Firm Location, Industrial Base, and the Gains from Trade

Sergey Nigai*

September, 2026

Abstract

I develop a framework in which a country's industrial base, the mass and productivity distribution of producers located there, is determined by firms' location choices. I embed this mechanism into canonical trade models and show that changes in the industrial base are a first-order determinant of the gains from trade. Using 1992-2019 data, I find that its welfare contribution is as large as the standard domestic-expenditure-share term and reverses welfare changes for many countries. In a counterfactual that undoes China's post-WTO productivity growth, I find positive welfare effects for many trading partners due to industrial base relocation, overturning conventional predictions.

Keywords: *Firm Location, Welfare Gains, Endogenous Heterogeneity, International trade, Market Access, Industrial Base*

JEL-codes: F10, F20.

*University of Colorado Boulder and CESIfo, sergey.nigai@colorado.edu.

This paper supersedes an earlier version circulated under the title"From Discrete Choice to Endogenous Distributions: Selection and Heterogeneity".

The author declares that he has no relevant or material financial interests that relate to the research described in this paper.

1 Introduction

Over the past few decades, the expansion of international trade has coincided with a large-scale relocation of global production. Against this backdrop, the concern that globalization may erode a country's industrial base has moved to the center of policy debates around the world, as exemplified by the recent resurgence of industrial policy.[1] Yet little is known about how changes in a country's industrial base induced by firms' location responses to globalization shape the aggregate gains from trade.

This paper develops a quantitative model of the world economy in which each country's industrial base is an equilibrium outcome of firms' location choices that respond to globalization. The central finding is that the welfare consequences of globalization depend not only on how much a country trades, but also on how trade changes the mass and productivity distribution of firms located there. A country may import more and consume cheaper foreign goods yet lose if globalization erodes its industrial base. Conversely, a country may import less and still gain if lower imports reflect an expanding industrial base rather than reduced access to foreign goods. Hence, from a policy perspective, it is important to understand when changes in the industrial base magnify the standard gains from trade and when they overturn them.

To fix ideas, consider the welfare formula in Arkolakis et al. (2012) based on the domestic-expenditure-share sufficient statistics, π_{ii}, written before imposing the restriction that the country-specific aggregate technology term, T_i, is exogenous and constant. Let Δ denote changes between two equilibria such that changes in welfare can be written as:

$$\Delta \ln W_i = -\frac{1}{\eta}\Delta \ln \pi_{ii} + \Delta \ln T_i,$$

where η is the trade elasticity parameter. In the canonical environments, T_i is treated as a primitive and held fixed, so $\Delta \ln T_i = 0$. In this case, the term $-(1/\eta)\Delta \ln \pi_{ii}$, which I refer to as the ACR term, summarizes the gains from trade. Under endogenous firm location, the aggregate term $\Delta \ln T_i$ is no longer absent and instead captures changes the mass and productivity distribution of firms that locate in country i, i.e., changes in its industrial base. These changes have direct welfare effects: firms located at home do more than supply the domestic market; they shape the country's effective

[1] For discussions of global production relocation, see Hanson (2012); for the recent resurgence of industrial policy, see Juhász et al. (2024).

aggregate productivity and hence real income.

The key advantage of the framework is that it delivers a tractable analytical mapping from producer-level location choices to the resulting mass and productivity distribution of firms in each country. I formulate location choice as a general profit-maximization problem in which payoffs depend on location-specific productivity, production costs, and market access. I specify conditions under which this problem can be embedded in different market structures, including perfect competition as in Eaton and Kortum (2002), Bertrand competition as in Bernard et al. (2003), and monopolistic competition in the spirit of Melitz (2003). I show that across these environments, the equilibrium system takes a common form. The framework therefore remains consistent with canonical models of international trade while changing the interpretation of firm heterogeneity within each country: the productivity distribution of active firms is itself an equilibrium object. Hence, an additional contribution of this paper is providing a microfoundation for location-specific firm heterogeneity as the outcome of selection across production locations.

I structurally estimate the key parameters and calibrate the model to global trade data from 1992 to 2019. The quantitative results show that the industrial-base margin is a first-order determinant of the gains from trade. Over this period, its welfare contribution is approximately as large as the ACR term, and the two are negatively correlated across countries: countries with larger declines in domestic expenditure shares also tend to experience losses in their industrial base. Accounting for this margin changes measured welfare gains substantially and reverses the sign of welfare changes for mode than a third of countries. I also consider a counterfactual exercise in the spirit of Eaton and Kortum (2002) and Caliendo et al. (2018) that quantifies how a productivity change in one location affects welfare in other countries through trade. Reversing China's aggregate technology to its pre-WTO accession level induces firms to leave China for other production locations. This relocation response overturns the standard conclusion about the benefits of foreign technology, turning welfare losses into gains for many of China's trading partners.

Relation to the Literature.

This paper is related to the literature on the *delocation* channel, through which changes in trade costs and policy affect welfare by reallocating production across

countries. This literature includes work on the home-market effect, which links production location to market size (Krugman 1980; Helpman and Krugman 1985), and work on how trade policy reallocates production across borders (Venables 1985, 1987). The welfare results in this paper also speak to the new economic geography literature, which studies how trade liberalization reshapes the spatial distribution of industrial production (Krugman and Venables 1995; Baldwin and Robert-Nicoud 2000; Baldwin et al. 2003).[2]Relative to this literature, I bring the delocation margin into quantitative trade by developing a unified endogenous firm-location problem that can be embedded in canonical gravity models. The framework preserves tractability and avoids common limitations in traditional economic geography, such as multiple equilibria, that complicate counterfactual analysis.

In terms of the broad research question, the paper relates to recent quantitative work that studies the interaction between trade policy and industrial base including homogeneous-firm environments (Ossa 2011; Lashkaripour and Lugovskyy 2023) and on heterogeneous-firm models in the spirit of Melitz (2003) and Chaney (2008) as in Haaland and Venables (2016); Spearot (2016); Caliendo et al. (2015, 2023); Caliendo and Parro (2019); Bagwell and Lee (2020) and others. The key distinction is that firms do not enter countries through separate free-entry conditions. Instead, each firm is characterized by a vector of productivities, one for each potential production location, and chooses where its cost-adjusted productivity and market access yield the highest profit. The within-country productivity distributions are therefore equilibrium objects determined by which firms choose to locate there.

On the modeling side, the paper contributes to a small literature on discrete location choice by heterogeneous firms. It generalizes Suárez Serrato and Zidar (2016) and Fajgelbaum et al. (2019), who study firm location in monopolistically competitive environments with independent location-specific productivity draws, in two directions.[3] First, I allow location-specific productivity draws from multivariate extreme-value and power-law distributions to be correlated across locations. Second, I generalize

[2]The paper also connects to empirical work on how globalization and trade policy shape the location of industrial activity (e.g., Hanson 1996, 2001; Flaaen et al. 2020), and offers a structural interpretation of why improved market access in one country can erode the industrial base elsewhere, the mechanism behind the manufacturing and local-labor-market effects documented in the China-shock literature (e.g., Autor et al. 2013).

[3]For empirical evidence on the connection between discrete-choice shares and average productivity of selected firms, also see Malgouyres et al. (2023) and Suárez Serrato and Zidar (2023).

the location problem so that it can be embedded in multiple market structures. Both generalizations matter because they make it possible to isolate environments in which endogenous location choice has aggregate consequences. For example, under perfect and Bertrand competition, independence makes the location margin disappear in the aggregate: the gain from attracting more firms is exactly offset by lower average productivity among selected firms, and the model collapses to the canonical Eaton and Kortum (2002) and Bernard et al. (2003) allocations. This helps explain why much of the canonical quantitative trade literature has operated without this margin. My empirical results point to strong cross-location dependence in firm productivity, which magnifies the role of the location margin in shaping the gains from trade across market structures.

This paper also provides an alternative to models in which complementarities and agglomeration forces shape the allocation of firms and workers across space. These include assignment models with supermodularity between agent and location characteristics (see Costinot and Vogel 2010; Behrens et al. 2014; Sampson 2014; Gaubert 2018), models of spatial sorting with agglomeration economies (Lucas and Rossi–Hansberg 2002; Baldwin and Okubo 2006; Okubo et al. 2010; Combes et al. 2012), and recent work on how agglomeration and scale economies affect international trade outcomes (see Kucheryavyy et al. 2023; Bartelme et al. 2025). The mechanism developed here does not rely on supermodularity or agglomeration externalities. Instead, the paper combines payoff-based discrete choice with multivariate productivity distributions, preserving tractability for both location shares and the productivity distributions induced by selection.[4]

Finally, the paper relates to the broader quantitative literature on the welfare gains from trade, the allocation of production, and economic development. On the one hand, it develops a sufficient-statistics approach to measuring the gains from trade, following Arkolakis et al. (2012) and subsequent work (Costinot and Rodríguez-Clare 2014; Melitz and Redding 2015; Allen et al. 2020). Relative to this strand, endogenous

[4]This methodological connection also relates the paper to tractable discrete-choice models following McFadden (1974, 1978) and to applications that use discrete choice in spatial, trade, and development settings (among many others, Artuç et al. 2010; Grogger and Hanson 2011; Allen and Arkolakis 2014; Dix-Carneiro 2014; Ahlfeldt et al. 2015; Holmes and Sieg 2015; Monte et al. 2018; Bryan and Morten 2019; Fajgelbaum and Gaubert 2020; Galle et al. 2023). The distributional results also relate to work with multivariate Fréchet heterogeneity (Lind and Ramondo 2023a,b) and multivariate Pareto heterogeneity (Arkolakis et al. 2018, 2025).

firm location generates a new channel for the gains from trade that is quantitatively as important as the standard ACR term. On the other hand, I connect to the literature on the spatial allocation of production factors, technology and macroeconomic development as determinants of long-run development (see Desmet and Rossi-Hansberg 2014; Desmet et al. 2018) by endogenizing local technology through the equilibrium allocation firms across space.

2 A Unified Firm Location Problem

This section develops a unified framework for endogenous firm location. The framework builds on Fajgelbaum et al. (2019) and generalizes it in two directions: it allows productivity draws to follow multivariate distributions with dependent components, and it specifies a location problem that can be embedded in several canonical market structures. As shown below, the dependence structure governs how firm allocation shapes the conditional productivity distribution in each location, which is the key object behind the industrial-base term in the welfare analysis.

To keep the results general, I impose a number of aggregate and firm-level restrictions. Consider a world economy composed of a finite set of countries $k \in \mathcal{K}$. There is a continuum of tradable goods indexed by ω. Labor is immobile across countries and is supplied exogenously, with $L_k > 0$ denoting the labor endowment of country k. Production in each country uses labor, with wage $w_k > 0$. Trade between countries is subject to bilateral iceberg trade costs $\tau_{kj} \geq 1$ for goods shipped from k to j as well as an ad valorem import tariff κ_{kj} imposed by destination country j.

Restriction 1. *In each destination j, the representative consumer has CES preferences over varieties with elasticity $\sigma > 1$.*

The endogenous firm location problem considered in this paper applies to models that impose CES preferences as a macro-level restriction, similar to Arkolakis et al. (2012). In each destination j, the representative consumer has preferences over the set of available varieties Ω_j given by

$$U_j = \left[\int_{\omega \in \Omega_j} q_j(\omega)^{\frac{\sigma-1}{\sigma}} d\omega \right]^{\frac{\sigma}{\sigma-1}}, \qquad \sigma > 1.$$

The associated CES price index is

$$P_j = \left[\int_{\omega \in \Omega_j} p_j(\omega)^{1-\sigma} d\omega \right]^{\frac{1}{1-\sigma}}.$$

As I show later, this restriction implies that bilateral sales are log-separable in exporter unit costs, bilateral barriers, and destination-market characteristics, which is central for the analytical tractability of the firm location problem.

Firms are indexed by productivity vectors $\boldsymbol{z}$ and can locate in any country $k \in \mathcal{K}$. Their productivities, however, need not be the same across locations and may depend on local conditions. I allow this dependence to be flexible and productivity draws across locations to range from independent to perfectly dependent. This is modeled using a multivariate distribution for the productivity vector $\boldsymbol{z} = (z_1, \ldots, z_K)$, which leads to the next restriction.

Restriction 2. *The productivity vector $\boldsymbol{z} = (z_1, \ldots, z_K)$ is distributed multivariate Fréchet with logistic dependence.*

The second restriction specifies the distribution of firm latent productivities across locations. While Eaton and Kortum (2002) use a special case of $F(\boldsymbol{z})$ under independence, they also discuss this cumulative distribution function with a general dependence structure. This distribution is also used in Lind and Ramondo (2023a,b) to examine correlated technologies. The results in this paper are complementary, but distinct: they imply that the productivity distribution of firms operating in a given country is not an exogenous primitive, but an equilibrium outcome shaped by firms' location choices. Hence, the mass of active firms and their productivity distribution are endogenous and jointly determined in equilibrium. This contrasts with standard approaches, which typically treat firm allocation and within-country heterogeneity as fixed primitives. Restriction 2 is not unique; I discuss alternative distributions in Section 6. I discuss properties of this distribution in Appendix A.

I assume that $\boldsymbol{z}$ follows a multivariate Fréchet distribution with logistic dependence:

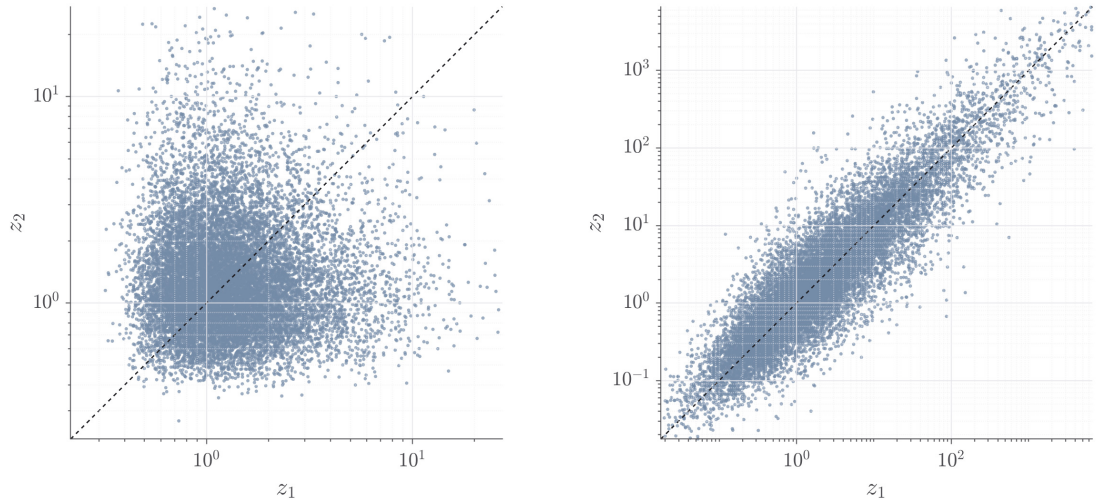
$$F(\boldsymbol{z}) = \exp\left(- \left(\sum_{k=1}^{K} \left[H_k z_k^{-\theta} \right]^{\rho} \right)^{1/\rho} \right), \qquad H_k > 0, \quad \theta > 0, \quad \rho \geq 1. \tag{1}$$

This distributional assumption guarantees several attractive properties. First, the unconditional marginal distribution of z_i is Fréchet with scale H_i and shape θ,

$$F_i(z) = \exp\left(-H_i z^{-\theta}\right),$$

so that the interpretation of H_i and θ is consistent with Eaton and Kortum (2002). The parameter H_i is a country-i primitive that shifts the distribution of potential productivity draws for firms locating in i. The parameter θ governs the dispersion of productivity draws within countries.

(a) Panel A: $\rho = 1$ (b) Panel B: $\rho = 3$



Notes: The figure illustrates 15,000 simulated draws from a bivariate Fréchet distribution with common scale parameters $H_k = 1$ and tail parameter $\theta = 2$. Panel A shows the independence case, $\rho = 1$; Panel B shows the case with positive dependence, $\rho = 3$. The dashed line is the 45° diagonal. Both panels are plotted on a log-log scale.

