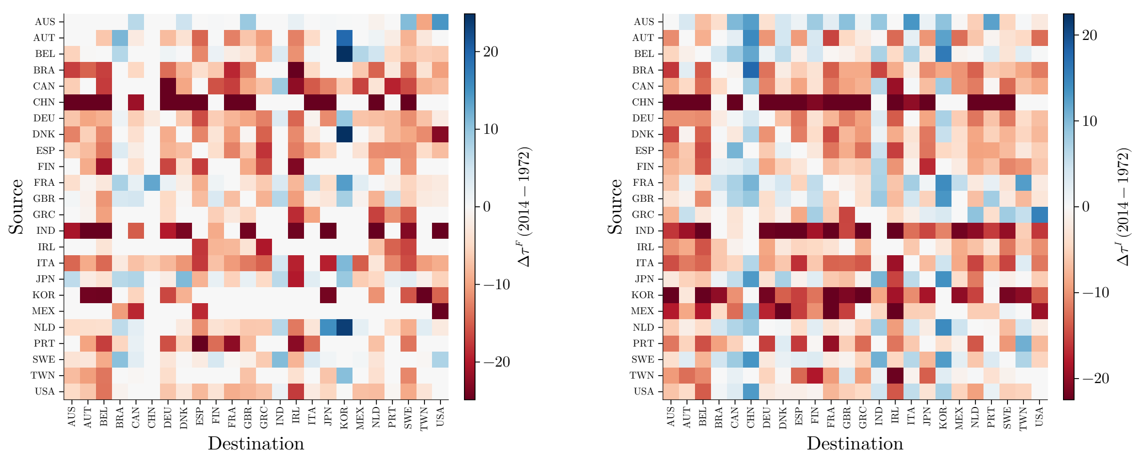

(a) Final goods

(b) Intermediate goods

Figure 15: Change in Bilateral Trade Costs: 1977–2014

Notes: Each cell reports the change in log $\hat{\tau}$ from 1967 to 2014 for a given bilateral pair.

Heterogeneous Homotopy

Figures 16 trace the change in GDP variance along the homotopy path from τ_{2014} to τ_{1977} for eight selected countries.

