Precautionary Sourcing and the Macroeconomics of De-Risking *

Ester Faia
Goethe University Frankfurt and CEPR

First draft: June 2026

Do no circulate without author's permission

Abstract

This paper assesses the efficiency and welfare costs of retrenchment from international suppliers due to tariff and geopolitical risk. Using a sector-level (Census data) and firm-level (Panjiva firm input-output data) exposure-weighted event-study analysis, I find marked retrenchment and reallocation that went beyond tariffed goods and, through an anticipation design using geopolitical-risk and trade-policy-uncertainty indices, can be attributed to heightened uncertainty rather than tariff levels. I build an input-output network model in which firms, facing *uninsurable productivity and currency risk*, optimally choose suppliers (severing very risky ones, albeit productive) and embed it in a dynamic macro model, where consumer-workers share the risk of the firms to which they are attached through a limited-commitment wage contract. I derive a risk-misallocation wedge that emerges only under an endogenous network and incomplete markets: it vanishes in the fixed-network benchmark (Baqaee and Farhi) or under complete markets. I simulate the model under a counterfactual tariff shock (to mimic the 2018 and 2025 events) and validate it against the empirical patterns. The efficiency cost of precautionary severing of the Chinese links is modest (0.065% of consumption) but scales with exposure to the full set of tariffed suppliers (0.31%); within the risk wedge the currency-uncertainty component is larger than the tariff-uncertainty component; and the consumption cost borne by workers sharing their firm's risk exceeds the pure misallocation cost by roughly three orders of magnitude and falls overwhelmingly on the poorest households, who bear some twenty-seven times the cost of the richest.

Keywords: *endogenous international production network, uninsurable risk, risk-misallocation, modified Hulten's theorem, distributional effects of tariffs, de-risking, re-trenchment, geoeconomics shocks.*

JEL Classification: *F41, F13, F44, D57, D52, D24.*

1 Introduction

This paper assesses the efficiency and welfare costs of retrenchment from foreign suppliers under rising geopolitical and trade risk. Using an exposure-weighted event study around the 2018

*Firm level input-output data from Panjiva and sales data from Compustat were accessed through the visiting scholar subscription at Harvard University. I thank Itay Saporta-Eksten, Oleg Itskhoki, Camilo Marchesini, Tommaso Monacelli and Esteban Rossi-Hansberg for useful conversations. Correspondence to: Ester Faia, Chair in Monetary and Fiscal Policy, Goethe University Frankfurt, Theodor W. Adorno Platz 3, 60323, Frankfurt am Main, Germany. E-mail: faia@wiwi.uni-frankfurt.de. Webpage: http://www.wiwi.uni-frankfurt.de/profs/faia/.

and 2025 tariff episodes, estimated at the sector level (using Census and Comtrade data) and the firm level (using input-output Panjiva data), I find evidence of significant retrenchment from suppliers subject to tariffs that goes beyond the goods subject to the Section 301 tariffs. Through an anticipation design, I find that this broad retrenchment from a sourcing country subject to tariffs is primarily driven by geopolitical and trade risk. I then build a model in which firms, facing uninsurable productivity and currency risk, optimally choose (or sever) their suppliers (domestic and foreign), with links that balance the productivity benefits of the sourcing organization against the risk of the nodes. I embed this production network in a dynamic macro model in which consumers-workers are exposed to the risk of the firms to which they are attached through a limited-commitment wage and self-insure. I derive analytically a risk-misallocation wedge that emerges only under an endogenous network and incomplete markets together: it vanishes in the fixed-network benchmark (Baqaee and Farhi [2019]) or under complete markets (the endogenous network in Kopytov et al. [2024]). The model's simulations under a counterfactual tariff shock (whose level and risk mimic the ones in 2018 and in 2025) show that: the efficiency cost of precautionary severing of the Chinese links is modest, about 0.065 percent of consumption, but scales with exposure to the full set of tariffed suppliers, to about 0.31 percent; within the risk wedge the currency-uncertainty component is larger than the tariff-uncertainty component; and the consumption cost borne by workers sharing their firm's risk exceeds the pure misallocation cost by roughly three orders of magnitude and falls overwhelmingly on the poorest households, who bear some twenty-seven times the cost of the richest.

We establish the facts with event studies around the 2018 and 2025 shocks, at the sector level from United States Census data and at the firm level from Panjiva importer input-output records. Sorting firms by prior China exposure, the retreat is proportional to that exposure and shifts toward allies, a pattern corroborated by sector-level concentration and a parallel contraction of sales-side linkages. Two features identify its nature. The retreat extends to untaxed goods as much as to taxed ones, so it is not a mechanical response to the tax. In an anticipation design, it moves with trade-policy uncertainty [Caldara et al., 2020] and geopolitical risk [Caldara and Iacoviello, 2022], with monetary-policy uncertainty as a placebo [Baker et al., 2016], so it is the uncertainty around the policy, not the tariff level, that drives it. The retreat is thus a broad, precautionary withdrawal from a risky origin, which is exactly the ex ante retrenchment the model formalizes.

The mechanism is precaution under commitment. Supply relationships take time to build, so a firm commits to its network before the period's productivity and exchange-rate shocks are realized and chooses input intensities only afterward. Under complete markets it discounts

states of the world with a stochastic discount factor common to everyone, and its network is efficient. Under incomplete markets, the firm, or the household that owns it and bears its risk, discounts with its own kernel, one that loads on its own risk, and so values a supplier not only for expected productivity but for how that supplier's price moves with the firm's own marginal utility, retrenching from suppliers that are productive but risky. Because that risk is diversifiable for the economy though not for the undiversified owner, the retreat is a misallocation relative to what a diversified owner would choose.

The model is an input-output production network (with domestic and foreign links) in which firms, facing uninsurable and productivity risk, choose (or sever) the links by optimally trading off the productivity benefits of the sourcing organization with the risk. Uninsurable risk leads to precautionary behavior: firms choose to sever links that are efficient but risky. In this model, I derive a risk-misallocation wedge, whose nature I discuss below, which vanishes in the fixed-network or complete-insurance benchmarks. I then embed this network in a dynamic macroeconomy in which the consumer-workers attached to each firm bear its risk through a limited-commitment wage contract: the contract insures them in normal states but, when the firm's participation constraint binds in bad states, passes more of the firm's shock into their wages. This microfounds the firm's pricing kernel: the degree of uninsurability that governs the wedge is an equilibrium object rather than a parameter. This contracting microfoundation also delivers a result of independent interest: the degree of uninsurability, and the earning process, is endogenously countercyclical, a fact well established empirically [Storesletten et al., 2004, Guvenen et al., 2014], but is usually imposed by assumption in incomplete market models; here it emerges from limited commitment in the tradition of Thomas and Worrall [1988] and Kocherlakota [1996]. The misallocation wedge inherits the pattern, near zero in normal states and large in downturns, when misallocation is most costly. Finally, in the model, the exchange rate dynamic is governed by the international risk-sharing condition of Backus and Smith [1993], so home bias in sourcing emerges as an inefficiency rather than a preference, and the tariff becomes a priced aggregate risk correlated with the real exchange rate.

Our central contribution is to isolate this distortion as the interaction of two departures from the frictionless benchmark, neither of which distorts on its own. With a fixed network, the discount factor cannot reshape an allocation that is not chosen, so incomplete markets alone are irrelevant; the benchmark of Baqaee and Farhi [2020]; with an endogenous network but complete markets, the firm reorganizes efficiently against the common kernel, the benchmark of Kopytov et al. [2024]. Put simply, a firm that could fully insure, or that faced complete markets, would choose its links on efficiency grounds alone; when it cannot hedge, it sources with precaution, and that precaution leads to misallocation. The latter needs both: the firm must be free to choose

its links, and unable to insure the risk they carry. Formally, the wedge is the interaction of the two. It takes a transparent form: when the kernel loads on the firm's own shock with intensity ϕ_i, the wedge on the link from supplier j to firm i is $\tau_{ij} = \phi_i[\mathcal{L}\Sigma]_{ji}$, the degree of uninsurability times the own risk times the supplier's network exposure through the Leontief inverse, and it aggregates into an economy-wide productivity loss by a network-corrected Hulten theorem. When the supplier is foreign, the same kernel prices the exchange-rate risk in the input's price, so the wedge carries a currency term alongside the productivity one, and the same mechanism distorts foreign sourcing.

Calibrated to United States input-output and productivity data and to the tariff schedule of the current conflict, and simulated under a counterfactual of a tariff shock in both its level and its variance (the latter proxying the heightened uncertainty of recent tensions), it reproduces the de-risking observed in the data. The level drives the reallocation of sourcing, but only modestly and, as in the data, largely off firms' prior exposure; the precautionary channel is where the action is. However, since measured productivity risk is small and China is low-beta (Chinese sourcing productivity is negatively correlated to US one in the data), only the tariff, and through it the dynamics of the dollar, supply a risk large and concentrated enough to matter. Indeed, although the reallocation triggered by the tariff level is modest, through the lens of the model the efficiency and welfare costs of uncertainty can be high. The efficiency cost of risk-driven misallocation in the firms' network is not the major component: it is modest on the China links, about 0.065 percent of consumption, but scales with exposure to roughly a third of a percent once all tariffed origins are included, and within the risk wedge the currency component exceeds the tariff component. The uninsurable consumption cost component is the largest. It is steeply regressive, running through the marginal propensity to consume, so that the poorest, liquidity-constrained households, who pass the shock almost fully into consumption, bear some twenty-seven times the cost of the richest.

The rest of the paper proceeds as follows. Section 2 presents sector and firm-level event-study evidence on the tariff episodes of 2018 and 2025. Section 3 develops the baseline endogenous input-output network model with firms facing uninsurable risk. Section 4 extends the model to a general equilibrium of the full economy, in which the firm's pricing kernel is given by the SDF of households who bear the firm's risk through the rent sharing implied by a limited commitment contract. Section 5 presents the calibration and the quantitative analysis of a counterfactual tariff event.

Related Literature. This paper sits at the intersection of five literatures, against which we compare its contribution. The paper builds first and foremost on the network-misallocation literature of Baqaee and Farhi [2019, 2020]. Baqaee and Farhi [2020] characterize how distor-

tions, summarized as wedges between prices and marginal costs, map into aggregate total factor productivity in general equilibrium, and Baqaee and Farhi [2019] give the nonlinear propagation of microeconomic shocks beyond Hulten [1978]. Our wedge lives in the same aggregation; the main difference is the wedge itself. In that literature, wedges are primitives: markups, taxes, or other exogenous frictions. In my model, the wedge is endogenous as it is the equilibrium price of uninsurable risk borne by a firm that commits its network before the shock. Furthermore, I examine misallocation costs in both a firm-level model and a general equilibrium model of the full economy.

The endogenous-network margin has been studied by Kopytov et al. [2024], and more broadly by Acemoglu and Azar [2020] and Oberfield [2018], in which firms choose suppliers and the network responds to shocks and policies. In these models, markets are complete, so the reorganization is efficient and the discount factor is allocationally irrelevant. My model departs by assuming incomplete markets with uninsurable risk: once the firm prices its committed network with its own kernel (conditional on the uninsurable risk realization) rather than a common one, the same endogenous reorganization becomes distortionary. We nest Kopytov et al. [2024] exactly as the complete-market benchmark: the firm-specific kernel that turns efficient reorganization into misallocation is the primary novelty. The paper also connects to a literature on empirically causally identified estimates of network diffusion (see Faia et al. [2026] for information diffusion, Alfaro et al. [2022] for preference diffusion in an endogenous network, Bilgin et al. [2024] for diffusion of technology adoption).

A parallel literature measures misallocation while accounting for general-equilibrium adjustment. Sraer and Thesmar [2023] show how to recover the aggregate, general-equilibrium effect of a firm-level policy from the distribution of treatment effects on output-to-capital ratios, without a full structural estimation. Our analysis shares the concern that partial-equilibrium wedges do not aggregate one-for-one. Still, we take a structural approach that embeds general-equilibrium effects through both network propagation and propagation in the full economy. The heterogeneous-firm tradition of Hopenhayn [1992] and Hopenhayn and Rogerson [1993] also examines the aggregate effects from firm-level decisions. Still, the propagation goes through the entry and exit margin, as well as through factor reallocation. Our primary propagation channel across firms is through a risk-misallocation wedge: a firm's uninsured precaution reshapes its sourcing, and the reshaping transmits through input-output linkages into the measured productivity of others.

This paper also relates to the trade literature that examines the welfare costs of globalization or de-globalization. The welfare cost of trade barriers is usually measured through their effect on real income, as in sufficient-statistic formulas that measure a tariff's cost from the import share

and the trade elasticity [Arkolakis et al., 2012]. That accounting is static and deterministic, and welfare refers to average consumption, without any risk component. A small set of papers examines trade models under uncertainty and views trade and tariffs as insurance: Helpman and Razin [1978] and di Giovanni and Levchenko [2012] show that sourcing and export patterns reflect insurance motives alongside comparative advantage, while Newbery and Stiglitz [1984] and Eaton and Grossman [1985] show that under incomplete markets free trade can be Pareto inferior and a tariff can insure. In that tradition, trade insures against domestic shocks. At the same time, my study reverses that logic, showing that a domestic policy (a tariff change) and its uncertainty sever external links that had supplied both efficiency and diversification against domestic productivity shocks. Furthermore, the propagation considered in my model is through the input-output network.

Finally, a fast-growing literature, inspired by the recent wave of US tariffs, studies various aspects of tariffs and supply. Here is a brief review focusing on the studies most closely related to my work. Baqaee and Farhi [2024] treat tariffs as wedge-like distortions in an international production network and quantify their societal cost with sufficient statistics; Liu [2019] studies policy in a distorted network; Gaubert and Itskhoki [2021] propose an import tariffs on large exporters to curb their power and Itskhoki and Mukhin [2025] show that the optimal tariff is lower for highly indebted governments; Aguiar et al. [2025] examines whether and how international financial claims can be honored without default when an intense trade war erupts; and a recent wave adds nominal rigidities and monetary policy [Antonova et al., 2025, Kalemli-Özcan et al., 2025, Monacelli, 2025, Guerrieri et al., 2025, Auclert et al., 2025] among others.[1] Closest in motive is the supply-chain-resilience literature: Grossman et al. [2024] study a general-equilibrium model with endogenous formation of supply links and endogenous investments in resilience, and characterize the subsidies that implement the first best. That work shares our concern with derisking the network. Still, the inefficiency there is a technological externality across firms under complete information, corrected by a Pigouvian subsidy on links and resilience. In contrast, ours is the private mispricing of risk under incomplete markets, which a subsidy on the link alone does not correct.

2 Sector and Firm Level Evidence

To motivate our model and to validate its broad patterns, we start by documenting the reallocation of sourcing origins following a discrete geoeconomic shock. The focus is on the tariff changes enacted recently by the US, namely the 2018 Section 301 tariffs and the 2025 escalation. Specif-

[1] None of those works features an input-output network except Kalemli-Özcan et al. [2025], who however does not feature uninsurable risk. Furthermore, Auclert et al. [2025] features uninsurable risk on households but has homogeneous firms and no network.

ically, I first use U.S. Census and Comtrade data and examine sector-level evidence through a diff-in-diff event study to identify which sectors U.S. firms retreated from and whether spillovers affected other parts of the economy. Next, I present firm-level evidence using high-quality data on business-to-business international transactions from Panjiva (a firm-level input-output network); specifically, we conduct an exposure-weighted event study analysis. Third, I present an anticipation design that identifies the level versus the uncertainty effects of the tariff events by matching Panjiva business-to-business data with uncertainty indicators (Caldara and Iacoviello [2022] and Baker et al. [2016]) Appendix B provides a detailed description of all the micro and sector data used (Census and Comtrade sector data, Panjiva firm-level input-output data, Compustat customer segment data) as well as the uncertainty indicators.

2.1 Sector-level Sourcing: the Retreat and its Concentration

I use quarterly United States goods imports by origin country and product at the four-digit level of the Harmonized System, from the third quarter of 2015 through the second quarter of 2020, from ITC Trade Map, which republishes United States Census Bureau and United Nations Comtrade statistics. I access them through the Deutsche Welle data team replication package [Deutsche Welle Data Team, 2020]. I flag each product subject to the United States Section 301 tariff list (Lists 1 through 4A) and aggregate to eight sectors (the choice maps into the calibration of the model presented below). The construction is described in Appendix B.2. For each quarter and product basket, I compute China's share of United States imports, Chinese import value over total United States import value in the basket.

I examine the impact of the tariff using a diff-in-diff event study. Let s_{pt} denote China's share of United States imports of product p in quarter t, let Tariffed_p indicate placement on a Section 301 list, and let Post_t indicate quarters from the third quarter of 2018 onward. The econometric specification reads as follows:

$$s_{pt} = \alpha_p + \delta_t + \beta \left(\text{Tariffed}_p \times \text{Post}_t\right) + \varepsilon_{pt}, \tag{1}$$

where α_p are product fixed effects and δ_t quarter effects. Figure 1 shows the results. China's share of total United States goods imports falls from about twenty-four percent in 2017 to about nineteen percent by the end of 2019, and the decline sets in around the 2018 tariff onset rather than earlier. It is sharply concentrated where United States buyers were most exposed: China's share falls by about six percentage points in electronics, from forty-three to thirty-seven percent, and by about six in machinery, from twenty-nine to twenty-three percent, while it barely moves in low-exposure sectors such as pharmaceuticals, metals, and motor vehicles. Sourcing shares shift to a small set of alternative origins: within the tariffed basket, China's share falls by about three and a half points between 2017 and 2019, while Mexico gains about one and a quarter

and Vietnam about eight tenths, with Taiwan, Malaysia, and Canada also gaining. This is **friend-shoring and near-shoring**, which the model will later rationalize as a retreat from riskier sourcing origins to safer ones. The shift clearly signals that a model that accounts for the observed pattern should feature an endogenous choice of links, which is absent from most input-output network models.

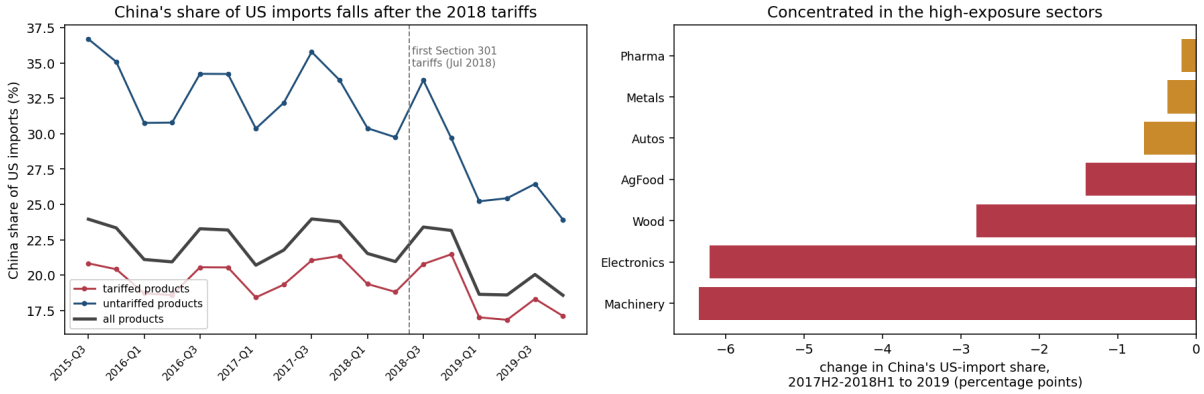


Figure 1

Sector-level de-risking around the 2018 Section 301 tariffs.

