



Firm entry, trade, and welfare in Zipf's world[☆]

Julian di Giovanni^a, Andrei A. Levchenko^{b,c,*}

^a International Monetary Fund, United States

^b University of Michigan, United States

^c NBER, United States

ARTICLE INFO

Article history:

Received 31 August 2011

Received in revised form 30 June 2012

Accepted 6 August 2012

Available online 21 August 2012

JEL classification:

F12

F15

Keywords:

Zipf's Law

Welfare

Entry costs

Trade barriers

ABSTRACT

Firm size follows Zipf's Law, a very fat-tailed distribution that implies a few large firms account for a disproportionate share of overall economic activity. This distribution of firm size is crucial for evaluating the welfare impact of economic policies such as barriers to entry or trade liberalization. Using a multi-country model of production and trade calibrated to the observed distribution of firm size, we show that the welfare impact of high entry costs is small. In the sample of the 50 largest economies in the world, a reduction in entry costs all the way to the U.S. level leads to an average increase in welfare of only 3.25%. In addition, when the firm size distribution follows Zipf's Law, the welfare impact of the extensive margin of trade – newly imported goods at or near the exporting cutoff – is negligible. The extensive margin of imports accounts for only about 5.2% of the total gains from a 10% reduction in trade barriers in our model. This is because under Zipf's Law, the large, infra-marginal firms have a far greater welfare impact than the much smaller firms that comprise the extensive margin in these policy experiments. The distribution of firm size matters for these results: in a counterfactual model economy that does not exhibit Zipf's Law the gains from a reduction in fixed entry barriers are an order of magnitude larger, while the gains from a reduction in variable trade costs are an order of magnitude smaller.

© 2012 Elsevier B.V. All rights reserved.

1. Introduction

An influential recent literature combines fixed costs of production and exporting with firm heterogeneity to study firm-level participation in international trade. Naturally, when the unit of the analysis is the firm, much of the emphasis has been placed on the entry decision into export markets – the so-called “extensive margin.” This literature is closely related to the research agenda in economic growth that documents the existence of large impediments to entry and cross-border trade, especially in developing countries.

This paper evaluates the importance of fixed costs of production and trade and the extensive margin of imports for welfare.¹ The key ingredient of our study is the observation that firm size follows Zipf's Law, a very fat-tailed distribution that implies a few large firms account for a disproportionate share of overall economic activity.² Our main result is that once Zipf's Law in firm size is accounted for, the impact of fixed costs and the extensive margin on welfare is vanishingly small.

The analysis is based on the workhorse multi-country model of international trade in the spirit of Melitz (2003) and Eaton et al. (2011). We show how this model can be calibrated to match Zipf's Law in firm size, and illustrate analytically how the shape of the firm size distribution affects the importance of fixed costs and the extensive

[☆] We are grateful to the editor (Pol Antràs), two anonymous referees, Roc Armenter, Jonathan Eaton, Stefania Garetto, Chris House, Amit Khandelwal, Francesc Ortega, Romain Rancière, David Romer, Matthew Shapiro, Linda Tesar, and seminar participants at the University of Michigan, Trinity College Dublin, International Monetary Fund, University of Houston, Boston University, Vanderbilt University, Tufts University, Tsinghua University, Hong Kong University of Science and Technology, CUNY-Queens College, 2009 NBER ITI Spring Meetings, 2009 CEPR ESSIM, the 2009 Philadelphia Area International Trade Workshop, and the 2011 AEA Meetings for helpful suggestions, and to Aaron Flaen for expert research assistance. Levchenko thanks the National Science Foundation for financial support under grant SES-0921971. The views expressed in this paper are those of the authors and should not be attributed to the International Monetary Fund, its Executive Board, or its management.

* Corresponding author.

URL's: JdiGiovanni@imf.org, <http://julian.digiovanni.ca> (J. di Giovanni), alev@umich.edu, <http://www.alevchenko.com> (A.A. Levchenko).

¹ As will become clear below, we analyze the extensive margin defined as the (dis) appearance of firms and exporters due to the changes in the production and exporting productivity cutoffs. This extensive margin is the focus of Chaney (2008), Arkolakis (2010), and Eaton et al. (2011), among many others. The longer NBER WP version of our paper (di Giovanni and Levchenko, 2010) also discusses the impact of the extensive margin coming from changes in the mass of potential firms in the economy, as in Krugman (1980) and Melitz (2003).

² This has been documented by Axtell (2001) for the census of U.S. firms, and by di Giovanni et al. (2011) for the census of French firms. Similar findings obtain for several European countries (Fujiwara et al., 2004) and Japan (Okuyama et al., 1999). Other phenomena known to follow power laws include city size, income and wealth, and CEO compensation (Gabaix, 2009).

margin of trade. Then, we calibrate the model to the 50 largest economies in the world, paying special attention to the observed variation in the fixed costs of starting a business or trading internationally. Paradoxically, when the canonical heterogeneous firm framework ideally suited to study the extensive margin of trade is actually calibrated to the observed degree of firm heterogeneity, the extensive margin ceases to matter.

In the quantitative exercise, we first simulate the welfare impact of a world-wide reduction in the fixed costs of entry and exporting all the way to the U.S. level – a 6-fold fall in fixed costs for the average country in the sample. Even such a sizeable improvement leads to an average increase in welfare of only 3.25%. Second, we reduce the variable (“iceberg”) trade costs by 10%, and decompose the welfare impact of this change into the intensive margin – existing exporters selling more at lower prices – and the extensive margin – new exporters entering markets. The results are striking: the extensive margin of foreign varieties accounts for only 5.2% of the total welfare gains in this policy experiment. By contrast, the intensive margin is responsible for 98% of the total welfare impact of the fall in the iceberg costs.³ Finally, we show that Zipf’s Law matters a great deal quantitatively. We carry out a counterfactual calibration in which the firm size distribution is instead not fat-tailed. Under this alternative, gains from a reduction in fixed costs are about 12 times *higher*, while total gains from the reduction in iceberg trade costs are 15 times *lower*. Predictably, in this counterfactual calibration the extensive margin of trade is also more important, accounting for 14.7% of the total welfare impact of a 10% fall in variable trade costs. Thus, the distribution of firm size matters a great deal for whether fixed or variable costs have a larger welfare impact. In fact, depending on whether the firm size distribution is fat-tailed, the conclusions are reversed: in Zipf’s world fixed costs matter little, while variable costs a great deal; the opposite is true in the counterfactual alternative calibration.

What is the intuition for these results? Changes in fixed costs affect only the behavior of marginal firms; similarly, the welfare impact of the extensive margin of international trade comes by definition from new, marginal exporters. The distribution of firm size contains information about the relative importance of the marginal compared to the infra-marginal firms for welfare. It is especially important to take this into account because Zipf’s Law – a power law with an exponent close to -1 – is a very fat-tailed distribution.⁴ Economically, Zipf’s Law implies that the marginal producers and exporters are far less productive, and therefore are much smaller and sell much less. As a result, their weight in the price index (this index corresponding roughly to the inverse of welfare) is extremely low. By contrast, the infra-marginal, extremely large firms sell a lot and carry a large weight in the price index. Therefore, what happens to the large firms has a first-order impact on welfare. Our calibration exercise allows us to make this mechanism quantitatively precise. In fact, we show analytically that in the limit as the model parameters approach Zipf’s Law, the welfare impact of the extensive margin of foreign trade goes to zero.

Ever since the influential work of Djankov et al. (2002), it has been known that cross-country differences in the cost of entry by firms are pronounced. These authors assemble data on the entry regulations in 85 countries, and document that the amount of time, the number of procedures, and the costs – in either dollar terms or as a percentage of per capita income – required to start a business vary widely between countries. The World Bank’s Doing Business Initiative collected data on regulations regarding obtaining licenses, registering property,

hiring workers, getting credit, and more. Almost invariably, the data show that the variation in these regulations across countries is considerable. In addition, in a cross section of countries entry barriers are robustly negatively correlated with per-capita income and other measures of welfare. However, using cross-country econometric models to quantify the size of the impact is difficult, if not impossible. Our paper presents an alternative approach to welfare analysis. We use the World Bank’s Doing Business Indicators database to calibrate the observed variation in fixed costs across countries, and show that a model-based welfare assessment reaches very different conclusions.

Parallel to the research on entry barriers, recent advances in international trade have focused attention on the role of individual firms, both in theory and empirics. Many stylized facts have emerged: most firms do not export, most exporters sell only small amounts abroad, while the bulk of exports at any one point in time is accounted for by a relatively small number of firms (see, e.g. Bernard et al., 2007). The very same model we analyze in this paper has been used in dozens of studies to examine the firm’s decision whether to export (e.g., Chaney, 2008), or how much to export (e.g., Arkolakis, 2010). Our analysis suggests that this literature’s emphasis on the marginal firms may have been misplaced, at least when it comes to aggregate welfare.

Arkolakis et al. (2008) and Arkolakis et al. (2012) show that in several classes of models, including the standard model of monopolistic competition with endogenous variety adopted in this paper, gains from trade are summarized by the overall trade volume relative to domestic absorption. These authors argue that the overall trade volume is a “sufficient statistic,” and thus information on the extensive margin is not necessary to estimate the total gains from trade. Atkeson and Burstein (2010) and Feenstra (2010) show that in two-country heterogeneous firm models with free entry and international trade, the welfare impact of newly imported varieties and existing firms’ productivity upgrading decisions is largely offset by the impact of changes in net entry, resulting in virtually no net welfare gains from variety.

Relative to these two results, our paper’s substantive point is complementary and distinct. In the sufficient statistic literature, the extensive margin “doesn’t matter” only in the sense that one need not observe it to estimate the gains from trade. The sufficient statistic analysis is silent on whether observed changes in the overall trade volumes, and therefore welfare, are due to the extensive or intensive margins. Thus, it cannot be used to determine which policy instruments – for instance, fixed or variable costs – have the greatest welfare impact. In our analysis, the extensive margin doesn’t matter for a very different, economic reason: the marginal firms are small. It is thus informative about the role of fixed versus variable costs in welfare. Our results complement Atkeson and Burstein’s (2010) and Feenstra’s (2010) by demonstrating that under Zipf’s Law, the welfare impact of not only the “net extensive margin” – foreign plus domestic – but also of the “gross extensive margin” – foreign and domestic individually – vanishes. In a sense, this is a stronger result as it does not depend on the two gross margins canceling out perfectly. Instead we show that they are both vanishingly small in absolute value. Finally, an additional contribution of this paper is quantitative: we present a systematic assessment of the role of both fixed entry costs and variable trade barriers for welfare in a calibrated multi-country model.

Neary (2010) and Bekkers and Francois (2008) depart from the monopolistic competition paradigm, and develop heterogeneous firm models that feature strategic interactions between the large firms. Since we show that under the empirically observed distribution of firm size the small firms are unimportant, our results are complementary to the research agenda that seeks a richer model of the interaction between the largest firms.

Before moving on to the description of the model, a caveat is in order for interpreting the results. Our quantitative exercise does not strictly speaking tell us that the extensive margin does not matter

³ The disappearing domestic varieties (the domestic extensive margin) have a correspondingly negative welfare impact.

⁴ A random variable generating a power law with an exponent between -1 and -2 has infinite variance. When the power law exponent is less than -1 in absolute value, the mean becomes infinite as well.

for welfare. As such, it is not in direct contradiction with the empirical studies that find a welfare impact of increased varieties (Broda and Weinstein, 2006; Goldberg et al., 2009, 2010). First, the extensive margin in the model need not coincide with the extensive margin measured in the data. For instance, the extensive margin as measured in the data corresponds to new varieties at any point of the distribution, rather than at the cutoff as in our baseline model. Second, our counterfactual scenario of a 10% reduction in iceberg trade costs need not be a good approximation of the shocks that affected the number of varieties in the existing empirical studies. What our results demonstrate is that if the extensive margin is to matter for welfare, it would be through channels not captured by the standard model in this paper. This is important because the literature so far has predominantly used this type of model for the study of the extensive margin.

The rest of the paper is organized as follows. Section 2 presents the theoretical framework. We show how the parameters of the model govern the distribution of firm size, and how they can be mapped into the empirical firm size distribution. We then derive a number of analytical results that foreshadow the conclusions from the quantitative exercise. Section 3 solves the model economy numerically and presents the quantitative results. Section 4 concludes.

