GRANULAR COMPARATIVE ADVANTAGE

Cecile Gaubert
Oleg Itskhoki

Granular Comparative Advantage
Cecile Gaubert and Oleg Itskhoki
NBER Working Paper No. 24807
July 2018
JEL No. D20,D43,F10,F40

ABSTRACT

Large firms play a pivotal role in international trade, shaping the export patterns of countries. We propose and quantify a granular multi-sector model of trade, which combines fundamental comparative advantage across sectors with granular comparative advantage embodied in outstanding individual firms. We develop an SMM-based estimation procedure, which takes full account of the general equilibrium of the model, to jointly estimate these fundamental and granular forces using French micro-data with information on firm domestic and export sales across manufacturing industries. We find that granularity accounts for about 20% of the variation in realized export intensity across sectors, and is more pronounced in the most export-intensive sectors. In turn, idiosyncratic firm dynamics accounts for a large share of the evolution of a country's comparative advantage over time. Governments face strong incentives to target trade policy at large individual foreign exporters, and to use lenient antitrust regulation at home to substitute for beggar-thy-neighbor trade policy.

Cecile Gaubert
Department of Economics
University of California, Berkeley
687 Evans Hall
Berkeley, CA 94720
and NBER
cecile.gaubert@berkeley.edu

Oleg Itskhoki
Department of Economics
Princeton University
Fisher Hall 306
Princeton, NJ 08544-1021
and NBER
itskhoki@princeton.edu

1 Introduction

Firms play a pivotal role in international trade. A significant share of exports is done by a small number of large firms, which enjoy substantial market power across destination countries.[1] The fates of these large firms shape, in part, the countries' trade patterns. For instance, Nokia in Finland or the Intel plant in Costa Rica have profoundly altered the specialization and export intensity of these countries.[2] The importance of large firms is also reflected in trade and industrial policies that are often so narrow that they appear tailor-made to target individual firms rather than industries. In particular, antitrust regulation, antidumping policies, and international sanctions all target large individual foreign firms.[3]

In this paper, we study the role of individual firms in determining the comparative advantage of a country. We decompose comparative advantage into two components: First, an inherent country-sector characteristic, which can stem from technological or endowment differences, and is, importantly, immune to the fate of individual firms. We call it *fundamental* comparative advantage. Second, specific expertise embodied in individual firms and entrepreneurs, which determines what we call *granular* comparative advantage. We set out to measure the contribution of granular comparative advantage to international trade flows. We then study the extent to which the granular origin of sectors shapes the evolution of their export performance over time. Finally, we analyze trade and antitrust policies targeted at granular sectors.

We thus revisit the fundamental questions in international trade: what goods do countries trade and what is the source of a country's comparative advantage? We find that granularity accounts for about 20% of the variation in realized export intensity across sectors, which suggests that a substantial share of international trade flows is intrinsically linked to the success of individual firms. Second, we show that idiosyncratic firm dynamics can account for the bulk of the evolution of a country's comparative advantage over time, contributing in an important way to the inter-sectoral reallocation of employment and production. Finally, we find that governments have an incentive to target trade policy at large individual foreign exporters,

---

[1]In their "Export superstars" paper, Freund and Pierola (2015) find that a single largest exporting firm accounts for 17% of total manufacturing exports, on average across 32 developing and middle-income countries in their dataset. In the French manufacturing dataset used in this paper, the largest firm accounts for 7% of all manufacturing exports, and within 4-digit industries the largest firm accounts on average for 28% of the industry exports.

[2]In Costa Rica, Intel decided to close its microchip plant and move it to Asia in 2014. The electronics sector represented a steady 27% of Costa-Rican exports until 2013, yet starting 2015 it fell to just 8%. In Finland, Nokia at its peak in the mid-2000s enjoyed a 25% share of total Finnish exports, a 3.7% share of Finnish GDP, and a 39% share of the global mobile phone market, before collapsing following the smartphone revolution launched by Apple, and being eventually bought-out by Microsoft in 2013.

[3]Recent examples of international antitrust regulations are the 2007 case of the European Commission (EC) against Microsoft Corporation and the 2017 fine imposed by the EC on Google. A very recent case of a granular trade war is the 292% tariff imposed by the US on a particular jet produced by the Canadian Bombardier. "Granular" tactics are particularly widespread in antidumping retaliation (see Blonigen and Prusa 2008) and international sanctions (as in the recent case of the US against the Chinese ZTE).

rather than entire industries. In addition, in granular sectors, they can use lenient antitrust policy as a substitute for beggar-thy-neighbor trade policy.

The international trade literature has shifted towards modeling individual firm exporting decisions, with a focus on firm heterogeneity and selection of the fittest firms into exporting. At the same time, the workhorse model maintains the assumption that sectors are comprised of a continuum of firms, each firm being infinitesimal. Under this continuum assumption, the productivity of any individual firm is inconsequential for sectoral trade patterns. Indeed, such *continuous models* are equivalent in the aggregate to a neoclassical Ricardian model that focuses on sector-level technologies and fully abstracts from modeling individual firms.

We propose and quantify an alternative multi-sector *granular model* of trade. We lift the assumption of a continuum of firms, acknowledging that a finite number of firms operate in each sector, with the largest firms often claiming a large share of the market. In our model, firms are small at the level of the whole economy, yet they are big within their individual industries. Under these circumstances, realized sectoral productivities, and hence the comparative advantage of a country, are shaped in part by the idiosyncratic productivity draws of individual firms, which do not average out at the sectoral level.

Our granular model builds on a number of familiar ingredients. Across sectors, it features classic Ricardian forces as in Dornbusch, Fischer, and Samuelson (1977; henceforth DFS). Within sectors, we maintain the firm heterogeneity assumption of the Melitz (2003) model, but relax the assumption of a continuum of firms, following Eaton, Kortum, and Sotelo (2012; henceforth EKS). Specifically, we assume that a finite number of firms in each sector draw productivities from a Pareto distribution, resulting in a distribution of sectoral firm sales which approximates Zipf's law, as observed in the data.[4] Our model is a multi-sector extension of EKS, which allows us to simultaneously nest fundamental and granular comparative advantage in a unified framework. We set up the model and study its analytical properties in Section 2.

In Section 3, we develop a simulated method of moments (SMM) estimation procedure which takes full account of the general equilibrium of our granular model. We estimate the model using firm-level data on domestic and export sales of French firms across 119 4-digit manufacturing industries. To disentangle the relative roles of fundamental and granular forces, we leverage information on the cross-sectional joint distribution of sectoral trade flows and domestic sales concentration within sectors. Intuitively, sectors in which export intensity is high due to granular comparative advantage are expected to feature outlying firms that are particularly large relative to other firms in the industry. We show that the estimated granular model, which features limited heterogeneity in parameters across sectors, is successful at

---

[4]Zipf's law is a fat-tailed Pareto distribution with a shape parameter of 1. The law of large numbers provides a particularly poor approximation when a large number of draws are taken from a fat-tailed distribution.

reproducing the rich empirical cross-sectoral distributions — such as those of the number of firms, of market shares of the largest firm, and of sectoral export intensity.

We use the estimated model in Section 4 to quantify the importance of granularity in shaping sectoral trade outcomes. In the model, as in the data, despite the fact that an average manufacturing sector features around 350 firms, the largest home firm accounts on average for 20% of the domestic-market sales by all home firms. This is a reflection of a fat-tailed distribution of firm sales, and hence of strong granular forces. We find that not only is a substantial share of trade flows of granular origin, but the contribution of granularity is particularly pronounced in the most export-intensive sectors, i.e. in the *export champions* of the country. Among the top 10% export-intensive sectors, nearly one third of exports is of granular origin. In addition, we show that conventional inference of fundamental sectoral productivities based on sectoral export shares leads to biased estimates in a granular world. Measurable proxies of granularity, such as the relative sales share of the top firm in the domestic market, predict trade flows and can be used to reduce this bias. We also develop a Bayesian filtering technique to predict the likelihood that exports in a given sector in the data are of granular origin.

Having established the role of granularity to shaping the sectoral trade patterns, we explore in Section 5 its dynamic implications. First, we study the role of granularity in explaining the evolution of comparative advantage over time. To this end, we extend our granular model to allow for industry dynamics driven by idiosyncratic firm productivity shocks, disciplined with the persistence of firm market shares in the data. We find that the model is simultaneously consistent with the hyper-specialization of countries in a few industries at any given point in time and with a relatively fast mean reversion in comparative advantage, emphasized in a recent paper by Hanson, Lind, and Muendler (2016; henceforth HLM). In particular, our granular model with firm-level productivity dynamics, but no dynamics in fundamental comparative advantage, can already account for over 60% of the mean reversion in comparative advantage documented in HLM for a panel of countries, and accounts for close to all of the mean reversion we observe in our French manufacturing dataset.[5] We illustrate further the dynamic implications of granularity for trade patterns by showing how the exit of a large firm can radically alter the export stance of a sector. In certain cases, sectors switch from being among the most export-intensive to becoming comparatively disadvantaged, as a disproportionate share of sales of the exiting granular firm is re-captured by foreign firms.

Finally, Section 6 studies some of the policy implications of granular comparative advantage. Some important dimensions of industrial and trade policy concern large individual firms, such as antitrust policy or trade policy targeted at individual firms or products. We use our

---

[5]Together with mean reversion in comparative advantage, firm dynamics in a granular open economy causes production and employment reallocation across sectors, which accounts for about 20% of total labor reallocation across firms, consistent with the empirical patterns documented in the macro-labor literature.

model to study the general equilibrium welfare and distributional implications of these policies in a granular open economy. First, we study the merger of two large domestic firms. The increase in monopoly power and the resulting destruction of consumer surplus make mergers undesirable in a closed economy. However, in an open economy, mergers can be welfare improving for the home country, as the merged home firm is able to capture a larger portion of the foreign consumer surplus. Thus, open economies might be excessively lenient towards mergers, especially in granular sectors, using them to manipulate the terms of trade.

Second, we study the incentives of a country to impose an import tariff on a single large foreign exporter. We find that in a granular world, a country prefers a tariff on the largest firm over a uniform tariff on all sectoral imports, particularly so in sectors where its trade partner enjoys a granular comparative advantage. The reason is that by taxing the largest foreign firm, a country takes advantage not only of the general-equilibrium terms-of-trade effect, operating via a reduction in the foreign wage rate, but also of the industry-level terms-of-trade improvement, due to a markup reduction by the large foreign firm.

Related literature   The term "granularity" has been coined in the macroeconomics literature, which following Gabaix (2011) has focused on the study of aggregate fluctuations driven by idiosyncratic productivity shocks (see e.g. Acemoglu, Carvalho, Ozdaglar, and Tahbaz-Salehi 2012, Carvalho and Gabaix 2013, Carvalho and Grassi 2014). The aggregate volatility consequences of granularity in an open economy have been studied by di Giovanni and Levchenko (2012) and di Giovanni, Levchenko, and Méjean (2014). Instead of aggregate volatility, we focus here on sectoral trade patterns, where the granular forces must be at least as prominent, since granularity is particularly pronounced within sectors.

In terms of modeling, we borrow from the recent trade literature, and in particular from EKS. EKS tackle a very different set of issues in the context of a single-sector model, such as explaining the prevalence of zeros in aggregate trade flows, while we develop a multi-sector environment to explore the implications of granularity for a country's comparative advantage.[6] In terms of the question studied, our paper therefore contributes to the empirical trade literature on the structure and evolution of comparative advantage, e.g. Chor (2010), Costinot, Donaldson, and Komunjer (2012), Freund and Pierola (2015), Sutton and Trefler (2016), Levchenko and Zhang (2016), and HLM.

For our analysis, we adopt a model of oligopolistic competition with variable markups, which has been used in a number of papers studying the behavior of markups, prices and market shares in an open economy (see e.g. Atkeson and Burstein 2008, Edmond, Midrigan, and Xu 2015, Amiti, Itskhoki, and Konings 2014, 2018, Hottman, Redding, and Weinstein 2015).

---

[6]In the context of import sourcing, Head, Jing, and Ries (2017) study the role of granularity of buyers in explaining hierarchy violations.

We follow Neary (2003, 2012, 2015) and Grossman and Rossi-Hansberg (2010) in studying an open economy oligopolistic environment with firms that are big in the small (at the sectoral level), but small in the big (at the economy-wide level). More generally, see Bernard, Jensen, Redding, and Schott (2018) for a recent review of the empirical and theoretical literature on the role of individual firms in international trade.

On the policy front, our study is related to the vast literature on trade policy and market structure, summarized in Helpman and Krugman (1989) and Bagwell and Staiger (2004).[7] We contribute to this literature by studying policy implications in a quantitative model of oligopolistic competition with many firms, which captures the salient features of the market structure of modern manufacturing industries.

2   A Granular Model of Comparative Advantage

This section sets up our granular model and discusses its properties. The model is a two-country multi-sector model, which combines a Ricardian DFS model across sectors with the EKS model of granular firms within each sector.[8] The model nests as special cases both the DFS-Melitz model, as firms become infinitesimal, as well as the Ricardian DFS model, as varieties of products become perfect substitutes and fixed costs tend to zero (see Appendix B.1).

Preferences   There is a unit continuum of sectors $z \in [0, 1]$. Households in each country have the same Cobb-Douglas preferences over the consumption of sectoral output $\{Q_z\}$:

$$Q = \exp\left\{\int_0^1 \alpha_z \log Q_z \, \mathrm{d}z\right\}, \tag{1}$$

where $\{\alpha_z\}_{z \in [0,1]}$ are non-negative preference parameters, which satisfy $\int_0^1 \alpha_z \mathrm{d}z = 1$, and determine the shares of household income spent on consumption across sectors.

Within each sector, there is a *finite* number of product varieties $i \in \{1, .., K_z\}$, which are combined into aggregate sectoral output using a CES aggregator:

$$Q_z = \left[\sum_{i=1}^{K_z} q_{z,i}^{\frac{\sigma-1}{\sigma}}\right]^{\frac{\sigma}{\sigma-1}}, \tag{2}$$

where $\sigma > 1$ is the within-sector elasticity of substitution, common across all sectors. The K_z product varieties available for consumption in the home market can be of both domestic and foreign origin. In the foreign market, there are K_z^* product varieties available for consumption,

---

[7]The early classics include Dixit (1984), Brander and Spencer (1984), Eaton and Grossman (1986); the examples of more recent work are Horn and Levinsohn (2001) and Costinot, Rodríguez-Clare, and Werning (2016).

[8]The EKS model is a granular version of the Melitz (2003) model, in its Chaney (2008) formulation.

which are in general different from the set of varieties marketed at home. In what follows, starred variables correspond to the foreign market.

With this demand structure, the home consumer expenditure on variety i in sector z is:

$$r_{z,i} \equiv p_{z,i}\, q_{z,i} = s_{z,i}\, \alpha_z Y \qquad \text{with} \qquad s_{z,i} \equiv \left(\frac{p_{z,i}}{P_z}\right)^{1-\sigma}, \tag{3}$$

where $p_{z,i}$ is the price and $s_{z,i}$ is the within-sector market share of the product variety, and Y is aggregate income (expenditure) in the home market. The expressions in (3) derive from the fact that with Cobb-Douglas preferences, consumers spend a constant share α_z of their income Y on purchasing varieties in sector z (i.e., $P_z Q_z = \alpha_z Y$), and within sector z the CES demand for variety i is given by $q_{z,i} = (p_{z,i}/P_z)^{-\sigma} Q_z$. The sectoral price index P_z satisfies:

$$P_z = \left[\sum_{i=1}^{K_z} p_{z,i}^{1-\sigma}\right]^{\frac{1}{1-\sigma}}. \tag{4}$$

The home and foreign households respectively supply L and L^* units of labor inelastically, with L/L^* measuring the relative size of the home country.

Production technology Each product variety is supplied by an individual firm with productivity $\varphi_{z,i}$ ($\varphi_{z,i}^*$, respectively, if the firm is foreign). Products are produced in their market of origin, and firms have access to a CRS production technology, which uses local labor, $y_{z,i} = \varphi_{z,i}\, \ell_{z,i}$. The output of the firm can be marketed domestically and exported. Exporting is associated with an iceberg trade cost $\tau \geq 1$, that is τ units of product need to be shipped for one unit to arrive at the foreign market. Therefore, the marginal cost of supplying the home market is constant and equal to:

$$c_{z,i} = \begin{cases} w/\varphi_{z,i}, & \text{if } i \text{ is a home variety,} \\ \tau w^*/\varphi_{z,i}^*, & \text{if } i \text{ is a foreign variety,} \end{cases} \tag{5}$$

where w and w^* are respectively the home and foreign wage rates. The marginal cost of serving the foreign market is defined symmetrically, and we denote it with $c_{z,i}^*$.

Furthermore, there is a fixed market access cost F in local units of labor, which is independent of the origin of the firm, i.e. applies both for local firms and exporters. As a result, the differential selection of domestic and foreign firms into the local market is driven by iceberg trade costs, rather than by a differential fixed access cost. In each market, we sort all potential entrants in the increasing order of marginal cost $c_{z,i}$ ($c_{z,i}^*$ in foreign, respectively). The index i refers to the marginal cost ranking of a firm in a given market, so that the same firm is in general represented by different indexes in different markets.

Productivity draws   We denote with M_z a potential (shadow) number of domestic products in sector z. M_z is the realization of a Poisson random variable with parameter $\bar{M}_z$, so that $\mathbb{E}(M_z) = \bar{M}_z$. Each of the M_z potential entrants takes an *iid* productivity draw from a Pareto distribution with a shape parameter θ and lower bound $\underline{\varphi}_z$.[9] We borrow this structure of productivity draws from the earlier work of Bernard, Eaton, Jensen, and Kortum (2003) and EKS. It results both in a tractable model environment and in a realistic cross-sectional distribution of firm revenues.

Given this productivity structure, the combined parameter:

$$T_z \equiv \bar{M}_z \cdot \underline{\varphi}_z^\theta \tag{6}$$

is a sufficient statistic that determines the expected productivity of a sector.[10] Intuitively, a sector is more productive either if there are more potential entrants (i.e., productivity draws), equal to $\bar{M}_z$ in expectation, or if the average productivity of a potential entrant is high, which is given by $\frac{\theta}{\theta-1} \underline{\varphi}_z$.

The pool of foreign potential products and the ensuing productivity draws are obtained in a symmetric way, with country-sector-specific parameters $\bar{M}_z^*$ and $\underline{\varphi}_z^*$, resulting in a sufficient statistic for the expected sectoral productivity $T_z^* = \bar{M}_z^* \cdot (\underline{\varphi}_z^*)^\theta$. The ratio T_z/T_z^* varies across sectors z and determines the expected relative productivity of the two countries, and thus is a measure of the home's fundamental comparative advantage. T_z/T_z^* is the only source of comparative advantage in the continuous DFS-Melitz limit of the model (see Appendix B.1).

Market structure   For a given set of K_z entrants, the firms play a Bertrand oligopolistic price setting game, similar to Atkeson and Burstein (2008). Specifically, firm $i \in \{1, .., K_z\}$ chooses its prices $p_{z,i}$, taking as given the prices of its competitors $\{p_{z,j}\}_{j \neq i}$, to maximize its profits from serving the home market:

$$\Pi_{z,i} = \max_{p_{z,i}} \left\{ \left(p_{z,i} - c_{z,i}\right) \frac{p_{z,i}^{1-\sigma}}{\sum_{j=1}^{K_z} p_{z,j}^{1-\sigma}} \alpha_z Y - wF \right\},$$

where we used the expressions for the market share of the firm (3) and the sectoral price index (4). While firms are large within their industries, and hence internalize their effect on

---

[9]Formally, the realized number of products M_z has the pdf $\mathbb{P}\{M_z = m\} = e^{-\bar{M}_z} \bar{M}_z^m / m!$ for $m = 0, 1, 2..$, while the cdf of productivity draws φ is given by $G_z(\varphi) = 1 - (\underline{\varphi}_z/\varphi)^\theta$.

[10]In particular, EKS show that the number of productivity draws above any given $\varphi > \underline{\varphi}_z$ is a Poisson random variable with a mean parameter $T_z \varphi^{-\theta}$, increasing in T_z and decreasing in φ. As long as the least efficient product stays inactive in equilibrium, the model is invariant to various combinations of $\bar{M}_z$ and $\underline{\varphi}_z$, which result in the same T_z. A convenient limiting case with $\bar{M}_z \to \infty$ and $\underline{\varphi}_z \to 0$ (holding T_z constant) ensures that there is always a sufficient number of draws and the least productive draw is necessarily inactive.

the sectoral price index (4), they are still infinitesimal at the level of the whole economy, since the model features a continuum of sectors, different from EKS. Therefore, firms take wage rates w and w^* as given, and hence treat $c_{z,i}$ as exogenous to their decisions.

The solution to this Bertrand-Nash competition game is a markup price setting rule:

$$p_{z,i} = \frac{\varepsilon_{z,i}}{\varepsilon_{z,i} - 1} \cdot c_{z,i}, \qquad \text{where} \qquad \varepsilon_{z,i} \equiv \varepsilon(s_{z,i}) = \sigma(1 - s_{z,i}) + s_{z,i}, \qquad (7)$$

with the market share of the firm $s_{z,i}$ defined in (3), and with $\varepsilon_{z,i} \in [1, \sigma]$ measuring the effective elasticity of residual demand for the product of the firm. This elasticity decreases, and hence the markup $\mu_{z,i} \equiv \frac{p_{z,i}}{c_{z,i}} = \frac{\varepsilon_{z,i}}{\varepsilon_{z,i}-1}$ increases, with the market share of the firm $s_{z,i}$. This contrasts with the constant-markup pricing under monopolistic competition in the continuous DFS-Melitz limit of the model.[11]

To summarize, given the set of entrants and their marginal costs $\{c_{z,i}\}_{i=1}^{K_z}$, the equilibrium in the Bertrand-Nash price setting game is a vector of prices and market shares $\{p_{z,i}, s_{z,i}\}_{i=1}^{K_z}$ and a sectoral price index P_z, which solve the fixed point defined by (7), (3) and (4). While there is no analytical characterization of the resulting prices and market shares, the equilibrium is unique and has the property that prices $p_{z,i}$ increase with marginal costs $c_{z,i}$, while markups $\mu_{z,i} = p_{z,i}/c_{z,i}$ and market shares $s_{z,i}$ decrease with $c_{z,i}$. Furthermore, the equilibrium firm profits from serving the home market are given by:

$$\Pi_{z,i} \equiv \Pi_z(s_{z,i}) = \frac{s_{z,i}}{\varepsilon(s_{z,i})} \alpha_z Y - wF. \qquad (8)$$

Indeed, operating profits are a fraction $\frac{1}{\varepsilon_{z,i}} = \frac{p_{z,i}-c_{z,i}}{p_{z,i}}$ of revenues (3), which equal the firm's share of the sectoral expenditure in the market, $s_{z,i}\alpha_z Y$. In equilibrium, firms with higher market shares command higher profits.

The price setting equilibrium in the foreign market is symmetric, resulting in prices, market shares and profits $\{p_{z,i}^*, s_{z,i}^*, \Pi_{z,i}^*\}_{i=1}^{K_z^*}$, given the set of entrants and their marginal costs $\{c_{z,i}^*\}_{i=1}^{K_z^*}$. Due to linearity of the production function, each firm's profit maximization problem is separable across markets, and hence can be considered one market at a time.

Entry An equilibrium of the entry game is achieved when for a subset of firms equilibrium profits given by (8) are non-negative, while for any additional entrant profits upon entry would

---

[11]Much of the earlier granularity literature (including Carvalho and Grassi 2014, di Giovanni and Levchenko 2012) adopts an *ad hoc* assumption of constant markups. The quantitative pricing-to-market literature following Atkeson and Burstein (2008) studies oligopolistic competition with variable markups, but adopts competition in quantities, which is qualitatively similar but results in greater markup variability (see discussion in Amiti, Itskhoki, and Konings 2018). We adopt a more natural case of oligopolistic competition in prices, following EKS, which results in a less pronounced quantitative difference from the constant markup case.

be negative. With a discrete number of potential entrants, there may exist multiple equilibria in the entry game. We therefore consider a sequential entry game in each market separately. Specifically, firms with lower marginal costs of serving a given market, $c_{z,i}$, move first. We assign the indexes i such that $c_{z,1} \leq c_{z,2} \leq ..$, and hence firms with lower indexes i choose whether to enter first.[12] With this equilibrium selection, the entry game has a unique cutoff equilibrium, so that only firms with marginal costs below some cutoff enter the market.

Formally, denote by $s_{z,i}^{K_z}$ the market share of firm $i \leq K_z$ resulting from the price-setting game when K_z firms choose to enter. The corresponding profits are given by $\Pi_{z,i}^{K_z} = \Pi_z(s_{z,i}^{K_z})$ defined in (8). We already know that, for a given K_z, $s_{z,i}^{K_z}$ is decreasing in i. Furthermore, it is easy to verify that $s_{z,i}^{K_z}$ is decreasing in K_z for all i, that is $s_{z,i}^{K_z} < s_{z,i}^{K_z+1}$ for all $i \leq K_z$. Intuitively, the entry of any additional firm reduces markups of all existing firms. Therefore, since $\Pi_z(s_{z,i}^{K_z})$ is a monotonically increasing function of $s_{z,i}^{K_z}$, there exists a unique K_z such that $\Pi_{z,i}^{K_z} \geq 0$ for all $i \leq K_z$ and $\Pi_{z,i}^{K} < 0$ for all $i > K_z$ and $K > K_z$. This K_z is the equilibrium number of entrants, and c_{z,K_z} is the cutoff cost level. Note that, due to monotonicity, it is sufficient to find the unqiue K_z such that $\Pi_{z,K_z}^{K_z} \geq 0$ and $\Pi_{z,K_z+1}^{K_z+1} < 0$.