Notes: Left panel: China's share of United States imports over time, for all products and for the tariffed and untariffed baskets separately, quarterly, with the pandemic-distorted quarters of 2020 omitted. Right panel: the change in China's share of United States imports by sector, from the year before 2019. Source: ITC Trade Map, republishing United States Census Bureau and United Nations Comtrade statistics; tariff-list assignments from the Office of the United States Trade Representative.

The sector data also reveal a spillover to other sectors. The retreat is not confined to the tariffed lines: the left panel of Figure 1 shows China's share falling in untariffed products as well. The spillovers suggest two additional transmission channels: a broad precautionary move away from a tariffed country and general equilibrium effects. Both are absent from most input-output network models, including endogenous network models, and from full-economy models that include them. In our model, firms will be heterogeneous due to uninsurable risk: the equilibrium of those models features general equilibrium spillover effects stemming from price adjustments in the market-clearing conditions. Those effects are absent in macro models with networks but complete markets or representative firms. Given this, the next section also presents evidence at the firm level.

2.2 Firm-Level Input-Output Event Study Analysis

The sector retreat has a firm-level counterpart in import transactions. The data are United States waterborne import records from Panjiva, at the individual-shipment level, naming the United States consignee, the foreign shipper, the country of origin, the arrival date, and the customs value and weight. We aggregate to the importer and year and construct, for each importer, its China import share, its share sourced from a set of alternative origins (Vietnam, Mexico, Taiwan, India, Thailand), and its count of distinct Chinese suppliers, in two panels,

one around each shock. The variables and their Panjiva field names are in Appendix B.1. The design is a within-importer event study estimated by two-way fixed effects, on the importer and year dimensions:

$$y_{it} = \alpha_i + \delta_t + \sum_{\tau \neq \bar{t}} \beta_\tau \left(\mathbf{1}[t = \tau] \times \text{Exposure}_i \right) + \varepsilon_{it}, \tag{2}$$

where α_i are importer fixed effects, the first of the two ways, absorbing any time-invariant difference across importers, δ_t are year fixed effects, the second, absorbing any common shock to all importers in a given year, Exposure_i is the importer's pre-shock China share, and $\bar{t}$ is the last pre-shock year, which is omitted. The coefficients β_τ are event-study estimates that trace the outcome of high-exposure relative to low-exposure importers year by year, and importer clusters standard errors.

The heterogeneous-timing critiques of Sun and Abraham [2021] and de Chaisemartin and D'Haultfœuille [2020] do not apply here. Those critiques concern settings where units are treated at different dates, so that the two-way estimator uses already-treated units as controls for newly treated ones and can assign negative weights to some comparisons, contaminating the average. In our setting, the treatment date is common to every importer, so no already-treated units serve as controls; the forbidden comparisons never enter, and each β_τ is a clean comparison of the treated cohort against its own pre-shock baseline and the never-more-treated variation in exposure. The negative-weighting problem therefore cannot arise by construction, not merely as an approximation.

Figure 2 reports the patterns of the event study for both the 2018 and the 2025 tariff episodes for a four-year window. Specifically, Figure 2 plots the event-study coefficients against the base year, so the pre-shock leads are the placebo coefficients, and in the short windows around each shock they sit close to zero. Furthermore, more-exposed importers reduce their China share and raise their ally share after the tariff, and the number of distinct Chinese suppliers falls. This confirms the sector-level evidence presented earlier. Taken at face value, the differential cut in the China share is large, on the order of twenty to thirty log points relative to an unexposed importer.

Several qualifications are worth mentioning, since they ultimately strengthen the map with the model's results. In Appendix B.1, we report additional event-study tests that extend the event window and include several pre-trend tests and pre-trend adjustments. Overall, pre-trends appear to explain the event study results partly. Economically, this suggests that the **level** effect of the tariff's incidence itself on reallocation from China to ally sourcing is modest and that most of it may come from an anticipation effect. But the latter is driven by expectation and uncertainty, which leads firms to sever risky links in advance as a precaution. This is precisely

the channel the model is after, and we return to these points in a final section that compares model-simulated results to a tariff counterfactual with the empirical results. The uncertainty channel would also generate a broader retreat from the origin country, rather than a surgical retreat from the tariffed goods. The next section supports this hypothesis with an anticipation design.

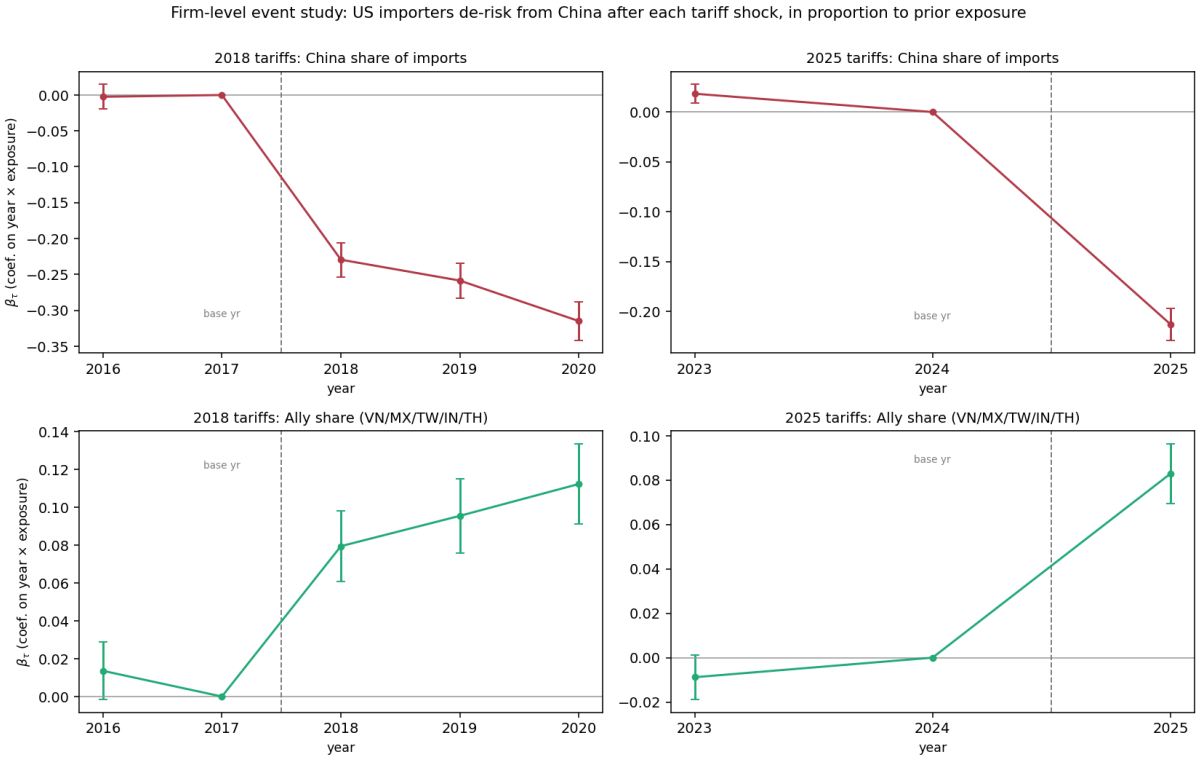


Figure 2

Firm-level event study of United States importers around the two tariff shocks, on the windows immediately around each shock.

Notes: Each panel plots the coefficients β_τ from (2) with 95 percent confidence intervals, clustered by importer, against the base year (2017 for the 2018 tariffs, 2024 for the 2025 tariffs). Left column, the 2018 tariffs; right column, the 2025 tariffs. Top row, the China share; bottom row, the ally share. We read this pattern as suggestive only: on extended windows, the pre-trend test rejects, and the break does not survive an exposure-specific trend, because firm-level China shares mean-revert (Appendix B.1). Source: Panjiva United States waterborne import records.

2.3 An Anticipation Design: the Breadth of the Retreat

The previous results suggest a significant retreat from tariffed countries, but this may partly reflect an anticipation effect. In this section, we devise an anticipation design to test how much of the retreat was due to tariff incidence, rather than a broad withdrawal from China as a risky origin in anticipation of policy and geopolitical uncertainty. The model makes a sharp prediction here. Because the network is committed before the shock, a firm re-sources when the expected risk of an origin rises, not when a particular product is taxed: model simulations show that the impact of tariff incidence is modest, but the impact of risk on precautionary severance is large.

We separate the two interpretations with a simple design that examines the impact on the

China import share in tariffed goods against the same share in untariffed goods, read alongside trade-policy uncertainty [Caldara et al., 2020] and China-specific geopolitical risk [Caldara and Iacoviello, 2022], with monetary-policy uncertainty as a placebo [Baker et al., 2016]. The logic is a discriminating restriction. Mechanical tariff incidence, the level channel, raises the price of taxed Chinese goods and predicts the retreat to concentrate in the tariffed basket, leaving untariffed goods untouched. A broad, precautionary withdrawal from China as an origin, the uncertainty channel the model describes, predicts the untariffed basket to retreat as well, since the firm is fleeing the origin's risk rather than the tax on any product. The untariffed China share is therefore the clean discriminator between a tax-incidence account and an origin-risk account. Formally, the design reads as follows:

$$s_{bt} = \gamma_b + \delta_t + \beta \left(\text{Tariffed}_b \times \text{Post}_t \right) + \varepsilon_{bt}, \tag{3}$$

Table 1 reports the China import share in each basket before the episode, 2015 to 2017, and after, in 2019. Both fall, but the untariffed share falls about three times as much as the tariffed share: 8.1 versus 2.7 percentage points, a 24 percent decline in untariffed China sourcing versus 13 percent in tariffed. This is the opposite of what tax incidence alone predicts: that the taxed basket should retreat more, and it is the signature of a broad, origin-level withdrawal from China. Figure 3 shows this retreat coincides with elevated trade-policy uncertainty and China geopolitical risk, which rise sharply from 2018, while the monetary-policy-uncertainty placebo tracks it less closely.

This design resolves the partly confounded analysis of the previous event study design, in which the tariff dummy may capture both the level incidence and the uncertainty effect, whose breadth spans taxed and untaxed products, as well as possible other pre-trends. The anticipation design resolves those issues since it is a within-period cross-product comparison rather than a timing comparison. A trend in a firm's Chinese sourcing that is common to its tariffed and untariffed goods moves both baskets together and does not open a gap between them, so the contrast differentiates it out, much as a control group absorbs common shocks. Only a force that treats the baskets differently can open the gap: a tax on the taxed lines, which would push the tariffed basket down more, or a broad withdrawal from China that happens to fall more on untaxed goods. The observed gap, with the untariffed basket retreating more, is inconsistent with the tax and points to the origin-level withdrawal. The one alternative the contrast does not rule out is that untariffed goods are a different product mix with reshoring dynamics of their own, which we cannot fully exclude with these data.

We are explicit about the limits of these data. The outcome is an aggregate quarterly series of eighteen observations, and the retreat is concentrated in a single year, 2019. Hence, the fall lags the 2018 uncertainty spike rather than clearly leading it at this frequency. Regressions of

Table 1

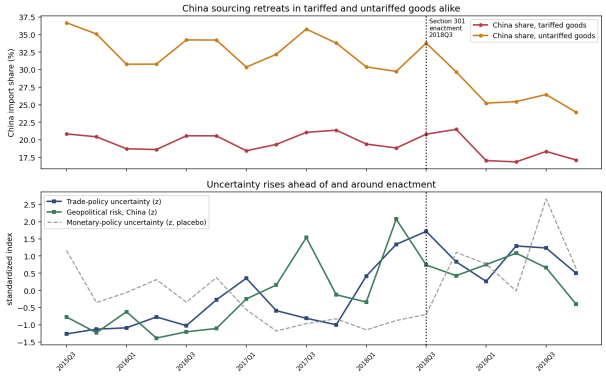
China import share.

China import share (%)	2015–2017	2019	change
Tariffed goods	20.0	17.3	−2.7
Untariffed goods	33.4	25.3	−8.1
All goods	22.6	19.0	−3.6

Notes: China import share by basket, before and after the 2018 episode. The untariffed basket retreats about three times as much as the tariffed one, the opposite of mechanical tariff incidence and the signature of a broad withdrawal from China as an origin.

Figure 3

Caption: China Shares by Tariff Flag and Geopolitical Uncertainty (Caldara and Iacoviello [2022])



Notes: The China sourcing retreat is broad and coincides with rising uncertainty. Top panel: the US China import share in tariffed and untariffed goods, quarterly, 2015Q3–2019Q4 (Panjiva). Both baskets are flat through 2018 and fall together in 2019, with the untariffed share falling further, the opposite of what a tax on the taxed lines alone would produce. Bottom panel: standardized trade-policy uncertainty [Caldara et al., 2020] and China-specific geopolitical risk [Caldara and Iacoviello, 2022], which rise from 2018, with monetary-policy uncertainty [Baker et al., 2016] shown as a placebo. The dotted line marks the first Section 301 enactment, 2018Q3. Indices are quarterly averages, standardized over the sample.

the quarterly share on the standardized uncertainty indices, with a linear trend, are imprecise in a sample this short. What the aggregate data do establish is the breadth of the retreat, the tariffed-untariffed contrast, which is robust and does not depend on the noisy time-series link to the indices. This is quite important for model selection.

The model's mechanism is that firms commit to a network ex ante and retreat from an origin whose expected risk rises, so the distortion is attached to the origin, not to the taxed product. The breadth of the observed retreat, untaxed Chinese goods withdrawing as much as taxed ones, is exactly this origin-level signature, and it is not what a model of tax incidence on committed but flexible quantities would produce. The evidence thus supports the premise the model builds on, a precautionary withdrawal from a risky origin, while leaving the welfare cost of that withdrawal, which no import-share design can measure, to the structural model.

2.4 Evidence on Sales from Compustat Data

The datasets above capture the buy side of the network, from which firms source. A complementary piece of evidence shows decoupling on the sales side as well. Using the customer-disclosure

data in Compustat[2] in which United States public firms report their major customers, we find that the number of reported relationships with Chinese customers falls from about ninety in 2017 to about fifty by 2021, a decline of roughly forty-five percent that begins after 2018. This evidence shows that US firms were also losing customers, signaling a broader precautionary retrenchment with spillovers across the board. Appendix B.3 describes the data and the count.

3 The Risk-Misallocation Wedge in an Endogenous Network Model with Uninsurable Risk

In this section, I lay out the baseline structure of the network model, in which firms that face uninsurable productivity risk and are therefore heterogeneous in their policy functions, choose the links to their suppliers and customers in anticipation of the future productivity realization. In doing so, they respond to uninsurable risk through precautionary behavior, which is reflected in the endogenous structure of the input-output relations. In the sections that follow, we extend this model by including capital accumulation and embedding it in a general equilibrium economy (see Section 4) where firms also face currency risk.

Given the presence of uninsurable risk, it is useful to clarify the timing of events. Firms choose their input-output relations, and therefore commit to a network, before the realization of the productivity shock and choose input quantities after it. Given the absence of a market for claims that insure across states, and given that the pre-committed network structure fixes how the firm's cost loads on the stochastic prices of its inputs, the network choice carries risk. In every period, when choosing or re-assessing their network links, firms do so in anticipation of this risk. In this set-up, the firm's own stochastic discount factor prices risk. The model has general equilibrium effects in that firms' policy functions depend on the input prices that clear the market, and the latter depend on individual firms' demand. However, we take the investors' pricing kernel exogenously; it will be microfounded in the full general equilibrium model considered in section 4.

In what follows, we set out the two-stage firm problem (input-output shares first, then input quantities), its optimality conditions, and derive analytically the fixed-point equilibrium that yields the price system clearing input markets.[3] Given the equilibrium solution and the resulting distributions, we characterize the aggregate wedge, the precautionary retrenchment it induces, and the productivity loss. To provide a benchmark comparison, I conclude by showing how this network nests the complete-market and fixed-network benchmarks, two cases extensively

[2]This was accessed through the subscription of Goethe University Frankfurt.

[3]The iterative process implies that, when choosing inputs, firms take the input clearing prices as given; the aggregation of their input demand then determines the clearing process until convergence. We show that under a constant-return-to-scale production function, this process has a simple closed-form solution.

studied in the literature, and I characterize the departures from those cases.

3.1 Environment

From now on, lowercase letters indicate log variables.

Sectors, firms, and goods. There are n sectors indexed by $i, j \in \{1, \ldots, n\}$. Each sector contains a continuum of monopolistically competitive firms producing differentiated varieties, which a constant-elasticity aggregator with elasticity $\sigma > 1$ combines into the sector good. The sector good is the node of the production network: what other sectors buy as an input and what households consume. In other words, firms within a sector share the same technology and make the same input-share choices. We return to the distinction between network- and firm-level decisions below. States of nature are $s \in S$ with probability $\pi(s)$.

Timing. The firm's problem has two stages. In stage one, before the shocks, the firm chooses its links α_i, the network structure before the shocks are realized (we will often refer to this as committed choice); this is the only choice under uncertainty, risky because it fixes how cost depends on the stochastic prices of committed inputs. In stage two, after the shocks, the firm chooses input quantities given realized shocks and prices. Demands are heterogeneous across firms and states, input markets clear, and this delivers the equilibrium price vector. We solve backward from stage two.

Two productivity shocks. A firm's productivity carries two shocks, both realized after the network is committed and neither insurable: a sector shock ε_i common to all firms in sector i, with $\varepsilon = (\varepsilon_1, \ldots, \varepsilon_n) \sim N(\mu, \Sigma)$, and an idiosyncratic firm shock u_{if}, mean zero and independent across firms. Only the sector shock exposes firms in the network to risk. Note that u_{if}, despite being uninsurable, is idiosyncratic; hence, it averages out across the continuum of firms. While ε_i is common across all firms in a sector, and hence it moves sector price indeces, whose risk in turn affects the individual firms' policy functions of the input-output relations. So the only committed-input risk is the sector covariance Σ, and the network is chosen to hedge that alone. To be clear, the shock u_{if} does affect the input total demand chosen in stage 2, but since it is idiosyncratic across firms in one sector, it averages to zero and affects neither the input shares nor the sector prices. In sum, it is not a sector risk against which firms insure with their ex ante choice of the links.[4]

Technology. A firm f in sector i produces with a Cobb-Douglas technology:

$$F(\alpha_i, L, X) = e^{\varepsilon_i + u_{if}} A_i(\alpha_i) \zeta(\alpha_i) L^{1-\sum_j \alpha_{ij}} \prod_j X_j^{\alpha_{ij}}, \tag{4}$$

where $\alpha_i = (\alpha_{i1}, \ldots, \alpha_{in})$ are the sector's input shares, chosen in stage one; L is labor, X_j

[4]It will reappear in the general-equilibrium model of Section 4, where the attached worker bears the firm's total shock $\varepsilon_i + u_{if}$ through the wage.

is the quantity of input j, and ζ is a normalization. Log productivity carries a third, deterministic term, $a_i(\alpha_i) \equiv \log A_i(\alpha_i)$: this is the productivity the sector engineers through its supplier weights, known at commitment and riskless. The choice of input shares trades off risk costs and benefits: the latter come from the productivity improvements that the optimal input organization provides. Total log productivity is $a_i(\alpha_i) + \varepsilon_i + u_{if}$, and the feasible set is $\mathcal{A}_i = \{\alpha_i \in [0,1]^n : \sum_j \alpha_{ij} \leq \bar{\alpha}_i\}$ with $\bar{\alpha}_i < 1$.[5] Firms in a sector share this technology and, at commitment, face the same environment with u_{if} not yet realized, so they are symmetric ex ante and commit to the same network; the network is a sector-level object, though firms differ ex post.