Figure 1: Simulated Draws from Bivariate Fréchet Distribution with Logistic Dependence

The parameter ρ governs the dependence structure across locations: $\rho = 1$ corresponds to independent productivity draws, while larger values of ρ allow for stronger positive dependence. In this sense, ρ provides a parsimonious measure of firms' potential mobility across countries. When $\rho = 1$, a high realization of z_i in country i carries no information about the firm's productivity in other countries. Location-specific productivity is then relatively important for the location decision. In the limit as $\rho \to \infty$, productivity is perfectly dependent across locations, after adjusting

for country-specific scale parameters, and the location problem approaches the deterministic case. This distinction is illustrated in Figure 1. Panel A shows independent productivity draws, while Panel B shows positively dependent draws in the bivariate case.

The third restriction is at the firm level and requires the profit-to-revenue ratio to be constant across firms. The framework can accommodate different market structures as long as they satisfy this condition. Let $\mathcal{S}$ denote the profit-to-revenue ratio, and let $\mathcal{M}_i$ denote the market-access term faced by firms located in i. The class of market structures considered in this section is summarized by the following restriction.

Restriction 3. *A market structure is a tuple $(\eta, \mathcal{S}, \{\mathcal{M}_i\})$ with η and $\mathcal{S}$ invariant across locations.*

In addition to productivity draws in $\boldsymbol{z}$, firms take wages w_i and market access $\mathcal{M}_i$ as given when choosing its location. Total revenue from locating in i satisfies:

$$R_i(\boldsymbol{z}) \propto \left(\frac{z_i}{w_i}\right)^{\eta} \mathcal{M}_i, \qquad \eta > 0, \tag{2}$$

where η is the elasticity of revenue with respect to cost-adjusted productivity z_i/w_i, and $\mathcal{M}_i$ is the market-access term. Total profits from locating in i are a constant fraction of revenue:

$$\Pi_i(\boldsymbol{z}) = \mathcal{S} R_i(\boldsymbol{z}).$$

Finally, I assume that firms' location decisions are equivalent to maximizing a strictly increasing transformation of total profits, or equivalently revenue, in this setting:

$$i^*(\boldsymbol{z}) = \arg\max_i \left\{ z_i w_i^{-1} \mathcal{M}_i^{1/\eta} \right\}. \tag{3}$$

Restrictions 1–3, together with the profit-maximization problem in equation (3), imply that firms choose locations based on productivity draws adjusted for production costs and destination sales opportunities summarized by market access. The next proposition uses this structure to characterize firm allocation across discrete alternatives and the conditional productivity distribution among firms selected into each location.

Proposition 1. *Suppose the vector of shocks $\boldsymbol{z} = (z_1, \ldots, z_K)$ has the joint cumulative*

distribution function F as in equation (1) and that agents choose the alternative k that maximizes the payoff in equation (3), then:

(i) The choice probabilities are given by:

$$\mu_i = \frac{H_i^{\rho} w_i^{-\theta\rho} \mathcal{M}_i^{\theta\rho/\eta}}{\sum_k H_k^{\rho} w_k^{-\theta\rho} \mathcal{M}_k^{\theta\rho/\eta}}.$$

(ii) The conditional distribution function of z_k given selection ($k^* = k$) is Fréchet:

$$F_k(z \mid k^* = k) = \exp\Big(-\tilde{T}_k z^{-\theta}\Big) \qquad \text{, where} \qquad \tilde{T}_k = H_k\, \mu_k^{-1/\rho}.$$

(iii) If $\theta > 1$, the expected shock conditional on choosing k satisfies:

$$\mathbb{E}[z_k \mid k^* = k] = H_k^{1/\theta} \mu_k^{-1/(\theta\rho)} \Gamma(1 - 1/\theta),$$

where $\Gamma(\cdot)$ denotes the Gamma function.

Proof. The proof is available in Appendix A. □

Proposition 1 shows how Fréchet heterogeneity within each country can arise endogenously from the firm location problem. Under a multivariate Fréchet distribution with logistic dependence, equilibrium location shares determine the scale of the conditional productivity distribution among firms that select into each country. In particular, the Fréchet scale in country k, $\tilde{T}_k = H_k \mu_k^{-1/\rho}$, combines the primitive technology term H_k with the endogenous selection term $\mu_k^{-1/\rho}$. This reinterpretation preserves the closed-form tractability of Fréchet-based trade models while endogenizing the allocation of firms, μ_k, across countries. Consequently, changes in fundamentals affect both the productivity distribution within countries and the equilibrium allocation of firms across locations.

Let T_k denote the industrial base of country k, interpreted as the supply shifter that governs country k's ability to serve every destination. It consists of the mass of firms that choose to locate in k with the productivity distribution of those firms. Using the

conditional productivity scale in Proposition 1(iii), this object can be written as

$$T_k = \underbrace{\mu_k}_{\text{mass of firms}} \times \underbrace{\left(H_k^{1/\theta}\mu_k^{-1/(\theta\rho)}\right)^{\eta}}_{\text{conditional revenue productivity}} . \tag{4}$$

The first term captures the measure of producers located in country k, and therefore the measure of varieties which k can supply to the world economy. The second term captures the productivity of the firms that select into k. Productivity enters with exponent η because this is the elasticity that maps firm-level productivity advantages into revenue and aggregate supply capacity and depends on the market structure. This term can be viewed as the conditional aggregate revenue productivity.

The expression for the industrial base in equation (4) shows how the mass-composition tradeoff depends on the market-structure parameter η and on the dependence parameter ρ. It also clarifies the special cases and the role of the generalization developed in this paper. First, under perfect competition and Bertrand competition with independent productivity draws, $\eta = \theta$ and $\rho = 1$. In this case, the gain from attracting more firms is exactly offset by the lower conditional productivity of selected firms, so $T_k = H_k$. The location margin then disappears in the aggregate, and the framework collapses to Eaton and Kortum (2002) and Bernard et al. (2003). Second, under monopolistic competition, $\eta = \sigma - 1$ as in Suárez Serrato and Zidar (2016) and Fajgelbaum et al. (2019) and the dependence parameter ρ governs how strongly selection affects the industrial base. The general framework makes it possible to estimate both elasticity parameters from the data and to quantify the role of the industrial base in shaping the gains from trade.

3 Market Structures

I now embed the unified firm-location problem in three market structures: perfect competition, Bertrand competition, and monopolistic competition. The purpose of this section is to show that these environments share the same location-choice structure once three market-structure-specific objects are identified: the revenue elasticity η, the market-access aggregator $\mathcal{M}_i$, and the profit-to-sales ratio $\mathcal{S}$. Given these objects, Proposition 1 immediately delivers the equilibrium location shares and the conditional productivity distributions of firms located in each country.

3.1 Perfect Competition and Bertrand Competition

I first embed the firm-location problem in two canonical models with Fréchet heterogeneity: the perfect-competition environment of Eaton and Kortum (2002) and the Bertrand-competition environment of Bernard et al. (2003). The new step is to introduce firm-level location choice and make country-level technology an endogenous outcome of these choices, while retaining the analytical tractability of the aggregate equations in the canonical models.

Introducing endogenous location choice in these environments requires some care because both frameworks rely on a continuum of atomistic producers and an aggregate representation of technology. Firm-level revenues must therefore be defined in a way that is consistent with both features. My strategy is to construct the continuum as the limit of economies with N firms, in which firm revenues are well defined and aggregate to country-level revenue. I then show that the atomistic limit delivers the revenue expression in Restriction 3, allowing me to apply Proposition 1 to characterize firms' location choices and the resulting selection into countries.

In the spirit of Klette and Kortum (2004), I treat each firm as owning a technology process that can be applied across different varieties. A firm located in country i is characterized by efficiency draws for all varieties and becomes country i's potential supplier of variety ω if it draws the highest efficiency among firms located there. The max-stability of the Fréchet distribution allows to represent competition among firms within a country as an additional nest inside the standard competition mechanism across countries.

Let me use superscript f to denote discrete firms. Conditional on locating in i, firm f's efficiency draws across varieties follow:

$$G_i^f(u) \equiv \Pr\left(\widetilde{z}_i^f(\omega) \leq u \mid z_i^f\right) = \exp\left(-[z_i^f]^\theta u^{-\theta}\right),$$

independently across firms and varieties. Thus, conditional on its location draw z_i^f, firm f is itself an Eaton-Kortum-type source with technology parameter $[z_i^f]^\theta$ and shape parameter θ.[5]

Aggregation from firms to countries follows directly from the max-stability of the

[5] Throughout this manuscript, as is standard in the literature, I assume that $\theta > \sigma - 1$.

Fréchet distribution. For any variety ω, country i's productivity is the best technology available among the firms located there. Conditional on firm capabilities and locations:

$$\Pr\left(\max_f \widetilde{z}_i^f(\omega) \leq u\right) = \prod_f G_i^f(u) \ = \exp\left[-\left(\sum_f [z_i^f]^\theta\right) u^{-\theta}\right].$$

Hence, for finite N, Fréchet max-stability property allows me to derive country-level Fréchet parameter as the summation of technology parameters of firms located in country i as well as the probability that firm f provides country i's best technology as follows:

$$\overline{T}_i = \sum_k [z_i^k]^\theta \quad \text{and} \quad s_i^f = \frac{[z_i^f]^\theta}{\overline{T}_i}.$$

This means that firm f's total revenues can be calculated as the share s_i^f of country i's total revenues such that it is proportional to $[z_i^f]^\theta$. Note that while the number of firms is finite, varieties remain a continuum and conditional on the realized country technology parameter $\overline{T}_i$, aggregate equations for trade shares can be derived following the same steps as in Eaton and Kortum (2002) and Bernard et al. (2003).

Let E_j denote total expenditure at tariff-inclusive consumer prices in destination country j. Further, suppose that conditional on supplying a variety, the distribution of prices is independent of the identity of the supplying firm. Then, firm f's total revenue is its share of country i's producer revenue in every destination:

$$R_i^f = \underbrace{s_i^f}_{\text{Firm } f\text{'s share}} \times \underbrace{\left(\sum_j \frac{\overline{T}_i(w_i\tau_{ij}\kappa_{ij})^{-\theta}}{\overline{\Phi}_j}\frac{E_j}{\kappa_{ij}}\right)}_{\text{Country } i\text{'s total producer revenue}},$$

where $\overline{\Phi}_j \equiv \sum_k \overline{T}_k(w_k\tau_{kj}\kappa_{kj})^{-\theta}$. Substituting the expression for f's share, the country technology parameter cancels and firm revenue becomes

$$R_i^f = \left(\frac{z_i^f}{w_i}\right)^\theta \overline{\mathcal{M}}_i, \qquad \overline{\mathcal{M}}_i \equiv \sum_j \tau_{ij}^{-\theta}\kappa_{ij}^{-(\theta+1)}\frac{E_j}{\overline{\Phi}_j},$$

where $\overline{\mathcal{M}}_i$ is the market-access term. Hence, at every finite N, firm-level revenue is well defined and takes the form required by Restriction 3 in equation (2) with

$\eta = \theta$. As N grows, $N \ln N \, \overline{\mathcal{M}}_i$ converges to the deterministic continuum market access $\mathcal{M}_i$, where $N \ln N$ is a normalization factor analogous to the normalization that arises for Zipf-tailed firm sizes in Gabaix (2011). Under this normalization, the finite-firm location problem converges to the continuum problem characterized by Proposition 1: realized aggregates converge to their continuum counterparts and individual firms become negligible.

In Appendix B, I show that:

$$\frac{N_i}{N} \xrightarrow{p} \mu_i, \qquad \frac{\overline{T}_i}{N \ln N} \xrightarrow{p} T_i, \qquad \max_{f \in \mathcal{F}_i} s_i^f \xrightarrow{p} 0,$$

jointly across countries, where N_i is the number of firms located in country i. The second limit implies $\overline{\Phi}_j / (N \ln N) \xrightarrow{p} \Phi_j \equiv \sum_k T_k (w_k \tau_{kj} \kappa_{kj})^{-\theta}$ and that the continuum expressions for revenues and market-access are as follows:

$$R_i(z_i) \propto \left(\frac{z_i}{w_i} \right)^{\theta} \mathcal{M}_i, \qquad \mathcal{M}_i = \sum_j \tau_{ij}^{-\theta} \kappa_{ij}^{-(\theta+1)} \frac{E_j}{\Phi_j}.$$

The profit expression is market-structure specific, but in both environments profits are a contant fraction of producer revenue $\Pi_i(z_i) = \mathcal{S} R_i(z_i)$, where $\mathcal{S}$ is common across origins, destinations, and firm types. Perfect competition and Bertrand competition differ in the constants governing markups, revenues, and profits, but not in the elasticity that maps firm productivity into sales or in the location statistic. In both environments, $\eta = \theta$. The two cases therefore imply the same revenue-ranking problem up to the market-structure-specific constant $\mathcal{S}$.

Under perfect competition, operating profits are zero and $\mathcal{S} = 0$. If the location objective is interpreted literally as profit maximization, the zero-profit case is degenerate because all locations yield zero profits. There are two alternatives to define firm-location problem in Eaton-Kortum-type models. The first is to assume that firms choose locations to maximize total producer revenue. The second is to view perfect competition as the limit of an environment with positive operating profits as $\mathcal{S} \to 0$. Under either interpretation, the firm-location problem is well defined and delivers the same allocation of firms across countries.

Bertrand competition leaves the fundamental cost comparison mechanism, and hence

$\eta = \theta$ and $\mathcal{M}_i$, unchanged. What differs is the markup and therefore the profit-to-sales ratio. Let $c_{(1)j}$ and $c_{(2)j}$ denote the lowest and second-lowest delivered costs for a variety offered in j. Under Bertrand competition the lowest-cost supplier serves the variety at the limit price:

$$p_j = \min\{c_{(2)j},\, \bar{m}\, c_{(1)j}\}, \qquad \bar{m} = \frac{\sigma}{\sigma - 1},$$

The markup over the winner's marginal cost is:

$$m_j \equiv \frac{p_j}{c_{(1)j}} = \min\left\{\frac{c_{(2)j}}{c_{(1)j}},\, \bar{m}\right\}.$$

In atomistic limit, the joint distribution of $(c_{(1)j}, c_{(2)j})$ conditional on the winning firm becomes independent of the firm's identity. Following Bernard et al. (2003), the unconstrained markup $c_{(2)j}/c_{(1)j}$ is Pareto with shape parameter θ. With the markup cap $\bar{m}$, the realized markup m_j has distribution:

$$\Pr(m_j \leq x) = 1 - x^{-\theta} \quad \text{for } 1 \leq x < \bar{m}, \qquad \Pr(m_j = \bar{m}) = \bar{m}^{-\theta}.$$

In this environment profits are a constant fraction of producer revenue:

$$\mathcal{S} = \frac{1}{1+\theta}.$$

Hence, Bertrand competition changes the level of operating profits relative to perfect competition, but because the profit-to-revenue ratio is common across locations, it does not change relative location payoffs conditional on the same revenue elasticity and market-access term.

Given the equilibrium effective Fréchet scales, T_i, defined in in equation (4) and using the result that $\eta = \theta$, international trade shares are characterized in the conventional form:

$$\pi_{ij} = \frac{T_i(w_i\tau_{ij}\kappa_{ij})^{-\theta}}{\sum_k T_k(w_k\tau_{kj}\kappa_{kj})^{-\theta}} = \frac{H_i\, \mu_i^{1-1/\rho}(w_i\tau_{ij}\kappa_{ij})^{-\theta}}{\sum_k H_k\, \mu_k^{1-1/\rho}(w_k\tau_{kj}\kappa_{kj})^{-\theta}}.$$

This expression shows why allowing for a multivariate distribution of location-specific productivities is important. The dependence parameter ρ governs how the mass of firms and their conditional productivity composition interact. When $\rho = 1$, produc-

tivity draws are independent across locations and the increase in the mass of firms selecting country i is exactly offset by a decline in the conditional Fréchet technology scale per selected firm. In this knife-edge case $T_i = H_i$, the location margin disappears from the aggregate allocation and the model collapses to the standard Eaton–Kortum-type structures with fixed country-level productivity distributions. When $\rho > 1$, productivities are positively dependent across locations, $T_i = H_i \mu_i^{1-1/\rho}$ varies with the equilibrium allocation of firms. The industrial-base channel therefore becomes active: countries that attract more firms become more competitive suppliers in world markets.