General equilibrium   is a vector of wage rates and incomes (w, w^*, Y, Y^*), such that labor markets clear in both countries and aggregate incomes equal aggregate expenditures. In particular, in the home country

$$Y = wL + \Pi, \tag{9}$$

where Π are aggregate profits of all home firms distributed to home households:

$$\Pi = \int_0^1 \left[\sum_{i=1}^{K_z} \iota_{z,i} \Pi_z(s_{z,i}) + \sum_{i=1}^{K_z^*} (1 - \iota_{z,i}^*) \Pi_z^*(s_{z,i}^*) \right] \mathrm{d}z, \tag{10}$$

with profit function $\Pi_z(s_{z,i})$ defined in (8) and $\iota_{z,i} \in \{0, 1\}$ denoting the indicator for whether firm i in sector z in the domestic market is of local origin, and by analogy $\iota_{z,i}^*$ for the foreign market. The equality between expenditure Y and income $wL+\Pi$ implies home budget balance and hence trade balance. We normalize $w = 1$ as numeraire and omit the foreign budget constraint by Walras' law.

Labor market clearing requires that the aggregate labor income wL equals the total expenditure of all firms on domestic labor:

$$wL = \int_0^1 \left[\alpha_z Y \sum_{i=1}^{K_z} \iota_{z,i} \frac{s_{z,i}}{\mu(s_{z,i})} + \alpha_z Y^* \sum_{i=1}^{K_z^*} (1 - \iota_{z,i}^*) \frac{s_{z,i}^*}{\mu(s_{z,i}^*)} + wFK_z \right] \mathrm{d}z. \tag{11}$$

---

[12]Note that index i is not a property of a firm, but rather of a firm-market pair. A firm is characterized by its origin and productivity draw φ, and a given firm in general has different indexes i in the two markets.

The three terms on the right-hand side of (11) correspond to expenditure on domestic labor for (i) production for domestic market, (ii) production for foreign market, and (iii) entry of firms in the domestic market, respectively. Note that $s_{z,i}\alpha_z Y/\mu(s_{z,i})$ is revenues from domestic sales (3) divided by markup $\mu(s_{z,i})$, which equals variable costs, i.e. expenditure on production labor. Recall that the markup $\mu(s_{z,i}) = \frac{\varepsilon_{z,i}}{\varepsilon_{z,i}-1}$ with $\varepsilon_{z,i}$ defined in (7). Furthermore, K_z is the total number of entrants, domestic and foreign, which all pay a fixed cost F in terms of domestic labor. A parallel market clearing condition to (11) holds in the foreign country.

Aggregate equilibrium conditions (9) and (11), together with their foreign counterparts, and under normalization $w = 1$, allow to solve for the aggregate equilibrium vector $\mathbf{X} \equiv (w, w^*, Y, Y^*)$, given the sectoral equilibrium vector $\mathbf{Z} \equiv \left\{ K_z, \{s_{z,i}\}_{i=1}^{K_z}, K_z^*, \{s_{z,i}^*\}_{i=1}^{K_z^*} \right\}_{z \in [0,1]}$.[13] In turn, given the aggregate equilibrium vector $\mathbf{X}$, the solution to the entry and price-setting game in each country-sector yields the sectoral equilibrium vector $\mathbf{Z}$. The resulting fixed point $(\mathbf{X}, \mathbf{Z})$ is the equilibrium in the granular economy.

Properties of the granular model   In the following sections, we use the granular model to quantify the role played by individual firms in shaping the comparative advantage of a country. To set the stage for this analysis, we now discuss the properties of one object of interest, which summarizes the comparative advantage of a country in a given sector, namely the *foreign share*. The sectoral foreign share is the cumulative market share of all foreign firms serving the home market in a given sector z, and we denote it by:[14]

$$\Lambda_z \equiv \frac{X_z^*}{\alpha_z Y} = \sum_{i=1}^{K_z} (1 - \iota_{z,i}) s_{z,i}, \tag{12}$$

where X_z^* is total home imports (foreign exports) and $\alpha_z Y$ is total home absorption in sector z. By analogy, the home export share in sector z is given by $\Lambda_z^* = X_z/(\alpha_z Y^*)$, where X_z denotes total home exports in sector z.

In the granular model, the realized foreign share is a random variable, which depends on the productivity of the home and foreign firms in sector z. These productivity draws are shaped, in turn, by the fundamental comparative advantage of the sector (T_z/T_z^*) and the idiosyncratic realizations of firm draws from the Poisson-Pareto process described above. The structure of the model provides a natural decomposition of the foreign share Λ_z into these fundamental and granular components. In particular, the *expected* foreign share, conditional

---

[13]One of the four aggregate equilibrium conditions is redundant by Walras Law, and is replaced by a numeraire normalization. Also note that in the closed economy conditions (9) and (11) are equivalent, and amount to $Y/w = \bar{\mu}[L - FK]$, where $K = \int_0^1 K_z \mathrm{d}z$ is the total number of firms serving the home economy and $\bar{\mu} = \left[\int_0^1 \alpha_z \sum_{i=1}^{K_z} s_{z,i}/\mu(s_{z,i}) \right]^{-1}$ is the (harmonic) average markup.

[14]One minus the foreign share, $1 - \Lambda_z$, is the *home share*, an object that features prominently in the gains from trade literature (see Arkolakis, Costinot, and Rodríguez-Clare 2012, henceforth ACR).

on fundamental comparative advantage of the sector T_z/T_z^*, is given by:[15]

$$\Phi_z \equiv \mathbb{E}_T\{\Lambda_z\} = \mathbb{E}\{\Lambda_z \,|\, T_z/T_z^*\} = \frac{1}{1 + (\tau\omega)^\theta \cdot T_z/T_z^*}. \tag{13}$$

The expected foreign share Φ_z decreases in all sectors in the trade cost τ and in the relative foreign wage rate $\omega \equiv w^*/w$. Across sectors, variation in Φ_z is shaped by the fundamental comparative advantage T_z/T_z^*. The expression in (13) is familiar from the quantitative trade literature, following Eaton and Kortum (2002), and it characterizes the *realized* trade shares in the continuous limit of our granular model (see Appendix B.1). In short, the granular model has, in *expectation*, the same sectoral trade shares as the continuous model.

Due to granularity, however, the realized trade shares Λ_z differ from their expectation Φ_z. We define the discrepancy between the realized and expected shares as the *granular residual*:

$$\Gamma_z \equiv \Lambda_z - \Phi_z, \qquad \text{such that} \qquad \mathbb{E}_T\{\Gamma_z\} = \mathbb{E}_T\{\Lambda_z - \Phi_z\} = 0. \tag{14}$$

Defined this way, the granular residual Γ_z is a *scalar sufficient statistic* for the effect of all idiosyncratic productivity draws within a sector, $\{\varphi_{z,j}\}_{i=j}^{M_z}$ and $\{\varphi_{z,j}^*\}_{j=1}^{M_z^*}$, on the sectoral trade pattern Λ_z relative to its expected value Φ_z. By construction, granular residuals have an expected value of zero and are uncorrelated with the fundamental comparative advantage Φ_z. This provides a convenient way to decompose the cross-sectional variation in the realized trade patterns Λ_z into the contribution of the fundamental and granular comparative advantage, Φ_z and Γ_z, as we do in Section 4.

Note that within-sector granularity does not create extra trade at the aggregate level, compared to the continuous benchmark with the same expected shares. Indeed, total imports are:[16]

$$X^* = \int_0^1 X_z^* \mathrm{d}z = Y \int_0^1 \alpha_z \big[\Phi_z + \Gamma_z\big] \mathrm{d}z = \Phi Y, \tag{15}$$

---

[15]This result applies despite the fact that market shares $s_{z,i}$ are complex non-linear transformation of firm productivity draws $\varphi_{z,i}$, which in particular depend on the endogenous markups $\mu_{z,i}$ that do not admit an analytical characterization. Nonetheless, due to the Poisson-Pareto productivity structure and the common entry cost F, the distribution of equilibrium market shares conditional on entry in a given market is the same for foreign and home firms. However, the expected number of entrants differs for foreign and home firms, and its ratio is given by Φ_z. The formal derivation of (13) is provided in Appendix B.2.

[16]Similarly, $X = \Phi^* Y^*$ is the aggregate value of exports. Due to local fixed costs, the trade balance in general is not $X = X^*$, but is instead

$$\Phi\left[Y - wFK\right] = \Phi^*\left[Y^* - w^* F^* K^*\right],$$

where K and K^* denote the total number of firms serving the two markets across all sectors. Indeeed, $[Y - wFK]$ are aggregate sales in the home market net of fixed entry costs, and a fraction Φ of these net sales is foreign income from exports. See Appendix B.2 for the derivation of (15) and the resulting simplification of the general equilibrium system (9)–(11).

where $\Phi \equiv \mathbb{E}\Phi_z = \int_0^1 \alpha_z \Phi_z \mathrm{d}z$ is the *aggregate* foreign share, which integrates Φ_z across sectors with different fundamental comparative advantage T_z/T_z^*. The aggregate amount of trade in a continuous model is also given by $X^* = \Phi Y$. While granularity does not create extra trade in the aggregate, it changes the distribution of trade flows across sectors, contributing to the patterns of a country's comparative advantage.

3   Estimation of the Granular Model

In a continuous trade model, the observed trade flows are assumed to be shaped entirely by the fundamental forces Φ_z in (13), making estimation and quantification of the continuous model particularly straightforward (see Eaton and Kortum 2002, and the vast quantitative literature it gave rise to). In contrast, the observed trade flows in a granular model confound both fundamental and idiosyncratic (granular) forces, $\Lambda_z = \Phi_z + \Gamma_z$. This poses an interesting identification challenge, which we address in this section. We first describe the data, and then detail our estimation procedure, present the estimation results and discuss the fit of the model.

3.1   Data

We estimate a two-country model with France as home and the rest of the world (ROW) as foreign. We use a dataset of French firms (BRN), which reports information on the balance sheets of firms declared for tax purposes. All firms with revenues over 730,000 euros are included. It reports in particular information on both domestic and export sales, as well as 4-digit industry classification, at the firm level. We use 2005 as our reference year for estimation. We match this data with international trade data from Comtrade, to get the aggregate imports and exports of France in each industry.[17] This leaves us with $\tilde{N} = 119$ manufacturing sectors at the 4-digit level with an average of about 350 French firms per sector.

We use tildes to denote the empirical variables that correspond to the theoretical objects that can be measured in the granular model of Section 2. The merged data allows us to construct French sectoral expenditure $\tilde{\alpha}_z \tilde{Y}$ as the sum of sectoral imports $\tilde{X}_z^*$ and domestic sales of all French firms $\{\tilde{r}_{z,j}\}_{j=1}^{\tilde{M}_z}$, where j is the rank of French firms and $\tilde{M}_z$ is the observed number of French firms in each sector $z = 1, ..., \tilde{N}$. Taking the ratio of sectoral imports to sectoral expenditure, we obtain the foreign share in the home market $\tilde{\Lambda}_z = \frac{\tilde{X}_z^*}{\tilde{\alpha}_z \tilde{Y}}$. Additionally, we introduce a measure of French export intensity:

---

[17]The industry classification used in the French data is the French NAF (based on European NACE classification), whereas the trade data uses ISIC rev3. We convert the French data into the ISIC rev3 classification using the crosswalk between NACE and ISIC available from UNstats.

$$\tilde{\Lambda}_z^{*\prime} = \frac{\tilde{X}_z}{\tilde{\alpha}_z \tilde{Y}}, \tag{16}$$

where we normalize exports by domestic expenditure. In the model, this measure is proportional to the French export share, $\Lambda_z^* = \Lambda_z^{*\prime} \frac{Y}{Y^*}$, but is easier to measure since we do not observe sectoral expenditure in the ROW.

Lastly, we construct the relative sales share of each French firm among all domestic firms in the domestic market:[18]

$$\tilde{s}_{z,j} = \tilde{r}_{z,j} \Big/ \sum_{j'=1}^{\tilde{M}_z} \tilde{r}_{z,j'}. \tag{17}$$

We rely on the dataset $\tilde{\Xi} \equiv \left\{ \tilde{\alpha}_z \tilde{Y}, \tilde{\Lambda}_z, \tilde{\Lambda}_z^{*\prime}, \tilde{M}_z, \{\tilde{s}_{z,j}\}_{j=1}^{\tilde{M}_z} \right\}_{z=1}^{\tilde{N}}$ to estimate the model using a simulated method of moments procedure (SMM), which we describe next.

3.2 Estimation procedure

For estimation, we parameterize the fundamental comparative advantage across sectors as draws from a log-normal distribution with parameters μ_T and σ_T, that is:

$$\log\left(T_z/T_z^*\right) \sim \mathcal{N}(\mu_T, \sigma_T^2). \tag{18}$$

Hanson, Lind, and Muendler (2016) show evidence that the distribution of measured comparative advantage across countries and sectors is well-approximated by a log-normal distribution, a property that we check holds in the estimated version of our granular model under the above distributional assumption.

With this assumption, we estimate the model parameters in two steps. In the first step, we calibrate Cobb-Douglas shares from the data as equal to the sectoral expenditure shares, and we report the histogram of the resulting α_z in Appendix Figure A1.[19] The relative sectoral productivity parameters T_z/T_z^* and the Cobb-Douglas shares α_z are the only parameters that vary across sectors, and in all other respects the sectoral outcomes emerge from a symmetric data generating process shaped by the remaining model parameters $(\sigma, \theta, \tau, F)$, which in the baseline case are assumed to be common across sectors and countries.[20]

---

[18]Note that the sales share $\tilde{s}_{z,j}$ differs from the home market share $s_{z,i}$, used for equilibrium characterization in Section 2. Specifically, $\tilde{s}_{z,j} = s_{z,i(j)}/(1 - \Lambda_z)$, where $i(j)$ is the rank of home firm j among all firms serving the home market in sector z.

[19]In the data, the largest Cobb-Douglas share is 7.5%, the 90th percentile is 1.7%, the median is 0.4% and the 10th percentile is 0.1%. By construction, the mean share is $1/\tilde{N} = 0.8\%$. In the model, we set $\alpha_z = \tilde{N}\tilde{\alpha}_z$, so that the average $\mathbb{E}\alpha_z = 1$, as is required by our model with a continuum of sectors.

[20]Our baseline model keeps cross-sectoral parametric heterogeneity to a minimum, yet the estimation procedure can be immediately extended to heterogeneity in other parameters (such as θ, σ and F) by analogy with heterogeneity in T_z/T_z^*, provided the availability of relevant empirical moments for identification. See Section 4 for our robustness analysis, which allows for heterogeneity in productivity parameters θ_z across sectors.

In the second step, we use a simulated method of moments (SMM) procedure to estimate the six parameters of the model $\Theta \equiv (\sigma, \theta, \tau, F, \mu_T, \sigma_T)$. The estimation is done in full general equilibrium, taking into acount in particular the impact of these parameters on firm entry, heterogeneous markups, and equilibrium prices and quantities. Our approach to statistical inference in this granular model leverages its multi-sector nature. We view each sector as a draw from the parametric data generating process (DGP) described in Section 2. The SMM procedure treats each sector as a (multi-dimensional) observation from the structural DGP, with parameters that need to be estimated. In particular, we treat the $\tilde{N}$ sectors in the data as a finite number of draws from a model with a continuum of sectors, and our statistical inference considers the asymptotics as $\tilde{N}$ increases unboundedly.

For a given parameter vector Θ, we simulate the model, compute a list of cross-sectoral moments $\mathcal{M}(\Theta)$, and contrast them with the equivalent moments in the data $\tilde{\mathbf{m}}$. We search for the parameter vector $\hat{\Theta}$ that minimizes the distance between the model and the empirical moments, according to the loss function $\mathcal{L}(\Theta) \equiv \big(\mathcal{M}(\Theta) - \tilde{\mathbf{m}}\big)' \mathbf{W} \big(\mathcal{M}(\Theta) - \tilde{\mathbf{m}}\big)$, where $\mathbf{W}$ is a weighting matrix. Specifically, we search for the best-fitting parameters on a series of coarse-to-fine grids, completed by a local minimum search starting from a subset of best-fitting points from the grid. The full SMM procedure is described in Appendix C.

Normalizations   In the model, home and foreign differ in labor endowments L and L^*. The model scales with L, as long as we keep L/L^* and L/F constant. In other words, L simply determines the units of labor, and hence we normalize $L = 100$, and estimate L^*/L and F/L. We calibrate $w/w^* = 1.13$, which corresponds to the ratio of wages in France to the average wage of its trading partners weighted by trade values. As we discuss below, this imposes a general equilibrium restriction on the other parameters, in particular the relative labor supplies L/L^*, which the procedure estimates along with the model parameters. Given the Cobb-Douglas preference structure, all variables of interest in the model scale with the common level of productivity, and therefore we normalize $T_z^* \equiv 1$ without loss of generality.[21]

Lastly, in our estimation, we find that the elasticity of substitution σ and the productivity parameter θ are weakly separately identified. Indeed, the moments tend to be sensitive to the ratio $\kappa \equiv \theta/(\sigma - 1)$, which approximately corresponds to the Pareto tail of the sales distribution across firms, but not to the values of θ and σ separately. Therefore, we choose to fix $\sigma = 5$ and estimate the constrained model with five parameters $\Theta' = (\theta, \tau, F, \mu_T, \sigma_T)$.[22] This reduces the parameter space and improves the precision of estimation for the remaining parameters.

---

[21]Note that if productivity in sector z doubles in both countries, the quantity in this sector doubles and the price halves, without any effect on market shares within or across sectors.

[22]The value of $\sigma = 5$ (within 4-digit sectors) is conventional in the trade literature (see Broda and Weinstein 2006). When we estimate the unrestricted model, we find $\sigma = 4.927$, yet imprecisely estimated.

3.3 Moments and identification

We target 15 empirical moments, which correspond to averages and standard deviations of sectoral outcomes. Table 2 summarizes the list of the moments. With 15 moments, the model with 5 parameters is overidentified, and variation in any of the parameters tends to affect all moments simultaneously. Nonetheless, some parameters are particularly sensitive to specific moments (see Andrews, Gentzkow, and Shapiro 2017). We provide here a discussion of the main forces ensuring identification.

First, we target the cross-sectoral mean and standard deviation of the log number of French firms operating in the domestic market, $\log \tilde{M}_z$. Intuitively, the fixed cost parameter F is particularly sensitive to the average number of firms, which is the key moment in identifying it. Second, we target the average and standard deviation across sectors of two measures of within-industry concentration — the relative size of the largest and top-3 largest French firms within-industry relative to other French firms, that is $\tilde{s}_{z,1}$ and $\sum_{j=1}^{3} \tilde{s}_{z,j}$, as defined in (17). The parameter combination $\kappa = \theta/(\sigma - 1)$ is particularly sensitive to these moments. Recall that κ is closely associated with the shape parameter of the sales distribution, and hence determines the average industry concentration.[23] Therefore, given the calibrated value of the elasticity of substitution σ, the concentration moments are key in identifying the productivity parameter θ, as we illustrate in the Appendix Figure A3.

Third, we target a set of five moments describing cross-sectoral patterns of international trade. In particular, we target the averages and standard deviations of import shares in the French market $\tilde{\Lambda}_z$, and of export sales of French firms relative to domestic expenditure $\tilde{\Lambda}_z^{*\prime}$, given in (16). These trade moments help identify the size of the trade cost τ and the average productivity advantage of France μ_T. Indeed, from (13), expected foreign shares (Φ_z and Φ_z^*) are both decreasing in τ, while one is decreasing and the other is increasing in T_z/T_z^*, which has a mean governed by μ_T. Furthermore, the value of σ_T, the standard deviation of the home country relative productivities, is particularly sensitive to the variation of trade shares across sectors. Finally, we target the fraction of French sectors in which export sales exceed the overall domestic sales of French firms. Due to trade costs, such sectors can emerge only when the ROW is larger than France, $Y^* > Y$. Therefore, this moment identifies the relative size of France, Y/Y^* and L/L^*, given the calibrated value of the relative wages $\omega = w^*/w$.[24]

The final four moments capture the correlations between French import and export shares,

---

[23]In a granular Zipf's world (i.e., when $\kappa \approx 1$), the number of firms in a sector (and hence the value of the fixed cost F) has only a small effect on the size of the largest firms. We use this correlation moment below as an external check of the fit of the estimated model.

[24]While this moment is directly informative of the relative size of the two economies, the model is also constrained by the general equilibrium trade balance restriction, $\Phi[Y - wFK] = \Phi^*[Y^* - w^*F^*K^*]$, which links the aggregate import shares, Φ and Φ^*, with the relative size of the two economies.

$\tilde{\Lambda}_z$ and $\tilde{\Lambda}_z^{*\prime}$, and the largest-French-firm relative sales shares, $\tilde{s}_{z,1}$ and $\sum_{j=1}^{3} \tilde{s}_{z,j}$. Specifically, we target the regression coefficients of $\tilde{\Lambda}_z$ and $\tilde{\Lambda}_z^{*\prime}$ separately on $\tilde{s}_{z,1}$ and $\sum_{j=1}^{3} \tilde{s}_{z,j}$, controlling in all four regressions for the size of the sector (log total domestic expenditure, $\log \tilde{\alpha}_z \tilde{Y}$). We denote these regression coefficients with $\hat{b}_j$ and $\hat{b}_j^*$ for $j \in \{1, 3\}$ respectively. These moments are instrumental for identifying the relative importance of fundamental versus granular forces in shaping trade shares. In a continuous model, trade flows are shaped exclusively by the expected productivity of a sector, and conditional on it should not be related to the moments of the relative size distribution of firm sales. In the data, however, we see a clear correlation pattern — sectors with more concentrated sales at the top have larger export shares, while there is no clear relationship for import shares. As we illustrate in the Appendix Figure A3, the model can capture this only with a particular combination of κ and σ_T, where the two parameters shape, respectively, the granular and fundamental forces.

3.4 Estimation results and model fit

Estimated parameters Table 1 reports the SMM estimates of the model parameters and their standard errors (described in Appendix C), along with the corresponding auxiliary variables implied by the general equilibrium of the estimated model. Overall, the parameters of the model are quite precisely estimated.

We point out a few features of the estimated parameters. First, $\kappa = \theta/(\sigma - 1)$ that controls the Pareto shape parameter of the sales distribution is estimated to equal 1.08, somewhat above 1, hence exhibiting *thinner* tails relative to Zipf's law (see Gabaix 2009). Since granular phenomena are stronger when distributions have thicker tails, our estimated model offers a conservative bound on the role of granularity in trade flows, as we further discuss below.

Next, we estimate μ_T to be positive, albeit small. A positive μ_T means that France has on average slightly better productivity draws relative to its trade partners, in line with the calibrated higher wage rate $w/w^* = 1.13$. The estimated value of $\sigma_T = 1.4$, the standard deviation of fundamental comparative advantage, is large. It suggests that in the cross-section of sectors, a one standard deviation increase in fundamental comparative advantage corresponds to a four-fold increase in the relative productivity of the two countries. Below, we discuss the relative role of $\kappa = 1.08$ and $\sigma_T = 1.4$ in generating the patterns of trade across sectors.

We find that the iceberg trade costs are $\tau = 1.34$, broadly in line with the estimates in the literature (see Anderson and van Wincoop 2004). Note that the estimated model implies that France is 1.7 times smaller than the rest of the world in terms of population (and 1.5 times smaller in terms of GDP). This is, of course, an abstraction of a two-country model with a common iceberg trade cost τ separating the two regions. The appropriate interpretation of L^*/L in the model is the relative size of the ROW, in which the individual countries are

Table 1: Estimated parameters

Parameter	Estimate	Std. error	Auxiliary variables	
σ	5	—	$\kappa = \frac{\theta}{\sigma-1}$	1.077
θ	4.307	0.246		
τ	1.341	0.061	w/w^*	1.130
F $(\times 10^5)$	0.946	0.252	L^*/L	1.724
μ_T	0.137	0.193	Y^*/Y	1.526
σ_T	1.422	0.232	Π/Y	0.211

discounted by their economic distance to France (i.e., if countries trade little with France, their population weight is heavily discounted). The model implies an aggregate share of profits in GDP (Π/Y) equal to 21%, broadly in line with the national income accounts, without being targeted in the estimation procedure.

Model fit   Table 2 reports the model-based values of the 15 moments used in estimation, and compares them with their empirical counterparts. The table also reports the percentage contribution of each moment to the overall loss function $\mathcal{L}(\hat{\Theta})$, as we describe in Appendix C. Overall, the model provides a reasonable fit to the data for the 15 moments targeted in estimation, as we now discuss. In addition, the Appendix Figure A4 shows the fit of the model over the whole distribution of sectoral outcomes, rather than just for the means and standard deviations reported in Table 2.

The model accurately matches the distribution of the number of firms across sectors. The median sector has around 350 French firms with a large variation across sectors: a sector at the 25th percentile has just over 100 firms and a sector at the 75th percentile has over 700 firms. The model also fits well the average size of the largest and top-3 largest French firms, which are about 20% and 35% of the overall French-firm sales, respectively. The ability of the model to closely replicate the distribution of the number of firms and the market shares of the large firms across sectors is important in the quantitative analysis of granularity. Furthermore, in the model, like in the data, average export and import shares across French manufacturing sectors are both around 35%.[25]

The regression coefficients of the sectoral trade share on either the top-firm or top-3 domestic sales shares in the data are 0.20–0.25 for exports and around zero for imports, and the model matches these patterns accurately. The table further reports the OLS standard errors for these regression coefficients, and the model is able to reproduce them as well, even though they are not targeted in estimation. In particular, the regression coefficients for the

---

[25]Note that trade is balanced in the model, which is not far from the small empirical manufacturing trade deficit that France ran in 2005.