2. Theoretical framework

The world is composed of N countries, indexed by $i, j = 1, \dots, N$. In country i , buyers (who could be final consumers or firms buying intermediate inputs) solve

$$\max \left[\int_{J_i} Q_i(k) \frac{\varepsilon-1}{\varepsilon} dk \right]^{\frac{\varepsilon}{\varepsilon-1}} \quad \text{s.t.} \quad \int_{J_i} p_i(k) Q_i(k) dk = X_i,$$

where $Q_i(k)$ is the quantity sold of good k in country i , $p_i(k)$ is the price of this good, X_i is total expenditure in the economy, and J_i is the mass of varieties consumed in country i , coming from all countries. It is well known that demand for variety k in country i is equal to

$$Q_i(k) = \frac{X_i}{P_i^{1-\varepsilon}} p_i(k)^{-\varepsilon}, \quad (1)$$

where P_i is the ideal price index in this economy,

$$P_i = \left[\int_{J_i} p_i(k)^{1-\varepsilon} dk \right]^{\frac{1}{1-\varepsilon}}. \quad (2)$$

Each country has a fixed number of potential (but not actual) entrepreneurs n_i , as in Chaney (2008), Arkolakis (2010), and Eaton et al. (2011).⁵ Each potential entrepreneur can produce a unique variety k , and faces downward-sloping demand given by Eq. (1). There are both fixed and variable costs of production and trade. Each entrepreneur's type is given by the unit input requirement $a(k)$. On the basis $a(k)$, each entrepreneur in country j decides which, if any, markets to serve. A firm based in country j that decides to sell to market i must pay a fixed cost f_{ij} , and an iceberg per-unit cost of $\tau_{ij} > 1$, with the iceberg cost of domestic sales normalized to one: $\tau_{jj} = 1$.

There is one factor of production, labor, with country endowments given by L_j , $j = 1, \dots, N$. Production uses both labor and intermediate inputs. In particular, the entrepreneur with unit input requirement $a(k)$ must use this many input bundles to produce one unit of output.

An input bundle has a cost $c_j = w_j p_j^{1-\beta}$, where w_j is the wage of workers in country j , and P_j is, as above, the ideal price index of all varieties available in j . Firm k from country j thus has a marginal cost $\tau_{ij} c_j a(k)$ of serving this market i . As is well known, the profit-maximizing price is a constant markup over marginal cost, $p_i(k) = \frac{\varepsilon}{\varepsilon-1} \tau_{ij} c_j a(k)$, the revenue is equal to $\frac{X_i}{P_i^{1-\varepsilon}} \left(\frac{\varepsilon}{\varepsilon-1} \tau_{ij} c_j a(k) \right)^{1-\varepsilon}$, and the variable profits are a constant multiple $1/\varepsilon$ of revenue.

The production structure of the economy is pinned down by the number of firms from each country that enter each market. In particular, there is a cutoff unit input requirement a_{ij} , above which firms in country j do not serve market i . The cutoff a_{ij} characterizes the entrepreneur in j who earns zero profits from shipping to country i :

$$a_{ij} = \frac{\varepsilon-1}{\varepsilon} \frac{P_i}{\tau_{ij} c_j} \left(\frac{X_i}{\varepsilon c_j f_{ij}} \right)^{\frac{1}{1-\varepsilon}}. \quad (3)$$

Closing the model involves finding expressions for a_{ij} , P_i , and w_i for all $i, j = 1, \dots, N$.

We make the standard assumption that productivity, $1/a$, is Pareto(b, θ), where b is the minimum value productivity can take, and θ regulates dispersion⁶:

$$\Pr(1/a < x) = 1 - \left(\frac{b}{x} \right)^{\theta}.$$

It is then straightforward to show that the unit input requirement, a , has a distribution function $G(a) = (ba)^{\theta}$. The price level then becomes, after plugging in the expressions for a_{ij} in Eq. (3):

$$P_i = \left\{ \sum_{j=1}^N \int_{J_{ij}} \left[\frac{\varepsilon}{\varepsilon-1} \tau_{ij} c_j a(k) \right]^{1-\varepsilon} dk \right\}^{\frac{1}{1-\varepsilon}} = \left(\sum_{j=1}^N n_j \int_0^{a_{ij}} \left[\frac{\varepsilon}{\varepsilon-1} \tau_{ij} c_j a \right]^{1-\varepsilon} dG(a) \right)^{\frac{1}{1-\varepsilon}} \quad (4)$$

$$= \frac{1}{b} \left[\frac{\theta}{\theta - (\varepsilon-1)} \right]^{-\frac{1}{\theta}} \frac{\varepsilon}{\varepsilon-1} \left(\frac{X_i}{\varepsilon} \right)^{-\frac{\theta - (\varepsilon-1)}{\theta(\varepsilon-1)}} \left(\sum_{j=1}^N n_j (\tau_{ij} c_j)^{-\theta} (c_j f_{ij})^{-\frac{\theta - (\varepsilon-1)}{\varepsilon-1}} \right)^{-\frac{1}{\theta}}. \quad (5)$$

We impose balanced trade for each country, and use the convenient property (originally noted by Eaton and Kortum, 2005) that total profits in the economy are a constant multiple of X_i : $\Pi_i = \frac{\varepsilon-1}{\theta\varepsilon} X_i$.⁷ Since total sales in the economy are equal to X_i , and the total profits are $\frac{\varepsilon-1}{\theta\varepsilon} X_i$, the total spending on inputs is $(1 - \frac{\varepsilon-1}{\theta\varepsilon}) X_i$. Labor receives a constant fraction β of the spending on inputs. Thus, total spending in country i is a constant multiple of its total labor income:

$$X_i = \frac{1}{\beta(1 - \frac{\varepsilon-1}{\theta\varepsilon})} w_i L_i. \quad (6)$$

The value of exports from country i to country j can be written as:

$$X_{ji} = \frac{X_j}{P_j^{1-\varepsilon}} \left(\frac{\varepsilon}{\varepsilon-1} \tau_{ji} c_i \right)^{1-\varepsilon} n_i \frac{b^{\theta} \theta}{\theta - (\varepsilon-1)} a_{ji}^{\theta - (\varepsilon-1)}.$$

⁶ The Pareto assumption is by far the most common distributional assumption made in the heterogeneous firm models. As we show below, it leads to a power law relationship in firm size. An alternative would be to assume a lognormal distribution with a high enough variance. However, Luttmer (2007) argues that a power law relationship fits the distribution of firm size significantly better than the lognormal distribution.

⁷ The NBER WP version of this paper (di Giovanni and Levchenko, 2010) presents the proof.

⁵ This is an appropriate description of the economy at a given point in time, and thus the comparative statics based on this model should be interpreted as short- to medium-run. The NBER WP version of this paper (di Giovanni and Levchenko, 2010) presents a model in which the mass of potential entrepreneurs can adjust, as in Krugman (1980) and Melitz (2003). The main quantitative results are similar.

Using the expression for a_{ji} in Eq. (3), and P_j in Eq. (5), total exports from i to j become:

$$X_{ji} = \frac{n_i (\tau_{ji} c_i)^{-\theta} (c_i f_{ji})^{-\frac{\theta-(\varepsilon-1)}{\varepsilon-1}}}{\sum_{l=1}^N n_l (\tau_{jl} c_l)^{-\theta} (c_l f_{jl})^{-\frac{\theta-(\varepsilon-1)}{\varepsilon-1}}} X_j. \quad (7)$$

Using the trade balance conditions, $X_i = \sum_{j=1}^N X_{ji}$ for each $i = 1, \dots, N$, the expression for X_i in Eq. (6), and the definition of c_i leads to the following system of equations in w_i :

$$w_i L_i = \sum_{j=1}^N \frac{n_i (\tau_{ji} w_i^\beta P_i^{1-\beta})^{-\theta} (w_i^\beta P_i^{1-\beta} f_{ji})^{-\frac{\theta-(\varepsilon-1)}{\varepsilon-1}}}{\sum_{l=1}^N n_l (\tau_{jl} w_l^\beta P_l^{1-\beta})^{-\theta} (w_l^\beta P_l^{1-\beta} f_{jl})^{-\frac{\theta-(\varepsilon-1)}{\varepsilon-1}}} w_j L_j, \quad (8)$$

$i = 1, \dots, N$. There are $N-1$ independent equations in this system, which can be solved for wages in $N-1$ countries given a numéraire wage in the remaining country. The wages and the price levels in all countries are determined jointly by Eq. (8) for wages and Eq. (5) for prices. We will solve these numerically in order to carry out the main quantitative exercise in this paper.

2.1. The distribution of firm size: model and data

Denote the sales of an individual firm k by $x(a(k))$. Firm sales x follow a power law if

$$\Pr(x > s) = cs^{-\zeta}. \quad (9)$$

When the exponent ζ is close to 1 in absolute value, the distribution is known as Zipf's Law. It has been argued that in the data, the firm sales distribution is quite well approximated by Zipf's Law. Thus, the empirical firm size distribution places a key restriction on important parameter values in models with heterogeneous firms.

Calibrating a trade model in which not all firms export to the firm size data requires some care. The conventional estimation approach is to fit a power law relationship in Eq. (9) to total firm sales (see, e.g. Axtell, 2001, for a prominent example). As argued at length by di Giovanni et al. (2011), in a model with selection into exporting total firm sales will not follow a power law, and thus there is no simple combination of parameters that can be set equal to the conventional empirical power law estimate. However, di Giovanni et al. (2011) also provide a solution to this problem, by suggesting two alternative approaches to calibrate heterogeneous firm models to the firm size data. The key is to notice that while in the model *total* sales do not follow a power law, *domestic* sales do. To see this, note that in the model, the domestic sales of a firm as a function of its unit input requirement are: $x(a) = Ca^{1-\varepsilon}$, where the constant C reflects market size, and we drop the country subscripts. Under the assumption that $1/a \sim \text{Pareto}(b, \theta)$, the power law follows:

$$\Pr(x > s) = \Pr(Ca^{1-\varepsilon} > s) = \Pr\left(\frac{1}{a} > \left(\frac{s}{C}\right)^{\frac{1}{1-\varepsilon}}\right) = \left(\frac{b^{-\varepsilon-1}C}{s}\right)^{\frac{\theta}{\varepsilon-1}},$$

satisfying Eq. (9) for $c = \left(b^{-\varepsilon-1}C\right)^{\frac{\theta}{\varepsilon-1}}$ and $\zeta = \frac{\theta}{\varepsilon-1}$.

The key point for connecting the model to the data is that in the model, *domestic sales* follow a power law with the slope of $\frac{\theta}{\varepsilon-1}$. Since this exponent can also be estimated in the data, what we observe in the data is informative about this combination of parameters. What do the data tell us about the domestic sales ζ ? Di Giovanni et al. (2011) report a range of estimates of the power laws in domestic sales, all of which are quite close to 1. The mid-range estimate in di

Giovanni et al. (2011) is 1.06, which is the value used throughout this paper.⁸

The question remains whether Zipf's Law obtains in the domestic sales distributions of many countries. Currently, no comprehensive set of results exists. Axtell (2001) estimates a power law for firm size in the U.S., and reports a range of estimates of ζ between 0.996 and 1.059, very precisely estimated with standard errors between 0.054 and 0.064. Evidence for a limited set of European countries is presented by Fujiwara et al. (2004) and for Japan by Okuyama et al. (1999). In Appendix B, we use ORBIS – the largest publicly available firm-level dataset covering a large number of countries – to show that firm size distributions are well approximated by a power law, with exponents quite close to -1 in most countries.⁹ In addition to the paucity of estimates based on total sales, there are no existing results that take explicit account of selection into exporting, and report estimates based on domestic rather than total sales. However, by comparing the power law estimates based on total and domestic sales, di Giovanni et al. (2011) show that the bias introduced by selection into exporting is typically not large for France, and thus empirical power law estimates based on total sales probably give a reasonably accurate ballpark estimate of the degree of dispersion in domestic sales as well.

To summarize, empirical estimates of the distribution of firm size put discipline on the parameters of the Melitz–Pareto model. In particular, existing estimates suggest that $\frac{\theta}{\varepsilon-1}$ is very close to 1. As we show in a series of exercises below, this has striking implications regarding gains from reductions in entry barriers and trade costs and the relative importance of intensive and extensive margins.

2.2. Entry costs, trade openness, and the magnitude of gains from trade

We now present a number of analytical results about the relative importance of fixed costs, trade openness, and the extensive margin for welfare. Real income per capita in country i is proportional to w_i/P_i , which is also a measure of welfare.¹⁰ It is possible to use trade shares to simplify the expression for the price level. Define $\pi_{ij} \equiv X_{ij}/X_i$ to be the share of total spending in country i on goods from country

⁸ Crozet and Koenig (2010, henceforth CK) use French firm-level export data and within-France geographical location to provide structural estimates of θ and ε at sector level, and find values that imply $\theta/(\varepsilon-1)$ much higher than 1. However, (i) their structural estimates are most likely inconsistent with other relationships implied by their own model, namely sales distributions; and (ii) there are important reasons to suspect the validity of these estimates because of inconsistencies in their estimation procedures. The structural model adopted by CK implies export sales to an individual destination follow a power law with exponent $\theta/(\varepsilon-1)$. The sector-level estimates of $\theta/(\varepsilon-1)$ based on the sales distributions in di Giovanni et al. (2011) reveal values close to 1. In our view the estimates reported in di Giovanni et al. (2011) paper are more reliable, because while CK's procedure is based on estimating three structural equations, and thus relies fully on three model-implied relationships, estimates in di Giovanni et al. (2011) are much more parsimonious, based on just one equation. In addition, the third estimating equation in CK is not internally consistent. It backs out the $\theta - (\varepsilon - 1)$ combination of parameters based on the cumulative sales of all firms with marginal cost less than a . To proxy for a , they use Olley–Pakes TFP estimates, which are revenue-based, whereas a is physical productivity. Revenue-based TFP, unlike true physical marginal cost a , is not distributed Pareto(θ), since more productive firms will charge lower prices. Thus, CK's third estimating equation is not precisely specified, and the coefficient in that equation should not be thought of as $\theta - (\varepsilon - 1)$. Indeed, it is clear that this equation is responsible for the divergence between di Giovanni et al.'s (2011) estimates and CK's results: Zipf's Law implies that $\theta - (\varepsilon - 1)$ is close to zero, whereas CK's coefficients in the third equation are between about 1.5 and 3.