The markup. The elasticity σ operates only across varieties within a sector, so every firm sets the same constant markup $\mathcal{M} = \sigma/(\sigma - 1)$ over unit cost, and the sector price index is that markup over the sector unit cost. The sector good faces unit-elastic Cobb-Douglas demand, so the network is committed at the sector level and sector sales are nonstochastic. The markup is a constant level shift over marginal cost, with the competitive economy as the limit when $\sigma \to \infty$ and $\mathcal{M} \to 1$. Markup profit accrues to the diversified fund.

The Households-Investors' Kernel and the Final Good Demand. The firm discounts state-contingent profits by its own kernel Λ_i, normalized so $\mathbb{E}[\Lambda_i] = 1$, with $\lambda_i \equiv \log \Lambda_i$. Markets are incomplete: the firm cannot trade a full set of state-contingent claims, so it commits to a single network rather than a state-by-state plan and cannot insure the committed-input risk. In this section,, we assume the kernel is exogenously given, while in 4 it will be microfounded through the SDF of workers who share the firm's uninsurable risk.

A representative household owns a diversified fund holding all firms, supplies one unit of labor inelastically, and has Cobb-Douglas preferences over the n goods with shares $\beta = (\beta_1, \ldots, \beta_n)$, $\beta_j > 0$, $\sum_j \beta_j = 1$, and constant relative risk aversion $\rho \geq 1$. Its budget is labor income plus fund profits, the latter now including the markup; with the wage as numeraire, nominal spending E is the wage bill plus fund profits, and final demand satisfies $P_j C_j = \beta_j E$. The household's diversified valuation grounds the benchmark kernel λ^c; the firm-specific kernels λ_i depart from it because firm risk is uninsurable.

Before proceeding with the firms' optimization, we introduce two definitions that we will use repeatedly later on.

Definition 1 (Leontief Inverse)**.** *the Leontief inverse* $\mathcal{L}(\alpha) \equiv (I - \alpha)^{-1} = \sum_{k \geq 0} \alpha^k$*, whose entry* $\mathcal{L}_{ij}$ *collects the direct input share* α_{ij} *and every indirect chain through which sector* j *enters sector* i*, so that* $\mathcal{L}$ *is the operator that propagates shocks and prices through the network*

[5] We will return on this assumption on the shares' caps later on; it serves to deliver equilibrium uniqueness.

Definition 2 (Domar weights)**.** *Log GDP is* $y = -\beta^\top p$*, and the Domar weight of sector* i *is its sales share in GDP,* $\omega_i = P_i Q_i/\text{GDP}$*; with the wage as numeraire nominal GDP equals one, so* $\omega_i = P_i Q_i$*, collected in the vector* $\omega^\top = \beta^\top \mathcal{L}$*, the solution of input market clearing in Lemma 1. Since* ω *is the vector of consumption shares* β *propagated through the Leontief inverse* $\mathcal{L}$*, the weight* ω_i *measures the importance of* i *as a direct and indirect supplier, and it is the first-order effect of* i*'s productivity on aggregate output. In the endogenous-network economy,* ω *is an equilibrium object, evaluated below at the efficient network* α^**.*

Next, we solve the two-stage firms' problem backward. Firm f in sector i has total log productivity $\varepsilon_i + u_{if}$, so its unit cost is $\theta_{if} = -(\varepsilon_i + u_{if}) - a_i(\alpha_i) + \sum_j \alpha_{ij} p_j$. Both idiosyncratic productivity shocks enter only as level shifters on the firm's own cost. The two differ only in the input prices p_j, which are sector price indices: the sector shock ε_i, being common to the sector, is contained in them and propagates through the network, whereas the idiosyncratic shock u_{if}, averages out of the sector index. We therefore suppress u_{if} in what follows.

Stage 2: input choice and market clearing. Given α_i and realized prices P, the firm meets demand Q_i at minimum cost:

$$\Theta_i(\alpha_i, P)\, Q_i = \min_{L,\, X_j \geq 0} \Big\{ L + \sum_j P_j X_j \;:\; F(\alpha_i, L, X) \geq Q_i \Big\}. \tag{5}$$

The cost-minimizing choices satisfy $P_j X_{ij} = \alpha_{ij}\, \Theta_i Q_i$ and $L_i = (1 - \sum_j \alpha_{ij})\, \Theta_i Q_i$, so cost shares equal input shares and the unit cost is:

$$\Theta_i(\alpha_i, P) = \frac{1}{e^{\varepsilon_i} A_i(\alpha_i)} \prod_j P_j^{\alpha_{ij}}, \qquad \theta_i = -\varepsilon_i - a_i(\alpha_i) + \sum_j \alpha_{ij} p_j, \tag{6}$$

with derivations in Appendix A and where, once more, small letters indicate logs. Monopolistic competition sets the sector price $P_i = \mathcal{M}\, \Theta_i(\alpha_i, P)$, one per sector, competitive at $\mathcal{M} \to 1$, and input markets clear when $\sum_i X_{ij} + C_j = Q_j$ for every j.

Stage 1: the network choice. In anticipation of stage two, the firm commits its links before the shock to maximize expected discounted profit:

$$\alpha_i^* \in \arg\max_{\alpha_i \in \mathcal{A}_i} \; \mathbb{E}\big[\Lambda_i\, Q_i\, (P_i - \Theta_i(\alpha_i, P))\big]. \tag{7}$$

Since only unit cost depends on α_i, the firm minimizes the kernel-weighted expected cost. Crucially, because markets are incomplete, the kernel is firm-specific. We characterize the input shares' choice below, once the equilibrium is defined.

Given the two optimization stages, we can now state the equilibrium definition. The main idea is that the two stages must be consistent: the shares each sector commits must be optimal against the price distribution and market clearing, and that distribution must be the one generated by every sector's committed shares through the price system below.

Definition 3 (Equilibrium). *An equilibrium is a network $\alpha^* = (\alpha_1^*, \ldots, \alpha_n^*)$ and state-contingent prices and quantities $\{P(s), L_i(s), X_{ij}(s), Q_i(s), C_j(s)\}$ such that: (i) given α^* and each realized s, input choices solve (5) at $P(s)$ and the household chooses consumption optimally at $P(s)$; (ii) prices satisfy $P_i(s) = \Theta_i(\alpha_i^*, P(s))$ for every i and s; (iii) goods and labor markets clear for every s; (iv) each α_i^* solves the stage-one problem (7) given the joint distribution of prices, demand, and the kernel, firms being atomistic and taking that distribution as given; and (v) the distribution in (iv) is the one induced by (i) to (iii) under α^* and $\varepsilon \sim N(\mu, \Sigma)$. Conditions (iv) and (v) are the fixed point between the committed network shares and the price distribution they generate.*

Under our assumptions of constant returns to scale and a Gaussian joint distribution of prices and kernels, the price system has a closed-form solution.

Lemma 1 (The price system). *Given a network α with $\sum_j \alpha_{ij} < 1$ for all i, the cost-price fixed point $p_i = \theta_i(\alpha_i, p)$ has a unique solution:*

$$p(s) = -\mathcal{L}(\alpha)\big(\varepsilon(s) + a(\alpha) - \log \mathcal{M}\,\mathbf{1}\big), \qquad \mathcal{L}(\alpha) = (I - \alpha)^{-1} = \sum_{k \geq 0} \alpha^k, \tag{8}$$

so prices are jointly Gaussian with $\mathbb{E}[p] = -\mathcal{L}(\mu + a - \log \mathcal{M}\,\mathbf{1})$ and $\mathrm{Var}(p) = \mathcal{L}\,\Sigma\,\mathcal{L}'$. Sales $\omega_i \equiv P_i Q_i$ solve input market clearing, $\omega_j = \beta_j + \sum_i \alpha_{ij}\,\omega_i$, that is $\omega^\top = \beta' \mathcal{L}(\alpha)$, and are nonstochastic, depending on α and β only.[6] Aggregate log output is $y(\alpha) = \omega(\alpha)'(\varepsilon + a(\alpha))$ with Domar weights $\omega(\alpha)' = \beta' \mathcal{L}(\alpha)$, where β is the vector of consumption shares.

Appendix A proves Lemma 1.

Note that the markup term $\log \mathcal{M}\,\mathcal{L}\mathbf{1}$ is a deterministic level shift with zero covariance, hence the variance of prices depends only on the variance of the uninsurable risk ϵ and how it transmits to the network through the Leontief matrix, $\mathcal{L}(\alpha)$.

Assumptions and Their Role for the Equilibrium. Let us summarize again how the assumptions made earlier determine the equilibrium structure. First, sellers determine the price system: they sell only when the price equals the mark-up plus the marginal cost, $P_i = \mathcal{M}\Theta(\alpha_i, P_i)$ (where P_i is the price vector, $\Theta(\alpha_i, P_i)$ is the marginal cost and $\mathcal{M}$ is the mark-up vector) and at that price they accommodate any quantity demanded. Because constant returns make unit cost independent of scale, supply is perfectly elastic at that price, so the price is determined from the cost side and quantities do not enter this system.

Second, each buyer determines the optimal demand by equating the marginal product to that common price, $P_j X_{ij} = \alpha_{ij} P_i Q_i$, and the sum of all demands clears the market for each

[6] With the markup, revenue-based input shares are $\alpha_{ij}/\mathcal{M}$ rather than α_{ij}, so revenue-based sales carry the uniform factor $\mathcal{M}$; since the markup is common across sectors this rescales sales uniformly and leaves the cost-based Domar weights $\omega^\top = \beta' \mathcal{L}(\alpha)$, the object relevant for productivity aggregation, and every relative quantity unchanged. Sales remain nonstochastic, which is all the network choice below uses.

good: this gives the sales system $\omega^\top = \beta' \mathcal{L}(\alpha)$.[7] Constant returns to scale make the equilibrium block recursive: the price block closes on itself and is solved first, and the quantity block is solved given those prices.

Finally, the equilibrium of stage 2 is unique since the assumption that α_i is below one makes the Leontief matrix $I - \alpha$ invertible and secures the inner fixed point through the spectral radius.[8] The equilibrium of stage 2 is guaranteed by the concavity of $a_i(\alpha_i$: this makes the firms' first-order conditions differentiable.[9]

This structure parallels the one in Kopytov et al. [2024]. It is not affected by incomplete markets: the latter just makes input demand conditional on shock realization, but equilibrium rests on the competitive structure and market clearing. It also parallels the structure in Hopenhayn [1992] and Hopenhayn and Rogerson [1993], where firms face uninsurable productivity risk: the price is read off a zero-profit condition with an infinite elastic supply, while quantity demanded clears the market.

Given the above, we can summarize the conditions that lead to existence and uniqueness.

Risk-adjusted prices. We now characterize the stage-one network choice. By relying on lognormality and the Cobb-Douglas production function, we can restate problem (7) as a risk-adjusted minimization. Derivations are in Appendix A.

Lemma 2 (Risk-adjusted prices)**.** *Problem* (7) *is equivalent to*

$$\alpha_i^* \in \arg\max_{\alpha_i \in \mathcal{A}_i} \; a_i(\alpha_i) - \sum_j \alpha_{ij}\, \mathcal{R}_{ij}, \qquad \mathcal{R}_{ij} \equiv \mathbb{E}[p_j] + \mathrm{Cov}(p_j, \lambda_i). \tag{9}$$

At an interior solution $\partial a_i / \partial \alpha_{ij} = \mathcal{R}_{ij}$*, and* $\partial \alpha_{ij} / \partial \mathcal{R}_{ik} = [H_i^{-1}]_{jk}$*, where* $[H_i^{-1}]_{jk}$ *is the entry of the inverse of firm* i*'s Hessian matrix; since the diagonal of* H_i^{-1} *is negative, a higher* $\mathcal{R}_{ij}$ *lowers* α_{ij}*. The link is severed once the implied share reaches the boundary of* $\mathcal{A}_i$*.*

Proof of 2 is in Appendix A. Note that, as we show in the derivations, the equivalence between the maximization problem in 5 and 9 is not an approximation, but an exact result. Intuitively, because returns to scale are constant, quantities and prices are non-stochastic, and risk is carried only through the kernel. The above clarifies the crucial difference with the complete market case. The risk-adjusted price $\mathcal{R}_{ij}$ is the expected log price of input j plus a loading equal to how j's price covaries with the buyer's kernel. This loading, $\mathrm{Cov}(p_j, \lambda_i)$, is specific to the buyer firm: the same customer firm, offering the same price, can be valued very differently by buyer firms with different kernels.

Let us now derive the optimality condition. The firm trades the productivity a network

[7] The labour market clears by Walras' law.

[8] Furthermore, we assumed that $\sum_j \alpha_{ij} \leq \bar{\alpha}_i < 1$, so labor is essential: this allows us to rule out the self-referential loop of pure intermediates that would lead to indeterminacy.

[9] The equilibrium in the system of input shares exists by the Brower's theorem.

delivers, $a_i(\alpha_i)$, against the risk-adjusted cost of the inputs it commits to, $\sum_j \alpha_{ij}\mathcal{R}_{ij}$. Differentiating in α_{ij} and setting the derivative to zero gives, at an interior optimum: $\frac{\partial a_i}{\partial \alpha_{ij}} = \mathcal{R}_{ij}$, so the firm leans on supplier j until the marginal productivity gain equals that supplier's risk-adjusted price. To examine uniqueness, we can totally differentiate the first-order condition $\partial a_i/\partial \alpha_i = \mathcal{R}_i$, which gives $H_i\, d\alpha_i = d\mathcal{R}_i$, where H_i is the Hessian of a_i, hence: $\frac{\partial \alpha_{ij}}{\partial \mathcal{R}_{ik}} = [H_i^{-1}]_{jk}$. Strict concavity of a_i makes H_i negative definite: a higher $\mathcal{R}_{ij}$ lowers α_{ij}. The firm sources less from an input that becomes more expensive or riskier, and once $\mathcal{R}_{ij}$ rises enough that the implied share reaches the boundary of $\mathcal{A}_i$, the link is severed.[10]

3.2 Comparison to Exogenous Network and Complete Markets

The endogenous network under incomplete markets derived so far also nests two special cases common in the past literature: the exogenous network (Baqaee and Farhi [2019]) and the endogenous network with complete markets (Kopytov et al. [2024]). This comparison also allows us to highlight the contribution of this paper.

Complete markets and the efficient network case. If $\lambda_i = \lambda^c$ for all i, then $\mathcal{R}_{ij} = \mathcal{R}_j^c \equiv \mathbb{E}[p_j] + \mathrm{Cov}(p_j, \lambda^c)$ is common to all buyers, every firm faces the same shadow-price vector, and the equilibrium network coincides with the solution of the planner's problem and is efficient. The firm solves the first-order conditions $\nabla a_i = \mathcal{R}^c$ as a single system in a common price vector, which is equivalent to the planner's first-order condition under the common kernel. This is the efficient endogenous network of the complete-market economy [Kopytov et al., 2024].

Fixed network: Domar sufficiency. Fix the network at any $\bar{\alpha}$. Then the equilibrium allocation is independent of the kernel,[11] and Domar weights $\omega = \beta' \mathcal{L}(\bar{\alpha})$ are a sufficient statistic for the first-order response of expected output to productivity, $\partial \mathbb{E}[y]/\partial \mu_i = \omega_i$. This allocation corresponds to the first-order accounting benchmark of Hulten [1978] and Baqaee and Farhi [2019], in which sales shares summarize the impact of productivity and the pricing kernel does not affect the real allocation.

Note that under a fixed network the discount factor is allocationally irrelevant, with or without complete markets. In other words, uncertainty or risk becomes relevant only to the extent that the choice of the input-output relations can be conditioned on that and changed.

[10] In the quantitative analysis we adopt a logit smoothing for the entry and exit events of a link, together with richer shock processes; it has the interior model as its zero-noise limit and is not needed for the results here.

[11] With $\bar{\alpha}$ fixed, prices $p = -\mathcal{L}(\bar{\alpha})(\varepsilon + a)$ depend only on $\bar{\alpha}$ and ε, and the quantities derived in stage 2 solve (5) at realized prices, which involves no kernel. From $\mathbb{E}[y] = \omega'(\mu + a)$, differentiating in μ_i gives ω_i.

3.3 The Risk-Misallocation Wedge and Precautionary Retrenchment

Decompose each firm's kernel as $\lambda_i = \lambda^c + \xi_i$, where $\xi_i \equiv \lambda_i - \lambda^c$ is idiosyncratic, absent under complete markets, and averages to zero across firms so that a diversified claim carries no loading on it.

Definition 4 (Wedge)**.** *The risk-misallocation wedge on the link from j to i is $\tau_{ij} \equiv \mathcal{R}_{ij} - \mathcal{R}_j^c = \mathrm{Cov}(p_j, \lambda_i - \lambda^c) = \mathrm{Cov}(p_j, \xi_i)$.*

Intuitively, the wedge is the distance that each firm faces, from the cost of risk between complete and incomplete markets. This distance is related to the correlation between the individual kernels and the input prices: **some firms are better than others at hedging**. This in turn leads to the wedge distribution.

Proposition 1 (Decomposition of precaution)**.** *The firm's precautionary loading on input j decomposes into:*

$$\underbrace{\mathrm{Cov}(p_j, \lambda_i)}_{\textit{total precaution}} = \underbrace{\mathrm{Cov}(p_j, \lambda^c)}_{\textit{efficient: aggregate risk}} + \underbrace{\tau_{ij}}_{\textit{distortionary: idiosyncratic risk}} . \tag{10}$$

The first component is the loading a diversified owner, or a planner, applies; the second is excess precaution against diversifiable risk. The latter carries no aggregate risk premium, and is zero for every firm under complete markets.

The identity 10 follows immediately from the linearity of the covariance applied to $\lambda_i = \lambda^c + \xi_i$. Note that $\tau_{ij} = 0$ for all i, j requires $\mathrm{Cov}(p_j, \xi_i) = 0$ for all j, since in equilibrium prices transmit the firms' shock to all other firms, and this holds if and only if $\xi_i = 0$ for all i, that is, under complete markets.