Table 2: Moments used in SMM estimation

	Moments		Data, $\tilde{\mathbf{m}}$	Model, $\bar{\mathcal{M}}(\hat{\Theta})$	Loss (%)
1.	Log number of firms, mean	$\log \tilde{M}_z$	5.631	5.624	0.1
2.	— st. dev.		1.451	1.222	7.9
3.	Top-firm sales share, mean	$\tilde{s}_{z,1}$	0.197	0.206	3.5
4.	— st. dev.		0.178	0.149	3.8
5.	Top-3 sales share, mean	$\sum_{j=1}^{3} \tilde{s}_{z,j}$	0.356	0.343	2.0
6.	— st. dev.		0.241	0.175	11.5
7.	Imports/dom. sales, mean	$\tilde{\Lambda}_z$	0.365	0.351	2.2
8.	— st. dev.		0.204	0.268	14.8
9.	Exports/dom. sales, mean	$\tilde{\Lambda}_z^{*\prime}$	0.328	0.350	6.0
10.	— st. dev.		0.286	0.346	6.5
11.	Fraction of sectors with exports>dom. sales	$\mathbb{P}\left\{ \frac{\tilde{X}_z>}{\tilde{Y}_z - \tilde{X}_z^*} \right\}$	0.185	0.092	37.9
Regression coefficients:[†]					
12.	export share on top-firm share	$\hat{b}_1^*$	0.215 (0.156)	0.243 (0.104)	2.6
13.	export share on top-3 share	$\hat{b}_3^*$	0.254 (0.108)	0.232 (0.090)	1.1
14.	import share on top-firm share	$\hat{b}_1$	-0.016 (0.097)	-0.020 (0.079)	0.0
15.	export share on top-3 share	$\hat{b}_3$	0.002 (0.074)	-0.005 (0.069)	0.1

Note: Last column reports the contribution of the moment to the loss function $\mathcal{L}(\hat{\Theta})$, as described in Appendix C.
[†]For moments 12-15, regressions include log domestic sectoral expenditure as a control for the size of the sector; OLS standard errors are reported in brackets.

export share are significant with t-statistics over 2, while the coefficients for import shares are well-estimated zeros with t-statistics close to zero.

In contrast, one moment where the fit of the model is not as good is the fraction of sectors with exports exceeding domestic sales: the model predicts 9% of such sectors against 18% in the data. Note that the presence of such sectors is only possible in a model with $Y^* > Y$, i.e. when France is smaller than the ROW. Our simplified two-country geography is likely the reason why the model has a hard time matching this moment. This is the only moment for which the model is off by a substantial amount, accounting for 38% of the loss function (the SMM objective), as can be seen in the last column of Table 2.

Moments not targeted in estimation   We consider here a series of overidentification checks by exploring the fit of the moments not targeted directly in our estimation procedure. First, we consider an additional measure of concentration by estimating the Pareto shape parame-

ter $\hat{\kappa}_z$ of domestic sales of French firms, industry-by-industry. Specifically, we follow Gabaix and Ibragimov (2011) and estimate by OLS:

$$\log(j - 0.5) = const - \hat{\kappa}_z \cdot \log \tilde{s}_{z,j} + \epsilon_{z,j}^{\kappa}, \tag{19}$$

$$\text{Data:} \quad \underset{[0.817, 1.208]}{1.015}$$

$$\text{Model:} \quad \underset{[1.010, 1.169]}{1.096}$$

where j is the domestic-sales rank of French firms in industry z, and we only use the top 25% of firms in each industry.[26] A lower $\hat{\kappa}_z$ corresponds to a more fat-tailed (concentrated) sales distribution. The display above reports the mean and the interquartile range of $\hat{\kappa}_z$, estimated in the data and in the model. On average across sectors, the distribution of domestic sales exhibits Zipf's law, i.e. the estimated Pareto shape parameter is equal to 1.015, close to 1. However, there is a substantial variation across sectors, with the interquartile range given by $[0.82, 1.21]$. The model somewhat overstates the mean of $\hat{\kappa}_z$, at 1.096, and understates the variation across sectors, with the interquartile range $[1.01, 1.17]$.[27] With a less fat-tailed sales distribution compared to the data, the model therefore offers a conservative bound for the role of granularity, as we explore in the following section.

Our second overidentification check considers the *joint* distribution of French-firm sales and the number of French firms across sectors. In particular, we regress the relative size of the largest French firm $\tilde{s}_{z,1}$ on the log domestic absorption $\tilde{Y}_z = \tilde{\alpha}_z \tilde{Y}$ and the log-number of French firms $\tilde{M}_z$ in the sector:

$$\tilde{s}_{z,1} = const + \gamma_M \cdot \log \tilde{M}_z + \gamma_Y \cdot \log \tilde{Y}_z + \epsilon_z^s$$

$$\text{Data:} \quad \underset{(0.008)}{-0.094} \qquad \underset{(0.008)}{0.018}$$

$$\text{Model:} \quad \underset{(0.007)}{-0.064} \qquad \underset{(0.006)}{0.025}$$

The display above reports the OLS-estimated semi-elasticities and their standard errors. In sectors with more French firms, the relative size of the largest French firm is smaller, however this relationship is not very steep. Conditional on the number of firms, the size of the sector (measured by domestic absorption) has a small positive correlation with the relative size of the largest firm. The model is successful at capturing these empirical patterns. Granular forces are likely to be shaped by both the number of firms $\tilde{M}_z$ and the relative size of the largest

---

[26]The results are similar for the sample of top 50% of firms.

[27]Recall that in the model the average shape parameter is closely related to $\kappa = \frac{\theta}{\sigma-1} = 1.077$, and is slightly higher (less fat-tailed) due to variable markups. Indeed, the markups are higher for larger firms, and hence the sales distribution is less concentrated than would be the case under constant markups.

firm $\tilde{s}_{z,1}$, and thus it is important that our model replicates not only their partial distributions (in Appendix Figure A4), but also their joint distribution across sectors.

From this analysis, we conclude that the model is capable of capturing the salient features of the cross-sector variation in the number of firms, top-firm market shares, trade shares and measures of concentration. The model is also successful at capturing the patterns of joint variation of some of these variables. This is, perhaps, surprising given the tightness of the model parameterization, which features only five parameters common across sectors and countries. The granularity mechanism results in rich variations across sectors that mimic the patterns observed in the data.

Equilibrium markups   We close by briefly commenting on the equilibrium markup variation across firms displayed in the estimated model. The oligopolistic competition in our granular model results in heterogeneous markups, with larger firms setting higher markups, as given by (7). However, under Bertrand competition, the equilibrium variation in markups is quite limited, as we illustrate in Appendix Figure A2. Indeed, only the largest firm in a sector charges a markup considerably above 1.25, which would be the value of the constant markup in a counterfactual continuous model with monopolistic competition ($\frac{\sigma}{\sigma-1} = 1.25$). The markup of the largest firm is 1.30 on average across sectors, and it is as high as 1.37 at the 90th percentile across sectors. In contrast, the third largest firm in a sector charges a markup just under 1.26 on average across sectors and with little cross-sectoral variation. This is almost indistinguishable from the monopolistic-competition markup. Therefore, the abstraction with constant markups used in much of the granularity literature is reasonable, except for the largest one or two firms in an industry. These firms, however, can be pivotal for sectoral outcomes, as well as for trade and industrial policies that we study below.

4   Quantifying Granular Trade

Armed with the estimated model, we now study the extent to which granularity shapes trade patterns. Recall from equation (12)–(14) that sectoral trade flows X_z are determined by three factors: (i) sectoral expenditure shares α_z, (ii) fundamental comparative advantage Φ_z, and (iii) granular comparative advantage, driven by outstanding firms and summarized by the granular residual Γ_z. That is, total sectoral exports can be expressed as follows:

$$X_z = \alpha_z \Lambda_z^* Y^* \qquad \text{and} \qquad \Lambda_z^* = \Phi_z^* + \Gamma_z^*.$$

Table 3 reports the decomposition of trade flows into the above three sources of variation, in the estimated model (column 1). The other columns of the table report robustness results, which we discuss below.

Table 3: Variance decomposition of trade flows

		Common θ		Sector-specific θ_z		
		(1)	(2)	(3)	(4)	(5)
Granular contribution	$\frac{\text{var}(\Gamma_z^*)}{\text{var}(\Lambda_z^*)}$	17.0%	22.3%	26.0%	28.4%	20.3%
Export share contribution	$\frac{\text{var}(\log \Lambda_z^*)}{\text{var}(\log X_z)}$	57.2%	59.2%	62.5%	63.9%	59.0%
Pareto shape parameter	$\kappa_z = \frac{\theta_z}{\sigma-1}$	1.08	1.00	**1.02**	0.96	1.15
Estimated Pareto shape	$\hat{\kappa}_z$	1.10	**1.02**	1.07	**1.02**	1.21
Top-firm sales share	$\tilde{s}_{z,1}$	**0.21**	0.25	0.26	0.29	**0.21**

Note: The lower panel reports averages across simulated sectors; $\hat{\kappa}_z$ is estimated as in (19); bold emphasizes the targeted moment in each specification. Specifications: (1) Baseline estimated model. (2) Counterfactual with $\sigma = 5.34$ (instead of $\sigma = 5$) to match the average estimated Pareto shape ($\hat{\kappa}_z$) equal to 1.02. In (3)–(5), $\sigma = 5$ and θ_z are sector-specific: (3) θ_z set such that $\theta_z/(\sigma - 1) = \hat{\kappa}_z$ in the data (the estimated Pareto shapes sector-by-sector). (4) like (3), but proportionally scaling θ_z *down* to match the average of $\hat{\kappa}_z$ (since $\theta_z/(\sigma - 1) > \hat{\kappa}_z$ due to variable markups). (5) like (3), but proportionally scaling θ_z *up* to match the average top-firm sales share $\tilde{s}_{z,1}$ (relative to other home firms in the home market; moment 3 in Table 2). See Appendix Figure A5 for the fit of the model across these specifications.

We first report the contribution of the granular residual Γ_z^* to the variation in export shares Λ_z^* across sectors, using the following variance decomposition:

$$\text{var}(\Lambda_z^*) = \text{var}(\Phi_z^*) + \text{var}(\Gamma_z^*). \tag{20}$$

By construction, Γ_z^* is a mean-zero granular residual, which is uncorrelated with the fundamental comparative advantage Φ_z^*, and hence this decomposition holds exactly without a covariance term. In our estimated model, we find that granularity shapes 17% of the variation in export shares across sectors, while the rest corresponds to fundamental comparative advantage. In turn, export shares Λ_z^* account for 57% of the variation in overall trade flows X_z, while the rest is accounted for by sectoral expenditure shares α_z.[28]

By construction, the granular contribution to trade flows is zero on average across sectors, and granularity does not create additional trade at the aggregate level. Instead, granularity creates additional trade flows in the granular sectors, which is compensated by missing trade in the non-granular sectors, as we investigate next. Here and below, we refer to a sector as *granular* if $\Gamma_z^* \gg 0$, while if $\Gamma_z^* < 0$ or $\Gamma_z^* \approx 0$ we label such sectors non-granular, even though *ex ante* all sectors are symmetric in terms of their expected granularity, as $\mathbb{E}\Gamma_z^* = 0$ for every z.

---

[28]We measure the contribution of export shares to the overall sectoral exports as $\text{var}(\log \Lambda_z^*)/\text{var}(\log X_z)$. The exact variance decomposition of X_z is $\text{var}(\log X_z) = \text{var}(\log \alpha_z) + \text{var}(\log \Lambda_z^*) + 2\text{cov}(\log \alpha_z, \log \Lambda_z^*)$. In the estimated model, $\text{cov}(\log \alpha_z, \log \Lambda_z^*) \approx 0$, and therefore there is no need to allocate the covariance term.
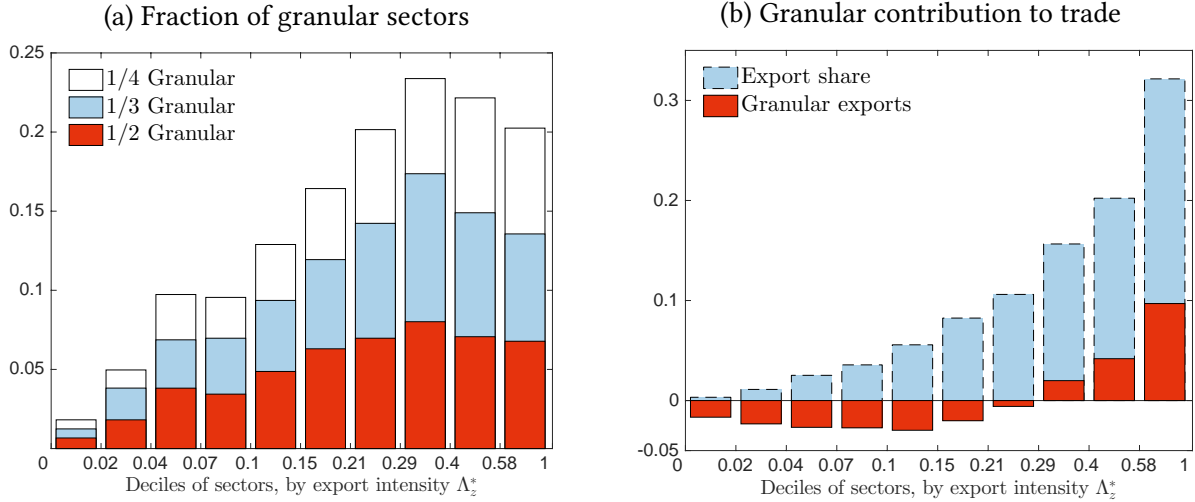
Figure 1: Export intensity and granularity

Note: All sectors are split into 10 deciles (bins of equal size in terms of the number of sectors) based on their export share, $\Lambda_z^* = X_z/(\alpha_z Y^*)$. The left panel plots for each decile the fraction of sectors for which $\Gamma_z^*/\Lambda_z^* \geq \vartheta \in \{1/2, 1/3, 1/4\}$. For example, the cumulative height of the blue and red bars corresponds to the fraction of sectors for which $\Gamma_z^* \geq \frac{1}{3}\Lambda_z^*$, or equivalently $\Gamma_z^* \geq \frac{1}{2}\Phi_z^*$. The right panel plots the contribution of deciles to aggregate trade (dashed blue bars) and the contribution of granular trade ($\Gamma_z^* \alpha_z Y^*$) to aggregate exports $X = \Phi^* Y^*$ (solid red bars), by deciles of sectors. The figures are based on a simulation of the estimated model with 10,000 sectors.

Figure 1 illustrates that the effects of granularity are particularly pronounced in the most export-intensive sectors, i.e. in the *export champions* of the country. This can be seen in two ways. Panel (a) illustrates that the likelihood of a sector being granular tends to increase with the export intensity of the sector Λ_z^*.[29] Panel (b) plots the corresponding export flows. As in panel (a), sectors are grouped by deciles of export intensity. The figure shows the contribution of each decile to the country's total exports (dashed blue bars), and highlights with red solid bars the contribution of granular trade to the total sectoral exports.[30] Note that the cumulative height of all blue bars is 1 (aggregate exports), while the cumulative height of all red bars is zero, as granularity does not change the aggregate amount of trade. The top three deciles of export-intensive sectors account for two thirds of the aggregate exports. These are exactly the sectors where the granular contribution to trade is positive on net, and accounts for a substantial fraction of trade flows. In all other bins of less-export-intensive sectors, the contribution of granular trade is negative, that is, these sectors would export slightly more in the continuous limit of the model.

---

[29]Specifically, the figure groups sectors into deciles of export intensity Λ_z^*, and reports for each decile the share of sectors where granular exports are important, representing more than 1/4 (1/3, or 1/2, respectively) of the total sectoral exports.

[30]Specifically, the figure plots $\frac{1}{X}\int_{z\in D_k}\alpha_z Y^*\Lambda_z^* dz$ and $\frac{1}{X}\int_{z\in D_k}\alpha_z Y^*\Gamma_z^* dz$, where $X = Y^*\int\alpha_z\Lambda_z^* dz$ are aggregate exports, and D_k is the kth decile of sectors sorted by export intensity Λ_z^*.

Overall, granularity shapes trade flows, and does so in a concentrated way among the most export-intensive sectors. An outstanding productivity draw in a sector (i.e., a very large firm) tends to have a major positive impact for production and exports in this sector, while the absence of such a draw in a sector (i.e., no outsized firm) tends to only have a moderate negative impact. This is balanced out by the fact that the presence of an outstanding draw is a rare outcome. Taken together, these forces add skewness to the distribution of export intensity across sectors in a granular economy.

Robustness Our baseline model features only two sources of ex-ante heterogeneity across sectors: the Cobb-Douglas expenditure shares α_z and fundamental productivities T_z/T_z^*, whereas in reality, sectors are likely to be heterogeneous in a number of different ways. One may thus worry that our results are sensitive to this simplifying assumption, and that we overstate the role of granularity by shutting down such heterogeneity. In particular, variation in firm size distribution is likely to be in part due to these other sources of heterogeneity, rather than driven by granularity alone.

To address this issue, we recalibrate the model by allowing for sector-specific θ_z, i.e. the parameters that govern the dispersion in firm productivity draws within sectors. In a continuous model, variation in this parameter is a natural way to obtain variation in firm-size distribution across sectors (see e.g. di Giovanni and Levchenko 2012, 2013). We discipline the distribution of θ_z across sectors in three alternative ways, with results reported in the last three columns of Table 3. First, we choose θ_z so that $\frac{\theta_z}{\sigma-1} = \hat{\kappa}_z$ for all sectors, where $\{\hat{\kappa}_z\}$ are the empirical estimates of the Pareto shapes of the firm size distribution in the data (from (19)). With a continuum of firms and constant markups, $\frac{\theta_z}{\sigma-1}$ exactly corresponds to the Pareto shape of the sales distribution, but a granular model with variable markups introduces a wedge between the two. Therefore, our second calibration proportionally scales down the distribution of θ_z to ensure that the mean value of the estimated Pareto shape parameters in the model, $\hat{\kappa}_z$, matches the one in the data. Third, since both of these calibrations overstate somewhat the average sales share of the largest firm relative to the data (and relative to the baseline model, which matches this moment), we scale up the distribution of θ_z to match the top sales share moment. The bottom panel of Table 3 and the Appendix Figure A5 illustrate the fit of different moments across these three specifications. In particular, the calibrated model can now accurately match the distribution of the estimated Pareto shape coefficients $\hat{\kappa}_z$ in the data.

Interestingly, Table 3 shows that the contribution of granularity *increases*, across all three specifications, compared to our baseline calibration with homogeneous θ across sectors. The contribution of granularity now ranges from 20% to 28%. Intuitively, the strength of granularity is largely determined by the market share of the largest firm in the sector (see Table A1 and

23

the discussion below). Having heterogeneous θ_z does not change the ability of the model to match the relative size of the largest firms. However, with heterogeneous θ_z, some sectors end up having smaller θ_z, and as a result even fatter-tailed sales distributions and larger top firms. This additional skewness increases the role of granularity in shaping the economic outcomes. Overall, this robustness exercise suggests that our baseline estimate of the role of granularity is, if anything, conservative.

As another robustness check, the second column of Table 3 reports the sensitivity of our results to the elasticity of substitution of demand, σ. We do this for two reasons. First, as we noted above, our estimation procedure is conservative in that we target the market share of the top firms, but understate the fatness of the tail of the sales distribution, as measured by the Pareto shape $\hat{\kappa}_z$. We report here what would be the outcome of a less conservative estimation procedure, which would target instead the measured Pareto shapes of the firm-size distribution (i.e., Zipf's law). Second, we note that the literature has been documenting an increase in concentration within industries (see e.g. Autor, Dorn, Katz, Patterson, and Van Reenen 2017, Gutiérrez and Philippon 2017). A common hypothesis is that it corresponds to an increased substitutability across products σ, for example due to the increased online competition.[31] Here we are interested in understanding the possible consequences of this increase for the role of granularity in shaping trade flows. We therefore consider a counterfactual with a larger elasticity of substitution $\sigma = 5.34$ (versus $\sigma = 5$ in the baseline). It allows the model to match exactly the average estimated Pareto shape parameter in the data (equal to 1.02; see (19)). We find that the contribution of granularity to sectoral export shares increases from 17% to 22%. Quite intuitively, the role of granularity increases as the economy becomes more concentrated.

Inference on sectoral comparative advantage   Having established the quantitative relevance of granularity in shaping the sectoral trade flows, we now ask the following question: Given the observed export stance of a sector, what inference can one make on its fundamental comparative advantage? In a conventional continuous model, there is a one-to-one mapping from the observed trade flows into the fundamental comparative advantage T_z/T_z^*, as $\Lambda_z^* = \Phi_z^*$ in this case, a feature that is used extensively in the quantitative trade literature following Eaton and Kortum (2002). The presence of granularity complicates this inference, as export shares Λ_z^* now reflect both fundamental and granular sources of comparative advantage. The ability to draw inference on this split is important if fundamental and granular comparative advantage have different implications, for example, for the dynamics of trade flows, as we explore in Section 5.

---

[31]A natural microfoundation for this mechanism is a frictional discrete choice model with decreasing search costs over time (see e.g. Hortacsu and Syverson 2014).

(a) Distribution of Λ_z^* conditional on Φ_z^* (b) Distribution of Φ_z^* conditional on Λ_z^*
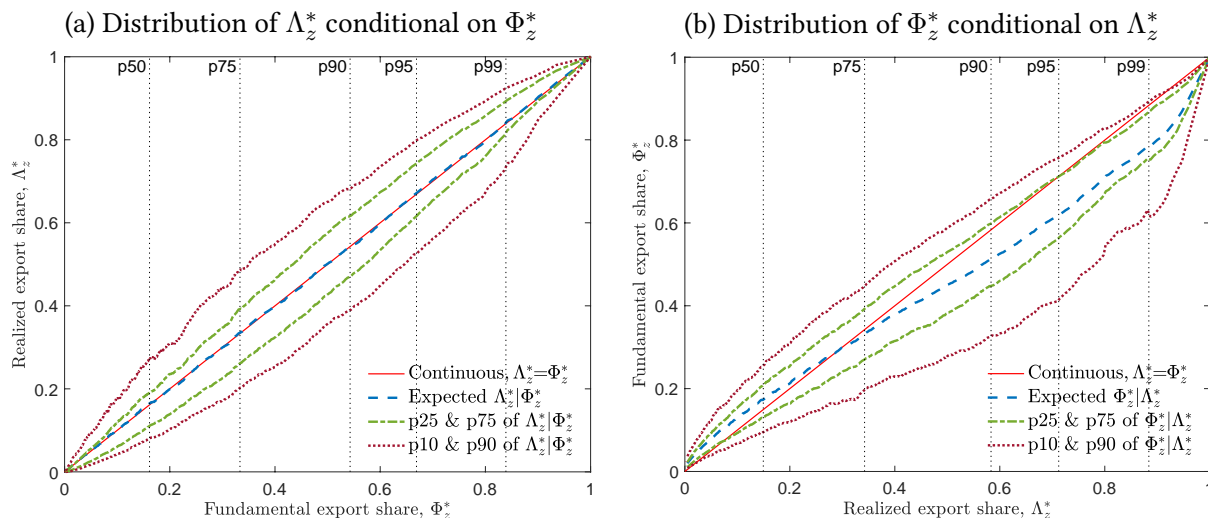
Figure 2: Comparative advantage and trade flows: distribution across realizations

Note: The figures plot moments and percentiles of the conditional distributions: $\Lambda_z^*|\Phi_z^*$ in the left panel and $\Phi_z^*|\Lambda_z^*$ in the right panel. In both figures, the red solid $45°$-line corresponds to $\Lambda_z^* = \Phi_z^*$, towards which the distributions degenerate in a continuous model. The vertical dotted lines plot the percentiles of the unconditional partial distribution of Φ_z^* in the left panel and of Λ_z^* in the right panel. The figures are based on a simulation of the estimated model with 100,000 sectors.

As a first step, we use the estimated model to plot, in the left panel of Figure 4, the distribution of realized export intensity Λ_z^* conditional on the fundamental comparative advantage of a sector Φ_z^*. The one-to-one deterministic mapping between the two in the continuous model is depicted with a red $45°$-line. In the granular model, export shares conditional on the fundamental forces are now random, reflecting the granular draws. Their conditional mean is depicted with a dashed blue line, which coincides with the red line. There is substantial variation in actual realizations, which is seen from the dotted lines that correspond to the percentiles of the conditional distribution of $\Lambda_z^*|\Phi_z^*$. The vertical departures from the $45°$-line correspond to the realizations of the sectoral granular residuals, $\Gamma_z^* = \Lambda_z^* - \Phi_z^*$.[32] This figure complements the decomposition in Table 3 in illustrating the contribution of granularity to sectoral trade shares.