⁹ Other related results also shed light on how fat-tailed size distributions are. For instance, it turns out that measures of Balassa revealed comparative advantage (Hinloopen and van Marrewijk, 2006) and highly disaggregated trade flows (Easterly et al., 2009) also follow power laws with an exponent close to -1 .

¹⁰ Welfare is proportional to the real wage even though in this economy there are profits. At noted above, profits are a constant multiple of the total expenditure, while due to the Cobb–Douglas functional form of the input bundle, the wage bill $w_i L_i$ is a constant multiple of total expenditure as well. Hence, the total profits in the economy are a constant multiple of the wage bill, making the total welfare proportional to the real wage. See Eq. (6).

j. Using Eq. (7), setting $i=j$ and rearranging yields the following relationship:

$$\sum_{l=1}^N n_l (\tau_{il} c_l)^{-\theta} (c_l f_{il})^{-1} = \frac{1}{\pi_{ii}} n_i (c_i)^{-\theta} (c_i f_{ii})^{-\frac{\theta-(\varepsilon-1)}{\varepsilon-1}}.$$

Plugging this expression into the price level (Eq. (5)) and rearranging, welfare in this economy can be written as:

$$\frac{w_i}{P_i} = \left\{ \frac{1}{b} \left[\frac{\theta}{\theta-(\varepsilon-1)} \right]^{-\frac{1}{\theta}} \frac{\varepsilon}{\varepsilon-1} n_i^{-\frac{1}{\theta}} \left(\frac{L_i}{f_{ii} \varepsilon \beta (1-\frac{\varepsilon-1}{\theta})} \right)^{-\frac{\theta-(\varepsilon-1)}{\theta(\varepsilon-1)}} \pi_{ii}^{\frac{1}{\theta}} \right\}^{-\frac{1}{B}} \quad (10)$$

where $B \equiv \beta - (1-\beta) \frac{\theta-(\varepsilon-1)}{\theta(\varepsilon-1)}$. This allows us to represent real income per capita in each country relative to the U.S. as a product of several components:

$$\frac{w_i/P_i}{w_{US}/P_{US}} = \left(\frac{n_i}{n_{US}} \right)^{\frac{1}{\theta}} \left(\frac{L_i}{L_{US}} \right)^{\frac{\theta-(\varepsilon-1)}{\theta(\varepsilon-1)}} \left(\frac{f_{ii}}{f_{US,US}} \right)^{\frac{\theta-(\varepsilon-1)}{\theta(\varepsilon-1)}} \left(\frac{\pi_{ii}}{\pi_{US,US}} \right)^{\frac{1}{\theta}}.$$

A special case of this expression is obtained if we adopt the assumption in Alvarez and Lucas (2007) and Chaney (2008) that the number of productivity draws in each country is proportional to its size: $n_i = \gamma L_i$, where γ is a constant. In that case, income differences can be decomposed as:

$$\frac{w_i/P_i}{w_{US}/P_{US}} = \left(\frac{L_i}{L_{US}} \right)^{\frac{1}{\theta}} \left(\frac{f_{ii}}{f_{US,US}} \right)^{-\frac{\theta-(\varepsilon-1)}{\theta(\varepsilon-1)}} \left(\frac{\pi_{ii}}{\pi_{US,US}} \right)^{-\frac{1}{\theta}}.$$

This expression is similar in spirit to Waugh (2010), with some key differences. The similarity is in the contribution of trade to income differences, which is summarized simply by the relative openness ($\pi_{ii}/\pi_{US,US}$). The difference is that in our model entry costs also matter (the $f_{ii}/f_{US,US}$ term), and there is a market size effect such that larger countries have lower price levels and higher real per-capita incomes, all else equal.

We can get a sense of the magnitudes involved by examining both the variation in the relative fixed costs and openness, as well as the exponents. We choose the parameter values as follows: $\beta=0.5$ from Jones (2011), $\varepsilon=6$ (Anderson and van Wincoop, 2004), and $\theta=5.3$, designed to match the power law exponent in domestic sales, $\frac{\theta}{\varepsilon-1}=1.06$ (di Giovanni et al., 2011). Then, the exponents in the expression above become:

$$\frac{w_i/P_i}{w_{US}/P_{US}} = \left(\frac{L_i}{L_{US}} \right)^{0.40} \left(\frac{f_{ii}}{f_{US,US}} \right)^{-0.02} \left(\frac{\pi_{ii}}{\pi_{US,US}} \right)^{-0.38}.$$

It is immediate that the relative fixed costs will matter far less than the other two terms. In a Zipf economy, what is really important for welfare is the presence of the large, very productive firms, which are inframarginal and not affected much by the level of fixed costs.

To make this more precise, we use the World Bank's Doing Business Indicators to measure variation in $f_{ii}/f_{US,US}$ present in the data, and compute how much per-capita income variation those can generate. It turns out that the country at the 95th percentile of the fixed cost distribution has an f_{ii} that is between 16 and 658 times the U.S. value, depending on the precise indicator we use. Plugging those ratios into the equation above, we get that the country at the 95th percentile of fixed entry costs has an income level between 0.86 and 0.94 that of the U.S., all else equal. In a Zipf economy, differences in fixed costs of entry cannot generate large per-capita income – and welfare – differences.

What about trade? In the sample of the 49 largest economies by total GDP, the ratio $\pi_{ii}/\pi_{US,US}$ for the economy in the 95th percentile of openness is 0.577. Taking that to the correct exponent implies

that this country has an income level 1.23 times that of the U.S. While the absolute variation in $\pi_{ii}/\pi_{US,US}$ in the data is far lower than the variation in fixed costs, the impact of trade openness on welfare is larger.

The distribution of firm size matters for these magnitudes. To see what happens when we depart from Zipf's Law, we set $\frac{\theta}{\varepsilon-1}$ equal to 2 (implying a value of $\theta=10$ given our chosen elasticity of substitution). When the exponent on the power law in firm size is greater than or equal to 2, the distribution of firm size has finite variance. Thus, in this alternative calibration we set the exponent on the power law in firm size to be the smallest such that the distribution still has a finite variance.

In a non-Zipf economy, the exponents change dramatically: on the $f_{ii}/f_{US,US}$ term, the exponent goes up from 0.02 to 0.22 in absolute value, a tenfold increase. By contrast, the exponent on the $\pi_{ii}/\pi_{US,US}$ term drops by almost half, from 0.38 to 0.22. This implies that the importance of fixed costs rises: now, a country in the 95th percentile of the f_{ii} distribution has an income level between 0.23 and 0.54 that of the U.S. By contrast, the contribution of trade drops by half: the country in the 95th percentile of trade openness has income per capita only about 1.12 times the U.S. level.

As a related point, the shape of the firm size distribution matters a great deal for the magnitude of gains from trade. In this model, gains from trade are equal to¹¹:

$$\pi_{ii}^{-\frac{1}{\theta}}. \quad (11)$$

At a given π_{ii} , gains from trade are decreasing in $\frac{\theta}{\varepsilon-1}$ as long as $\beta\varepsilon > 1$.¹² In other words, the closer is the economy to Zipf's Law, the larger are the gains from trade. This is intuitive: in the world dominated by ultra-productive firms, the big gains from trade come from having access to those extremely productive foreign varieties. Using the values of β , ε , and θ described above, in the sample of 50 largest economies in the world, average gains from trade are 13%, with a standard deviation of 11% across countries. Assuming instead that $\frac{\theta}{\varepsilon-1}=2$ (the firm size distribution is not fat-tailed) reduces the estimated mean gains almost in half (to 7%), and the variation across countries in half as well (standard deviation of 6%).

We note the connections between our results and others in the literature. Bernard et al. (2003) also use firm/plant data to calibrate a quantitative trade model with firm-level heterogeneity. Although they use a different framework, by calibrating to firm-level data they also obtain a lower value of θ than what is implied by alternative estimation procedures, and in turn, larger overall welfare gains from trade. Something similar is taking place in our analysis: when calibrating the model using firm-level data, we find a value of θ close to 5, which is at the lower end of the estimates from the gravity literature. On a related note, the choice of $\varepsilon=6$ and $\frac{\theta}{\varepsilon-1}$ ranging from 1 to 2 corresponds to θ between 5 and 10, which is the same range of θ s considered in Arkolakis et al. (2012).

2.3. Extensive vs. intensive margins

The Zipf economy is one dominated by few large producers, that are not likely to be “marginal” exporters. Intuitively, this suggests that the distribution of firm size will also affect the relative importance of intensive versus extensive margins for welfare. In this subsection we examine analytically the importance of the two margins. The conclusion is striking: as the firm-size distribution converges to Zipf's Law, the welfare impact of the extensive margin of exports (or indeed domestic production) goes to zero.

¹¹ The gains from trade can be calculated by observing that autarky welfare is given by Eq. (10) when $\pi_{ii}=1$.

¹² This latter condition is likely to be satisfied in the data. Typical estimates of ε range from 3 to 10, while β is on the order of 0.5 (Jones, 2011).

The price level, Eq. (4), can be rewritten as a function of the extensive margin as follows:

$$P_i = \left(\frac{\varepsilon}{\varepsilon-1} b^{\varepsilon-1} \frac{\theta}{\theta-(\varepsilon-1)} \sum_{j=1}^N n_j (\tau_{ij} c_j)^{1-\varepsilon} G(a_{ij})^{\frac{\theta-(\varepsilon-1)}{\theta}} \right)^{\frac{1}{1-\varepsilon}}. \quad (12)$$

Here, the price level is expressed in terms of the share of firms from country j supplying country i , $G(a_{ij})$, which is precisely the extensive margin: a change in $G(a_{ij})$ is exactly the increase in the number (mass) of firms supplying market i . To derive the analytical result in the simplest way, let us assume that the countries are symmetric: $L_i = L$, $n_i = n$, $f_{ii} = f \forall i$, and $\tau_{ij} = \tau$, $f_{ij} = f_X \forall i, j$, $j \neq i$. In that case, wages are the same in all countries, and we normalize them to 1. The price levels are the same in all countries as well, and thus dropping the country subscripts we obtain:

$$P = \left\{ \frac{\varepsilon}{\varepsilon-1} b^{\varepsilon-1} \frac{\theta}{\theta-(\varepsilon-1)} n \left(G(a_D)^{\frac{\theta-(\varepsilon-1)}{\theta}} + (N-1) \tau^{1-\varepsilon} G(a_X)^{\frac{\theta-(\varepsilon-1)}{\theta}} \right) \right\}^{\frac{1}{\varepsilon(1-\varepsilon)}}, \quad (13)$$

where a_D is the cutoff for domestic production, and a_X is the cutoff for exporting. These are of course the same across all countries as well.

Note that since wages are normalized to 1, the total welfare in this economy is simply $W = 1/P$. We are now ready to evaluate the relative importance of the extensive and intensive margins. Imagine that there is a reduction in trade costs τ . This reduction will affect both the prices that existing exporters charge in the domestic market, given by $p(k) = \frac{\varepsilon}{\varepsilon-1} \tau c a(k)$, and the mass of firms serving the market, $G(a_X)$. The elasticity of welfare with respect to the extensive and the intensive margins is described in the following proposition.

Proposition 1. *With symmetric countries,*

$$\lim_{\tau \rightarrow 1} \frac{d \log W}{d \log G(a_X)} = 0$$

and

$$\lim_{\tau \rightarrow 1} \frac{d \log W}{d \log p} > 0.$$

Proof. From the expression for the price level (Eq. (13)), it is immediate that the elasticity of welfare with respect to the extensive margin is equal to:

$$\frac{d \log W}{d \log G(a_X)} = \frac{1}{\beta} \frac{\theta-(\varepsilon-1)}{\theta} \frac{(N-1) \tau^{1-\varepsilon} G(a_X)^{\frac{\theta-(\varepsilon-1)}{\theta}}}{G(a_D)^{\frac{\theta-(\varepsilon-1)}{\theta}} + (N-1) \tau^{1-\varepsilon} G(a_X)^{\frac{\theta-(\varepsilon-1)}{\theta}}}.$$

As the economy approaches Zipf's Law – $\theta \rightarrow (\varepsilon-1)$ – the welfare impact of the extensive margin goes to zero: $\frac{d \log W}{d \log G(a_X)} \rightarrow 0$.