An Example. When the kernel gap between complete and incomplete markets is proportional to the firm's own productivity shock, one can derive a closed form for the firm-level wedge. Suppose for instance that $\xi_i = -\phi_i \varepsilon_i$, where $\phi_i \geq 0$ indexes the degree of uninsurability. Since $\mathrm{Cov}(p_j, \varepsilon_i) = -[\mathcal{L}\Sigma]_{ji}$ from $p(s) = -\mathcal{L}(\alpha)\big(\varepsilon(s) + a(\alpha) - \log \mathcal{M}\,\mathbf{1}\big)$, the wedge becomes:

$$\tau_{ij} = \phi_i \, [\mathcal{L}\Sigma]_{ji}, \tag{11}$$

which under diagonal Σ is $\phi_i \mathcal{L}_{ji} \Sigma_{ii}$. The wedge scales with the degree of uninsurability, the amount of own risk, and supplier j's network exposure to firm i, given by the Leontief inverse.[12] The firm weights states by Λ_i when committing to the network; it therefore tilts away from inputs whose prices are high in the states the firm discounts most, namely its own bad states. It does so partly to hedge against aggregate risk, with a loading on the complete market kernel λ^c,

[12] Equation **??** is also the reduced form taken to the data to calibrate the full general equilibrium model presented in Section 5.1.

and partly due to a precautionary motive related to the firm's idiosyncratic risk distribution, τ_{ij}. Typically, the precaution is proportional to the variance of the shocks; in a network, the covariance with shocks from other firms matters too.

Proposition 2 (Irrelevance of uninsurable risk under a fixed network)**.** *Suppose the network is fixed at any $\bar{\alpha}$. The equilibrium allocation under firm-specific kernels $\{\lambda_i\}$ coincides, state by state, with the allocation under the common kernel λ^c. Incomplete markets cause no production misallocation when the network cannot adjust.*

The proof is straightforward. The kernel enters only through the network choice in stage 1, but when $\alpha_i = \bar{\alpha}_i$ is fixed, that choice is not made. The input quantities chosen in stage 2 solve (5) at realized prices, and prices (equation (8)) depend only on $\bar{\alpha}$ and ε. Therefore, every price and quantity is identical to the complete-market allocation, regardless of $\{\lambda_i\}$. This makes uninsurable risk irrelevant under a fixed network.

Let us now examine the consequences of risk for the network structure or for individual links. Let α^* solve $\nabla a_i = \mathcal{R}^c$ (the efficient network) and α^{IM} solve $\nabla a_i = \mathcal{R}^c + \tau_i$ (the incomplete-market network). The following proposition holds.

Proposition 3 (Retrenchment)**.** *To first order, holding the price distribution at its efficient-network value, $\alpha_i^{IM} - \alpha_i^* = H_i^{-1}\tau_i + o(\|\tau_i\|)$. Since H_i is negative definite, $\tau_{ij} > 0$ lowers α_{ij}: the firm sources less from suppliers whose prices rise in its own bad states, and severs a link active under α^* once the wedge drives its share to the boundary. The general-equilibrium adjustment of the price distribution, itself first order in τ, is accounted for in Proposition 4.*

The proof of 3 is also straightforward. The efficient and distorted networks solve $\nabla a_i(\alpha_i^*) = \mathcal{R}_i^c$ and $\nabla a_i(\alpha_i^{IM}) = \mathcal{R}_i^c + \tau_i$. Subtracting and taking the second derivative, $H_i(\alpha_i^{IM} - \alpha_i^*) = \tau_i$. The jth diagonal component is $[H_i^{-1}]_{jj}\tau_{ij} < 0$ when $\tau_{ij} > 0$; when the displacement carries α_{ij} to zero, the interior condition fails, and the link is dropped.

Firms' heterogeneous ability to hedge against risk leads to misallocation of inputs. We therefore now define and construct a misallocation wedge. To parallel the concept in Hsieh and Klenow [2009], we construct the wedge by first defining a Pareto benchmark, which we identify with the choice of a diversified investor or the planner. Note that since the distortion comes from risk, we assume the diversified investor (or the planner) has some degree of risk aversion, ρ. This makes risk costly. The planner is going to maximize $\mathcal{W}(\alpha) = \mathbb{E}[y(\alpha)] - \frac{1}{2}(\rho - 1)\,\mathrm{Var}[y(\alpha)]$ where α is stochastic since is chosen taking into account the amount of risk in the network. First of all, making use of the equilibrium output derived in Lemma 1, $y(\alpha) = \omega(\alpha)'(\varepsilon + a(\alpha))$ with

Domar weights $\omega(\alpha)' = \beta'\mathcal{L}(\alpha)$, we can re-write the welfare in terms of Domar's weights:[13]

$$\mathcal{W}(\omega) = \omega^\top \mu + \bar{a}(\omega) - \tfrac{1}{2}(\rho - 1)\,\omega^\top \Sigma \omega, \tag{12}$$

with $\bar{a}$ the aggregate productivity shifter. Rewriting welfare in terms of Domar's weights is useful for several reasons. First, based on Hulten's theorem,, the Domar weight provides the pass-through from firms' shocks or wedges to aggregate output. Furthermore, note that the Domar weights already capture the network propagation since they are a function of the Leontief matrix, $\omega(\alpha)' = \beta'\mathcal{L}(\alpha)$: as such, they already embed how the firms' distortions propagate through the network. Third, those weights are non-stochastic since they depend only on technology and preferences; this makes the pass-through of the wedge into the aggregate independent of the wedge itself (which depends on the extent of risk). Lastly, using Domar's weight allows me to nest the welfare results from Baqaee and Farhi [2020].

To measure the aggregate welfare cost of a shift in the firms' wedge, it is useful to define the following objects. First, the Hessian of 12, $\mathcal{H} = \nabla^2 \bar{a}(\omega^*) - (\rho - 1)\Sigma$: this measures the curvature of welfare in the sales allocation; its diagonal is how costly it is to move weight onto a single sector, with firms trading off at the margin the productivity benefits from the re-organization with the risk-costs of it. Its off-diagonal is how the two sectors interact in welfare, joint technological substitutability plus their output covariance. Second, let us $\mathcal{T}$ as the aggregate wedge shift, the firm wedges $\{\tau_i\}$ aggregated into Domar-weight space.

Proposition 4 (Misallocation)**.** *Relative to the efficient network, the incomplete-market network shifts the Domar weights and lowers welfare by:*

$$\omega^{IM} - \omega^* = -\mathcal{H}^{-1}\mathcal{T} + o(|\mathcal{T}|), \qquad \mathcal{W}(\omega^*) - \mathcal{W}(\omega^{IM}) = \tfrac{1}{2}\,\mathcal{T}^\top\big(-\mathcal{H}^{-1}\big)\mathcal{T} + \Delta^{\text{ext}} + o(|\mathcal{T}|^2), \tag{13}$$

a positive-definite quadratic form in the wedges (the intensive margin) plus a non-negative discrete term Δ^{ext} from severed links (the extensive margin). The matrix $\mathcal{H}^{-1}$ propagates each firm's wedge through the network: its off-diagonal entries, the global substitution patterns, transmit a wedge at i to the Domar weights, and hence the misallocation, of firms connected to i. When these cross-firm effects are switched off, so that $\mathcal{H}$ is block-diagonal with blocks $\omega_i H_i$, the intensive-margin loss reduces to the Domar-weighted sum $\frac{1}{2}\sum_i \omega_i\,\tau_i^\top(-H_i^{-1})\tau_i$. The loss is zero if and only if $\tau = 0$, that is, under complete markets, or the network is fixed. Measured aggregate productivity falls with the wedge, since the reorganization routes inputs away from cost-minimizing suppliers.

The proof of 4 is given in Appendix A.

[13] Since $\mathbb{E}[\varepsilon] = \mu$ and $a(\alpha)$ is deterministic given the network,$\mathbb{E}[y] = \omega'\,\mathbb{E}[\varepsilon + a] = \omega'(\mu + a)$.. Writing $\bar{a}(\omega) \equiv \omega' a$ for the Domar-weighted productivity shifter, the mean is $\omega'\mu + \bar{a}(\omega)$, the first two terms of the welfare expression. Let us now derive the variance. The only random part of y is $\omega'\varepsilon$, since $a(\alpha)$ is deterministic given the network, so$\text{Var}(y) = \text{Var}(\omega'\varepsilon) = \omega'\Sigma\,\omega$. Multiplying by the risk coefficient, the penalty is $\frac{1}{2}(\rho - 1)\,\omega'\Sigma\,\omega$, the third term.

Summary: the Role of Uninsurable Risk in Optimizing Network Choices for Misallocation. We have shown that risk misallocation, that is, a distribution of wedges proportional to the risk loading onto individual firms, arises when two conditions materialize: the network is formed as a result of a firm's optimization problem, and firms face uninsurable risk. The fact that firms re-optimize their links in the face of shocks is a very plausible and realistic scenario. Uninsurable risk is equally plausible. In this case, we also show that the firms' discount factor acquires an allocational role, as it determines the wedge distribution.

Before moving to the full general equilibrium economy, we provide an intermediate step showing how the misallocation wedge would change if we were to extend the network to include foreign suppliers, and hence exchange rate risk.

3.4 Introducing Exchange Rate Risk

We now extend the wedge to include exchange rate risk being transmitted through an international input-output network with foreign suppliers. Since the model described in this section takes the kernel and investors' valuation as exogenously given, we also introduce the exchange rate process exogenously. In Section 4, we make the exchange rate microfounded through the Backus and Smith [1993] international risk-sharing condition. To fix ideas and derive a closed-form solution of the wedge for this case, we rely on the same affine parametric form that we introduced in 11 for the gap between the fully diversified and the undiversified kernel. We extend it to include exchange rate risk.:

$$\xi_i = -\phi_i, \varepsilon_i + \psi_i, (e - \mathbb{E}e), \tag{14}$$

with $\phi_i \geq 0$ the degree of uninsurability and $\psi_i \geq 0$ the loading on the exchange rate. As explained earlier, and still true here, even with imperfect insurance pass-through, $\phi_i \geq 0$ and $\psi_i \geq 0$ are taken as given. They will be microfounded in 4 through a limited-commitment contract between the firm and the worker, which implies some partial degree of rent and risk sharing. When $\psi_i = 0$, this nests the closed economy case and the wedge is: $\tau_{ij} = \phi_i, [\mathcal{L}, \Sigma]ji$. In the open economy, $\psi_i > 0$ for firms that source abroad, and the wedge on a foreign link carries an additional currency risk component. Since the effective price of a foreign input in the buyer's currency is $p_j + e$ we can re-derive the wedge as follows:

$$\tau ij^F = \mathrm{Cov}(p_j + e, , \xi_i) = \underbrace{\phi_i, [\mathcal{L}, \Sigma]ji}\text{ node risk} + \underbrace{\psi_i, \mathrm{Var}(e)}_{\text{currency wedge}} + \psi_i, \mathrm{Cov}(p_j, e - \mathbb{E}e), \tag{15}$$

4 General Equilibrium Model with Endogenous Input-Output Network and Uninsurable Risk

Our research question is to measure the cost of misallocation. That should account for all general equilibrium spillovers beyond those stemming from input-market equilibria. We therefore embed models with heterogeneous firms interacting through input-output relations into a full general equilibrium model with households/workers, who are also exposed to firms' risk through a limited-commitment contract. Formalizing the firms' investors' maximization problem will allow us to provide microfoundations for the firm-specific kernel assumed so far.

The model's backbone features two types of households. Some work in the firm and bear the firm's wage risk through rent sharing via a limited-commitment contract. Those households cannot fully diversify this wage risk: we assume they invest their savings in a pension fund that, due to regulatory restrictions, cannot fully diversify. They thus solve a dynamic consumption-saving problem through which they self-insure, and the resulting SDF is conditional on the uninsurable shock realization. A second type of household, the capitalists, owns the firms' shares through a diversified fund. Given this, the firms' discount factor is a combination of pricing kernels, with the firm's specific component given by the workers' discount factors. This intermediation structure, in which investors with different risk attitudes sort into funds with different characteristics, is again supported by findings from the asset-demand literature (see Koijen and Yogo [2019]).

Below we write down the household/workers' dynamic optimization problems, the contract that sets the wage, and the intermediation setting. We then examine how the misallocation wedges and their distribution change in this context. Lastly, we simulate the model, which now also accommodates an aggregate shock, and provide a quantitative analysis geared toward assessing how a specific counterfactual policy, a trade tariff, may affect the wedge distribution. In doing so, we also validate the model against the empirical predictions laid down in section 2.1 and 2.2. Derivations are in Appendix A.

Households-Workers and Partial Insurance within the Firm. A worker attached to firm i supplies one unit of labor inelastically, earns the wage w_i set by the contract below, and saves in the pension fund. Human capital is non-tradable, so the worker cannot issue claims on the wage. With pension holdings a_i^p, pension return R^p, and aggregate state Z that includes the cross-sectional distribution of wealth, the worker solves the following optimization problem:

$$V_i^w(a_i^p, Z) = \max_{c_i^w \geq 0,\, a_i^{p\prime} \geq 0} u(c_i^w) + \beta\, \mathbb{E}\big[V_i^w(a_i^{p\prime}, Z') \mid Z, \varepsilon_i\big] \quad \text{s.t.} \quad c_i^w + a_i^{p\prime} = w_i(Z, \varepsilon_i) + R^p(Z)\, a_i^p, \tag{16}$$

which leads to the Euler equation:

$$u'(c_i^w) = \beta\, \mathbb{E}\big[R^p(Z')\, u'(c_i^{w\prime}) \mid Z, \varepsilon_i\big] \tag{17}$$

and the stochastic discount factor $\Lambda_i^w = \beta\, u'(c_i^{w\prime})/u'(c_i^w)$. The wage loads on firm i's productivity, which is uninsurable, and the pension return cannot hedge that risk, so consumption, and the discount factor, load on the firm's productivity. The worker is attached to an individual firm, so the relevant shock is the firm's total productivity, the sector component ε_i plus the idiosyncratic component u_{if}; the worker's kernel loads on this firm-level shock, which is more volatile than the sector shock alone. Let us make two remarks. First, note that the firm shock, u_{if}, which did not propagate through the network since it did not affect the sector prices, does affect households' saving behavior. Second, the markup profit accrues to the diversified fund and so does not enter the worker's kernel. Given the lack of full insurability, c_i^w is strictly increasing in ε_i and Λ_i^w loads on ε_i, differing from λ^c by a firm-specific component; saving mitigates but does not eliminate this component, since no traded asset spans ε_i.

Capitalists. A capitalist owns equity through the capital fund and does not work. Given the return that the capitalist receives from the fund, R^c, and given the holdings a^c, the capitalist solves the following optimization problem:

$$V^c(a^c, Z) = \max_{c\geq 0,\, a^{c\prime}} \; u(c) + \beta\, \mathbb{E}\big[V^c(a^{c\prime}, Z') \mid Z\big] \quad \text{s.t.} \quad c + a^{c\prime} = R^c(Z)\, a^c, \tag{18}$$

with Euler equation $u'(c^c) = \beta\, \mathbb{E}[R^c(Z')\, u'(c^{c\prime}) \mid Z]$ and discount factor $\Lambda^c = \beta\, u'(c^{c\prime})/u'(c^c)$. Since the fund holds the market portfolio, its payoff is measurable with respect to aggregate variables and risk by a law of large numbers: hence the capitalists' consumption and SDF, Λ^c, loads only on aggregate risk.

Funds Structure. Saving is intermediated by two funds. The pension fund, in which workers invest their savings, invests a share ϕ_w in the market portfolio and the rest in a safe bond, so its return $R^p = (1-\phi_w)R^f + \phi_w R^m$ loads only on aggregates. The pension holds only aggregate assets, so it does not hedge a worker's exposure to their own employer; the worker's labor income therefore carries firm-specific risk that no traded asset offsets.

The capital fund, held by capitalists, holds the market portfolio and is the marginal equity investor. Market clearing requires pension assets to equal workers' savings, the safe bond to be in fixed supply, and equity demand to equal the claims firms issue.

This intermediation encodes the non-diversifiability of labor income: households cannot sell claims on their own wages, and own-employer equity positions are both limited and, being a doubling of labor-income risk, welfare-reducing. The partition into worker and capitalist funds is a model-based implementation of that friction.

Wage Contract. The wage is set by a limited-commitment contract between firm i and its workers, following Kocherlakota [1996], signed before the shock and self-enforcing state by state so that no promise-keeping state is carried across periods. Full commitment, under which the wage is fully insured and the wedge vanishes, is the polar case; relative to positing uninsurable income risk directly, as in Aiyagari [1994], the contract derives the uninsurability and its dependence on the current shock from the friction rather than assuming them.[14]

Given the firms' profits, $\Pi_i(s) = P_i(s)Q_i(s) - \Theta_i(s)Q_i(s) - w_i(s)L_i$, and given the workers' consumption $c_i^w(s)$, obtained from the budget constraint defined in 17, the constrained-efficient contract is defined so that the wage schedule $\{w_i(s)\}$ and the network α_i are chosen to maximize the fund's valuation of the equity payoff subject to the worker's ex ante participation, the worker's state-by-state limited commitment, and the firm's state-by-state limited commitment[15]:

$$\max_{\{w_i(s)\},\,\alpha_i} \quad \mathbb{E}\big[\Lambda^c\,\Pi_i\big] \tag{19}$$

$$\text{s.t.} \quad \mathbb{E}\big[u\big(c_i^w(s)\big)\big] \geq \bar{U}_i \qquad (\nu_i) \tag{20}$$

$$u\big(c_i^w(s)\big) \geq \Omega_i^w(s) \qquad (\mu_i^w(s)) \quad \text{for all } s \tag{21}$$

$$\Pi_i(s) \geq \Omega_i^f(s) \qquad (\mu_i^f(s)) \quad \text{for all } s. \tag{22}$$

The constraint (20) is the worker's ex ante participation constraint: the contract must deliver expected utility at least the worker's reservation value $\bar{U}_i$, the value of not entering the relationship, and its multiplier is ν_i. Constraints (21) and (22) are the state-by-state limited-commitment constraints: in each state the worker must prefer to stay, with continuation value $\Omega_i^w(s)$, and the firm must prefer to honor the contract rather than take its outside value $\Omega_i^f(s)$, with multipliers $\mu_i^w(s)$ and $\mu_i^f(s)$. The multiplier ν_i is the shadow price of the ex ante promise and is constant across states; the multipliers $\mu_i^w(s)$ and $\mu_i^f(s)$ are the shadow prices of the state constraints and are positive only where those constraints bind. Together they govern how much of the firm's risk passes into the wage. Given the above the first-order condition with respect tom wages is:

$$\Lambda^c(s) + \mu_i^f(s) = \big(\nu_i + \mu_i^w(s)\big)\,u'\big(c_i^w(s)\big). \tag{23}$$

When both state constraints are slack, only ν_i is active and the wage first-order condition (23) reduces to $\Lambda^c(s) = \nu_i\,u'(c_i^w(s))$, so the worker's marginal utility is proportional to the fund's kernel and the wage is fully insured. Pass-through arises only in states where the worker's

[14] The contract is two-sided limited commitment, with participation constraints on both the firm and the worker: the wage is fully insured where neither constraint binds, and tracks the firm's productivity only where one binds. Signing it afresh each period and requiring it to be self-enforcing within the period removes the promised-utility state that a fully dynamic contract would carry across periods, so the worker's exposure is a time-invariant function of the current shock rather than of the history of past binding constraints; if productivity is persistent the exposure still moves over time through the current shock, but the mapping itself does not depend on history. This is what yields a time-invariant loading in closed form, at the cost of not modeling promised-utility dynamics.

