Granular Comparative Advantage

Cecile Gaubert

University of California Berkeley

Oleg Itskhoki

University of California Los Angeles

Large firms play a pivotal role in international trade. We develop a multisector granular model of trade where sectors host a finite number of firms. Both aggregate and granular firm-level forces shape comparative advantage. The model is estimated using French microdata on firm domestic and export sales. We find that granularity accounts for about 20% of the sectoral variation in export intensity and is more pronounced in highly export-intensive sectors. An extension to a dynamic environment with both idiosyncratic and aggregate shocks reveals that firm dynamics plays a central role in shaping comparative advantage reversals observed in the data.

Correction: This article was reposted on February 18, 2021, with two corrections. In subsection IV.A.4, in the display equation in the first paragraph, the exponent in the numerator of the fraction was changed to $-\sigma$. *In subsection IV.A.5, in the fourth sentence of the second paragraph ("Furthermore . . ."), the inequality after "that is" was changed from less than to greater than.*

We thank Ariel Burstein, Lorenzo Caliendo, Francis Kramarz, Andrei Levchenko, Mathieu Parenti, and Jon Vogel for insightful discussions; David Atkin, Kyle Bagwell, Kamran Bilir, Matilde Bombardini, Vasco Carvalho, Jonathan Eaton, Xavier Gabaix, Gene Grossman, Sam Kortum, Ulrich Müller, Ezra Oberfield, Andrés Rodríguez-Clare, Esteban Rossi-Hansberg, Jesse Shapiro, Chris Sims, Sebastian Sotelo, Bob Staiger, Mark Watson, and seminar and conference participants for insightful comments; Joe Abadi, Dima Mukhin, Mark Razhev, and in particular Max Vogler for excellent research assistance; and Davin Chor for sharing his data. Gaubert gratefully acknowledges financial support from the Clausen Center at University of California Berkeley. Data are provided as supplementary material online.

Electronically published January 25, 2021
[*Journal of Political Economy*, 2021, vol. 129, no. 3]

I. 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 countries. We aim to measure what part of comparative advantage can be traced to characteristics that are common to all firms in a given sector—such as the availability of specific human capital, infrastructure, and technology—versus idiosyncratic contribution of individual firms, driven by their specific know-how and managerial talent. We call the former fundamental comparative advantage and the latter granular comparative advantage. We anchor our analysis around a multisector model with a finite number of firms operating in each sector, whereby very large firms have the potential to shape sectoral outcomes. This approach is in stark contrast to much of the literature that assumes a continuum of infinitesimal firms, leaving no room for individual firms to affect sectoral aggregates.

Using this conceptual framework, we set out to measure the contribution of individual firms to international trade flows, as summarized by

[1] In their paper "Export Superstars," 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 data set. In the French manufacturing data set used in this paper, the largest firm accounts for 7% of all manufacturing exports, and within four-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 in 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 gross domestic product (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 against Microsoft and the 2017 fine imposed by the European Commission on Google. A very recent case of a granular trade war is the 292% tariff imposed by the United States 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 United States against the Chinese Zhongxing Telecommunications Equipment). For a recent theoretical and empirical analysis of granular international lobbying, see Blanga-Gubbay, Conconi, and Parenti (2020).

granular comparative advantage, thus revisiting the fundamental questions in international trade: what goods do countries trade and what is the source of a country's comparative advantage? The decomposition between fundamental comparative advantage and granular comparative advantage is perhaps best illustrated in terms of counterfactuals. Suppose, for instance, that a given firm and its technological know-how disappear. How does the export stance of the sector change? If comparative advantage is shaped by only sector-level characteristics, it does not change—other domestic firms in the sector expand, or enter, to absorb the market share of the exiting firm. This is the neoclassical benchmark, which abstracts from individual firms altogether and focuses on sectoral technologies and supplies of factors. However, if comparative advantage is in part driven by the performance of individual firms, then the export stance of the sector will change as the firm disappears with its specific strengths. In this world, the firm's market share is taken over by other domestic firms that contribute differently to export patterns or even by foreign firms. More generally, fundamental comparative advantage and granular comparative advantage have different implications for the evolution of comparative advantage over time. When comparative advantage comes in part from granular comparative advantage, standard firm dynamics—whereby individual firms gain and lose market shares against other domestic firms—result in changing aggregate sectoral exports, even without sectoral shocks.

We begin in section II by formalizing the concept of granular comparative advantage and discussing possible approaches to its identification in the data. In section III, we then present suggestive empirical evidence that granularity may be at play in shaping sectoral trade flows. To proxy for sectoral granularity, we adopt a measure of concentration of domestic sales among domestic firms. This measure identifies sectors with unusually large home firms without being directly affected by the international competitiveness of the sector. We show that this measure of within-sector domestic concentration is systematically correlated with aggregate sectoral exports both in the cross section and in the time series. Furthermore, it is also predictive of future changes in sectoral exports out of sample: granular sectors (with stronger concentration at the top) have a greater tendency to see their exports mean revert—an empirical regularity consistent with a granular firm dynamics model, as we explore below.

To go further and quantify the importance of granular comparative advantage, we adopt a structural identification approach. In section IV, we develop and characterize a model of granular trade, which we later quantify and use for counterfactual analysis. Our model of granular trade contrasts with the bulk of the international trade literature, which maintains the assumption that sectors are comprised of a continuum of heterogeneous firms, that is, that every firm is 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 (as demonstrated in Arkolakis, Costinot, and Rodríguez-Clare 2012).

We propose an alternative multisector granular model of trade, which acknowledges a finite number of firms operating in each sector, with the largest firm often claiming a massive share of the market. In the model, each firm draws a productivity realization from a sector-specific distribution. Given that there is a finite number of firms in each sector, with some of them very large, 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. Formally, our model combines Ricardian comparative advantage across sectors, as in Dornbusch, Fischer, and Samuelson (1977; henceforth DFS), with the Melitz (2003) model of firm heterogeneity within sectors, in which we relax the assumption of a continuum of firms, following Eaton, Kortum, and Sotelo (2012).[4] This allows the model to simultaneously nest fundamental and granular comparative advantage in a unified framework.

We estimate the model in section V, using firm-level data on domestic and export sales of French firms across 119 four-digit manufacturing industries. We use a simulated method of moments (SMM) estimation procedure, which takes full account of the general equilibrium of our granular model. The extent of firm concentration—hence the potential for granularity—is disciplined by targeting moments on the number of firms and the market share of top firms across industries. Comparative advantage is revealed using export and import intensity of sectors. To disentangle the relative roles of fundamental and granular forces in driving comparative advantage, we use moments of the joint distribution of sectoral trade flows and within-sector domestic sales concentration. Intuitively, sectors in which export intensity is high because of granular comparative advantage are expected to feature firm outliers relative to other domestic firms, large enough to drive aggregate sectoral productivity. Despite its parsimony, the estimated granular model is successful at reproducing the rich cross-sectoral heterogeneity of the data.

We use the estimated model in section VI to quantify the importance of granularity in shaping sectoral trade outcomes, using counterfactual analysis. We find a significant part of trade flows to be of a granular origin (around 20%) and that the contribution of granularity is particularly pronounced in the most export-intensive sectors—the export champions of the country. Among the top 10% of export-intensive sectors, our results

[4] Specifically, a random integer number of firms draw productivities from a fat-tailed Pareto distribution, with a mean productivity parameter that varies by country and sector, capturing fundamental Ricardian forces.

suggest that nearly one-third of exports is of granular origin. We also show that in a granular world, conventional inference of fundamental sectoral productivities based on sectoral export shares leads to biased estimates.

Importantly, we establish the robustness of our results to alternative parameterizations of the model and distributional assumptions. We find that parameterizations that still match our key identifying moments lead to similar quantitative conclusions on the importance of granularity. Conversely, alternative parameterizations, such as ones with a thinner-tailed firm productivity distribution, are unable to match the data. This emphasizes the importance of the empirical moments, relative to the assumed functional forms, in delivering the identification in our structural framework.

Having established the contribution of granularity to long-run trade patterns using our static model, we explore in section VII its dynamic implications. In particular, we study the ability of granular forces to explain and predict the evolution and reversals in comparative advantage of countries. To this end, we extend our granular model to feature industry dynamics, driven simultaneously by sector-level and firm-level productivity shocks. We calibrate the dynamic productivity processes to match the mean reversion of domestic firm sales shares and sectoral exports. We find that granular forces account for 25% of the year-to-year changes in sectoral export shares. Furthermore, the dynamic granular model is consistent with both the hyperspecialization of countries in a few industries at any given point in time and a relatively fast mean reversion in comparative advantage over time, emphasized in a recent paper by Hanson, Lind, and Muendler (2018). We find that idiosyncratic firm productivity dynamics alone accounts for about one-half of comparative advantage reversals for the most export-intensive sectors. Finally, we show that the dynamic model captures accurately the empirical patterns we documented in section III and, in particular, that empirical proxies of sectoral granularity are predictors of future mean reversion in sectoral exports. Thus, we view these empirical patterns as strongly suggestive of the granular mechanism of the type modeled in this paper. We conclude the paper with a discussion of granular policies in section VIII and final remarks in section IX.

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 et al. 2012; Carvalho and Gabaix 2013; Grassi 2017; Carvalho and Grassi 2020).[5] 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).

[5] Building on these insights, Gabaix and Koijen (2020) propose a way to construct a granular instrumental variable, which can be used to isolate the causal effects of aggregate economic shocks.

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 Eaton, Kortum, and Sotelo (2012). They 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 multisector 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; Levchenko and Zhang 2016; Sutton and Trefler 2016; Hanson, Lind, and Muendler 2018).

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; Amiti, Itskhoki, and Konings 2014, 2019; Edmond, Midrigan, and Xu 2015; Hottman, Redding, and Weinstein 2016). Grassi (2017) also studies oligopolistic competition in a granular setting. We follow Neary (2010, 2016) 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 et al. (2018) for a recent review of the empirical and theoretical literature on the role of individual firms in international trade. Our study is also related to the vast literature on trade policy and market structure, summarized in Helpman and Krugman (1989) and Bagwell and Staiger (2004).

II. Granular Comparative Advantage

In order to quantify fundamental and granular comparative advantage empirically, we introduce a formal definition of these concepts and discuss possible approaches to their identification.

A. *Definition*

Consider the export intensity of a sector z that is the ratio of sectoral exports X_z to domestic expenditure (absorption) Y_z. We denote it by $\Lambda_z \equiv X_z/Y_z$. It is an intuitive measure of comparative advantage that maps into

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

alternative definitions across a range of international trade models.[7] Mechanically, it can be decomposed into the sum of the contributions to exports of all firms in the sector:

$$\Lambda_z = \sum_{i=1}^{N_z} s_{z,i}\lambda_{z,i} = \mathbf{s}_z'\boldsymbol{\lambda}_z,$$

where N_z is the number of home firms in the sector, $s_{z,i} \equiv d_{z,i}/Y_z$ is the firm-level domestic market share, $\lambda_{z,i} \equiv x_{z,i}/d_{z,i}$ is the firm-level export intensity (the ratio of firm exports to domestic sales), and $(\boldsymbol{\lambda}_z, \mathbf{s}_z)$ is the corresponding vector notation. We treat the observed market shares and export intensities in sector z as a realization of a stochastic data-generating process: $(\boldsymbol{\lambda}_z, \mathbf{s}_z) \sim \mathcal{F}_z(\cdot)$.[8] The distribution function $\mathcal{F}_z$ embodies the characteristics of sector z in a given country that are a priori accessible to any firm in the sector.

We are interested in decomposing the aggregate outcome Λ_z, itself a random variable, into its expected value on the basis of sectoral characteristics and the contribution of idiosyncratic firm draws around this expected value:

$$\Lambda_z = \Phi_z + \Gamma_z, \text{ where } \Phi_z \equiv \mathbb{E}_z\{\Lambda_z\} = \int (\mathbf{s}'\boldsymbol{\lambda})\mathrm{d}\mathcal{F}_z(\boldsymbol{\lambda}, \mathbf{s}). \tag{1}$$

Thus, we define *fundamental comparative advantage*, denoted Φ_z, as the population mean of the export intensity given the sectoral characteristics embodied in $\mathcal{F}_z$. In turn, *granular comparative advantage*, denoted $\Gamma_z \equiv \Lambda_z - \Phi_z$, is a granular residual that captures departures from the population mean for the observed realization of a sector.

Conventional theoretical models assume a continuum of (heterogeneous) firms and hence feature no granular residual ($\Gamma_z \equiv 0$). Indeed, the law of large numbers applies within sectors, and therefore all firm-level idiosyncrasies wash away in the aggregate. This knife-edge result disappears as soon as the model acknowledges that the actual number of firms N_z is a finite integer (in general, a random variable determined by the data-generating process $\mathcal{F}_z$), so that firms are not infinitesimal. Under these circumstances, idiosyncratic firm outcomes have the potential to

[7] Throughout the paper, we study cross-sectional variation of Λ_z, which maps one for one with variation in comparative advantage (see sec. IV.B). Formally, measures of comparative advantage usually take double ratios of exports across sectors and destinations (see, e.g., Costinot, Donaldson, and Komunjer 2012). Here, Λ_z is a simple ratio, which delivers the variations in comparative advantage in which we are interested, given that our analysis is cross-sectional within a country and hence holds aggregate quantities constant.

[8] Note the relationship to the macroeconomics granularity literature, which focuses on the decomposition of the aggregate growth rate G_t into idiosyncratic growth rates $g_{i,t}$ (in place of the sectoral and firm-level export intensities, Λ_z and $\lambda_{z,i}$, respectively) and usually treats market shares $s_{i,t}$ as exogenously given.

be detectable after aggregation and, in this sense, influence sectoral outcomes: sector-level realizations differ from their expected values, and this difference is captured by the granular residual.

A few further remarks are in order. First, fundamental comparative advantage Φ_z and granular comparative advantage Γ_z are well defined, in general, as the integral in equation (1) cumulates market shares that are bounded by construction.[9] Note that market shares are well behaved even when the underlying firm sales distribution is fat tailed and has unbounded mean and/or variance. Second, the population mean Φ_z is not, in general, equivalent to the asymptotic value of Λ_z as $N_z \to \infty$.[10] Consequently, our inference approach below builds on repeated observations of Λ_z across sectors with a finite number of firms rather than on hypothetical sequences of sectors with N_z increasing toward infinity. Third, nontrivial differences between Λ_z and Φ_z, which constitute the granular residual Γ_z, may arise both in small sectors with few firms (e.g., $N_z \approx 10$) as well as in sectors with a large number of firms ($N_z > 100$ or even $>1{,}000$), so long as the distribution of market shares $s_{z,i}$ is sufficiently skewed. This latter possibility constitutes the empirically relevant case on which we focus.

Last, note that the granular residual Γ_z is mean zero by construction as a difference between a random variable Λ_z and its conditional mean Φ_z. As a consequence, fundamental comparative advantage Φ_z and granular comparative advantage Γ_z are orthogonal in the cross section of sectors z, and the following variance decomposition of export intensity across sectors holds:

$$\mathrm{var}(\Lambda_z) = \mathrm{var}(\Phi_z) + \mathrm{var}(\Gamma_z).$$

One goal of this paper is to quantify this variance decomposition and hence the respective contribution of granular and fundamental forces to comparative advantage both in levels and in changes over time.

B. Identification

If one observed many realizations of a sector drawn from a given data-generating process $\mathcal{F}_z$, Φ_z could be estimated by the average (across realizations) export intensity of the sector, and then the granular residual

[9] Note that $s_{z,i}\lambda_{z,i} = x_{z,i}/Y_z = s^*_{z,i} \cdot (Y^*_z/Y_z)$, where $s^*_{z,i} = x_{z,i}/Y^*_z$ is the foreign market share of a domestic exporter, which is bounded between 0 and 1. We make the mild assumption that the relative sectoral aggregates across countries, Y^*_z/Y_z, are bounded and have a finite expectation.

[10] Nevertheless, the structural model of sec. IV has the convenient property that Φ_z is not only the population mean of Λ_z with finite N_z but also its continuous limit as the number of firms N_z becomes infinite, with an appropriately chosen decreasing sequence of fixed entry costs.

Γ_z could be measured realization by realization. Thus, in principle, one could make an inference on sectoral fundamental comparative advantage and granular comparative advantage directly from the data. In practice, however, such inference is rarely feasible, as $\mathcal{F}_z$ likely varies both in the cross section of country-sectors as well as within the country-sector over time ($\mathcal{F}_{z,t}$ in this case). Separating small-sample deviations from changes in fundamentals $\mathcal{F}_z$ across sectors or over time is challenging empirically, absent a natural experiment that shifts idiosyncratic firm productivities without affecting fundamental sectoral characteristics.[11]

To make progress on this issue, we adopt a structural approach. That is, we reduce the dimensionality of the problem by parameterizing the data-generating process $\mathcal{F}_z$ and its variation across sectors. Specifically, we use a general equilibrium economic model for market shares and export intensities, $\{\lambda_{z,i}, s_{z,i}\}$, as described in section IV. Our model is closely related to workhorse international trade models, with the important difference that we relax the assumption that firms in each sector form a continuum, thereby allowing individual firms to impact sectoral outcomes. We estimate the model using moments from the data that summarize variation across sectors in both sector-level and firm-level outcomes. Once estimated, the model can be simulated any number of times with the same data-generating process, which allows us to quantify the relative importance of granular and fundamental forces as discussed above as well as to carry out static and dynamic counterfactuals. Importantly, we show that this quantification is robust to a set of alternative parametric assumptions, so long as the model matches the set of identifying moments in the data.

III. Empirical Regularities

Before delving into the structural framework in section IV, we look at the data for patterns that are suggestive of the granular mechanism. The empirical regularities documented here are then used in the structural estimation of section V to help identify the model parameters that govern the relative intensity of granular and fundamental comparative advantage as well as in the dynamic analysis of section VII.

If granular forces are at play, one would expect that the presence of a few unusually large firms within a sector (relative to other domestic firms) correlates with the aggregate export intensity of the sector. To detect such patterns, we regress log sectoral exports on the concentration

[11] Possible empirical strategies along these lines could be, e.g., to take the death of a chief executive officer as a shock to a firm (see, e.g., Bennedsen, Perez-Gonzalez, and Wolfenzon 2010), but these strategies appear too limited in scope to conduct the analysis systematically, as we set out to do here.

ratio of domestic sales among the top three domestic firms, controlling for the size of the domestic market:

$$\log X_z = \alpha + \beta \sum_{i=1}^{3} \tilde{s}_{z,(i)} + \log D_z + \varepsilon_z, \qquad (2)$$

where D_z are total domestic sales of all domestic firms in sector z and $\tilde{s}_{z,(i)} \equiv d_{z,(i)}/D_z$ is the relative domestic sales share of the ith largest firm in the domestic market. We use the top three concentration ratio, $\Sigma_{i=1}^{3} \tilde{s}_{z,(i)}$, as our baseline proxy for sectoral granularity.[12] Note that this measure is not mechanically correlated with comparative advantage, as it relies solely on the relative sales of domestic firms in the domestic market, in particular relative to other domestic firms.