The same is not true for the intensive margin. The price p that each exporter charges in the domestic market is proportional to τ . Therefore, the elasticity of welfare with respect to the intensive margin equals:

$$\frac{d \log W}{d \log p} = \frac{1}{\beta} \frac{(N-1) \tau^{1-\varepsilon} G(a_X)^{\frac{\theta-(\varepsilon-1)}{\theta}}}{G(a_D)^{\frac{\theta-(\varepsilon-1)}{\theta}} + (N-1) \tau^{1-\varepsilon} G(a_X)^{\frac{\theta-(\varepsilon-1)}{\theta}}}.$$

The welfare impact of the intensive margin clearly does not converge to zero as $\theta \rightarrow (\varepsilon-1)$. $\square$

What is the intuition for these results? In a Zipf economy the most productive firms are vastly better than the marginal firms. As a result, most of the welfare impact of trade is driven by what happens to

these best firms, rather than by whether trade liberalization leads to new entry. That is, a reduction in trade costs impacts welfare mainly because the “major brands” – Sony, Panasonic, etc. – become cheaper, rather than because many additional inferior brands of television sets become available.

The statement of the proposition may be somewhat unconventional in the sense that the comparative statics are with respect to the extensive margin of imports $G(a_X)$, which is itself an endogenous object and a function of model parameters. This is done for expositional simplicity, in order to convey the point that the welfare impact *through the extensive margin* of any parameter change that we could consider is negligible in the limit. To be completely rigorous, we could have stated the results with respect to the impact of each parameter in the model ξ (which could be fixed or variable trade costs, L , n , etc.) on welfare through the extensive margin, $\frac{d \log W}{d \log G(a_X)} \frac{d \log G(a_X)}{d \log \xi}$. It is easily verified that for any exogenous parameter, $\frac{d \log G(a_X)}{d \log \xi}$ is finite.¹³ Thus, the proposition that the derivative $\frac{d \log W}{d \log G(a_X)} \frac{d \log G(a_X)}{d \log \xi}$ goes to zero in the limit as $\theta \rightarrow (\varepsilon-1)$ holds.

This discussion shows that the conclusions about the impact of entry barriers, international trade, and the extensive margin are very sensitive to the assumption about the shape of the firm size distribution. All else equal, the fixed costs of production and exporting and the extensive margin matter less as the economy approaches Zipf's Law.

Before proceeding to the quantitative assessment of the importance of entry costs and the intensive and extensive margins of trade in a multi-country calibrated model, it is worth making two additional remarks regarding our modeling approach to fixed costs. Arkolakis (2010) develops a framework in which the fixed costs of entry are replaced by smoother market penetration costs, and firms choose not just whether to enter markets, but also what share of consumers to serve in each market. Appendix A.1 presents the model with market penetration costs, and shows that proportional changes in welfare obtained in that model are identical to those in a simple fixed costs model of the main text. This result holds for all parameter values that govern the distribution of firm size and the curvature of market penetration costs. In addition, we show that as the distribution of firm size converges to Zipf's Law, the level of welfare in that model also becomes identical to the baseline model. This is because under Zipf's Law, what matters most for welfare are the very large firms, which are least affected by the introduction of the market penetration margin. The large firms choose to penetrate markets fully, making their sales nearly the same as what they would be in a simple fixed cost model.

Second, as pointed out by Eaton et al. (2011) among others, there is evidence that not all exporters face the same fixed costs of entering foreign markets. Appendix A.2 develops an extension of the model with heterogeneous export entry costs, and shows that the results are unchanged by that alternative modeling approach as well. For these reasons, we choose to adopt the standard formulation of fixed costs of entry in our analysis.

3. Quantitative assessment

In order to implement the model numerically, we find the wages and price levels for each country, w_i and P_i that satisfy the system of equations given by Eqs. (5) and (8). To solve this system, we calibrate the values of L_i , n_i , τ_{ij} , and f_{ij} for each country and country pair, as well as the parameters common to all countries.

¹³ The closed-form expression for the extensive margin in terms of exogenous parameters is: $G(a_X) = \left(\frac{\varepsilon-1}{\varepsilon} b \right)^{(1-\Omega)\theta} \left(\frac{\theta-(\varepsilon-1)}{\theta} \right)^\Omega \left(\tau f_X^{\frac{1}{\varepsilon}} \right)^{-\theta} n^{-\Omega} \left[\beta \varepsilon \left(1 - \frac{\varepsilon-1}{\varepsilon \theta} \right) L \right]^{\frac{\theta}{\varepsilon-1} - \left(\frac{\theta}{\varepsilon-1} - 1 \right) \Omega}$
 $\left[f^{\frac{\theta-(\varepsilon-1)}{\varepsilon-1}} + (N-1) \tau^{-\theta} f_X^{\frac{\theta-(\varepsilon-1)}{\varepsilon-1}} \right]^{-\Omega}$, where $\Omega = \frac{\theta(\beta \varepsilon - 1)}{\theta(\beta \varepsilon - 1) + (1-\beta)(\varepsilon-1)}$.

3.1. Calibration

The elasticity of substitution is $\varepsilon=6$. Anderson and van Wincoop (2004) report available estimates of this elasticity to be in the range of 3 to 10, and we pick a value close to the middle of the range. The key parameter is θ , as it governs the slope of the power law. As described above, in this model domestic sales follow a power law with the exponent equal to $\frac{\theta}{\varepsilon-1}$. In the data, domestic sales follow a power law with the exponent close to 1. Di Giovanni et al. (2011) report the value of 1.06, which we use to find θ given our preferred value of ε : $\theta=1.06 \times (\varepsilon-1)=5.3$. As mentioned above, we set the share of intermediates $\beta=0.5$, following Jones (2011).

For finding the values of L_i , we follow the approach of Alvarez and Lucas (2007). First, we would like to think of L not as population per se, but as “equipped labor,” to take explicit account of TFP and capital endowment differences between countries. To obtain the values of L that are internally consistent in the model, we start with an initial guess for L_i for all $i=1, \dots, N$, and use it to solve the model. Given the vector of equilibrium wages, we update our guess for L_i for each country in order to match the ratio of total GDPs between each country i and the U.S. Using the resulting values of L_i , we solve for the new set of wages, and iterate to convergence (for more on this approach, see Alvarez and Lucas, 2007). Thus, our procedure generates vectors w_i and L_i in such a way as to match exactly the relative total GDPs of the countries in the sample. In practice, the results are close to simply equating L_i to the relative GDPs of the countries. In this procedure, we must normalize the population of one of the countries. We thus set L_{US} to its actual value of 291 million as of 2003, and compute L_i of every other country relative to this U.S. value. Finally, we set n_i in proportion to L_i . That is, the country's endowment of entrepreneurs is simply proportional to its “equipped labor” endowment. An important consequence of this assumption is that countries with higher TFP and capital abundance will have a greater number of potential productivity draws, all else equal. This is an assumption adopted by Alvarez and Lucas (2007) and Chaney (2008). We set $n_{US}=10,000,000$, that is, there are ten million potential firms in the U.S. In this calibration it implies that there are about 9,500,000 operating firms there. According to the 2002 U.S. Economic Census, there were 6,773,632 establishments with a payroll in the United States. There are an additional 17,646,062 business entities that are not employers, but they account for less than 3.5% of total shipments. Thus, choosing $n_{US}=10,000,000$ gets the correct order of magnitude for the number of firms.

Next, we must calibrate the values of τ_{ij} for each pair of countries. To do that we use the set of gravity estimates from the empirical model of Helpman et al. (2008). That is, we combine geographical characteristics such as bilateral distance, common border, common language, whether the two countries are in a currency union and others, with the coefficient estimates reported by Helpman et al. (2008) to calculate values of τ_{ij} for each country pair. Note that in this formulation, $\tau_{ij}=\tau_{ji}$ for all i and j .

Finally, we must take a stand on the values of f_{ii} and f_{ij} . The Doing Business Indicators database (The World Bank, 2007a) collects information on the administrative costs of setting up a firm – the time it takes, the number of procedures, and the monetary cost – in a large sample of countries in the world. The particular variable we use is the amount of time required to set up a business. We favor this indicator compared to others that measure entry costs either in dollars or in units of per capita income, because in our model f_{ii} is a quantity of inputs rather than value. As we must normalize f_{ii} for one country, we set the absolute level of $f_{US,US}$ to ensure an interior solution for the domestic production cutoff.¹⁴ Then, if according to the Doing Business

Indicators database it takes 10 times longer to register a business in country i than in the U.S., $f_{ii}=10 \times f_{US,US}$.¹⁵

To measure the fixed costs of international trade, we use the Trading Across Borders module of the Doing Business Indicators. This module provides the costs of exporting a 20-foot dry-cargo container out of each country, as well as the costs of importing the same kind of container into each country. Parallel to our approach to setting the domestic cost f_{ii} , the indicators we choose are the amount of time required to carry out these transactions. This ensures that f_{ii} and f_{ij} are measured in the same units. We take the bilateral fixed cost f_{ij} to be the sum of the two: the cost of exporting from country j plus the cost of importing into country i .¹⁶ The foreign trade costs f_{ij} are on average about 40% of the domestic entry costs f_{ii} . This is sensible, as it presumably is more difficult to set up production than to set up a capacity to export.

We carry out the analysis on the sample of the largest 49 countries by total GDP, plus the 50th that represents the rest of the world.¹⁷ These 49 countries together cover 97% of world GDP. We exclude entrepôt economies of Hong Kong and Singapore, both of which have total trade well in excess of their GDP, due to significant re-exporting activity. Thus, our model is not intended to fit these countries. (We do place them into the rest-of-the-world category.) The country sample, sorted by total GDP, is reported in Table A1.

3.2. Model fit

As described above, our iterative procedure ensures that the ratio of total GDPs in the model for any two countries matches exactly the ratio of the total GDPs in the data. Since one of the goals of the paper is to examine the role of trade openness in welfare, it is also important that the model produces bilateral and overall trade volumes that are close to the data.

Fig. 1a reports the scatterplot of bilateral trade ratios $\pi_{ij}=X_{ij}/X_i$. On the horizontal axis is the natural log of π_{ij} that comes from the model, while on the vertical axis is the corresponding value of that bilateral trade flow in the data. Hollow dots represent exports from one country to another, π_{ij} , $i \neq j$. Solid dots, at the top of the scatterplot, represent sales of domestic firms as a share of domestic absorption, π_{ii} . For convenience, we added a 45-degree line. It is clear that the trade volumes implied by the model match the actual data well.


¹⁵ In our calibration, individuals differ across countries in their efficiency units of labor. Taking the Doing Business Indicator data literally in light of differences in efficiency units of labor would compress the variation in f_{ii} and f_{ij} across countries, and lead to an even smaller welfare impact of the reductions in f 's, making our results even more dramatic. This is because countries with low efficiency units of labor are also the countries with high f_{ii} 's. As an example, suppose that the Doing Business Indicators database says that it takes 10 days to start a business in Mexico, and 1 day in the U.S. And suppose that when we calibrate efficiency units of labor, we find that one U.S. worker (workday) is 10 times as productive as one Mexican worker (workday). In that case, it takes the same amount of efficiency units of labor – one U.S. worker/workday – to start a business in the U.S. as in Mexico. Thus, the patterns in the data are such that accounting for differences in efficiency units of labor will only reduce the variation in f_{ii} 's across countries. The variation in efficiency units of labor observed in the data does not apply as much to the values in the Doing Business Indicators. While for a variety of reasons – adoption of technologies, quantity and quality of capital, etc. – workers in different countries have very different efficiencies in production of output, it is much less clear that efficiency differences should be equally large for the activities that the Doing Business Indicators measure, such as visiting government offices, filling out forms, and getting permits.

¹⁶ An earlier version of the paper carried out the analysis setting the bilateral fixed cost to be the sum of domestic costs of starting a business in the source and destination countries: $f_{ij}=f_{ii}+f_{jj}$. This approach may be preferred if fixed costs of exporting involved more than just shipping, and required, for instance, the exporting firm to create a subsidiary for the distribution in the destination country. The results were virtually identical.

¹⁷ We set the parameters, such as τ_{ij} and f_{ij} , for the rest-of-the-world category as the average values among the remaining countries in the world.

¹⁴ That is, we set $f_{US,US}$ to a level just high enough that $a_{ij}<1/b$ for all $i,j=1, \dots, N$ in all the baseline and counterfactual exercises, with $1/b$ being the upper limit of the distribution of $a(k)$.