The CES price index is as follows:

$$P_j = \mathcal{B}(\theta, \sigma)\, \Phi_j^{-1/\theta}, \qquad \Phi_j = \sum_k H_k\, \mu_k^{1-1/\rho} (w_k \tau_{kj} \kappa_{kj})^{-\theta},$$

where $\mathcal{B}(\theta, \sigma)$ is a constant that depends on the market structure but not on country-level fundamentals.

3.2 Monopolistic Competition

I next embed endogenous firm location in a monopolistic-competition environment. Unlike the perfect- and Bertrand-competition cases, where each firm owns a technology process spanning the variety space, each firm here produces a single differentiated variety. The environment is close in spirit to Melitz (2003), but I restrict attention to specifications in which export participation can be summarized by the same market-access term. This includes both the case without fixed export costs and the isomorphic case with stochastic fixed export costs proportional to destination-market size. Under these restrictions, fixed export costs affect the level of market access but not the structure of the location problem, and can therefore be suppressed.

Under Restriction 1, a firm located in i with productivity z_i that serves destination j charges the consumer price according to the standard constant mark-up rule:

$$p_{ij}(z_i) = \frac{\sigma}{\sigma - 1} \frac{w_i \tau_{ij} \kappa_{ij}}{z_i}.$$

Producer revenue, excluding tariff payments, is as follows:

$$R_{ij}(z_i) = z_i^{\sigma-1}\left(\frac{\sigma}{\sigma-1}w_i\tau_{ij}\right)^{1-\sigma}\kappa_{ij}^{-\sigma}P_j^{\sigma-1}E_j.$$

Under monopolistic competition with CES demand, operating profit is a constant fraction of producer revenue:

$$\Pi_{ij}(z_i) = \frac{1}{\sigma}R_{ij}(z_i).$$

Thus, total operating profit of firm z conditional on locating in i is:

$$\Pi_i(z_i) = \frac{1}{\sigma}\left(\frac{\sigma}{\sigma-1}\right)^{1-\sigma}\left(\frac{z_i}{w_i}\right)^{\sigma-1}\mathcal{M}_i,$$

where the market-access term is:

$$\mathcal{M}_i \equiv \sum_j \tau_{ij}^{1-\sigma}\kappa_{ij}^{-\sigma}\frac{E_j}{P_j^{1-\sigma}}. \tag{5}$$

Hence, monopolistic competition is characterized by

$$\eta = \sigma - 1, \qquad \mathcal{S} = \frac{1}{\sigma},$$

with market access given by equation (5). As before, the mass of firms located in i is μ_i, with the total mass of firms normalized to one. I can now apply the results in Proposition 1 such that conditional on selection into i, firm productivity is Fréchet with scale:

$$\tilde{T}_i = H_i\mu_i^{-1/\rho}.$$

Thus, $\tilde{T}_i$ is the per-firm productivity scale among firms located in i, while the aggregate source term also depends on the selected mass μ_i.

Integrating producer revenue over the conditional productivity distribution of firms located in i, average bilateral producer revenue in destination j is

$$\int R_{ij}(z)\,dF_i(z \mid i^* = i) = \mathcal{C}(\sigma,\theta)\,\tilde{T}_i^{(\sigma-1)/\theta}w_i^{1-\sigma}\tau_{ij}^{1-\sigma}\kappa_{ij}^{-\sigma}P_j^{\sigma-1}E_j,$$

where $\mathcal{C}(\sigma,\theta)$ is a constant.

Consumer expenditure on varieties produced in i is κ_{ij} times producer revenue. Hence, international trade shares can be written as:

$$\pi_{ij} = \frac{H_i^{(\sigma-1)/\theta} \mu_i^{1-(\sigma-1)/(\rho\theta)} (w_i \tau_{ij} \kappa_{ij})^{1-\sigma}}{\sum_k H_k^{(\sigma-1)/\theta} \mu_k^{1-(\sigma-1)/(\rho\theta)} (w_k \tau_{kj} \kappa_{kj})^{1-\sigma}}.$$

The CES price index satisfies:

$$P_j^{1-\sigma} = \sum_i \mu_i \int p_{ij}(z)^{1-\sigma} dF_i(z \mid i^* = i). \tag{6}$$

Using the conditional Fréchet distribution from Proposition 1, it can be written as follows:

$$P_j^{1-\sigma} = \mathcal{D}(\sigma, \theta) \sum_i H_i^{(\sigma-1)/\theta} \mu_i^{1-(\sigma-1)/(\rho\theta)} (w_i \tau_{ij} \kappa_{ij})^{1-\sigma},$$

where $\mathcal{D}(\sigma, \theta)$ is a constant.

4 Aggregate Mappings and Equilibrium

The previous section shows that the endogenous firm-location problem can be embedded in perfect competition, Bertrand competition, and monopolistic competition in a unified way. Although these market structures differ in their microeconomic foundations, each is summarized by the same aggregate objects: a revenue elasticity η, a profit-to-revenue ratio $\mathcal{S}$, and a market-access term $\mathcal{M}_i$. Once these objects are specified, location shares, trade shares, price indices, and welfare can be written in an isomorphic form. This section collects these aggregate mappings and defines the equilibrium system conditional on the market-structure parameters.

Given the allocation of firms in Proposition 1, destination j's consumer-price expenditure share on goods sourced from i is

$$\pi_{ij} = \frac{H_i^{\eta/\theta} \mu_i^{1-\eta/(\theta\rho)} (w_i \tau_{ij} \kappa_{ij})^{-\eta}}{\sum_k H_k^{\eta/\theta} \mu_k^{1-\eta/(\theta\rho)} (w_k \tau_{kj} \kappa_{kj})^{-\eta}}. \tag{7}$$

The market-access term relevant for the location decision is

$$\mathcal{M}_i = \sum_j \frac{\tau_{ij}^{-\eta} \kappa_{ij}^{-(\eta+1)}}{\sum_k H_k^{\eta/\theta} \mu_k^{1-\eta/(\theta\rho)} (w_k \tau_{kj} \kappa_{kj})^{-\eta}} E_j. \tag{8}$$

Finally, the aggregate CES price index satisfies

$$P_j^{-\eta} = \mathcal{A} \sum_i H_i^{\eta/\theta} \mu_i^{1-\eta/(\theta\rho)} (w_i \tau_{ij} \kappa_{ij})^{-\eta}, \tag{9}$$

where $\mathcal{A}$ is a market-structure constant.

Given the aggregate variables $\{\mu_i, \pi_{ij}, \mathcal{M}_i, P_i\}$, it remains to specify aggregate expenditure and the equilibrium condition that pins down wages. Aggregate expenditure in country i consists of labor income, redistributed profits, tariff revenue, and an exogenous trade deficit D_i:

$$E_i = w_i L_i + \Pi_i + TR_i + D_i. \tag{10}$$

Since E_j is final expenditure at consumer prices while producers receive revenue net of tariffs, the producer-price trade flow from i to j is as follows:

$$X_{ij} = \frac{\pi_{ij}}{\kappa_{ij}} E_j. \tag{11}$$

Tariff revenue collected by destination j is therefore:

$$TR_j = \sum_i (\kappa_{ij} - 1) X_{ij} = (1 - \Lambda_j) E_j, \tag{12}$$

where the wedge Λ_j is the share of destination j's tariff-inclusive expenditure that is received by producers at producer prices:

$$\Lambda_j = \sum_i \frac{\pi_{ij}}{\kappa_{ij}}. \tag{13}$$

Equivalently, $1 - \Lambda_j$ is the share of expenditure collected as tariff revenue. When $\kappa_{ij} = 1$ for all sources, $\Lambda_j = 1$ and tariff revenue is zero.

Producer revenue in country i is the sum of producer-price sales across destinations:

$$Y_i = \sum_j X_{ij} = w_i L_i + \Pi_i.$$

This holds because in equilibrium producer revenue is distributed to domestic labor as factor income and profits. Combining this identity with the budget constraint and

the expression for tariff revenue gives:

$$E_i = Y_i + (1 - \Lambda_i)E_i + D_i.$$

Thus the equilibrium trade-balance condition is:

$$\Lambda_i E_i = Y_i + D_i. \tag{14}$$

This allows me to define equilibrium.

Definition 1. For a fixed market structure $(\eta, \mathcal{S})$, technology parameters $\{H_i\}$, common parameters (θ, ρ), labor endowments $\{L_i\}$, trade costs $\{\tau_{ij}\}$, tariffs $\{\kappa_{ij}\}$, and deficit parameters $\{d_i\}$, an equilibrium is a vector of wages $\{w_i\}$, final expenditures $\{E_i\}$, producer revenues $\{Y_i\}$, firm-location shares $\{\mu_i\}$, market-access terms $\{\mathcal{M}_i\}$, trade shares $\{\pi_{ij}\}$, and price indices $\{P_i\}$ such that:

(i) *Trade and prices.* Trade shares, market access, and price indices satisfy (7), (8), and (9).

(ii) *Location.* Firm-location shares satisfy the expression in Proposition 1.

(iii) *Income and expenditure.* Producer revenues and final expenditures satisfy (10)-(14).

5 Welfare

This section derives welfare implications of the unified representation using sufficient statistics. First, as in the standard gains-from-trade formula of Arkolakis et al. (2012), welfare changes depend on changes in the domestic expenditure share and the trade elasticity. Second, as in Demidova and Rodríguez-Clare (2009), Felbermayr et al. (2013), and Lashkaripour (2021), tariff revenues will also shape welfare. Most importantly, welfare depends on an additional sufficient statistic: the change in country i's industrial base.

I measure welfare in country i by real final expenditure per worker defined as:

$$W_i \equiv \frac{E_i}{P_i L_i}.$$

I assume throughout that the trade deficit is fixed relative to producer income ac-

cording to the following country-specific constants:

$$d_i = \frac{D_i}{Y_i}.$$

For a discussion of how this assumption affects welfare estimates, see Costinot and Rodríguez-Clare (2014). Using the trade-balance condition in (14) and the factor-income identity $w_i L_i = (1 - \mathcal{S})Y_i$, welfare can be written as follows:

$$W_i = \underbrace{\frac{w_i}{P_i}}_{\text{real wage}} \times \underbrace{\frac{1}{\Lambda_i}}_{\text{tariff wedge}} \times \underbrace{\frac{1 + d_i}{1 - \mathcal{S}}}_{\text{country constant}}.$$

Thus welfare depends on the real wage w_i/P_i, the tariff-revenue wedge Λ_i, and country-specific constant term that consists of the deficit-to-income ratio d_i and the profit share $\mathcal{S}$. Together with the expression for the intra-trade shares, this representation leads to the following proposition that derives sufficient-statistic welfare formula.

Proposition 2. *Let $\Delta \ln x_i \equiv \ln x_i' - \ln x_i$ denote the log change between two equilibria. Holding H_i, d_i, and the market structure fixed, the change in welfare is as follows:*

$$\Delta \ln W_i = -\frac{1}{\eta} \Delta \ln \pi_{ii} + \left(\frac{1}{\eta} - \frac{1}{\theta \rho} \right) \Delta \ln \mu_i - \Delta \ln \Lambda_i. \tag{15}$$

Proof. See Appendix C. □

The gains from trade formula in equation (15) isolates three relevant welfare channels. The first term $-\frac{1}{\eta} \Delta \ln \pi_{ii}$ is the ACR statistics with the trade elasticity η determined by the market structure: $\eta = \theta$ under cost competition and $\eta = \sigma - 1$ under monopolistic competition. The second term is the industrial-base term, which is the novel channel that appears due to the endogenous-location mechanism. Finally, the term $-\Delta \ln \Lambda_i$ is the tariff-revenue channel.

The gains-from-trade formula in equation (15) isolates three welfare channels. The first term, $-\frac{1}{\eta} \Delta \ln \pi_{ii}$, is the ACR statistic, with the trade elasticity η determined by market structure: $\eta = \theta$ under cost competition and $\eta = \sigma - 1$ under monopolistic competition. The second term is the industrial-base channel, which arises from endogenous location. Under cost competition its coefficient reduces to $(1 - 1/\rho)/\theta$,

which is precisely the contribution of $\Delta \ln T_i$ with $T_i = H_i \mu_i^{1-1/\rho}$; the general coefficient $1/\eta - 1/(\theta\rho)$ extends this to monopolistic competition. The third term, $-\Delta \ln \Lambda_i$, is the tariff-revenue channel.

The domestic-expenditure-share and tariff-revenue terms are directly observable, whereas the allocation of producers across countries, μ_i, may be harder to observe directly. The structure of the model relates μ_i to an observable sufficient statistic. Summing X_{ij} in equation (11) across destinations and using (8) gives $Y_i = H_i^{\eta/\theta} \mu_i^{1-\eta/(\theta\rho)} w_i^{-\eta} \mathcal{M}_i$. Them using the location-share equation in Proposition 1 and imposing $\sum_i \mu_i = 1$ gives the following:

$$\mu_i = \frac{Y_i}{\sum_k Y_k}. \tag{16}$$

Hence, in equilibrium the share of producers located in country i therefore equals its share of world output. The intuition behind this identify is a no-arbitrage condition – in equilibrium, the expected profit of a producer is the same in every country. Attractive locations draw more producers, but the marginal entrant is less capable than the incumbents, so a larger producer base comes with lower average productivity. This continues until average profit per producer is equalized across countries, and since profits are a common fraction of revenue, so is revenue per producer: $Y_i/\mu_i = \sum_k Y_k$. A country's output share therefore coincides with its producer share, with differences in market access and fundamentals absorbed by the number and composition of producers.

Next, I consider several special cases of the general formula in equation (15). In case of perfect and Bertrand compeition, $\eta = \theta$, so (15) becomes:

$$\Delta \ln W_i = -\frac{1}{\theta} \Delta \ln \pi_{ii} + \frac{\rho - 1}{\theta\rho} \Delta \ln \mu_i - \Delta \ln \Lambda_i. \tag{17}$$

The location coefficient $\frac{\rho-1}{\theta\rho}$ vanishes when technology draws are independent across locations, $\rho = 1$. Moreover, under the assumption of balanced trade and no tariffs, (17) collapses to the standard ACR sufficient statistic.

Under monopolistic competition, $\eta = \sigma - 1$, and the coefficient on the location term is:

$$\frac{1}{\eta} - \frac{1}{\theta\rho} = \frac{1}{\sigma - 1} - \frac{1}{\theta\rho}.$$

This coefficient is strictly positive at $\rho = 1$ under the standard restriction $\theta > \sigma - 1$.

Hence, unlike in the perfect-competition and Bertrand cases, the endogenous-location channel remains active under independent draws in monopolistic competition. Its quantitative strength, however, depends on ρ and must therefore be disciplined empirically.

6 Alternative Distributions

The selection mechanism developed in this paper is not unique to the multivariate Fréchet distribution. The key requirement is that equilibrium location shares tractably determine the conditional productivity distribution of firms selecting into each location. The Fréchet specification is convenient because it preserves standard closed-form aggregation, but other distributions can generate the same selection logic.

I also consider an alternative distributional foundation that connects the framework to the literature on power-law heterogeneity (see, e.g., Gabaix 2009). Proposition 3 studies a multivariate Dagum distribution with Clayton dependence. This distribution delivers the same economic mechanism: firm location choices induce conditional productivity distributions, and the dependence structure governs how the mass and composition of firms interact in aggregate productivity. So, while the main quantitative analysis uses the Fréchet formulation because it nests directly in canonical quantitative trade models, the power-law case shows that the mechanism is not specific to extreme-value heterogeneity.