To estimate equation (2), we use French firm-level manufacturing data on domestic sales and exports, $\{d_{z,i}, x_{z,i}\}$. We aggregate firm-level data to obtain sectoral domestic sales D_z and sectoral exports X_z. The data include 300 four-digit sectors according to the Statistical Classification of Economic Activities in the European Community (Nomenclature Statistique des Activités Économiques dans la Communauté Européenne [NACE]), with an average of 290 French firms per sector.[13]

We report the results of several specifications in table 1. Columns 1 and 2 use the 2005 cross section (our benchmark year for estimation) and show that a 10 percentage point greater top three sales share in the domestic market is associated with a 9% (log points) increase in aggregate sectoral exports. This relationship holds controlling for two-digit sectoral fixed effects (col. 2). This relationship also holds more generally in the panel between 1997 and 2007, with year fixed effects and with and without two-digit sectoral fixed effects (cols. 3, 4). The resulting estimates are almost the same as in the 2005 cross section. In all cases, we control for the size of the sectors, as sectors with greater total domestic sales tend to have greater total exports with nearly a unitary elasticity.

These cross-sectional results are consistent with granularity shaping—in part, trade patterns—as the relative size of the largest firms within sectors is predictive of aggregate sectoral exports. However, they leave room for alternative interpretations. For instance, it could be that a greater within-sector productivity dispersion is associated with a greater sectoral

[12] In table A1, we report all empirical results using top firm sales share $\tilde{s}_{z,(1)}$ instead of $\Sigma_{i=1}^{3} \tilde{s}_{z,(i)}$, resulting in the same quantitative patterns with somewhat less precisely estimated coefficients.

[13] These data do not contain information on imports. In our structural estimation, in sec. V, we merge these data with Comtrade to ensure that the estimated model matches simultaneously the corresponding patterns for both exports and imports. Because of different sectoral definitions between data sets, the matching can be done only at a higher level of aggregation, leaving us with 119 sectors. The qualitative patterns we focus on here are robust to this aggregation.

TABLE 1
GRANULARITY AND EXPORTS

	Cross Section, 2005		Panel, 1997–2007		Dynamic Regressions		
$\log X_z$	(1)	(2)	(3)	(4)	(5)	(6)	(7)
$\Sigma_{i=1}^{3} \tilde{s}_{z,(i)}$	.860***	.908***	.903***	.918***	.405***	.495***	.493***
	(.297)	(.303)	(.308)	(.309)	(.131)	(.137)	(.138)
$\log D_z$	.897***	.938***	.907***	.952***			
	(.051)	(.053)	(.052)	(.054)			
Sector fixed effects		2-digit		2-digit	4-digit		2-digit
Year fixed effects			Yes	Yes	Yes		
Observations	300	300	3,300	3,300	3,300	3,000	3,000
R^2_{adj}	.518	.620	.520	.649	.950	.009	.007

NOTE.—Data and variables are described in the text, with standard errors in parentheses. Specification 5 contains a full set of 300 sectoral fixed effects and uses time series variation over 1997–2007; specifications 6 and 7 are in first differences, namely, regressing $\Delta \log X_{z,t}$ on $\Delta(\Sigma_{i=1}^{3} \tilde{s}_{z,(i),t})$ over 1998–2007.

*** Significant at the 1% level.

trade share because of standard Melitz (2003) selection forces or that sectors with a high dispersion in productivity across firms happen to be those in which France has a comparative advantage (for a microfoundation, see, e.g., Bonfiglioli, Crinò, and Gancia 2018). We partly address this concern by controlling for two-digit fixed effects aimed to group together sectors with similar technological and skill requirements.

To explore this further, we run three specifications that exploit within-sector variation over time. Column 5 reruns the fixed effects panel specification from column 4, replacing the two-digit industry fixed effects with a full set of four-digit industry fixed effects, which in particular absorb sectoral characteristics that are stable over time. Columns 6 and 7 estimate the specification in first differences, regressing the change in log sectoral exports $\Delta \log X_{z,t} \equiv \log X_{z,t} - \log X_{z,t-1}$ on the change in the top three share $\Delta(\Sigma_{i=1}^{3} \tilde{s}_{z,(i),t})$ year to year in our panel. We again find a notable statistical association between the granular proxy and aggregate sectoral exports, now in the time series dimension. Sectors where top French firms increase their relative dominance in the domestic market also see an improvement in their aggregate export stance: a 10 percentage point increase in the top three concentration ratio is associated with a 5% (log points) increase in the sectoral exports.

Predictive regressions.—Finally, we ask whether empirical proxies of granularity can predict future time series evolution of sectoral trade patterns. In other words, beyond the time series comovement we documented earlier, we investigate whether micro-level measures of sectoral granularity have predictive power for the evolution (out of sample) of aggregate trade flows. Table 2 reports the results. It uses the 1997 cross section of the French firm-level data to predict the evolution of sectoral trade flows over the next 10 years, up to 2007. Column 1 establishes the presence of (partial) mean reversion in export patterns: a sector with greater exports X_z, controlling for its size (domestic sales, D_z), is expected to export less over time. Specifically, a sector that exported 10% (log points) more in 1997 exhibited a 1.3% (log points) decline in its exports over the next 10 years.

Columns 2 and 3 further show that greater firm concentration at the top in the initial year is associated with a stronger mean reversion of sectoral exports over time. We find that a 10 percentage point higher top three share in 1997 predicts a 5.5% (log points) decline in sectoral exports over the next 10 years, or a corresponding 4.2% decline after controlling for the initial export level. These effects are large quantitatively.

The results in columns 1–3 of table 2 suggest that it is not just high exports that predict future reversals in sectoral comparative advantage but rather exports that are high for nonfundamental granular reasons. We verify this hypothesis more formally in column 4, where we consider a two-stage regression, first projecting log sectoral exports on the top three

TABLE 2
MEAN REVERSION IN EXPORTS

$\log X_{z,+10} - \log X_z$	OLS (1)	OLS (2)	OLS (3)	TWO STAGE (4)
$\log X_z$	−.130***		−.111***	−.468***
	(.039)		(.040)	(.183)
$\Sigma_{i=1}^{3} \tilde{s}_{z,(i)}$		−.551***	−.421***	
		(.200)	(.203)	
$\log D_z$	.118**	−.044	.067	.424**
	(.048)	(.037)	(.054)	(.169)
Sector fixed effects	2-digit	2-digit	2-digit	2-digit
Observations	300	300	300	300
R^2_{adj}	.082	.070	.093	. . .

NOTE.—Data and variables are as in table 1. Column 4 reports the second-stage regression, where in the first stage $\log X_z$ is projected on $\Sigma_{i=1}^{3} \tilde{s}_{z,(i)}$ in the initial year (1997), similar to specification 2 in table 1.
** Significant at the 5% level.
*** Significant at the 1% level.

sales share (as in col. 2 of table 1) and then regressing the change in log exports over the next 10 years on the predicted value for initial sectoral exports. This magnifies the coefficient on log exports nearly fourfold compared with the ordinary least squares (OLS) specification—from −0.13 in column 1 to −0.47 in column 4—emphasizing a much stronger pattern of mean reversion in exports in sectors that export for granular reasons. We view these results as strongly suggestive reduced-form evidence of the granular mechanism, which we explore structurally in the rest of the paper.[14]

IV. A Granular Model of Trade

This section sets up the granular trade model, giving rise to the structural data-generating process discussed in section III. We use this model to explore the quantitative implications of granularity for trade patterns. It is a two-country multisector model, which features Ricardian DFS forces across sectors and an Eaton, Kortum, and Sotelo (2012) model of granular firms within each sector.[15] It therefore combines fundamental comparative advantage of sectors and granular comparative advantage arising from idiosyncratic productivity draws of individual firms within sectors. In the limit where firms become infinitesimal, granularity is shut off and the model converges to a multisector Melitz (2003) model (see app. B1). We

[14] Our quantitative model reproduces the empirical patterns we document here, albeit with somewhat smaller point estimates relative to table 2 (see table A2 and the discussion in sec. VII).

[15] The Eaton, Kortum, and Sotelo (2012) model is a granular version of the Melitz (2003) model in its Chaney (2008) formulation.

start with a static version of the model and make the model dynamic in section VII.

A. *Model Setup*

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 $\{Q_z\}$:

$$Q = \exp\left\{\int_0^1 \alpha_z \log Q_z \mathrm{d}z\right\}, \tag{3}$$

where $\{\alpha_z\}_{z\in[0,1]}$ are nonnegative 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 constant elasticity of substitution (CES) aggregator:

$$Q_z = \left[\sum_{i=1}^{K_z} q_{z,i}^{(\sigma-1)/\sigma}\right]^{\sigma/(\sigma-1)}, \tag{4}$$

where $\sigma > 1$ is the within-sector elasticity of substitution, common across 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, 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

$$r_{z,i} \equiv p_{z,i}q_{z,i} = s_{z,i}\alpha_z Y, \text{ with } s_{z,i} \equiv \left(\frac{p_{z,i}}{P_z}\right)^{1-\sigma}, \tag{5}$$

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 equation (5) 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

$$P_z = \left[\sum_{i=1}^{K_z} p_{z,i}^{1-\sigma}\right]^{1/(1-\sigma)}. \tag{6}$$

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

2. Production Technology

Each product variety is supplied by an individual firm with productivity $\varphi_{z,i}$ ($\varphi^*_{z,i}$ if the firm is foreign). Products are produced in their market of origin, and firms have access to a constant returns to scale 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

$$c_{z,i} = \begin{cases} \dfrac{w}{\varphi_{z,i}} & \text{if } i \text{ is a home variety,} \\[2ex] \dfrac{\tau w^*}{\varphi^*_{z,i}} & \text{if } i \text{ is a foreign variety,} \end{cases} \tag{7}$$

where w and w^* are the home and foreign wage rates, respectively. 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; that is, it applies for both 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). 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.

3. Productivity Draws

We denote with M_z a potential (shadow) number of domestic products in sector z. The term M_z is the realization of a Poisson random variable with parameter $\overline{M}_z$, so that $\mathbb{E}(M_z) = \overline{M}_z$. Each of the M_z potential entrants takes an independently and identically distributed (i.i.d.) productivity draw from a Pareto distribution with a shape parameter θ and lower bound $\underline{\varphi}_z$.[16] Lower θ corresponds to a more dispersed and skewed

[16] Formally, the realized number of products M_z has the probability density function $\mathbb{P}\{M_z = m\} = e^{-\overline{M}_z}\overline{M}_z^m/m!$ for $m = 0, 1, 2 \ldots$, while the cumulative distribution function of productivity draws φ is given by $G_z(\varphi) = 1 - (\underline{\varphi}_z/\varphi)^\theta$.

distribution of productivity draws, which increases the strength of the granular forces in the model, as we discuss below. We borrow this structure of productivity draws from the earlier work of Bernard et al. (2003) and Eaton, Kortum, and Sotelo (2012). It results both in a tractable model environment and in a realistic cross-sectional distribution of firm sales. In section VI, we explore robustness to an alternative lognormal statistical process for firm productivity draws.

With the Poisson-Pareto productivity structure, the combined parameter

$$T_z \equiv \overline{M}_z \cdot \underline{\varphi}_z^{\theta} \tag{8}$$

is a sufficient statistic that determines the expected productivity of a sector.[17] Intuitively, a sector is more productive either if there are more potential entrants (i.e., productivity draws) equal to $\overline{M}_z$ in expectation or if the average productivity of a potential entrant is high, which is given by $[\theta/(\theta-1)]\underline{\varphi}_z$.

The pool of foreign potential products and the ensuing productivity draws are obtained in a symmetric way, with country-sector-specific parameters $\overline{M}_z^*$ and $\underline{\varphi}_z^*$, resulting in a sufficient statistic for the expected sectoral productivity $T_z^* = \overline{M}_z^* \underline{\varphi}_z^{*\theta}$. The ratio T_z/T_z^* varies across sectors z and determines the expected relative productivity of the two countries, and it is thus a measure of the home's fundamental comparative advantage. The term T_z/T_z^* is the only source of comparative advantage in the continuous DFS-Melitz limit of the model (see app. B1).

4. 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 $i \in \{1, \ldots, K_z\}$ 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:

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

where we used the expressions for the market share of the firm (5) and the sectoral price index (6). While firms are large within their industries and hence internalize their effect on the sectoral price index (6), they are still

[17] In particular, Eaton, Kortum, and Sotelo (2012) 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 $\overline{M}_z$ and $\underline{\varphi}_z$, which result in the same T_z. A convenient limiting case with $\overline{M}_z \to \infty$ 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.

infinitesimal at the level of the whole economy, since the model features a continuum of sectors different from Eaton, Kortum, and Sotelo (2012). 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:

$$p_{z,i} = \frac{\varepsilon_{z,i}}{\varepsilon_{z,i} - 1} \cdot c_{z,i}, \text{ where } \varepsilon_{z,i} \equiv \varepsilon(s_{z,i}) = \sigma(1 - s_{z,i}) + s_{z,i}, \tag{9}$$

with the market share of the firm $s_{z,i}$ defined in equation (5) 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 $\mu_{z,i} \equiv p_{z,i}/c_{z,i} = \varepsilon_{z,i}/(\varepsilon_{z,i} - 1)$ 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.[18]

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 equations (5), (6), and (9). 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

$$\Pi_{z,i} \equiv \Pi_z(s_{z,i}) = \frac{s_{z,i}}{\varepsilon(s_{z,i})} \alpha_z Y - wF. \tag{10}$$

Indeed, operating profits are a fraction $1/\varepsilon_{z,i} = (p_{z,i} - c_{z,i})/p_{z,i}$ of revenues (5), 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 $\{c^*_{z,i}\}_{i=1}^{K^*_z}$. Because of 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.

[18] Much of the earlier granularity literature (including di Giovanni and Levchenko 2012; Carvalho and Grassi 2020) 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 2019). We adopt a more natural case of oligopolistic competition in prices, following Eaton, Kortum, and Sotelo (2012), which results in a less pronounced quantitative difference from the constant markup case.

5. Entry

An equilibrium of the entry game is achieved when for a subset of firms equilibrium profits given by equation (10) are nonnegative, while for any additional entrant profits upon entry would 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 $c_{z,1} \leq c_{z,2} \leq \ldots$, and hence firms with lower indexes i choose whether to enter first.[19] 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 equation (10). 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 market shares (and hence 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 because of monotonicity, it is sufficient to find the unique K_z such that $\Pi_{z,K_z}^{K_z} \geq 0$ and $\Pi_{z,K_z+1}^{K_z+1} < 0$.

6. General Equilibrium

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,

$$Y = wL + \Pi, \tag{11}$$

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

$$\Pi = \int_0^1 \left[\sum_{i=1}^{K_z} \iota_{z,i} \Pi_z(s_{z,i}) + \sum_{i=1}^{K_z^*} (1 - \iota_{z,i}^*) \Pi_z^*(s_{z,i}^*) \right] \mathrm{d}z, \tag{12}$$

with profit function $\Pi_z(s_{z,i})$ defined in equation (10) and $\iota_{z,i} \in \{0, 1\}$ denoting the indicator for whether firm i in sector z in the domestic market

[19] Note that index i is a property not 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.

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's law.

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

$$wL = \int_0^1 \left[\alpha_z Y \sum_{i=1}^{K_z} \iota_{z,i} \frac{s_{z,i}}{\mu(s_{z,i})} + \alpha_z Y^* \sum_{i=1}^{K_z^*} (1 - \iota^*_{z,i}) \frac{s^*_{z,i}}{\mu(s^*_{z,i})} + wFK_z \right] \mathrm{d}z. \quad (13)$$

The three terms on the right-hand side of equation (13) correspond to expenditure on domestic labor for (1) production for domestic market, (2) production for foreign market, and (3) 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 (5) divided by markup $\mu(s_{z,i})$ and hence equals variable costs, that is, expenditure on production labor. Recall that the markup $\mu(s_{z,i}) = \varepsilon_{z,i}/(\varepsilon_{z,i} - 1)$, with $\varepsilon_{z,i}$ defined in equation (9). 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 (13) holds in the foreign country.

Aggregate equilibrium conditions (11) and (13), together with their foreign counterparts and under normalization $w = 1$, allow us to solve for the aggregate equilibrium vector $\mathbf{X} \equiv (w, w^*, Y, Y^*)$, given the sectoral equilibrium vector $\mathbf{Z} \equiv \{K_z, \{s_{z,i}\}_{i=1}^{K_z}, K_z^*, \{s^*_{z,i}\}_{i=1}^{K_z^*}\}_{z\in[0,1]}$.[20] 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.

We note that in the baseline model, home and foreign differ only ex ante in their relative population size and in the mean productivity of their sectors, while the sectors differ only ex ante in their Cobb-Douglas expenditure shares. The countries and sectors are otherwise characterized by identical structural parameters for parsimony. We relax these assumptions and allow for more asymmetry and ex ante heterogeneity when we explore the robustness of our findings in section VI.B.

B. Properties of the Granular Model

We note the relationship between this model and the conceptual granularity framework laid out in section II. Our structural model maps

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

the statistical process for productivity draws $\{\varphi_{z,j}\}_{i=j}^{M_z}$ and $\{\varphi^*_{z,j}\}_{j=1}^{M^*_z}$ into firm-level and sectoral economic outcomes, such as firm market shares $s_{z,i}$ and export shares $\lambda_{z,i}$.[21] Hence, given the productivity distribution and model parameters, this framework characterizes the structural data-generating process denoted $\mathcal{F}_z(\cdot)$ in section II.