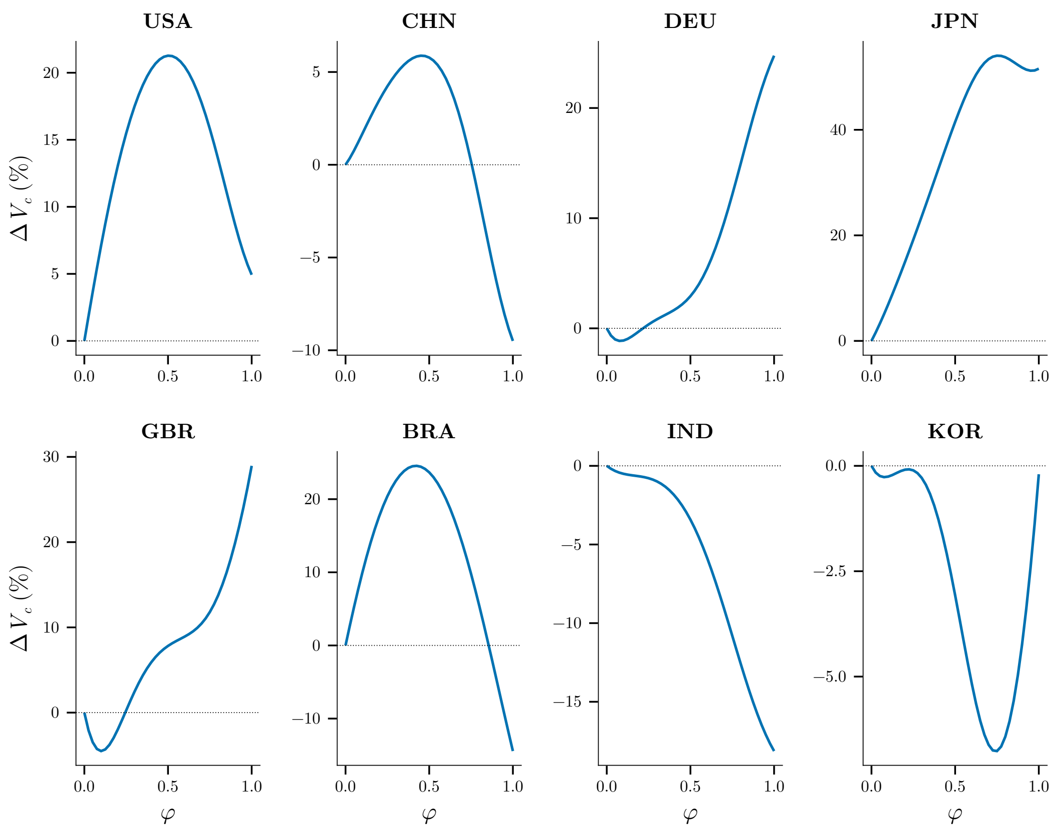

Figure 16: Country-Level Variance Paths: τ_{2014} to τ_{1977}

Notes: Each panel shows $\Delta V_c(\%) = 100 \times (\mathrm{Var}(\mathrm{GDP}_c)_\varphi - \mathrm{Var}(\mathrm{GDP}_c)_{\varphi=0})/\mathrm{Var}(\mathrm{GDP}_c)_{\varphi=0}$ along the homotopy from τ_{2014} ($\varphi = 0$) to τ_{1977} ($\varphi = 1$).

The paths are heterogeneous. For Germany, the United Kingdom, and Japan, raising trade costs increases variance nearly monotonically over most of the homotopy, ending well above the 2014 baseline (+25% for DEU, +29% for GBR and +50% for JPN at $\varphi = 1$): these countries are in the diversification-dominated portion of the variance-openness curve at 2014 trade costs, so increases in barriers steadily erode their diversification benefits. For the United States, the relationship is non-monotone but ends positive: variance rises sharply with moderate trade-cost increases (peaking near +21% around $\varphi = 0.5$) and then partially recedes as the fragmentation effect kicks in at very high barriers, ending at +5% above the 2014 baseline. For China, Brazil, and Korea, the path is hump-shaped and ends below baseline: variance first rises with moderate trade cost increases (peaks of roughly +6% for CHN and +25% for BRA, with KOR essentially flat near zero before turning down), but at high enough barriers variance falls below the 2014 level (−9% for CHN, −14% for BRA, and slightly below zero for KOR), meaning that lower trade costs in 2014 raised their variance relative to the 1977 counterfactual. For India, the path is nearly monotonic decreasing, steadily falling as the barriers rise to about −17% below the baseline at $\varphi = 1$.

Robustness: 2007 vs. 2014 Endpoint

Figure 17 compares the change in GDP variance from the historical liberalization counterfactual evaluated at two endpoints: $\tau_{2014} \rightarrow \tau_{1977}$ and $\tau_{2007} \rightarrow \tau_{1977}$.