Proposition 3. *Suppose the vector of shocks $\boldsymbol{z} = (z_1, \ldots, z_K)$ follows a Dagum distribution with Clayton-type dependence:*

$$F(\boldsymbol{z}) = \left(1 + \sum_{k=1}^{K} \left[\frac{z_k}{H_k}\right]^{-\xi}\right)^{-1/\delta}, \qquad H_k > 0,\ \xi > 0,\ \delta > 0,$$

where parameter ξ governs the power-law tail of post-selection heterogeneity, while δ governs dependence across alternatives. Agents choose the alternative k that maximizes the payoff in equation (3). Then:

(i) The choice probabilities are given by:

$$\mu_i = \frac{H_i^{\xi}\, w_i^{-\xi}\, \mathcal{M}_i^{\xi/\eta}}{\sum_k H_k^{\xi}\, w_k^{-\xi}\, \mathcal{M}_k^{\xi/\eta}}.$$

(ii) The conditional distribution of z_k given selection $(k^* = k)$ is Dagum:

$$F_k(z \mid k^* = k) = \left[1 + \left(\frac{z}{\tilde{D}_k}\right)^{-\xi}\right]^{-1/\delta}, \qquad , \text{ where} \qquad \tilde{D}_k = H_k \, \mu_k^{-1/\xi}.$$

(iii) If $\xi > 1$, the expected shock conditional on selection is

$$\mathbb{E}[z_k \mid k^* = k] = \frac{\Gamma(1 - 1/\xi)\,\Gamma(1/\delta + 1/\xi)}{\Gamma(1/\delta)} \, H_k \, \mu_k^{-1/\xi},$$

where $\Gamma(\cdot)$ denotes the Gamma function.

Proof. The proof is available in Appendix D. □

The Fréchet and Dagum specifications are tractable for the same economic reason. In both cases, the productivity vector can be interpreted as having two components:

$$z_k = \bar{Z} \, Z_k.$$

The common component $\bar{Z}$ raises or lowers productivity in all locations for a given firm, while the idiosyncratic component Z_k determines which location is relatively most attractive. Since $\bar{Z}$ is common across locations, it cancels from the location decision. The firm's choice is therefore governed by the relative draws Z_k, while $\bar{Z}$ determines how strongly productivity draws are correlated across locations.

The two distributions differ in how dependence affects relative draws. In the logistic-Fréchet case, stronger dependence changes the effective tail of the idiosyncratic component that ranks locations. The dependence parameter therefore enters the location shares directly. In the Clayton–Dagum case, dependence operates differently. It changes the joint distribution and the moments conditional on selection, but it does not change the tail parameter that governs relative choice probabilities. As a result, the dependence parameter appears in the conditional productivity distribution only through terms that are common across alternatives. Taken together, the two cases cover the two main distributional foundations for firm heterogeneity used in the literature: extreme-value heterogeneity and power-law heterogeneity.

7 Estimation

This section quantifies the importance of the industrial-base channel generated by endogenous firm location in otherwise canonical models of international trade. Proposition 2 shows that, relative to standard sufficient-statistics approach based on bilateral trade flows X_{ij}, the model requires two additional inputs: measures of each country's industrial base, μ_i, and the structural parameters $\{\eta, \theta, \rho\}$.

I combine several data sources to structurally estimate the firm-location shares μ_i and the parameters $\{\eta, \theta, \rho\}$ that jointly determine the quantitative importance of the industrial-base channel. International trade flows and domestic production are from the International Trade and Production Database described in Borchert et al. (2021) and Larch et al. (2025). I focus on manufacturing trade flows over 1992-2019 and exclude dependent territories as well as retain economies with at least one positive recorded domestic or international manufacturing flow in 1992. This restriction establishes a common geographical baseline for gravity estimation and the subsequent construction of market access and instruments. International tariff data are from Feodora Teti's Global Tariff Database described in Teti (2024). Bilateral distance is from the CEPII database described in Mayer and Zignago (2011). In estimation, I also use country-level controls from the Penn World Table v.11 described in Feenstra et al. (2015).

I proceed in two steps. The first step estimates the trade elasticity η from tariff variation in a structural gravity equation. This step also recovers exporter-year fixed effects, importer-year fixed effects, and the bilateral component of trade costs. The second step uses these recovered gravity components to construct market-access terms and estimate the composite parameter that governs the strength of the industrial-base channel in the welfare formula.

While the model is static, the estimation uses panel variation in the data. Let t index years. The empirical counterpart of the producer-price trade flow in equation (11) is estimated by Poisson Pseudo-Maximum Likelihood:

$$X_{ij,t} = \exp\left(\upsilon_{i,t} + \log b_{ij,t} + \chi_{j,t}\right) + e_{ij,t}, \tag{18}$$

where $\upsilon_{i,t}$ is an exporter-year fixed effect, $\chi_{j,t}$ is an importer-year fixed effect, and $b_{ij,t}$ is the bilateral component that captures all trade barriers including tariffs parame-

terized as follows:

$$\log b_{ij,t} = \psi_{ij} + \varphi_{q(ij)}t + \beta \log \kappa_{ij,t},$$

where ψ_{ij} is a pair fixed effect, $q(ij)$ denotes the log-distance decile of pair ij such that $\varphi_{q(ij)}t$ allows for different trends across distance bins. Since in the data $X_{ij,t}$ is measured in CIF terms but net of tariffs, the model implies:

$$\beta = -(\eta + 1).$$

Thus, the trade elasticity parameter can be recovered as $\widehat{\eta} = -\widehat{\beta} - 1$ reported in Table 1. The tariff coefficient is -3.6, implying $\widehat{\eta} = 2.6$. This estimate sits between the two main strands of the literature: above the lower long-run trade elasticity of 1.75–2.25 estimated by Boehm et al. (2023), and below the value of 3.6 used in Eaton and Kortum (2002) and Bernard et al. (2003).

Table 1: First-Stage Structural Gravity Estimation

	Coefficient
$\log \kappa_{ij,t}$	-3.59
	(0.28)
Implied $\widehat{\eta}$	2.59
Exporter-year fixed effects	Yes
Importer-year fixed effects	Yes
Pair fixed effects	Yes
Distance-decile trends	Yes
Observations	745,108
Pair clusters	26,611

Notes: The regression is estimated by PPML and includes exporter-year, importer-year, and pair fixed effects as well as log-distance-decile-specific linear trends. Robust standard errors clustered by country pair are reported in parentheses.

As derived by Fally (2015), estimating the gravity equation via PPML automatically satisfies the adding up constraints, which together with the theoretical identities of exporter and importer fixed effects allow to recover a direct measure of market access:

$$\widehat{\mathcal{M}}_{i,t} = \sum_j \widehat{b}_{ij,t} \exp(\widehat{\chi}_{j,t}).$$

Next, I log-linearize the expression for firms' allocation shares from Proposition 1 and measure $\mu_{i,t}$ according to equation (16) as i's share in total world output in year t.

Then, I estimate the following panel model:

$$\mu_{i,t} = \text{country}_i + \text{year}_t + \gamma \log \widehat{\mathcal{M}}_{i,t} + \mathbf{W}'_{i,t}\Gamma + u_{i,t}, \tag{19}$$

where country_i and year_t are exporter and year fixed effects. The coefficient on log market access identifies $\gamma = \left(\frac{\theta\rho}{\eta}\right)$. Note that under Bertrand and perfect competiti-ton, $\eta = \theta$ such that ρ can be inferred from γ. The vector $\mathbf{W}_{i,t}$ contains country-level controls intended to proxy for variation in $H_{i,t}$ and $w_{i,t}$ not absorbed by country and year fixed effects. These controls comprise log GDP per capita, log capital per worker, the human capital index, the log output price level, the labor share, log total factor productivity, and log terms of trade, all constructed from Penn World Table data.

Table 2: Second-Stage OLS and IV Estimation

	OLS		IV (I)		IV (II)	
	(1)	(2)	(3)	(4)	(5)	(6)
Coefficient γ	1.64	1.74	1.38	1.91	1.59	1.99
Standard errors	(0.13)	(0.24)	(0.43)	(0.31)	(0.72)	(0.49)
Exporter fixed effects	Yes	Yes	Yes	Yes	Yes	Yes
Year fixed effects	Yes	Yes	Yes	Yes	Yes	Yes
Controls	No	Yes	No	Yes	No	Yes
Observations	4,704	2,972	4,704	2,972	4,704	2,972
Exporter clusters	168	107	168	107	168	107
Kleibergen–Paap F	–	–	28.41	37.45	14.79	12.78

Notes: The dependent variable is $\log \mu_{i,t}$ is country i's (log) share of world manufacturing output, computed from total sales across all destinations, including domestic sales. Columns 1–2 report OLS estimates. Columns 3–4 instrument market access with predicted foreign access constructed from the full estimated bilateral component, including distance-bin trends, and base-year importer effects. Columns 5–6 use MFN-tariff-predicted foreign access and base-year importer effects, excluding estimated distance-bin trends from the instrument. Both instruments exclude domestic sales. Controls include log GDP per capita, log capital per worker, human capital, log output price level, the labor share, log TFP, and log terms of trade. Differences in sample sizes reflect the availability of controls. Standard errors in parentheses are clustered by exporter.

I first estimate the panel model in equation (19) by OLS, with and without controls, and report the results in Table 2. Depending on the specification, the coefficient on $\log \widehat{\mathcal{M}}_{i,t}$ varies between 1.64 and 1.74. Under the assumption $\eta = \theta$, these coefficients directly estimate ρ. Both estimates are relatively precise and indicate positive dependence in location-specific productivity draws, exceeding the independence benchmark $\rho = 1$.

OLS estimates, however, may be biased for two reasons. First, measurement error in estimated market access, $\widehat{\mathcal{M}}_{i,t}$, may attenuate the coefficient toward zero. Second, contemporaneous shocks may jointly affect market access and output shares, $\mu_{i,t}$. To address these concerns, I use two model-consistent instruments based on predicted variation in foreign market access. Identification requires these instruments to be orthogonal to the structural disturbance and to measurement error in market access, conditional on the included controls and fixed effects.

The first instrument fixes destination attractiveness at the start of the panel and relies on the time-series variation in the estimated bilateral component:

$$Z_{i,t}^{(I)} = \sum_{j \neq i} \exp\Big(\widehat{\psi}_{ij} + \widehat{\varphi}_{q(ij)}(t - t_0) - (\widehat{\eta} + 1)\log(\kappa_{ij,t})\Big) \exp(\widehat{\chi}_{j,t_0}). \tag{20}$$

The instrument excludes the domestic market and removes contemporaneous destination-demand movements from market access. Its time-series variation comes from the estimated evolution of bilateral trade costs, including distance-decile trends and tariff changes, weighted by destination attractiveness in the base year. The identifying assumption is that the resulting variation in $Z_{i,t}^{(I)}$ affects the recovered exporter-year component only through market access.

The second instrument is more conservative. It replaces actual bilateral tariffs with MFN tariffs and omits the estimated distance-decile trend:

$$Z_{i,t}^{(II)} = \sum_{j \neq i} \exp\Big(\widehat{\psi}_{ij} - (\widehat{\eta} + 1)\log(\kappa_{ij,t}^{MFN})\Big) \exp(\widehat{\chi}_{j,t_0}). \tag{21}$$

The intuition follows Boehm et al. (2023). While destination j may be able to target tariffs toward a specific partner i, changes in MFN rates apply to many partners and are therefore less likely to reflect bilateral shocks specific to i. The second instrument therefore uses only pair-specific exposure, base-year destination attractiveness, and MFN-tariff-predicted changes in foreign market access. This removes contemporaneous destination-demand variation, excludes the domestic market, and eliminates variation coming from preferential bilateral tariffs.

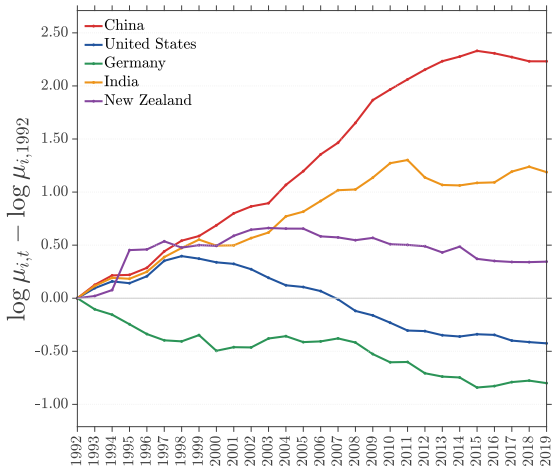
Using $Z_{i,t}^{(I)}$, which combines the full estimated bilateral component, including distance-bin trends, with base-year importer effects, the estimated coefficient on market access

is 1.38 without controls and 1.91 with controls. The corresponding Kleibergen–Paap first-stage F-statistics are 28.41 and 37.45. Both estimates are positive and precise. The MFN-based instrument $Z_{i,t}^{(II)}$, which excludes estimated distance-bin trends, yields coefficients of 1.60 without controls and 1.98 with controls, with corresponding Kleibergen–Paap F-statistics of 14.79 and 12.78. Thus, across all four IV specifications, the first-stage F-statistics exceed 10 and the estimated γ is positive and statistically significant at the 5 percent level.

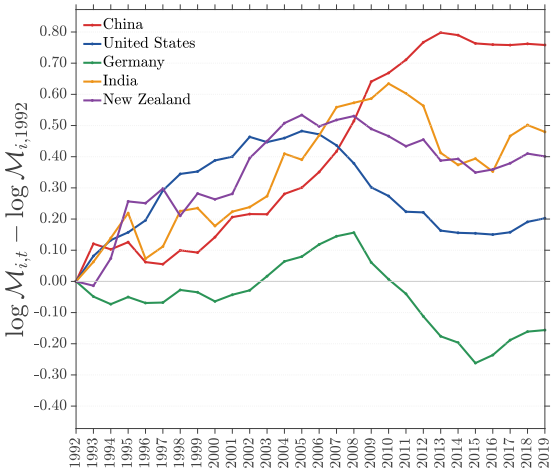
7.1 Evolution of Market Access and Producer Shares

The estimation procedure recovers market access, $\widehat{\mathcal{M}}_{i,t}$, while the model's equilibrium identity between producer shares and output shares provides an empirical measure of $\mu_{i,t}$. Figure 2 presents their evolution for selected countries over 1992–2019. Panel A plots manufacturing output shares interpreted through the model as producer-location shares, and Panel B plots estimated market access. Both panels report log changes relative to each country's 1992 value. Together, they illustrate how changes in market access accompany the geographical redistribution of manufacturing activity.

(a) Panel A: Industrial Base



(b) Panel B: Market Access



Notes: Each series is expressed as a log change relative to 1992. The left panel reports changes in the producer-location share, $\ln \mu_{i,t} - \ln \mu_{i,1992}$. The right panel reports changes in recovered market access, $\ln \widehat{\mathcal{M}}_{i,t} - \ln \widehat{\mathcal{M}}_{i,1992}$.

Figure 2: Changes in Mass of Producers and Market Access

China and India experience substantial increases in both market access and their recovered location shares, consistent with the model's prediction that better market

access attracts manufacturing producers. China's expansion is evident in the 1990s and continues through the following decade, suggesting that WTO accession in 2001 might have reinforced an existing process by improving access to foreign markets. By 2019, China's location share is 2.23 log points above its 1992 level, alongside a market-access increase of 0.76 log points. Its location share continues rising until 2015, two years after market access peaks, before declining modestly. India's corresponding long-run increases are 1.19 and 0.48 log points. These patterns illustrate the economic distinction between fixed industrial bases and endogenous producer location. With an exogenously fixed allocation of producers, $\mu_{i,t}$ would not respond to changes in market access; with endogenous location, better market access raises the attractiveness of producing in a country, with the strength of this channel governed by $\theta\rho/\eta$, holding other determinants constant. The observed co-movement is consistent with this mechanism, although the figure cannot separate its contribution from contemporaneous changes in productivity, production costs, and other conditions affecting manufacturing location.

Among the developed economies shown, market access deteriorates around the global financial crisis and accompanies a longer-run decline in the manufacturing location shares of the United States and Germany that was already underway before the crisis. By 2019, US market access remains 0.20 log points above its 1992 level, while its location share is 0.43 log points lower; Germany records declines of 0.16 and 0.80 log points, respectively. These differences highlight the relative nature of producer location: improved sales opportunities need not preserve a country's share of manufacturing activity when competing locations become more attractive through changes in their own market access, productivity, or production costs. New Zealand experiences a smaller immediate decline in market access during the crisis and retains gains in both variables over the full period, with market access and its location share respectively 0.40 and 0.34 log points above their 1992 levels in 2019. These contrasting trajectories motivate the structural analysis, which provides a framework for studying systematically how market access and production fundamentals jointly shape the international allocation of manufacturing activity.