In the following sections, we use this 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 the sectoral export share—a measure of the country's comparative advantage in a given sector. The sectoral export share is the cumulative market share of all home firms in the foreign market in a given sector z, and we denote it by[22]

$$\Lambda^*_z \equiv \frac{X_z}{\alpha_z Y^*} = \sum_{i=1}^{K^*_z}(1 - \iota^*_{z,i})s^*_{z,i}, \tag{14}$$

where X_z is total home exports and $\alpha_z\, Y^*$ is total foreign absorption in sector z. By analogy, the foreign export share in sector z is given by $\Lambda_z = X^*_z/(\alpha_z Y)$, where X^*_z denotes total home imports (foreign 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 on fundamental comparative advantage of the sector T_z/T^*_z, is given by[23]

[21] Note that firm i's export share is simply the ratio of its sales in the two markets, $\lambda_{z,i} = (s^*_{z,i}Y^*)/(s_{z,i}Y)$.

[22] One minus the export (or foreign) share, $1 - \Lambda^*_z$, is the *home share,* which features prominently in the gains from trade literature (see Arkolakis, Costinot, and Rodríguez-Clare 2012). Note the slight difference here with sec. II, where we normalized home exports by home absorption rather than foreign absorption. The two measures of export shares differ by Y/Y^*, which is constant across sectors and hence does not affect cross-sectional variance decomposition.

[23] This result applies despite the fact that market shares $s_{z,i}$ are complex nonlinear transformations 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, because of 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. At the same time, the expected number of entrants differs for foreign and home firms, and its ratio is given by Φ_z. The formal derivation of eq. (15) is provided in app. B2.

$$\Phi_z \equiv \mathbb{E}_T \Lambda_z = \mathbb{E}\left\{\Lambda_z \,\middle|\, \frac{T_z}{T_z^*}\right\} = \frac{1}{1 + (\tau\omega)^\theta \cdot T_z/T_z^*}, \tag{15}$$

and symmetrically $\Phi_z^* \equiv \mathbb{E}_T \Lambda_z^* = [1 + (\tau/\omega)^\theta \cdot T_z^*/T_z]^{-1}$ is the expected export share. 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 one-to-one with fundamental comparative advantage T_z/T_z^*. Note that this makes clear why it is natural to use Λ_z as a measure of comparative advantage in the cross section of sectors. The expression in equation (15) 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 app. B1). In short, the granular model has, in expectation, the same sectoral trade shares as the continuous model.

Because of 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*:

$$\Gamma_z \equiv \Lambda_z - \Phi_z, \text{ such that } \mathbb{E}_T \Gamma_z = \mathbb{E}_T\{\Lambda_z - \Phi_z\} = 0. \tag{16}$$

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, offering 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.

By construction, the within-sector granularity does not create extra trade at the aggregate level as compared with the continuous benchmark. Indeed, total imports are[24]

$$X^* = \int_0^1 X_z^* \mathrm{d}z = Y \int_0^1 \alpha_z [\Phi_z + \Gamma_z] \mathrm{d}z = \Phi Y, \tag{17}$$

where $\Phi \equiv \mathbb{E}\{\Phi_z\} = \int_0^1 \alpha_z \Phi_z \mathrm{d}z$ is the aggregate foreign share. 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.

[24] Similarly, $X = \Phi^* Y^*$ is the aggregate value of exports. Because of local fixed costs, the trade balance in general is not $X = X^*$ but is instead $\Phi[Y - wFK] = \Phi^*[Y^* - w^*F^*K^*]$, where K and K^* denote the total number of firms serving the two markets across all sectors. Indeed, $Y - wFK$ 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 app. B2 for the derivation of eq. (17) and the resulting simplification of the general equilibrium system (11)–(13).

V. 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 equation (15), making the 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, after describing the data used in estimation.

A. Data

Our empirical analysis is based on France as home and the rest of the world (ROW) as foreign. We use a data set of French firms (Bordereau de Référence Nominative), 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, $\tilde{d}_{z,j}$ and $\tilde{x}_{z,j}$, as well as four-digit industry classification at the firm level. We use 2005 as our reference year for estimation. We match these data with international trade data from Comtrade to get the aggregate imports and exports of France in each industry. This leaves us with $\tilde{N} = 119$ manufacturing sectors at the four-digit level, with an average of about 350 French firms per sector.[25]

We use tildes to denote the empirical variables that correspond to the theoretical objects that can be measured in the granular model of section VI. The merged data allow us to construct French sectoral expenditure $\tilde{Y}_z = \tilde{\alpha}_z \tilde{Y}$ as the sum of sectoral imports $\tilde{X}_z^*$ (from Comtrade) and domestic sales of all French firms $\tilde{D}_z = \Sigma_{j=1}^{\tilde{M}_z} \tilde{d}_{z,j}$, where j is the rank of French firms and $\tilde{M}_z$ is the observed number of French firms in each sector $z = 1, \ldots, \tilde{N}$. Taking the ratio of sectoral imports to sectoral expenditure, we obtain the foreign share in the home market $\tilde{\Lambda}_z = \tilde{X}_z^* / \tilde{Y}_z$. We also construct a measure of French export intensity as $\tilde{\Lambda}_z^{*\prime} = \tilde{X}_z / \tilde{Y}_z$, where we normalize exports, $\tilde{X}_z = \Sigma_{j=1}^{\tilde{M}_z} \tilde{x}_{z,j}$, by domestic expenditure.[26]

[25] The industry classification used in the French data is the Nomenclature des Activités Françaises (based on the European NACE classification), whereas the trade data use the International Standard Industrial Classification (ISIC) revision 3. We convert the French data into the ISIC revision 3 classification using the crosswalk between NACE and ISIC available from the United Nations Statistical Division. Although the French data provide a finer level of industry aggregation, $\tilde{N} = 119$ is the finest level of aggregation at which both the French data and Comtrade overlap. This precludes estimating the granular model at a finer level of aggregation. In matching the data sets, we have to aggregate some of the smaller French sectors, which explains the difference with the 300 four-digit sectors that we use in tables 1 and 2, where we do not need to match to the Comtrade data.

[26] In the model, this measure is proportional to the French export share, $\Lambda_z^* = \Lambda_z^{*\prime}(Y/Y^*)$, but it is easier to measure in the data since we do not observe sectoral expenditure in the rest of the world.

Last, we construct the relative sales share of French firms in the domestic market:

$$\tilde{s}_{z,j} = \frac{\tilde{d}_{z,j}}{\tilde{D}_z}. \tag{18}$$

Note that the sales share $\tilde{s}_{z,j}$ is different from the market share $s_{z,j}$, as it is calculated only among the domestic firms and hence excludes import sales (in particular, $\tilde{s}_{z,j} = s_{z,j}/(1 - \Lambda_z)$). This is important for identification, as $\{\tilde{s}_{z,j}\}$ are not directly affected by sectoral comparative advantage. To summarize, the data set used in estimation is $\{\tilde{M}_z, \{\tilde{s}_{z,j}\}_{j=1}^{\tilde{M}_z}, \tilde{\Lambda}_z, \tilde{\Lambda}_z^{*\prime}, \tilde{D}_z, \tilde{X}_z, \tilde{X}_z^*\}_{z=1}^{\tilde{N}}$.

B. *Estimation Procedure*

1. Model Parameterization

The strength of fundamental versus granular forces depends on the relative extent of heterogeneity in sectoral productivity levels T_z versus firm productivity draws $\varphi_{z,i}$. To capture the empirical properties of the firm sales distribution, we assume that $\varphi_{z,i}$ are drawn from a Pareto distribution with shape parameter θ, which determines the potential strength of the granular forces. In turn, we parameterize sectoral heterogeneity as being drawn from a lognormal distribution with parameters μ_T and σ_T; that is,

$$\log\left(\frac{T_z}{T_z^*}\right) \sim \mathcal{N}(\mu_T, \sigma_T^2). \tag{19}$$

While μ_T controls the home's absolute advantage, σ_T is the key parameter that determines the strength of the fundamental comparative advantage. In adopting this lognormal assumption as our baseline, we follow the evidence of Hanson, Lind, and Muendler (2018), who show that the distribution of measured comparative advantage across countries and sectors is well approximated by a lognormal distribution. We check that this property is also true in our granular model, quantified under the above distributional assumption. In section VI, we explore robustness to alternative distributional assumptions for both $\varphi_{z,i}$ and T_z.

2. Estimation Strategy

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.[27] In the second step, we use the SMM procedure to

[27] We report the histogram of the resulting α_z in fig. A1. 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 average 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.

estimate the six parameters of the model, $\Theta \equiv (\sigma, \theta, \tau, F, \mu_T, \sigma_T)$. Importantly, our approach to statistical inference in this granular model leverages its multisector nature. We view each sector as a draw from the parametric data-generating process described in section IV. The SMM procedure treats each sector as a (multidimensional) observation from the structural data-generating process, with parameters, common across sectors, that need to be estimated. That is, we treat the $\tilde{N}$ sectors in the data as a finite number of draws from a model with a continuum of sectors. Our statistical inference considers the asymptotics as $\tilde{N}$ increases unboundedly. We note that our baseline model keeps cross-sectoral parametric heterogeneity to a minimum, but the estimation procedure can be readily extended to heterogeneity in other parameters, provided relevant empirical moments are available for identification. See section VI for our robustness analysis, which allows for heterogeneity in productivity parameters θ_z across sectors.

The estimation proceeds as follows: 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 $\mathcal{L}(\Theta) \equiv (\mathcal{M}(\Theta) - \tilde{\mathbf{m}})'\mathbf{W}(\mathcal{M}(\Theta) - \tilde{\mathbf{m}})$, 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.

3. 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 $T_z^* \equiv 1$ for all z without loss of generality.[28]

[28] 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.

Last, 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, $\Theta' = (\theta, \tau, F, \mu_T, \sigma_T)$.[29] This reduces the parameter space and improves the precision of estimation.

C. *Moments and Identification*

We target 15 empirical moments, which correspond to averages and standard deviations of sectoral outcomes, as summarized in table 4. With five parameters, the model 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.

We choose to target moments that are informative about (1) the prevalence of large firms in domestic sectoral sales, (2) the intensity of sectoral exports, and (3) the joint distribution of these two characteristics. This way, we ensure that the model can replicate the heterogeneity across sectors in top firm concentration, export stance, and, importantly, the extent to which the two are correlated, capturing granular forces at play in shaping sectoral outcomes (recall the suggestive evidence in sec. II). We discuss these moments in turn.

1. Domestic Sales Distribution

We target the average and standard deviation across sectors of two measures of within-industry concentration—the relative sales shares of the largest and the top three largest French firms within industry relative to other French firms, that is, $\tilde{s}_{z,1}$ and $\Sigma_{j=1}^{3}\tilde{s}_{z,j}$, as defined in equation (18). These moments are important to make sure the model replicates the sales size distribution in the right tail, a key prerequisite of granularity. The combined parameter $\kappa \equiv \theta/(\sigma - 1)$, which determines the shape of the sales distribution, is particularly sensitive to these moments of industry concentration. Given the calibrated value of the elasticity of substitution σ, these moments are key in identifying the productivity dispersion parameter θ, as we illustrate in figure A3*B*. In addition, we target the

[29] The value of $\sigma = 5$ (within four-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.

average (log) number of French firms operating within sectors as well as its standard deviation. This ensures that the model captures simultaneously the large number of firms operating in French sectors with the high concentration of sales—a reflection of the thick right tail of the productivity distribution rather than high barriers to entry. Intuitively, the fixed cost parameter F is particularly sensitive to the average number of firms, as can be seen in figure A3*A*.

2. Sectoral Trade Patterns

We target a set of five moments describing sectoral trade patterns, a key object of our interest. Specifically, we match the average and standard deviation of foreign shares in the French market $\tilde{\Lambda}_z$ and the French export intensity $\tilde{\Lambda}_z^{*\prime}$, as defined above. These trade moments help inform the estimation of the size of the trade cost τ and the average productivity advantage of France μ_T. Indeed, from equation (15), 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^*, whose mean is governed by μ_T (see illustration in fig. A3*C*). We also target the fraction of French sectors in which export sales exceed the overall domestic sales of French firms. Because of trade costs, such sectors can emerge only when the rest of the world 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 $\omega = w^*/w$.

3. Firm Sales Shares and Sectoral Trade Shares

Finally, we match four moments describing the correlation between French import and export shares, $\tilde{\Lambda}_z$ and $\tilde{\Lambda}_z^{*\prime}$, and the sectoral sales concentration at home, $\tilde{s}_{z,1}$ and $\Sigma_{j=1}^{3}\tilde{s}_{z,j}$. Specifically, we target the regression coefficients of $\tilde{\Lambda}_z$ and $\tilde{\Lambda}_z^{*\prime}$ separately on $\tilde{s}_{z,1}$ and $\Sigma_{j=1}^{3}\tilde{s}_{z,j}$, controlling in all four regressions for the size of the sector (log total domestic expenditure, $\log \tilde{Y}_z$). We denote these regression coefficients with $\hat{b}_\ell$ and $\hat{b}_\ell^*$ for $\ell \in \{1, 3\}$, respectively. Note that the export regressions are related to the evidence reported in section III, with the difference that here we focus on export shares rather than log exports. These moments are instrumental in identifying the relative importance of fundamental and granular forces in shaping trade patterns. In the data, we see a clear correlation pattern—sectors with more concentrated domestic sales at the top have larger export shares, while there is no relationship with import shares.[30]

[30] Note that this asymmetry in the import and export coefficients is at odds with an alternative candidate explanation based on Melitz (2003) selection forces, discussed in sec. III, whereby greater skewness (and hence concentration) in sectoral sales shares should be positively associated with both sectoral exports and imports.

In the model, given the firm productivity parameter θ, these correlations are particularly sensitive to σ_T, which governs the strength of the fundamental comparative advantage, as we illustrate in figure A3*D*. When comparative advantage is dominated by fundamental forces (higher σ_T), the correlation between sectoral exports and top firm sales share tends to be small and can even turn negative. The intuition is as follows: sectors with high fundamental productivity tend to have more domestic firms that enter the market because of their superior productivity draws. More entry leads to lower market shares at the top, all else equal. Therefore, with strong fundamental forces, higher export intensity sectors tend to have lower market shares at the top. In contrast, with granular forces, large market shares at the top are a sign of a strong granular draw—a source of the sector's comparative advantage over and above its fundamental characteristics. Granularity ensures that the correlation we target in our estimation is positive in the model, with its specific value shaped by the interplay of the granular and fundamental forces.

D. Estimation Results and Model Fit

1. Estimated Parameters

Table 3 reports the SMM estimates of the model parameters and their standard errors (as described in app. 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 1.096, significantly above 1, hence exhibiting thinner tails relative to Zipf's law (see Gabaix 2009). Next, we estimate μ_T to be positive, albeit small. A positive μ_T means that France has slightly better productivity draws relative to its average trade partner, in line with the calibrated higher wage rate $w/w^* = 1.13$. The estimated value of $\sigma_T = 1.39$, the standard deviation of fundamental comparative advantage, is large. It suggests that in the cross section of sectors, a 1 standard deviation increase in fundamental comparative advantage corresponds

TABLE 3
Estimated Parameters

Parameter	Estimate	Standard Error	Auxiliary Variable	Estimate
σ	5	. . .	$\kappa = \theta/(\sigma - 1)$	1.096
θ	4.382	.195		
τ	1.342	.101	w/w^*	1.130
F ($\times 10^5$)	1.179	.252	L^*/L	1.932
μ_T	.095	.150	Y^*/Y	1.710
σ_T	1.394	.190	Π/Y	.180

to a fourfold increase in the relative productivity T_z/T_z^*. Below, we discuss the relative role of $\kappa = 1.096$ and $\sigma_T = 1.39$ 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 about two times smaller than the rest of the world in terms of population. 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 rest of the world, in which the individual countries are 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 18%, broadly in line with the national income accounts, without being targeted in the estimation procedure.

2. Model Fit

Table 4 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, 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 4.

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 three largest French firms, which are 20% and 35%, respectively, of the overall domestic sales of all French firms. The ability of the model to closely replicate the distribution of the number of firms and the market shares of the largest firms across sectors is important for 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%.[31]

The regression coefficients of the sectoral trade shares on either the top firm or the top three domestic concentration ratios are 0.20–0.25 in the data for exports and around zero for imports. The model matches these patterns accurately. The table further reports the OLS standard

[31] 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 4
MOMENTS USED IN SMM ESTIMATION

Moments	(1)	Data, $\tilde{\mathbf{m}}$ (2)	Model, $\overline{\mathcal{M}}(\hat{\Theta})$ (3)	Loss (%) (4)
1. Log number of firms, mean	$\log \tilde{M}_z$	5.631	5.429	1.9
2. Log number of firms, SD		1.451	1.230	3.9
3. Top firm sales share, mean	$\tilde{s}_{z,1}$	.197	.205	3.0
4. Top firm sales share, SD		.178	.148	4.5
5. Top three sales share, mean	$\Sigma_{j=1}^{3}\tilde{s}_{z,j}$	.356	.343	2.0
6. Top three sales share, SD		.241	.176	12.2
7. Imports/domestic sales, mean	$\tilde{\Lambda}_z$	.365	.354	1.5
8. Imports/domestic sales, SD		.204	.266	15.2
9. Exports/domestic sales, mean	$\tilde{\Lambda}_z^{*\prime}$	.328	.345	3.9
10. Exports/domestic sales, SD		.286	.346	7.2
11. Fraction of sectors with exports > domestic sales	$\mathbb{P}\{\tilde{X}_z > \tilde{D}_z\}$	.185	.095	39.7
Regression coefficients:				
12. Export share on top firm share	$\hat{b}_1^*$	.215	.240	2.2
		(.156)	.104	
13. Export share on top three share	$\hat{b}_3^*$	.254	.222	2.6
		(.108)	(.090)	
14. Import share on top firm share	$\hat{b}_1$	−.016	−.011	.1
		(.097)	(.079)	
15. Import share on top three share	$\hat{b}_3$	.002	.008	.1
		(.074)	(.069)	

NOTE.—Column 4 reports the contribution of the moment to the loss function $\mathcal{L}(\hat{\Theta})$, as described in app. C. Moments 12–15 are regression coefficients of $\tilde{\Lambda}_z^{*\prime}$ and $\tilde{\Lambda}_z$ on $\tilde{s}_{z,1}$ and $\Sigma_{j=1}^{3}\tilde{s}_{z,j}$, controlling in all cases for the size of the sector with the log domestic sectoral expenditure $\tilde{Y}_z$; OLS standard errors are in parentheses. SD = standard deviation.

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 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 9.5% of such sectors against 18.5% in the data. Note that the presence of such sectors is possible only in a model with $Y^* > Y$, that is, when France is smaller than the rest of the world. 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 column 4 of table 4.