(a) Bilateral Trade Shares



(b) Trade Openness




Fig. 1. Bilateral trade shares and trade openness: data and model predictions. Notes: Figure (a) reports the scatterplot of domestic output (π_{ii}) and bilateral trade (π_{ij}), both as a share of domestic absorption. Solid dots represent observations of π_{ii} , while hollow dots represent bilateral trade observations (π_{ij}). Both axes are in log scale. Figure (b) reports the scatterplot of total imports as a share of domestic absorption ($1 - \pi_{ii}$). In both figures, the values implied by the model are on the horizontal axis; actual values are on the vertical axis, and the line through the data is the 45-degree line. Source: International Monetary Fund (2007).

Most observations are quite close to the 45-degree line. It is especially important that we get the overall trade openness ($1 - \pi_{ii}$) right, since that will drive the gains from trade in each country. Fig. 1b plots the actual values of ($1 - \pi_{ii}$) against those implied by the model, along with a 45-degree line. We can see that though the relationship is not perfect, it is close.

Table 1 compares the means and medians of π_{ii} and π_{ij} 's for the model and the data, and reports the correlations between the two. The correlation between domestic shares π_{ii} in the model and in the data for this sample of countries is around 0.49. The means and the medians look very similar as well, with the countries in the model slightly more open on average than the data. The correlation between export shares, π_{ij} , is actually higher at 0.72.

Overall, the model fits bilateral trade data well. This is of course not surprising, as the Helpman et al. (2008) coefficient estimates are based on the gravity relationship, which is well known to fit trade data quite well. Nonetheless, since our calibration procedure does not use explicit information on actual trade flows, we must

Table 1

Bilateral trade shares: data and model predictions for the 50-country sample. Source: International Monetary Fund (2007).

	Model	Data
Domestic sales as a share of domestic absorption (π_{ii})		
Mean	0.7070	0.7555
Median	0.7086	0.7982
Corr(model, data)	0.4900	
Export sales as a share of domestic absorption (π_{ij})		
Mean	0.0060	0.0047
Median	0.0027	0.0011
Corr(model, data)	0.7171	

Notes: This table reports the means and medians of domestic output (top panel), and bilateral trade (bottom panel), both as a share of domestic absorption, in the model and in the data.

check the fit to actual trade. We now turn to the analysis of welfare gains from reduction in entry costs and trade barriers implied by the model.

3.3. Counterfactual I: reduction in entry costs

Using the calibrated model above, the first counterfactual we perform is a reduction in the fixed costs of entry f_{ii} and f_{ij} . We simulate a complete harmonization of entry costs across the world, such that entry costs everywhere are the same as in the U.S.¹⁸ This is a substantial improvement. As first shown by Djankov et al. (2002), the differences in these fixed costs are substantial across countries. In our sample of the world's 49 largest economies, it takes on average 6 times longer to start a business compared to the U.S. For a country at the 75th percentile of the distribution, it takes almost 8 times longer, and the country with the highest entry costs in this sample – Brazil – it takes 25 times longer than in the U.S. This experiment also entails a substantial drop in the fixed costs of cross-border trade. The average exporting cost in this sample is 3 times higher than in the U.S., and the average importing cost is 4 times higher.

Table 2 reports the associated welfare gains. The top panel presents the baseline calibration, in which the firm-size distribution is set to match Zipf's Law. The welfare gains are small. Even a dramatic drop (6-fold on average) in the fixed costs of production and exporting improves welfare by only 3.26% on average. It could be that this average number is hiding a lot of heterogeneity, since different countries are experiencing a different size reduction in trade costs. In parentheses below the average value, we report the range of welfare gains in the entire sample. We can see that even in the country that gains the most from this institutional improvement, the gain is only about double the average, at 7.32%.¹⁹ Zipf's Law matters a great deal for this conclusion. The bottom panel reports the alternative counterfactual calibration, in which $\theta/(\varepsilon - 1) = 2$. The welfare gain from the same reduction in entry barriers is on average 40.87% in the non-Zipf world, 12 times higher than in the Zipf's Law

¹⁸ To be precise, we set the cost of setting up a firm, $f_{ii} = f_{US,US}$, and then we set the cost of importing to any country i to the cost of importing to the U.S., and set the cost of exporting from any country i to the cost of exporting from the U.S. Thus the fixed cost of exporting between any two countries f_{ij} in the counterfactual becomes equal to the fixed cost of exporting from the U.S. plus the fixed cost of importing into the U.S.

¹⁹ Without input linkages ($c_i = w_i$), since by construction f_{ij} affects entry but not the variable costs of existing firms, we can think of these welfare gains as coming from the extensive margin. In the presence of input–output linkages, this is not strictly speaking correct because of general equilibrium effects: entry affects the price levels in each country, which in turn enter the marginal costs of each operating firm, and thus can have an impact through the intensive margin as well. Separating out the positive intensive margin effect would leave an even smaller direct extensive margin effect of reductions in entry costs.

Table 2
Welfare gains.

Counterfactual	Total change in welfare	Intensive margin	Extensive margin	
			Foreign	Domestic
$\frac{\theta}{\varepsilon-1} = 1.06$ (Zipf's world)				
Complete harmonization of entry costs	3.26 (0.03, 7.32)			
10% reduction in τ	4.33 (0.28, 8.26)	4.23 (0.28, 8.06) 0.980	0.21 (0.02, 0.34) 0.052	−0.11 (−0.14, −0.03) −0.032
$\frac{\theta}{\varepsilon-1} = 2$				
Complete harmonization of entry costs	40.87 (0.00, 104.78)			
10% reduction in τ	0.28 (0.01, 1.72)	0.27 (0.01, 1.64) 0.968	0.038 (0.000, 0.194) 0.147	−0.025 (−0.114, −0.001) −0.115

Notes: This table reports the welfare changes, in percentage points, due to each counterfactual experiment. The numbers in parentheses indicate the range across the 50 countries in the sample. The numbers in bold give the share of each margin (intensive, foreign extensive, and domestic extensive) in the total welfare impact.

calibration. The range is also greater: the country gaining the most more than doubles its welfare.²⁰

The intuition for this result is that the distribution of firm size contains information on the relative importance of the marginal and the inframarginal varieties. Under Zipf's Law, the inframarginal varieties – the very large firms – are overwhelmingly more important than the marginal varieties. Thus, since the high entry costs do not affect the entry decision of the very large firms, they do not have much impact on welfare. As our quantitative exercise demonstrates, this is true even in a model with a substantial intermediate input multiplier. As we move away from Zipf's Law, the distribution of firm size becomes flatter. As a result, entry of the marginal firms, and consequently the fixed costs of entry, become more important for welfare.

3.4. Counterfactual II: reduction in variable trade barriers

Consider a global reduction in trade costs τ_{ij} . How will it affect welfare, and what will be the relative importance of the intensive and the extensive margins? We know that welfare in this model is proportional to real income, $W_i = w_i/P_i$. From Eq. (12), welfare can be expressed, up to a constant that is the same in all countries and trade regimes, as follows:

$$W_i = \left[\sum_{j=1}^N n_j \left(\tau_{ij} \frac{c_j}{c_i} \right)^{1-\varepsilon} G(a_{ij})^{\frac{\theta-(\varepsilon-1)}{\theta}} \right]^{\frac{1}{\beta(\varepsilon-1)}}. \quad (14)$$

A reduction in trade costs will impact the intensive margin, by making existing goods cheaper. That is captured by the $\tau_{ij} \frac{c_j}{c_i}$ term. Additionally, welfare will increase due to the extensive margin, by leading to a greater number of varieties. This is captured by the

$G(a_{ij})$ term. Using a Taylor expansion, we can write the proportional increase in welfare as a function of the two margins:

$$\frac{\Delta W_i}{W_i} \approx \frac{1}{\beta} \sum_{j=1}^N \varphi_{ij} \left[\underbrace{-\frac{\Delta \left(\tau_{ij} \frac{c_j}{c_i} \right)}{\tau_{ij} \frac{c_j}{c_i}}}_{\text{Intensive Margin}} + \underbrace{\frac{\theta-(\varepsilon-1)}{\theta(\varepsilon-1)} \frac{\Delta G(a_{ij})}{G(a_{ij})}}_{\text{Extensive Margin}} \right], \quad (15)$$

where φ_{ij} is the weight of country j in country i 's price level:

$$\varphi_{ij} \equiv \frac{n_j \left(\tau_{ij} c_j \right)^{1-\varepsilon} G(a_{ij})^{\frac{\theta-(\varepsilon-1)}{\theta}}}{\sum_{l=1}^N n_l \left(\tau_{il} c_l \right)^{1-\varepsilon} G(a_{il})^{\frac{\theta-(\varepsilon-1)}{\theta}}}.$$

It is immediate from Eq. (15) that the extensive margin does not have much of a chance to impact welfare. Any given change in the mass of new firms, $\frac{\Delta G(a_{ij})}{G(a_{ij})}$, while it may be large, is pre-multiplied by the term $\frac{\theta-(\varepsilon-1)}{\theta(\varepsilon-1)}$, which goes to zero as the economy approaches Zipf's Law. The calibrated value of this combination of parameters is about 0.01.

Table 2 reports the quantitative results for our sample of countries. A 10% reduction in trade barriers leads to an average increase in welfare of about 4.3%, with a range between 0.28% and 8.26%. Notably, this is somewhat higher than welfare gain we saw following a complete harmonization of entry barriers across countries. It turns out that the intensive margin accounts for 98% of the overall welfare gain. The table breaks down the extensive margin into the component coming from the new foreign varieties, and the component due to the disappearance of some domestic ones. The foreign extensive margin contributes 5.2% of the total welfare gain. It is partially undone by the domestic extensive margin, whose welfare impact is negative. As we can see, in Zipf's world, the extensive margin plays a minimal role relative to the intensive one.

It is important to emphasize that this result is not due to a small increase in the number of foreign varieties. In this experiment, the 10% reduction in τ_{ij} leads to an average 28% increase in the number of imported foreign varieties in this set of countries. The extensive margin, as measured by the number of varieties, is quantitatively important. However, its contribution to welfare is not.

The bottom panel reports these results with the alternative, non-fat-tailed calibration. Two features are most striking. First, the overall gains from a 10% reduction in τ_{ij} are tiny compared to the baseline calibration. The average gains are only 0.28% (less than one third of 1%), with a maximum of 1.7%. This is 15 times lower than the same reduction in trade costs in the baseline calibration. Second, the overall importance of the intensive margin is almost the same as in the baseline calibration, 96.8%. At first glance this is surprising. But it turns out that the welfare impact of the foreign extensive margin is indeed much bigger than in the baseline calibration, as expected. The foreign extensive margin contributes 14.7% of the total welfare gain, almost 3 times greater than in the baseline calibration. However, the domestic extensive margin is also more important for welfare, contributing −11.5% of the total impact. That is, the disappearance of existing domestic varieties that accompanies the drop in trade costs also has a greater (negative) welfare impact compared to the Zipf case. The two partially cancel out, leaving the relative importance of the intensive margin roughly unchanged.

The main results are presented graphically in Fig. 2. On the x-axis is the power law exponent in firm size, $\theta/(\varepsilon-1)$, which varies from 1.06 (Zipf's Law calibration) to 2. The lines display the welfare impact of the two counterfactual experiments we consider: a 10% reduction in τ_{ij} (solid line) and the complete harmonization in f_{ij} to their U.S. level. The figure illustrates the importance of the firm size distribution for our conclusions about welfare. In particular, it is clear that changes in

²⁰ An interesting question is how large is the role of international trade in generating this welfare gain. To get a sense of this, we calculated the gains from the same reduction in fixed costs of entry under the assumption that each country is in autarky. It turns out that the magnitude of the autarky gains is very similar. For instance, in the Zipf calibration the autarky gain is 3.47%, compared to 3.26% in the baseline open economy case. We conjecture that the average autarky percentage gain is slightly higher because in the absence of the possibility of importing, it is more important to have access to the most domestic varieties.




Fig. 2. The welfare impact of reductions in fixed and variable costs and the size distribution of firms. Notes: This figure reports the percentage changes in welfare due to a reduction in iceberg trade costs (solid line, left axis) and a reduction in fixed costs of entry (dashed line, right axis), as a function of the distribution of firm size.

variable costs matter *more* for welfare as the economy approaches Zipf's Law, while changes in fixed costs matter *less*.

3.5. Robustness

This section assesses the robustness of the main results in several dimensions: (i) alternative calibration of iceberg trade costs; (ii) alternative values of the elasticity of substitution; and (iii) the importance of the continuum of goods assumption.

The baseline calibration uses estimates of τ_{ij} from Helpman et al. (2008), and does not match bilateral or overall trade flows perfectly. A natural question is whether the results are affected by this approach. To check for this possibility, we implement an alternative solution procedure, that instead of matching relative country GDPs, picks values of τ_{ij} to match perfectly all bilateral trade flows. To be precise, given the values of L_i implied by our baseline procedure, we find values of w_i , P_i , and $\tau_{ij} \forall i, j$ such that the equilibrium market clearing conditions (5) and (8) hold and all the trade shares (7) in the model match perfectly the data values.²¹

The results are presented in the first row of Table 3. The change in welfare due to the fall in f_{ij} of 3.24% is virtually the same as the baseline value of 3.26%. The welfare impact of a 10% reduction in τ_{ij} is somewhat lower, at 2.57%, but the extensive margin of exports contributes 0.13% to welfare, or 5% of the total. This share of the extensive margin is virtually identical to the baseline finding.