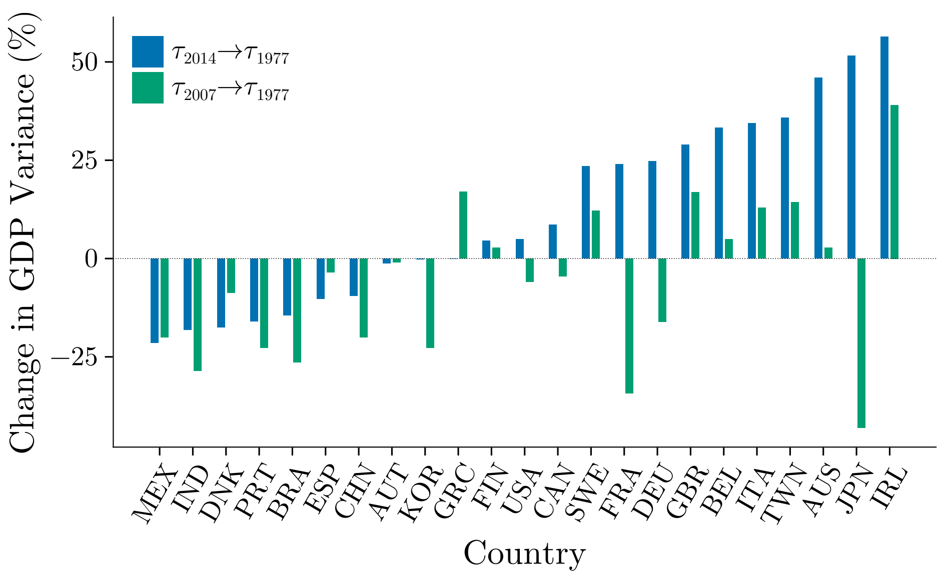

Figure 17: Change in GDP Variance: 2007 vs. 2014 Endpoint

Notes: Each bar shows $\Delta\text{Var} = (\text{Var}(\tau_{1977}) - \text{Var}(\tau_T))/\text{Var}(\tau_T) \times 100$, with $T \in \{2007, 2014\}$. Positive: reverting to 1977 raises variance (liberalization reduced it). Negative: reverting to 1977 lowers variance (liberalization raised it).

At the 2007 endpoint, the centrality channel is already active for the majority economies: 14 of 23 countries display $\Delta\text{Var} < 0$, which means that liberalization between 1977 and 2007 increased their aggregate variance. The greatest variance increases are concentrated in JPN (−41%), FRA (−33%), IND (−28%), BRA (−27%), KOR (−23%), CHN (−22%) and PRT (−22%). Only a handful of countries show net diversification gains at the 2007 endpoint, led by IRL (+39%), GRC (+17%), GBR (+17%) and TWN (+14%); the cross-country average is mildly negative. Extending the endpoint to 2014 reverses this pattern for many countries: the additional liberalization between 2007 and 2014 pushes the European hubs (FRA, DEU, GBR, BEL, ITA), Japan and several Asian-Pacific economies firmly into the region reducing variance, so that 14 of 23 countries now display $\Delta\text{Var} > 0$. The difference between the two endpoints reflects changes in trade cost between 2007 and 2014, not changes in the uncertainty environment (Figure 18).

To determine whether the difference between the 2007 and 2014 exercises reflects changes in trade costs or changes in the uncertainty environment, Figure 18 adds a hybrid counterfactual: $\tau_{2007} \longrightarrow \tau_{1977}$ evaluated on the post-crisis covariance matrix Σ_{2014}. The hybrid closely tracks the 2007 baseline rather than the 2014 baseline. Swapping the uncertainty environment has little effect on the variance pattern; the difference between the two endpoints is driven by the trade cost changes between 2007 and 2014, not by the elevated post-crisis shock variances.

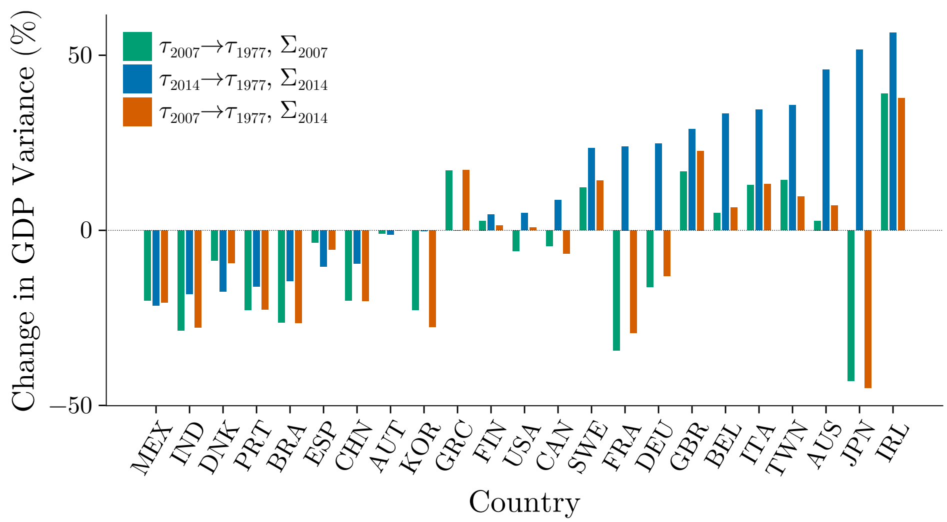

Figure 18: Hybrid Counterfactual: Trade Costs vs. Uncertainty

Notes: Three counterfactuals are shown. *Baseline 2007*: $\tau_{2007} \longrightarrow \tau_{1977}$ with Σ_{2007}. *Baseline 2014*: $\tau_{2014} \longrightarrow \tau_{1977}$ with Σ_{2014}. *Hybrid*: $\tau_{2007} \longrightarrow \tau_{1977}$ with Σ_{2014} (2007 trade costs, 2014 uncertainty). The hybrid tracks the 2007 baseline, indicating that the difference between the two endpoints is driven by trade costs, not by the shock covariance matrix.