3. Moments Not Targeted in Estimation

We consider here a series of overidentification checks by exploring the fit of the model for moments not targeted in estimation. First, we estimate the Pareto shape parameter $\hat{\kappa}_z$ of domestic sales of French firms industry by industry. Following Gabaix and Ibragimov (2011), we run the following regression on the top 25% of firms in each industry (the results are similar for the sample of top 50% of firms):

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

$$\begin{array}{rl} \text{Data:} & 1.015 \\ & [0.817, 1.208] \\ \\ \text{Model:} & 1.121 \\ & [1.001, 1.201] \end{array} \tag{20}$$

where j is the domestic-sales rank of French firms in industry z and the display reports the mean and the interquartile range of $\hat{\kappa}_z$ across sectors. This provides an additional measure of within-sector concentration, with a lower $\hat{\kappa}_z$ corresponding to a more fat-tailed (concentrated) sales distribution. On average, the distribution of domestic sales exhibits Zipf's law in the data, that is, the estimated Pareto shape parameter is equal to 1.015, close to 1. The model somewhat overstates the mean of $\hat{\kappa}_z$, at 1.12.[32] With a less fat-tailed sales distribution compared with the data, the model therefore offers a conservative bound for the role of granularity, as we explore in section VI.B.

Consider now the joint distribution of the number of French firms and their concentration ratio across sectors. Recall that in the estimation, we match their properties separately. As an overidentification check, we regress the relative size of the largest French firm $\tilde{s}_{z,1}$ on the log number of French firms $\tilde{M}_z$, controlling for the log domestic absorption $\tilde{Y}_z$ in the sector. We report the OLS-estimated semielasticities, γ_M and γ_Y, and their standard errors:

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

$$\begin{array}{rcc} \text{Data:} & -0.094 & 0.018 \\ & (0.008) & (0.008) \\ \\ \text{Model:} & -0.124 & 0.079 \\ & (0.010) & (0.010) \end{array}$$

[32] Recall that in the model, the average shape parameter is closely related to $\kappa = \theta/(\sigma - 1) = 1.096$ and is slightly higher (less fat tailed) because of variable markups. Indeed, the markups are higher for larger firms, and hence the sales distribution is less concentrated than it would be under constant markups.

In the data, sectors with more French firms, tend to have relatively smaller largest firms; however, this relationship is not very steep. Furthermore, conditional on the number of firms, the size of the sector (measured by domestic absorption) correlates positively (albeit very weakly) with the relative size of the largest firm. Both of these patterns are in line with the predictions of the estimated granular model.

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, as well as their joint covariation. This is perhaps surprising given the parsimony of the model's parameterization, which features only five parameters. Granular forces are instrumental in generating these patterns of variation across sectors, mimicking the patterns observed in the data.

4. 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 equation (9). However, under Bertrand competition, the equilibrium variation in markups is quite limited, as we illustrate in 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 ($\sigma/(\sigma - 1) = 1.25$). 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 can be a useful simplification in some applications, yet this is not the case in general. Top firms are pivotal for a range of sectoral outcomes, and their variable markups are at the core of the optimal trade and industrial policies, as we briefly discuss at the end of our analysis.

VI. Quantifying Granular Trade

A. Granular Contribution to Comparative Advantage

Armed with the estimated model, we now study the extent to which granularity shapes trade patterns. Recall from equations (14)–(16) that sectoral trade flows X_z are determined by three factors: (1) sectoral expenditure shares α_z, (2) fundamental comparative advantage Φ_z, and

(3) 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:

$$X_z = Y^* \alpha_z \Lambda_z^*,$$
$$\Lambda_z^* = \Phi_z^* + \Gamma_z^*.$$

Table 5 reports the decomposition of trade flows into the above three sources of variation in the estimated model (col. 1). Columns 2–5 and R1–R3 report robustness results, which we discuss below.

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:

$$\text{var}(\Lambda_z^*) = \text{var}(\Phi_z^*) + \text{var}(\Gamma_z^*). \tag{21}$$

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 18% of the variation in export shares across sectors, while the rest corresponds to fundamental comparative advantage. In turn, export shares Λ_z^* account for 54% of the variation in overall trade flows X_z, while the rest is accounted for by the (expenditure) size of the sectors α_z.[33]

Since the granular contribution to trade flows is zero on average across sectors, granularity does not create additional trade at the aggregate level. Instead, granularity creates additional trade flows in granular sectors, which are compensated by missing trade in nongranular sectors, as we investigate next. To clarify our terminology, by convention, 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 *nongranular*, even though ex ante all sectors are symmetric in terms of their expected granularity, as $\mathbb{E}\Gamma_z^* = 0$ for every z.

Figure 1 illustrates that the effects of granularity are particularly pronounced in the most export-intensive sectors, that is, 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^*. Panel B plots the corresponding export flows; that is, it shows the contribution of each group of sectors to the country's total exports (dashed light gray bars) and highlights with dark gray solid bars the granular contribution. Note that the cumulative height of all

[33] 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 also features a covariance term between Λ_z^* and α_z, which, however, happens to be close to zero and therefore does not affect quantitatively the results of the variance decomposition.

TABLE 5
VARIANCE DECOMPOSITION OF TRADE FLOWS

	COMMON θ		SECTOR-SPECIFIC θ_z			ROBUSTNESS		
	(1)	(2)	(3)	(4)	(5)	(R1)	(R2)	(R3)
Granular contribution (%)	17.8	25.8	27.1	31.8	22.1	17.8	1.2	18.2
Λ_z^* contribution (%)	54.4	60.6	59.4	62.9	56.2	48.5	85.0	54.8
Top firm sales share, $\tilde{s}_{z,1}$	**.21**	.26	.27	.32	**.21**	**.21**	**.17**	**.21**
Estimated Pareto shape, $\hat{\kappa}_z$	1.12	**1.02**	1.12	**1.02**	1.24	1.09	1.09	1.12

NOTE.—Granular contribution to sectoral export shares is $\text{var}(\Gamma_z^*)/\text{var}(\Lambda_z^*)$; Λ_z^* contribution to sectoral exports is $\text{var}(\log \Lambda_z^*)/\text{var}(\log X_z)$ (see decomposition of X_z in the text). We report the moments of firm concentration, with the targeted moment in bold; $\hat{\kappa}_z$ is estimated as in eq. (20). Specifications are as follows: (1) baseline estimated model; (2) counterfactual with $\sigma = 5.5$ to match average $\hat{\kappa}_z$; (3) $\theta_z = (\sigma - 1)\hat{\kappa}_z$, where $\hat{\kappa}_z$ are Pareto shapes estimated in the data sector by sector; (4) similar to specification 3 but proportionally scaling all θ_z down to match the average $\hat{\kappa}_z$ in the data; (5) similar to specification 3 but proportionally scaling θ_z up to match $\tilde{s}_{z,1}$; (R1) fat-tailed T_z/T_z^*; (R2) lognormal $\varphi_{z,i}$; (R3) nongranular foreign. Table A3 describes the moment fit of alternative robustness specifications.

light gray bars is 1 (aggregate exports), while the cumulative height of all dark gray 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 they account for a substantial fraction of trade flows. In all other bins of less export-intensive sectors, the contribution of granular trade is negative;

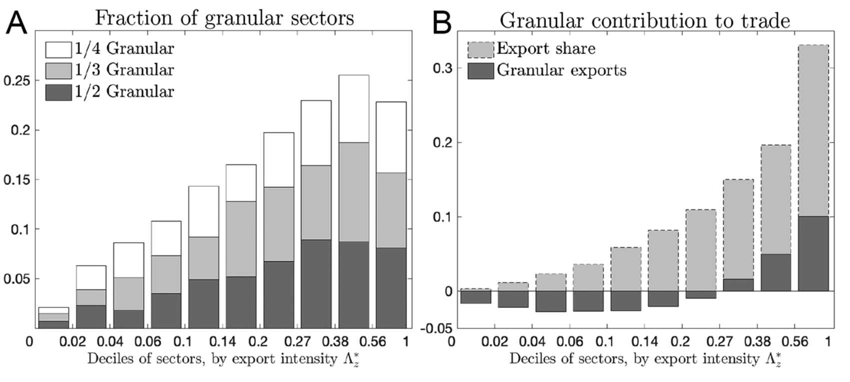

FIG. 1.—Export intensity and granularity. All sectors are split into 10 deciles (bins of equal size in the number of sectors) based on their export share, $\Lambda_z^* = X_z/\alpha_z Y^*$. Panel A plots for each decile the fraction of sectors for which $\Gamma_z^* \geq \vartheta\Lambda_z^*$ for $\vartheta \in \{1/2, 1/3, 1/4\}$. For example, the cumulative height of the light and dark gray bars corresponds to the fraction of sectors with $\Gamma_z^* \geq (1/3)\Lambda_z^*$ or, equivalently, $\Gamma_z^* \geq (1/2)\Phi_z^*$. Panel B plots the contributions of deciles to aggregate trade (dashed light gray bars) and of granular trade ($\Gamma_z^*\alpha_z Y^*$) to aggregate exports $X = \Phi^* Y^*$ (solid dark gray bars) by deciles of sectors. A color version of this figure is available online.

that is, these sectors would export slightly more in the continuous limit of the model.

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 have only 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.

Inference on sectoral comparative advantage.—We next explore what inference one can make on the fundamental comparative advantage of a sector given the observed export stance of a sector Λ_z^*. 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 in this case $\Lambda_z^* = \Phi_z^*$, 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 below.

We use the estimated model to plot, in panel A of figure 2, the distribution of realized export intensity Λ_z^* conditional on the fundamental

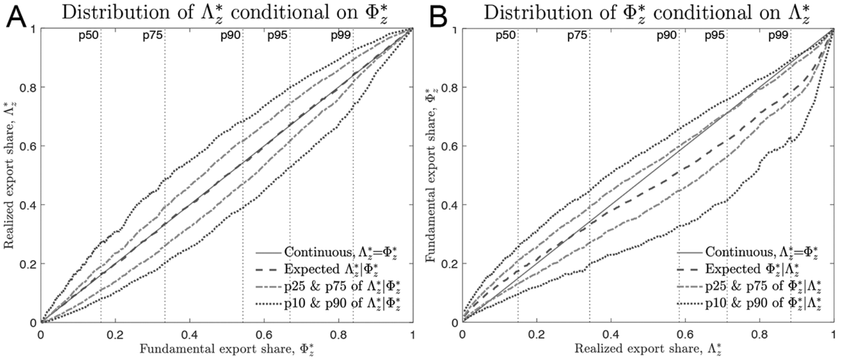

FIG. 2.—Comparative advantage and trade flows: distribution across realizations. The panels plot moments and percentiles of the conditional distributions: $\Lambda_z^*|\Phi_z^*$ (*A*) and $\Phi_z^*|\Lambda_z^*$ (*B*). In both panels, the solid diagonal line corresponds to $\Lambda_z^* = \Phi_z^*$, toward which the distributions degenerate in a continuous model. The vertical dotted lines plot the percentiles of the unconditional partial distribution of Φ_z^* (*A*) and Λ_z^* (*B*). A color version of this figure is available online.

comparative advantage of a sector Φ_z^*. The one-to-one deterministic mapping between the two in the continuous model is depicted with a solid diagonal 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 line, which coincides with the solid diagonal 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 diagonal correspond to the realizations of the sectoral granular residuals, $\Gamma_z^* = \Lambda_z^* - \Phi_z^*$.[34] This figure complements the decomposition in table 5 in illustrating the contribution of granularity to sectoral trade shares.

Panel B of figure 2 describes instead the conditional distribution of Φ_z^* given Λ_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, panel B switches the axes of panel A. The continuous model is again represented by the solid 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 dashed line, which unlike in panel B now departs from the solid diagonal line. In other 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 panel A. 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 to have positive granular residuals. Therefore, sectors with the largest realized export shares have systematically lower expected export shares, $\Phi_z^* < \Lambda_z^*$, that is, a lower fundamental comparative advantage than a continuous model would predict.[35]

To summarize, 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. An estimated granular model can be used to correct for this bias on average.[36]

[34] 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^*.

[35] This corresponds to a classical selection (or reversion to the mean) effect: a sector outlier is shaped only in part by fundamental forces, and the less so the more of an outlier it is. From panel B of fig. 3, 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, $\mathbb{E}\Phi_z^* = \mathbb{E}\Lambda_z^*$).

[36] In an early working paper (Gaubert and Itskhoki 2018), we discuss a Bayesian inference procedure of the probability that exports in a given sector z are of a significant granular origin, e.g., that $\Gamma_z^* \geq \vartheta\Lambda_z^*$ for a given cutoff $\vartheta \in (0, 1)$, conditional on the sectoral observables.

B. *Robustness*

Since our approach to identifying granular contribution relies on a parametric structural model, we now consider a variety of alternative parameterizations to verify the robustness of our quantitative findings. In particular, we relax the assumption of a common productivity parameter θ across sectors, and we explore the robustness of our results to alternative distributional assumptions for sectoral and firm-level productivity draws.

1. Matching the Pareto Shape

Column 2 of table 5 reports the sensitivity of our baseline results (in col. 1) to alternative values of the elasticity of substitution of demand, σ. We do this for two reasons. First, as we note 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 estimated Pareto shape $\hat{\kappa}_z$ (see eq. [20]). We report here what would be the outcome of a less conservative calibration 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., Gutiérrez and Philippon 2017; Autor et al. 2020). One common hypothesis is that it corresponds to an increased substitutability across products σ, for example, because of the increased online competition.[37] 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 $\sigma = 5.5$, which allows the model to match exactly the average estimated Pareto shape parameter $\hat{\kappa}_z$ in the data (equal to 1.02) while overstating somewhat the size of the largest firms. We find that the contribution of granularity to sectoral export shares increases from 18% to 26%—intuitively, the role of granularity increases in a more concentrated economy. Our baseline estimation can therefore be viewed as conservative. We conclude that a reasonable contribution of granularity to sectoral trade is around 20%.

2. Heterogeneity in θ

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 heterogeneous in a number of different ways. One may thus worry that our results

[37] A natural microfoundation for this mechanism is a frictional discrete choice model with decreasing search costs over time (see, e.g., Hortaçsu and Syverson 2004).

are sensitive to this simplifying assumption and, specifically, that we overstate the role of granularity by shutting down such heterogeneity. In particular, variation in the firm size distribution across sectors 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, that is, the parameters that govern the dispersion in firm productivity draws within sectors.[38] 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 columns 3–5 of table 5. First, we choose θ_z so that $\theta_z/(\sigma - 1) = \hat{\kappa}_z$ sector by sector, where $\hat{\kappa}_z$ are the empirical estimates of the Pareto shapes of the firm size distribution in the data (using eq. [20]). With a continuum of firms and constant markups, $\theta_z/(\sigma - 1)$ exactly corresponds to the Pareto shape of the sales distribution and hence justifies this calibration approach. Variable markups, however, introduce a gap between $\theta_z/(\sigma - 1)$ and the Pareto shape of firm sales in the model. To match the Pareto shape of the data, our second calibration therefore proportionally scales down the distribution of θ_z while preserving its heterogeneity across sectors 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 the average sales share of the largest firm relative to the data, we proportionally scale up the distribution of θ_z to match the top sales share moment, as does the baseline model. Table 5 and 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 $\hat{\kappa}_z$.

Interestingly, table 5 shows that the contribution of granularity increases across all three specifications with heterogeneous θ_z compared with our baseline with homogeneous θ. The contribution of granularity now ranges from 22% to 32%. Intuitively, the strength of granularity is largely determined by the market share of the largest firm in the sector. Having heterogeneous θ_z does not change the ability of the model to match the relative size of the largest firms on average across sectors. 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 ends up reinforcing the aggregate role of granularity—even if some sectors end up being less granular—because granularity is inherently an infrequent right-tail outcome, which is favored by this

[38] Another realistic source of heterogeneity is variation in the fixed cost of entry, F_z. However, in the model, which is estimated to simultaneously fit the fat-tailed sales distribution and the large number of entrants, marginal firms are very small. Therefore, variation in F_z has a very limited effect on the key moments of interest.

increased asymmetry (recall fig. 1*B*). Overall, this robustness exercise confirms that our baseline estimate of the role of granularity is, if anything, conservative.

3. Fat-Tailed Distribution of Sector-Level Productivity

In our baseline, we assume that relative sectoral productivity levels T_z/T_z^*—which shape fundamental comparative advantage (recall eq. [15])—are lognormally distributed. As we demonstrated, this assumption allows the model to match the data well. However, a natural concern is that assuming that this distribution is thin tailed (lognormal) could mechanically limit the extent of heterogeneity in fundamental comparative advantage, and it thus leads us to overstate the important of granularity. To address this concern, we replace the lognormal distribution for T_z/T_z^* with a fat-tailed double-sided Pareto, with μ'_T and σ'_T still parametrizing the mean and the standard deviation of $\log T_z/T_z^*$.[39] With this assumption, sector-level productivity draws and firm-level productivity draws are on equal footing in terms of producing potential right-tail outcomes.[40] We keep the remaining parametric assumptions unchanged and estimate this version of the model using the same procedure, described in section V. We report the results of this robustness estimation in column R1 of table 5. Perhaps surprisingly, we find that both models are nearly identical in their ability to fit the set of identifying moments in table 4, with the fat-tailed counterfactual model slightly underperforming on moments 11–15.[41] In turn, we find that, reassuringly, this alternative assumption leads to estimating exactly the same granular contribution as in the baseline (namely, 17.81% vs. 17.76%), with a slightly smaller role played by the export intensity in shaping the total exports (48.5% vs. 54.4%). This result suggests that our baseline parametrization is not mechanically constrained by the distributional assumptions to imply a high granular contribution but that this result is rather driven by the model's ability to fit data moments.

4. Thin-Tailed Distribution of Firm-Level Productivity

We next turn to assumptions on the distribution of firm-level productivity. We go back to our baseline specification and now relax the assumption

[39] Formally, we assume $\log T_z/T_z^* \sim Laplace(\mu'_T, \sigma'_T/2^{1/2})$, where $Laplace(a, b)$ is a two-sided exponential with density $f(x|a, b) = (1/2b)\exp\{-|x - a|/b\}$.

[40] An alternative strategy to calibrate the sectoral productivities T_z/T_z^* could be to match directly the export shares sector by sector. This, however, mechanically implies $\Lambda_z = \Phi_z$ and no role for granularity $\Gamma_z = 0$. As discussed in sec. V.C, such an approach would be rejected by the data, as it would fail to match the correlation between sectoral exports and concentration at the top.