The right panel of Figure 4 describes instead the conditional distribution of $\Phi_z^*|\Lambda_z^*$, that is, the inference one can make on the fundamental Φ_z^* conditional on observing a realized export share Λ_z^*. To that end, the right panel switches the axes of the left panel. The continuous model is again represented by the solid red diagonal line. In the granular model, inference is very different. The conditional expectation of Φ_z^* given the observed Λ_z^* is depicted with a blue dashed line, which unlike in the left panel now departs from the red diagonal. In other

---

[32]For example, at the 75th percentile of $\Phi_z^* = 0.33$, the interquartile range of $\Lambda_z^*|\Phi_z^*$ is $[0.27, 0.40]$, and its 90th percentile is 0.49, corresponding to almost the 90th percentile of the unconditional distribution of Λ_z^*.

words, the sectoral $\Phi_z^*|\Lambda_z^*$ is not symmetric or even centered around Λ_z^*, as was the case for $\Lambda_z^*|\Phi_z^*$ in the left panel. This reflects the pattern we already observed in Figure 1, namely that sectors with small realized export shares tend to have negative granular residuals and sectors with large realized export shares tend have positive granular residuals. Therefore, sectors with the largest realized export shares have systematically lower expected export shares, $\Phi_z^* < \Lambda_z^*$, i.e. a lower fundamental comparative advantage than a continuous model would predict.[33] In other words, using a continuous model to estimate fundamental sectoral productivities in a granular world would lead to a systematic positive bias for high export-intensity sectors.

Inference on sectoral granularity   We next use the estimated model to study the covariation properties of the granular trade share Γ_z^*, which is unobservable in the data, with variables that are, in contrast, observable in the data. The goal of this analysis is to establish effective measurable proxies for granularity. This acts as a check of our identification strategy in Section 3 and provides input for further empirical investigation of granular trade flows. In particular, we use these insights below to estimate the likelihood that sectoral exports have a large granular component, for each French sector in our dataset.

By construction, the granular share Γ_z^* is uncorrelated with sectoral fundamental productivity advantage captured by Φ_z^*. We find, furthermore, that Γ_z^* is largely uncorrelated with a number of sector-level variables, including the size of the sector, measured by total expenditure $\log(\alpha_z Y)$, and the total number of French firms serving the sector, $\log \tilde{M}_z$. This emphasizes that granularity is not merely a reflection of a small number of firms. In contrast, the variable that is most closely associated with Γ_z^* is the relative size of the largest French firm, as captured by its domestic sales share relative to other French firms in the sector, $\tilde{s}_{z,1}$.[34] A sector in which the largest French firm has a 10 percentage points higher sales share relative to other French firms has on average a 3.5 percentage points higher export share to the foreign market, entirely of granular origin. Note that $\tilde{s}_{z,1}$ does not contain any direct information on the firm's advantage over its foreign rivals, yet it signals the likelihood that the best domestic firm is a granular outlier, and hence predicts exports.

To conclude, simple statistics characterizing within-sector firm-size distribution, such as $\tilde{s}_{z,1}$, offer powerful proxies for the overall role of granularity in a sector and the aggregate sectoral outcomes, such as sectoral exports. This equilibrium statistical association was at the core of

---

[33]This corresponds to a classical selection (or reversion-to-the-mean) effect: a sector-outlier is only in part shaped by fundamental forces, and the less so the more of an outlier it is. From the right panel of Figure 4, note that over 70% of sectors (with smallest Λ_z^*) have $\mathbb{E}\{\Phi_z^*|\Lambda_z^*\} > \Lambda_z^*$, and it is only the most export-intensive sectors that share the reverse feature (indeed, unconditionally, $\mathbb{E}\Phi_z^* = \mathbb{E}\Lambda_z^*$).

[34]This variable alone can explain about 35% of variation in Γ_z^*, and when it is combined with a foreign counterpart $\tilde{s}_{z,1}^*$, they together explain over a half of granular trade. Additional variables, including other measures of concentration and skewness in the domestic sales shares, improve the fit only marginally. We summarize these results in Appendix Table A1.
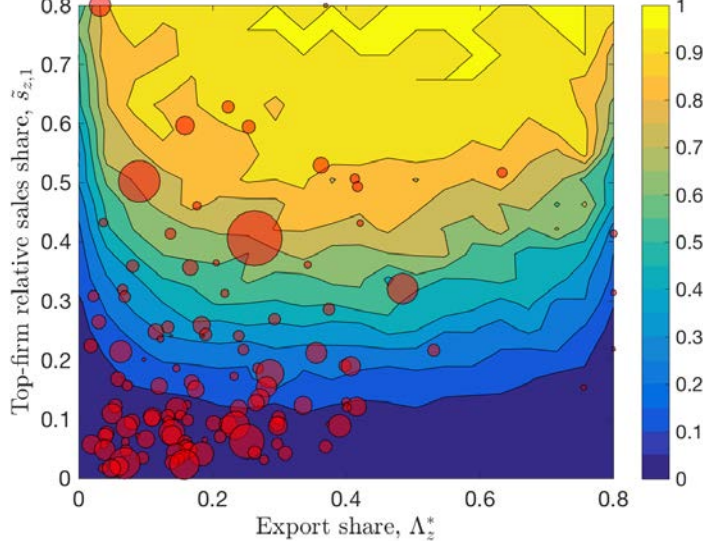
Figure 3: Likelihood of granularity, $\mathbb{P}\left\{\Gamma_z^* > \frac{1}{4}\Lambda_z^* \,\middle|\, \Lambda_z^*, \tilde{s}_{z,1}\right\}$

Note: Colors (see scale on the right) reflect the probability that at least 25% of the export share Λ_z^* in a sector with an observed outcome $(\Lambda_z^*, \tilde{s}_{z,1})$ is of granular origin, i.e. $\Gamma_z^* > \vartheta\Lambda_z^*$ for $\vartheta = 1/4$ (*cf* Figure 1). Red circles depict the $\tilde{N} = 119$ sectors in our French dataset, with their sizes proportional to the sectoral expenditure shares α_z. See Appendix C for the details of estimation.

our empirical identification strategy in Section 3.2, as captured by the last four moments in Table 2 (see also Navaretti, Bugamelli, Forlani, and Ottaviano 2016, who explore the predictive ability of skewness in the firm-sales distribution for sectoral trade flows).

Lastly, we use the insights from this analysis, and briefly address the possibility of identifying specific granular sectors in the data. This is a challenging task, as neither Φ_z^* nor Γ_z^* are observable, and the inference must be made using only the overall trade share Λ_z^* and the observable proxies for granularity, such as the relative size of the largest firm $\tilde{s}_{z,1}$. Nonetheless, one can make a Bayesian inference of the probability that exports in a given sector z are of a significant granular origin, e.g. $\Gamma_z^* \geq \vartheta\Lambda_z^*$ for some cutoff $\vartheta \in (0,1)$, as we defined in Figure 1. We plot the estimated granular probabilities for the French sectors in Figure 3 and detail the estimation procedure in Appendix C.

5   Dynamics of Comparative Advantage

Having established the implications of granularity for cross-sectional trade patterns, we now extend our granular model to allow for industry dynamics driven by idiosyncratic firm productivity shocks, and study its implications for the evolution of a country's comparative advantage over time. This line of analysis brings us closer to the macro granularity literature that, following Gabaix (2011), focuses on aggregate fluctuations.

5.1 Dynamic model

We introduce dynamics by assuming that the productivity of each firm evolves over time according to a random growth process. As a consequence, individual firm market shares change over time, as well as the granular comparative advantage and the resulting sectoral trade patterns. Note that in a continuous model, in contrast, such firm-level idiosyncratic volatility would have no impact on sectoral trade shares.

We assume that the productivity $\varphi_{z,i,t}$ of firm i in sector z at period t evolves according to a geometric random walk with downward drift $\mu < 0$ and a reflecting barrier at the lower bound $\underline{\varphi}_z$. Specifically, if the resulting $\varphi_{z,i,t} \geq \underline{\varphi}_z$, then

$$\log \varphi_{z,i,t} = \mu + \log \varphi_{z,i,t-1} + \nu \varepsilon_{z,i,t}, \qquad \varepsilon_{z,i,t} \sim iid\,\mathcal{N}(0,1), \qquad (21)$$

otherwise, $\log \varphi_{z,i,t}$ is reflected symmetrically away from the lower bound $\underline{\varphi}_z$.[35] The negative drift term $\mu = -\theta\nu^2/2$ ensures stationarity of the cross-sectional productivity distribution, which is Pareto with shape parameter θ (see Gabaix 2009), like in the static model. This productivity process applies to all shadow firms in the economy. There is no entry or exit: the pool of shadow firms remains constant over time. Shadow firms decide each period whether to pay a per-period fixed cost and be active, or be inactive. Since firms do not incur sunk costs of entry, their choice of operating in each period is a static one. As a result, the equilibrium of this dynamic model is a sequence of static equilibria: each period, firms play the static entry and price-setting game described in Section 2, given the current productivity distribution. This allows us to embed, in a tractable way, a persistent dynamic process for firm-level productivity that, in particular, affects the evolution of sectoral trade in the granular economy over time.

Quantitatively, we discipline the idiosyncratic productivity process using the evolution of firm market shares in the French data over 11 years, 1997 to 2007. The only new parameter that needs to be calibrated is the standard deviation of productivity innovations ν. Importantly, this parameter does not affect any of the static moments used in estimation. We choose to target two moments of the domestic sales shares $\tilde{s}_{z,i,t}$ dynamics: namely, the (inverse) short-run persistence measured by $\mathrm{std}(\Delta\tilde{s}_{z,i,t+1})$ and the long-run persistence measured by $\mathrm{corr}(\tilde{s}_{z,i,t+10}, \tilde{s}_{z,i,t})$, where we average across sectors z (and firms i, for the former moment), and one period corresponds to a year. Targeting these moments leads to $\nu = 0.05$.[36]

---

[35]Formally, if $\mu + \log\varphi_{z,i,t-1} + \nu\varepsilon_{z,i,t} < \log\underline{\varphi}_z$, then $\log\varphi_{z,i,t} = \log\underline{\varphi}_z + \left[\log\underline{\varphi}_z - (\mu + \log\varphi_{z,i,t-1} + \nu\varepsilon_{z,i,t})\right]$. Recall that $\underline{\varphi}_z$ is also the lower bound for the initial productivity draws, which are obtained according to the same Poisson-Pareto process as in the static model. $\underline{\varphi}_z$ is normalized to be a small number that never binds in the entry game; this parameter is of no consequence otherwise.

[36]With $\nu = 0.05$, the model slightly overstates the short-run persistence and slightly understates the long-run persistence (see Table 4). Setting $\nu = 0.053$ allows to match $\mathrm{std}(\Delta\tilde{s}_{z,i,t+1}) = 0.0018$, while setting $\nu = 0.045$ allows to match $\mathrm{corr}(\tilde{s}_{z,i,t+10}, \tilde{s}_{z,i,t}) = 0.86$. The quantitative results change little with this variation: namely, the 20-year turnover moment in Table 4 changes from 72% to 70% as we increase ν from 0.045 to 0.053.

Note that this calibration of dynamics does not rely on any trade-related moments.

We use the calibrated dynamic model to simulate the dynamic equilibrium path of the economy over an extended period of time, in order to explore the implications of granularity for the dynamics of comparative advantage.

5.2 Granular dynamics of comparative advantage

In a recent paper, Hanson, Lind, and Muendler (2016; henceforth, HLM) study the dynamic evolution of the comparative advantage of countries. We take their study as a benchmark. They emphasize two striking patterns:

  (i) *hyper-specialization of exports*: a single sector accounts for 21% of a country's total exports on average across countries, while the three top sectors account for over 40%. For France, these two concentration ratios are somewhat lower, equal to 17% and 30% respectively. Our estimated granular model reproduces this, as we report in Table 4.[37]

 (ii) *high turnover of comparative advantage*: a sector in the top-5 of sectors by export intensity has about a 50-50 chance of staying in the top-5 two decades later.

The combination of these two facts is indeed intriguing: countries appear to exhibit extreme specialization, yet their comparative advantage tends to change significantly in the medium run. We argue here that granular effects can rationalize these patterns. To show this, we quantify how much turnover in the comparative advantage of sectors is generated by our quantified granular model with firm dynamics.

We find that firm dynamics in the granular model generate mean reversion in sectoral export patterns. A sector in the top 5% in terms of export intensity only has a 71% chance to remain there two decades later. This is exclusively due to idiosyncratic shocks to firm productivity: there is no change over time in the fundamental comparative advantage of sectors. Therefore, quantitatively, this granular force alone can explain over 60% of the sectoral turnover documented by HLM. Furthermore, it can explain almost all the turnover we observe in the French data, where a sector among the top-5 in terms of export intensity in the beginning of our sample has an 80% chance to remain in the top 5 by the end of the sample, 10 years later. In comparison, our simulated model predicts a 79% chance over a 10-year horizon.[38]

---

[37]HLM split all products into 113 sectors, so we interpret top-1, top-3 and top-5 respectively as top-1%, top-3% and top-5% of sectors in the context of our model with a large number (continuum) of sectors. HLM show that small developing countries exhibit more extreme patterns of both specialization and mean reversion, and anecdotally FDI likely plays an important role in this (e.g., Intel operations in Costa Rica). Since we focus on France, a large developed country, we report the moments from HLM which correspond to the subsample of developed countries. Table 4 also reports the corresponding moments in our French data with 119 manufacturing sectors.

[38]Appendix Figure A6 plots the measure of turnover for different horizons $T \geq 0$ in the simulated model. The probability of staying in the top 5% of export-intensive sectors decreases to 65% after 50 years, and it is just above 50% in the very long-run, reflecting the role of fundamental forces at keeping some sectors at the top of the export intensity distribution.

Table 4: Firm dynamics and comparative advantage

Moment	Data		Model
	HLM	France	
SR persistence $\mathrm{std}(\Delta \tilde{s}_{z,i,t+1})$	—	0.0018	0.0017
LR persistence $\mathrm{corr}(\tilde{s}_{z,i,t+10}, \tilde{s}_{z,i,t})$	—	0.86	0.83
Share in aggregate exports:			
— top-1% of sectors	21%	17%	18%
— top-3% of sectors	43%	30%	33%
Turnover of comparative advantage:			
— remain in top-5% after 20 years	52%	—	71%
— remain in top-5% after 10 years	—	80%	79%

Note: Empirical moments from HLM for developed countries and the French data, where available. Top-1% (3% and 5%) of sectors corresponds to 1 (3 and 5) sectors in the data (total number of sectors in HLM is 113 and 119 in the French dataset) and 1% (3% and 5%) of sectors in the model-simulated data (with 25,000 sectors). The export share moments are based on sectoral exports $\Lambda_z^* \alpha_z Y^*$; the turnover moments are based on export shares Λ_z^*.

Recall that we do not target any trade-related moments in the calibration of dynamics, as we only use the persistence of sales shares of the French firms, among other French firms, in the domestic market. These results suggest that idiosyncratic firm dynamics is important quantitatively in shaping the observed turnover among the most export-intensive sectors.

More generally, what are the implications of granularity for the dynamics of comparative advantage across the full range of sectors? To illustrate the role of granularity, we rank sectors by their granular trade shares Γ_z^* in the initial period of the simulation, and group them in 10 bins of equal size (deciles). The left panel of Figure 4 plots the average changes in export shares $\Delta \Lambda_z^*$ by decile, over 20 and 50 years. The strongest mean-reversion forces are at play in the most granular sectors, which tend to lose export shares over time. The non-granular sectors (with $\Gamma_z^* < 0$ or $\Gamma_z^* \approx 0$) are expected to gain export shares over time, but at a much slower pace, emphasizing again the skewness inherent in granular outcomes (compare with Figure 1).

The right panel of Figure 4 shows that, in addition to mean reversion, the most granular sectors also exhibit very volatile export patterns over time. This figure plots the standard deviation of annual changes in export shares $\Delta \Lambda_z^*$, again by deciles of sectors. By this measure, the most granular sectors are over five times more volatile than the least granular sectors, which resonates with the results on aggregate fluctuations in the macro-granularity literature. This volatility is an order of magnitude larger than the annualized rates of mean reversion in comparative advantage. That is, while granular sectors tend to mean revert on average and over longer horizons (panel a), they can also gain considerably in comparative advantage over a medium run (panel b). This volatility, intrinsic to granular outcomes, emphasizes the need

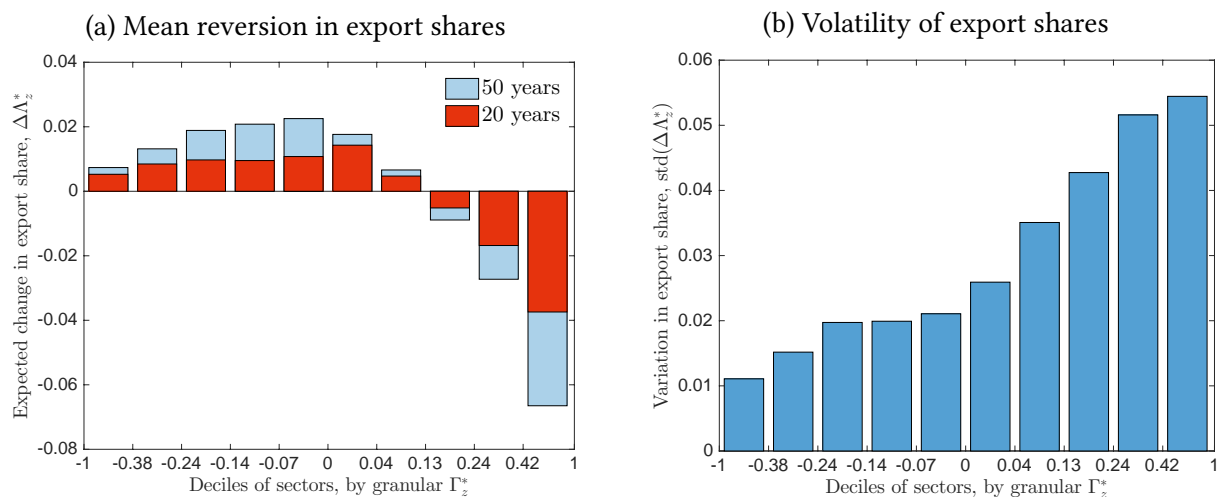| (a) Mean reversion in export shares | (b) Volatility of export shares |

Figure 4: Evolution of granular comparative advantage

Note: 25,000 simulated sectors are split into 10 deciles based on their granular export share, $\Gamma_z^* = \Lambda_z^* - \Phi_z^*$, in the initial period. The left panel plots the average (by bin) change in export shares, $\Delta\Lambda_z^*$, over the next 20 and 50 years respectively. The right panel plots the standard deviation of the time-series annual changes in $\Delta\Lambda_z^*$, averaged within deciles of sectors.

for a model to study the granular patterns, which would be very difficult to measure directly in the data with a sufficient degree of precision.

Granularity, openness and sectoral reallocation A natural corollary of what we have established so far is that, in a granular open economy, firm-level shocks generate production and labor reallocation, and do so not only between firms, but also across sectors. This allows the open economy granular model to speak to an important macroeconomic feature of the data, namely the equilibrium intensity of labor reallocation across industries. Note that, in contrast, the continuous counterpart of the model predicts no inter-sectoral labor reallocation triggered by firm-level industry dynamics.

In our simulated dynamic model with idiosyncratic firm productivity shocks, the annual job creation and job destruction rates are about 12%. Importantly, 19% of this turnover is due to inter-sectoral job reallocation, reflecting the shifts in production across sectors in response to granular changes in the country's comparative advantage. This extent of job turnover, both within and across sectors, is in line with the empirical patterns documented by Davis and Haltiwanger (1999; see their Tables 1, 2 and 5). Furthermore, the share of inter-sectoral labor reallocation in the overall job flows is very sensitive to the degree of openness of the economy, and it falls to just 3% in an economy virtually closed to international trade. We provide further details in Appendix Table A2, which also shows that similar patterns of inter-sectoral reallocation hold over longer horizons. These results emphasize the interaction between granularity
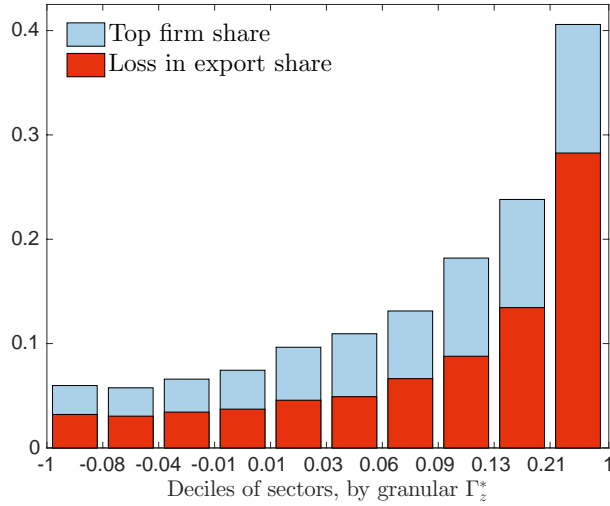
Figure 5: Trade effects of individual firm exit

Note: *Top firm share* (blue bars) is the foreign market share of the largest home exporter. *Loss in export share* (red bars) equals $-\Delta\Lambda_z^*$ associated with the exit of the largest exporter. The figure plots averages by deciles of sectors in terms of granular exports $\Gamma_z^* = \Lambda_z^* - \Phi_z^*$, for the subset of 30% most export intensive sectors in terms of Λ_z^*, which account for 67% of total exports. Appendix Figure A7 reproduces the same figure for all sectors.

and trade openness in contributing to the increased volatility of resource allocation in open economies (*cf* Rodrik 1998).

Individual firm exit   Another way to quantify the important role played by individual firms in shaping comparative advantage is to consider a counterfactual in which the largest firm in a sector exits the industry for an exogenous reason. We study the impact of this shock on the export performance of the sector, which would be nil in a continuous model. In our quantified model, the largest exporter accounts on average for over a quarter of total sectoral exports and over 6% of the foreign market. If this large exporter exits, its market share is redistributed towards existing home and foreign firms in the sector, as well as towards new entrants at home and abroad. The reallocation of this lost market share towards foreign firms reflects a loss in comparative advantage.

In Figure 5, we look at the comparative-advantage sectors (top 30% of sectors in terms of export shares Λ_z^* that account for over two-thirds of total exports) and group them into deciles by granular trade shares Γ_z^*. We plot the average foreign market share of the largest domestic exporter (blue bars) and the average loss in sectoral export share when this largest exporter fails (red bars). In non-granular sectors, the impact of a firm exit is limited. On average, about half of the exporter's sales are recaptured by domestic exporters, while the remaining half is reallocated towards foreign firms. For most sectors this leads only to a mild loss in export intensity and comparative advantage.

The situation is very different, however, for the most granular sectors. In the top decile, the largest exporter commands as much as a 40% foreign market share, and, importantly, over 70% of it is lost to foreign firms if this largest exporter fails. This results in a sharp loss of comparative advantage. That is, not only do granular firms matter because they hold a large share of exports, but also because loosing these firms would generate a disproportionate loss of sectoral export shares and comparative advantage. In certain cases, sectors may drop entirely from being among the most export-intensive sectors (top 5%) into becoming comparative disadvantaged (bottom 50%).

6 Granular Policies in an Open Economy

Having established the role of granularity in shaping static and dynamic trade patterns, we finally turn to exploring its consequences for policy. A range of policies specifically target large firms. An obvious example is antitrust policy that regulates mergers of firms with significant market power. Merger policy is often viewed as part of a toolkit that policymakers use to affect foreign market access (see e.g. Bagwell and Staiger 2004, Chapter 9). Further, countries may be interested in targeting large foreign firms directly, for example, as part of a trade war. What impact do these policies have on trade flows and welfare? This question cannot be analyzed using standard continuous models where, even in the presence of heterogeneity, every firm is infinitesimal. In contrast, our quantitative granular model is well-suited to analyze the consequences of and the economic motivation behind such policies. In what follows, we explore in turn two policies: a merger between two large firms in a given sector of an open granular economy and a granular import tariff imposed on a single large foreign exporter. Before presenting the results, we start by outlining the general methodology we follow to compute and decompose the welfare effects of policies.

Welfare decomposition   In our model, the welfare of a representative consumer at home is given by $\mathbb{W} = Y/P$, where Y is aggregate home income and $P = \exp\left\{\int_0^1 \alpha_z \log P_z \mathrm{d}z\right\}$ is the price index. In general, aggregate income can be decomposed as $Y = wL + \Pi + TR$, where wL is labor income, Π is aggregate profits defined in (10), and TR is government policy revenues distributed lump-sum to workers. Since labor is supplied inelastically and we adopt the home wage as numeraire, the log-change in home welfare in response to a policy can be expressed as follows:[39]

$$\hat{\mathbb{W}} \equiv \mathrm{d}\log\frac{Y}{P} = \frac{\mathrm{d}\Pi}{Y} + \frac{\mathrm{d}TR}{Y} - \int_0^1 \alpha_z \mathrm{d}\log P_z \mathrm{d}z, \tag{22}$$

---

[39]Note that the change in the real wage is fully accounted for by the changes in the price level P since nominal wage $w = 1$ by our choice of the numeraire; otherwise, there would be an additional term $\frac{wL}{Y}\mathrm{d}\log w$.

The three components in (22) correspond to the respective changes in the producer surplus, government revenues, and consumer surplus.

We are interested in the general equilibrium impact of policies targeted at large firms, and in particular, in contrasting the effects they have in granular versus non-granular sectors. To this end, we study a given policy change in a positive measure of sectors Z that have similar levels of granularity.[40] More concretely, we bin sectors into quintiles by granular export share Γ_z^* defined in (14). We then compute the corresponding welfare impact $\hat{\mathbb{W}}_Z = d \log(Y/P)$ of the policy in bin Z, and report its *average* aggregate welfare effect, normalized by the size of set Z, defined by:[41]

$$\hat{W}_Z = \frac{1}{\int_{z \in Z} \alpha_z dz} \hat{\mathbb{W}}_Z. \tag{23}$$

We also consider the decomposition of the average welfare effects $\hat{W}_Z$ into the contributions of the changes in the consumer and producer surplus, according to (22).