Figure 19 traces network concentration as trade costs are progressively raised from 2014 toward their 1977 levels. After a small initial dip near $\varphi \approx 0.2$, HHI rises with φ, reaching about 0.0160 at $\varphi = 1$. Combined with Figure 7, in which HHI also rises as trade costs are lowered from 2014 toward free trade, the two homotopies trace out a shallow U-shape in trade costs with the minimum near the 2014 baseline. Concentration rises in both directions: at high barriers, the network fragments into regional clusters in which firms concentrate sourcing on a small set of accessible suppliers; at low barriers, it reconcentrates around globally central, productive suppliers.

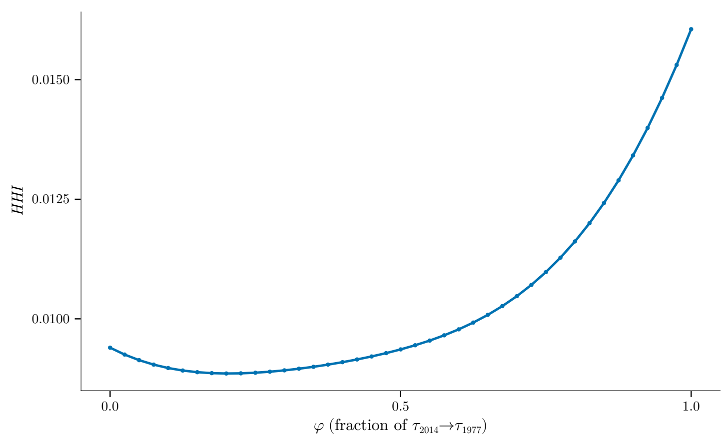

Figure 19: Network Concentration: Homotopy from τ_{2014} to τ_{1977}

Notes: φ is the fraction of the path from τ_{2014} ($\varphi = 0$) to τ_{1977} ($\varphi = 1$). HHI of Domar weights at each step.

Variance Effects Across Endpoint Years

Figure 20 trace the cross-country distribution of the variance effect as the counterfactual endpoint year varies from 1972 to 2014. For each endpoint year T, the figure plots $\Delta\text{Var}(\text{GDP}_c) = \text{Var}(\tau_{1977}) - \text{Var}(\tau_T)$: a positive value means that reverting to 1977 trade costs raises variance, so liberalization *lowered* that country's variance. From 1975 to the mid-1990s, the cross-country mean and median hover near zero (between roughly −5% and +2%), indicating that early in the sample liberalization had small and roughly offset effects across countries. Around 1995, both lines shift sharply upward: from then on, the mean fluctuates between +5% and +19% and the median between 0% and +8%, reflecting growing diversification gains for the typical country. The interquartile range also widens after 1995, with the 75th percentile rising above +25% in some years, while the 25th percentile dips below −20%, indicating substantial cross-country heterogeneity. The heterogeneous pattern in the main text is therefore not idiosyncratic to the 2014 endpoint, but a feature of the post-1995 liberalization episode.