[41] We report the fit of alternative models in table A3, where we also report the overall loss function, which is 0.259 for the baseline and slightly bigger at 0.297 for this robustness counterfactual.

that firm productivity draws are fat tailed (Pareto). Instead, we now assume that their distribution is thin tailed (lognormal). This checks whether imposing a fat-tailed distribution on the firm-level productivity draws could lead us mechanically to find a sizable granular contribution.[42] We again reestimate the model under the new distributional assumption. This time, conclusions are different. In short, the model is unable to match the data in this case. In particular, it fails to jointly match the number of firms per sector and the size of the largest firm—understating both moments by about 30% (see table A3). Furthermore, it also misses on our key identifying moments 12–15, failing to capture even qualitatively the patterns of correlation between the size of the largest firm and sectoral trade flows (see discussion in sec. V). We conclude that the lognormal distribution does not have enough skewness to reproduce the salient empirical patterns we highlight. Unsurprisingly, as a corollary of this failure to fit the data, this counterfactual model features a negligible granular contribution, reported in column R2 of table 5.

To summarize, this robustness check along with the previous one suggests that our main results on the importance of granularity are not driven by specific choices of functional forms or distributional assumptions. Instead, these choices allow the baseline model to match the moments of the data that speak to granularity and also help discriminate between parametric models. To the extent that these moments are matched, our quantitative conclusions on the importance of granularity are robust to alternative modeling choices.

5. Nongranular Foreign

Our last robustness considers the case where the rest of the world (being large) is assumed to be nongranular, while France still is. We keep the estimated parameters the same as in the baseline but replace the foreign productivity draws $\{\varphi^*_{z,i}\}$ with deterministic values so that the counterfactual model features the same Φ_z and Φ^*_z as the granular model yet shuts down the uncertainty regarding foreign productivity draws.[43] We find that this alternative has little impact on the moments for home (France) and, in particular, changes little our conclusions about the importance of granular comparative advantage. As we report in column R3 of table 5, the granular contribution to French export intensity Λ^*_z increases marginally to 18.2%.

[42] Formally, we take a fixed number $M_z = \alpha_z M$ of potential entrants who draw productivity such that $\log \varphi_{z,i} \sim \mathcal{N}(\mu_z, \theta'^2)$, where $\mu_z = \log T_z \sim \mathcal{N}(\mu_T, \sigma_T^2)$ is expected (fundamental) sectoral productivity and θ' parameterizes the dispersion of the productivity draws. We provide formal details in app. C.

[43] Specifically, we fix the number of firms per sector at $\overline{M}^*_z$ and set $\varphi^*_{z,i}$ to equal the $\overline{M}^*_z$ percentiles of the *Pareto*$(\underline{\varphi}^*_z; \theta)$ distribution, with $\overline{M}^*_z \underline{\varphi}^{*\theta}_z = T^*_z$ to keep the average realization of $\varphi^*_{z,i}$ as in the baseline.

VII. Dynamics of Comparative Advantage

Having established the implications of granularity for cross-sectional trade patterns, we now extend our model to a dynamic environment and study the implication of granularity for the evolution of a country's comparative advantage over time. We are interested in the contribution of granular forces to both the volatility of sectoral export patterns as well as the predictability of comparative advantage reversals observed in the data.

A. *Dynamic Model*

We introduce dynamics by assuming that firm-level productivity evolves over time according to a random growth process subject to both aggregate (sectoral) and idiosyncratic (firm-level) shocks. As a consequence, both fundamental and granular comparative advantage change over time along with the within-sector distributions of firm sales shares. In this granular model of industry dynamics, firm-level volatility contributes to shaping the dynamics of sectoral trade flows and the evolution of comparative advantage.

Consistent with our cross-sectional model, we assume that there exist a Poisson-distributed number of shadow firms M_z in each sector with productivity $\varphi_{z,i,t}$ that evolves over time such that at each date, (1) within each sector, the cross-sectional distribution of relative firm productivities remains stable and distributed Pareto with shape θ; and (ii) across sectors, the distribution of expected sector-level productivities T_z remains stable lognormal with parameters μ_T and σ_T. Stability over time requires mean reversion for both firm-level and sectoral productivity components. To achieve this, we assume that productivity $\varphi_{z,i,t}$ of firm i in sector z at period t, relative to a reference (cutoff) sectoral productivity $\underline{\varphi}_{z,t}$, evolves according to a geometric random walk with a negative drift μ and a reflecting barrier at 1 (0 in logs). Therefore, $\underline{\varphi}_{z,t}$ is the lower bound of the firm productivity distribution, and a change in $\underline{\varphi}_{z,t}$ shifts proportionally the entire productivity distribution in sector z, thus capturing shocks to fundamental comparative advantage. We assume that $\log \underline{\varphi}_{z,t}$ follows an autoregressive process with parameter ρ.

Formally, we set

$$\begin{cases} \log\left(\dfrac{\varphi_{z,i,t}}{\underline{\varphi}_{z,t}}\right) = \left|\mu + \log\left(\dfrac{\varphi_{z,i,t-1}}{\underline{\varphi}_{z,t-1}}\right) + \alpha_u u_{z,i,t}\right|, \\ \log\left(\dfrac{\underline{\varphi}_{z,t}}{\underline{\varphi}^o_z}\right) = \rho \log\left(\dfrac{\underline{\varphi}_{z,t-1}}{\underline{\varphi}^o_z}\right) + \alpha_v v_{z,t}, \end{cases} \tag{22}$$

where $u_{z,i,t}, v_{z,t} \sim iid\mathcal{N}(0, 1)$ are idiosyncratic and aggregate innovations, respectively, and $\underline{\varphi}^o_z$ are exogenous long-run sectoral means. The parameters

α_u and α_v capture the magnitude of idiosyncratic and fundamental shocks, respectively, in driving productivity changes. Note that when $\alpha_v = 0$, the model features only idiosyncratic shocks and no change in the fundamental comparative advantage over time. In contrast, when $\alpha_u = 0$, there are no idiosyncratic shocks and, in particular, the relative sales shares of domestic firms remain stable over time.

We rely on the following result to ensure that the properties we describe above hold:[44]

LEMMA 1. Let $\{\underline{\varphi}_{z,t}, \varphi_{z,i,t}\}$ follow equation (22). If $\mu = -\theta\alpha_u^2/2$, $\rho = (1 - \theta^2\alpha_v^2/\sigma_T^2)^{1/2}$, $\underline{\varphi}_z^o = (1/\theta)(\mu_T - \log \overline{M}_z)$, and under suitable initial conditions, we have that in every time period t,

$$\begin{aligned} \varphi_{z,i,t} &\sim iid Pareto(\underline{\varphi}_{z,t};\theta), \\ T_{z,t} &= \overline{M}_z \underline{\varphi}_{z,t}^\theta \sim iid \log \mathcal{N}(\mu_T, \sigma_T^2). \end{aligned} \tag{23}$$

We adopt the parametric restrictions in the lemma to ensure the stability of the cross-sectional distributions of productivities in every time period. Finally, there is no entry or exit, but firms decide each period whether to pay a per-period fixed cost and be active or stay inactive. Since firms do not incur sunk costs, their choice is static. That is, each period, firms play the static entry game described in section IV, given the current realized productivity distribution, which gradually evolves over time according to equation (22). This offers a tractable way to extend our granular model to a dynamic environment with persistent productivity processes at both the sector and the firm levels: every cross section of the model for $t \in \{0, 1, 2, ...\}$ is consistent with the static model in section IV.

Dynamic calibration.—By lemma 1, every cross section of the dynamic model is consistent with the static model, and therefore we can directly adopt our cross-sectional parameter estimates from section V. Furthermore, the dynamics are driven by two additional parameters that drive the productivity process (22), α_u and α_v. These parameters govern both the volatility and persistence of the sector- and firm-level productivity process. We discipline them by matching the time series properties of firm-level and sectoral sales.

Specifically, using the panel of French firms in our data from 1997 to 2007, we target the 10-year mean reversion coefficients for sectoral log exports, $\log X_{z,t}$, and for firm-level relative sales shares in the domestic

[44] In our simulation, $\log \overline{M}_z = m$ is constant across sectors, and we choose m large enough to ensure that the least productive firms at $\underline{\varphi}_{z,t}$ always stay inactive. As a result, we have $\underline{\varphi}_z^o = (\mu_T - m)/\theta$ also constant across z, and therefore we use $\log(\varphi_{z,0}/\underline{\varphi}_z^o) = (\log T_{z,0} - \mu_T)/\theta \sim iid \log \mathcal{N}(0, \sigma_T^2/\theta^2)$ and $\varphi_{z,i,0} \sim iid\text{Pareto}(\underline{\varphi}_{z,0};\theta)$ as initial conditions. The negative drift term $\mu = -\theta\alpha_u^2/2 < 0$ ensures stationarity of the relative productivity distribution, $\varphi_{z,i,t}/\underline{\varphi}_{z,t}$, which is Pareto with shape parameter θ (see Gabaix 2009). Similarly, for aggregate shocks, the relationship between mean reversion ρ and volatility α_v^2 ensures that the dispersion of $\log T_{z,t}$, $\sigma_T^2 = \theta^2\alpha_v^2/(1 - \rho^2)$, remains constant in every cross section.

TABLE 6
Dynamics of Comparative Advantage

	Data		Model	
Moment	Hanson, Lind, and Muendler 2018	France	Baseline	$\alpha_v = 0$
Target moments:				
β_X, 10-year mean reversion in log $X_{z,t}$		−.106 (.034)	−.106	−.025
β_s, 10-year mean reversion in $\tilde{s}_{z,i,t}$		−.108 (.003)	−.108	−.097
Granular contribution to var $(\Delta_k \Lambda_{z,t})$ (%):				
Year-to-year changes ($k = 1$)			24.4	
10-year changes ($k = 10$)			23.0	100
50-year changes ($k = 50$)			22.4	
Share in aggregate exports (%):				
Top 1% of sectors	21	17	17.2	
Top 3% of sectors	43	30	32.3	
Turnover of comparative advantage (%):				
Remain in top 5% after 10 years		80	76	90
Remain in top 5% after 20 years	52		62	87

Note.—Empirical moments are from Hanson, Lind, and Muendler (2018) for developed countries and from our French data set, where available. The export share moments are based on sectoral exports $\Lambda_z^* \alpha_z Y^*$; the turnover moments are based on export shares Λ_z^*.

market, $\tilde{s}_{z,i,t}$, as we report in table 6.[45] In the model, aggregate shocks affect the former moment but not the latter, which captures the properties of the relative firm productivities within a sector and hence is not affected by the sectoral comparative advantage.[46] This ensures identification: the firm-level moment is essentially sensitive only to α_u, while the sectoral trade moment is sensitive to both α_v and α_u. Importantly, this means that the idiosyncratic productivity process is not identified from the properties of the trade flows.

[45] Specifically, we run $\log(X_{z,t+10}/X_{z,t}) = \alpha_X + \beta_X \log X_{z,t} + \gamma_X \log D_{z,t} + \varepsilon^X_{z,t}$, where $D_{z,t}$ is the control for the size of the market (domestic sales) and $\tilde{s}_{z,t+10} - \tilde{s}_{z,t} = \alpha_s + \beta_s \tilde{s}_{z,i,t} + \varepsilon^s_{z,i,t}$. We target the two mean reversion coefficients, β_X and β_s. We use all 300 four-digit sectors in our data, as in tables 1 and 2, since these regressions do not rely on the match with the Comtrade database, and we aim to obtain the most precise possible estimate of β_X for sectoral exports. We use a balanced panel of 43,882 French firms that survive throughout our sample to estimate β_s (and an equivalent procedure in the model-simulated data), resulting in a 10-year autocorrelation of within-industry firm sales shares equal to $1 + \hat{\beta}_s = 0.892$.

[46] This is an exact analytical result with constant markups and applies approximately in our environment, as the behavior of most firms is accurately approximated by a constant markup rule (recall fig. A2). We check quantitatively in the estimated model that the firm-level moment β_s is not sensitive to the volatility of the aggregate shocks α_v.

Matching the empirical mean reversion moments requires setting $\alpha_v = 0.034$ and $\alpha_u = 0.050$, which correspond to the annual volatility of sectoral and idiosyncratic productivity shocks, respectively.[47] Since the annual volatility of the aggregate shocks, $\alpha_v = 0.034$, is much smaller than the cross-sectional dispersion of comparative advantage, $\sigma_T = 1.39$, the model requires a very high value of $\rho = 0.995$ (from the formula of lemma 1) to reconcile the dynamics with the cross section. In other words, fundamental comparative advantage is highly persistent, albeit mean reverting over long horizons.

B. Granular Dynamics of Comparative Advantage

Equipped with our quantitative dynamic model, we now study the contribution of granularity to the evolution of comparative advantage over time. We start with the dynamic counterpart to the variance decomposition in section VI. Note that this is a distinctly different decomposition. While the static and dynamic analyses both rely on the same general structure of the granular model and, in particular, the same dispersion of firm-level outcomes shaped by θ, granular forces can play different roles in the cross section and in the time series. The long-run steady-state properties of the model are shaped by the cross-sectional dispersion of fundamental comparative advantage, σ_T, while the short- to medium-run outcomes depend on the relative volatility of aggregate and idiosyncratic shocks, α_v and α_u.

1. Variance Decomposition

We begin with the variance decomposition of changes in export intensity across sectors, $\text{var}(\Delta_k\Lambda^*_{z,t})$, where $\Delta_k\Lambda^*_{z,t} \equiv \Lambda^*_{z,t+k} - \Lambda^*_{z,t}$ is the k period forward difference. Given that $\Lambda^*_{z,t} = \Phi^*_{z,t} + \Gamma^*_{z,t}$, where $\Phi^*_{z,t}$ evolves together with $T_{z,t}/T^*_{z,t}$ according to equation (15), we decompose

$$\text{var}(\Delta_k\Lambda^*_{z,t}) = \text{var}(\Delta_k\Phi^*_{z,t}) + \text{var}(\Delta_k\Gamma^*_{z,t}), \tag{24}$$

and we are interested in the granular contribution, $\text{var}(\Delta_k\Gamma^*_{z,t})/\text{var}(\Delta_k\Lambda^*_{z,t})$. Compare this decomposition with equation (21) in section VI.[48]

[47] Our calibration matches the long-run mean reversion of the firm sales shares for both small and large firms, which are approximately the same in the data. At the same time, we somewhat understate the extent of year-to-year volatility in the firm sales shares for both small and large firms. When we target either of these latter moments, we recover a larger α_u and, correspondingly, a smaller α_v, which implies a greater role for granularity. Therefore, from the point of view of the counterfactuals below, our choice of calibration targets is conservative. In addition to this conservative baseline, we also consider a robustness calibration, which puts weight on both the long-run mean reversion β_s and the short-run volatility $\text{std}(\Delta\tilde{s}_{z,i,t})$, resulting in a somewhat larger $\alpha_u = 0.060$.

[48] Note that the covariance term is zero in both cases, as by construction $\Gamma^*_{z,t}$ is orthogonal to $\Phi^*_{z,t}$ in every cross section, and in addition $\Gamma^*_{z,t+k}$ is orthogonal to $\Phi^*_{z,t}$ for $k \gtrless 0$, so that $\text{cov}(\Delta_k\Phi^*_{z,t}, \Delta_k\Gamma^*_{z,t}) = 0$.

Table 6 reports the results. In our calibrated dynamic model, the contribution of the granular component is over 24% for annual changes and 23% for 10-year changes, relative to 18% in the cross section (recall table 5). Therefore, the contribution of granularity to the dynamics of comparative advantage is greater than to its long-run variation across sectors. This is because the granular component, driven by relatively large idiosyncratic firm-level shocks, moves faster than the fundamental component, which is driven by less volatile sectoral shocks.[49] As we increase the horizon of the variance decomposition, the role of the granular component gradually declines but still stays above 22% even 50 years out.

2. Reversals in Comparative Advantage

Beyond the variance decomposition of changes in export intensity, one can study the predictive ability of granularity for the future changes in export patterns. Motivated by the recent empirical findings of pronounced reversals in comparative advantage and by the reduced-form predictive patterns we document in our data in section III, we explore in particular the relative contribution of aggregate and idiosyncratic shocks to mean reversion in comparative advantage.

In a recent paper, Hanson, Lind, and Muendler (2018) emphasize two striking patterns: (1) hyperspecialization of exports—with about 100 sectors in their data, a single top sector accounts for 21% of a country's total exports on average across countries, while the top three sectors account for 43%; and (2) high turnover of comparative advantage—a top five sector in terms of export intensity has about a 50–50 chance of staying among the top five 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 over the medium run. We explore here the extent to which our granular model can capture this combination of cross-sectional and dynamic patterns.

We report the results in table 6, where we also summarize the stylized facts from Hanson, Lind, and Muendler (2018) and the corresponding moments in our French data. Note that for France, the empirical patterns are somewhat less extreme than for an average country in Hanson, Lind, and Muendler (2018).[50] First, the top 1% and top 3% shares in aggregate exports are somewhat lower, equal to 17% and 30%, respectively. Second, the turnover ratio over the 10 years available in our panel is also

[49] The dynamic contribution of granularity is even greater in the robustness calibration (see n. 47), explaining more than one-third of the overall dynamics of comparative advantage.

[50] Hanson, Lind, and Muendler (2018) show that small, developing countries exhibit more extreme patterns of both specialization and mean reversion, and foreign direct investment likely plays an important role in this (e.g., the closure of the Intel plant in Costa Rica).

somewhat more moderate: 80% of the sectors in the top 5% stay in the top 5% a decade later. Nonetheless, qualitatively, the patterns are similar.

The dynamic granular model fits well both the cross-sectional and time series patterns observed in the French data. Specifically, the granular model reproduces the high concentration of sectoral exports as well as accounts for fast turnover of top comparative advantage sectors observed in France. A sector in the top 5% in terms of export intensity has a 76% chance to remain in the top 5% one decade later and only a 62% chance after another decade. This goes a substantial way toward reconciling the findings in Hanson, Lind, and Muendler (2018).

To quantify the contribution of granularity to these patterns, we rerun our model with granular shocks only—that is, we shut down the sectoral shocks by setting $\alpha_v = 0$. First, note that this counterfactual model still fits well the dynamics of firm-level sales shares, as captured by the mean reversion coefficient β_s, yet falls short on mean reversion in sectoral exports, β_X. Nonetheless, this model goes a considerable way in explaining the turnover at the very top of export-intensive sectors. Specifically, granular dynamics alone accounts for a 10% (13%) probability of dropping out of the top 5% of export-intensive sectors over a 10-year (20-year) horizon, that is, a half of the total turnover that we observe in the French data.