[15] The profits in the contract are weighted by the SDF of the diversified funds, since the principal of the contract represents the capitalists' interests.

constraint binds, $\mu_i^w(s) > 0$, in bad states, or the firm's binds, $\mu_i^f(s) > 0$, in good states. The wage then tracks the shock. This is the source of the **countercyclical uninsurability**: insurance holds in normal states and breaks at the tails.

Let us now examine the first order condition of the network. Since the wage is separately optimal, the envelope theorem lets us price the network by its direct effect on the objective alone, which is the firm's profit weighted by the fund's kernel Λ^c plus the firm's participation constraint weighted by μ_i^f. We normalize this sum to the effective kernel $\Lambda_i^{\text{eff}} \propto \Lambda^c + \mu_i^f$. Its value depends on the state. When the firm is not tempted to renege its constraint is slack, $\mu_i^f = 0$, and $\Lambda_i^{\text{eff}} \propto \Lambda^c$: the network is priced by the diversified kernel, the efficient benchmark. When the firm is tempted, $\mu_i^f > 0$, and in those same states the worker's constraint is slack, $\mu_i^w = 0$, since the two are tempted in opposite states, the firm in good states and the worker in bad. Setting $\mu_i^w = 0$ in the wage condition (23) gives $\Lambda^c + \mu_i^f = \nu_i\, u'(c_i^w)$, so there the effective kernel is proportional to the worker's marginal utility, $\Lambda_i^{\text{eff}} \propto \Lambda_i^w$. The sourcing decision that emerges is priced by an effective kernel blending the two households' discount factors,

$$\Lambda_i^{\text{eff}} = \gamma_i^c\, \Lambda^c + \gamma_i^w\, \Lambda_i^w, \qquad \gamma_i^c + \gamma_i^w = 1, \tag{24}$$

with the worker's weight γ_i^w increasing in the limited-commitment multipliers, zero under full commitment and positive under limited commitment.

The Wedge. Given 24, the wedge then follows from Definition 4, and the level identity:

$$\tau_{ij} = \text{Cov}\left(p_j,\, \lambda_i^{\text{eff}} - \lambda^c\right) = \gamma_i^w\ \text{Cov}\left(p_j,\, \lambda_i^w - \lambda^c\right), \tag{25}$$

Full commitment sets $\mu_i^f \equiv 0$ and, with a fully insured wage, $\lambda_i^w = \lambda^c$, so $\tau_{ij} = 0$.

The same contract determines who bears which risk. Aggregate risk is borne by capitalists through the capital fund and by workers to the extent $\phi_w > 0$ through the pension fund. Firm-specific risk is borne by the attached worker in the fraction γ_i^w and is retained by the wage. The residual passes into equity, where the fund diversifies it away, so the capitalist is fully insured from idiosyncratic risk. Workers smooth smooth the retained risk through precautioonary saving or self-insurance over time but not across states. Hence consumption loads on ε_i with an intensity that falls in wealth but does not reach zero, while the capitalist's consumption loads on aggregates alone. The idiosyncratic risk the worker carries is the gap $\lambda_i^w - \lambda^c$ that enters the wedge 25.

Definition 5 (Equilibrium.)**.** *A recursive competitive equilibrium consists of value functions and policies for workers and capitalists, a wage contract and a network for each firm, asset returns and rentals, and a law of motion for Z, such that households solve (16) and (18), the contract solves (19), firms solve their Stage-1 and Stage-2 problems, the funds price and clear the equity and bond markets, all input and goods markets clear, and the wealth distribution evolves*

consistently.

The Channels Relatively to the Complete Market Case. Two features distinguish the solution from the complete-market model. First, the aggregate state contains the cross-sectional distribution of wealth, so the model is a production network built on an incomplete-markets household side. Second, the effective kernels (24) are firm-specific and the equilibrium is a genuine fixed point in the discount factors and prices. Therefore any shocks that affects the network spills over to the rest of the economy through two sources: the input prices that affect firms decisions (the channel present also in the firm model considered before) and the effect on households. The latter include the impact of risk on precautionary savings and distributional effects, since the risk is perceived differently along the income and wealth distribution. We will examine this aspects in more detail in the numerical assessment.

Endogenous Counter-cyclical Earning Risk. The earning process and risk are determined in this model endogenously from the limited commitment contract. Hence, one question that arises is whether the implied dynamic is in line with empirical evidence (see for instance [Storesletten et al., 2004, Guvenen et al., 2014]) The primitive that in our model drives the earning process is the limited-commitment contract (see [Thomas and Worrall, 1988, Kocherlakota, 1996]). In a bad aggregate state the firm's profit and its outside option deteriorate, the firm's participation constraint binds, and the constrained-efficient contract can no longer sustain full insurance. Therefore the firm passes more of its shock into the wage. The worker's marginal utility then loads more heavily on the firm shock, and the degree of uninsurability rises. This leads to countercyclical earning risk. A contribution of this model is that the countercyclicality is not imposed, as it is often the case in incomplete market models, but is the resulting of the contracting arrangement. This also bears consequences for the objects of interest to this paper. The wedge inherits the countercyclicality, hence it is near zero in normal states, where the contract insures and $\phi(z) = 0$, and large in bad states, where the constraint binds and ϕ rises.

International Value Chains and Exchange Rate Risk. We now extend the macro model to accommodate domestic and foreign production linkages and exchange rate risk for the latter. Some of a firm's potential suppliers are located abroad: their prices are denominated in a foreign currency and carry an exchange rate risk. The origin of the supplier is an exogenous attribute of a node in a foreign economy: in this case firm's endogenous choice of the links (and its severance) takes into account a price that embeds exchange rate risk.

A foreign input's price is converted at the log exchange rate e, so a cross-border link contributes $p_j + e$ rather than p_j, weighted by each firm's net cross-border exposure. This is the single additive term $\mathcal{L}(\alpha)\chi(\alpha)e$ in the price system (28). Other than that, the same Leontief inversion applies with one extra column driven by e. That e is pinned by the Backus and Smith

[1993] condition. If the home and foreign funds trade a complete set of internationally traded state-contingent claims, their first-order conditions equate state prices:

$$\Lambda^c(s)/\mathbb{E}\Lambda^c = Q(s)\Lambda^*(s)/\mathbb{E}\Lambda^* \tag{26}$$

with $Q(s)$ the level real exchange rate, and taking logs gives $e(s) = \lambda^c(s) - \lambda^*(s)$. It follows that the exchange rate is priced by the single diversified kernel λ^c, the marginal international investor. Nodes are partitioned into domestic $\mathcal{D}$ and foreign $\mathcal{F}$. Each firm prices in its own currency, and the price a buyer i pays for input j, in the buyer's currency, is $p_j + c_{ij}\, e$, where e is the log exchange rate, home currency per unit of foreign, hence:

$$c_{ij} = \begin{cases} +1 & i \in \mathcal{D},\ j \in \mathcal{F}, \\ -1 & i \in \mathcal{F},\ j \in \mathcal{D}, \\ \ \ 0 & \text{same country.} \end{cases} \tag{27}$$

A cross-border edge carries $+e$ when a home firm buys abroad and $-e$ when a foreign firm buys at home. The firm's unit cost gains one term, $\chi_i(\alpha_i)\, e$, where $\chi_i(\alpha_i) = \sum_j \alpha_{ij} c_{ij}$ is its net cross-border input exposure. The price system is then the Leontief inversion of Lemma 1 with one additive term.

Lemma 3 (Price system, open economy)**.** *With cross-border conversion, the equilibrium log-price vector is*

$$p(s) = -\mathcal{L}(\alpha)\big(\varepsilon(s) + a\big) + \mathcal{L}(\alpha)\,\chi(\alpha)\,e(s), \tag{28}$$

where $\chi(\alpha) = (\chi_i(\alpha_i))$ is the vector of net cross-border exposures, propagated through the network by the Leontief inverse.

The proof of the lemma above is straightforward. With cross-border conversion the unit cost is $\theta_i = -\varepsilon_i - a_i + \sum_j \alpha_{ij}(p_j + c_{ij}e) = -\varepsilon_i - a_i + \sum_j \alpha_{ij} p_j + \chi_i e$, where $\chi_i = \sum_j \alpha_{ij} c_{ij}$. Competitive pricing $p_i = \theta_i$ leads to the system in matrix notation $(I - \alpha)p = -(\varepsilon + a) + \chi e$. Inverting by $\mathcal{L}(\alpha)$ gives (28). Note that since the exchange rate is pinned down by the [Backus and Smith, 1993] condition, rather than from market clearing, it does not enter the fixed point iteration that characterizes the equilibrium.

The Stage-1 first-order condition for a home firm's foreign link is $\partial a_i/\partial \alpha_{ij} = \mathbb{E}[p_j + e] + \mathrm{Cov}(p_j + e, \lambda_i)$, and the wedge on that link splits into a node component and a currency component. This leads to the next proposition.

Proposition 5 (The foreign currency wedge)**.** *For a home firm i sourcing from a foreign node j, the wedge is:*

$$\tau^F_{ij} = \underbrace{\mathrm{Cov}(p_j, \xi_i)}_{\textit{node risk}} + \underbrace{\mathrm{Cov}(e, \xi_i)}_{\textit{currency wedge}}, \tag{29}$$

where $\xi_i = \lambda_i^{\text{eff}} - \lambda^c$ is the firm's idiosyncratic kernel gap. All previous results on precaution,

misallocation and retrenchment carry similarly.

The proof of 5 is in Appendix A.

The currency wedge is where the open economy sharpens the mechanism. The worker attached to a foreign-sourcing firm bears exchange-rate risk through the wage: a depreciation raises the firm's imported-input costs, lowers its profit, and lowers the wage, so the worker's marginal utility is high exactly when e is high. The fund that prices equity is internationally diversified and does not bear this comovement. Hence the worker's kernel loads on e more than the diversified kernel does, the currency wedge is positive, and the firm retrenches from foreign suppliers.

Corollary 1 (Retrenchment)**.** *If the fund pricing equity is internationally diversified while the firm-attached worker is not, then* $\mathrm{Cov}(e, \xi_i) > 0$*, so the incomplete-market network places lower shares on foreign suppliers than the efficient network and severs foreign links whose currency wedge is large enough. The* ***home bias is inefficient*** *because the currency risk it avoids is diversifiable at the level of the fund but not at the level of the worker.*

5 Quantitative Analysis: The Tariff Counterfactual

5.1 Numerical Methods and Calibration

We now provide a quantitative investigation of a counterfactual policy experiment in our theoretical environments to assess changes in the wedge distribution and related costs. We run the simulations in both versions of the model: the one with an exogenously given pricing kernel, and the full general equilibrium version, both featuring currency risk. The comparison between the two also lets us discuss the role of general equilibrium effects that go beyond changes in input prices through market clearing but spill over into workers'/households' asset accumulation. We use the model to assess the response to a counterfactual change in tariff policy that involves both a change in the level and the uncertainty component of the process. This lets us validate the model's responses and patterns against the event study analysis we presented in the empirical session. We devote particular attention to the calibration of the model parts directly related to the transmission channels of interest (the production structure, the pricing kernel, and the productivity distribution) using micro data on US firms importing from abroad (OECD Inter-Country Input-Output Database to calibrate the network matrix and NBER-CES manufacturing data and OECD-STAN to calibrate the firms' productivity distributions and processes). Next, we describe the simulation method and briefly account for the calibration. Detailed explanations of the calibration are reported in Appendix B.5. Before discussing the results we briefly present the solution method and briefly describe the calibration strategy.

One note worth mentioning, before describing the solution method and calibration, that to make the model's calibration more realistic one of the inputs is calibrated as physical capital. The choice of capital is made in exactly the same way as all other inputs: shares chosen before the shock realized, and quantities chosen after. Hence, there is no material change in the structure of the model, except that the capital share also bears a precautionary motive with respect to future productivity, tariff and currency risk.

Numerical Methods. I start by describing the numerical method for the model with heterogeneous firms interacting in an input-output production relation with an exogenous kernel. In this case, the equilibrium is the fixed point of the market-clearing system: one clearing condition per input market, eight in all, with labor redundant by Walras' law. The shares and prices are pinned jointly by this system. Given prices, the sourcing first-order condition of Lemma 2 returns the origin shares; given the shares, the markup price system returns the prices, since the network and the productivity shifter both depend on the shares through the Leontief inverse $\mathcal{L}(\alpha)$. These two conditions are mutually dependent and solved together. Everything else is block-recursive rather than a further fixed point: under the constant returns, the same inversion that returns prices returns the Domar weights $\omega^\top = \beta' \mathcal{L}(\alpha)$, and the price covariances, the kernel covariance, and the wedge are then linear algebra on the node system, which comprises the eight home sectors and the Chinese and safe-bloc counterparts of the tradable sectors. Goods-market clearing is thus embedded in the price block, and the market-clearing fixed point reduces to the sourcing condition against the price system.[16]

The general-equilibrium model adds one equilibrium fixed point to the heterogeneous-firm solution. It is the iteration between households' saving decisions and the interest rate that clears the asset market. Taking the interest rate as given, each household's saving policy solves its consumption-saving problem on the state (assets, firm TFP). Aggregating those policies over the stationary wealth distribution gives aggregate asset demand. The asset-market-clearing condition sets the interest rate at which that demand meets asset supply, here a target liquid-wealth-to-income ratio. In turn, the interest rate feeds back into the policies. The saving policy at each candidate rate is solved by the endogenous grid method of Carroll [2006].[17] In the model, there are two states (assets, firm TFP) and an occasionally borrowing constraint.[18]

The network fixed point, the loop between sourcing shares and prices with goods-market clearing embedded, is unchanged. The two fixed points are solved sequentially: the household

[16] In practice we solve this by iterating the two conditions to convergence, on the order of thirty steps: given a network, the intraperiod block is a single matrix inversion with no iteration inside it, and the outer loop updates the sourcing shares against the resulting prices until the shares stop moving.

[17] The method places a grid on end-of-period assets, evaluates expected next-period marginal utility at each savings level, and inverts the marginal-utility function to obtain current consumption and the beginning-of-period assets consistent with it.

[18] The policy function is interpolated back onto the fixed asset grid to iterate the wealth distribution.

loop is solved first and produces the degree of uninsurability ϕ. The latter then enters the network loop, which is solved next.

The simulations assess two policy counterfactuals represented by changes in tariff regimes in 2018 and 2025. We change both the level and the variance of the shock (the latter to proxy changes in uncertainty) either separately or jointly. In response to the shock to the permanent level, we compare the two stochastic steady states solved with the fixed points described above, and the transitional dynamics. In response to the variance shock, we compute an impulse response function.

Calibration. We describe the calibration briefly here and provide more details on how it is settled in Appendix B.5. First, we use data from OECD Inter-Country Input-Output to calibrate the baseline network shares. The second part uses the NBER-CES Manufacturing Industry Database and the OECD STAN database to calibrate the productivity variance-covariance matrix. BIS data are used to estimate exchange-rate risk, while estimates from the literature (see Guiso et al. [2005] and Storesletten et al. [2004]) are used to calibrate the parameter that determines the extent to which firms' risk passes through to workers through limited-commitment contracts. The calibration also parametrizes the functional form of the productivity shifter, $a_i(\alpha)$; since the latter also determines changes in the optimal input shares, we choose its functional form to obtain smooth changes in the input-output shares. The firms' TFP process is an AR(1) with persistence of 0.95 and is leptokurtotic. The tariff process is described below in the next paragraph. Together with the calibration of ϕ, which regulates how much of the firm's risk passes through to workers, this yields a worker earnings process with the properties estimated in past literature (see Guiso et al. [2005] or Storesletten et al. [2004]). Most importantly, since *phi* depends on the extent to which firms' participation constraint in the limited-commitment contract binds across states, the resulting earnings process is endogenously countercyclical.

The tariff process. The policy experiment treats the log tariff on a foreign link as a random variable,

$$t_{ij} = \bar{t}_{ij} + \theta_{o(j)}\, T + u_{ij}, \tag{30}$$

a common tariff-regime factor T with origin-specific loadings and an idiosyncratic term. Four features of the 2025-2026 episode discipline it. The mean $\bar{t}_{ij}$ is the realized schedule of Section 5.1, about 23 percent on China and below 5 percent on the allies. The loadings concentrate the regime factor on China, $\theta_{\text{China}} = 1$ against $\theta_{\text{safe}} = 0.3$ and zero at home, matching the fact that the actions targeted China. The regime standard deviation σ_T is set so the implied volatility of the China tariff matches the size of the episode's swings; since it enters the two exercises at different points, we use $\sigma_T = 0.20$ in the reduced-form heterogeneous-firm exercise and $\sigma_T = 0.35$ at

the general-equilibrium steady state, with the variance impulse response peaking at 0.60. The idiosyncratic term has standard deviation $\sigma_u = 0.05$. The remaining parameter is the correlation of the tariff regime with the real exchange rate: the realized episode, tariffs rising as the dollar depreciated, makes $\mathrm{Cov}(T, e) > 0$ and the tariff a risk, which we take as the baseline, while a deficit-trigger rule would make it negative and the tariff a partial hedge, which we carry as the alternative.

5.2 Results

I present the results of the tariff-shock counterfactual starting from the model with an exogenous kernel: first, the impact of the shift on the moments that characterize the stochastic steady state, and then impulse response analysis. At the end, we comment on the distributional consequences of the tariff and how they match the empirical evidence presented earlier. To make the two models comparable, they are simulated as nested versions of each other: in the first case the parameter, ϕ, that makes firm risk pass through onto households is fixed. Throughout, the wedges and losses we report are those on the links to China, the tariffed origin we track; the aggregate welfare cost, however, is broader, since it must sum the misallocation across all tariffed origins in the exogenous-kernel model and add the consumption-risk cost of the general-equilibrium closure. The welfare costs are presented first for the China-links only and then they are scaled up to the full tariff-exposure countries.

We start with the model with the exogenous kernel. Table 2 reports the two stochastic steady states of the heterogeneous-firm model, holding the degree of uninsurability fixed at $\phi = 0.10$. The move raises the tariff level and its variance together. The initial steady state has the tariff at zero variance, the final one has positive variance (calibrated as before). The top three rows report the foreign and China input shares and the output loss from misallocation; the last four decompose the distortion on China's sourcing into the policy wedge (the mean tariff), the productivity wedge (precaution against China's productivity risk), the tariff-risk wedge (precaution against the tariff being uncertain), and the currency wedge (precaution against the exchange-rate risk of foreign inputs), each pushing firms toward China if negative and away if positive.