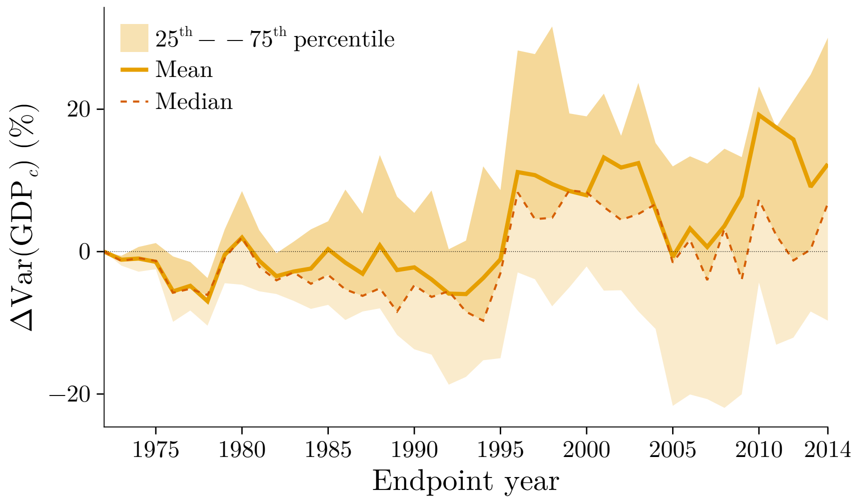

Figure 20: Variance Effect of Trade Liberalization Across Endpoint Years

Notes: For each endpoint year T, $\Delta\text{Var}(\text{GDP}_c)(\%) = \text{Var}(\tau_{1977}) - \text{Var}(\tau_T)$. Negative: reverting to 1977 lowers variance (liberalization raised it). Positive: reverting raises variance (liberalization lowered it). Solid: mean. Dashed: median. Shaded: 25th–75th percentile.

Sensitivity Analysis

This section examines the sensitivity of the main quantitative results to the main parameters ψ (Frisch elasticity), γ (risk aversion) and κ (adjustment cost). In each exercise, one parameter varies while the others are held at their baseline values; trade costs τ and shock variances Σ are re-estimated at each parameter value. The dashed vertical line marks the baseline calibration.

Figure 21 plots the average change in GDP variance from the 1967–2014 trade liberalization counterfactual as a function of each parameter.

Adjustment cost κ (Panel a). This is the parameter to which the results are most sensitive. At baseline $\kappa = 5$, the average variance change is around 26.5%, rising slightly to a peak of about 27.5% near $\kappa \approx 5.5$. From there the variance effect declines steadily as the network becomes more rigid, falling to roughly 20% at $\kappa = 8.5$, 18% at $\kappa = 9$, 15% at $\kappa = 10$, and about 10.5% at $\kappa = 12$. The decline is approximately linear over $\kappa \in [6, 12]$ rather than convex.

The direction is intuitive: high κ anchors firms near their ideal network m^0, preventing the endogenous concentration that drives the variance result.

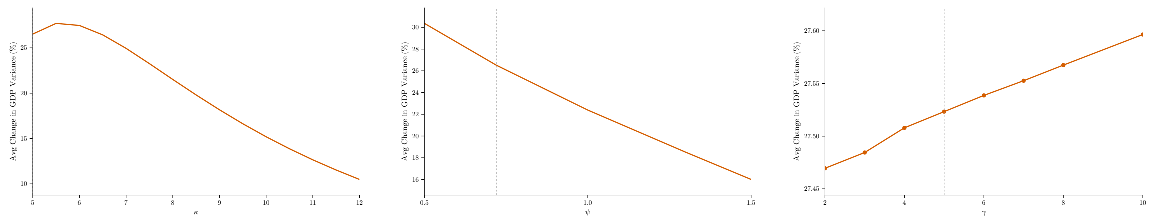

(a) vs. κ (b) vs. ψ (c) vs. γ

Figure 21: Sensitivity of Variance Counterfactual to Parameters

Notes: Average change in $\text{Var}(\text{GDP}_c)$ from the 1967–2014 counterfactual. Each point re-estimates the model at a different parameter value, holding the other two at baseline. Dashed vertical line: baseline value.

Frisch elasticity ψ (Panel b); A higher elasticity of the labor supply dampens the effect of variance. At $\psi = 0.5$ the mean variance change is roughly 31%; at the beginning of $\psi = 0.72$ it is 28%; at $\psi = 1.5$ it falls to 18%. The mechanism operates through the re-estimation of shock variances: since Σ is re-estimated at each ψ, a larger Leontief inverse (from higher $\psi/(1+\psi)$) requires a geometrically smaller primitive shock variance σ_g^2 to match the same empirical data variance. Therefore, the effective risk premium $(\gamma-1)(\nu_{\hat{c}k}-\bar{\nu})\nu_c, \sigma_g^2$ declines, dampening the reallocation mechanism and the variance effect of trade liberalization. The relationship is approximately linear in ψ.

Risk aversion γ (Panel c). The change in the average variance is nearly insensitive to γ, varying by less than half a percentage point in $\gamma \in [2, 10]$. This is because γ governs how much firms tend to insurance (low-covariance suppliers) versus low prices, but the aggregate variance effect is dominated by the network structure (Leontief amplification) rather than the composition of individual portfolios. The result is reassuring: the main quantitative finding does not depend on the precise value of risk aversion.