Next, we assess the sensitivity of the results to our choice of ε , by implementing the model under two alternative values, while keeping the calibration to Zipf's law, that is, $\theta/(\varepsilon-1)=1.06$ throughout. We try a lower value of $\varepsilon=4$ (and thus $\theta=3.18$), and then a higher value of $\varepsilon=8$ (and thus $\theta=7.42$). Table 3 presents the results. As expected, the gains under a lower ε are higher, and vice versa, but the main conclusions are unchanged. With respect to the harmonization in entry costs, the gains are 2.42% under $\varepsilon=8$ and 4.86% under

Table 3
Robustness.

	Complete harmonization of entry costs	10% reduction in τ	
		Total	Foreign extensive
Perfect match to trade flows	3.24	2.57	0.13
$\varepsilon=4$	4.86	10.55	0.20
$\varepsilon=8$	2.42	0.91	0.08
Discrete number of firms	5.00	4.61	0.22

Notes: This table reports the welfare changes, in percentage points, under alternative calibration approaches.

$\varepsilon=4$, compared to the baseline of 3.26%. The range of gains from a 10% reduction in τ_{ij} is somewhat larger, from 0.9% to 10.55%, but in either case the contribution of the extensive margin of exports is small, 0.08% and 0.2% respectively (compared to the baseline of 0.21%).

Finally, we check how important the continuum of firms assumption is for the results. The main mechanism in our paper is that under Zipf's Law, the largest firms are extremely large, and thus matter much more for welfare than the marginal firms. Under the continuum of firms assumption and the Pareto distribution, the upper support of firm productivity is infinite. However, in the real world of course no actual firm has infinite productivity, instead there is a discrete and finite number of firms in each country (the implications of this for international trade flows have recently been explored by Eaton et al., 2012). To assess how this affects the results, we simulate the model under a discrete number of firms. That is, we randomly draw n_i productivities for each country i , and given the model solution for the cutoffs a_{ij} , check whether each discrete firm produces domestically and serves each export market. For each set of global productivity draws we perform the main experiments of lowering f_{ij} and lowering τ_{ij} . We compute the price indices directly as the discrete equivalents of Eq. (2), and changes in the extensive margin $\frac{\Delta G(a_{ij})}{G(a_{ij})}$ in Eq. (15) directly as the proportional changes in the numbers of active firms from each market j in each i . We draw a worldwide set of random productivities 1001 times, and take the median welfare change in each country for each counterfactual experiment.²²

The results, reported in Table 3, are very similar to the baseline. As expected, the impact of reductions in f_{ij} and τ_{ij} is slightly higher than under the continuum of firms. The discrete simulation effectively cuts out the infinitely productive upper tail, thus making the marginal firms slightly more important. The difference is not large, however. The impact of the extensive margin of exports, in both the absolute terms and relative to the total increase in welfare, is virtually identical. We conclude from this exercise that the assumption of a continuum of firms does not drive the quantitative results.

4. Conclusion

The world economy and world trade flows are dominated by very large firms. This paper studies the implications of this stylized fact for two related aspects of the economy: entry costs and the extensive margin of exports. The conclusions about the welfare impact of higher entry barriers and the extensive margin of trade are very sensitive to

²¹ We cannot match both relative country sizes (as in the baseline calibration) and trade flows perfectly at the same time. We conjecture that this is because the model is too simple: it has one factor of production, no explicit non-traded sector, and no aggregate trade imbalances. The calibration that matches trade flows perfectly produces relative country sizes that are very close to the data, with the correlation in country size between the model and the data of 0.92. Among the 2500 possible unidirectional bilateral trade flows, 18 are zeros. Since the model solution cannot handle zero trade flows, in the matching procedure we replace each data zero trade share with the value of the smallest non-zero trade share found in the sample, which has order of magnitude 10^{-6} .

²² See di Giovanni and Levchenko (2011) for a more detailed description of the model and simulation with a discrete number of firms. Note that this procedure does not represent a complete solution to the model under each individual set of random productivity draws, because we do not re-calculate the cutoffs a_{ij} and relative wages w_i for each draw. Finding the complete solution to the model under a discrete number of firms would not be feasible, as it would require for each vector of random productivity draws to solve for a fixed point in all the cutoffs, prices, and aggregate expenditures. Since we draw 10 million productivities just for the U.S., and thus dozens of millions of productivity draws for the world overall, looking for the fixed point in a_{ij} 's, w_i 's and X 's for the 50 countries would be impractical. Thus, our simulation uses values of a_{ij} 's computed in the solution of the model with a continuum of firms. This procedure is clearly an approximation, but nonetheless it can give us a sense of how the continuum assumption affects the main mechanisms of the model.

the assumptions on the size distribution of firms. In a model calibrated to match the observed firm-size distribution, the welfare costs of entry barriers are low, and the extensive margin accounts for only 5% of the overall gains from a reduction in iceberg trade costs.

What should we take away from this exercise? Quantitative evidence cannot be used to argue that entry costs and the extensive margin of trade are not important for welfare. We can establish, however, that the canonical model of production and trade with endogenous variety cannot generate a significant welfare impact of entry barriers and the extensive margin, while at the same time matching both the empirically observed distribution of firm size and trade volumes. If these matter, it must be through some other channel. Uncovering the conditions under which the costs of entry into domestic and foreign markets matter more remains a fruitful avenue for future research.

Appendix A. Alternative specifications of fixed costs

A.1. Market penetration costs

A recent contribution by Arkolakis (2010) emphasizes that the model with simple fixed costs of accessing markets is too stark. Instead, Arkolakis (2010) proposes a model in which firms choose not only whether to enter a particular market, but what share of the consumers in that market to serve. Arkolakis (2010) and Eaton et al. (2011) demonstrate that modeling entry costs in this more continuous way is important to account for the empirical regularity that many firms export only small amounts abroad.

In this appendix, we extend the baseline model to feature market penetration costs instead of fixed entry costs, and demonstrate that the total welfare in such a model differs from the baseline only by a constant. As a result, in any policy experiment the market penetration costs model produces welfare changes that are identical to the baseline fixed costs model.

Our functional form assumption follows Eaton et al. (2011). Assume that rather than paying the fixed cost $f_{ij}c_j$ to gain access to all consumers in market i , a firm in country j incurs a cost

$$f_{ij}c_j \frac{1-(1-s)^{1-\frac{1}{\lambda}}}{1-\frac{1}{\lambda}}$$

to reach a share s of consumers in that market. Given the demand for its variety by the consumer reached in country i , the firm with unit input requirement $a(k)$ from country j maximizes its profits by choosing both its price and market penetration $s_i(k)$ optimally. The profits are given by:

$$\pi_i(k) = [p_i(k) - \tau_{ij}c_j a(k)] \left(\frac{p_i(k)}{P_i} \right)^{-\varepsilon} s_i(k) X_i - f_{ij}c_j \frac{1-(1-s)^{1-\frac{1}{\lambda}}}{1-\frac{1}{\lambda}},$$

where the price index, P_i , now aggregates over the prices of varieties available to a typical consumer in i , and not over all the varieties that are sold in that country. It is easily verified that the price is still a constant markup over the marginal cost. Optimal market penetration for a firm with unit input requirement $a(k)$ is given by:

$$s_i(k) = 1 - \left[\frac{X_i}{\varepsilon c_j f_{ij}} \left(\frac{\tau_{ij}c_j a(k)}{P_i} \right)^{1-\varepsilon} \right]^{-\lambda}. \quad (\text{A.1})$$

Finally, the firm will only enter market i if at zero market penetration, profits are increasing in s : $\frac{\partial \pi_i(k)}{\partial s}|_{s=0} > 0$. It turns out that the cutoff a_{ij} for positive sales from j to i has the exact same form as in the baseline model, and is given by Eq. (3). That expression can be combined with Eq. (A.1) to write the sales of a firm with unit input requirement $a(k)$ from country j to country i as:

$$\left[1 - \left(\frac{a(k)}{a_{ij}} \right)^{\lambda(\varepsilon-1)} \right] \left(\frac{\tau_{ij}c_j a(k)}{P_i} \right)^{1-\varepsilon} X_i.$$

As first observed by Arkolakis (2010), the baseline model with simple fixed costs provides the best approximation to the sales of the largest firms: as the unit input requirement $a(k)$ decreases, $s_i(k) = \left[1 - \left(\frac{a(k)}{a_{ij}} \right)^{\lambda(\varepsilon-1)} \right]$ approaches 1 and the firm penetrates the entire market. This result does not rely on the Zipf's Law assumption: the market penetration ratio $s_i(k)$ does not depend on the combination of parameters $\frac{\theta}{\varepsilon-1}$. As we argue at the end of this section, Zipf's Law does imply that the large firms are the ones most important for welfare, and thus the assumption of simple fixed costs adopted in the main text will not substantially affect our conclusions.

Under the Pareto distribution of productivity draws, the expression for the price level in country i is given by:

$$\begin{aligned} P_i^{mp} &= \left(\sum_{j=1}^N n_j \int_0^{a_{ij}} \left[\frac{\varepsilon}{\varepsilon-1} \tau_{ij}c_j a(k) \right]^{1-\varepsilon} s_j(k) dG(a(k)) \right)^{\frac{1}{1-\varepsilon}} \\ &= \frac{1}{b} \left[\frac{\theta}{\theta-(\varepsilon-1)} - \frac{\theta}{\theta-(\varepsilon-1)(1-\lambda)} \right]^{-\frac{1}{\theta}} \frac{\varepsilon}{\varepsilon-1} \left(\frac{X_i}{\varepsilon} \right)^{-\frac{\theta-(\varepsilon-1)}{\theta(\varepsilon-1)}} \\ &\quad \times \left(\sum_{j=1}^N n_j (\tau_{ij}c_j)^{-\theta} (c_j f_{ij})^{-\frac{\theta-(\varepsilon-1)}{\varepsilon-1}} \right). \end{aligned} \quad (\text{A.2})$$

Comparing Eqs. (5) and (A.2), it is clear that the price levels in the baseline model and the market penetration cost model differ only by a constant. The rest of the solution is unchanged. In particular, it is straightforward to show that total profits in each economy are still a constant multiple of X_i , and that the wages are still determined by Eq. (8). Thus, the solution to the market penetration costs model proceeds to find w_i^{mp} and P_i^{mp} for all $i = 1, \dots, N$ that solve the system of equations given by Eqs. (8) and (A.2). We now state the main result of this section.

Proposition 2. Let the vectors $[w_1, \dots, w_N]$ and $[P_1, \dots, P_N]$ jointly be a solution to the system of equations defining the equilibrium in the baseline fixed costs model, Eqs. (5) and (8). Then, the vectors

$$[w_1^{mp}, \dots, w_N^{mp}] = [w_1, \dots, w_N] \quad (\text{A.3})$$

and

$$[P_1^{mp}, \dots, P_N^{mp}] = \delta [P_1, \dots, P_N] \quad (\text{A.4})$$

are a solution to the system of Eqs. (A.2) and (8) that define the equilibrium in the market penetration costs model.