The tariff level does the reallocation. The higher mean tariff pushes firms out of China, so the foreign input share falls from 0.0279 to 0.0265, the China share from 0.0064 to 0.0055, and the policy wedge on China rises almost tenfold, from 0.014 to 0.154 log points. This is the largest change in the table and the de-risking the analytical results predict. In absolute terms the reallocation on the China links is small, since the calibrated China sourcing share is small; this is a statement about China alone, and we return below to the full tariffed exposure.

The tariff variance induces the precaution effect, through the risk wedges. In the initial stochastic steady state the tariff-risk wedge is zero and the only active risk wedges are the productivity and currency wedges. The productivity wedge is small and negative, 0.00001, and would if anything pull firms toward China, since China is low-beta: its productivity moves little with the rest of the economy, so sourcing there is a hedge and precaution mildly rewards a hedge. Switching on the tariff variance, σ_T, in the move to the final steady state creates a tariff-risk wedge of 0.00027, positive because the tariff uncertainty is concentrated on China,[19] so it flips precaution from favoring China to penalizing it, and the misallocation loss rises with it, from 4.2×10^{-9} to 1.3×10^{-8} of output on the China links. This confirms the role of tariff uncertainty for misallocation.

The same undiversified loading also prices a currency wedge on the foreign links, which the decomposition in Table 2 reports separately. It has two parts: the exchange-rate variance $\psi \operatorname{Var}(e)$, present whenever foreign inputs are held, and the tariff-currency covariance $\psi\theta \operatorname{Cov}(T, e)$, which the tariff variance switches on. Because the recent tariffs rose while the dollar depreciated, the $\operatorname{Cov}(T, e)$ is positive, so the currency channel adds to rather than hedges the tariff risk. In the final steady state the currency wedge on China, 0.00062 log points, is the largest of the risk wedges, about twice the tariff-risk wedge and far above the productivity wedge.

Figure 4 turns to the dynamics. The left panel is the transition to the permanent level change (variance held off) for the capital stock. Given the calibration the tariff's effect on aggregate productivity and hence on the steady-state capital stock is minor, so the capital response is a fraction of a percent. The middle panel is the response to a temporary spike in the tariff variance σ_T at the 2026 steady state, with the China wedge split into its regular part (productivity plus tariff-risk) and its currency part: both jump and decay with the variance, and the currency wedge is higher then the regular one throughout. The remaining panels show the risk-misallocation loss, which jumps and decays with the wedge, and the China share. The latter barely moves since the mean tariff dominates the sourcing decision even as the risk shifts. The variance moves the price of risk, both for currency and tariff risk, so the wedges and the loss respond sharply while the share does not.

Four things follow from the table and the figure together. First, the tariff level is the larger mover of the quantity of reallocation, more than any risk term, yet the reallocation it produces is small on the China links. Below, we will reconsider the numbers by scaling up the exposure to all tariffed countries. Second, precaution against productivity risk is small and favorable, since

[19] The calibration delivers $\theta_{China} = 1$ and $\theta_{safe} = 0.3$, since the 2025-2026 tariff actions were aimed at China, with the allies treated more lightly through carve-outs and lower rates.

China is low-beta. Third, precaution against tariff uncertainty is positive and concentrated on China, and it flips the sign of precaution once the variance is on. Fourth, the currency wedge is the largest risk wedge on China, driven by the positive tariff-currency correlation, so exchange-rate risk is the dominant precautionary distortion here. Across all of them the pattern is the same: the risk wedges move the price of risk and the loss, while the quantity of reallocation stays small. In this respect the model is consistent with the data: the level drives the reallocation the data can see, while the risk wedges drive the cost, which we can only read through the lens of the model.

Table 2

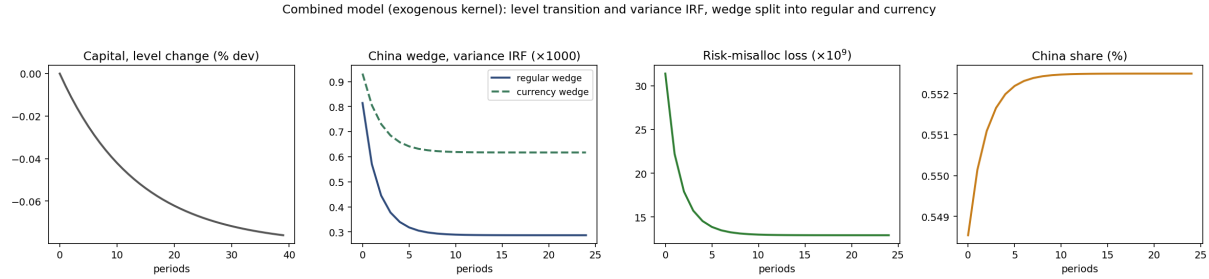
Two stochastic steady states: aggregates and the China wedge decomposition, with the currency wedge shown separately.

	SS0 (pre-2025)	SS1 (2026)	change
Domar-weighted foreign share	0.0279	0.0265	−0.0014
Domar-weighted China share	0.0064	0.0055	−0.0009
risk-misallocation loss	4.2e-09	1.3e-08	×3
China policy wedge (mean tariff)	0.0138	0.1540	+0.1402
China productivity wedge	0.00001	0.00001	0
China tariff-risk wedge	0.00000	0.00027	+0.00027
China currency wedge	0.00018	0.00062	+0.00044

Notes: Heterogeneous-firm model with capital, $\phi = \psi = 0.10$. The regular wedges are the policy wedge (mean tariff), the pure productivity wedge $\phi[\mathcal{L}\Sigma]$, and the pure tariff-risk wedge (systematic, idiosyncratic, and cycle terms); the currency wedge is $\psi \operatorname{Var}(e)$ plus the tariff-currency covariance $\psi\theta \operatorname{Cov}(T, e)$. Wedges are China-import-weighted, in log points.

Figure 4

Transitional Dynamics (left panel) and Impulse Response Function to a Tariff Shock (to the level and the variance). Model with Exogenous Kernel



Notes: Left: capital impulse response to the permanent level change (with the variance of the tariff process sett to zero). Middle panels: China tariff-risk wedge in response to a tariff-variance shock and the risk-misallocation loss in response to the same shock. Right panel: response of the China share.

The Tariff Shock in the General Equilibrium Model. Next, we discuss the same counterfactual in the full general equilibrium model, which includes costs for consumers/workers exposed to the firms that are retrenching. Before discussing the counterfactual, we verify the earlier claim on the endogenous countercyclicality of the earnings process. Solved numerically, ϕ is zero across the insured region and rises to about 0.77 in the worst firm state. This also implies that the risk-misallocation wedge is a tail phenomenon that switches on in downturns,

when misallocation is most costly. The tariff experiment is the same as before; here we report only what the endogenous kernel adds.

Table 3 compares the two steady states. The degree of uninsurability is about 0.23 in the normal regime and rises to about 0.55 in the high-tariff-variance regime, as firms are pushed into the region where the participation constraint of the limited-commitment contract binds. The foreign and China shares move much as in the exogenous-kernel model, since the tariff level drives them. What differs is the cost: the risk-misallocation loss rises by roughly an order of magnitude between the two steady states, because the endogenous kernel amplifies every wedge by raising the degree of uninsurability in the worst states. This amplification applies to the currency wedge as well, which the table again reports separately: the currency wedge is the larger part of the China wedge, and it grows with ϕ from 0.00041 to 0.00337 log points, so the exchange-rate channel is amplified by the endogenous kernel in the same way as the tariff channel.

Figure 5 traces the response to a temporary spike in the tariff variance at the 2026 steady state. The variance shock raises ϕ, which then amplifies the China wedge, both its regular and its currency parts. As a consequence firms de-risk more sharply than under the exogenous kernel. There is therefore a countercyclical amplification: an aggregate risk shock raises uninsurability, which raises the wedge. This goes on top of the direct effect of the shock on the tariff bill and the exchange-rate exposure.

Table 3

General-equilibrium tariff counterfactual with the currency wedge shown separately.

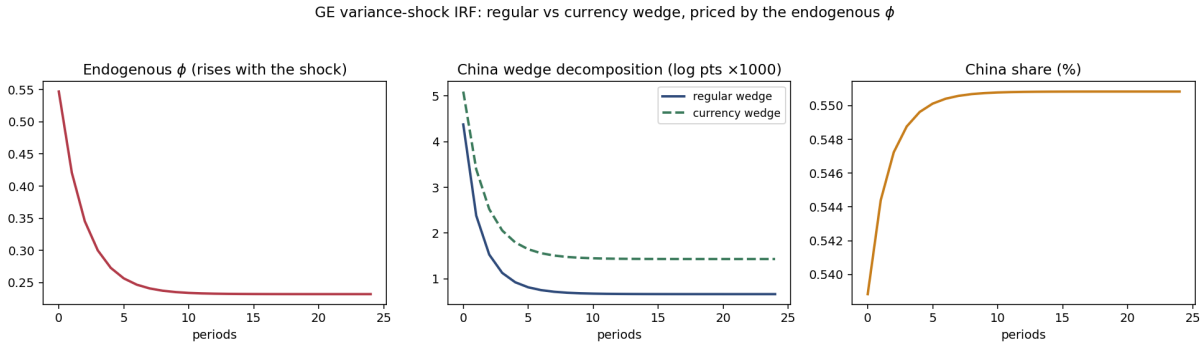
	SS0 (pre-2025)	SS1 (2026)	change
Degree of uninsurability ϕ	0.232	0.547	+0.315
Domar-weighted foreign share	0.0278	0.0263	−0.0015
Domar-weighted China share	0.0064	0.0055	−0.0009
China regular wedge (log pts)	0.00003	0.00155	+0.00152
China currency wedge (log pts)	0.00041	0.00337	+0.00296
Risk-misallocation loss	2.3e-08	3.9e-07	×17

Notes: Two stochastic steady states, the degree of uninsurability ϕ endogenous in each. The China wedge is decomposed into the regular part (productivity propagation $\phi[\mathcal{L}\Sigma]$ plus the tariff systematic, idiosyncratic, and cycle terms) and the currency part ($\psi \operatorname{Var}(e)$ plus the tariff-currency covariance $\psi\theta \operatorname{Cov}(T, e)$), both priced by the same endogenous kernel ($\psi = \phi$) and weighted by China import value.

The Decomposition of Welfare Cost: Misallocation versus Consumption Costs. The welfare cost of tariff uncertainty, expressed as a consumption equivalent, has two components: the network misallocation attributable to the tariff variance, and the additional consumption risk that workers bear when ϕ rises. The second dominates by orders of magnitude. Tariff uncertainty is costly in this model mainly because it raises uninsurable consumption risk for the workers employed in the exposed firms. The network distortion is second order. On

Figure 5

General-equilibrium variance-shock impulse response.



Notes: Left: the endogenous degree of uninsurability ϕ rises with the tariff-variance shock and decays. Middle: the China tariff wedge, amplified by the higher ϕ. Right: the China input share, which dips as firms de-risk and recovers as ϕ falls back.

the China links alone the aggregate cost is modest, since the calibrated United States sourcing from China is small. But China is only one tariffed origin, so the economy-wide cost must add the misallocation on every tariffed link and the consumption-risk cost of the general-equilibrium closure, which we do next.

Scaling up: the Full Tariff Exposure. The numbers above are computed for the United States' exposure to China, whose Domar weight is small, about 0.006. The 2025 tariffs fall on a far broader set of origins. Because the consumption-risk cost scales with the tariffed input exposure, applying the same tariff uncertainty to the full foreign input exposure, a Domar weight of about 0.026 and roughly five times the China exposure, raises the consumption-equivalent cost from 0.065 percent to about 0.31 percent of consumption, while the network-rewiring loss stays an order of magnitude smaller. Table 4 collects these figures. The scaling is linear in exposure and holds the degree of uninsurability fixed, so it understates the cost if a broader tariff also raises aggregate profit risk and hence ϕ.

Table 4

Consumption-equivalent welfare cost of tariff uncertainty: China links and the full tariffed exposure.

	China only	Full tariffed exposure
Tariffed Domar exposure	0.006	0.026
Network-rewiring cost (% output)	0.00003	0.00018
Consumption-risk cost (% consumption)	0.065	0.31
Total welfare cost (% consumption)	0.065	0.31

Notes: General-equilibrium consumption-equivalent cost of the tariff-variance shock. The consumption-risk cost dominates; the network-rewiring cost is the misallocation loss attributable to the tariff variance. The full-exposure column applies the China-intensity tariff uncertainty to the full foreign input exposure and scales the consumption-risk cost by the exposure ratio, holding the degree of uninsurability fixed.

The Distributional Consequences of the Tariff. The model features also a wage distribution and households heterogeneity and it is useful also to assess the distributional con-

sequences of the tariff shock. Table 5 splits the consumption-equivalent cost by wealth quintile, normalized to the richest. The cost is steeply regressive: the poorest quintile bears about twenty-seven times the cost of the richest, and the burden is concentrated in the bottom two quintiles, while the top three, who can self-insure, bear little. The mechanism is driven by the marginal propensity to consume. The poorest households are almost entirely hand-to-mouth, with a marginal propensity to consume near 0.82, so they pass the tariff-induced income shock almost fully into consumption, whereas the richest, with a propensity near 0.18, smooth it through savings. The same limited-commitment friction that makes uninsurability countercyclical in the aggregate thus makes its incidence fall on the households least able to absorb it, so the aggregate cost is borne mainly by liquidity-constrained low-wealth workers.

Table 5

Distributional welfare cost of tariff uncertainty by wealth group.

Wealth group	Hand-to-mouth (%)	Mean MPC	Welfare cost (× richest)
Q1 (poorest)	91	0.82	26.9
Q2	0	0.36	7.4
Q3	0	0.19	1.1
Q4	0	0.18	1.0
Q5 (richest)	0	0.18	1.0

Notes: The cost is the consumption-equivalent cost of the extra consumption risk that tariff uncertainty imposes, normalized to the richest quintile. The pattern is driven by the marginal propensity to consume: hand-to-mouth households pass the shock into consumption, wealthier households smooth it.

Model and data. It is useful to take stock of how the model and the evidence line up, and where each is informative. Both speak to the tariff's effect on the reallocation of sourcing across firms and origins, and they agree that this reallocation is small: the firm-level event study finds only a modest retreat from China, much of which reflects firms' pre-existing exposure rather than the tariff itself (Appendix B.1), and the model likewise produces only a small reallocation from the tariff level, the China input share moving by fractions of a percentage point. More than the magnitude, the two agree on the channel. The anticipation design shows the retreat is broad, falling on untaxed Chinese goods as much as on taxed ones and coinciding with rising trade-policy and geopolitical uncertainty. This is the sign of a precautionary withdrawal from China as a risky origin rather than a mechanical response to the tax. This is exactly the model's mechanism, and it matches the model's decomposition, in which the tariff level moves quantities only modestly while the precautionary response to the tariff's uncertainty carries the wedge. The empirical event study on import shares can only measure the quantity of de-risking but not its welfare consequence. The structural model, that matched the empirical patterns, also supply the welfare calculations, as well as their decomposition. The losses from the misallocatioon within the network is second order, while the one in the uninsurable consumption risk that the countercyclical contract passes to workers, is large. Furthermore, a cost that is modest on

the China links is sizable once summed over all tariffed origins. At last, the losses are steeply regressive, borne mainly by the poorest, liquidity-constrained households.

6 Conclusions

The change in the geopolitical landscape has compelled economists to reconsider the role and costs of global shocks on the global value chain and the macroeconomy. A full-fledged assessment of such costs requires two key ingredients missing from models currently available in the literature: a risk that firms cannot insure away and the possibility for firms to retrench and reactivate trade links elsewhere. This paper proposes a model that advances on both fronts: firms, facing uninsurable productivity risk, interact in input-output relations that they choose to reconsider every period. Through the lens of the model, I show how the combination of uninsurable risk and endogenous retrenchment drives efficiency costs. I use the model to examine a counterfactual tariff policy calibrated on the 2025 episode: the main takeaway is that the extent of reallocation out of the countries hit the most by the tariff is small, but its misallocation and welfare costs may be high.

A Derivations and Proofs

Derivations: cost minimization problem in (5). The Lagrangian of (5) is $L_i + \sum_j P_j X_{ij} - \eta_i(F(\alpha_i, L_i, X_i) - Q_i)$, where η_i is the Lagrange multiplier. The first-order conditions are: $1 = \eta_i \partial F/\partial L_i$ and $P_j = \eta_i \partial F/\partial X_{ij}$. Under the Cobb-Douglas production function, (4), $\partial F/\partial X_{ij} = \alpha_{ij} F/X_{ij}$ and $\partial F/\partial L_i = (1 - \sum_j \alpha_{ij})F/L_i$, so $P_j X_{ij} = \alpha_{ij}\eta_i F$ and $L_i = (1 - \sum_j \alpha_{ij})\eta_i F$. Summing all those input costs shares, and since $F = Q_i$, implies that the multiplier η_i is equivalent to the unit cost of production, Θ_i. We use this equivalence $\eta_i = \Theta_i$ since now on. Substituting the optimal input demands into the production function, (4), and solving for η_i gives $\Theta_i = e^{-\varepsilon_i} A_i (\alpha_i)^{-1} \prod_j P_j^{\alpha_{ij}}$ after using the normalization ζ, and cost shares equal α_{ij} and $1 - \sum_j \alpha_{ij}$. Taking logs gives (6).

Proof of Lemma 1. This provides the analytical derivations of the equilibrium variable, prices and quantities, which under the knife-edge conditions of constant return to scale and jointly Gaussian shocks, prices and kernels, features a closed form solution. We rely on the block recursivity and solve prices first.

Solution to equilibrium prices. Using markup pricing and the unit cost of production, (6), and using logs give $p_i = \log \mathcal{M} - \varepsilon_i - a_i + \sum_j \alpha_{ij} p_j$, which can be written as: $(I - \alpha)p = -(\varepsilon + a - \log \mathcal{M}\,\mathbf{1})$. The markup is a constant, so it shifts the price level, but leaves every covariance below unchanged. Since $\sum_j \alpha_{ij} < 1$, the nonnegative matrix α has row sums below one, so its spectral radius is below one. Hence $I - \alpha$ is invertible, the Leontief matrix, $\mathcal{L} = (I - \alpha)^{-1} = \sum_k \alpha^k$, converges, and the solution to the price system, equation (8), is unique. Normality of the price distribution follows from that fact that prices are affine (hence linear) functions of ε (see here: $(I - \alpha)p = -(\varepsilon + a - \log \mathcal{M}\,\mathbf{1})$), and the moments can be computed from the price system (8).