7.2 Model-Implied Location Shares versus Data

The natural object for external validation is $\mu_{i,t}$, the model-implied share of producers that choose country i as their location and hence a key determinant of its industrial

base. Among the objects recovered from the model, $\mu_{i,t}$ is most directly tied to the endogenous-location mechanism and admits an independent empirical counterpart. Other recovered objects, such as market access, are latent gravity constructs or aggregates that can arise even in models without endogenous location choice. I therefore validate the mechanism by asking whether the model's predictions that location shares, $\mu_{i,t}$, calculated as the share of world output of i in year t align with independent measures of the producer base.

I construct the external counterpart to $\mu_{i,t}$ from the UNIDO INDSTAT data on the number of manufacturing establishments, which can be considered the closest empirical analogue to the model's measure of producers. I harmonize ISIC Revision 3 and ISIC Revision 4 into a single panel data sample to maximize coverage. Let $N_{i,t}$ denote the resulting external measure of the number of manufacturing establishments in country i and year t. I convert this object into an empirical location share by normalizing across countries within each year:

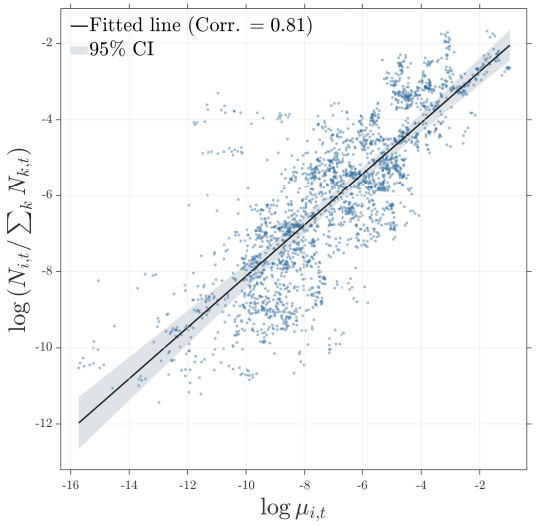
$$\mu_{i,t}^{\text{ext}} = \frac{N_{i,t}}{\sum_k N_{k,t}}.$$

This normalization matches the model object because $\mu_{i,t}$ is the within-period share of the world mass of producers located in country i. The external measure is constructed entirely from data that are neither used nor targeted in the estimation and is therefore suitable for external validation.

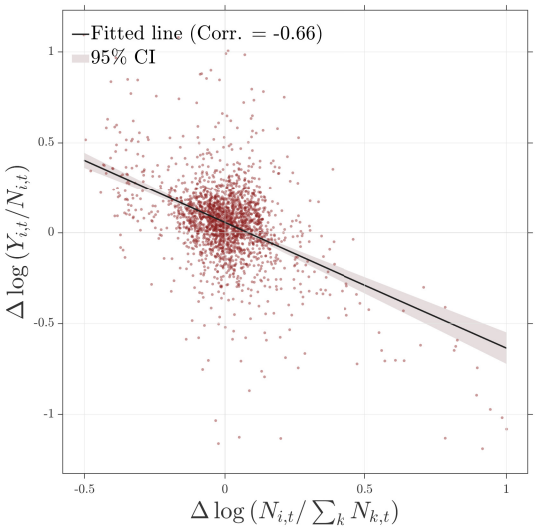
Panel A in Figure 3 plots model-consistent $\mu_{i,t}$ versus externally calculated $\mu_{i,t}^{\text{ext}}$. The relationship is strongly positive, with a correlation of 0.81, suggesting that the mass of producers predicted by the model captures meaningful variation in the independent establishment-based measure of where producers are located. This relationship provides meaningful external validation of the endogenous-location mechanism, since the recovered production-base variable is constructed from trade flows while the external measure is constructed independently from establishment data.

Panel B of Figure 3 provides a second external validation exercise. It uses the model prediction that average revenue productivity among firms located in country i declines with the location share $\mu_{i,t}$. From equation (4), under the assumption that year-to-year changes in the fundamental component $H_{i,t}$ are absent, the relevant conditional

(a) Panel A: Mass of Producers



(b) Panel B: Revenue Productivity



Notes: Panel A compares (log) model-implied location shares with (log) observed manufacturing-firm shares. Panel B relates annual (log) changes in revenue per firm to changes in firm shares. Firm counts are from UNIDO. The sample covers 1992–2019, follows the baseline gravity country restrictions, and normalizes firm counts over reporting countries each year. Panel B uses consecutive-year changes and excludes observations outside either variable's first and ninety-ninth percentiles from the displayed sample; its reported correlation uses the full sample. Solid lines are pooled OLS fits to displayed observations; shaded areas are pointwise 95% confidence intervals clustered by country.

Figure 3: External Validation

(log) revenue-productivity term is as follows:

$$\Delta \ln \mathcal{R}_{i,t} = -\frac{\eta}{\theta\rho}\Delta \ln \mu_{i,t}.$$

The intuition is that as more firms choose to locate in country i, selection becomes weaker and the marginal entrants are less productive than the firms already located there. Average revenue productivity therefore falls as the domestic firm base expands. To assess this prediction, I use average manufacturing output per firm as an empirical proxy for average revenue productivity and plot its annual log change against the annual log change in the country's share of the global manufacturing firms. To keep the figure legible, observations outside the first and ninety-ninth percentiles of either variable are omitted from the scatterplot, while the reported correlation is computed using the full sample. The correlation is strongly negative at -0.66, consistent with the model's prediction that expansions of the firm base are accompanied by lower average productivity among located firms.

7.3 Quantifying the Welfare Gains from Trade

The estimation procedure recovers β and γ, identifying the trade elasticity $\eta = -\beta - 1$ and the composite parameter $\gamma = \theta\rho/\eta$. Together, these parameters suffice to calculate welfare changes without separately identifying θ and ρ. Holding productivity fundamentals and deficit-to-income ratios fixed, Proposition 2 implies:

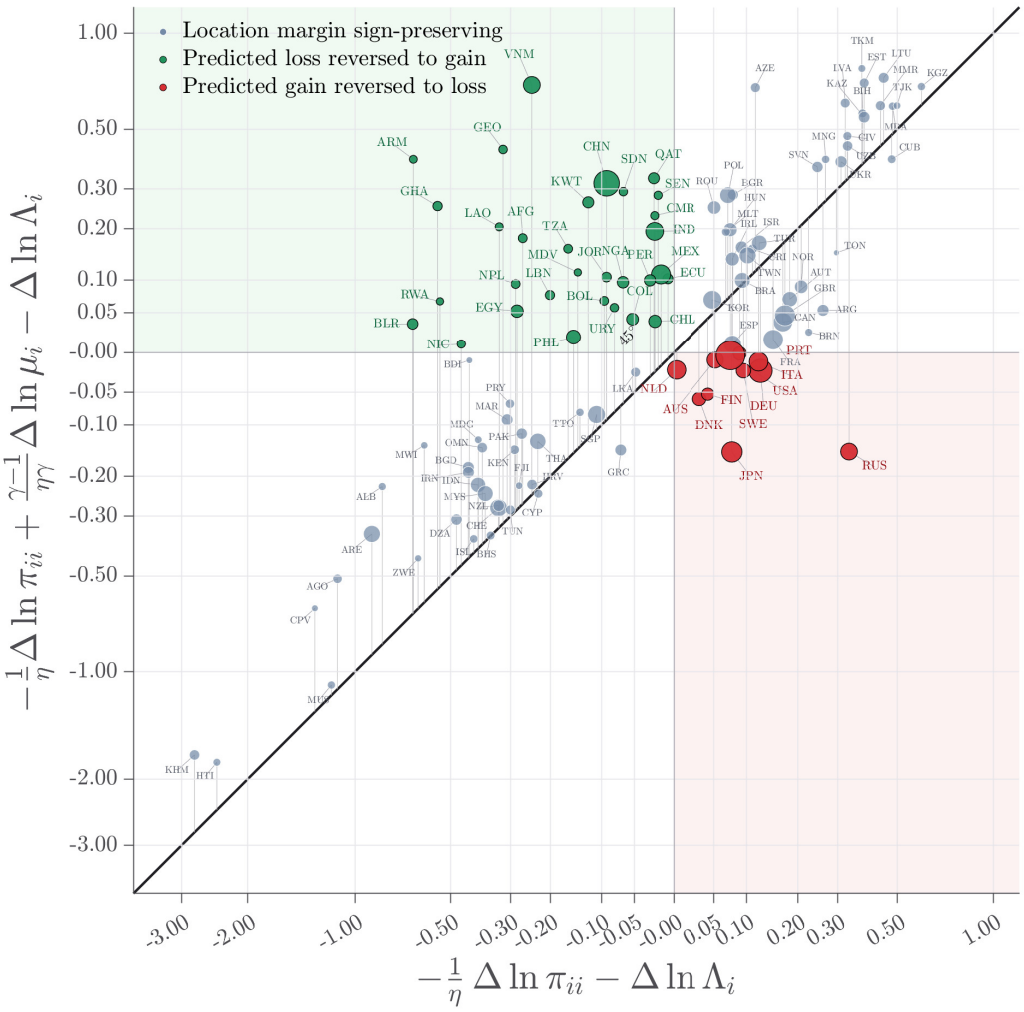
$$\Delta \ln W_i = \underbrace{-\frac{1}{\eta}\Delta \ln \pi_{ii}}_{\text{domestic share}} + \underbrace{\frac{\gamma - 1}{\eta\gamma}\Delta \ln \mu_i}_{\text{industrial base}} - \underbrace{\Delta \ln \Lambda_i}_{\text{tariff wedge}} .$$

Welfare changes therefore depend on the observable changes $\{\Delta \ln \pi_{ii}, \Delta \ln \mu_i, \Delta \ln \Lambda_i\}$ and the estimated parameters $\{\eta, \gamma\}$. The same sufficient-statistics formula applies under perfect competition, Bertrand competition, and monopolistic competition.

I compute model-implied welfare changes and their three components using long differences between 1992 and 2019 for 114 countries with the required data at both endpoints. Trade and tariff data determine changes in domestic expenditure shares and tariff wedges. As before, conssitent with the mode, the industrial-base share is measured as $\widehat{\mu}_{i,t} = Y_{i,t} / \sum_\ell Y_{\ell,t}$, where $Y_{i,t}$ denotes total manufacturing sales. Figure 4 plots full welfare changes against those calculated using only the ACR term and the tariff wedge. The difference between these two welfare measures is exactly the novel industrial-base contribution term developed in this paper.

First, the industrial-base contribution is comparable in magnitude to the conventional sufficient statistics. Across the 114 countries, the unweighted arithmetic mean of country-level percentage welfare changes is approximately 12.0%, and the median is 5.7%. The mean absolute ACR contribution is 0.27 log points, compared with 0.24 log points for the industrial-base contribution and 0.05 log points for the tariff wedge. In 57 countries, the industrial-base contribution exceeds the ACR contribution in absolute value. For the United States, the contraction of the industrial base contributes -0.078 log points, slightly outweighing the combined 0.074-log-point gain from the ACR and tariff-wedge terms. For China, the expansion of the industrial base contributes 0.41 log points, more than offsetting a combined loss of 0.09 log points through the other two channels.

Second, accounting for the industrial base reverses the sign of the measured welfare change in 42 countries, or more than one-third of the sample. In 31 countries, the



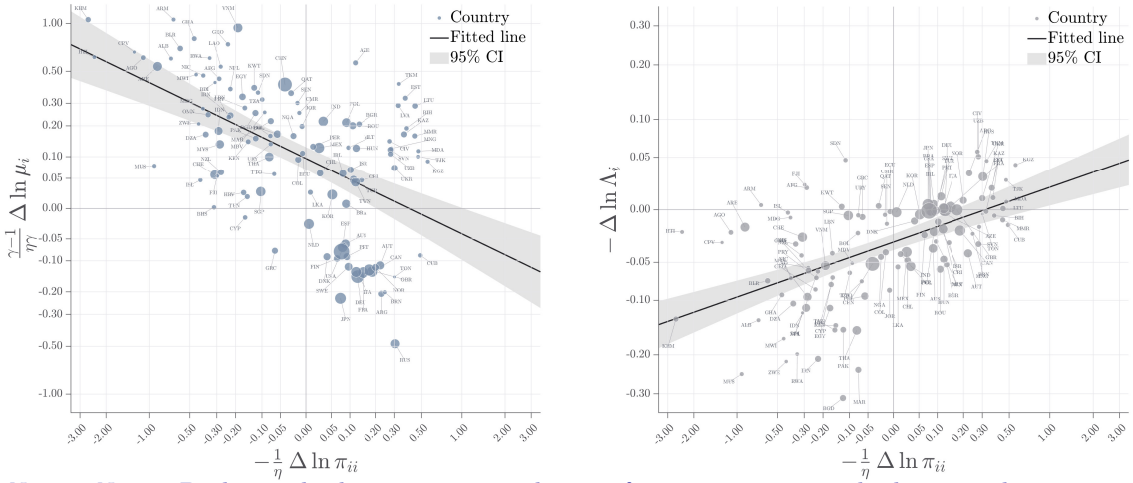
Notes: Welfare change with versus without the endogenous location channel. The horizontal axis omits the industrial-base term, whereas the vertical axis is the full change in welfare. Distance from the 45° line is the industrial-base contribution. Marker size is proportional to 2019 imports and axes use a symlog scale.

Figure 4: Welfare Gains from Trade 1992-2019

standard components imply a welfare loss, whereas the full measure implies a gain. These countries appear in the upper-left region of Figure 4 and include China, Mexico, Peru, and Vietnam. In another 11 countries, the standard components imply a gain, but the contraction of the industrial base produces a loss under the full measure. These countries appear in the lower-right region and include the United States, Japan, Germany, and Sweden. Changes in the domestic industrial base therefore matter for both the magnitude and the direction of the model-implied welfare effects.

The industrial-base component can reverse the sign of measured welfare changes because it often dominates the ACR component. A decline in the domestic expenditure share generates a positive ACR contribution, but when accompanied by a contraction in the country's share of world manufacturing output, it also generates a negative industrial-base contribution. However, the two margins do not necessarily move in opposite directions – a country can expand its industrial base while becoming more reliant on imports. Panel A of Figure 5 examines this relationship by plotting the industrial-base contribution against the ACR contribution. Panel B plots the tariff-wedge contribution against the ACR contribution.

(a) Panel A: Industrial Base Term (b) Panel B: Tariff Wedge Term



Notes: Notes: Both panels show 1992–2019 changes for 114 countries. The horizontal axis reports the ACR contribution, $-(1/\eta)\Delta \ln \pi_{ii}$. The vertical axis reports the industrial-base contribution, $(\gamma-1)/(\eta\gamma)\Delta \ln \mu_i$, in Panel A and the tariff-wedge contribution, $-\Delta \ln \Lambda_i$, in Panel B. Marker size is proportional to 2019 imports. Axes use a symlog scale.

Figure 5: Welfare Decomposition Channels Across Countries

Panel A shows a negative correlation of -0.59 between the industrial-base and ACR contributions. Across the 114 countries observed in both 1992 and 2019, these contributions have opposite signs in 79 countries. In 58 countries, including China and Vietnam, the domestic expenditure share rises alongside an expanding industrial-base share, producing a negative ACR contribution and a positive industrial-base contribution. Conversely, 21 countries, including France, Japan, and the United States, experience declining domestic expenditure shares alongside shrinking industrial-base shares. Yet greater import penetration does not necessarily coincide with industrial-base contraction: 33 countries gain on both margins, including Chile, Mexico, Poland,

and Turkey. These patterns illustrate how accounting for producer location changes the welfare interpretation of movements in domestic expenditure shares.