6.1 Welfare implications of a merger

We study here the welfare consequences of a merger between the two top domestic firms in a sector, with a binary policy option of whether or not to allow the merger. Mergers and acquisitions are sought by firms for a range of reasons, including increasing their market power, realizing cost synergies and transferring core competencies and best practices. We capture these channels in the following way. First, we assume that, upon merging, the single entity continues to produce the two distinct product lines previously produced by the two separate firms, but that it now sets markups to maximize the total profit of the merged entity. As a consequence, the new firm's market power and markups increase.[42] Second, we assume that the merged firm pays only one fixed cost rather than two. We note however that this assumption is largely inconsequential quantitatively, as fixed costs are a very small fraction of revenues for the largest firms. Finally, we allow the merger to generate productivity spillovers between the merged entities. Specifically, we allow for the less-productive product line to inherit some of the efficiency of the most productive one, with the strength of the spillover

---

[40]Note that in our model with a continuum of sectors, we need to consider a positive measure of sectors with a policy change, since changes in a single sector have no aggregate effects. The direct effect of a policy change in sector z can be captured by $(d\Pi_z + dTR_z)/Y - \alpha_z d \log P_z$, but the indirect general equilibrium effects have impact on all other sectors $z' \in [0, 1]$ and need to be taken into account.

[41]With this definition, in the limit as sets Z become tight around individual sectors z, the aggregate welfare change $\hat{\mathbb{W}}$ from a sectoral policy vector $\{\varsigma_z\}_{z \in [0,1]}$ can be decomposed as a cross-sectoral weighted-average of the welfare effects $\hat{W}_z$ of the sectoral policies ς_z, namely $\hat{\mathbb{W}} = \int_0^1 \alpha_z \hat{W}_z dz$.

[42]Given CES demand, the optimal markups are the same for both products and depend on their cumulative market share $s'_{z,1} + s'_{z,2}$ in the new equilibrium, according to the same functional relationship as in (7).

governed by the parameter $\varrho \in [0, 1]$:

$$\varphi'_{z,2} = \varrho\varphi_{z,1} + (1 - \varrho)\varphi_{z,2},$$

where $\varphi_{z,i}$ is the productivity of the pre-merger firm i and $\varphi'_{z,i}$ is the productivity of the corresponding post-merger product line. Given this post-merger market structure and productivity distribution, we solve for the new entry game and price-setting equilibrium in each sector. To get at the full welfare effect of a merger, we simulate it for a subset of sectors $z \in Z$, and recompute the new general equilibrium. This allows us to quantify and decompose the average general equilibrium welfare effect of a merger, $\hat{W}_Z$, as described above.

To illustrate our findings, we report in Figure 6 the welfare effects of a merger in a baseline experiment in which the productivity spillover is set at a high value of $\varrho = 0.5$, and the Appendix Figure A8 provides the results for an alternative low level of productivity spillover $\varrho = 0.1$. The figure plots the welfare implication of a top merger, as a function of the level of granularity of the sector captured by the granular trade share Γ^*_z, and for different levels of trade openness captured by the variable trade cost τ. A stark picture emerges. First, mergers are much less welfare detrimental, and instead can increase domestic welfare, when the economy is more open to international trade, even if the productivity spillover ϱ is small. Second, when the productivity spillover is sufficiently strong or if the economy is sufficiently open, mergers are particularly beneficial in the most granular sectors.[43] This interaction between granularity and trade openness is driven by the fact that the merged entity enjoys increased market power and greater profits. The more open the economy, the more these profits are made at the expense of the foreign country, by extracting its consumer surplus and improving the terms of trade for the home country.

The right panel of Figure 6 dissects the overall welfare effect of a merger into the contributions from consumer and producer surplus, $-\mathrm{d}\log P$ and $\mathrm{d}\Pi/Y$. First, note that a merger, by increasing monopoly power, leads the price level to increase and hence the real wage to decline, reducing consumer surplus. Even a substantial productivity spillover ($\varrho = 0.5$) cannot make up for the deterioration in industry competitiveness coming from an increase in markups. This decline in consumer surplus is counterbalanced by an increase in producer surplus, i.e. by the profits of domestic firms, at home and abroad. In sufficiently open economies and granular sectors, this increase in producer surplus dominates the reduction in consumer surplus, resulting in a net increase in domestic welfare.[44]

---

[43]With a high level of productivity spillover ($\varrho = 0.5$), mergers are detrimental to welfare only in sectors with little granularity and when the economy is closed to international trade. With low productivity spillovers ($\varrho = 0.1$), mergers are only beneficial when the economy is very open to international trade and in the most granular sectors. Compare Figures 6 and A8.

[44]The loss of consumer surplus may be lower or higher in more granular sectors, depending on the strength
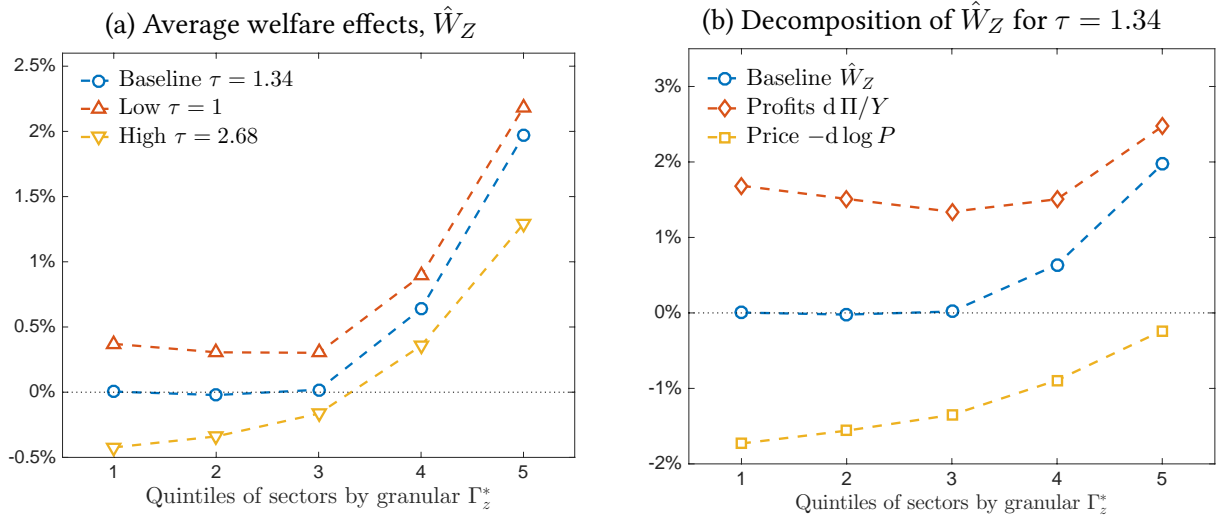
Figure 6: Welfare effects of a merger

Note: The figure plots the average welfare effects $\hat{W}_Z$ of a merger between the first two firms in an industry with a high productivity spillover $\varrho = 0.5$ (see description in the text), where sets Z are quintiles of sectors by granular export share Γ_z^*. The left panel plots the overall average welfare effects $\hat{W}_Z$ in economies with different levels of trade costs τ (estimated baseline $\tau = 1.34$, free trade $\tau = 1$, and double the baseline $\tau = 2.68$ corresponding to a high 168% trade cost). The right panel plots the decomposition of the overall effect $\hat{W}_Z$ into the contributions from consumer and producer surplus, following (22), for the baseline level of trade cost $\tau = 1.34$ (the blue circled lines are the same in the two panels). Appendix Figure A8 reports the alternative case with low spillover $\varrho = 0.1$.

This analysis suggests a rationale for governments in open countries to be more lenient towards mergers, especially in more granular industries. Of course, mergers that are beneficial from the perspective of the home economy can be detrimental for foreign welfare, as they lead to a deterioration in foreign terms of trade. A lenient antitrust policy can, thus, substitute for a beggar-thy-neighbor trade policy. This suggests an important role for international cooperation over M&A policies in open economies to avoid excessive build-up of market power.

6.2 Granular import tariff

The second policy we consider is a granular import tariff $\varsigma_{z,1}$ imposed on the largest foreign exporter. Narrow trade restrictions and antidumping duties that target individual firms have been regularly emphasized in the policy debate. We contrast this granular tariff with a uniform import tariff $\bar{\varsigma}_z$ in the same industry that generates the same tariff revenue.[45] The government may consider a granular tariff over a sectoral one for two reasons. First, it might be more

---

of productivity spillovers ϱ, which leads to greater efficiency gain in more granular sectors (*cf* Figures 6 and A8). The gain in producer surplus is larger in more granular sectors, conditional on fundamental productivity Φ_z^* (note that a single sort on Γ_z^* in the figures induces a negative selection on Φ_z^*, explaining the slight non-monotonicity).

[45]In particular, we set $\bar{\varsigma}_z$ as a uniform 1% tariff on all imports, and then search for the value of $\varsigma_{z,1}$ on the top foreign firm, which yields the same tariff revenue at the industry level. See Appendix B.2 for full modeling details.
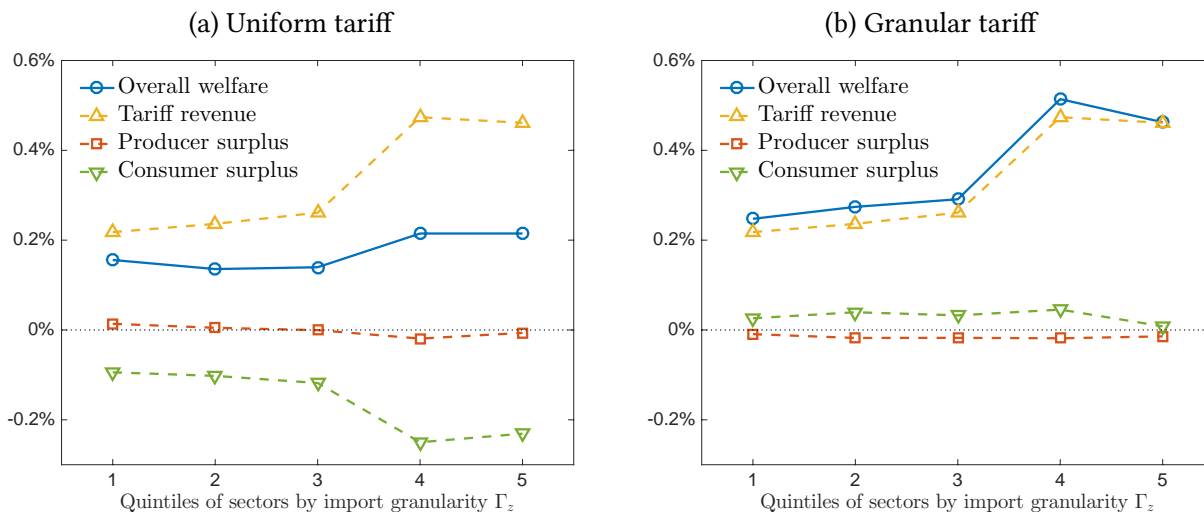
Figure 7: Welfare effects of granular and uniform tariffs

Note: GE average welfare effects $\hat{W}_Z$ (and their decomposition according to (22)) of a 1% uniform import tariff $\bar{\varsigma}_z$ and a corresponding granular tariff on the top foreign exporter $\varsigma_{z,1} \geq \bar{\varsigma}_z$, which yields the same tariff revenue (hence, yellow-triangled lines coincide in the two subplots, providing a benchmark for comparison). Subsets Z correspond to quintiles of sectors sorted by *foreign* granularity Γ_z. Tariff revenue is TR/Y, the change in the producer surplus is $d\Pi/Y$ and the change in the consumer surplus is $-d\log P$, which together add up to the overall welfare change $\hat{W}_Z$.

attractive in terms of domestic political economy, though perhaps more complex to impose legally. Second, it might be a more effective policy at extracting surplus from foreign producers and improving the home country's terms of trade. As we shall shortly see, this latter consideration is indeed the case in our granular model with oligopolistic competition.

Figure 7 describes the general equilibrium welfare consequences of the two alternative import tariffs and its decomposition based on (22). Again, we consider the policy applied to a subset of sectors $z \in Z$, which correspond to the quintiles of sectors, but now in terms of *foreign* granularity Γ_z. We find that home benefits from both types of import tariffs, but always more from imposing a granular tariff, especially in sectors where foreign enjoys granular comparative advantage. While a 1% uniform import tariff leads to average welfare gains of around 0.2%, the gains from a corresponding granular tariff can be more than doubled in sectors with outstanding foreign firms.

Turning to the decomposition of these effects, we first note that, by construction, tariff revenues are the same under both policies. Furthermore, Figure 7 shows that the post-tariff change in the total home firm profits is very small and nearly identical under the two alternative tariffs. Hence, the main difference between the two policies lies in the extent of consumer surplus destruction, i.e. increase in the home price level triggered by the tariffs. A uniform tariff has a strong negative effect on the consumer surplus, which offsets a large part of the

gains from tariff revenue. On net, the general equilibrium terms of trade improvement, triggered by a reduction in foreign wages, leads home to enjoy an increase in real income and welfare, largely independently from the extent of foreign granularity of the sector.[46]

In contrast, a granular tariff achieves an improvement in terms of trade without destroying the surplus of home consumers. As can be seen in the right panel of Figure 7, perhaps surprisingly, the average prices faced by the home consumers even fall slightly, a version of the *Metzler paradox* (see also Helpman and Krugman 1989, Chapters 4 and 7). The reason is that in response to a granular tariff, the large foreign exporter choses to reduce its markup in an attempt to maintain its home market share. This incomplete pass-through of the tariff into prices leads to a partial-equilibrium terms of trade improvement, which reinforces the general-equilibrium one via a declining foreign wage.[47] A further implication of this mechanism is that a granular tariff leads to a much smaller loss in the volume of trade, reducing the import shares in the targeted sectors by a small percentage. This is in contrast with a uniform import tariff, which leads to a large reduction in trade flows, as we illustrate in the Appendix Figure A9. To conclude, this analysis suggests that a granular tariff is an effective way of extracting surplus from foreign producers and converting it into home tariff revenues.

7   Conclusion

Granular firms play a pivotal role in international trade. The goal of this paper is to contribute to our understanding of the granular features of the global economy, with a particular focus on international trade flows, and to develop tools to analyze them. To this end, we propose and quantify a granular multi-sector model of trade, which combines fundamental comparative advantage across sectors with granular comparative advantage embodied in outstanding individual firms. We adopt the workhorse model of international trade with heterogeneous firms, in which we relax the standard assumption of a continuum of firms, and hence allow individual firms to shape aggregate trade patterns. The model, estimated using a rich set of sectoral and firm-level moments, suggests that granularity account for about 20% of the variation in realized export intensity across sectors. Moreover, granularity contributes markedly to skewness in aggregate outcomes, as it is most pronounced in the most export-intensive sectors. As a consequence, idiosyncratic firm dynamics accounts for a large share of the evolution of a country's comparative advantage over time, as well as its mean reversion across sectors.

---

[46]Since foreign granular residual Γ_z correlates positively with home import share $\Lambda_z = \Phi_z + \Gamma_z$, a 1% uniform import tariff results in greater tariff revenues, but also in larger loss in consumer surplus, in bins of sectors with high Γ_z. These two opposing effects net out, leaving the overall welfare effect of a uniform tariff rather stable.

[47]The pass-through is more complete for a uniform tariff for two reasons: (a) it constitutes a more aggregate shock, affecting a larger number of firms, and (b) smaller firms exhibit larger cost pass-through (in the model, as well as in the data; see Amiti, Itskhoki, and Konings 2018).

The granular structure of the world economy offers powerful incentives for governments to adopt trade and industrial policies targeted at individual firms.

By relying on the conventional modeling assumption of exogenous productivity draws, our model abstracts from an important question, the one of the origins of outstanding firms. We see this line of analysis as very fruitful for future research. In particular, it would help better understand whether government policies can and should promote the growth of "national champions". Another mechanism we assume away in this paper are productivity spillovers between independent firms. Such spillovers may be important in practice, especially for firms that are located close together, as the literature on agglomeration economies suggests. Analyzing the role of granular firms and their location decisions in determining the productivity and growth trajectories of individual cities (e.g., the decisions of Microsoft to move from Albuquerque to Seattle in 1979) is another fascinating question that we leave for future research.
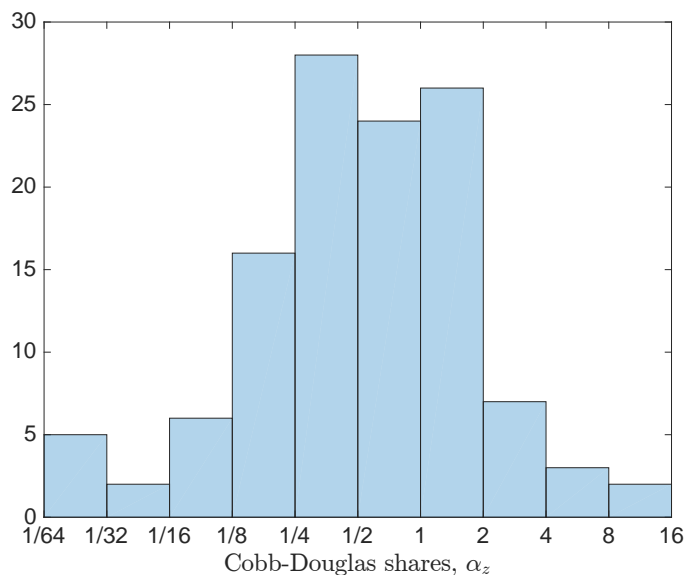
A   Additional Figures and Tables

Figure A1: Sectoral Cobb-Douglas shares in the data

Note: $\alpha_z = \tilde{N}\tilde{\alpha}_z$ so that $\mathbb{E}\alpha_z = 1$, as required by a model with a continuum of sectors.
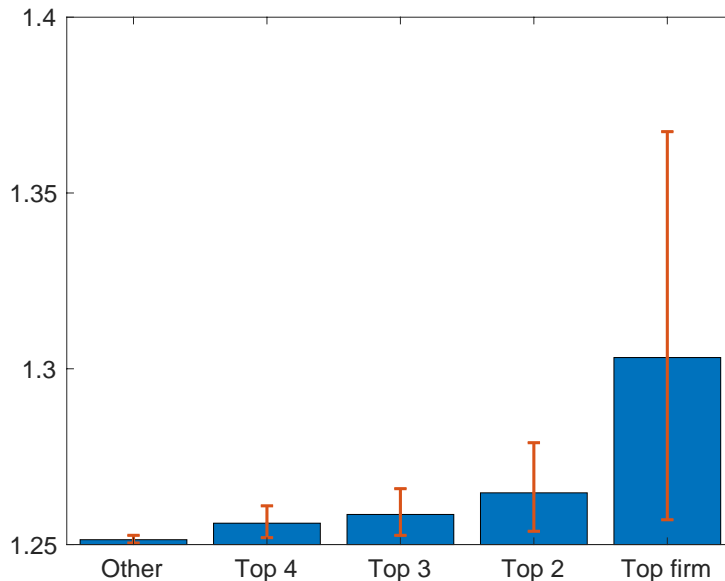

Figure A2: Equilibrium markups

Note: The bars in the figure correspond to markups for the four largest French firms in each sector and for the residual fringe of French firms, averaged across sectors, while the intervals correspond to the 10–90 percentiles across sectors. Markups under monopolistic competition with continuum of firms equal $\frac{\sigma}{\sigma-1} = 1.25$ for all firms, and this constitutes the lower bound for all markups in our oligopolistic model.
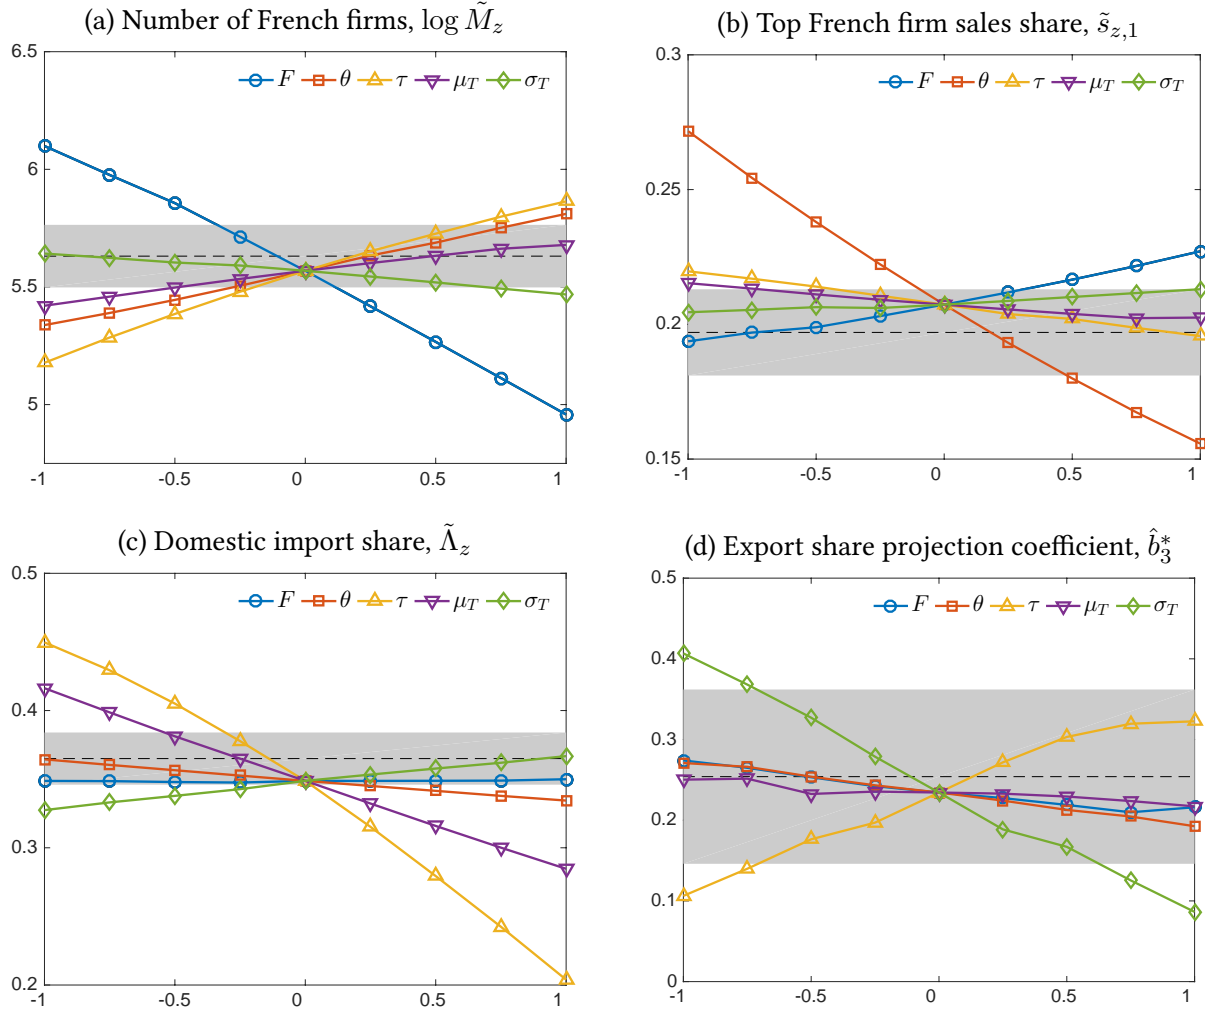
Figure A3: Identification plots