Proof. It is immediate from examining Eq. (8) that the vector $[w_1, \dots, w_N]$ that solves Eq. (8) is the same under $[P_1, \dots, P_N]$ and $[P_1^{mp}, \dots, P_N^{mp}]$ when the latter is defined by Eq. (A.4), since δ cancels out from the numerator and the denominator. We now show that as long as Eq. (A.3) is satisfied, Eq. (A.4) holds as well for some constant δ . The vector $[P_1^{mp}, \dots, P_N^{mp}]$ provides a solution to the market penetration costs model if $\forall i$, Eq. (A.2) holds. We check directly whether the vector $\delta [P_1, \dots, P_N]$ satisfies that condition:

$$\begin{aligned} P_i^{mp} = \delta P_i &= \frac{1}{b} \left[\frac{\theta}{\theta-(\varepsilon-1)} - \frac{\theta}{\theta-(\varepsilon-1)(1-\lambda)} \right]^{-\frac{1}{\theta}} \frac{\varepsilon}{\varepsilon-1} \left(\frac{w_i L_i}{\varepsilon \beta \left(1 - \frac{\varepsilon-1}{\theta \varepsilon} \right)} \right)^{-\frac{\theta-(\varepsilon-1)}{\theta(\varepsilon-1)}} \\ &\quad \times \left(\sum_{j=1}^N n_j \left(\tau_{ij} w_j^\beta (\delta P_j)^{1-\beta} \right)^{-\theta} \left(w_j^\beta (\delta P_j)^{1-\beta} f_{ij} \right)^{-\frac{\theta-(\varepsilon-1)}{\varepsilon-1}} \right)^{-\frac{1}{\theta}}. \end{aligned} \quad (\text{A.5})$$

After rearranging it becomes:

$$\delta^{\beta-(1-\beta)\frac{\theta-(\varepsilon-1)}{\theta(\varepsilon-1)}} P_i = \frac{1}{b} \left[\frac{\theta}{\theta-(\varepsilon-1)} - \frac{\theta}{\theta-(\varepsilon-1)(1-\lambda)} \right]^{-\frac{1}{\theta}} \frac{\varepsilon}{\varepsilon-1} \left(\frac{w_i L_i}{\varepsilon \beta \left(1 - \frac{\varepsilon-1}{\theta \varepsilon}\right)} \right)^{\frac{\theta-(\varepsilon-1)}{\theta(\varepsilon-1)}} \\ \times \left(\sum_{j=1}^N n_j (\tau_{ij} w_j^{\beta} P_j^{1-\beta})^{-\theta} (w_j^{\beta} P_j^{1-\beta} f_{ij})^{-\frac{\theta-(\varepsilon-1)}{\varepsilon-1}} \right)^{\frac{1}{\theta}},$$

which is the same as Eq. (5) for δ satisfying $\left[\frac{\theta}{\theta-(\varepsilon-1)} \right]^{-\frac{1}{\theta}} = \left[\frac{\theta}{\theta-(\varepsilon-1)} - \frac{\theta}{\theta-(\varepsilon-1)(1-\lambda)} \right]^{-\frac{1}{\theta}} \delta^{-(\beta-(1-\beta)\frac{\theta-(\varepsilon-1)}{\theta(\varepsilon-1)})}$. Since the vector $[P_1, \dots, P_N]$ satisfies Eq. (5), we have shown that $\delta[P_1, \dots, P_N]$ satisfies Eq. (A.5), which completes the proof. $\square$

The main consequence of Proposition 2 is that the total welfare in the market penetration costs model differs from the welfare in the basic fixed costs model only by a constant: $w_i^{mp}/P_i^{mp} = (1/\delta)w_i/P_i$. This implies that any percentage change in welfare calculated in this model will be identical to the baseline in the main text.

One additional remark is worth making on the relationship between the market penetration costs model and this paper. Straightforward rearranging yields the following expression for δ :

$$\delta = \left[\frac{\lambda(\varepsilon-1)}{\theta-(\varepsilon-1)(1-\lambda)} \right]^{-\frac{1}{\theta} \frac{1}{\beta-(1-\beta)\frac{\theta-(\varepsilon-1)}{\theta(\varepsilon-1)}}}.$$

Setting $\lambda=1$ the expression in the square brackets becomes $(\varepsilon-1)/\theta$.²³ Therefore, it is immediate that as we approach Zipf's Law, $\delta \rightarrow 1$ and the welfare level in the market penetration cost model converges exactly to the welfare level in the simple fixed costs model. This is intuitive: under Zipf's Law, what matters the most for welfare are the biggest firms, for which the market penetration margin matters the least, since they choose to serve the entire market.

A.2. Stochastic fixed costs

Another concern with respect to fixed costs is that they may not be the same for all firms. For instance, the model with non-stochastic fixed costs implies a strict hierarchy of entry into markets, which is known not to hold perfectly (Eaton et al., 2011). This section presents an extension of the model with stochastic export entry costs, and shows that it is largely isomorphic to the basic model in the main text.

Suppose that a firm in market j faces a fixed cost $f_{ij}\eta$ to enter market $i \neq j$, where f_{ij} is the component common across firms and $\eta > 0$ is a random firm-specific shock. Following standard practice (see, e.g. Eaton et al., 2011), we assume that η is distributed independently of firm unit input requirement a . The cutoff for entering market i by market j firms is now contingent on the realization of η :

$$a_{ij}(\eta) = \frac{\varepsilon-1}{\varepsilon} \frac{P_i}{\tau_{ij} c_j} \left(\frac{X_i}{\varepsilon c_j f_{ij} \eta} \right)^{\frac{1}{\varepsilon-1}}. \quad (\text{A.6})$$

The price level then becomes

$$P_i = \left(\sum_{j=1}^N n_j \int_{\eta} \int_0^{a_{ij}} \left[\frac{\varepsilon}{\varepsilon-1} \tau_{ij} c_j a \right]^{1-\varepsilon} dG(a) dF_{ij}(\eta) \right)^{\frac{1}{1-\varepsilon}} \\ = \frac{1}{b} \left[\frac{\theta}{\theta-(\varepsilon-1)} \right]^{-\frac{1}{\theta}} \frac{\varepsilon}{\varepsilon-1} \left(\frac{X_i}{\varepsilon} \right)^{-\frac{\theta-(\varepsilon-1)}{\theta(\varepsilon-1)}} \left(\sum_{j=1}^N n_j \tilde{\eta}_{ij} (\tau_{ij} c_j)^{-\theta} (c_j f_{ij})^{-\frac{\theta-(\varepsilon-1)}{\varepsilon-1}} \right)^{-\frac{1}{\theta}},$$

where $F_{ij}(\eta)$ is the CDF of η that can vary by both source and destination, and $\tilde{\eta}_{ij} \equiv \int_{\eta} \eta^{-\frac{\theta-(\varepsilon-1)}{\theta(\varepsilon-1)}} dF_{ij}(\eta)$. Since we assume that the costs of entering the domestic market are non-stochastic, trivially $\tilde{\eta}_{ii} = 1 \forall i$.

It is immediate that the expression for welfare in Eq. (10), and the analysis in Section 2.2 is unchanged. When it comes to the analysis of the impact of the extensive margin on trade in Section 2.3, straightforward steps lead to the following modified expression for the price level, and hence welfare:

$$P = \left\{ \frac{\varepsilon}{\varepsilon-1} b^{\varepsilon-1} \frac{\theta}{\theta-(\varepsilon-1)} n \left(G(a_D)^{\frac{\theta-(\varepsilon-1)}{\theta}} + (N-1) \tau^{1-\varepsilon} \int_{\eta} G(a_X(\eta))^{\frac{\theta-(\varepsilon-1)}{\theta}} dF(\eta) \right) \right\}^{\frac{1}{\theta(1-\varepsilon)}}, \quad (\text{A.8})$$

where now $a_X(\eta)$ is the cutoff unit input requirement necessary for exporting for firms that drew the entry cost disturbance equal to η . Differentiating inside the integral, the elasticity of welfare with respect to the extensive margin $G(a_X(\eta))$ for each level of η is equal to:

$$\frac{d \log W}{d \log G(a_X(\eta))} = \frac{1}{\beta} \frac{\theta-(\varepsilon-1)}{\theta} \frac{(N-1) \tau^{1-\varepsilon} G(a_X(\eta))^{\frac{\theta-(\varepsilon-1)}{\theta}}}{G(a_D)^{\frac{\theta-(\varepsilon-1)}{\theta}} + (N-1) \tau^{1-\varepsilon} \int_{\eta} G(a_X(\eta))^{\frac{\theta-(\varepsilon-1)}{\theta}} dF(\eta)}.$$

As the economy approaches Zipf's Law – $\theta \rightarrow (\varepsilon-1)$ – the welfare impact of the extensive margin at each level of η goes to zero: $\frac{d \log W}{d \log G(a_X(\eta))} \rightarrow 0$.

Introducing stochastic fixed exporting costs has virtually no impact on the results because our main mechanism operates at each level of fixed costs. That is, as long as the distribution of sales is fat-tailed conditional on any realization of entry costs $f_{ij}\eta$, it is still the case that the inframarginal exporters among that set of firms are much more productive than the marginal ones. To build intuition, suppose that for some exporter j , η can take only two values, η_H and η_L . This setup is roughly equivalent to facing two exporting countries each with non-stochastic fixed exporting costs, one with $f_{ij}\eta_H$, the other with $f_{ij}\eta_L$, and all of the results in the main text apply.²⁴ Assuming a continuum of η 's instead of two discrete values leaves the basic intuition unchanged.

Appendix B. Power laws in firm size in the ORBIS database

This appendix uses a large cross-country firm-level database to assess whether Zipf's Law approximates well the distribution of firm size in a large sample of countries. Though we use the largest available non-proprietary firm-level database in this analysis, the results should be interpreted with caution: coverage is quite uneven across countries and years, implying that power law estimates may not be reliable or comparable across countries. Nonetheless, as we describe below, Zipf's Law provides a good approximation for the firm size distribution in most countries in this sample.

²³ This is the value of λ preferred by Arkolakis (2010). Using Simulated Method of Moments, Eaton et al. (2011) indeed estimate a value of $\lambda = 0.91$ with a standard error of 0.12. This type of value for λ implies a fair amount of curvature to the market penetration costs, and thus many firms that choose to penetrate only a small share of the export market. The fixed cost model obtains instead when $\lambda = \infty$.

²⁴ The reason this setup is not exactly equivalent to having two countries is that labor markets have to clear and a single wage has to prevail in exporter j . This general equilibrium effect does not appreciably affect the heuristic argument here.

ORBIS is a multi-country database published by Bureau van Dijk that contains information on more than 50 million companies worldwide.²⁵ The data come from a variety of sources, including, but not limited to, registered filings and annual reports. Coverage varies by world region: there are data on some 17 million companies in the U.S. and Canada, 22 million companies in the 46 European countries, 6.2 million companies from Central and South America, 5.3 million from Asia, but only 260,000 from Africa and 45,000 from the Middle East. Importantly, the database includes both publicly traded and privately held firms. While in principle data are available going back to mid-1990s for some countries, coverage improves dramatically for more recent years. Thus, for each country we use the year with the largest number of firms to generate power law estimates. In practice, this implies using more recent years, 2006 to 2008. The main variable used in the analysis is total sales. It has been observed that in some instances a power law is only a good fit for the size distribution above a certain minimum cutoff. This is potentially an even more serious problem in this database, as the likely undersampling of smaller firms will bias the power law estimates towards zero. Following standard practice (Gabaix, 2009), we plot the data for all firms for each country, and select the minimum size cutoff by looking for a “kink” in the distribution above which the relationship between log rank and log size is approximately linear.²⁶ We restrict our empirical analysis to countries that have sales figures for at least 1000 firms. The final sample includes 44 countries.

In order to obtain reliable estimates, this paper uses three standard methods of estimating the slope of the power law ζ . The first method, based on Axtell (2001), makes direct use of the definition of the power law (Eq. (9)), which in natural logs becomes:

$$\log(\Pr(x > s)) = \log(c) - \zeta \log(s). \quad (\text{B.1})$$

For a grid of values of sales s , the estimated probability $\Pr(x > s)$ is simply the number of firms in the sample with sales greater than s divided by the total number of firms. We then regress the natural log of this probability on $\log(s)$ to obtain our first estimate of ζ . Following the typical approach in the literature, we do this for the values of s that are equidistant from each other on log scale. This implies that in absolute terms, the intervals containing low values of s are narrower than the intervals at high values of s . This is done to get a greater precision of the estimates: since there are fewer large firms, observations in small intervals for very high values of s would be more noisy.

The second approach starts with the observation that the cdf in Eq. (9) has a probability density function

$$f(s) = c\zeta s^{-(\zeta+1)}. \quad (\text{B.2})$$

To estimate this pdf, we divide the values of firm sales into bins of equal size on the log scale, and compute the frequency as the number of firms in each bin divided by the width of the bin. Since in absolute terms the bins are of unequal size, we regress the resulting frequency observations on the value of s which is the geometric mean of the endpoints of the bin (this approach follows Axtell, 2001). Note that the resulting coefficient is an estimate of $-(\zeta+1)$.²⁷

²⁵ The well-known AMADEUS database of European firms is the precursor of ORBIS, which contains all of AMADEUS plus information on non-European countries. Thus, AMADEUS is a strict subset of ORBIS.

²⁶ This is a conservative approach. The estimates obtained without imposing the minimum size cutoff yield power law coefficients even lower in absolute value, implying an even more fat-tailed distribution of firm size.

²⁷ Finally, we also regressed $\log(\text{rank} - 1/2)$ of each firm in the sales distribution on log of its sales. This is the estimator suggested by Gabaix and Ibragimov (2011), which delivers very similar results. If anything, the power law exponents implied by this estimator are even lower in absolute value than those reported in this appendix.

Table A2 reports the results. The left panel reports estimates of Eq. (B.1), the right panel, Eq. (B.2). (Note that the right panel's estimates are of $-(\zeta+1)$, thus they should differ from the right panel by about -1 .) The columns report the power law coefficient, the R^2 , and the p -value of the test that the coefficient differs from -1 (-2 in the right panel). Several things are worth noting about these results. First, the power law approximates the data well: the median R^2 is 0.99, with the minimum R^2 of 0.95. Second, most of the power law coefficients are very close to -1 in absolute terms, and many are not statistically different from -1 . Those that are statistically different from -1 tend to be lower in absolute value, implying that if the firm size distribution follows a power law in those countries, it is even more fat-tailed than Zipf. The least fat-tailed country, Serbia, has the power law exponent of about -1.18 or -1.16 , still quite far from -2 and thus comfortably within the Zipf's Law range. Finally, the country sample is diverse: it includes major European economies (France, Germany, Netherlands), smaller E.U. accession countries (Czech Republic, Estonia), major middle income countries (Brazil, Argentina), as well as the two largest emerging markets (India and China). All in all, in this sample of 44 countries with very different characteristics, the distributions of firm size are remarkably consistent with Zipf's Law.