Panel B shows a positive correlation of 0.36 between the tariff-wedge and ACR contributions. One accounting mechanism can generate this association: holding the import-weighted tariff rate fixed at a positive value, greater import penetration lowers the domestic expenditure share and raises the fraction of expenditure subject to tariffs, producing positive contributions through both terms. The two contributions have the same sign in 72 countries, so the tariff-wedge term more often reinforces than offsets the ACR term. Its quantitative importance remains smaller, with a mean absolute contribution of 0.05 log points, compared with 0.27 for the ACR term.

Next, I repeat the decomposition exercise over alternative time windows that correspond to different phases of globalization: the benchmark 1992–2019 period, as well as 2001–2019, 2008–2019, 2012–2019, and 2015–2019. Table 3 decomposes welfare changes into the ACR, industrial-base, and tariff-wedge channels for each horizon. For each period, I report the average welfare change and the cross-country standard deviation. As expected, both statistics decline over shorter horizons.

Table 3: Welfare Decomposition in Different Time Periods

		$\Delta \ln \widehat{W}_i$		Channel contribution			Sign
Period	N	Average	Std. Dev.	Trade	Industrial base	Tariff	Reversals
1992–2019	114	+0.050	0.388	−0.092	+0.184	−0.042	36.8%
2001–2019	114	+0.076	0.147	+0.068	+0.017	−0.010	31.6%
2008–2019	114	−0.032	0.107	−0.020	−0.006	−0.007	42.1%
2012–2019	114	−0.030	0.198	−0.027	+0.002	−0.005	27.2%
2015–2019	114	+0.008	0.127	−0.008	+0.019	−0.003	24.6%

Notes: All windows use the same 114 countries and end in 2019. Entries report unweighted cross-country means and the sample standard deviation of log welfare changes. Channel means sum to the total. Parameters are $\eta = 2.59$ and $\gamma = 1.9$, with industrial-base weight $(\gamma - 1)/(\eta\gamma)$. "Sign reversals" is the percentage of all countries for which the full welfare change and the ACR-plus-tariff baseline have opposite signs.

The comparison across different time horizons shows that the positive industrial-base contribution is concentrated in the earlier part of the sample. For the same 114 countries, its average contribution is 0.184 log points over 1992–2019, compared with 0.017 over 2001–2019. The corresponding contributions are −0.006 over 2008–2019, 0.002 over 2012–2019, and 0.019 over 2015–2019. Because the country sample and parameters are held fixed, these comparisons imply that most of the positive

industrial-base contribution to the full-period mean accumulated during the 1990s.

Despite these smaller average contributions over later time windows, the industrial-base channel continues to change the direction of measured welfare effects for many countries. It reverses the sign of the welfare change implied by the ACR and tariff-wedge terms alone for 31.6% of countries over 2001–2019, 42.1% over 2008–2019, 27.2% over 2012–2019, and 24.6% over 2015–2019. Thus, the importance of the industrial-base channel for the sign of measured welfare changes therefore persists across all the horizons examined.

7.4 Counterfactual: The Benefits of Foreign Technology

In this section, I return to the counterfactual in Eaton and Kortum (2002) and Caliendo et al. (2018) who study how changes in technology propagate across space through trade. For that, I use the hat algebra approach as in Dekle et al. (2007), writing the equilibrium in proportional changes $\widehat{x} = x'/x$ relative to an observed baseline.

I consider an exogenous shock to the technology primitive in China, H_{CHN}. The counterfactual asks how much other countries would gain or lose if Chinese productivity were hit by a negative shock that returned it to its pre-WTO-accession level. To discipline the magnitude of this shock, I use the estimate in Brandt et al. (2017), who find that the tariff reductions associated with China's WTO accession raised Chinese manufacturing productivity by 10.8%. I take this estimate as the target for the counterfactual change in China's technology primitive and study its reversal: how would welfare in the rest of the world change if China's accession-era productivity gains were undone?

In this counterfactual exercise, I hold all other primitives at their observed values. In hat notation, this means setting changes in trade costs, tariffs, labor endowments as follows:

$$\widehat{\tau}_{ij} = 1, \qquad \widehat{\kappa}_{ij} = 1, \qquad \widehat{L}_i = 1,$$

while holding deficit ratios d_i fixed.

Specifically, the hat-algebra system solves for the equilibrium changes in wages and location shares, $\{\widehat{w}_i, \widehat{\mu}_i\}$. The baseline expenditure shares π_{ij}, location shares μ_i, bilateral flows X_{ij}, and levels Y_i and E_j are taken from the benchmark equilibrium.

I also use the baseline export shares:

$$s_{ij} \equiv \frac{X_{ij}}{Y_i}, \qquad \sum_j s_{ij} = 1.$$

To express the counterfactual system in terms of the estimated parameters η and γ, define the composite technology shock as $\widehat{A}_i \equiv \widehat{H}_i^{\eta/\theta}$. This is the technology shifter that enters trade shares and price indices. Because the primitive productivity scale is proportional to $H_i^{1/\theta}$, undoing China's 10.8% productivity gain implies $\widehat{H}_{\text{CHN}} = 1.108^{-\theta}$ such that:

$$\widehat{A}_{\text{CHN}} = \widehat{H}_{\text{CHN}}^{\eta/\theta} = 1.108^{-\eta}, \qquad \widehat{A}_i = 1 \quad \text{for } i \neq \text{CHN}.$$

Given these baseline objects and the counterfactual shocks, the counterfactual equilibrium solves conditions (i)–(vii):

$$(i) \quad \widehat{\pi}_{ij} = \frac{\widehat{A}_i \widehat{\mu}_i^{(\gamma-1)/\gamma} \left(\widehat{w}_i \widehat{\tau}_{ij} \widehat{\kappa}_{ij}\right)^{-\eta}}{\widehat{P}_j^{-\eta}}$$

$$(ii) \quad \widehat{P}_j^{-\eta} = \sum_k \pi_{kj} \widehat{A}_k \widehat{\mu}_k^{(\gamma-1)/\gamma} \left(\widehat{w}_k \widehat{\tau}_{kj} \widehat{\kappa}_{kj}\right)^{-\eta}$$

$$(iii) \quad \widehat{\mathcal{M}}_i = \sum_j s_{ij} \widehat{\tau}_{ij}^{-\eta} \widehat{\kappa}_{ij}^{-(\eta+1)} \widehat{P}_j^{\eta} \frac{E'_j}{E_j}$$

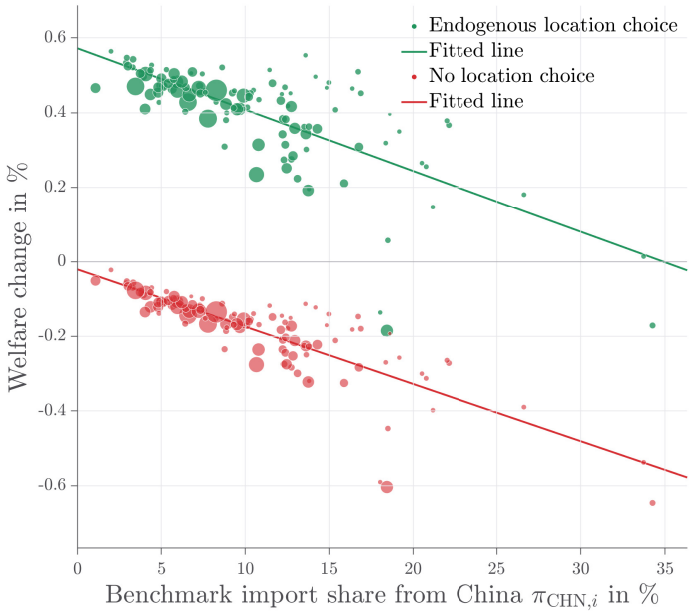
$$(iv) \quad \widehat{\mu}_i = \frac{\widehat{A}_i^{\gamma} \widehat{w}_i^{-\eta\gamma} \widehat{\mathcal{M}}_i^{\gamma}}{\sum_k \mu_k \widehat{A}_k^{\gamma} \widehat{w}_k^{-\eta\gamma} \widehat{\mathcal{M}}_k^{\gamma}}$$

$$(v) \quad \widehat{\Lambda}_i = \frac{1}{\Lambda_i} \sum_k \frac{\pi_{ki} \widehat{\pi}_{ki}}{\kappa_{ki} \widehat{\kappa}_{ki}}$$

$$(vi) \quad E'_i = \frac{(1+d_i) \widehat{w}_i Y_i}{\widehat{\Lambda}_i \Lambda_i}$$

$$(vii) \quad \widehat{w}_i = \frac{1}{Y_i} \sum_j \frac{\pi_{ij} \widehat{\pi}_{ij}}{\kappa_{ij} \widehat{\kappa}_{ij}} E'_j.$$

To isolate the contribution of the location margin, I compute each country's welfare change twice: once in the full model, in which firms relocate in response to the shock ($\widehat{\mu}_i$ endogenous), and once in a restricted version that holds the location distribution



Notes: Welfare change from reversing China's productivity gain. Red markers omit the industrial-base term; green markers report the full change. Solid lines are OLS fits. Marker size is proportional to the benchmark shares μ_i.

Figure 6: Welfare and China Technology Shock

fixed ($\widehat{\mu}_i = 1$), which corresponds to the sufficient-statistic computation standard in the prior literature. The difference between the two is the industrial-base contribution.

Figure 6 plots each country's benchmark import share from China against the welfare effect of reversing China's WTO-accession-era productivity gain. The figure highlights the importance of the endogenous-location channel. In the no-location counterfactual, shown in red, the shock lowers welfare for most countries, with larger losses among those that initially import more from China. This is the standard foreign-technology channel: a negative productivity shock in China raises the cost of Chinese goods and reduces welfare by worsening access to foreign varieties.

Once firm location is endogenous, however, the welfare effects are largely reversed. The negative shock to China makes it a less attractive production location and reallocates firms toward the rest of the world. For most countries, the resulting expansion of the industrial base more than offsets the loss from higher import prices, generating positive welfare effects in the full model, shown in green. The contrast between the red and green series underscores the central mechanism of the paper: the welfare

consequences of foreign productivity growth depend not only on import exposure, but also on the induced reallocation of production across countries.

8 Conclusions

This paper develops a tractable model of endogenous firm location and embeds it in canonical quantitative trade frameworks. Firms draw location-specific productivities from a multivariate distribution and choose where to produce. A country's industrial base, defined by the mass and productivity distribution of producers located there, is therefore an equilibrium object rather than a primitive. This location margin introduces a new channel for the gains from trade through changes in the domestic industrial base. I estimate the model structurally to quantify the importance of this channel.

The industrial-base margin is quantitatively first-order. Over 1992 to 2019, its contribution is comparable in magnitude to the ACR term, and the two channels often move in opposite directions. Countries whose domestic expenditure shares fall may gain through the ACR term yet lose through a contraction of the industrial base. Conversely, countries that attract production may gain through the industrial-base channel even when the standard terms imply losses. Accounting for endogenous firm location therefore changes not only the magnitude but also the sign of measured welfare effects for many countries, both in the historical decomposition and in the counterfactual analysis.

The broader implication is that trade affects welfare not only through what countries consume, but also through where production takes place. Domestic expenditure shares alone cannot distinguish gains from improved access to foreign goods from losses caused by an erosion of the domestic production base. This distinction is central to interpreting globalization, evaluating foreign productivity shocks, and designing trade and industrial policy. A natural extension is to introduce forward-looking location decisions in a dynamic setting, which I leave for future research.

References

Ahlfeldt, G. M., S. J. Redding, D. M. Sturm, and N. Wolf (2015). The economics of density: Evidence from the berlin wall. *Econometrica 83*(6), 2127–2189.

Allen, T. and C. Arkolakis (2014, 05). Trade and the topography of the spatial economy *. *The Quarterly Journal of Economics 129*(3), 1085–1140.

Allen, T., C. Arkolakis, and Y. Takahashi (2020). Universal gravity. *Journal of Political Economy 128*(2), 393–433.

Arkolakis, C., A. Costinot, and A. Rodríguez-Clare (2012, February). New trade models, same old gains? *American Economic Review 102*(1), 94–130.

Arkolakis, C., N. Ramondo, A. Rodríguez-Clare, and S. Yeaple (2018, August). Innovation and production in the global economy. *American Economic Review 108*(8), 2128–73.

Arkolakis, C., A. Rodríguez-Clare, and Z. Su (2025). A multivariate distribution with pareto tails and pareto maxima. Working paper.

Artuç, E., S. Chaudhuri, and J. McLaren (2010, June). Trade shocks and labor adjustment: A structural empirical approach. *American Economic Review 100*(3), 1008–45.

Autor, D. H., D. Dorn, and G. H. Hanson (2013, October). The china syndrome: Local labor market effects of import competition in the united states. *American Economic Review 103*(6), 2121–68.

Bagwell, K. and S. H. Lee (2020). Trade policy under monopolistic competition with firm selection. *Journal of International Economics 127*, 103379.

Baldwin, R., R. Forslid, P. Martin, G. Ottaviano, and F. Robert-Nicoud (2003). *Economic Geography and Public Policy*. Princeton, NJ: Princeton University Press.

Baldwin, R. E. and T. Okubo (2006, 06). Heterogeneous firms, agglomeration and economic geography: spatial selection and sorting. *Journal of Economic Geography 6*(3), 323–346.

Baldwin, R. E. and F. Robert-Nicoud (2000). Free trade agreements without delocation. *Canadian Journal of Economics 33*(3), 766–786.

Bartelme, D., A. Costinot, D. Donaldson, and A. Rodríguez-Clare (2025). The textbook case for industrial policy: Theory meets data. *Journal of Political Economy 133*(5), 1527–1573.

Behrens, K., G. Duranton, and F. Robert-Nicoud (2014). Productive cities: Sorting, selection, and agglomeration. *Journal of Political Economy 122*(3), 507–553.

Bernard, A. B., J. Eaton, J. B. Jensen, and S. Kortum (2003, September). Plants and productivity in international trade. *American Economic Review 93*(4), 1268–1290.

Boehm, C. E., A. A. Levchenko, and N. Pandalai-Nayar (2023, April). The long and short (run) of trade elasticities. *American Economic Review 113*(4), 861–905.

Borchert, I., M. Larch, S. Shikher, and Y. V. Yotov (2021, None). The international trade and production database for estimation (itpd-e). *International Economics 166*(C), 140–166.

Brandt, L., J. Van Biesebroeck, L. Wang, and Y. Zhang (2017, September). Wto accession and performance of chinese manufacturing firms. *American Economic Review 107*(9), 2784–2820.

Bryan, G. and M. Morten (2019). The aggregate productivity effects of internal migration: Evidence from indonesia. *Journal of Political Economy 127*(5), 2229–2268.

Caliendo, L., R. C. Feenstra, J. Romalis, and A. M. Taylor (2015, December). Tariff reductions, entry, and welfare: Theory and evidence for the last two decades. Working Paper 21768, National Bureau of Economic Research.

Caliendo, L., R. C. Feenstra, J. Romalis, and A. M. Taylor (2023). A second-best argument

for low optimal tariffs on intermediate inputs. *Journal of International Economics 145*, 103824.

Caliendo, L. and F. Parro (2019). The quantitative effects of trade policy on industry and labor locations. 2019 Meeting Papers 1575, Society for Economic Dynamics.

Caliendo, L., F. Parro, E. Rossi-Hansberg, and P.-D. Sarte (2018, 10). The impact of regional and sectoral productivity changes on the u.s. economy. *The Review of Economic Studies 85*(4), 2042–2096.

Chaney, T. (2008). Distorted gravity: The intensive and extensive margins of international trade. *The American Economic Review 98*(4), 1707–1721.

Combes, P.-P., G. Duranton, L. Gobillon, D. Puga, and S. Roux (2012). The productivity advantages of large cities: Distinguishing agglomeration from firm selection. *Econometrica 80*(6), 2543–2594.

Costinot, A. and A. Rodríguez-Clare (2014). Chapter 4 - trade theory with numbers: Quantifying the consequences of globalization. In G. Gopinath, E. Helpman, and K. Rogoff (Eds.), *Handbook of International Economics*, Volume 4 of *Handbook of International Economics*, pp. 197–261. Elsevier.

Costinot, A. and J. Vogel (2010). Matching and inequality in the world economy. *Journal of Political Economy 118*(4), 747–786.