Note: The lines in the plots trace out the effects of a change in one parameter at a time on select moments used in estimation: (a) mean log number of French firms, $\log \tilde{M}_z$ (moment 1 in Table 2); (b) mean top French firm domestic sales share relative to all French firms, $\tilde{s}_{z,1}$ (moment 3); (c) average foreign share in the home market, $\tilde{\Lambda}_z$ (moment 7); and (d) regression coefficient of export share $\tilde{\Lambda}_z^{*\prime}$ on top-3 firms relative sales share in the home market ($\sum_{j=1}^{3} \tilde{s}_{z,j}$), b_3^* (moment 13). Black dashed horizontal lines correspond to the empirical values of the respective moments, and the shaded areas plot a bootstrap standard error band, which characterizes the degree of empirical uncertainty about the value of the moment. The x-axis is the normalized grid for the values of the parameters, where 0 corresponds to the estimated parameter vector $\hat{\Theta}$: (i) for F we use a log grid on $[\hat{F}/2, 2\hat{F}]$; (ii) for θ we use a linear grid such that $\kappa = \theta/(\sigma-1)$, where $\sigma = 5$, ranges on $\hat{\theta}/(\sigma-1) \pm 0.125 \approx [0.95, 1.2]$; (iii) for $\tau - 1$ we use a log grid such it varies on $[(\hat{\tau}-1)/2, 2(\hat{\tau}-1)] \approx [0.15, 0.7]$; (iv-v) for μ_T and σ_T we use linear grids on $\hat{\mu}_T \pm 0.4$ and $\sigma_T \pm 0.4$ respectively. See the text in Section 3.3 for interpretation.
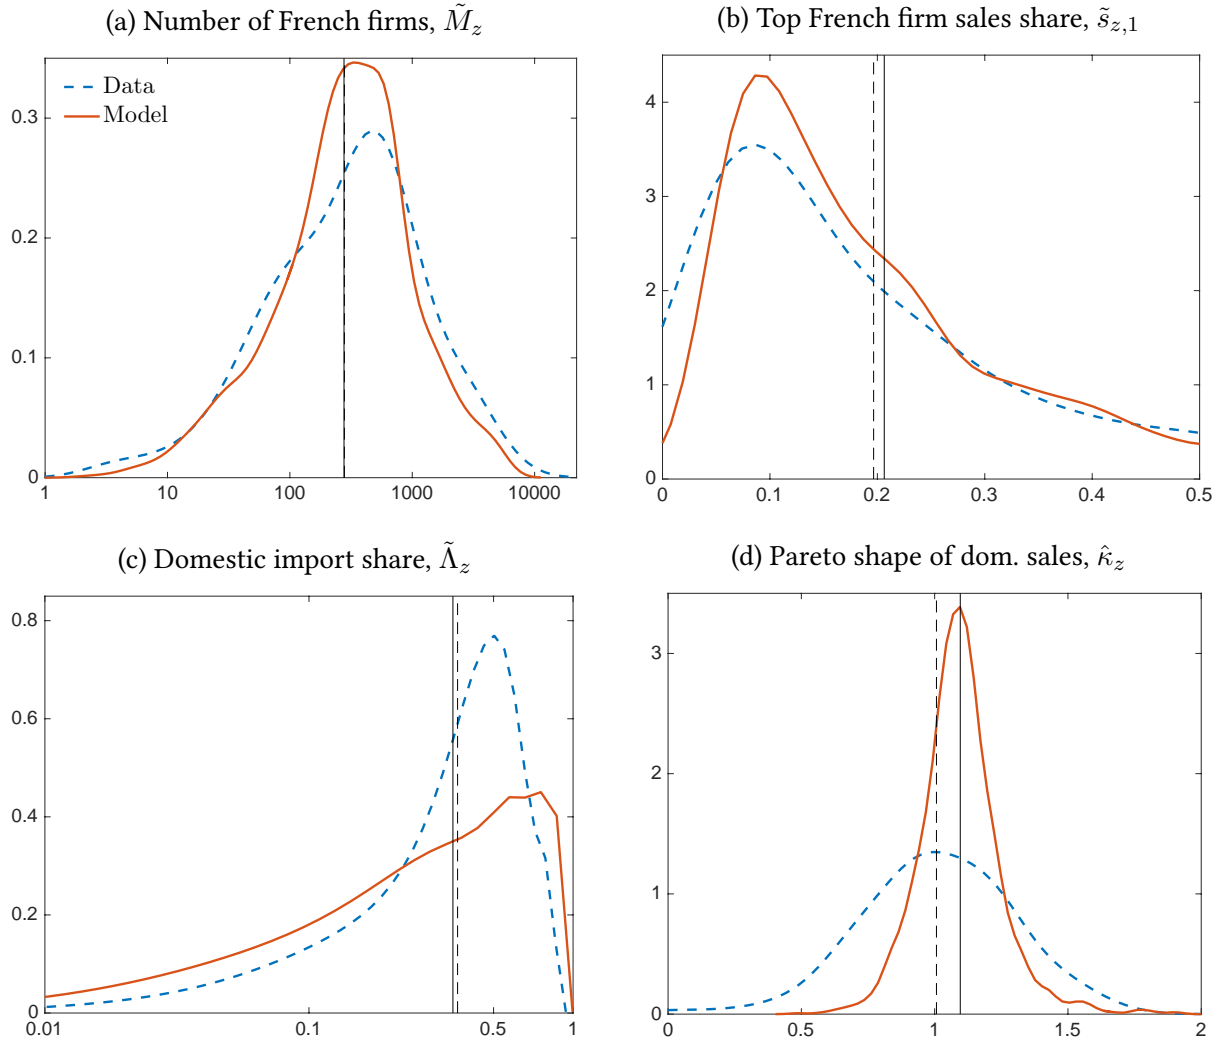
Figure A4: Distributions across sectors: model and data

Note: (a) corresponds to moments 1–2 in Table 2; (b) corresponds to moments 3–4; (c) corresponds to moments 7–8; while the moments in (d) are not directly target in the baseline estimation (see Table 3). In (b), top French firm market share is relative to other French firms in the domestic market. Pareto shapes $\hat{\kappa}_z$ are estimated according to (19) for firms above the 75th percentile in terms of domestic sales within sector. The vertical lines indicate the means of the respective distributions (dashed for data and solid for the model).

(a) Pareto shape parameters, $\kappa_z = \frac{\theta_z}{\sigma-1}$

(b) Estimated Pareto shapes, $\hat{\kappa}_z$

(c) Top French firm sales share, $\tilde{s}_{z,1}$

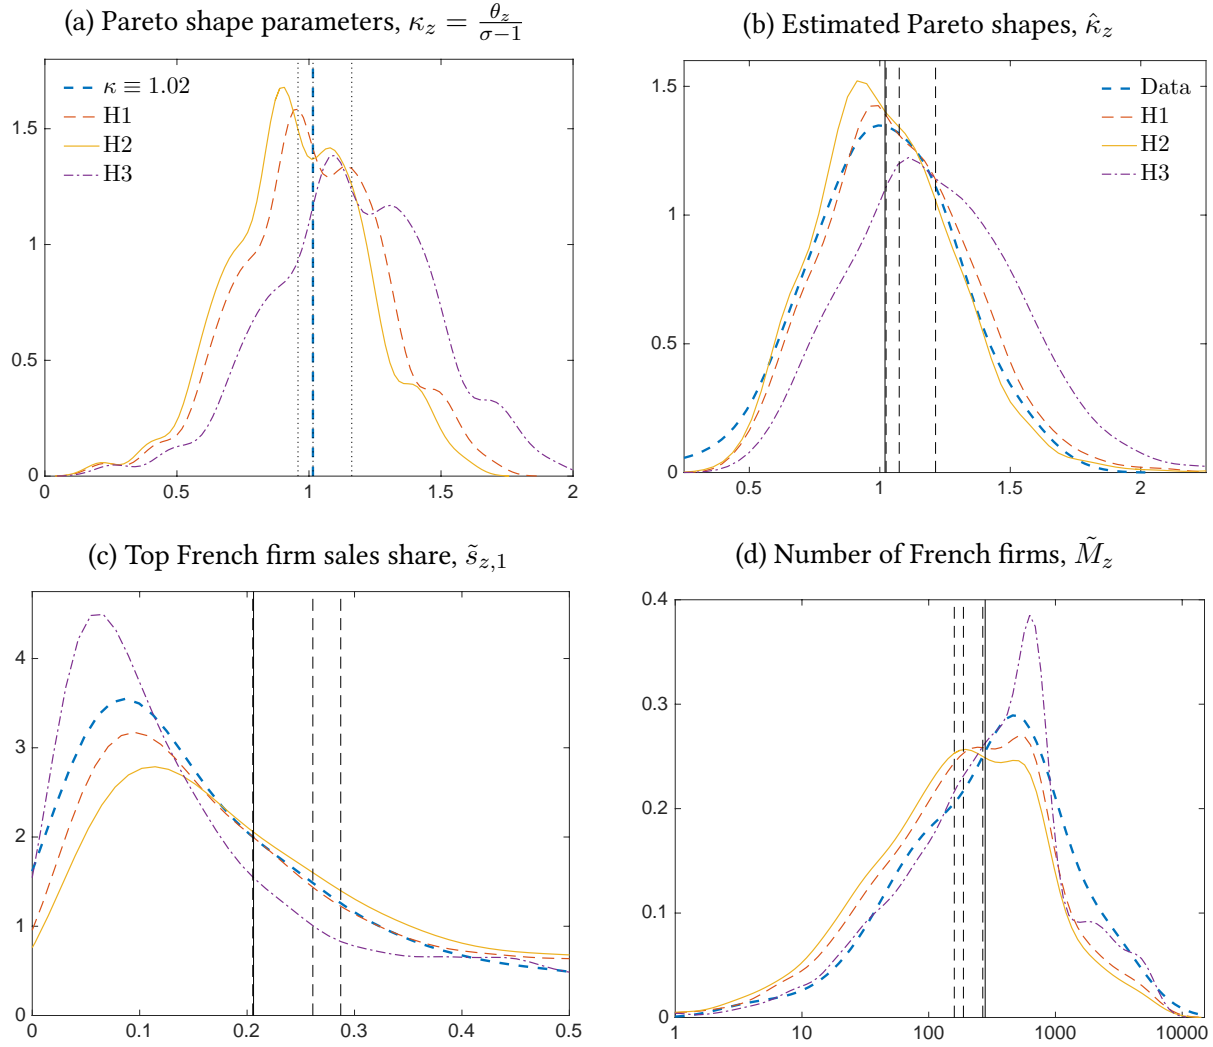(d) Number of French firms, $\tilde{M}_z$

Figure A5: Distributions across sectors: different model specifications with heterogeneous θ_z

Note: Panels (b)–(d) correspond to panels (d), (b) and (a) in Figure A4. Panel (a) plots the kernel densities of the model parameter $\kappa_z = \frac{\theta_z}{\sigma-1}$. Each plot considers three specifications with heterogeneous sector-specific θ_z, as described in Table 3, which we denote H1–H3 respectively. H1 matches average $\kappa_z = 1.02$. H2 matches average estimated Pareto shapes $\hat{\kappa}_z = 1.02$. H3 matches average top market share $\tilde{s}_{z,1} = 0.21$.
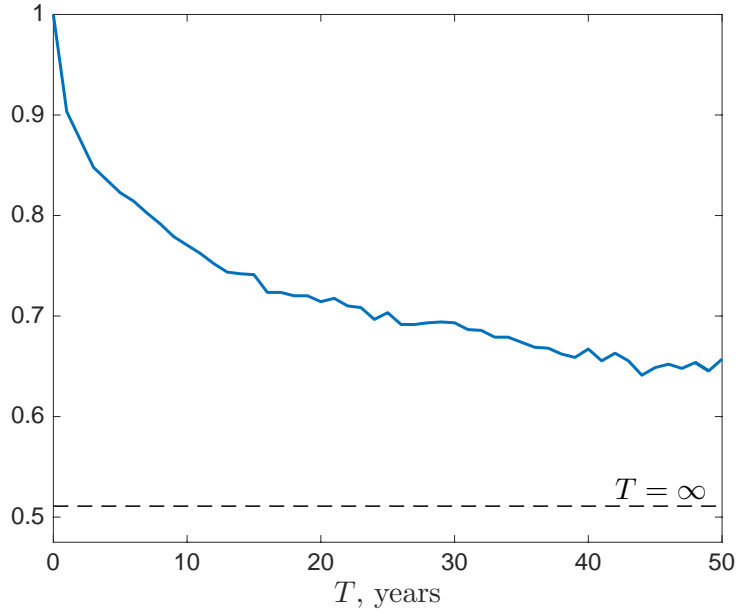
Figure A6: Probability a sector remains among top-5% of export-intensive sectors, over time

Note: Model-simulated data for 25,000 sectors (as in Table 4 and Figure 4). Sectors are sorted by export intensity Λ_z^*. $T = \infty$ corresponds to a 10,000 year-horizon.

(a) All sectors, deciles of Γ_z^*

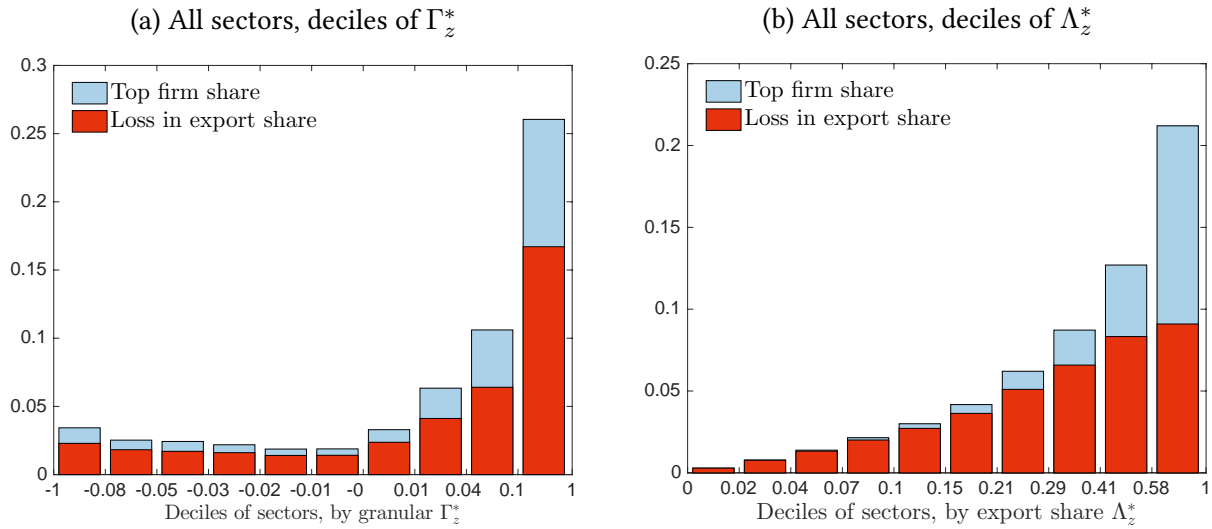(b) All sectors, deciles of Λ_z^*

Figure A7: Trade effects of individual firm exit: additional patterns

Note: The left plot is the same as Figure 5 in the main text, but done for all sectors, and not just for the top 30% of sectors in terms of export intensity Λ_z^*. In this case, there is a non-monotonic pattern across deciles of Γ_z^* because of the strong selection effect for the very bottom deciles of Γ_z^* (as they must be very high Φ_z^* sectors since $\Lambda_z^* = \Phi_z^* + \Gamma_z^* \geq 0$ by construction). The right panel is the same as the left panel, but deciles are sorted by the overall initial export intensity of the sectors Λ_z^*. Bottom deciles by Λ_z^* are low on both Φ_z^* and Γ_z^*, and hence the destroyed market share of the top home firm is fully captured by foreign firms.
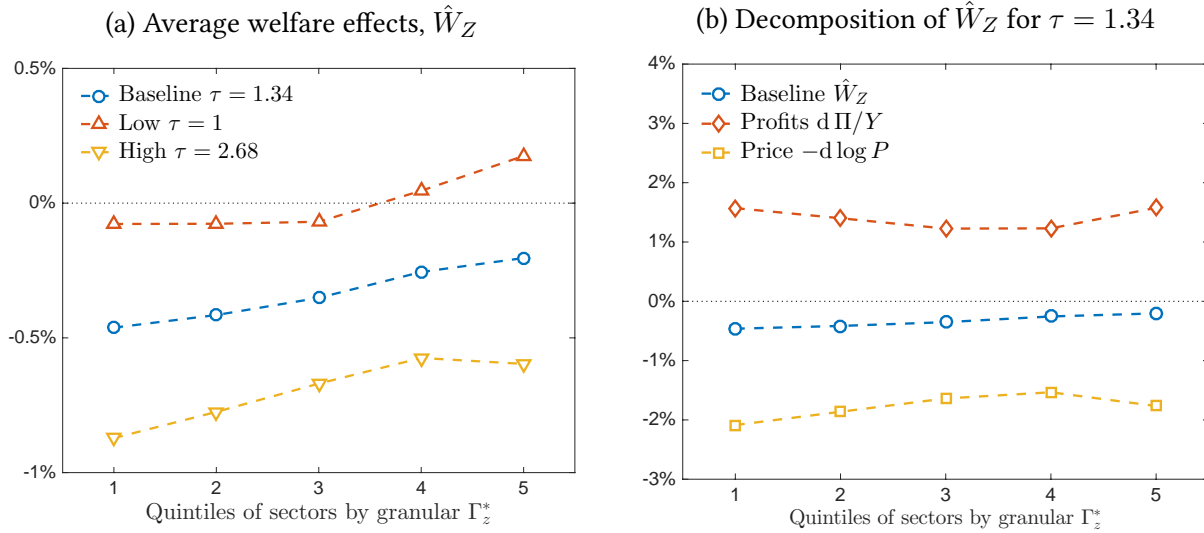
44

(a) Average welfare effects, $\hat{W}_Z$

(b) Decomposition of $\hat{W}_Z$ for $\tau = 1.34$

Figure A8: Welfare effects of a merger: low productivity spillover ($\varrho = 0.1$)

Note: This figure reproduces Figure 6 from the text for an alternative low level of productivity spillover $\varrho = 0.1$.

(a) Average welfare effects, $\hat{W}_Z$

(b) Change in import share $\hat{\Lambda}_Z$
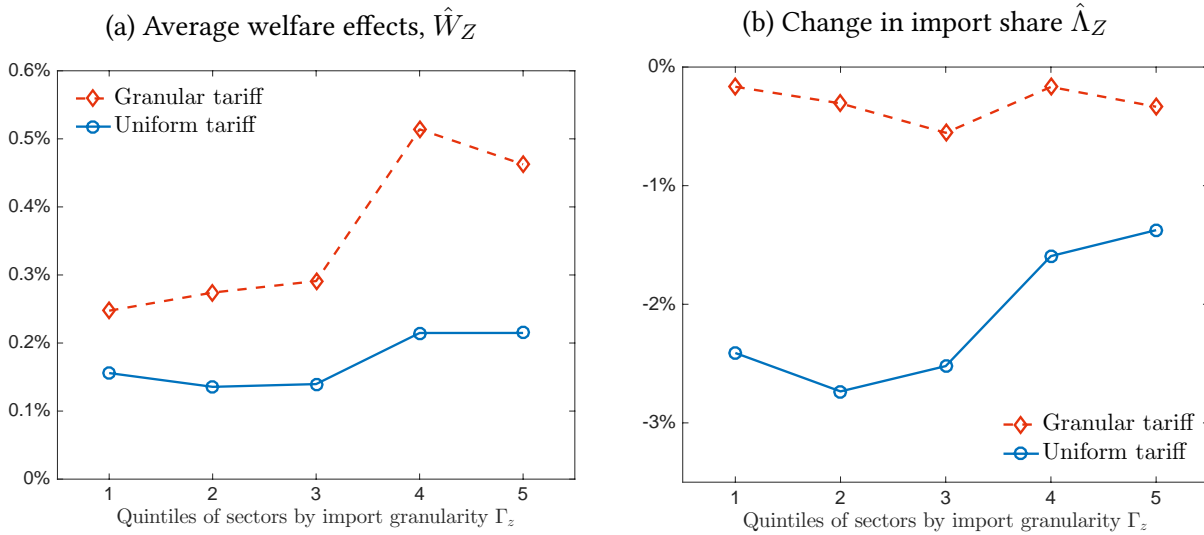
Figure A9: Granular versus uniform import tariff

Note: Welfare and trade effects of a 1% uniform import tariff $\bar{\varsigma}_z$ and a corresponding granular tariff on the top foreign exporters $\varsigma_{z,1} \geq \bar{\varsigma}_z$, which yields the same tariff revenue. $\hat{W}_Z$ are average general equilibrium effects for a subset of sectors $z \in Z$, which correspond to quintiles of sectors sorted by *foreign* granularity Γ_z. See the decomposition of $\hat{W}_Z$ in Figure 7. $\hat{\Lambda}_Z$ is the percent change in import share in the impacted sectors $z \in Z$.

Table A1: Properties of granular exports Γ^*_z

	(1)	(2)	(3)	(4)	(5)	(6)
$\tilde{s}_{z,1}$		0.335	0.373	0.379	0.357	0.354
$\tilde{s}^*_{z,1}$					−0.254	−0.268
$\log \tilde{M}_z$	−0.008		0.012	0.016		−0.011
$\log(\alpha_z Y)$				−0.005		0.013
Φ^*_z				0.004		0.073
R^2	0.013	0.353	0.375	0.376	0.520	0.539

Note: The table reports coefficients and R^2 from the projections of granular trade share $\Gamma^*_z = \Lambda^*_z - \Phi^*_z$ on various observable variables (and a constant). $\tilde{M}_z$ is the number of French firms in the French market, $\tilde{s}_{z,1}$ is the relative size of the largest French firm in the French market relative to other French firms, and $\tilde{s}^*_{z,1}$ is a corresponding variable for foreign firms. Projections are done using 10,000 counterfactual sectors drawn from the estimated model, and thus standard errors on the projection coefficients are nearly zero. The partial correlations of Γ^*_z with $\log(\alpha_z Y)$ and Φ^*_z are close to zero, and we omit these results for brevity (compare the incremental R^2 in columns 2–4 and 5–6). We have also experimented with two other proxies for granularity — concentration ratio of the three largest domestic firms (relative to other domestic firms) and the ratio of the top to median market shares of the domestic firms; both of these variables have similar, albeit slightly lower, explanatory power as $\tilde{s}_{z,1}$, but they do not improve the fit when combined together.

Table A2: Granularity, openness and labor reallocation

	Variable trade cost				
	$\tau = 6.70$	$\tau = 2.68$	$\tau = 2.00$	$\tau = 1.34$	$\tau = 1$
A. Annual changes:					
Total labor reallocation, TLR	0.120	0.119	0.117	0.115	0.115
Sectoral labor reallocation, SLR	0.004	0.007	0.013	0.022	0.025
SLR/TLR	0.034	0.063	0.109	0.187	0.218
B. 10-year changes:					
Total labor reallocation, TLR	0.346	0.342	0.335	0.327	0.325
Sectoral labor reallocation, SLR	0.010	0.017	0.030	0.050	0.062
SLR/TLR	0.028	0.051	0.091	0.152	0.191

Note: dynamics simulations of the estimated granular economy with 50,000 sectors over 11 years, for different values of variable trade cost τ. Total labor reallocation TLR (average of job creation and job destruction rates) and sectoral labor reallocation SLR (net sectoral employment change) are defined formally in (A15) in Appendix B.2 (see also Davis and Haltiwanger 1999). Top (bottom) panel reports the results for one-year (ten-year) employment changes: similar patterns hold over longer horizons, with inter-sectoral reallocation accounting for a slightly lower share of the total labor reallocation.

B Theory Appendix

B.1 Continuous DFS-Melitz model

We review here the *continuous* model, which serves as a benchmark in our *granular* analysis. Consider a two-country multi-sector extension of the Melitz (2003) model, with Ricardian comparative advantage across a unit continuum of sectors indexed by $z \in [0, 1]$, as in Dornbusch, Fischer, and Samuelson (1977).[48] We refer to this benchmark economy as DFS-Melitz. More specifically, within each sector z we consider the Chaney (2008) version of the Melitz model without free entry, in which an exogenous mass of firms $\bar{M}_z$ are present and their productivities are drawn from a Pareto distribution with a sector-specific lower bound $\underline{\varphi}_z$ and a shape parameter θ common across all sectors. We show below that in this model, the overall sectoral productivity is determined by $T_z = \bar{M}_z \cdot \underline{\varphi}_z^\theta$, as in (6). The two countries differ in the sectoral productivity measures, $\{T_z\}$ at home and $\{T_z^*\}$ in foreign, which is the source of the Ricardian comparative advantage across sectors.

Households are as described in Section 2 with the exception that, instead of (2), the sectoral CES consumption bundles aggregate over a continuum of individual varieties ω:

$$Q_z = \left[\int_{\omega \in \Omega_z} q_z(\omega)^{\frac{\sigma-1}{\sigma}} \mathrm{d}\omega \right]^{\frac{\sigma}{\sigma-1}}, \tag{A1}$$

where Ω_z is the set of varieties available for consumption in sector z at home, and the resulting price index is $P_z = \left[\int_{\omega \in \Omega_z} p_z(\omega)^{1-\sigma} \mathrm{d}\omega \right]^{1/(1-\sigma)}$. The foreign demand structure is symmetric, with Ω_z^* replacing Ω_z.

Firms and productivity are also as described in Section 2, with the exception that $\bar{M}_z$ is a deterministic mass of existing shadow firms in each sector, with individual productivities $\varphi_z(\omega) \sim iid \, \mathrm{Pareto}(\theta, \underline{\varphi}_z)$ with $\mathbb{P}\{\varphi_z(\omega) \leq \varphi\} = 1 - (\underline{\varphi}_z/\varphi)^\theta$ representing the realized productivity frequencies. A continuous model requires a parameter restriction $\theta > \sigma - 1$.