Finally, we find that in the calibrated model, the extent of granularity can help predict future comparative advantage reversals, as we have observed earlier in the data (recall table 2). To show this, we compute the changes in sectoral export shares $\Delta_k\Lambda^*_{z,t}$ over time (over $k = 20$ and 50 years) and compare it with the initial strength of granularity in these sectors at $t = 0$. We report the results in figure 3, grouping sectors by deciles of initial granular residual $\Gamma^*_{z,0}$. Panel A reports the average 20- and

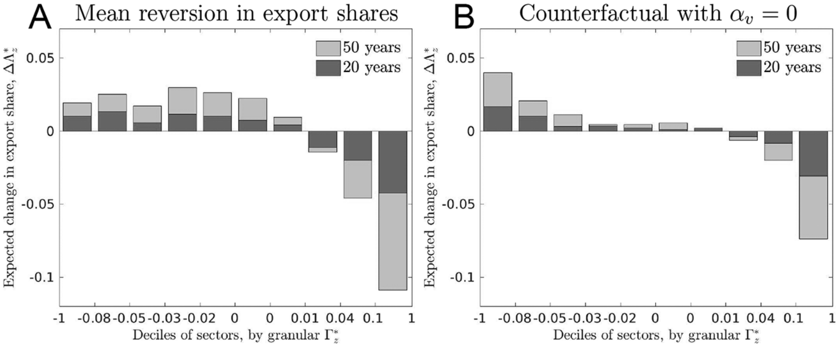

FIG. 3.—Comparative advantage reversals. Simulated equilibrium path of the calibrated granular model (*A*) and its counterfactual version with $\alpha_v = 0$ (*B*). Sectors are sorted into deciles on $\Gamma_{z,0}$ in the initial period, and within-decile averages are reported for $\Lambda_{z,k} - \Lambda_{z,0}$, $k = 20$ and 50 (years). A color version of this figure is available online.

50-year changes in sectoral export intensities. Clearly, the strongest mean reversion forces are at play in the most granular sectors, which tend to lose export shares over time.[51] Panel B of figure 3 plots the same predicted mean reversion patterns in the counterfactual model with granular shocks alone. The bulk of the mean reversion predicted by the full model is due to the granular forces. In particular, in the top decile of sectors, where mean reversion is most pronounced, granularity accounts for 75% of mean reversion over 20 years and 70% over 50 years. The dynamic granular model therefore rationalizes our empirical findings from table 2 that sectors with a stronger concentration at the top tend to mean revert faster in their aggregate exports.[52] This suggests that granular firm dynamics is a key contributor to comparative advantage reversals at the aggregate level.

3. Additional Consequences of Granularity

We finish this section with a brief discussion of two addition dynamic implications of granular comparative advantage, and we refer the reader to the earlier draft for details (see Gaubert and Itskhoki 2018). First, we point out that granularity has implications for intersectoral reallocation of labor. In a granular open economy, firm-level shocks generate production and labor reallocation not only within sectors but also across sectors. In our quantified model, the annual job creation and job destruction rates are about 12%, of which about one-fifth (2.4%) is due to intersectoral job reallocation, reflecting shifts 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). It is also interesting to note that the share of intersectoral labor reallocation in the overall job flows is very sensitive to the degree of openness of the economy: in particular, it falls sixfold when the economy goes to autarky. The interaction between granularity and openness thus contributes to the increased volatility of equilibrium reallocation, which may be costly in frictional economies (cf. Rodrik 1998).

Second, we note that in the granular model, the exit of a single firm in a given industry can have a marked impact on the country's comparative advantage—an effect absent in models with a continuum of firms. In our

[51] An average sector in the top decile of granularity, $\Gamma^{*}_{z,0}$, is expected to lose on average about 4 (11) percentage points of export intensity over 20 (50) years. While large, these patterns are highly volatile, with a typical standard deviation of around 13 (18) percentage points; over 20 years, there is a 10% chance that a sector in the top decile loses 21 percentage points or more of export intensity or gains 12 percentage points or more.

[52] Table A2 reproduces the empirical regression from tables 1 and 2 on a model-simulated data set for the baseline and robustness dynamic calibrations. The model captures accurately these empirical patterns, with somewhat smaller point estimates relative to table 2.

quantified model, the largest exporter accounts on average for over a quarter of total sectoral exports. If such a large exporter fails and exits, its market share is redistributed to both home and foreign firms. The reallocation of this lost market share toward foreign firms reflects a loss in comparative advantage. In the most granular sectors, our model predicts that over 70% of a firm's sales share is lost to the foreign competitors, resulting in a sharp loss of comparative advantage. A single large firm leaving the industry can flip the sector from comparative advantage into disadvantage.

VIII. Granular Policies in an Open Economy

We conclude with a brief outline of the consequences of granularity for policy, which we explore further in Gaubert, Itskhoki, and Vogler (2020). 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 tool kit that policy makers use to affect foreign market access (see, e.g., Bagwell and Staiger 2004, chap. 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 economic motivation and the international spillovers of such policies.

Consider, first, mergers and acquisitions between large firms. They are sought by firms for a range of reasons. Some of them tend to be welfare destroying (by increasing market power and distortions in the economy), while other are desirable from a welfare standpoint (because of positive spillovers, such as cost synergies and transfer of best practices). In a closed economy, the antitrust authority maximizing domestic welfare will allow only mergers with sufficient positive spillovers to offset the increase in markups. Matters are different, however, in an open economy. The increase in markups following a merger is in part passed along to the foreign consumer, creating a negative spillover for the foreign country akin to a terms-of-trade manipulation. This suggests a rationale for policy makers to be excessively lenient toward domestic mergers in export-oriented sectors at the expense of foreign countries. Our estimated model suggests that these negative spillover effects are significant quantitatively and are particularly pronounced in the most granular and open sectors, emphasizing the need for international cooperation over mergers and acquisitions policies to avoid excessive buildup of market power.

Next, consider narrow trade restrictions and antidumping duties that target individual firms. They have been regularly emphasized in the policy debate. To capture this type of policies, we use the estimated model to study the effect of a granular import tariff imposed on the largest foreign exporter as opposed to a uniform industry-wide tariff imposed on all imports. Granular tariffs may be more attractive to policy makers because of domestic political economy considerations, though perhaps more complex to impose legally. Our quantified model suggests that granular tariffs are also more effective at extracting surplus from foreign producers and improving the home country's terms of trade. The intuition behind this result is that a granular tariff on a single large foreign exporter achieves the desired terms-of-trade improvement with a minimal associated reduction in the domestic consumer surplus due to higher import prices. Indeed, much of the granular tariff is absorbed by a reduction in the markup of the foreign exporter, as it aims to maintain its market share.[53] 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 share in the targeted sector by a small percentage and hence offering an effective way to extract producer surplus from the dominant foreign firm.

IX. 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 multisector model of trade, which combines fundamental comparative advantage across sectors with granular comparative advantage embodied in outstanding individual firms. The model, estimated using a rich set of sectoral and firm-level moments, suggests that granularity accounts 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. Extending the analysis to a dynamic setting shows that idiosyncratic firm dynamics account for a large share of the evolution of a country's comparative advantage over time, with a strong predictive ability for comparative advantage reversals observed in the data. 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 of

[53] Perhaps surprisingly, the average prices faced by the home consumers can even fall slightly, if the general equilibrium effect from the fall in relative foreign wages is stronger than the direct effect of the tariff pass-through into import prices—a version of the Metzler paradox (see Helpman and Krugman 1989, chap. 4, 7).

the origin of outstanding firms. We see this line of analysis as very fruitful for future research. In particular, it would help us 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.

Appendix A

Additional Figures and Tables

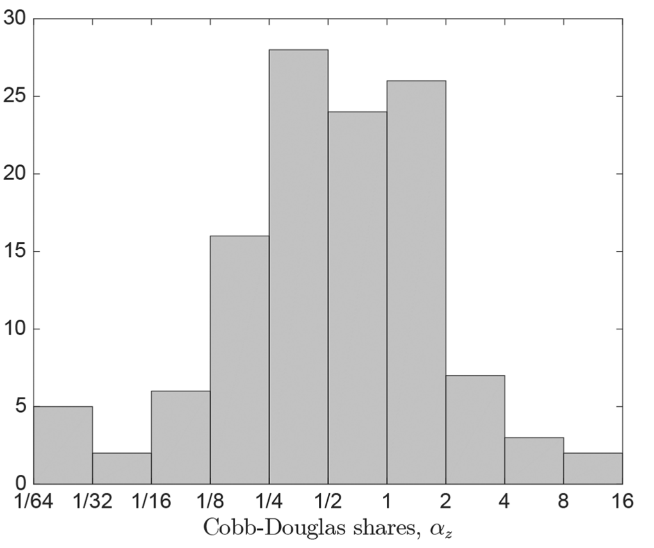

FIG. A1.—Sectoral Cobb-Douglas shares in the data, with $\alpha_z = \tilde{N}\tilde{\alpha}_z$ so that $\mathbb{E}\alpha_z = 1$, as required by a model with a continuum of sectors. A color version of this figure is available online.

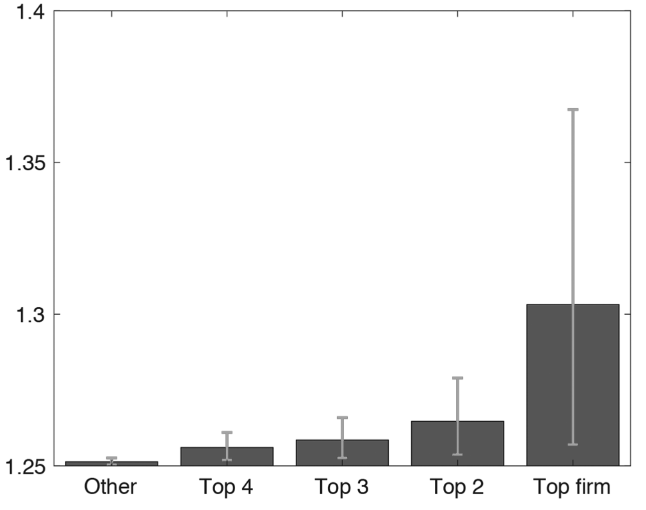

FIG. A2.—Equilibrium markups. The bars correspond to markups for the four largest French firms in each sector and for the residual fringe of French firms averaged across sectors, while the intervals correspond to the 10th–90th percentiles across sectors. Markups under monopolistic competition with continuum of firms equal $\sigma/(\sigma - 1) = 1.25$ for all firms, and this constitutes the lower bound for all markups in our oligopolistic model. A color version of this figure is available online.

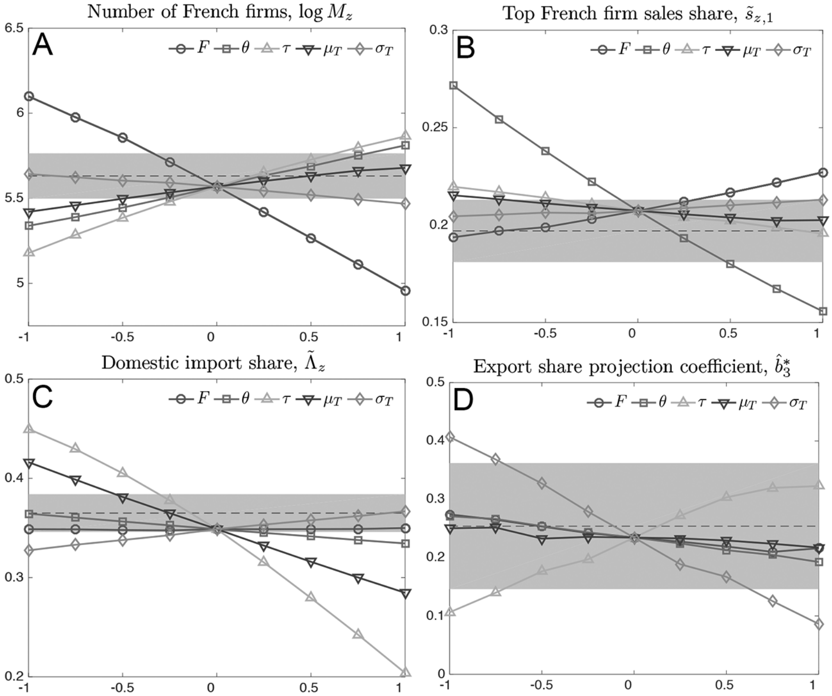

Fig. A3.—Identification plots. The lines trace out the effects of a change in one parameter at a time on select moments used in estimation: mean log number of French firms, $\log \tilde{M}_z$ (moment 1 in table 4; *A*); mean top French firm domestic sales share relative to all French firms, $\tilde{s}_{z,1}$ (moment 3; *B*); average foreign share in the home market, $\tilde{\Lambda}_z$ (moment 7; *C*); and regression coefficient of export share $\tilde{\Lambda}_z^{*\prime}$ on top three firms relative sales share in the home market $(\Sigma_{j=1}^{3}\tilde{s}_{z,j})$, b_3^* (moment 13; *D*). 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 $\hat{\Theta}$: (1) for F we use a log grid on $[\hat{F}/2, 2\hat{F}]$; (2) for θ we use a linear grid such that $\kappa = \theta/(\sigma - 1)$, where $\sigma = 5$, ranges on $\hat{\theta}/(\sigma - 1) \pm 0.125 \approx [0.95, 1.2]$; (3) for $\tau - 1$ we use a log grid such that it varies on $[(\hat{\tau} - 1)/2, 2(\hat{\tau} - 1)] \approx [0.15, 0.7]$; (4), (5) for μ_T and σ_T we use linear grids on $\hat{\mu}_T \pm 0.4$ and $\hat{\sigma}_T \pm 0.4$, respectively. See section V.C for interpretation. A color version of this figure is available online.

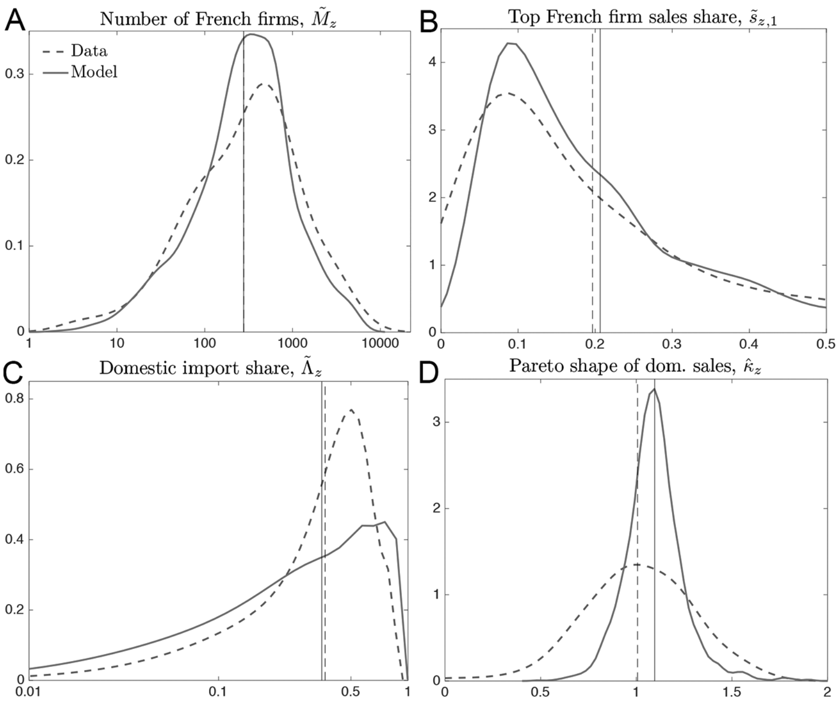

FIG. A4.—Distributions across sectors: model and data. *A*, Corresponds to moments 1 and 2 in table 4. *B*, Corresponds to moments 3 and 4. *C*, Corresponds to moments 7 and 8. *D*, Moments are not directly targeted in the baseline estimation (see table 5). In *B*, the top French firm market share is relative to other French firms in the domestic market. Pareto shapes $\hat{\kappa}_z$ are estimated according to equation (20) 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 color version of this figure is available online.

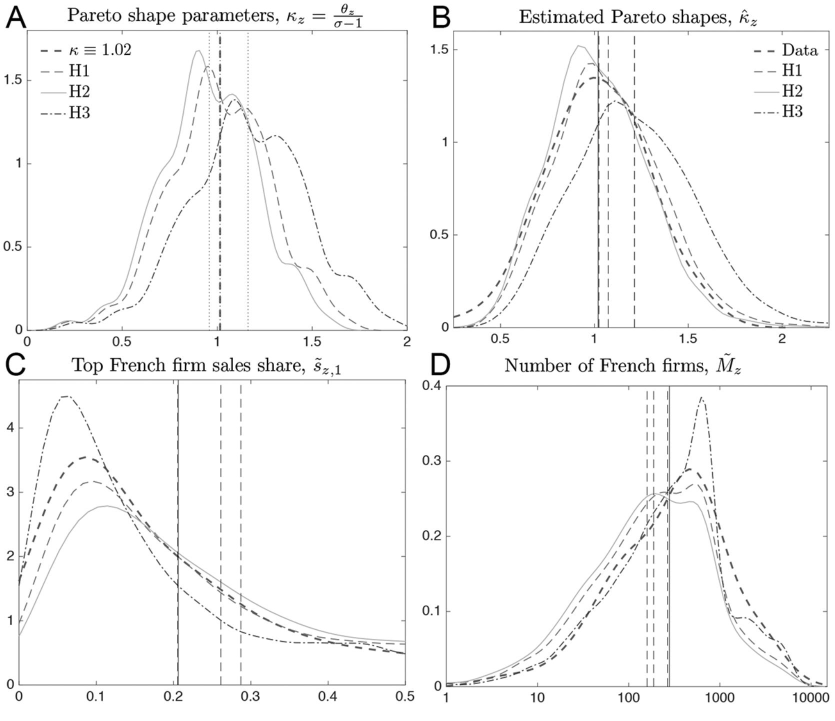

FIG. A5.—Distributions across sectors: different model specifications with heterogeneous θ_z. *B–D* correspond to *D*, *B*, and *A* in figure A4. *A*, Kernel densities of the model parameter $\kappa_z = \theta_z/(\sigma - 1)$. Each plot considers three specifications with heterogeneous sector-specific θ_z, as described in table 5, which we denote H1–H3. H1 matches average $\kappa_z = 1.02$. H2 matches average estimated Pareto shapes $\hat{\kappa}_z = 1.02$. H3 matches average top market share $\tilde{s}_{z,1} = 0.21$. A color version of this figure is available online.