Dekle, R., J. Eaton, and S. Kortum (2007, May). Unbalanced trade. *American Economic Review 97*(2), 351–355.

Demidova, S. and A. Rodríguez-Clare (2009). Trade policy under firm-level heterogeneity in a small economy. *Journal of International Economics 78*(1), 100–112.

Desmet, K., D. K. Nagy, and E. Rossi-Hansberg (2018). The geography of development. *Journal of Political Economy 126*(3), 903–983.

Desmet, K. and E. Rossi-Hansberg (2014, April). Spatial development. *American Economic Review 104*(4), 1211–43.

Dix-Carneiro, R. (2014). Trade liberalization and labor market dynamics. *Econometrica 82*(3), 825–885.

Eaton, J. and S. Kortum (2002). Technology, geography, and trade. *Econometrica 70*(5), 1741–1779.

Fajgelbaum, P., E. Morales, J. C. S. Serrato, and O. Zidar (2019). State taxes and spatial misallocation. *Review of Economic Studies 86*, 333–376.

Fajgelbaum, P. D. and C. Gaubert (2020, 01). Optimal spatial policies, geography, and sorting*. *The Quarterly Journal of Economics 135*(2), 959–1036.

Fally, T. (2015). Structural gravity and fixed effects. *Journal of International Economics 97*(1), 76–85.

Feenstra, R. C., R. Inklaar, and M. P. Timmer (2015, October). The next generation of the penn world table. *American Economic Review 105*(10), 3150–82.

Felbermayr, G., B. Jung, and M. Larch (2013). Optimal tariffs, retaliation, and the welfare loss from tariff wars in the melitz model. *Journal of International Economics 89*(1), 13–25.

Feller, W. (1971). *An Introduction to Probability Theory and Its Applications* (2 ed.), Volume II. New York: Wiley.

Flaaen, A., A. Hortaçsu, and F. Tintelnot (2020, July). The production relocation and price effects of us trade policy: The case of washing machines. *American Economic Review 110*(7), 2103–27.

Gabaix, X. (2009). Power laws in economics and finance. *Annual Review of Economics 1*(Volume 1, 2009), 255–294.

Gabaix, X. (2011). The granular origins of aggregate fluctuations. *Econometrica 79*(3), 733–772.

Galle, S., A. Rodríguez-Clare, and M. Yi (2023). Slicing the pie: Quantifying the aggregate and distributional effects of trade. *Econometrica 91*(1), 1–47.

Gaubert, C. (2018, November). Firm sorting and agglomeration. *American Economic Review 108*(11), 3117–53.

Grogger, J. and G. H. Hanson (2011). Income maximization and the selection and sorting of international migrants. *Journal of Development Economics 95*(1), 42–57. Symposium on Globalization and Brain Drain.

Haaland, J. I. and A. J. Venables (2016). Optimal trade policy with monopolistic competition and heterogeneous firms. *Journal of International Economics 102*, 85–95.

Hanson, G. H. (1996). Economic integration, intraindustry trade, and frontier regions. *European Economic Review 40*(3), 941–949. Papers and Proceedings of the Tenth Annual Congress of the European Economic Association.

Hanson, G. H. (2001). Scale economies and the geographic concentration of industry. *Journal of Economic Geography 1*(3), 255–276.

Hanson, G. H. (2012, May). The rise of middle kingdoms: Emerging economies in global trade. *Journal of Economic Perspectives 26*(2), 41–64.

Helpman, E. and P. R. Krugman (1985). *Market Structure and Foreign Trade: Increasing Returns, Imperfect Competition, and the International Economy*. Cambridge, MA: MIT Press.

Holmes, T. J. and H. Sieg (2015). Chapter 2 - structural estimation in urban economics. In G. Duranton, J. V. Henderson, and W. C. Strange (Eds.), *Handbook of Regional and Urban Economics*, Volume 5 of *Handbook of Regional and Urban Economics*, pp. 69–114. Elsevier.

Juhász, R., N. Lane, and D. Rodrik (2024). The new economics of industrial policy. *Annual Review of Economics 16*(Volume 16, 2024), 213–242.

Klette, T. and S. Kortum (2004). Innovating firms and aggregate innovation. *Journal of Political Economy 112*(5), 986–1018.

Krugman, P. and A. J. Venables (1995). Globalization and the inequality of nations. *Quarterly Journal of Economics 110*(4), 857–880.

Krugman, P. R. (1980). Scale economies, product differentiation, and the pattern of trade. *American Economic Review 70*(5), 950–959.

Kucheryavyy, K., G. Lyn, and A. Rodríguez-Clare (2023, April). Grounded by gravity: A well-behaved trade model with industry-level economies of scale. *American Economic Journal: Macroeconomics 15*(2), 372–412.

Larch, M., S. Shikher, and Y. Yotov (2025, May). The international trade and production database for estimation - release 3 (itpd-e-r03). Working Papers 202524, Center for Global Policy Analysis, LeBow College of Business, Drexel University.

Lashkaripour, A. (2021, 10). Can trade taxes be a major source of government revenue? *Journal of the European Economic Association 19*(5), 2399–2428.

Lashkaripour, A. and V. Lugovskyy (2023, October). Profits, scale economies, and the gains from trade and industrial policy. *American Economic Review 113*(10), 2759–2808.

Lind, N. and N. Ramondo (2023a, December). Global innovation and knowledge diffusion. *American Economic Review: Insights 5*(4), 494–510.

Lind, N. and N. Ramondo (2023b, February). Trade with correlation. *American Economic Review 113*(2), 317–53.

Lucas, R. E. and E. Rossi–Hansberg (2002). On the internal structure of cities. *Economet-*

rica 70(4), 1445–1476.

Malgouyres, C., T. Mayer, and C. Mazet-Sonilhac (2023, August). Who benefits from state corporate tax cuts? a local labor markets approach with heterogeneous firms: Comment. *American Economic Review 113*(8), 2270–86.

Mayer, T. and S. Zignago (2011, December). Notes on CEPII's distances measures: The GeoDist database. Working Papers 2011-25, CEPII research center.

McFadden, D. (1974). Conditional logit analysis of qualitative choice behavior. In P. Zarembka (Ed.), *Fontiers in Econometrics*, pp. 105–142. New York: Academic press.

McFadden, D. (1978). Modeling the choice of residential location. In A. Karlqvist, L. Lundqvist, F. Snickars, and J. Weibull (Eds.), *Spatial Interaction Theory and Planning Models*, Amsterdam, pp. 75–96. North-Holland.

Melitz, M. J. (2003). The impact of trade on intra-industry reallocations and aggregate industry productivity. *Econometrica 71*(6), 1695–1725.

Melitz, M. J. and S. J. Redding (2015, March). New trade models, new welfare implications. *American Economic Review 105*(3), 1105–46.

Monte, F., S. J. Redding, and E. Rossi-Hansberg (2018). Commuting, migration, and local employment elasticities. *American Economic Review 108*(12), 3855–3890.

Nelsen, R. B. (2006). *An Introduction to Copulas* (Second ed.). New York: Springer.

Okubo, T., P. M. Picard, and J.-F. Thisse (2010). The spatial selection of heterogeneous firms. *Journal of International Economics 82*(2), 230–237.

Ossa, R. (2011). A "new trade" theory of GATT/WTO negotiations. *Journal of Political Economy 119*(1), 122–152.

Sampson, T. (2014, August). Selection into trade and wage inequality. *American Economic Journal: Microeconomics 6*(3), 157–202.

Spearot, A. (2016, May). Unpacking the long-run effects of tariff shocks: New structural implications from firm heterogeneity models. *American Economic Journal: Microeconomics 8*(2), 128–67.

Suárez Serrato, J. C. and O. Zidar (2016, September). Who benefits from state corporate tax cuts? a local labor markets approach with heterogeneous firms. *American Economic Review 106*(9), 2582–2624.

Suárez Serrato, J. C. and O. Zidar (2023, December). Who benefits from state corporate tax cuts? a local labor market approach with heterogeneous firms: Reply. *American Economic Review 113*(12), 3401–10.

Teti, F. A. (2024). Missing tariffs. Technical Report 11590, CESifo Working Papers. Feodora Teti's Global Tariff Database (v_beta1-2024-12).

Train, K. E. (2009). *Discrete Choice Methods with Simulation*. Cambridge University Press.

Venables, A. J. (1985). Trade and trade policy with imperfect competition: The case of identical products and free entry. *Journal of International Economics 19*(1–2), 1–19.

Venables, A. J. (1987). Trade and trade policy with differentiated products: A chamberlinian–ricardian model. *Economic Journal 97*(387), 700–717.

Appendix A: Multivariate Fréchet with Logistic Dependence

In this appendix I construct the multivariate Fréchet distribution with logistic (Gumbel–Hougaard) dependence used in the text and prove Proposition 1. Throughout, z_k is the productivity shock at location k, with shape $\theta > 1$, scale $H_k > 0$, and dependence $\rho \geq 1$; the deterministic return components are the wage w_k and market access $\mathcal{M}_k$, and η is the market-access elasticity.

A1: Construction of Multivariate Fréchet with Logistic Dependence

Begin with the marginal law of each shock z_k, Fréchet with shape $\theta > 0$ and scale H_k:

$$F_k(z) = \exp\left(-H_k z^{-\theta}\right).$$

Define $u_k \equiv F_k(z_k)$, uniform on $[0,1]$ by the probability integral transform, and impose the logistic copula as discussed in Nelsen (2006):

$$C(u_1, \ldots, u_K) = \exp\left(-\left[\sum_{j=1}^{K}(-\log u_j)^{\rho}\right]^{1/\rho}\right), \qquad \rho \geq 1. \tag{A1}$$

At $\rho = 1$ this is the independence copula $\prod_j u_j$ and as $\rho \to \infty$ it approaches comonotonicity. By Sklar's theorem $F(\mathbf{z}) = C(F_1(z_1), \ldots, F_K(z_K))$. Since $-\log u_k = -\log F_k(z_k) = H_k z_k^{-\theta}$, one can establish the following:

$$\sum_{j=1}^{K}(-\log u_j)^{\rho} = \sum_{j=1}^{K}\left(H_j z_j^{-\theta}\right)^{\rho},$$

leading to the joint law used in Proposition 1:

$$F(z_1, \ldots, z_K) = \exp\left(-\left[\sum_{j=1}^{K}\left(H_j z_j^{-\theta}\right)^{\rho}\right]^{1/\rho}\right).$$

A2: Proof of Proposition 1

The idea behind the proof of Proposition 1 maps the firm location problem into a logistic generalized extreme-value model and then applies standard GEV results for choice probabilities and the joint distribution of the maximizing alternative and the maximized payoff.

The productivity vector $\boldsymbol{z} = (z_1, \ldots, z_K)$ follows (1), with $H_j > 0$, $\theta > 1$, and $\rho \geq 1$. Agents choose k^* to maximize

$$V_k = W_k z_k, \qquad W_k \equiv w_k^{-1} \mathcal{M}_k^{1/\eta} > 0.$$

A2.1: Log-payoff transformation

Let me reformulate the payoff problem in logs as follows:

$$U_k = \ln V_k = \ln W_k + \ln z_k.$$

This is a random-utility model with deterministic component $v_k = \ln W_k$ and stochastic component $\widetilde{\epsilon}_k = \ln z_k$. Recall that the marginal distribution of z_k if Fréchet(H_k, θ) such that the following holds:

$$F_{\widetilde{\epsilon}_k}(y) = \Pr\left(z_k \leq e^y\right) = \exp\left\{-\exp\left(\ln H_k - \theta y\right)\right\}.$$

Hence, $\widetilde{\epsilon}_k$ follows a type-I extreme-value distribution with location parameter $\psi_k = \ln H_k/\theta$ and common scale $\varsigma = 1/\theta$.

A2.2: Reduction to the logistic GEV model

To reduce the problem to the logistic GEV, let me define the following variables:

$$u_k = \frac{v_k + \psi_k}{\varsigma} = \theta \ln W_k + \ln H_k, \qquad \nu_k = \frac{\widetilde{\epsilon}_k - \psi_k}{\varsigma}.$$

Then the log payoff function can be rewritten as:

$$U_k = \varsigma(u_k + \nu_k).$$

Since $\varsigma > 0$, maximizing U_k is equivalent to solving the following problem:

$$k^* = \arg\max_k \{u_k + \nu_k\}.$$

For any threshold vector $\boldsymbol{x} \in \mathbb{R}^K$:

$$\{\nu_j \leq x_j \text{ for all } j\} = \{z_j \leq \exp(\psi_j + \varsigma x_j) \text{ for all } j\}.$$

Evaluating equation (1) at these thresholds gives

$$F_{\boldsymbol{\nu}}(\boldsymbol{x}) = \exp\left\{-\left[\sum_j \left(H_j \exp\left[-\theta(\psi_j + \varsigma x_j)\right]\right)^{\rho}\right]^{1/\rho}\right\}.$$

Combining the identities that $\theta\psi_j = \ln H_j$ and $\theta\varsigma = 1$ allows to express $H_j \exp\left[-\theta(\psi_j + \varsigma x_j)\right]$ as e^{-x_j} such that the following holds:

$$F_{\boldsymbol{\nu}}(\boldsymbol{x}) = \exp\left\{-\left[\sum_j e^{-\rho x_j}\right]^{1/\rho}\right\}.$$

This is the symmetric logistic GEV distribution:

$$F_{\boldsymbol{\nu}}(\boldsymbol{x}) = \exp\left\{-G(e^{-x_1}, \ldots, e^{-x_K})\right\}, \quad \text{where} \quad G(y_1, \ldots, y_K) = \left(\sum_j y_j^{\rho}\right)^{1/\rho},$$

with dependence parameter $1/\rho$ (McFadden 1978; Train 2009). Then, the implied choice probabilities are:

$$\mu_k \equiv \Pr(k^* = k) = \frac{e^{\rho u_k}}{\sum_j e^{\rho u_j}}.$$

Recall that $e^{u_k} = H_k W_k^{\theta}$ and the expression for the location shares can be rewritten as:

$$\mu_k = \frac{H_k^{\rho} w_k^{-\theta\rho} \mathcal{M}_k^{\theta\rho/\eta}}{\sum_j H_j^{\rho} w_j^{-\theta\rho} \mathcal{M}_j^{\theta\rho/\eta}},$$

which proves part (i) of Proposition 1.

A2.3: Conditional distribution of the selected shock

By the standard GEV result established above, $V^* \equiv \max_j V_j$ is independent of the identity of the maximizing location, k^*. Moreover, conditional on $k^* = k$, the winning payoff satisfies $V_k = V^*$ such that the following holds:

$$F_{V_k|k^*=k}(v) = \Pr(V^* \leq v \mid k^* = k) = \Pr\left(z_j \leq \frac{v}{W_j} \text{ for all } j\right) = \exp\left(-\Xi v^{-\theta}\right),$$

where the aggregate term is defined as:

$$\Xi \equiv \left[\sum_j \left(H_j W_j^{\theta} \right)^{\rho} \right]^{1/\rho}.$$

For $x > 0$, use the identity $V_k = W_k z_k$:

$$F_{z_k|k^*=k}(x) = \Pr\left(z_k \leq x \mid k^* = k\right) = \Pr\left(V_k \leq W_k x \mid k^* = k\right) = \exp\left(-\Xi W_k^{-\theta} x^{-\theta}\right).$$

Using the expressions for μ_k and Ξ this can be rewritten as:

$$F_{z_k|k^*=k}(x) = \exp\left(-H_k \mu_k^{-1/\rho} x^{-\theta}\right),$$

which proves part (ii) of Proposition 1.