Each firm is infinitesimal in the markets it serves. Therefore, upon entry, firms compete according to monopolistic competition in each market. They set a constant markup $\sigma/(\sigma-1)$ over their marginal costs. This implies that the firm's operating profit in each market equals $1/\sigma$ of its revenues, and the overall profit of the firm can be written as:

$$\pi_z(\omega) = \left[\left(\frac{\sigma}{\sigma-1} \frac{w/P_z}{\varphi_z(\omega)} \right)^{1-\sigma} \frac{\alpha_z Y}{\sigma} - wF \right]^+ + \left[\left(\frac{\sigma}{\sigma-1} \frac{\tau w/P_z^*}{\varphi_z(\omega)} \right)^{1-\sigma} \frac{\alpha_z Y^*}{\sigma} - w^* F^* \right]^+,$$

---

[48]This model extends Melitz (2003) in a multi-sector way, the same way Costinot, Donaldson, and Komunjer (2012) extend the Eaton and Kortum (2002) model. Other papers which considered a multi-sector DFS-Melitz environment, albeit under somewhat different formulation, are Okubo (2009) and Fan, Lai, and Qi (2015).

where we substituted the markup pricing rule over the marginal cost into the expression for revenues (3), and we use the notation $[x]^+ \equiv \max\{0, x\}$.[49] Firms with sufficiently high productivities profitably enter the home and the foreign markets respectively, as is conventional in the Melitz model. We denote with $\varphi_{h,z}$ and $\varphi_{f,z}$ the productivity cutoffs for a domestic firm to enter the home and foreign markets respectively in sector z, and rewrite profits as:

$$
\begin{cases}
\pi_z(\omega) = wF \left[\left(\frac{\varphi_z(\omega)}{\varphi_{h,z}} \right)^{\sigma-1} - 1 \right]^+ + w^* F^* \left[\left(\frac{\varphi_z(\omega)}{\varphi_{f,z}} \right)^{\sigma-1} - 1 \right]^+, \\
\varphi_{h,z} = \frac{\sigma}{\sigma-1} \frac{w}{P_z} \left(\frac{\sigma wF}{\alpha_z Y} \right)^{1/(\sigma-1)} \quad \text{and} \quad \varphi_{f,z} = \frac{\sigma}{\sigma-1} \frac{\tau w}{P_z^*} \left(\frac{\sigma w^* F^*}{\alpha_z Y^*} \right)^{1/(\sigma-1)}.
\end{cases}
\tag{A2}
$$

The foreign firms are symmetric, and we denote with $\pi_z^*(\omega)$ their profits, and with $\varphi_{h,z}^*$ and $\varphi_{f,z}^*$ their productivity cutoffs for entry into the home and foreign markets respectively.

Sectoral equilibrium Using the definition of the price index, the markup pricing rules, the cutoff definitions in (A2), and the Pareto productivity distribution, we can integrate to solve for the price index in sector z in the home market:

$$
P_z = \frac{\sigma}{\sigma - 1} w \left[\frac{\kappa}{\kappa - 1} \frac{T_z}{1 - \Phi_z} \right]^{-1/\theta} \left(\frac{\sigma wF}{\alpha_z Y} \right)^{(\kappa-1)/\theta}, \tag{A3}
$$

where $\kappa \equiv \theta/(\sigma - 1)$ and Φ_z is the *foreign share*, as defined in (13).[50] The sectoral price index in (A3) increases in the local wage rate and in the relative fixed cost of entry $(wF)/(\alpha_z Y)$, and decreases in sectoral productivity T_z and in the foreign share Φ_z, which reflects the gains from trade (see ACR). Using (A3), we can express all sectoral variables as functions of the general equilibrium vector (w, w^*, Y, Y^*) and exogenous parameters of the model, completing the description of the sectoral equilibrium.

The definition of the foreign share Φ_z, and its symmetric counterpart in the foreign country Φ_z^*, makes it straightforward to calculate sectoral exports of home and foreign countries respectively:

$$
X_z = \alpha_z \Phi_z^* Y^* \qquad \text{and} \qquad X_z^* = \alpha_z \Phi_z Y, \tag{A4}
$$

and sectoral net exports is $NX_z = X_z - X_z^*$. In addition, we also characterize the allocation

---

[49]Specifically, a home firm sets $p_z(\omega) = \frac{\sigma}{\sigma-1} \frac{w}{\varphi(\omega)}$ in the home market, which results in revenues $(p_z(\omega)/P(z))^{1-\sigma} \alpha_z Y$, according to (3), and the operating profits equal fraction $1/\sigma$ of these revenues due to constant markup pricing. Net profits are operating profits net of the fixed entry cost. Symmetric characterization applies to profits in the foreign market, with the difference that the marginal cost of delivering a good abroad is augmented by iceberg trade cost τ.

[50]We note that the foreign share in (12) does not depend on the fixed costs since both domestic and foreign firms are assumed to face the same fixed costs of entry into the home market. As a result, fixed costs in this framework have little effect on the key variables which characterize equilibrium, apart from the price indexes P_z and P_z^*, which increase with the fixed cost of entry into the market, thereby reducing local welfare.

of aggregate labor supply to sector z, which in the home market satisfies:

$$wL_z = \alpha_z Y \left(\tfrac{\sigma\kappa-1}{\sigma\kappa}[1 - \Phi_z] + \tfrac{\kappa-1}{\sigma\kappa}\Phi_z \right) + \alpha_z Y^* \tfrac{\sigma-1}{\sigma}\Phi_z^*. \tag{A5}$$

The last term is labor used in production of goods for foreign market, while the first two terms are labor used for production and entry costs in the home market.[51] Combining (A4) and (A5), with (A8) below, we obtain the relationship between sectoral net exports and labor allocation:

$$\frac{L_z}{L} = \alpha_z + \frac{\theta}{\sigma\kappa - 1}\frac{NX_z}{Y}. \tag{A6}$$

In autarky, $L_z = \alpha_z L$ due to the Cobb-Douglas preferences, yet in the open economy labor reallocates towards the sectors with comparative advantage.

General equilibrium requires balanced current account and labor market clearing in both countries, which (together with our choice of numeraire $w^* = 1$) allow us to solve for (w, w^*, Y, Y^*). These three conditions also imply countries' budget balances (9) by Walras Law.

Balanced current account can, in general, be different from the balanced trade in this model, as exporting requires paying a fixed cost in the destination market. Nonetheless, the two coincide in the continuous model with a Pareto distribution. The total home income obtain from exports in sector z equals the value of exports $X_z = \alpha_z\Phi_z^* Y^*$ net of the fixed cost of entry into the foreign market $\tfrac{\kappa-1}{\sigma\kappa}\alpha_z\Phi_z^* Y^*$, which is proportional to exports $X_z = \alpha_z\Phi_z^* Y^*$, with a constant factor $\tfrac{\sigma\kappa-\kappa+1}{\sigma\kappa}$ in front. Aggregating across sectors and equalizing with the foreign export income, we obtain the balanced current account (and trade balance condition):

$$Y \int_0^1 \alpha_z\Phi_z \mathrm{d}z = Y^* \int_0^1 \alpha_z\Phi_z^* \mathrm{d}z. \tag{A7}$$

Next, aggregating sectoral labor demand in (A5) across z and using trade balance (A7), we obtain aggregate labor market clearing:

$$wL = \tfrac{\sigma\kappa-1}{\sigma\kappa}Y \qquad \text{and} \qquad w^*L^* = \tfrac{\sigma\kappa-1}{\sigma\kappa}Y^*. \tag{A8}$$

Therefore, total labor income is a constant share of GDP (total income), with the complementary share coming from firm profits. Combining (A7) with (A8) and normalizing $w = 1$, allows to solve for (w^*, Y, Y^*), completing the description of the general equilibrium.[52]

---

[51]A fraction $\tfrac{\sigma-1}{\sigma}$ of revenues goes to cover variable production labor costs (in the country of production). Integrating across firms, a fraction $\tfrac{\kappa-1}{\sigma\kappa}$ of revenues goes to cover entry labor costs (in the country of entry). Note that the first term in (A5) can be decomposed as $\tfrac{\sigma\kappa-1}{\sigma\kappa} = \tfrac{\sigma-1}{\sigma} + \tfrac{\kappa-1}{\sigma\kappa}$. The remaining $\tfrac{1}{\sigma\kappa}$ share is net profits.

[52]Taking the ratio of the two equations in (A8), we have $Y/Y^* = (wL)/(w^*L^*)$, which together with (A7) allows to solve for both relative wage w/w^* and relative incomes Y/Y^*, as in the DFS model. Recall from (12), that Φ_z and Φ_z^* can be written as function of relative wages w/w^* and the exogenous parameters of the model.

DFS limit The continuous DFS-Melitz benchmark admits as a limiting case the classical DFS formulation when within-sector firm heterogeneity collapses. Specifically, the DFS model emerges as a limit of the DFS-Melitz model when $\theta, \sigma \to \infty$, $F \to 0$, while at the same time holding constant $\kappa = \theta/(\sigma - 1)$, σF and the following productivity parameters: $a_z \equiv T_z^{1/\theta}$ and $a_z^* \equiv (T_z^*)^{1/\theta}$. In the DFS limit, the foreign shares Φ_z and Φ_z^* in (12) become step functions, defined by two cutoffs $\underline{z}, \bar{z} \in [0, 1]$. Specifically, we rank all sectors $z \in [0, 1]$ such that $a_z/a_z^* = (T_z/T_z^*)^{1/\theta}$ is a monotonically increasing function of z, and define the cutoffs to satisfy:

$$\frac{a_{\underline{z}}}{a_{\underline{z}}^*} = \frac{w}{\tau w^*} \qquad \text{and} \qquad \frac{a_{\bar{z}}}{a_{\bar{z}}^*} = \frac{\tau w}{w^*}, \tag{A9}$$

which implies $\underline{z} < \bar{z}$. For sectors $z \in [0, \underline{z})$, foreign is the only supplier of the good on both domestic and foreign markets, goods $z \in (\underline{z}, \bar{z})$ are non-traded and produced in both countries, and for goods $z \in (\bar{z}, 1]$ home is the only world supplier.

Continuous limit Lastly, we discuss how the granular model of Section 2 admits the continuous DFS-Melitz limit described above. We introduce a scaler $M > 0$, and rewrite the price index in (7) and the market share in (3) as follows:

$$P_z = \left[\frac{1}{M} \sum_{i=1}^{K_z} p_{z,i}^{1-\sigma} \right]^{1/(1-\sigma)} \qquad \text{and} \qquad s_{z,i} = \frac{1}{M} \left(\frac{p_{z,i}}{P_z} \right)^{1-\sigma},$$

where the granular model of Section 2 corresponds to the case with $M = 1$. Note that $\sum_{i=1}^{K_z} \tilde{s}_{z,i} = 1$ for any $M > 0$. We also rewrite the utility in (2) as $\tilde{Q}_z = \left[\frac{1}{M} \sum_{i=1}^{K_z} \tilde{q}_{z,i}^{\frac{\sigma-1}{\sigma}} \right]^{\frac{\sigma}{\sigma-1}}$, where $\tilde{q}_{z,i} = M q_{z,i}$ are the new consumption units. Lastly, the derived productivity parameter in (6) is generalized as $T_z = \frac{\bar{M}_z}{M} \cdot \underline{\varphi}_z^\theta$.

With this generalization to an arbitrary $M > 0$, we can now take the following limit: $M, \bar{M}_z \to \infty$ and $F \to 0$, such that $\bar{M}_z/M = const$ for all z and $MF = const$, and holding constant the other parameters of the model, including the location of the productivity distribution $\underline{\varphi}_z$. This keeps T_z unchanged. Furthermore, $\bar{M}_z/M$ now represents the relative measure of shadow firms in sector z. The ratio $K_z/\bar{M}_z$ tends to a constant related to productivity cutoffs (A2) in the continuous model; the price index P_z tends to a constant, the price level in the continuous model (A3); the market shares $s_{z,i} \to 0$ so that the elasticity in (7) $\varepsilon_{z,i} \to \sigma$ and markups become constant equal to $\sigma/(\sigma - 1)$; and the non-negativity of profits in (8) with $F \to 0$ at the same rate as $s_{z,i} \to 0$ now corresponds to the cutoff condition in (A2). All sums (redefined to feature $1/M$ or $s_{z,i}$ weights) converge to corresponding integrals in the continuous model, which is the direct counterpart to the granular model of Section 2.

B.2 Derivations and proofs for the granular model

Foreign share Consider the foreign share Λ_z defined in (12). We reproduce

$$\Lambda_z = \sum_{i=1}^{K_z} (1 - \iota_{z,i}) s_{z,i},$$

where $\iota_{z,i}$ is an indicator for whether the firm is of home origin. There is no analytical characterization for the distribution of $s_{z,i}$, which are complex transformation of the realized productivity vector, which relies both on the price setting and entry outcomes (e.g., see (3), (7) and (8)). Nonetheless, following EKS, we can prove that the conditional distributions of $s_{z,i}|\iota_{z,i} = 1$ and $s_{z,i}|\iota_{z,i} = 0$ are the same, i.e. the distribution of $s_{z,i}$ is symmetric for firms of home and foreign origin, and hence the expectation of Λ_z simply equals the unconditional expectation that any entrant is of foreign origin (i.e., the relative extensive margin of entry into the home market).

The formal argument proceeds in two steps (all expectations $\mathbb{E}_T\{\cdot\}$ are conditional on the realization of fundamental productivity T_z and T_z^*, which are hence treated as parameters):

1. For any $s > 0$, $\mathbb{E}_T\{\iota_{z,i}|s_{z,i} > s\} = \mathbb{P}_T\{\iota_{z,i} = 1|s_{z,i} > s\} = \frac{T_z w^\theta}{T_z w^\theta + T_z^* (\tau w^*)^\theta} = 1 - \Phi_z$, as defined in (13). Hence, $\mathbb{E}_T\{\iota_{z,i}|s_{z,i} > s\}$ does not depend on s, and $\mathbb{E}_T\{\iota_{z,i}|s_{z,i}\} = \mathbb{E}_T\iota_{z,i}$. See a sketch of a proof below.

2. $\mathbb{E}_T\Lambda_z = \sum_{i=1}^{K_z} \mathbb{E}_T\{(1-\iota_{z,i})s_{z,i}\} = \sum_{i=1}^{K_z} \mathbb{E}_T\{s_{z,i} \cdot \mathbb{E}_T\{1-\iota_{z,i}|s_{z,i}\}\} = \Phi_z \sum_{i=1}^{K_z} \mathbb{E}_T s_{z,i} = \Phi_z$, since $\mathbb{E}_T\{\sum_{i=1}^{K_z} s_{z,i}\} = \mathbb{E}_T\{1\} = 1$, and where the third equality uses property 1.

Property 1 follow from the Poisson-Pareto productivity draw structure and the application of the Bayes' formula. Indeed, in a given sectoral equilibrium, $s_{z,i}$ decreases with the cost of the firm $c_{z,i}$, which in turn decreases with the firm productivity ($\varphi_{z,i}$ if the firm is home and $\varphi_{z,i}^*$ if the firm is foreign; see (5)). Given the productivity draw structure, the number of home firms with productivity above φ is a Poisson random variable with parameter $\varphi^{-\theta} T_z$, and symmetrically for the foreign firms. Consequently, the number of home and foreign firms with a cost below c are independent Poisson random variables with parameters $(w/c)^{-\theta} T_z$ and $(\tau w^*/c)^{-\theta} T_z^*$, respectively. Therefore, we can calculate:

$$\mathbb{P}_T\{\iota_{z,i} = 1|s_{z,i} > s\} = \mathbb{P}_T\{\iota_{z,i} = 1|c_{z,i} < c\}$$
$$= \frac{\mathbb{P}_T\{c_{z,i} < c, \iota_{z,i} = 1\}}{\sum_{\iota \in \{0,1\}} \mathbb{P}_T\{c_{z,i} < c, \iota_{z,i} = \iota\}} = \frac{(w/c)^{-\theta} T_z}{(w/c)^{-\theta} T_z + (\tau w^*/c)^{-\theta} T_z^*} = 1 - \Phi_z.$$

Therefore, we conclude that indeed $\mathbb{E}_T\Lambda_z = \Phi_z$, and the granular residual $\Gamma_z = \Lambda_z - \Phi_z$ is zero in expectation for any sector z (see (13) and (14)).

Equilibrium system We reproduce here the full general equilibrium system of the granular model, which consists of the aggregate budget constraints and labor market clearing in both countries. Using (8) and (10), we write the home country budget $Y = wL + \Pi$ constraint as:

$$Y = wL + Y(1 - \Lambda)\frac{\bar{\mu}_H - 1}{\bar{\mu}_H} - wFK_H + Y^*\Lambda^*\frac{\bar{\mu}_H^* - 1}{\bar{\mu}_H^*} - w^*F^*K_H^*, \qquad (A10)$$

where

$$K_H = \int_0^1 \left[\sum_{i=1}^{K_z} \iota_{z,i}\right] dz,$$

$$K_H^* = \int_0^1 \left[\sum_{i=1}^{K_z^*} (1 - \iota_{z,i}^*)\right] dz,$$

$$1 - \Lambda = \int_0^1 \alpha_z(1 - \Lambda_z)dz = \int_0^1 \alpha_z \left[\sum_{i=1}^{K_z} \iota_{z,i}s_{z,i}\right] dz,$$

$$\Lambda^* = \int_0^1 \alpha_z\Lambda_z^*dz = \int_0^1 \alpha_z \left[\sum_{i=1}^{K_z^*} (1 - \iota_{z,i}^*)s_{z,i}^*\right] dz,$$

$$\frac{1}{\bar{\mu}_H} = \frac{1}{1 - \Lambda} \int_0^1 \alpha_z \left[\sum_{i=1}^{K_z} \iota_{z,i}\frac{s_{z,i}}{\mu(s_{z,i})}\right] dz,$$

$$\frac{1}{\bar{\mu}_H^*} = \frac{1}{\Lambda^*} \int_0^1 \alpha_z \left[\sum_{i=1}^{K_z^*} (1 - \iota_{z,i}^*)\frac{s_{z,i}^*}{\mu(s_{z,i}^*)}\right] dz,$$

where $\mu(s) = \frac{\varepsilon(s)}{\varepsilon(s)-1}$ and $\varepsilon(s) = \sigma(1 - s) + s$, as defined in (7). Note that:

- K_H and K_H^* are the total numbers of the home firms selling in the home and foreign markets respectively, across all industries;

- $1 - \Lambda$ and Λ^* are the average shares of the home firm sales in aggregate home and foreign expenditure Y and Y^* respectively;

- $\bar{\mu}_H$ and $\bar{\mu}_H^*$ are the (harmonic) average markups of the home firms in the home and foreign markets respectively, and hence $(\bar{\mu}_H - 1)/\bar{\mu}_H$ and $(\bar{\mu}_H^* - 1)/\bar{\mu}_H^*$ are the average shares of operating profits in aggregate revenues of the home firms in the home and foreign markets respectively, since $\frac{\mu(s_{z,i})-1}{\mu(s_{z,i})} = \frac{p_{z,i}-c_{z,i}}{p_{z,i}}$ for a firm with market share $s_{z,i}$.

A similar equation defines foreign budget $Y^* = w^*L^* + \Pi^*$, which we write as:

$$Y^* = w^*L^* + Y^*(1 - \Lambda^*)\frac{\bar{\mu}_F^* - 1}{\bar{\mu}_F^*} - w^*F^*K_F^* + Y\Lambda\frac{\bar{\mu}_F - 1}{\bar{\mu}_F} - wFK_F, \qquad (A11)$$

with K_F^*, K_F, $\bar{\mu}_F^*$ and $\bar{\mu}_F^*$ defined by analogy with the respective variables for home firms.

Now consider the home labor market clearing condition in expenditure terms (11), which

we write as:

$$wL = wFK + Y(1 - \Lambda)\frac{1}{\bar{\mu}_H} + Y^*\Lambda^*\frac{1}{\bar{\mu}_H^*}, \tag{A12}$$

where

$$K = K_H + K_F = \int_0^1 K_z \mathrm{d}z$$

is the total entry of firms in the home market across all sectors. A symmetric labor market clearing condition for foreign is:

$$w^*L^* = w^*F^*K^* + Y^*(1 - \Lambda^*)\frac{1}{\bar{\mu}_F^*} + Y\Lambda\frac{1}{\bar{\mu}_F}, \tag{A13}$$

where $K^* = K_H^* + K_F^*$ is the total entry of firms in the foreign market across all sectors.

It is immediate to verify that the equilibrium system (A10)–(A13) has the following properties:

1. It is linear in the general equilibrium vector (w, w^*, Y, Y^*) conditional on the vector

$$(\Lambda, \Lambda^*, K_H, K_H^*, K_F, K_F^*, K, K^*, \bar{\mu}_H, \bar{\mu}_H^*, \bar{\mu}_F, \bar{\mu}_F^*),$$

which depends on the outcome of the partial equilibrium $\left\{ K_z, K_z^*, \{s_{z,i}\}_{i=1}^{K_z}, \{s_{z,i}^*\}_{i=1}^{K_z^*} \right\}_{z \in [0,1]}$.

2. It is linearly dependent, so that any of the four equations follow from the other three. Normalizing $w = 1$ and dropping any of the equations (for example (A11)) results in a linearly independent system of three equations in three unknown (w^*, Y, Y^*) with a unique solution.

3. Substituting in labor market clearing (A12) into the budget constraint (A10) (or equivalently (A13) into (A11)) results in the current account balance condition (which in general differs from the trade balance $NX = \Lambda^*Y^* - \Lambda Y$):

$$\Lambda Y - wFK_F = Y^*\Lambda^* - w^*F^*K_H^*. \tag{A14}$$

The equilibrium system can be represented by system of three linearly independent equations (A12)–(A14). Note the similarity and differences of this equilibrium system with a corresponding system in the continuous model (A7)–(A8). In particular, due to discreteness and variable markups, the shares of labor income and profits in aggregate income are no longer constants $(\sigma\kappa - 1)/(\sigma\kappa)$ and $1/(\sigma\kappa)$.

Finally, using the same strategy we used to prove that $\mathbb{E}_T\Lambda_z = \Phi_z$ above, we can show that

$$\Lambda = \frac{K_F}{K_H + K_F} = \Phi = \int_0^1 \alpha_z\Phi_z\mathrm{d}z \qquad \text{and} \qquad \Lambda^* = \frac{K_H^*}{K_H^* + K_F^*} = \Phi^* = \int_0^1 \alpha_z\Phi_z^*\mathrm{d}z,$$

where the integrals of Φ_z and Φ_z^* can be viewed as expectations taken over the joint distribution of $(\alpha_z, T_z/T_z^*)$. As α_z and T_z/T_z^* are assumed independent, the values of Φ and Φ^* depend only on the parameters θ, τ and (μ_T, σ_T) of the distribution of T_z/T_z^*. Using this result, we can simplify the equilibrium system. For example, conditions (A10) and (A14) can be rewritten as:

$$Y = wL + (1 - \Phi)\left[Y\frac{\bar{\mu}_H - 1}{\bar{\mu}_H} - wFK\right] + \Phi^*\left[Y^*\frac{\bar{\mu}_H^* - 1}{\bar{\mu}_H^*} - w^*F^*K^*\right],$$

$$\Phi\big[Y - wFK\big] = \Phi^*\big[Y^* - w^*F^*K^*\big],$$

which corresponds to the expression in footnote 16. Lastly, note that in a closed economy $\Phi = \Phi^* = 0$, and therefore the country budget constraint (A10) becomes $Y = \bar{\mu}w[L - FK]$, as we have it in footnote 13.

Granular tariff   Consider firm-specific tariffs $\{\varsigma_{z,i}\}$ imposed by the home government on foreign firms i in sector z. In particular, if a foreign firm generates revenues $r_{z,i} = s_{z,i}\alpha_z Y$ in the home market, it needs to pay $\varsigma_{z,i}r_{z,i}$ to the home government, and takes home $(1 - \varsigma_{z,i})r_{z,i}$.

Then the foreign firm's profit maximization in the home market is:

$$\Pi_{z,i} = \max_{p_{z,i}} \left[(1 - \varsigma_{z,i})p_{z,i} - c_{z,i}\right]p_{z,i}^{-\sigma}\frac{\alpha_z Y}{\sum_{j=1}^{K_z} p_{z,j}^{1-\sigma}} - wF,$$

with the solution for prices and markups as if its costs were increased to $c'_{z,i} = c_{z,i}/(1 - \varsigma_{z,i})$, or equivalently productivity draw reduced to $\varphi'_{z,i} = \varphi_{z,i}(1 - \varsigma_{z,i})$. We denote the resulting market shares $\{s'_{z,i}\}$, and the resulting profits for foreign firms:

$$\Pi'_{z,i} = (1 - \varsigma_{z,i})\alpha_z Y \frac{s'_{z,i}}{\varepsilon(s'_{z,i})} - wF,$$

where $\varepsilon(s) = s + \sigma(1 - s)$ is as before.[53]

The expenditure on foreign goods in the home market is still given by $s'_{z,i}\alpha_z Y$, and the foreign share is still $\Lambda'_z = \sum_{i=1}^{K_z}(1 - \iota'_{z,i})s'_{z,i}$. However now, the home government collects $TR_z = \alpha_z Y \sum_{i=1}^{K_z}(1 - \iota'_{z,i})\varsigma_{z,i}s'_{z,i}$, while the rest $(\Lambda'_z\alpha_z Y - TR_z)$ is the revenue of foreign firms, which are split between production labor $\alpha_z Y \sum_{i=1}^{K_z}(1 - \iota'_{z,i})(1 - \varsigma_{z,i})\frac{s'_{z,i}}{\mu(s'_{z,i})}$, fixed costs $wF\sum_{i=1}^{K_z}(1 - \iota'_{z,i})$, and profits $\sum_{i=1}^{K_z}(1 - \iota'_{z,i})\Pi'_{z,i}$, where $\mu(s) = \frac{\varepsilon(s)}{\varepsilon(s)-1}$.

---

[53]Note that a non-uniform tax creates a computational challenge for the entry game, as the effective condition for entry becomes $\alpha_z Y \frac{s'_{z,i}}{\varepsilon(s'_{z,i})} \geq \frac{wF}{1-\varsigma_{z,i}}$, and ranking firms on $c'_{z,i}$ (and hence $s'_{z,i}$ does not guarantee monotonicity of $\Pi'_{z,i}$. We assume, however, that for a small enough $\varsigma_{z,i}$ (as is the case in our simulation), the approximation $F/(1 - \varsigma_{z,i}) \approx F$ is sufficient accurate in the entry game. Indeed, recall that entry is a discrete zero-one decision, in which most entering firms are inframarginal, with $\Pi'_{z,i} \gg 0$ due to the Zipf's law.