TABLE A1

Granularity and Exports: Robustness with $\tilde{s}_{z,(1)}$

	Cross Section, 2005		Panel, 1997–2007		Dynamic Regressions			Out of Sample, $\log X_{z,+10} - \log X_z$		
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(2′)	(3′)	(4′)
$\log X_z$									−.119***	−.650***
									(.039)	(.338)
$\tilde{s}_{z,(1)}$	.490	.531	.695*	.664	.496***	.435***	.432***	−.599***	−.489***	
	(.371)	(.372)	(.407)	(.407)	(.111)	(.112)	(.112)	(.256)	(.255)	
$\log D_z$	.867***	.901***	.880***	.917***				−.029	.085	.588*
	(.050)	(.053)	(.052)	(.053)				(.036)	(.051)	(.308)
Sector fixed effects		2-digit		2-digit	4-digit		2-digit	2-digit	2-digit	2-digit
Year fixed effects			Yes	Yes	Yes					
Observations	300	300	3,300	3,300	3,300	3,000	3,000	300	300	300
R^2_{adj}	.507	.610	.510	.641	.950	.009	.008	.063	.091	. . .

Note.—Shown is the robustness table, which replaces $\Sigma_{i=1}^{3}\tilde{s}_{z,(i)}$ with $\tilde{s}_{z,(1)}$ and reruns the specifications in tables 1 and 2. Columns 1–7 correspond to cols. 1–7 of table 1, and cols. 2′–4′ correspond to cols. 2–4 of table 2. The right-hand-side variable is $\log X_z$ in cols. 1–5, $\Delta \log X_z$ in cols. 6 and 7 (with right-hand-side variable in first differences as well), and 10-year forward difference $\log X_{z,+10} - \log X_z$ in cols. 2′–4′.

* Significant at the 10% level.

*** Significant at the 1% level.

TABLE A2

GRANULARITY AND EXPORTS IN MODEL: SIMULATED DATA

	Table 1 Panel: log X_z				Table 2 out of Sample: $\log X_{z,+10} - \log X_z$			
	(1)	(3)	(5)	(6)	(1)	(2)	(3)	(4)
	A. Baseline: $\alpha_v = .034$, $\alpha_u = .050$							
log X_z					−.115		−.110	−.140
					[−.143, −.061]		[−.145, −.069]	[−.377, −.036]
$\Sigma_{i=1}^{3} \tilde{s}_{z,(i)}$	1.083	1.144	.601	.653		−.168	−.031	
	[.909, 1.351]	[.940, 1.342]	[.394, .845]	[.525, .871]		[−.396, .030]	[−.214, .132]	
log D_z	1.274	1.282			.111	−.043	−.101	−.123
	[1.255, 1.311]	[1.255, 1.318]			[.034, .166]	[−.080, −.015]	[.031, .046]	[−.062, .380]
	B. Robustness: $\alpha_v = .029$, $\alpha_u = .060$							
log X_z					−.100		−.092	−.188
					[−.133, −.056]		[−.127, −.079]	[−.438, .079]
$\Sigma_{i=1}^{3} \tilde{s}_{z,(i)}$	1.058	1.126	1.017	1.078		−.227	−.153	
	[.904, 1.321]	[.945, 1.340]	[.844, 1.189]	[1.006, 1.248]		[−.462, −.089]	[−.292, .039]	
log D_z	1.275	1.280			.099	−.036	.075	.200
	[1.256, 1.306]	[1.251, 1.315]			[.032, .151]	[−.077, −.013]	[.013, .118]	[−.083, .461]

NOTE.—Empirical regressions from tables 1 and 2 are implemented on the model-simulated data set (column numbers correspond to the respective empirical specifications). The table reports median point estimates over 20 simulated data sets, with [10%, 90%] ranges of point estimates across simulations in square brackets. For panel regressions in table 1, specification 5 includes sectoral fixed effects, and specification 6 is in first differences over time; for out-of-sample predictive regressions in table 2, specification 4 is an instrumental variable second stage, where the first stage is equivalent to the cross-sectional regression 1 in table 1. Panel A corresponds to the baseline calibration (see table 6); panel B corresponds to the robustness calibration (see n. 47).

TABLE A3
Robustness to Distribution Assumptions: Additional Details

	Data	Baseline	(R1)	(R2)	(R3)
Granular accounting (%):					
Granular contribution		17.8	17.8	1.2	18.2
Export share contribution		54.4	48.5	85.0	54.8
Estimated parameters:					
θ (θ' in R2)		4.382	4.302	.532	
τ		1.342	1.281	1.560	
F ($\times 10^5$)		1.179	.771	.310	
μ_T		.095	.377	.154	
σ_T (σ'_T in R1)		1.394	1.467	.801	
L^*/L		1.932	1.447	1.377	
Moments:					
1. Log number of firms, $\log \tilde{M}_z$	5.631	5.429	5.709	**5.348**	5.470
3. Top firm sales share, $\tilde{s}_{z,1}$	.197	.205	.207	**.174**	.201
5. Top three sales share, $\Sigma_{j=1}^{3} \tilde{s}_{z,j}$	.356	.343	.344	.326	.336
7. Imports/domestic sales, $\tilde{\Lambda}_z$	.365	.354	.350	.364	.335
9. Exports/domestic sales, $\tilde{\Lambda}_z^{*\prime}$	.328	.345	.348	.348	.314
11. Fraction of sectors with $\tilde{X}_z > \tilde{D}_z^*$	.185	.095	.059	.122	.040
Regression coefficients:					
12. Export share on top firm share, $\hat{b}_1^*$	.215	.234	.257	**.012**	.235
13. Export share on top three share, $\hat{b}_3^*$	.254	.222	.219	**.013**	.228
14. Import share on top firm share, $\hat{b}_1$	−.016	−.011	.010	**.139**	−.044
15. Import share on top three share, $\hat{b}_3$	.002	.008	.040	**.146**	−.032
Overall loss function	. . .	.067	.077	.170	.086

Note.—Additional results behind the robustness checks are in table 5 (which reproduces the counterfactual variance decompositions), where R1 is fat-tailed T_z/T_z^*, R2 is log-normal $\varphi_{z,i}$, and R3 is nongranular foreign (under the baseline parametrization). We report the parameter estimates (as in table 3) and select moment fits (as in table 4; numbered rows). The overall loss function is in units of the average proportional deviation from the empirical moments, namely, $[(1/15)\mathcal{L}(\Theta)]^{1/2}$, for a given model specification, with $\mathcal{L}(\Theta)$ defined in app. C; e.g., the baseline specification misses the average moment by 6.7%, while R2 misses by 17%. In col. R2, boldface indicates the poor-fit moments by the corresponding specification of the model.

Appendix B

Theory Appendix

B1. Continuous DFS-Melitz Model

We review here the continuous model, which serves as a benchmark in our granular analysis. Consider a two-country multisector 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).[54] 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

[54] This model extends Melitz (2003) in a multisector way, the same way that Costinot, Donaldson, and Komunjer (2012) extend the Eaton and Kortum (2002) model. Other papers that considered a multisector DFS-Melitz environment, albeit under somewhat different formulation, include Okubo (2009) and Fan, Lai, and Qi (2019).

which an exogenous mass of firms $\overline{M}_z$ are present and their productivities are drawn from a Pareto distribution with a sector-specific lower bound $\underline{\phi}_z$ and a shape parameter θ common across all sectors. We show below that in this model, the overall sectoral productivity is determined by $T_z = \overline{M}_z \cdot \underline{\varphi}_z^\theta$, as in equation (8). 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 IV, with the exception that instead of equation (4), the sectoral CES consumption bundles aggregate over a continuum of individual varieties ω:

$$Q_z = \left[\int_{\omega\in\Omega_z} q_z(\omega)^{(\sigma-1)/\sigma} \mathrm{d}\omega\right]^{\sigma/(\sigma-1)}, \tag{B1}$$

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 IV, with the exception that $\overline{M}_z$ is a deterministic mass of existing shadow firms in each sector, with individual productivities $\varphi_z(\omega) \sim iid\text{Pareto}(\theta, \underline{\varphi}_z)$ with $\mathbb{P}\{\varphi_z(\omega) \le \varphi\} = 1 - (\underline{\varphi}_z/\varphi)^\theta$ 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

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

where we substituted the markup pricing rule over the marginal cost into the expression for revenues (5), and we use the notation $[x]^+ \equiv \max\{0, x\}$.[55] Firms with sufficiently high productivities profitably enter the home and the foreign markets, 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 we rewrite profits as

$$\begin{cases} \pi_z(\omega) = wF\left[\left(\dfrac{\varphi_z(\omega)}{\varphi_{h,z}}\right)^{\sigma-1} - 1\right]^+ + w^*F^*\left[\left(\dfrac{\varphi_z(\omega)}{\varphi_{f,z}}\right)^{\sigma-1} - 1\right]^+, \\ \varphi_{h,z} = \dfrac{\sigma}{\sigma-1}\dfrac{w}{P_z}\left(\dfrac{\sigma wF}{\alpha_z Y}\right)^{1/(\sigma-1)}, \varphi_{f,z} = \dfrac{\sigma}{\sigma-1}\dfrac{\tau w}{P_z^*}\left(\dfrac{\sigma w^*F^*}{\alpha_z Y^*}\right)^{1/(\sigma-1)}. \end{cases} \tag{B2}$$

[55] Specifically, a home firm sets $p_z(\omega) = [\sigma/(\sigma-1)](w/\varphi(\omega))$ in the home market, which results in revenues $(p_z(\omega)/P(z))^{1-\sigma}\alpha_z Y$, according to eq. (5), and the operating profits equal fraction $1/\sigma$ of these revenues because of 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

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.

B1.1. Sectoral Equilibrium

Using the definition of the price index, the markup pricing rules, the cutoff definitions in equation (B2), and the Pareto productivity distribution, we can integrate to solve for the price index in sector z in the home market:

$$P_z = \frac{\sigma}{\sigma - 1} w \left[\frac{\kappa}{\kappa - 1} \frac{T_z}{1 - \Phi_z}\right]^{-1/\theta} \left(\frac{\sigma w F}{\alpha_z Y}\right)^{(\kappa-1)/\theta}, \tag{B3}$$

where $\kappa \equiv \theta/(\sigma - 1)$ and Φ_z is the foreign share, as defined in equation (15).[56] The sectoral price index in equation (B3) 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 Arkolakis, Costinot, and Rodríguez-Clare 2012). Using equation (B3), 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:

$$X_z = \alpha_z \Phi_z^* Y^*, X_z^* = \alpha_z \Phi_z Y, \tag{B4}$$

and sectoral net exports is $NX_z = X_z - X_z^*$. In addition, we also characterize the allocation of aggregate labor supply to sector z, which in the home market satisfies

$$wL_z = \alpha_z Y \left[\frac{\sigma\kappa - 1}{\sigma\kappa}\left(1 - \Phi_z\right) + \frac{\kappa - 1}{\sigma\kappa}\Phi_z\right] + \alpha_z Y^* \frac{\sigma - 1}{\sigma}\Phi_z^*. \tag{B5}$$

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.[57] Combining equations (B4) and (B5) with equation (B8) below, we obtain the relationship between sectoral net exports and labor allocation:

$$\frac{L_z}{L} = \alpha_z + \frac{\theta}{\sigma\kappa - 1}\frac{NX_z}{Y}. \tag{B6}$$

[56] We note that the foreign share in eq. (14) 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 that 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.

[57] A fraction $(\sigma - 1)/\sigma$ of revenues goes to cover variable production labor costs (in the country of production). Integrating across firms, a fraction $(\kappa - 1)/\sigma\kappa$ of revenues goes to cover entry labor costs (in the country of entry). Note that the first term in eq. (B5) can be decomposed as $(\sigma\kappa - 1)/\sigma\kappa = [(\sigma - 1)/\sigma] + [(\kappa - 1)/\sigma\kappa]$. The remaining $1/\sigma\kappa$ share is net profits.

In autarky, $L_z = \alpha_z L$ because of the Cobb-Douglas preferences, yet in the open economy, labor reallocates toward 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 that countries' budget balances equation (11) by Walras's 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 $[(\kappa - 1)/\sigma\kappa]\alpha_z \Phi_z^* Y^*$, which is proportional to exports $X_z = \alpha_z \Phi_z^* Y^*$, with a constant factor $(\sigma\kappa - \kappa + 1)/\sigma\kappa$ in front. Aggregating across sectors and equalizing with the foreign export income, we obtain the balanced current account (and trade balance condition):

$$Y \int_0^1 \alpha_z \Phi_z \mathrm{d}z = Y^* \int_0^1 \alpha_z \Phi_z^* \mathrm{d}z. \tag{B7}$$

Next, aggregating sectoral labor demand in equation (B5) across z and using trade balance (B7), we obtain aggregate labor market clearing:

$$wL = \frac{\sigma\kappa - 1}{\sigma\kappa} Y, \; w^* L^* = \frac{\sigma\kappa - 1}{\sigma\kappa} Y^*. \tag{B8}$$

Therefore, total labor income is a constant share of GDP (total income), with the complementary share coming from firm profits. Combining equation (B7) with equation (B8) and normalizing $w = 1$ allows us to solve for (w^*, Y, Y^*), completing the description of the general equilibrium.[58]

B1.2. 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 equation (14) 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

$$\frac{a_{\underline{z}}}{a_{\underline{z}}^*} = \frac{w}{\tau w^*}, \; \frac{a_{\bar{z}}}{a_{\bar{z}}^*} = \frac{\tau w}{w^*}, \tag{B9}$$

[58] Taking the ratio of the two equations in (B8), we have $Y/Y^* = (wL)/(w^* L^*)$, which together with eq. (B7) allows us to solve for both relative wage w/w^* and relative incomes Y/Y^*, as in the DFS model. Recall from eq. (14) that Φ_z and Φ_z^* can be written as a function of relative wages w/w^* and the exogenous parameters of the model.

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 nontraded and produced in both countries, and for goods $z \in (\bar{z}, 1]$, home is the only world supplier.

B1.3. Continuous Limit

Last, we discuss how the granular model of section IV admits the continuous DFS-Melitz limit described above. We introduce a scaler $M > 0$ and rewrite the price index in equation (9) and the market share in equation (5) as follows:

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

where the granular model of section IV corresponds to the case with $M = 1$. Note that $\Sigma_{i=1}^{K_z}\tilde{s}_{z,i} = 1$ for any $M > 0$. We also rewrite the utility in equation (4) as

$$\tilde{Q}_z = \left[\frac{1}{M}\sum_{i=1}^{K_z}\tilde{q}_{z,i}^{(\sigma-1)/\sigma}\right]^{\sigma/(\sigma-1)},$$

where $\tilde{q}_{z,i} = Mq_{z,i}$ are the new consumption units. Last, the derived productivity parameter in equation (8) is generalized as $T_z = \overline{M}_z/M \cdot \underline{\varphi}_z^{\theta}$.

With this generalization to an arbitrary $M > 0$, we can now take the following limit: $M, \overline{M}_z \to \infty$ and $F \to 0$, such that $\overline{M}_z/M = const$ for all z and $MF = const$, holding constant the other parameters of the model, including the location of the productivity distribution $\underline{\phi}_z$. This keeps T_z unchanged. Furthermore, $\overline{M}_z/M$ now represents the relative measure of shadow firms in sector z. The ratio $K_z/\overline{M}_z$ tends to a constant related to productivity cutoffs (B2) in the continuous model; the price index P_z tends to a constant, the price level in the continuous model (B3); the market shares $s_{z,i} \to 0$ so that the elasticity in equation (9), $\varepsilon_{z,i} \to \sigma$, and markups become constant equal to $\sigma/(\sigma - 1)$; and the nonnegativity of profits in equation (10) with $F \to 0$ at the same rate as $s_{z,i} \to 0$ now corresponds to the cutoff condition in equation (B2). 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 IV.

B2. Derivations and Proofs for the Granular Model

B2.1. Foreign Share

Consider the foreign share Λ_z defined in equation (14). 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 on both the price-setting and the entry outcomes (e.g., see eqq. [5], [9], [10]). Nonetheless, following Eaton, Kortum, and Sotelo (2012), 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; that is, 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 $s > 0$, $\mathbb{E}_T\{\iota_{z,i}|s_{z,i} > s\} = \mathbb{P}_T\{\iota_{z,i} = 1|s_{z,i} > s\} = T_z w^\theta/[T_z w^\theta + T_z^*(\tau w^*)^\theta] = 1 - \Phi_z$, as defined in equation (15). 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. $\mathbb{E}_T\Lambda_z = \Sigma_{i=1}^{K_z}\mathbb{E}_T\{(1-\iota_{z,i})s_{z,i}\} = \Sigma_{i=1}^{K_z}\mathbb{E}_T\{s_{z,i}\cdot\mathbb{E}_T\{1-\iota_{z,i}|s_{z,i}\}\} = \Phi_z\Sigma_{i=1}^{K_z}\mathbb{E}_T s_{z,i} = \Phi_z$, since $\mathbb{E}_T\{\Sigma_{i=1}^{K_z}s_{z,i}\} = \mathbb{E}_T\{1\} = 1$, and where the third equality uses property 1.

Property 1 follows 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 eq. [7]). 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 the following:

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

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 eqq. [15], [16]).

B2.2. 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 equations (10) and (12), we write the home country budget $Y = wL + \Pi$ constraint as

$$Y = wL + Y(1-\Lambda)\frac{\bar{\mu}_H - 1}{\bar{\mu}_H} - wFK_H + Y^*\Lambda^*\frac{\bar{\mu}_H^* - 1}{\bar{\mu}_H^*} - w^*F^*K_H^*, \qquad \text{(B10)}$$

where

$$
\begin{aligned}
K_H &= \int_0^1 \left[\sum_{i=1}^{K_z} \iota_{z,i}\right] \mathrm{d}z, \\
K_H^* &= \int_0^1 \left[\sum_{i=1}^{K_z^*} (1 - \iota_{z,i}^*)\right] \mathrm{d}z, \\
1 - \Lambda &= \int_0^1 \alpha_z (1 - \Lambda_z) \mathrm{d}z = \int_0^1 \alpha_z \left[\sum_{i=1}^{K_z} \iota_{z,i} s_{z,i}\right] \mathrm{d}z, \\
\Lambda^* &= \int_0^1 \alpha_z \Lambda_z^* \mathrm{d}z = \int_0^1 \alpha_z \left[\sum_{i=1}^{K_z^*} (1 - \iota_{z,i}^*) s_{z,i}^*\right] \mathrm{d}z, \\
\frac{1}{\bar{\mu}_H} &= \frac{1}{1 - \Lambda} \int_0^1 \alpha_z \left[\sum_{i=1}^{K_z} \iota_{z,i} \frac{s_{z,i}}{\mu(s_{z,i})}\right] \mathrm{d}z, \\
\frac{1}{\bar{\mu}_H^*} &= \frac{1}{\Lambda^*} \int_0^1 \alpha_z \left[\sum_{i=1}^{K_z^*} (1 - \iota_{z,i}^*) \frac{s_{z,i}^*}{\mu(s_{z,i}^*)}\right] \mathrm{d}z,
\end{aligned}
$$

where $\mu(s) = \varepsilon(s)/(\varepsilon(s) - 1)$ and $\varepsilon(s) = \sigma(1 - s) + s$, as defined in equation (9). Note the following:

1. K_H and K_H^* are the total numbers of the home firms selling in the home and foreign markets, respectively, across all industries;
2. $1 - \Lambda$ and Λ^* are the average shares of the home firm sales in aggregate home and foreign expenditure Y and Y^*, respectively; and
3. $\bar{\mu}_H$ 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 $(\mu(s_{z,i}) - 1)/\mu(s_{z,i}) = (p_{z,i} - c_{z,i})/p_{z,i}$ for a firm with market share $s_{z,i}$.