We now solve quantities. Multiply the market-clearing condition $\sum_i X_{ij} + C_j = Q_j$ by P_j. We can now substitute the cost-minimizing expenditures $P_j X_{ij} = \alpha_{ij}\Theta_i Q_i = \alpha_{ij} P_i Q_i/\mathcal{M}$, using $P_i = \mathcal{M}\Theta_i$, so revenue-based shares carry the uniform factor $1/\mathcal{M}$ that cancels in the relative Domar weights. Household demand is given by $P_j C_j = \beta_j$, which follows from Cobb-Douglas preferences and the budget $\sum_j P_j C_j = 1$. Given $\omega_j \equiv P_j Q_j$ (which is the sector j sales or the revenues) this gives $\omega_j = \beta_j + \sum_i \alpha_{ij}\omega_i$, which we can write in vector notation as $\omega = \beta + \alpha^\top \omega$. The solution is $\omega^\top = \beta^\top \mathcal{L}$ and is unique for the same argument above of the spectral radius. Note that the equation that determines the sales, ω_i, depends only on α and β, hence they are nonstochastic. Labor clears by Walras' law: $\sum_i L_i = \sum_i (1 - \sum_j \alpha_{ij})\omega_i = \omega^\top(\mathbf{1} - \alpha\mathbf{1}) = \beta^\top \mathcal{L}(I - \alpha)\mathbf{1} = \beta^\top \mathbf{1} = 1$.

At last we solve for output in this economy. Real output is the Cobb-Douglas consumption

aggregate $Y = \prod_j C_j^{\beta_j}$. Taking logs:

$$y = \sum_j \beta_j c_j. \tag{31}$$

Household demand, $C_j = \beta_j / P_j$, in logs is $c_j = \log \beta_j - p_j$, so

$$y = \sum_j \beta_j \big(\log \beta_j - p_j \big) = \text{const} - \beta' p, \tag{32}$$

and dropping the constant $\sum_j \beta_j \log \beta_j$, a units normalization, gives $y = -\beta' p$. Using the price system $p = -\mathcal{L}(\varepsilon + a)$ and the Domar weights $\omega' = \beta' \mathcal{L}$ derived above, we get:

$$y = -\beta' p = \beta' \mathcal{L}(\varepsilon + a) = \omega'(\varepsilon + a). \tag{33}$$

□

Proof of Lemma 2. Since only unit cost depends on α_i, maximizing expected discounted profit in (7) is equivalent to minimizing the kernel-weighted cost of producing Q_i,

$$\min_{\alpha_i \in \mathcal{A}_i} \ \mathbb{E}\big[\Lambda_i \, Q_i \, \Theta_i(\alpha_i, P)\big] = \min_{\alpha_i \in \mathcal{A}_i} \ \mathbb{E}\big[e^{\lambda_i + q_i + \theta_i}\big], \tag{34}$$

where lowercase letters are logs, $\lambda_i = \log \Lambda_i$, $q_i = \log Q_i$, and $\theta_i = \log \Theta_i$.

We start by computing the expectation of the log. Minimizing a positive quantity is equivalent to minimizing its logarithm. The exponent $\lambda_i + q_i + \theta_i$ is jointly normal, since λ_i is log-normal by assumption and θ_i is affine in the Gaussian prices. For a normal random variable x, $\log \mathbb{E}[e^x] = \mathbb{E}[x] + \frac{1}{2} \operatorname{Var}(x)$, so

$$\log \mathbb{E}\big[e^{\lambda_i + q_i + \theta_i}\big] = \mathbb{E}[\lambda_i + q_i + \theta_i] + \tfrac{1}{2} \operatorname{Var}(\lambda_i + q_i + \theta_i). \tag{35}$$

We now drop terms which do not depend on α_i). The kernel λ_i is taken as given, so its mean and variance do not depend on α_i. Sales are nonstochastic and were pinned down by market clearing as functions of only α and β (see proof to Proof of Lemma 1). Hence q_i does not depend on α_i and drops from the risk-adjusted problem. So $\mathbb{E}[\lambda_i + q_i]$ drops out. Expanding the variance of the three-term sum:

$$\operatorname{Var}(\lambda_i + q_i + \theta_i) = \operatorname{Var}(\lambda_i) + \operatorname{Var}(q_i) + \operatorname{Var}(\theta_i) + 2\operatorname{Cov}(\lambda_i, \theta_i) + 2\operatorname{Cov}(\lambda_i, q_i) + 2\operatorname{Cov}(q_i, \theta_i). \tag{36}$$

Only θ_i depends on α_i, so most of the variance and covariance terms drop out. The two terms that involve θ_i are the own variance $\operatorname{Var}(\theta_i)$ and the covariance $\operatorname{Cov}(\lambda_i, \theta_i)$. The own variance is second order in the shocks, while the first-order risk term, linear in the network, is the covariance, which is: $\frac{1}{2} \cdot 2 \operatorname{Cov}(\lambda_i, \theta_i) = \operatorname{Cov}(\lambda_i, \theta_i)$. Collecting all terms that depend on α_i:

$$\text{objective} = \mathbb{E}[\theta_i] + \operatorname{Cov}(\theta_i, \lambda_i) + \tfrac{1}{2} \operatorname{Var}(\theta_i) + \text{const}. \tag{37}$$

Nest, we substitute the unit cost. From (6), $\theta_i = -\varepsilon_i - a_i(\alpha_i) + \sum_j \alpha_{ij} p_j$, whose only

stochastic, α_i-dependent part is $\sum_j \alpha_{ij} p_j$. Hence:

$$\mathbb{E}[\theta_i] = -a_i(\alpha_i) + \sum_j \alpha_{ij}\, \mathbb{E}[p_j] + \text{const}, \qquad \text{Cov}(\theta_i, \lambda_i) = \sum_j \alpha_{ij}\ \text{Cov}(p_j, \lambda_i). \tag{38}$$

Substituting, and discarding constants and the second-order own-variance term $\frac{1}{2}\text{Var}(\theta_i)$, the cost to minimize is:

$$-a_i(\alpha_i) + \sum_j \alpha_{ij} \big(\mathbb{E}[p_j] + \text{Cov}(p_j, \lambda_i)\big). \tag{39}$$

Minimizing the previous line is equivalent to maximizing its negative, which gives: (9):

$$\alpha_i^* \in \arg\max_{\alpha_i \in \mathcal{A}_i}\ a_i(\alpha_i) - \sum_j \alpha_{ij}\, \mathcal{R}_{ij}, \qquad \mathcal{R}_{ij} \equiv \mathbb{E}[p_j] + \text{Cov}(p_j, \lambda_i). \tag{40}$$

The object $\mathcal{R}_{ij}$ is the risk-adjusted price of input j: its expected log price plus a penalty for how that price covaries with the firm's own kernel. □

Proof of Proposition 4. In Domar-weight space the efficient allocation satisfies: $\nabla\mathcal{W}(\omega^*) = \nabla\bar{a}(\omega^*) + \mu - (\rho-1)\Sigma\omega^* = 0$, and $\nabla^2\mathcal{W} = \nabla^2\bar{a} - (\rho-1)\Sigma = \mathcal{H}$, negative definite since $\bar{a}$ is concave, Σ positive definite, and $\rho \geq 1$. The firms' wedges leave a residual gradient at the distorted allocation, $\nabla\mathcal{W}(\omega^{IM}) = -\mathcal{T}$, which defines the aggregate wedge $\mathcal{T}$. Linearizing the gradient around ω^* gives $\mathcal{H}(\omega^{IM} - \omega^*) = -\mathcal{T} + o(\|\mathcal{T}\|)$, hence:

$$\omega^{IM} - \omega^* = -\mathcal{H}^{-1}\mathcal{T} + o(\|\mathcal{T}\|). \tag{41}$$

Since ω^* maximizes $\mathcal{W}$, a second-order expansion and substitution of this displacement give:

$$\mathcal{W}(\omega^*) - \mathcal{W}(\omega^{IM}) = -\tfrac{1}{2}\,(\omega^{IM} - \omega^*)^\top \mathcal{H}\,(\omega^{IM} - \omega^*) + o(\|\mathcal{T}\|^2) = \tfrac{1}{2}\,\mathcal{T}^\top\big(-\mathcal{H}^{-1}\big)\mathcal{T} + o(\|\mathcal{T}\|^2), \tag{42}$$

positive since $-\mathcal{H}^{-1}$ is positive definite. The discrete term $\Delta^{\text{ext}} \geq 0$ collects the losses from severed links, each non-negative because it removes a strictly optimal interior margin.

Block-diagonal specialization. Hold the general-equilibrium response of Domar weights and prices fixed, so that only firms' own shares vary. The productivity contribution is $\omega^\top a(\alpha) = \sum_i \omega_i a_i(\alpha_i)$, so the own-shares block of $\nabla^2\mathcal{W}$ in firm i is $\omega_i H_i$, that is, $\mathcal{H}$ is block-diagonal with blocks $\{\omega_i H_i\}$. Writing $\delta_i = \alpha_i^{IM} - \alpha_i^* = H_i^{-1}\tau_i$ and substituting into $\frac{1}{2}\sum_i \delta_i^\top(\omega_i H_i)\delta_i$ gives:

$$\mathcal{W}(\alpha^{IM}) - \mathcal{W}(\alpha^*) = \tfrac{1}{2}\sum_i \omega_i\, \tau_i^\top H_i^{-1}\tau_i, \tag{43}$$

negative because H_i^{-1} is negative definite; negating gives the Domar-weighted intensive-margin loss $\frac{1}{2}\sum_i \omega_i\, \tau_i^\top(-H_i^{-1})\tau_i$. The single power of ω_i survives because the curvature block carries $\omega_i H_i$ while the displacement carries H_i^{-1}.

Finally, the loss is zero exactly when $\mathcal{T} = 0$, that is, when every $\tau_i = 0$, which holds under complete markets by Proposition 1, or when the network is fixed so that ω cannot respond, $\omega^{IM} = \omega^*$, by Proposition 2. Measured aggregate productivity is $\omega^\top a(\alpha)$ at the realized network;

since α^{IM} routes inputs away from the cost-minimizing suppliers of α^*, this term falls. See Kopytov et al. [2024] for the aggregation and Baqaee and Farhi [2020] for the general-equilibrium accounting. □

Proof of Proposition 5. The effective price of foreign input j for home firm i, in the home firm's currency, is $p_j + e$. Its risk-adjusted price is $\mathbb{E}[p_j + e] + \mathrm{Cov}(p_j + e, \lambda_i)$, and the wedge relative to the diversified benchmark is $\mathrm{Cov}(p_j + e, \lambda_i - \lambda^c)$, which by linearity of covariance equals $\mathrm{Cov}(p_j, \xi_i) + \mathrm{Cov}(e, \xi_i)$ with $\xi_i = \lambda_i^{\text{eff}} - \lambda^c$. The node component is Definition 4 applied to the foreign supplier's price p_j, itself endogenous through (28); the currency component is $\mathrm{Cov}(e, \xi_i)$. Because a foreign node enters the firm's problem exactly as a domestic node with its price augmented by e, the first-order conditions and the arguments of Propositions 1 to 4 apply verbatim with the wedge (29). □

Proof of Corollary 1. The worker's wage loads on the firm's profit, which for a foreign-sourcing firm falls with a depreciation through the imported-input bill $\chi_i e$, so c_i^w falls and $u'(c_i^w)$ rises when e rises, giving $\mathrm{Cov}(e, \lambda_i^w) > 0$. The fund is internationally diversified, so its kernel is hedged against the exchange rate and $\mathrm{Cov}(e, \lambda^c)$ is small; hence $\mathrm{Cov}(e, \xi_i) = \gamma_i^w \big(\mathrm{Cov}(e, \lambda_i^w) - \mathrm{Cov}(e, \lambda^c) \big) > 0$. By Proposition 3 applied to the foreign wedge (29), a positive τ_{ij}^F lowers α_{ij} and severs the link at the boundary. The inefficiency follows as in Proposition 1, since the currency risk avoided carries no social premium at the diversified fund. □

B Data and Methods for the Motivating Evidence

This appendix documents the three datasets behind Section 2, the construction of the variables, and the econometric detail for the firm-level event study. A replication package with the code and the raw and intermediate data for each of the three parts accompanies the paper.

B.1 Firm-level import data (Panjiva)

The firm-level evidence uses United States waterborne import records distributed by Panjiva, a product of S&P Global Market Intelligence, which are derived from the vessel manifest data collected by United States Customs and Border Protection and released under the disclosure provisions for waterborne cargo. Each record is a shipment, pairing a United States consignee with a foreign shipper. We pulled two windows, 2016 to 2020 for the 2018 tariff event and 2023 to March 2026 for the 2025 event, restricting to shipments whose origin is China or one of a set of comparison origins. The first window contains about 590,000 shipments and 113,000 distinct importers, the second about 566,000 shipments and 99,000 importers.

Table 6 lists the fields we retain, their Panjiva names, what each measures, and why we use it. The unit of analysis is the consignee, identified by the stable Panjiva consignee identifier, which lets us follow the same United States buyer across years; the shipper identifier lets us count a buyer's distinct suppliers; the origin field defines whether a shipment is from China or an alternative origin; the arrival date places the shipment in a year; and value and weight measure its size.

Table 6

Panjiva import fields used in the firm-level event study.

Panjiva field	Measures	Why we use it
`conpanjivaid`	consignee identifier	unit of analysis; follows a buyer over time
`conname`	consignee name	identifies the United States importer
`concountry`	consignee country	confirms the buyer is United States
`constateregion`	consignee state	location, robustness
`shppanjivaid`	shipper identifier	counts a buyer's distinct suppliers
`shpname`	shipper name	identifies the foreign supplier
`shpcountry`	shipper country	supplier location
`shpmtorigin`	country of origin	defines China versus alternative origins
`arrivaldate`	United States arrival date	places the shipment in a year
`billofladingnumber`	bill of lading number	shipment identifier, deduplication
`portofunlading`	United States port	robustness
`valueofgoodsusd`	customs value, dollars	magnitude of the shipment (primary)
`weightkg`	weight, kilograms	magnitude (fallback where value is missing)

From the shipment records we build an importer-by-year panel. For each importer and year we sum shipment magnitude, using the customs value where present and the weight where the value is missing; value is populated for about 96 percent of records and weight for all of them. We then form three outcomes: the China share, the value sourced from China divided by the importer's total; the ally share, the value from the comparison origins (Vietnam, Mexico, Taiwan, India, Thailand) divided by the total; and the count of distinct Chinese suppliers, the number of distinct shipper identifiers with a Chinese origin. China excludes Hong Kong. We keep importers present in every year of a window, giving a balanced panel of about 5,000 importers for the 2018 event and 9,400 for the 2025 event, and we drop the partial year 2026 from the estimation panel. Exposure is the importer's average China share over the pre-shock years, 2016 to 2017 for the first event and 2023 to 2024 for the second.

Robustness Checks and Their Implications. In this paragraph, we examine some robustness checks extending the event-study windows to examine the possible role of pre-trends in the exposure measure. We find evidence for this, which motivates the anticipation design shown in the main draft, geared precisely to identify the role of uncertainty in the dynamics of the exposure measures.

The event study in the main draft follows an established convention for the windows, given the frequency of the data. They show no pre-trend. Still, some trends in exposure may

materialize if we extend the event-study windows. Testing this possibility constitutes not only a robustness check, but, as we argue below, also helps us further nail down the drivers of the retrenchment.

When we extend the event-study window to two years, we find that firm-level China shares mean-revert. We also run two tests. The first is a joint pre-trend F-test on the clean pre-years, those not used to define exposure, which rejects the null of no pre-trend for every outcome in both events. Second, we allow an exposure-specific linear trend and test whether a break survives on top of it, estimating $y_{it} = \alpha_i + \delta_t + \gamma\,(\text{Exposure}_i \times t) + \sum_{\tau \geq \bar{t}+1} \beta_\tau (\text{Exposure}_i \times \mathbf{1}[t = \tau]) + \varepsilon_{it}$, with exposure taken from the earliest clean pre-year. Figure 6 plots the extended-window paths and Table 7 reports both tests. The China-share trend is negative, about twelve to seventeen log points per year, and the post-shock coefficients net of that trend are positive, meaning the decline decelerates after the tariff rather than accelerating. The ally-share breaks net of trend carry the wrong sign. Finally, the supplier-count breaks are insignificant net of trend. One conclusion we draw from this, and explore further in the main draft, is that much of the retrenchment is driven by trends in exposure. For this reason, we move to examine the anticipation design that allows us to unpack the effects of uncertainty and anticipatory expectations.

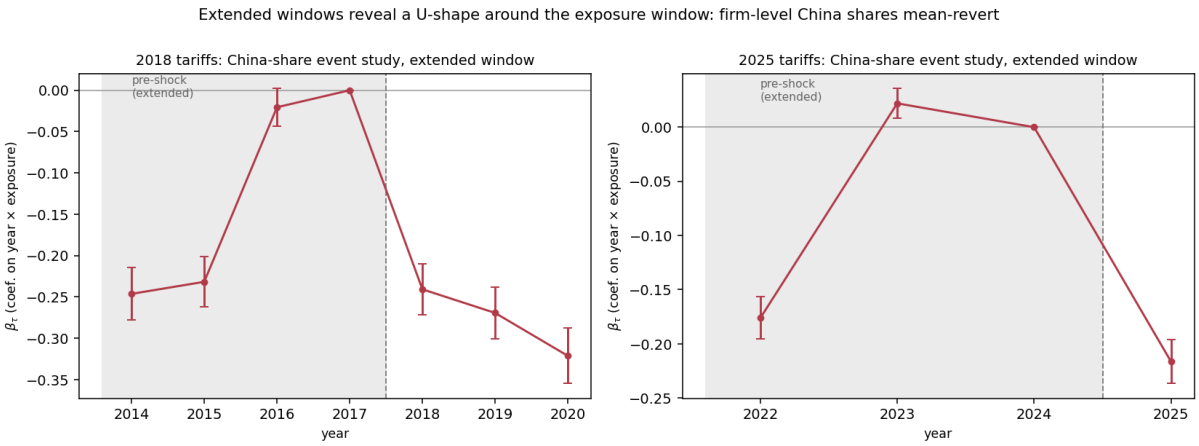


Figure 6

China-share event study on the extended windows, exposure defined as in the main specification. Coefficients β_τ from (2) with 95 percent confidence intervals, clustered by importer; the shaded region is the pre-shock period. In both events the coefficients dip below the baseline in the pre-shock years, return to zero at the years that define exposure, and dip again afterward, the U-shape that signals mean reversion in firm-level China shares rather than a break at the tariff. Source: Panjiva United States waterborne import records.

B.2 Sector-level trade data (Census and Comtrade via ITC Trade Map)

The sector evidence uses quarterly United States goods imports by origin country and product, from the third quarter of 2015 through the second quarter of 2020, at the four-digit level of the Harmonized System. The series come from ITC Trade Map, which republishes the official

Table 7

Firm-level pre-trend diagnostics on the extended windows. Column (1) is the p-value of the joint F-test that the interaction coefficients at the clean pre-shock years (those not used to define exposure) are zero. Column (2) is the estimated exposure-specific linear trend per year. Column (3) is the break in the final post-shock year net of that trend, with its p-value; $*$ marks significance at five percent. Exposure for the trend specification is taken from the earliest clean pre-year. Standard errors clustered by importer.