A2.4: Conditional expectation of the selected shock

Once the distribution of z_k conditional on $(k^* = k))$ is known to be Fréchet with shape θ and scale $s_k = \left(H_k \mu_k^{-1/\rho}\right)^{1/\theta} = H_k^{1/\theta} \mu_k^{-1/(\theta\rho)}$, it immediately follows that for $\theta > 1$:

$$\mathbb{E}[z_k \mid k^* = k] = s_k \, \Gamma\left(1 - \tfrac{1}{\theta}\right) = H_k^{1/\theta} \, \mu_k^{-1/(\theta\rho)} \, \Gamma\left(1 - \tfrac{1}{\theta}\right),$$

which is part (iii) of Proposition 1. □

Appendix B: Consistency of the Finite-Firm Construction

In this Appendix I prove the three limits in the main text:

$$(i) \frac{N_i}{N} \xrightarrow{p} \mu_i, \qquad (ii) \frac{\overline{T}_i}{N \ln N} \xrightarrow{p} T_i, \qquad (iii) \max_{f \in \mathcal{F}_i} s_i^f \xrightarrow{p} 0, \tag{B1}$$

Start with N firms drawing capability vectors $\{\boldsymbol{z}^1, \ldots, \boldsymbol{z}^N\}$ independently from the multivariate logistic–Fréchet distribution and assign each firm according to $i^f \equiv i^*(\boldsymbol{z}^f)$. Let $\mathcal{F}_i \equiv \{f : i^f = i\}$ denote the set of N_i firms located in country i.

(i) Since the productivity draws are i.i.d. and the location rule is a function of a firm's own vector, the indicators $\mathbf{1}\{i^f = i\}$ are i.i.d. Bernoulli with mean $\Pr(i^*(\boldsymbol{z}) = i) = \mu_i$,

the continuum location share of Proposition 1. The law of large numbers gives

$$\frac{N_i}{N} = \frac{1}{N}\sum_{f=1}^{N} \mathbf{1}\{i^f = i\} \xrightarrow{p} \mu_i,$$

the first limit in equation (B1).

(ii) Using the result from *(i)*, fix i with $\mu_i > 0$. Because each pair $(\boldsymbol{z}^f, i^f)$ has the same law as the generic continuum pair $(\boldsymbol{z}, i^*(\boldsymbol{z}))$, Proposition 1 gives, for every f and every finite N:

$$\Pr\left(z_i^f \le z \mid i^f = i\right) = \exp\left(-\widetilde{T}_i z^{-\theta}\right), \qquad \widetilde{T}_i = H_i \mu_i^{-1/\rho}. \tag{B2}$$

Define $x_i^f \equiv \mathbf{1}\{i^f = i\}[z_i^f]^\theta$, so that $\overline{T}_i = \sum_{f=1}^{N} x_i^f = \sum_{f\in\mathcal{F}_i}[z_i^f]^\theta$ and $\{x_i^f\}_{f=1}^{N}$ are i.i.d. such that equation (B2) gives the following:

$$\Pr(x_i^f > y) = \mu_i\left[1 - \exp\left(-\frac{\widetilde{T}_i}{y}\right)\right] = \frac{T_i}{y} + O\left(y^{-2}\right), \qquad y \to \infty, \tag{B3}$$

where $T_i \equiv \mu_i \widetilde{T}_i = H_i \mu_i^{1-1/\rho}$. Thus x_i^f has a Pareto upper tail with index one and $\mathbb{E}[x_i^f] = \infty$. The sum therefore obeys the weak law of large numbers for regularly varying tails with index one Feller (1971), the case that Gabaix (2011) analyzes under Zipf's law.

For a truncation point b, define $m_i(b) \equiv \mathbb{E}[x_i^f \mathbf{1}\{x_i^f \le b\}]$ as follows:

$$m_i(b) = \int_0^b \Pr(x_i^f > y)\, dy - b\Pr(x_i^f > b) = T_i \ln b + O(1), \qquad b \to \infty,$$

where the second equality splits the integral at a fixed $y_0 > 0$, applies the expansion on $[y_0, b]$, and uses $b\Pr(x_i^f > b) = T_i + O(b^{-1})$. Set $b_N \equiv N\ln N$ and $Y_{i,N}^f \equiv x_i^f \mathbf{1}\{x_i^f \le b_N\}$. Since $\ln b_N / \ln N \to 1$, the following holds:

$$\frac{N\mathbb{E}[Y_{i,N}^f]}{N\ln N} = \frac{m_i(N\ln N)}{\ln N} = T_i \frac{\ln(N\ln N)}{\ln N} + O\left(\frac{1}{\ln N}\right) \longrightarrow T_i, \tag{B4}$$

which identifies $N\ln N$ as the proper normalization scale.

Concentration follows from the global bound implied by equation (B3) and $1-e^{-a} \le a$,

$$\Pr(x_i^f > y) \le \frac{T_i}{y} \qquad \text{for every } y > 0, \tag{B5}$$

which together with $\Pr(x_i^f > y) \le 1$ gives $y \Pr(x_i^f > y) \le \min\{y, T_i\}$. Hence $\mathbb{E}[(Y_{i,N}^f)^2] \le 2\int_0^{b_N} y \Pr(x_i^f > y)\, dy \le 2\int_0^{b_N} \min\{y, T_i\}\, dy = O(b_N)$, and, by independence,

$$\operatorname{Var}\left[\frac{1}{N \ln N} \sum_{f=1}^{N} Y_{i,N}^f\right] \le \frac{N\mathbb{E}[(Y_{i,N}^f)^2]}{N^2(\ln N)^2} = O\left(\frac{1}{\ln N}\right) \longrightarrow 0. \tag{B6}$$

Equations (B4) and (B6) with Chebyshev's inequality give $(N \ln N)^{-1} \sum_f Y_{i,N}^f \xrightarrow{p} T_i$. The truncation is immaterial: by the union bound and equation (B5),

$$\Pr\left(\max_{1 \le f \le N} x_i^f > b_N\right) \le N \Pr(x_i^f > b_N) \le \frac{T_i}{\ln N} \longrightarrow 0,$$

so the truncated and untruncated sums coincide with probability approaching one such that the following holds:

$$\frac{\overline{T}_i}{N \ln N} \xrightarrow{p} T_i = H_i \mu_i^{1-1/\rho},$$

the second limit in equation (B1).

(iii) Maintain $\mu_i > 0$. For any $\varepsilon > 0$, the union bound and equation (B5) give $\Pr\left(\max_f x_i^f > \varepsilon N \ln N\right) \le T_i/(\varepsilon \ln N) \to 0$, so $\max_f x_i^f/(N \ln N) \xrightarrow{p} 0$. On the event $\overline{T}_i > 0$, dividing by $\overline{T}_i/(N \ln N) \xrightarrow{p} T_i > 0$ and applying the continuous-mapping theorem gives $\max_f x_i^f/\overline{T}_i \xrightarrow{p} 0$. Because productivities are strictly positive, $\overline{T}_i = 0$ only if $\mathcal{F}_i = \varnothing$, an event of probability $(1-\mu_i)^N \to 0$. Hence, with $s_i^f = x_i^f/\overline{T}_i$ the within-country technology share and the maximum defined as zero on $\{\mathcal{F}_i = \varnothing\}$,

$$\max_{f \in \mathcal{F}_i} s_i^f = \frac{\max_{1 \le f \le N} x_i^f}{\overline{T}_i} \xrightarrow{p} 0,$$

the third limit in equation (B1). Because the set of countries is finite, the convergence results in parts *(i)–(iii)* hold jointly across countries.

Appendix B: Proof of Proposition 2

This appendix proves the results in Proposition 2. The real wage can be expressed in terms of the domestic expenditure share. With $\tau_{ii} = \kappa_{ii} = 1$, equations (7) and (9) imply:

$$\pi_{ii} = \mathcal{A}\, H_i^{\eta/\theta} \mu_i^{1-\eta/(\theta\rho)} \left(\frac{P_i}{w_i}\right)^{\eta}.$$

Solving for w_i/P_i leads to the following equations:

$$\frac{w_i}{P_i} = \mathcal{A}^{1/\eta} H_i^{1/\theta} \mu_i^{1/\eta - 1/(\theta\rho)} \pi_{ii}^{-1/\eta}. \tag{C1}$$

Substituting (C1) into (15) gives the welfare representation in logs:

$$\ln W_i = -\frac{1}{\eta}\ln \pi_{ii} + \left(\frac{1}{\eta} - \frac{1}{\theta\rho}\right)\ln \mu_i - \ln \Lambda_i + \frac{1}{\theta}\ln H_i + \ln(1+d_i) + C_m, \tag{C2}$$

where $C_m = \frac{1}{\eta}\ln \mathcal{A} - \ln(1-\mathcal{S})$ is constant within a market structure. Since d_i is held fixed across equilibria, differencing (C2) leaves only π_{ii}, μ_i, and Λ_i, giving the expression in equation (15).

Appendix D: Multivariate Dagum with Clayton Dependence

In this appendix, I construct the multivariate Dagum distribution with Clayton dependence and prove Proposition 3. Throughout, z_k is the productivity shock at location k, with tail parameter $\xi > 0$, scale $H_k > 0$, and dependence $\delta > 0$; the deterministic return components are the wage w_k and market access $\mathcal{M}_k$, with market-access elasticity η.

D1: Construction of Multivariate Dagum with Clayton Dependence

Let me start with the marginal distribution of each shock z_k, specified as Dagum with parameters $\xi > 0$, $1/\delta > 0$, and $H_k > 0$:

$$F_k(z) = \left[1 + \left(\frac{z}{H_k}\right)^{-\xi}\right]^{-1/\delta}. \tag{D1}$$

Define $u_k \equiv F_k(z_k)$, uniform on $[0,1]$, and impose the Clayton copula (Nelsen 2006):

$$C(u_1,\ldots,u_K) = \left(\sum_{j=1}^{K} u_j^{-\delta} - (K-1)\right)^{-1/\delta}, \qquad \delta > 0.$$

As $\delta \to 0$ this is the independence copula; as $\delta \to \infty$ it approaches comonotonicity. By Sklar's theorem $F(\mathbf{z}) = C(F_1(z_1), \ldots, F_K(z_K))$. Since $u_k^{-\delta} = 1 + (z_k/H_k)^{-\xi}$,

$$\sum_{j=1}^{K} u_j^{-\delta} - (K-1) = K + \sum_{j=1}^{K} \left(\frac{z_j}{H_j}\right)^{-\xi} - (K-1) = 1 + \sum_{j=1}^{K} \left(\frac{z_j}{H_j}\right)^{-\xi},$$

which leads to the joint distribution in Proposition 3:

$$F(z_1, \ldots, z_K) = \left[1 + \sum_{j=1}^{K} \left(\frac{z_j}{H_j}\right)^{-\xi}\right]^{-1/\delta}. \tag{D2}$$

D2: Proof of Proposition 3

The productivity vector $\boldsymbol{z} = (z_1, \ldots, z_K)$ follows (D2), with $\xi > 1$, $\delta > 0$, and $H_j > 0$. Firms choose k^* to maximize

$$V_k = W_k z_k, \qquad W_k \equiv w_k^{-1} \mathcal{M}_k^{1/\eta} > 0.$$

The proof proceeds in three steps. First, I separate the common dependence component from independent Fréchet shocks. Second, I derive the joint distribution of the selected location and its shock, from which the location shares and the conditional shock distribution follow simultaneously. Third, I recover the conditional productivity distribution.

D2.1: Shock Transformation

To prove the results, I first transform the shocks. Let $\zeta_1, \ldots, \zeta_K$ be i.i.d. standard Fréchet random variables with shape parameter ξ:

$$F_\zeta(z) = \exp\left(-z^{-\xi}\right), \qquad z > 0.$$

Further, let $Q \sim \Gamma(1/\delta, 1)$ be independent of $\boldsymbol{\zeta}$, and define:

$$z_k = H_k Q^{1/\xi} \zeta_k.$$

Conditional on $Q = q$, independence of the ζ_k implies

$$\Pr(z_j \leq e_j \text{ for all } j \mid Q = q) = \exp\left[-q \sum_j \left(\frac{e_j}{H_j}\right)^{-\xi}\right].$$

Integrating over Q leads to the distribution in equation (D2):

$$\Pr(z_j \leq e_j \text{ for all } j) = \left[1 + \sum_j \left(\frac{e_j}{H_j}\right)^{-\xi}\right]^{-1/\delta},$$

and letting all but one argument tend to infinity yields the marginal distribution in (D1).

Because $Q^{1/\xi} > 0$ is common across locations, it drops out of the location choice problem as follows:

$$k^* = \arg\max_k \left\{W_k H_k Q^{1/\xi} \zeta_k\right\} = \arg\max_k \left\{\phi_k \zeta_k\right\},$$

where

$$\phi_k \equiv W_k H_k = H_k w_k^{-1} \mathcal{M}_k^{1/\eta}.$$

D2.2: Location shares and conditional selection

The joint probability that location k is selected and its standard shock does not exceed z is as follows:

$$\Pr(\zeta_k \leq z, k^* = k) = \int_0^z \Pr\left(\phi_j \zeta_j \leq \phi_k t \text{ for all } j \neq k\right) dF_\zeta(t) = \int_0^z \prod_{j \neq k} F_\zeta\left(\frac{\phi_k}{\phi_j} t\right) dF_\zeta(t).$$

Using that $dF_\zeta(t) = \xi t^{-\xi-1} \exp\left(-t^{-\xi}\right) dt$,, I can rewrite it as follows:

$$\Pr(\zeta_k \leq z, k^* = k) = \int_0^z \xi t^{-\xi-1} \exp\left[-t^{-\xi} \sum_j \left(\frac{\phi_j}{\phi_k}\right)^{\xi}\right] dt.$$

Let me define as:

$$A_k \equiv \sum_j \left(\frac{\phi_j}{\phi_k}\right)^{\xi} = \frac{\sum_j \phi_j^{\xi}}{\phi_k^{\xi}}.$$

Then, the following holds:

$$\Pr(\zeta_k \leq z, k^* = k) = A_k^{-1} \exp\left(-A_k z^{-\xi}\right).$$

because:

$$\frac{d}{dt} \exp\left(-A_k t^{-\xi}\right) = A_k \xi t^{-\xi-1} \exp\left(-A_k t^{-\xi}\right),$$

Next, let $z \to \infty$ to obtain the location share as follows:

$$\mu_k \equiv \Pr(k^* = k) = A_k^{-1} = \frac{\phi_k^{\xi}}{\sum_j \phi_j^{\xi}}.$$

Or equivalently, substituting for ϕ_k gives:

$$\mu_k = \frac{H_k^{\xi} w_k^{-\xi} \mathcal{M}_k^{\xi/\eta}}{\sum_j H_j^{\xi} w_j^{-\xi} \mathcal{M}_j^{\xi/\eta}},$$

which proves part (i) of Proposition 3.

Next, using $A_k = \mu_k^{-1}$, the joint probability can be rewritten as

$$\Pr(\zeta_k \leq z, k^* = k) = A_k^{-1} \exp\left(-A_k z^{-\xi}\right) = \mu_k \exp\left(-\mu_k^{-1} z^{-\xi}\right).$$

Dividing by $\Pr(k^* = k) = \mu_k$ gives the conditional distribution

$$F_{\zeta_k | k^* = k}(z) = \Pr(\zeta_k \leq z \mid k^* = k) = \exp\left(-\mu_k^{-1} z^{-\xi}\right),$$

which remains Fréchet with shape parameter ξ and exponent-measure coefficient μ_k^{-1}. Defining the conditional scale parameter $\tilde{D}_k \equiv H_k \mu_k^{-1/\xi}$. and following the same derivations as before, the conditional distribution can then be written as:

$$F_{z_k | k^* = k}(x) = \left[1 + \left(\frac{x}{\tilde{D}_k}\right)^{-\xi}\right]^{-1/\delta},$$

which establishes part (ii) of Proposition 3.

D2.4: Conditional expectation of the selected shock

For a Dagum random variable with distribution:

$$F(x) = \left[1 + \left(\frac{x}{D_k}\right)^{-\xi}\right]^{-1/\delta},$$

the m-th moment exists for $m < \xi$ and equals:

$$\mathbb{E}[x^m] = D_k^m \frac{\Gamma(1/\delta + m/\xi)\Gamma(1 - m/\xi)}{\Gamma(1/\delta)}.$$

Because $\xi > 1$, the conditional expectation exists and is

$$\mathbb{E}[z_k \mid k^* = k] = H_k \mu_k^{-1/\xi} \frac{\Gamma(1 - 1/\xi)\Gamma(1/\delta + 1/\xi)}{\Gamma(1/\delta)},$$

which proves part (iii) of Proposition 3. □