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

$$
Y^* = w^* L^* + Y^*(1 - \Lambda^*) \frac{\bar{\mu}_F^* - 1}{\bar{\mu}_F^*} - w^* F^* K_F^* + Y\Lambda \frac{\bar{\mu}_F - 1}{\bar{\mu}_F} - wFK_F, \qquad \text{(B11)}
$$

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 (13), which we write as

$$
wL = wFK + Y(1 - \Lambda) \frac{1}{\bar{\mu}_H} + Y^* \Lambda^* \frac{1}{\bar{\mu}_H^*}, \qquad \text{(B12)}
$$

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

$$w^*L^* = w^*F^*K^* + Y^*(1-\Lambda^*)\frac{1}{\bar{\mu}_F^*} + Y\Lambda\frac{1}{\bar{\mu}_F}, \tag{B13}$$

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 (B10)–(B13) 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 $\{K_z, K_z^*, \{s_{z,i}\}_{i=1}^{K_z}, \{s_{z,i}^*\}_{i=1}^{K_z^*}\}_{z\in[0,1]}$.
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 (e.g., eq. [B11]) 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 (B12) into the budget constraint (B10) (or, equivalently, [B13] into [B11]) results in the current account balance condition (which in general differs from the trade balance $NX = \Lambda^*Y^* - \Lambda Y$):
$$\Lambda Y - wFK_F = Y^*\Lambda^* - w^*F^*K_H^*. \tag{B14}$$
The equilibrium system can be represented by the system of three linearly independent equations (B12)–(B14). Note the similarity and differences of this equilibrium system with a corresponding system in the continuous model (B7) and (B8). In particular, because of 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

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

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 (B10) and (B14) can be rewritten as

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

which corresponds to the expression in footnote 24. Last, note that in a closed economy $\Phi = \Phi^* = 0$, and therefore the country budget constraint (B10) becomes $Y = \bar{\mu}w[L - FK]$, as we have it in footnote 20.

Appendix C

Estimation Appendix

The detailed estimation procedure is as follows:

1. For given parameter values of μ_T and σ_T, we draw N relative sectoral productivities T_z from the lognormal distribution (recall our normalization $T_z^* \equiv 1$).[59] 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 $z = 1 \ldots N$, we draw productivities of potential entrants $\{\varphi_{z,j}\}_{j=1}^{M_z}$ in a manner consistent with the distributional assumptions of the model.[60] 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 $\mathbf{Z} \equiv \{K_z, \{s_{z,i}\}_{i=1}^{K_z}, K_z^*, \{s_{z,i}^*\}_{i=1}^{K_z^*}\}$, as described in section IV.[61]

[59] The numerical procedure necessarily simulates a finite number of sectors and not a continuum, so that the law of large numbers applies only approximately across sectors. Increasing the number of simulated sectors helps limit small-sample deviation from the model-based deterministic quantities. We increased the number of sectors in the simulated sample from 119 to four times that (476) to limit the dependency of general equilibrium quantities to the small-sample randomness of draws, striking a balance with computational feasibility of estimation. In particular, with 476 sectors, the variation in the general equilibrium quantities across simulations is less than 1% of their median values. We check that the average moments calculated using this sample approximate closely their population means. When computing conditional moments by finer bins of sectors to build counterfactual figs. 1–3, we draw a larger number of sectors (10,000) in order to increase precision.

[60] Specifically, we follow Eaton, Kortum, and Sotelo (2012) 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 Frechet (θ, 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 $U_{z,1}$, $(U_{z,2} - U_{z,1}),(U_{z,3} - U_{z,2})$, ... are i.i.d. exponential with cumulative distribution function $G_U(u) = 1 - e^{-u}$. 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 the rest of the world. For smaller sectors (in terms of Cobb-Douglas shares), we use 700 and 1,400 draws 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.

[61] 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 eq. [9] 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 $\hat{K}_z$

iii. Given **Z** and the normalization $L = 100$, use the general equilibrium conditions (11) and (13)) to solve for the new values of Y and Y^*. Note that these equations are 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 (13) (namely, eq. [B13]) 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 $z = 1 \dots N$, 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 4, and the associated SMM loss function:

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

where $\overline{\mathcal{M}}(\Theta) \equiv (1/N)\Sigma_{z=1}^{N}\mathcal{M}_z(\Theta)$, $\tilde{\mathbf{m}}$ are the values of the moments in our empirical data set, and $\mathbf{W} = \text{diag}\{\mathbf{w}\}$ is the weighting matrix, which we chose to be diagonal and inversely proportional to $\tilde{\mathbf{m}}^2$.[62] Table 4 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 $w_k(\overline{\mathcal{M}}_k(\hat{\Theta}) - \tilde{m}_k)^2/\mathcal{L}(\hat{\Theta})$, 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 efficiently 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 encompass 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, that is, 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)

firms, we then solve for additional entry ΔK_z, assuming that 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.

[62] 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 4 for the list of moments). For moments 14 and 15, with empirical values close to zero, **W** uses the values of the symmetric moments 12 and 13. Finally, we downweight 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.

the point of local convergence with the lowest loss function, $\hat{\Theta} = \text{argmin}_\Theta \mathcal{L}(\Theta)$.

C1. 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 $m_i(\Theta) = \overline{\mathcal{M}}(\Theta) - \widetilde{\mathbf{m}}_i$ is the moment function such that $(1/\tilde{N})\Sigma_{i=1}^{\tilde{N}} m_i(\Theta) = \overline{\mathcal{M}}(\Theta) - \widetilde{\mathbf{m}}$, 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:

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

Furthermore, note that $\overline{\mathcal{M}}(\Theta)$ are model-evaluated moments, which do not contribute to the sample variation in $m_i(\Theta)$.[63] Thus, all sample variation emerges from the empirically measured moments $\widetilde{\mathbf{m}}_i$ 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:

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

where V_Θ is the asymptotic sandwich-form variance matrix, $J = \mathbb{E}\{\partial m_i(\Theta)/\partial\Theta\}$ 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 $\widetilde{\mathbf{m}}$ 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 $\overline{\mathcal{M}}(\Theta)$, evaluated around $\Theta = \hat{\Theta}$; that is, $\hat{J} = \partial\overline{\mathcal{M}}(\hat{\Theta})/\partial\Theta$. The variance of moments matrix H depends on both $\overline{\mathcal{M}}(\hat{\Theta})$ and the data, and we calculate its estimate as follows:

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

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 $\hat{\Theta}$:

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

The standard errors for parameters in table 4 are then calculated as $s.e.(\hat{\Theta}) = (\text{diag}(\hat{V}_\Theta/\tilde{N}))^{1/2}$.

C2. Robustness to Pareto Firm Productivity Draws

We describe here the procedure for the robustness check of section VI.B, in which we replace the Pareto distribution for $\varphi_{z,i}$ draws with a thinner-tailed lognormal:

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

1. The number of shadow firms in sector z at home is a deterministic constant $M_z = \text{round}(M\alpha_z)$ proportional to expenditure size of the sector α_z for some large constant $M \gg 1$. In foreign, $M_z^* = kM_z$ for some factor $k > 1$ (reflecting the relative size of the foreign, L^*/L). These details are important only to the extent that we need to ensure that M and k are large enough that the least productive firms are never active, as in the baseline. In practice, we set $k = 1.5$ and $M = 350$ (recall that the average of α_z is one).
2. We set $\mu_z^* = 0$ for all z as a normalization and choose $\mu_z = \log(T_z/T_z^*) \sim \mathcal{N}(\mu_T, \sigma_T^2)$ for each sector in the simulation.
3. Draw firm productivities according to $\varphi_{z,i} \sim \log\mathcal{N}(\mu_z, \theta'^2)$ for M_z home shadow firms and $\varphi_{z,i}^* \sim \log\mathcal{N}(\mu_z^*, \theta'^2)$ for M_z^* foreign shadow firms.
4. Given these lognormal productivity draws, carry out the rest of the numerical solution and estimation procedure, as described above, to estimate (θ', σ_T, μ_T, τ, F). Check that least productive firms are inactive in every sector-country (adjust k and M if needed).

References

Acemoglu, Daron, Vasco M. Carvalho, Asuman Ozdaglar, and Alireza Tahbaz-Salehi. 2012. “The Network Origins of Aggregate Fluctuations.” *Econometrica* 80 (5): 1977–2016.

Amiti, Mary, Oleg Itskhoki, and Jozef Konings. 2014. “Importers, Exporters, and Exchange Rate Disconnect.” *A.E.R.* 104 (7): 1942–78.

———. 2019. “International Shocks, Variable Markups and Domestic Prices.” *Rev. Econ. Studies* 86 (6): 2356–402.

Anderson, James E., and Eric van Wincoop. 2004. “Trade Costs.” *J. Econ. Literature* 42 (3): 691–751.

Andrews, Isaiah, Matthew Gentzkow, and Jesse M. Shapiro. 2017. “Measuring the Sensitivity of Parameter Estimates to Estimation Moments.” *Q.J.E.* 132 (4): 1553–92.

Arkolakis, Costas, Arnaud Costinot, and Andrés Rodríguez-Clare. 2012. “New Trade Models, Same Old Gains?” *A.E.R.* 102 (1): 94–130.

Atkeson, Andrew, and Ariel Burstein. 2008. “Trade Costs, Pricing-to-Market, and International Relative Prices.” *A.E.R.* 98 (5): 1998–2031.

Autor, David, David Dorn, Lawrence F. Katz, Christina Patterson, and John Van Reenen. 2020. “The Fall of the Labor Share and the Rise of Superstar Firms.” *Q.J.E.* 135 (2): 645–709.

Bagwell, Kyle, and Robert W. Staiger. 2004. *The Economics of the World Trading System.* Cambridge, MA: MIT Press.

Bennedsen, Morten, Francisco Perez-Gonzalez, and Daniel Wolfenzon. 2010. “Do CEOs Matter?” Working Manuscript, Graduate School Bus., Columbia Univ.

Bernard, Andrew B., Jonathan Eaton, J. Bradford Jensen, and Samuel Kortum. 2003. “Plants and Productivity in International Trade.” *A.E.R.* 93 (4): 1268–90.

Bernard, Andrew B., J. Bradford Jensen, Stephen J. Redding, and Peter K. Schott. 2018. “Global Firms.” *J. Econ. Literature* 56 (2): 565–619.

Blanga-Gubbay, Michael, Paola Conconi, and Mathieu Parenti. 2020. “Globalization for Sale.” Working Manuscript, Dept. Econ., Univ. Libre Bruxelles.

Blonigen, Bruce A., and Thomas J. Prusa. 2008. “Antidumping.” In *Handbook of International Trade*, edited by E. Kwan Choi and James Harrigan, 251–84. Malden, MA: Wiley-Blackwell.

Bonfiglioli, Alessandra, Rosario Crinò, and Gino Gancia. 2018. "Betting on Exports: Trade and Endogenous Heterogeneity." *Econ. J.* 128 (609): 612–51.

Broda, Christian, and David Weinstein. 2006. "Globalization and the Gains from Variety." *Q.J.E.* 121 (2): 541–85.

Carvalho, Vasco M., and Xavier Gabaix. 2013. "The Great Diversification and Its Undoing." *A.E.R.* 103 (5): 1697–727.

Carvalho, Vasco M., and Basile Grassi. 2020. "Large Firm Dynamics and the Business Cycle." *A.E.R.* 109 (4): 1375–425.

Chaney, Thomas. 2008. "Distorted Gravity: The Intensive and Extensive Margins of International Trade." *A.E.R.* 98 (4): 1707–21.

Chor, Davin. 2010. "Unpacking Sources of Comparative Advantage: A Quantitative Approach." *J. Internat. Econ.* 82 (2): 152–67.

Costinot, Arnaud, Dave Donaldson, and Ivana Komunjer. 2012. "What Goods Do Countries Trade? A Quantitative Exploration of Ricardo's Ideas." *Rev. Econ. Studies* 79 (2): 581–608.

Davis, Steven J., and John Haltiwanger. 1999. "Gross Job Flows." In *Handbook of Labor Economics*, vol. 3, edited by O. Ashenfelter and D. Card, 2711–805. New York: Elsevier.

di Giovanni, Julian, and Andrei A. Levchenko. 2012. "Country Size, International Trade, and Aggregate Fluctuations in Granular Economies." *J.P.E.* 120 (6): 1083–132.

———. 2013. "Firm Entry, Trade, and Welfare in Zipf's World." *J. Internat. Econ.* 89 (2): 283–96.

di Giovanni, Julian, Andrei A. Levchenko, and Isabelle Méjean. 2014. "Firms, Destinations, and Aggregate Fluctuations." *Econometrica* 82 (4): 1303–40.

Dornbusch, Rudiger, Stanley Fischer, and Paul A. Samuelson. 1977. "Comparative Advantage, Trade, and Payments in a Ricardian Model with a Continuum of Goods." *A.E.R.* 67 (5): 823–39.

Eaton, Jonathan, and Samuel Kortum. 2002. "Technology, Geography, and Trade." *Econometrica* 70 (5): 1741–79.

———. 2010. "Technology in the Global Economy: A Framework for Quantitative Analysis." Working Manuscript.

Eaton, Jonathan, Samuel Kortum, and Sebastian Sotelo. 2012. "International Trade: Linking Micro and Macro." Working Paper no. 17864, NBER, Cambridge, MA.

Edmond, Chris, Virgiliu Midrigan, and Daniel Yi Xu. 2015. "Competition, Markups, and the Gains from International Trade." *A.E.R.* 105 (10): 3183–221.

Fan, Haichao, Edwin L.-C. Lai, and Han Steffan Qi. 2019. "Trade Liberalization and Firms Export Performance in China: Theory and Evidence." *J. Comparative Econ.* 47 (3): 640–68.

Freund, Caroline, and Martha Denisse Pierola. 2015. "Export Superstars." *Rev. Econ. and Statis.* 97 (5): 1023–32.

Gabaix, Xavier. 2009. "Power Laws in Economics and Finance." *Ann. Rev. Econ.* 1 (1): 255–94.

———. 2011. "The Granular Origins of Aggregate Fluctuations." *Econometrica* 79 (3): 733–72.

Gabaix, Xavier, and Rustam Ibragimov. 2011. "Rank − 1/2: A Simple Way to Improve the OLS Estimation of Tail Exponents." *J. Bus. and Econ. Statis.* 29 (1): 24–39.

Gabaix, Xavier, and Ralph S. J. Koijen. 2020. "Granular Instrumental Variables." Working Manuscript, Dept. Econ., Harvard Univ.

Gaubert, Cecile, and Oleg Itskhoki. 2018. “Granular Comparative Advantage.” Working Paper no. 24807, NBER, Cambridge, MA.

Gaubert, Cecile, Oleg Itskhoki, and Maximilian J. Vogler. 2020. “Government Policies in a Granular Global Economy.” Working Manuscript.

Grassi, Basile. 2017. “IO in I–O: Size, Industrial Organization, and the Input-Output Network Make a Firm Structurally Important.” Working Manuscript, Dept. Econ., Bocconi Univ.

Grossman, Gene M., and Esteban Rossi-Hansberg. 2010. “External Economies and International Trade Redux.” *Q.J.E.* 125 (2): 829–58.

Gutiérrez, Germán, and Thomas Philippon. 2017. “Declining Competition and Investment in the U.S.” Working Paper no. 23583, Cambridge, MA.

Hanson, Gordon H., Nelson Lind, and Marc-Andreas Muendler. 2018. “The Dynamics of Comparative Advantage.” Working Manuscript, Dept. Econ., Univ. California San Diego.

Head, Keith, Ran Jing, and John Ries. 2017. “Import Sourcing of Chinese Cities: Order versus Randomness.” *J. Internat. Econ.* 105 (C): 119–29.

Helpman, E., and P. R. Krugman. 1989. *Trade Policy and Market Structure*. Cambridge, MA: MIT Press.

Hortaçsu, Ali, and Chad Syverson. 2004. “Product Differentiation, Search Costs, and Competition in the Mutual Fund Industry: A Case Study of S&P 500 Index Funds.” *Q.J.E.* 119 (2): 403–56.

Hottman, Colin, Stephen J. Redding, and David E. Weinstein. 2016. “Quantifying the Sources of Firm Heterogeneity.” *Q.J.E.* 131 (3): 1291–364.

Levchenko, Andrei A., and Jing Zhang. 2016. “The Evolution of Comparative Advantage: Measurement and Welfare Implications.” *J. Monetary Econ.* 78 (C): 96–111.

Melitz, Marc J. 2003. “The Impact of Trade on Intra-Industry Reallocations and Aggregate Industry Productivity.” *Econometrica* 71 (6): 1695–725.

Neary, J. Peter. 2010. “Two and a Half Theories of Trade.” *World Econ.* 33 (1): 1–19.

———. 2016. “International Trade in General Oligopolistic Equilibrium.” *Rev. Internat. Econ.* 24 (4): 669–98.

Okubo, Toshihiro. 2009. “Firm Heterogeneity and Ricardian Comparative Advantage within and across Sectors.” *Econ. Theory* 38 (3): 533–59.

Rodrik, Dani. 1998. “Why Do More Open Economies Have Bigger Governments?” *J.P.E.* 106 (5): 997–1032.

Sutton, John, and Daniel Trefler. 2016. “Capabilities, Wealth, and Trade.” *J.P.E.* 124 (3): 826–78.