Event	Outcome	(1) pre-trend F, p	(2) exposure × trend/yr	(3) break net of trend
2018	China share	< 0.001	-0.123	$+0.374^{*}$
2018	Ally share	< 0.001	$+0.037$	-0.098^{*}
2018	Chinese suppliers	0.032	-0.136	$+0.224$
2025	China share	< 0.001	-0.172	$+0.188^{*}$
2025	Ally share	< 0.001	$+0.056$	-0.049^{*}
2025	Chinese suppliers	< 0.001	-0.127	$+0.014$

trade statistics of the United States Census Bureau and the United Nations Comtrade database, accessed through the replication package of the Deutsche Welle data team [Deutsche Welle Data Team, 2020]. To each four-digit product we attach an indicator for placement on one of the United States Section 301 tariff lists, Lists 1 through 4A, using the assignments published by the Office of the United States Trade Representative and distributed in the same package. We aggregate the four-digit products to the eight sectors of the model by Harmonized System chapter, so the sector evidence is on the calibration's partition. For each quarter and product basket we compute the share of United States imports originating in China, Chinese import value over total United States import value in the basket, and the analogous shares for the comparison origins. We omit the two quarters of 2020 from the figures, since the onset of the pandemic distorts trade in a way unrelated to the mechanism.

B.3 Firm-level customer data (Compustat)

The sales-side evidence uses the customer-segment disclosures in Compustat, in which United States public firms report, under the relevant accounting standard, customers accounting for a material share of their sales. We take the historical customer-segment file for 2015 to 2021 and the annual fundamentals file for firm identifiers and locations. Each customer record names a customer and classifies it, and we retain the records classified as company customers, of which there are about seventy-five thousand across roughly forty-seven hundred reporting firms. We identify Chinese customers by name, matching the major Chinese firms that appear (Huawei, Lenovo, Xiaomi, China Mobile, Alibaba, Hon Hai and Foxconn, Geely, and others), and count the reported relationships with a Chinese customer by year. The count falls from about ninety in 2017 to about fifty in 2021. This is the sales-side counterpart to the sourcing evidence: it captures whom firms sell to rather than whom they source from, and the China sample is thin, so it is complementary rather than central. Compustat customer segments contain no supplier

field, so the sourcing side that the model concerns is not observable there, which is why the firm-level sourcing evidence comes from the import records instead.

B.4 Uncertainty indices and the anticipation design

The anticipation design of Section 2.3 uses three publicly available, monthly time-series measures of uncertainty. Trade-policy uncertainty is the index of Caldara et al. [2020], a text-based measure built from the frequency of joint mentions of trade policy and uncertainty in newspaper coverage. Geopolitical risk is the index of Caldara and Iacoviello [2022], constructed analogously from coverage of geopolitical tensions; we use both the overall index and its China-specific component, which isolates coverage of geopolitical risk involving China and is the most relevant to sourcing from that origin. As a placebo we use the monetary-policy category of the economic-policy-uncertainty index of Baker et al. [2016]: monetary-policy uncertainty should not move firms' sourcing away from China, so a null association there checks that the design is not capturing generic uncertainty.[20]

We match the indices to the sourcing outcome at quarterly frequency: each monthly index is averaged within the quarter and standardized over the estimation window, then merged to the quarterly China import share by basket of Appendix B.2. The aggregate contrast (3) reads the tariffed and untariffed China shares against these standardized indices around the Section 301 enactment date.

The firm-level version of the design, which identifies the uncertainty channel rather than only the breadth of the retreat, matches the same standardized indices to the firm-by-quarter Panjiva panel of Appendix B.1. Each importer's pre-episode China share is its exposure, the dose, and the specification interacts that exposure with the standardized index at a chosen lead, with importer and quarter fixed effects and standard errors clustered by importer:

$$\Delta s_{it} = \alpha_i + \delta_t + \beta \left(\text{Exposure}_i \times U_{t-k}\right) + \varepsilon_{it}, \tag{44}$$

where Δs_{it} is the change in firm i's China sourcing share, U_{t-k} is standardized trade-policy uncertainty or China geopolitical risk at lag k, and a lead ($k > 0$) tests whether more-exposed firms cut sourcing before enactment. The quarter effects δ_t absorb the aggregate level of the index, so identification comes from the interaction of the time-varying uncertainty with the cross-section of exposure; the placebo replaces U with monetary-policy uncertainty, and a safe-origin (Mexico and Vietnam) share serves as a reallocation check. This firm-level design requires the licensed Panjiva panel and is provided in the replication package.

[20] The Caldara trade-policy-uncertainty index measures trade-policy uncertainty directly, so it, rather than the trade category of the economic-policy-uncertainty index, is our primary treatment.

B.5 Calibration

We calibrate the heterogeneous-firm block of the model to the United States and take the tariff experiment to the mid-2026 schedule. Every primitive except the degree of uninsurability is disciplined with data. The uninsurability parameter is set from the estimation literature and swept.

The baseline network is calibrated onto the 2017[21] table of the OECD Inter-Country Input-Output database[22], which records how much each industry buys from every other industry. From the columns referring to the United States I take the entries of the input-output matrix by dividing intermediate flows by gross output. The fifty industries available are then aggregated to eight sectors, some of which have been hit by the tariffs (autos, metals, electronics, and so on) and some have not.[23] For each sector, its input purchases are split by where they come from: bought at home, bought from China, or bought from the safe bloc (USMCA plus the EU).

The calibrated IO table delivers four things. The intermediate share by sector (the fraction of a sector's costs spent on intermediate inputs rather than labor), which ranges from 0.40 in services to 0.75 in autos. The eight-by-eight matrix of category weights, which is how much each of the eight sectors buys from each of the eight, hence the α shares. The split of the origin country of the sourcing firm, which is broken down in three blocks: home, China, and safe bloc. And the final-demand shares, that is how household consumption is divided across sectors (about 0.84 on services). The one primitive the table does not pin down is the vector of match qualities in the productivity shifter, the $A_i(\alpha)$ terms that make a given network optimal. We set these indirectly, by choosing the values at which the incomplete-market network, solved under the 2017 low-tariff regime, reproduces the observed 2017 table exactly, so that the 2017 network is an equilibrium of the model at baseline tariffs.[24] The match qualities are then held fixed across all the policy counterfactuals, so any change in the network comes from the tariff and the risk mechanism and not from re-tuning parameters.

The risk block. The first object to calibrate is the cross-sectional variance covariance ma-

[21] The year 2017 is a normal pre-trade-war benchmark, so the baseline regime is a genuine low-tariff steady state.

[22] Taken from the 2025 edition.

[23] The tariffs are the Section 232 sectoral levers on metals and autos and the Section 301 broad China layer.

[24] The within-sector elasticity σ sets the constant markup $\mathcal{M} = \sigma/(\sigma - 1)$, which raises the price of each home-produced good over its unit cost while imports, priced at the world price plus tariff, carry no markup. This shift in the home price levels is absorbed one for one by the match-quality inversion: solving for the z that reproduce the 2017 shares under the markup returns the same shares as under none, with z reinterpreted as the match quality net of the markup. We verified in the model that a markup of twenty percent leaves the calibrated shares, the wedge, and the misallocation loss unchanged to the reported precision, moving only z (by up to 0.36 in log points). The markup is thus not separately identified from the match qualities off the sourcing shares, so we set σ rather than estimate it and let the inversion absorb it; since the wedge and the loss depend on covariances and Domar weights and the sourcing response is governed by ν, none of the reported results depends on the value chosen.

trix, Σ, across sector, which enters the wedge. The covariance matrix Σ combines two sources on a common scale. The variance of the productivity innovations for the six manufacturing sectors is the annual growth of five-factor total factor productivity from the NBER-CES Manufacturing Industry Database, aggregated from six-digit NAICS with fixed value-added weights. The co-movement structure, meaning each node's loading on a common factor and the sign of that loading, is estimated from real value-added growth in the OECD STAN database, which is internally consistent across the United States and the foreign origins. The reason for keep the scale-free co-movement structure from STAN and set every variance on a total-factor-productivity-equivalent scale is that usually value-added growth is three to six times more volatile than total factor productivity, since value-added growth mixes the technology residual with cyclical factor-input movements, hence using it for the variance would overstate the risk-wedge.

China appears in neither STAN nor NBER-CES, so its part of the covariance matrix is not estimated directly. The six Chinese tradable nodes are instead built by rescaling the corresponding safe-bloc nodes, which are themselves calibrated from STAN value-added variances placed on the total-factor-productivity-equivalent scale. Two empirical premia do the rescaling, sector by sector. The first is a volatility premium: from the value-added rows of seven ICIO tables (1995 to 2020) we take demeaned annualised value-added growth for the six tradable sectors, pool them, and compare the pooled standard deviation of China with that of the safe bloc. China is about 1.28 times as volatile, so each Chinese node's variance is set to that multiple of its safe-bloc counterpart. The second is a co-movement premium, the object that decides whether China is risky to source from: using WIOD real value-added growth (2000 to 2014) we find that Chinese sectors correlate with the US cycle at about $+0.15$, against $+0.50$ for the safe bloc. Therefore each Chinese node actually features a smaller loading on the common US factor than its safe-bloc counterpart. China therefore enters the matrix as more volatile but less correlated with the United States, a low-beta supplier whose fluctuations are largely idiosyncratic to it.

The variance of the exchange rate that enters the currency wedge (15) is the variance of annual log changes in the BIS real broad effective exchange rate for the United States on the 1995-onward window that matches the rest of Σ and is equal to an annual standard deviation of about 4.2 percent. The degree of uninsurability ϕ is the pass-through from firms' risk to workers' marginal utility and is calibrated as the product of two objects obtained from the earnings literature. The first is the pass-through of firm-specific shocks into the incomes of those who bear the firm's risk, which Guiso et al. [2005] estimate to be between 0.05 and 0.15 for permanent shocks. The second is the partial-insurance coefficient: Blundell et al. [2008] estimate how much of a permanent income shock passes through to consumption. That pass-through is about 0.6, meaning roughly 60 percent of a permanent shock shows up in consumption

and the other 40 percent is insured. The product of those two parameters is roughly 0.10, which we take as the baseline and check robustness with numbers from 0.05 to 0.30. The currency loading ψ is set to the same scale, so a single undiversified-stakeholder loading factor applies to productivity, currency, and the tariff shock that we consider further below in the counterfactual.

The network step rests on one functional-form choice, the shape of the productivity shifter $a_i(\alpha)$. We take the entropy form, under which the first-order condition of the optimal input share leads to an expression for the optimal shares that reads as a logit in the risk-adjusted prices:

$$s_{io} \;\propto\; \exp\big((z_{io} - \mathcal{R}_{io})/\nu\big), \tag{45}$$

with z_{io} the match quality of the supplier's location, o, for sector i. The economically meaningful object is the elasticity of substitution across suppliers' locations, that dictates how fast a firm shifts sourcing between home, China, and the safe bloc when their relative costs move. In the logit this elasticity is $1/\nu$, and the same ν is the curvature of a_i that enters the analytical results as the Hessian H_i. A small ν makes the shares respond sharply to relative prices, so the firm reallocates readily, and a large ν makes the network rigid. We set the elasticity to about four, the standard value in the trade literature, so $\nu = 0.25$. The same form governs the extensive margin, since severing a supplier implies that the optimal share moves smoothly toward zero. Furthermore, note that a calibration fixed point searches over the match qualities z until the equilibrium network reproduces the observed 2017 table, converging to a share error below 10^{-10}.

Next, the tariffs are the mid-2026 United States schedule, the Section 232 sectoral levers and the Section 301 broad layer, with the origin dimension applying close to the full sectoral rate to China and a small fraction to the safe bloc, matching realized effective rates near twenty-three percent on China and below five percent on the allies.

Finally, we need to calibrate the firm level TFP process, which in the general equilibrium model affects the households earning process through the rent sharing parameter ϕ. We choose a 0.95 persistence and a process that leads to a leptokrtotic earning distribution (which resembles the data equivalent; see Storesletten et al. [2004]).

References

Daron Acemoglu and Pablo D. Azar. Endogenous production networks. *Econometrica*, 88(1): 33–82, 2020.

Mark A Aguiar, Manuel Amador, and Doireann Fitzgerald. Tariff wars and net foreign assets. Technical report, National Bureau of Economic Research, 2025.

S. Rao Aiyagari. Uninsured idiosyncratic risk and aggregate saving. *Quarterly Journal of Economics*, 109(3):659–684, 1994.

Laura Alfaro, Ester Faia, Nora Lamersdorf, and Farzad Saidi. Health externalities and policy: The role of social preferences. *Management Science*, 68(9):6751–6761, 2022.

Anastasiia Antonova, Luis Huxel, Mykhailo Matvieiev, and Gernot J. Müller. The propagation of tariff shocks via production networks. Working Paper 11917, CESifo, 2025.

Costas Arkolakis, Arnaud Costinot, and Andrés Rodríguez-Clare. New trade models, same old gains? *American Economic Review*, 102(1):94–130, 2012.

Adrien Auclert, Matthew Rognlie, and Ludwig Straub. The macroeconomics of tariff shocks. Technical report, National Bureau of Economic Research, 2025.

David K. Backus and Gregor W. Smith. Consumption and real exchange rates in dynamic economies with non-traded goods. *Journal of International Economics*, 35(3-4):297–316, 1993.

Scott R Baker, Nicholas Bloom, and Steven J Davis. Measuring economic policy uncertainty. *The quarterly journal of economics*, 131(4):1593–1636, 2016.

David Baqaee and Emmanuel Farhi. Networks, barriers, and trade. *Econometrica*, 92(2):505–541, 2024.

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

David Rezza Baqaee and Emmanuel Farhi. Productivity and misallocation in general equilibrium. *Quarterly Journal of Economics*, 135(1):105–163, 2020.

Nuriye Melisa Bilgin, Ester Faia, and Gianmarco IP Ottaviano. Technology spillovers, diffusion and rivalry in firm networks. Technical report, Centre for Economic Policy Research, 2024.

Richard Blundell, Luigi Pistaferri, and Ian Preston. Consumption inequality and partial insurance. *American Economic Review*, 98(5):1887–1921, 2008.

Dario Caldara and Matteo Iacoviello. Measuring geopolitical risk. *American economic review*, 112(4):1194–1225, 2022.

Dario Caldara, Matteo Iacoviello, Patrick Molligo, Andrea Prestipino, and Andrea Raffo. The economic effects of trade policy uncertainty. *Journal of monetary economics*, 109:38–59, 2020.

Christopher D Carroll. The method of endogenous gridpoints for solving dynamic stochastic optimization problems. *Economics letters*, 91(3):312–320, 2006.

Clément de Chaisemartin and Xavier D'Haultfœuille. Two-way fixed effects estimators with heterogeneous treatment effects. *American Economic Review*, 110(9):2964–2996, 2020.

Deutsche Welle Data Team. Trade dispute: Data and analysis on US-China trade diversion. Replication repository, ITC Trade Map data (US Census and UN Comtrade), 2020. https://github.com/dw-data/trade-dispute.

Julian di Giovanni and Andrei A. Levchenko. The risk content of exports: A portfolio view of international trade. *NBER International Seminar on Macroeconomics*, 8(1), 2012.

Jonathan Eaton and Gene M. Grossman. Tariffs as insurance: Optimal commercial policy when domestic markets are incomplete. *Canadian Journal of Economics*, 18(2):258–272, 1985.

Ester Faia, Maximilian Mayer, and Vincenzo Pezone. The value of firm networks: a natural experiment on board connections. *The Review of Corporate Finance Studies*, page cfag011, 2026.

Cecile Gaubert and Oleg Itskhoki. Granular comparative advantage. *Journal of Political Economy*, 129(3):871–939, 2021.

Gene M. Grossman, Elhanan Helpman, and Alejandro Sabal. Optimal resilience in multitier supply chains. *Quarterly Journal of Economics*, 139(4):2377–2425, 2024.

Veronica Guerrieri, Guido Lorenzoni, and Ivan Werning. Tariffs as cost-push shocks: Implications for optimal monetary policy. Technical report, National Bureau of Economic Research, 2025.

Luigi Guiso, Luigi Pistaferri, and Fabiano Schivardi. Insurance within the firm. *Journal of Political Economy*, 113(5):1054–1087, 2005.

Fatih Guvenen, Serdar Ozkan, and Jae Song. The nature of countercyclical income risk. *Journal of Political Economy*, 122(3):621–660, 2014.

Elhanan Helpman and Assaf Razin. *A Theory of International Trade under Uncertainty*. Academic Press, New York, 1978.

Hugo Hopenhayn and Richard Rogerson. Job turnover and policy evaluation: A general equilibrium analysis. *Journal of Political Economy*, 101(5):915–938, 1993.

Hugo A. Hopenhayn. Entry, exit, and firm dynamics in long run equilibrium. *Econometrica*, 60 (5):1127–1150, 1992.

Chang-Tai Hsieh and Peter J Klenow. Misallocation and manufacturing tfp in china and india. *The Quarterly journal of economics*, 124(4):1403–1448, 2009.

Charles R. Hulten. Growth accounting with intermediate inputs. *Review of Economic Studies*, 45(3):511–518, 1978.

Oleg Itskhoki and Dmitry Mukhin. The optimal macro tariff. Technical report, National Bureau of Economic Research, 2025.

Şebnem Kalemli-Özcan, Cansu Soylu, and Muhammed A. Yıldırım. Global networks, monetary policy and trade. Working Paper 33686, National Bureau of Economic Research, 2025.

Narayana R. Kocherlakota. Implications of efficient risk sharing without commitment. *Review of Economic Studies*, 63(4):595–609, 1996.

Ralph S. J. Koijen and Motohiro Yogo. A demand system approach to asset pricing. *Journal of Political Economy*, 127(4):1475–1515, 2019.

Alexandr Kopytov, Bineet Mishra, Kristoffer Nimark, and Mathieu Taschereau-Dumouchel. Endogenous production networks under supply chain uncertainty. *Econometrica*, 92(5):1621–1659, 2024.

Ernest Liu. Industrial policies in production networks. *Quarterly Journal of Economics*, 134(4): 1883–1948, 2019.

Tommaso Monacelli. Tariffs and monetary policy. Technical report, CEPR DP20142, 2025.

David M. G. Newbery and Joseph E. Stiglitz. Pareto inferior trade. *Review of Economic Studies*, 51(1):1–12, 1984.

Ezra Oberfield. A theory of input-output architecture. *Econometrica*, 86(2):559–589, 2018.

David Sraer and David Thesmar. How to use natural experiments to estimate misallocation. *American Economic Review*, 113(4):906–938, 2023.

Kjetil Storesletten, Christopher I. Telmer, and Amir Yaron. Cyclical dynamics in idiosyncratic labor market risk. *Journal of Political Economy*, 112(3):695–717, 2004.

Liyang Sun and Sarah Abraham. Estimating dynamic treatment effects in event studies with heterogeneous treatment effects. *Journal of Econometrics*, 225(2):175–199, 2021.

Jonathan Thomas and Tim Worrall. Self-enforcing wage contracts. *Review of Economic Studies*, 55(4):541–554, 1988.