Therefore, there are changes to the three general equilibrium conditions (A10), (A11) and (A13). In particular, (A10) becomes:

$$Y = wL + \Pi + TR, \qquad \text{where} \qquad TR = Y \int_0^1 \alpha_z \left[\sum_{i=1}^{K_z'} (1 - \iota_{z,i}') \varsigma_{z,i} s_{z,i}' \right] \mathrm{d}z,$$

and where the profits of home firms Π is still expressed as in (A10). Foreign income (A11) is still $Y^* = w^* L^* + \Pi^*$, but now the profits from the home market need to be adjusted for tariffs:

$$\Pi^* = Y^*(1-\Lambda^*)\frac{\bar{\mu}_F^* - 1}{\bar{\mu}_F^*} - w^* F^* K_F^* + Y\Lambda\frac{\bar{\mu}_F - 1}{\bar{\mu}_F} - wFK_F - Y\int_0^1 \alpha_z \left[\sum_{i=1}^{K_z'} (1 - \iota_{z,i}')\frac{\varsigma_{z,i} s_{z,i}'}{\varepsilon(s_{z,i}')} \right] \mathrm{d}z.$$

Finally, the foreign labor market clearing (A13) also needs to be adjusted as follows:

$$w^* L^* = w^* F^* K^* + Y^*(1 - \Lambda^*)\frac{1}{\bar{\mu}_F^*} + Y\Lambda\frac{1}{\bar{\mu}_F} - Y\int_0^1 \alpha_z \left[\sum_{i=1}^{K_z'} (1 - \iota_{z,i}')\frac{\varsigma_{z,i} s_{z,i}'}{\mu(s_{z,i}')} \right] \mathrm{d}z.$$

Lastly, the current account balance (A14) becomes:

$$\Lambda Y - wFK_F - TR = Y^* \Lambda^* - w^* F^* K_H^*,$$

as now the foreign income from exporting is reduced by TR.

Labor reallocation   In an open granular economy, idiosyncratic firm productivity shocks affect sectoral comparative advantage and hence the allocation of production and employment across sectors.[54] We calculate the share of the total labor reallocation across firms (TLR) accounted for by the reallocation of labor across sectors (SLR). Formally, we define:

$$SLR = \frac{1}{2L} \int_0^1 \left| \Delta L_{z,t} \right| \mathrm{d}z \qquad \text{and} \qquad TLR = \frac{1}{2L} \int_0^1 \sum_{i=1}^{M_z} \left| \Delta L_{z,i,t} \right| \mathrm{d}z, \qquad \text{(A15)}$$

where $\Delta L_{z,i,t}$ is the annual employment change of firm i in sector z and $\Delta L_{z,t}$ is the annual employment change in sector z, and $|\cdot|$ takes the absolute values of these changes.[55] By construction $0 \leq SLR \leq TLR \leq 1$. TLR captures the overall extent of labor reallocation across firms in the economy and SLR captures the extent of the inter-sectoral labor reallocation. We report SLR, TLR and SLR/TLR in Appendix Table A2, for both annual and 10-year employment changes, and for different counterfactual values of trade openness.

---

[54]In contrast, in a continuous model, sectoral labor allocation is given by (A6), and just like trade patterns is not sensitive to idiosyncratic firm productivity shocks.

[55]SLR corresponds to the net employment change and TLR corresponds to the average of job creation and job destruction rates in Davis and Haltiwanger (1999). Note that $L_{z,t} = \sum_{i=1}^{M_z} L_{z,i,t}$ and $L = \int_0^1 L_{z,t}\mathrm{d}z$, and therefore $\int_0^1 \Delta L_{z,t}\mathrm{d}z = \int_0^1 \sum_{i=1}^{M_z} \Delta L_{z,i,t}\mathrm{d}z = 0$, and with a continuum of sectors SLR and TLR are constant over time. We calculate a home firm's domestic employment as $L_{z,i,t} = \frac{\alpha_z Y}{w}\frac{s_{z,i,t}}{\mu(s_{z,i,t})} + \frac{\alpha_z Y^*}{w}\frac{s_{z,i,t}^*}{\mu(s_{z,i,t}^*)} + F\mathbf{1}_{\{s_{z,i,t}>0\}}.$

C   Estimation Appendix

Detailed estimation procedure:

1. For given parameter values of μ_T and σ_T, we draw N relative sectoral productivities T_z from the log-normal distribution (recall our normalization $T_z^* \equiv 1$).[56] We keep the seed of all random draws constant throughout estimation.

2. For given values of parameter θ and realization of T_z in each sector $z = 1..N$, we draw productivities of potential entrants $\{\varphi_{z,j}\}_{j=1}^{M_z}$ in a manner consistent with the distributional assumptions of the model.[57] We obtain foreign productivity draws $\{\varphi_{z,j}^*\}_{j=1}^{M_z^*}$ in the same manner.

3. With the calibrated value of the relative wage rate w/w^* and normalization $w = 1$, and given the productivity draws and the remaining model parameters (σ, τ, F), we implement the following fixed point procedure:

    (i) Take an initial guess for (Y, Y^*), which completes the general equilibrium vector $\mathbf{X} = (w, w^*, Y, Y^*)$.

    (ii) Given $\mathbf{X}$, solve for sectoral equilibrium in each sector and each country, characterizing $\mathbf{Z} \equiv \left\{ K_z, \{s_{z,i}\}_{i=1}^{K_z}, K_z^*, \{s_{z,i}^*\}_{i=1}^{K_z^*} \right\}$, as described in Section 2.[58]

    (iii) Given $\mathbf{Z}$ and the normalization $L = 100$, use the general equilibrium conditions (9) and (11) to solve for the new values of Y and Y^*. Note that these equations are

---

[56]We use four replications of the French sectors. The simulated economy has therefore $N = 476$ $(= 119 \times 4)$ sectors in total. Without such replication, the model is sensitive to the randomness in the match between sectoral shares α_z and productivities T_z, while with four replications the effect of this randomness is already negligible. In each of the four replications, the empirical Cobb-Douglas shares and reshuffled randomly across sectors.

[57]Specifically, we follow EKS in using rank-order statistics for the Poisson-Pareto data generating process to directly draw the productivity of the most productive firm, which follows a $\mathrm{Frechet}(\theta, T_z)$ distribution, and each firm thereafter, with spacings following an exponential distribution. Specifically, denote $U_{z,j} \equiv T_z \varphi_{z,j}^{-\theta}$, where j is the rank of domestic firms in industry z. Eaton and Kortum (2010) show that $U_{z,1}$, $(U_{z,2} - U_{z,1})$, $(U_{z,3} - U_{z,2}), \ldots$ are iid exponential with cdf $G_U(u) = 1 - e^{-u}$. We use the transformation to convert the exponential draws into productivity draws $\{\varphi_{z,j}\}$. We draw enough shadow firms in each sector to ensure that the least productive ones never enter the market. Specifically, we use 5,000 firm draws by sector for France and 10,000 for ROW. For smaller sectors (in terms of Cobb-Douglas shares), we use 700 and 1,400 draws respectively, to reduce computing time. We check that with these number of draws, and over the relevant range of parameter values used in estimation, it is never the case that all shadow firms enter in any of the sectors.

[58]Solving for exact equilibrium values of K_z and K_z^* is computationally costly, therefore, we adopt the following approximation procedure. We solve for equilibrium $\hat{K}_z$ under the counterfactual assumption of constant markup equal to $\hat{\mu} = \sigma/(\sigma - 1)$, which is a simple analytical problem. It is easy to show that $\hat{K}_z$ is a lower bound for equilibrium K_z with variable markups (since from (7) equilibrium markups are strictly higher than $\hat{\mu}$, and hence price level is higher, yielding room for additional entry). We solve for oligopolistic equilibrium markups and market shares given $\hat{K}_z$. Given these markups for the first $\hat{K}_z$ firms, we then solve for additional entry ΔK_z, assuming the marginal entrants charge constant markup $\hat{\mu}$. We then set $K_z = \hat{K}_z + \Delta K_z$, and recalculate the oligopolistic equilibrium markups and market shares for this K_z. We check numerically that this procedure recovers a K_z which differs from the exact solution by at most one or two firms. Given that a typical French sector has over 300 firms, we view this approximation error as small.

linear in (Y, Y^*), and hence this is done by simple inversion.

(iv) Update the initial values of (Y, Y^*) taking a half step between the initial vector from step (i) and the new vector from step (iii), and loop over until convergence.

Upon convergence of (Y, Y^*), we use the foreign counterpart to labor market clearing condition (11) (namely, (A13)) to recover the value of L^*, which is consistent with the general equilibrium relative wage w/w^*, given parameter vector Θ.

(v) Upon convergence of the equilibrium vector $(\mathbf{X}, \mathbf{Z})$, simulate the model and calculate the moment vector $\mathcal{M}_z(\Theta)$ for all sectors $z = 1..N$, corresponding to parameter vector $\Theta = (\sigma, \theta, \tau, F, \mu_T, \sigma_T)$.

4. On a grid for parameters Θ with 20,000 points, evaluate the moment function $\mathcal{M}_z(\Theta)$, with moments described in Table 2, and the associated SMM loss function:

$$\mathcal{L}(\Theta) \equiv \left(\bar{\mathcal{M}}(\Theta) - \tilde{\mathbf{m}}\right)' \mathbf{W} \left(\bar{\mathcal{M}}(\Theta) - \tilde{\mathbf{m}}\right) = \mathbf{w}' \left(\bar{\mathcal{M}}(\Theta) - \tilde{\mathbf{m}}\right)^2,$$

where $\bar{\mathcal{M}}(\Theta) \equiv \frac{1}{N} \sum_{z=1}^{N} \mathcal{M}_z(\Theta)$, $\tilde{\mathbf{m}}$ are the values of the moments in our empirical dataset $\tilde{\Xi}$, and $\mathbf{W} = \text{diag}\{\mathbf{w}\}$ is the weighting matrix, which we chose to be diagonal and inversely proportional to $\tilde{\mathbf{m}}^2$.[59] Table 2 also reports the relative contribution of each moment k to the overall loss function, which with a diagonal weighting matrix is straightforward to calculate as $w_k \left(\bar{\mathcal{M}}_k(\hat{\Theta}) - \tilde{m}_k\right)^2 / \mathcal{L}(\hat{\Theta})$, where subindex k refers to the kth entry of the respective vector.

We use a Halton sequence to define the grid points, so that it covers the whole parameter space more efficienctly than if points were regularly spaced.

5. With the results from the first Halton grid, we recompute a second finer Halton grid of 20,000 points. We restrict this grid to be wide enough to ecompass the 50 best fitting parameter values of the previous grid, but exclude the regions with the highest loss function. We iterate this procedure five times. After five iterations, the procedure converges to a narrow region of the parameter space.

6. We take the best 20 of all the evaluated grid points, i.e. the ones that correspond to the lowest value of the loss function, and start local minimizers from each of them. We take as our estimate (the global minimizer) the point of local convergence with the lowest loss function, $\hat{\Theta} = \arg\min_{\Theta} \mathcal{L}(\Theta)$.

---

[59]We use this weighting to express the moment fit in percentage-deviation terms, apart for the first moment $\log \tilde{M}_z$, which is already in relative (log) terms (see Table 2 for the list of moments). For moments 14 and 15, with empirical values close to zero, $\mathbf{w}$ uses the values of the symmetric moments 12 and 13. Finally, we down-weight all standard deviation moments relative to the mean moments by a factor of 3, to emphasize the greater importance of matching the average patterns relative to the patterns of variation across sectors.

Standard errors (asymptotic inference)   We use the standard SMM asymptotics (as the number of sectors increases unboundedly) to calculate the standard errors for our estimator $\hat{\Theta}$. Rewrite the moment conditions as $\mathbb{E}m_i(\Theta) = 0$, where $m_i(\Theta) = \bar{\mathcal{M}}(\Theta) - \tilde{\mathbf{m}}_i$ is the moment function such that $\frac{1}{\tilde{N}} \sum_{i=1}^{\tilde{N}} m_i(\Theta) = \bar{\mathcal{M}}(\Theta) - \tilde{\mathbf{m}}$, where i correspond to one of $\tilde{N}$ sectors we observe in the data. With this, we express our SMM estimator $\hat{\Theta}$ minimizing $\mathcal{L}(\Theta)$ as a conventional extremum estimator:

$$\hat{\Theta} = \arg\min_{\Theta} \frac{1}{\tilde{N}} \sum_{i=1}^{\tilde{N}} m_i(\Theta)' \mathbf{W} \frac{1}{\tilde{N}} \sum_{i=1}^{\tilde{N}} m_i(\Theta).$$

Furthermore, note that $\bar{\mathcal{M}}(\Theta)$ are model-evaluated moments, which do not contribute to the sample variation in $m_i(\Theta)$.[60] Thus, all sample variation emerges from the empirically measured moments $\tilde{\mathbf{m}}_i$ over a finite sample of $\tilde{N}$ sectors. This gives rise to the standard errors of SMM estimation, which we compute according to the conventional asymptotic theory for an extremum estimator:

$$\sqrt{\tilde{N}} \cdot (\hat{\Theta} - \Theta) \to \mathcal{N}(0, V_\Theta) \quad \text{with} \quad V_\Theta \equiv (J'\mathbf{W}J)^{-1} J'\mathbf{W}H\mathbf{W}J(J'\mathbf{W}J)^{-1},$$

where V_Θ is the asymptotic sandwich-form variance matrix with $J = \mathbb{E}\left\{\frac{\partial m_i(\Theta)}{\partial \Theta}\right\}$ is the Jacobian and $H = \mathbb{E}\{m_i(\Theta)m_i(\Theta)'\}$ is the variance of moments, both in population under the true parameter vector Θ. Note that with our SMM moment structure, the effects of the data $\tilde{\mathbf{m}}$ and the model parameters Θ separate inside the moment function $m_i(\Theta)$, and hence the Jacobian J does not depend at all on the data. Hence we calculate J by numerical differentiation using the model-generated moment function $\bar{\mathcal{M}}(\Theta)$, evaluated around $\Theta = \hat{\Theta}$, that is $\hat{J} = \frac{\partial \bar{\mathcal{M}}(\hat{\Theta})}{\partial \Theta}$. The variance of moments matrix H depends on both $\bar{\mathcal{M}}(\hat{\Theta})$ and the data, and we calculate its estimate as follows:

$$\hat{H} = \frac{1}{\tilde{N}} \sum_{i=1}^{\tilde{N}} m_i(\hat{\Theta})m_i'(\hat{\Theta}) = \frac{1}{\tilde{N}} \sum_{i=1}^{\tilde{N}} \left(\bar{\mathcal{M}}(\hat{\Theta}) - \tilde{\mathbf{m}}_i\right)\left(\bar{\mathcal{M}}(\hat{\Theta}) - \tilde{\mathbf{m}}_i\right)'.$$

We combine $\hat{H}$ and $\hat{J}$, and the weighting matrix $\mathbf{W}$, to calculate the estimate of the variance matrix for the estimated parameter vector $\hat{\Theta}$:

$$\hat{V}_\Theta = \left(\hat{J}'\mathbf{W}\hat{J}\right)^{-1} \hat{J}'\mathbf{W}\hat{H}\mathbf{W}\hat{J}\left(\hat{J}'\mathbf{W}\hat{J}\right)^{-1}.$$

The standard errors for parameters in Table 2 are then calculated as $s.e.(\hat{\Theta}) = \sqrt{\text{diag}\left(\hat{V}_\Theta/\tilde{N}\right)}$.

---

[60]We simulate a sufficient number of sectors in the model, so that this assumption is indeed accurate.

Likelihood of granularity   We are interested in evaluating the probability that export share Λ_z^* in a given sector z, with a vector of additional observables $\mathbf{r}_z$, has a substantial granular component, namely $\Gamma_z^* \geq \vartheta \Lambda_z^*$ for some $\vartheta \in (0, 1)$. Using the model, with the estimated parameter vector $\hat{\Theta}$, this probability can be evaluated as follows:

$$\mathbb{P}_{\hat{\Theta}}\{\Gamma_z^* \geq \vartheta \Lambda_z^* | \Lambda_z^*, \mathbf{r}_z\} = \frac{\int_{\Lambda_z^* - \Phi^* \geq \vartheta \Lambda_z^*} g(\Phi_z^*, \Lambda_z^*, \mathbf{r}_z | \hat{\Theta}) \mathrm{d}\Phi_z^*}{\int_0^1 g(\Phi_z^*, \Lambda_z^*, \mathbf{r}_z | \hat{\Theta}) \mathrm{d}\Phi_z^*},$$

where $g(\cdot | \hat{\Theta})$ is the probability density of $(\Phi_z^*, \Lambda_z^*, \mathbf{r}_z)$ realizations across sectors in the estimated model, which we characterize numerically by simulating a large number of sectors using the estimated model. Note that we integrate over possible realizations of $\Phi_z^* = \Lambda_z^* - \Gamma_z^*$ using Bayes' rule, in particular that $g_{\Lambda, \mathbf{r} | \Phi}(\Lambda_z^*, \mathbf{r}_z | \Phi_z^*) g_\Phi(\Phi_z^*) = g(\Phi_z^*, \Lambda_z^*, \mathbf{r}_z)$, and where we omit the $\hat{\Theta}$ notation for brevity. Figure 3 plots the resulting probabilities $\mathbb{P}\{\Gamma_z^* \geq \vartheta \Lambda_z^* | \Lambda_z^*, \mathbf{r}_z\}$ for the case $\vartheta = 1/4$ and $\mathbf{r}_z = \tilde{s}_{z,1}$.

References

ACEMOGLU, D., V. M. CARVALHO, A. OZDAGLAR, AND A. TAHBAZ-SALEHI (2012): "The Network Origins of Aggregate Fluctuations," *Econometrica*, 80(5), 1977–2016.

AMITI, M., O. ITSKHOKI, AND J. KONINGS (2014): "Importers, Exporters, and Exchange Rate Disconnect," *American Economic Review*, 7(104), 1942–1978.

———— (2018): "International Shocks, Variable Markups and Domestic Prices," http://www.princeton.edu/~itskhoki/papers/DomesticPrices.pdf.

ANDERSON, J. E., AND E. VAN WINCOOP (2004): "Trade Costs," *Journal of Economic Literature*, 42(3), 691–751.

ANDREWS, I., M. GENTZKOW, AND J. M. SHAPIRO (2017): "Measuring the Sensitivity of Parameter Estimates to Estimation Moments*," *The Quarterly Journal of Economics*, 132(4), 1553–1592.

ARKOLAKIS, C., A. COSTINOT, AND A. RODRÍGUEZ-CLARE (2012): "New Trade Models, Same Old Gains?," *American Economic Review*, 102(1), 94–130.

ATKESON, A., AND A. BURSTEIN (2008): "Trade Costs, Pricing-to-Market, and International Relative Prices," *American Economic Review*, 98(5), 1998–2031.

AUTOR, D., D. DORN, L. F. KATZ, C. PATTERSON, AND J. VAN REENEN (2017): "The Fall of the Labor Share and the Rise of Superstar Firms," NBER Working Papers No. 23396.

BAGWELL, K., AND R. W. STAIGER (2004): *The Economics of the World Trading System*. The MIT Press.

BERNARD, A. B., J. EATON, J. B. JENSEN, AND S. KORTUM (2003): "Plants and Productivity in International Trade," *American Economic Review*, 93(4), 1268–1290.

BERNARD, A. B., J. B. JENSEN, S. J. REDDING, AND P. K. SCHOTT (2018): "Global Firms," *Journal of Economic Literature*, 56(2), 565–619.

BLONIGEN, B. A., AND T. J. PRUSA (2008): "Antidumping," in *Handbook of International Trade*, ed. by E. Choi, and J. Harrigan, chap. 9, pp. 251–284. Wiley-Blackwell.

BRANDER, J. A., AND B. J. SPENCER (1984): "Trade warfare: Tariffs and cartels," *Journal of International Economics*, 16(3–4), 227–242.

BRODA, C., AND D. WEINSTEIN (2006): "Globalization and the Gains from Variety," *Quarterly Journal of Economics*, 121(2), 541–85.

CARVALHO, V. M., AND X. GABAIX (2013): "The Great Diversification and Its Undoing," *American Economic Review*, 103(5), 1697–1727.

CARVALHO, V. M., AND B. GRASSI (2014): "Firm Dynamics and the Granular Hypothesis," http://www.crei.cat/~carvalho/research.html.

CHANEY, T. (2008): "Distorted Gravity: The Intensive and Extensive Margins of International Trade," *American Economic Review*, 98(4), 1707–1721.

CHOR, D. (2010): "Unpacking sources of comparative advantage: A quantitative approach,"

Journal of International Economics, 82(2), 152–167.

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

Costinot, A., A. Rodríguez-Clare, and I. Werning (2016): "Micro to Macro: Optimal Trade Policy with Firm Heterogeneity," NBER Working Papers No. 21989.

Davis, S. J., and J. Haltiwanger (1999): "Gross job flows," in *Handbook of Labor Economics*, ed. by O. Ashenfelter, and D. Card, vol. 3, chap. 41, pp. 2711–2805. Elsevier.

di Giovanni, J., and A. A. Levchenko (2012): "Country Size, International Trade, and Aggregate Fluctuations in Granular Economies," *Journal of Political Economy*, 120(6), 1083–1132.

——— (2013): "Firm entry, trade, and welfare in Zipf's world," *Journal of International Economics*, 89(2), 283–296.

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

Dixit, A. (1984): "International Trade Policy for Oligopolistic Industries," *Economic Journal*, 94(376a), 1–16.

Dornbusch, R., S. Fischer, and P. A. Samuelson (1977): "Comparative Advantage, Trade, and Payments in a Ricardian Model with a Continuum of Goods," *American Economic Review*, 67(5), 823–839.

Eaton, J., and G. M. Grossman (1986): "Optimal Trade and Industrial Policy Under Oligopoly," *The Quarterly Journal of Economics*, 101(2), 383–406.

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

——— (2010): *Technology in the Global Economy: A Framework for Quantitative Analysis*. unpublished manuscript.

Eaton, J., S. Kortum, and S. Sotelo (2012): "International Trade: Linking Micro and Macro," NBER Working Papers No. 17864.

Edmond, C., V. Midrigan, and D. Y. Xu (2015): "Competition, markups, and the gains from international trade," *The American Economic Review*, 105(10), 3183–3221.

Fan, H., E. L.-C. Lai, and H. S. Qi (2015): "Heterogeneous Firms and Trade Liberalization: Theory and Evidence," http://ihome.ust.hk/~elai/.

Freund, C., and M. D. Pierola (2015): "Export superstars," *Review of Economics and Statistics*, 97(5), 1023–1032.

Gabaix, X. (2009): "Power Laws in Economics and Finance," *Annual Review of Economics*, 1(1), 255–294.

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

GABAIX, X., AND R. IBRAGIMOV (2011): "Rank - 1 / 2: A Simple Way to Improve the OLS Estimation of Tail Exponents," *Journal of Business and Economic Statistics*, 29(1), 24–39.

GROSSMAN, G. M., AND E. ROSSI-HANSBERG (2010): "External Economies and International Trade Redux," *The Quarterly Journal of Economics*, 125(2), 829–858.

GUTIÉRREZ, G., AND T. PHILIPPON (2017): "Declining Competition and Investment in the U.S.," NBER Working Papers No. 23583.

HANSON, G. H., N. LIND, AND M.-A. MUENDLER (2016): "The Dynamics of Comparative Advantage," https://gps.ucsd.edu/faculty-directory/gordon-hanson.html.

HEAD, K., R. JING, AND J. RIES (2017): "Import sourcing of Chinese cities: Order versus randomness," *Journal of International Economics*, 105(C), 119–129.

HELPMAN, E., AND P. KRUGMAN (1989): *Trade Policy and Market Structure*. The MIT Press.

HORN, H., AND J. LEVINSOHN (2001): "Merger Policies and Trade Liberalisation," *The Economic Journal*, 111(April), 244–276.

HORTACSU, A., AND C. SYVERSON (2014): "Product Differentiation, Search Costs, and the Welfare Effects of Entry: A Case Study of S&P 500 Index Funds," *Quarterly Journal of Economics*, 119(4), 403–456.

HOTTMAN, C., S. J. REDDING, AND D. E. WEINSTEIN (2015): "Quantifying the Sources of Firm Heterogeneity," *Quarterly Journal of Economics*, forthcoming.

LEVCHENKO, A. A., AND J. ZHANG (2016): "The evolution of comparative advantage: Measurement and welfare implications," *Journal of Monetary Economics*, 78(C), 96–111.

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

NAVARETTI, G. B., M. BUGAMELLI, E. FORLANI, AND G. I. P. OTTAVIANO (2016): "It takes (more than) a moment: Revisiting the link between firm productivity and aggregate exports," http://www.lse.ac.uk/economics/people/faculty/gianmarco-ottaviano.

NEARY, J. P. (2003): "Presidential Address: Globalization and Market Structure," *Journal of the European Economic Association*, 1(2-3), 245–271.

——— (2012): "Superstar Firms and Trade," Nottingham Lectures in International Economics.

——— (2015): "International Trade in General Oligopolistic Equilibrium," Economics series working papers, University of Oxford, Department of Economics.

OKUBO, T. (2009): "Firm heterogeneity and Ricardian comparative advantage within and across sectors," *Economic Theory*, 38(3), 533–559.

RODRIK, D. (1998): "Why do More Open Economies Have Bigger Governments?," *Journal of Political Economy*, 106(5), 997–1032.

SUTTON, J., AND D. TREFLER (2016): "Capabilities, Wealth, and Trade," *Journal of Political Economy*, 124(3), 826–878.