It is important to note that these results do not establish that the distribution of firm size in these countries follows a power law, as opposed to some other distribution. Indeed, as noted by Gabaix (2009), with more parameters (allowing for more curvature), one will always fit the data better. Rather, Gabaix (2009) suggests that what is important is whether a power law provides a good fit to the data, which appears to be the case in our results.

Table A1

Top 49 countries and the rest of the world in terms of 2004 GDP.
Source: The World Bank (2007b).

Country	GDP/world GDP	Country	GDP/world GDP
United States	0.300	Indonesia	0.006
Japan	0.124	South Africa	0.006
Germany	0.076	Norway	0.006
France	0.054	Poland	0.005
United Kingdom	0.044	Finland	0.005
Italy	0.041	Greece	0.004
China	0.028	Venezuela, RB	0.004
Canada	0.026	Thailand	0.004
Brazil	0.021	Portugal	0.003
Spain	0.020	Colombia	0.003
India	0.017	Nigeria	0.003
Australia	0.016	Algeria	0.003
Russian Federation	0.015	Israel	0.003
Mexico	0.015	Philippines	0.003
Netherlands	0.015	Malaysia	0.002
Korea, Rep.	0.011	Ireland	0.002
Sweden	0.010	Egypt, Arab Rep.	0.002
Switzerland	0.010	Pakistan	0.002
Belgium	0.009	Chile	0.002
Argentina	0.008	New Zealand	0.002
Saudi Arabia	0.007	Czech Republic	0.002
Austria	0.007	United Arab Emirates	0.002
Iran, Islamic Rep.	0.007	Hungary	0.002
Turkey	0.007	Romania	0.002
Denmark	0.006	Rest of the world	0.027

Notes: Ranking of top 49 countries and the rest of the world in terms of 2004 U.S.\$ GDP. We include Hong Kong, POC, and Singapore in Rest of the world.

Table A2

Country-by-country estimates of power laws in firm size.

Country	CDF estimation			PDF estimation		
	PL coef.	R ²	p-Value	PL coef.	R ²	p-Value
Argentina	−1.046**	0.988	0.243	−2.039**	0.994	0.466
Australia	−0.992**	0.986	0.838	−1.905**	0.994	0.076
Austria	−0.695**	0.963	0.000	−1.677**	0.989	0.000
Belgium	−0.972**	0.999	0.011	−1.956**	0.998	0.150
Bosnia and Herzegovina	−1.022**	0.990	0.508	−2.036**	0.992	0.550
Brazil	−0.918**	0.966	0.162	−1.892**	0.991	0.096
Bulgaria	−0.981**	0.979	0.686	−2.007**	0.992	0.908
Canada	−0.888**	0.989	0.004	−1.913**	0.995	0.069
China	−1.117**	0.976	0.060	−2.091**	0.996	0.061
Croatia	−1.094**	0.988	0.034	−2.120**	0.992	0.074
Czech Republic	−1.083**	0.992	0.020	−2.072**	0.998	0.031
Denmark	−0.776**	0.950	0.003	−1.684**	0.987	0.001
Estonia	−1.017**	0.986	0.674	−2.067**	0.987	0.389
Finland	−0.869**	0.989	0.001	−1.879**	0.997	0.006
France	−0.886**	0.999	0.000	−1.894**	1.000	0.000
Germany	−0.853**	0.999	0.000	−1.960**	0.981	0.653
Greece	−0.992**	0.997	0.620	−1.951**	0.998	0.089
Hungary	−0.953**	0.995	0.050	−1.987**	0.996	0.741
India	−0.975**	0.988	0.476	−1.954**	0.995	0.319
Ireland	−0.761**	0.998	0.000	−1.718**	0.999	0.000
Italy	−1.030**	0.996	0.172	−2.037**	0.999	0.093
Japan	−0.955**	0.990	0.177	−1.985**	0.996	0.716
Latvia	−1.118**	0.989	0.011	−2.054**	0.995	0.281
Lithuania	−1.153**	0.992	0.001	−2.151**	0.996	0.009
Macedonia	−1.109**	0.999	0.000	−2.095**	0.990	0.176
Netherlands	−0.906**	0.994	0.002	−1.917**	0.995	0.082
Norway	−1.045**	0.970	0.454	−1.975**	0.997	0.516
Poland	−1.086**	0.987	0.051	−2.125**	0.995	0.028
Portugal	−0.919**	0.996	0.001	−1.924**	0.999	0.001
Korea	−0.880**	0.999	0.000	−1.860**	1.000	0.000
Romania	−1.002**	0.990	0.956	−2.047**	0.995	0.349
Russia	−1.039**	0.996	0.086	−2.027**	0.998	0.384
Serbia	−1.181**	0.989	0.001	−2.163**	0.996	0.004
Singapore	−0.888**	0.979	0.021	−1.825**	0.995	0.002
Slovakia	−1.139**	0.990	0.003	−2.124**	0.996	0.018
Slovenia	−0.993**	0.986	0.846	−1.998**	0.989	0.981
Spain	−0.978**	1.000	0.005	−2.011**	0.997	0.769
Sweden	−0.884**	0.997	0.000	−1.895**	0.998	0.002
Switzerland	−0.791**	0.990	0.000	−1.760**	0.996	0.000
Taiwan POC	−0.889**	0.989	0.003	−1.863**	0.991	0.031
Thailand	−0.956**	0.976	0.381	−1.953**	0.994	0.358
Ukraine	−1.058**	0.991	0.102	−2.007**	0.999	0.802
United Kingdom	−1.010**	0.975	0.856	−2.017**	0.992	0.775

Notes: This table reports the estimated of power laws in firm size across countries. Column “PL coef.” reports the coefficient on the power law for each country, the second column reports the R², the third column reports the p-value of the test that the power law coefficient is statistically different from −1 (−2 in the right panel). The estimates are based on firm-level sales data from ORBIS. Variable definitions, sources, and estimation techniques are described in detail in the text.

** Significant at the 1% level.

References

- Alvarez, Fernando, Lucas Jr., Robert E., 2007. General equilibrium analysis of the Eaton–Kortum model of international trade. *Journal of Monetary Economics* 54 (6), 1726–1768 (September).
- Anderson, James, van Wincoop, Eric, 2004. Trade costs. *Journal of Economic Literature* 42 (3), 691–751.
- Arkolakis, Costas, 2010. Market penetration costs and the new consumers margin in international trade. *Journal of Political Economy* 118 (6), 1151–1199 (December).
- Arkolakis, Costas, Demidova, Svetlana, Klenow, Peter J., Rodríguez-Clare, Andrés, 2008. Endogenous variety and the gains from trade. *American Economic Review Papers and Proceedings* 98 (2), 444–450 (May).
- Arkolakis, Costas, Costinot, Arnaud, Rodríguez-Clare, Andrés, 2012. New trade models, same old gains? *American Economic Review* 102 (1), 94–130 (February).

- Atkeson, Andrew, Burstein, Ariel, 2010. Innovation, firm dynamics, and international trade. *Journal of Political Economy* 3, 433–484 (June).
- Axtell, Robert L., 2001. Zipf distribution of U.S. firm sizes. *Science* 293 (5536), 1818–1820 (September).
- Bekkers, Eddy, Francois, Joseph, 2008. Heterogeneous firms, the structure of industry, and trade under oligopoly. Working Paper 0811. Department of Economics, Johannes Kepler University of Linz. August.
- Bernard, Andrew B., Eaton, Jonathan, Bradford Jensen, J., Kortum, Samuel, 2003. Plants and productivity in international trade. *American Economic Review* 93 (4), 1268–1290 (September).
- Bernard, Andrew, Bradford Jensen, J., Redding, Stephen, Schott, Peter, 2007. Firms in international trade. *Journal of Economic Perspectives* 21, 105–130.
- Broda, Christian, Weinstein, David, 2006. Globalization and the gains from variety. *Quarterly Journal of Economics* 121 (2), 541–585 (May).
- Chaney, Thomas, 2008. Distorted gravity: the intensive and extensive margins of international trade. *American Economic Review* 4, 1707–1721 (September).
- Crozet, Matthieu, Koenig, Pamina, 2010. Structural gravity equation with extensive and intensive margins. *Canadian Journal of Economics* 43 (1), 41–62 (February).
- di Giovanni, Julian, Levchenko, Andrei A., 2010. Firm entry, trade, and welfare in Zipf’s world. NBER Working Paper No. 16313. August.
- di Giovanni, Julian, Levchenko, Andrei A., 2011. Country size, international trade, and aggregate fluctuations in granular economies. NBER Working Paper No. 17335. August.
- di Giovanni, Julian, Levchenko, Andrei A., Rancière, Romain, 2011. Power laws in firm size and openness to trade: measurement and implications. *Journal of International Economics* 85 (1), 42–52 (September).
- Djankov, Simeon, La Porta, Rafael, Lopez-de-Silanes, Florencio, Shleifer, Andrei, 2002. The regulation of entry. *Quarterly Journal of Economics* 117 (1), 1–37.
- Easterly, William, Reshef, Ariell, Schwenkenberg, Julia, 2009. The power of exports. World Bank Policy Research Working Paper 5081. October.
- Eaton, Jonathan, Kortum, Samuel S., 2005. Technology in the global economy: a framework for quantitative analysis. Book Manuscript. New York University and University of Chicago.
- Eaton, Jonathan, Kortum, Samuel S., Kramarz, Francis, 2011. An anatomy of international trade: evidence from French firms. *Econometrica* 5, 1453–1498 (September).
- Eaton, Jonathan, Kortum, Samuel S., Sotelo, Sebastian, 2012. International trade: linking micro and macro. NBER Working Paper No. 17864. February.
- Feenstra, Robert C., 2010. Measuring the gains from trade under monopolistic competition. *Canadian Journal of Economics* 43 (1), 1–28 (February).
- Fujiwara, Yoshi, Aoyama, Hideaki, Di Guilmi, Corrado, Souma, Wataru, Gallegati, Mauro, 2004. Gibrat and Pareto–Zipf Revisited with European firms. *Physica A: Statistical Mechanics and its Applications* 344 (1–2), 112–116.
- Gabaix, Xavier, 2009. Power laws in economics and finance. *Annual Reviews of Economics* 1, 255–293.
- Gabaix, Xavier, Ibragimov, Rustam, 2011. Rank-1/2: a simple way to improve the OLS estimation of tail exponents. *Journal of Business and Economic Statistics* 29 (1), 24–39 (January).
- Goldberg, Pinelopi, Khandelwal, Amit, Pavcnik, Nina, Topalova, Petia, 2009. Trade liberalization and new imported inputs. *American Economic Review Papers and Proceedings* 99 (2), 494–500 (May).
- Goldberg, Pinelopi, Khandelwal, Amit, Pavcnik, Nina, Topalova, Petia, 2010. Imported intermediate inputs and domestic product growth: evidence from India. *Quarterly Journal of Economics* 125 (4), 1727–1767 (November).
- Helpman, Elhanan, Melitz, Marc, Rubinstein, Yona, 2008. Estimating trade flows: trading partners and trading volumes. *Quarterly Journal of Economics* 123 (2), 441–487 (May).
- Hinloopen, Jeroen, van Marrewijk, Charles, 2006. Comparative advantage, the rank-size rule, and Zipf’s law. Tinbergen Institute Discussion Paper 06-100/1.
- International Monetary Fund, 2007. Directions of Trade Statistics.
- Jones, Charles, 2011. Intermediate goods and weak links in the theory of economic development. *American Economic Journal: Macroeconomics* 3 (2), 1–28 (April).
- Krugman, Paul, 1980. Scale economies, product differentiation, and the pattern of trade. *American Economic Review* 70 (5), 950–959.
- Luttmer, Erzo, 2007. Selection, growth, and the size distribution of firms. *Quarterly Journal of Economics* 122, 1103–1144.
- Melitz, Marc J., 2003. The impact of trade on intra-industry reallocations and aggregate industry productivity. *Econometrica* 6, 1695–1725 (November).
- Neary, J. Peter, 2010. Two and a half theories of trade. *The World Economy* 33 (1), 1–19 (January).
- Okuyama, K., Takayasu, M., Takayasu, H., 1999. Zipf’s law in income distribution of companies. *Physica A: Statistical Mechanics and its Applications* 269 (1), 125–131.
- The World Bank, 2007a. Doing Business Indicators.
- The World Bank, 2007b. World Development Indicators.
- Waugh, Michael, 2010. International trade and income differences. *American Economic Review* 100 (5), 2093–2124 (December).