

NBER WORKING PAPER SERIES

FIRM ENTRY, TRADE, AND WELFARE IN ZIPF'S WORLD

Julian di Giovanni
Andrei A. Levchenko

Working Paper 16313
<http://www.nber.org/papers/w16313>

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
August 2010

We are grateful to Roc Armenter, Jonathan Eaton, Chris House, Amit Khandelwal, Francesc Ortega, Romain Ranciere, 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, 2009 NBER ITI Spring Meetings, 2009 CEPR ESSIM, and the Philadelphia Area International Trade Workshop for helpful suggestions. 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, its management, or the National Bureau of Economic Research.

NBER working papers are circulated for discussion and comment purposes. They have not been peer-reviewed or been subject to the review by the NBER Board of Directors that accompanies official NBER publications.

© 2010 by Julian di Giovanni and Andrei A. Levchenko. All rights reserved. Short sections of text, not to exceed two paragraphs, may be quoted without explicit permission provided that full credit, including © notice, is given to the source.

Firm Entry, Trade, and Welfare in Zipf's World
Julian di Giovanni and Andrei A. Levchenko
NBER Working Paper No. 16313
August 2010
JEL No. F12,F15

ABSTRACT

Firm size follows Zipf's Law, a very fat-tailed distribution that implies a few large firms account for 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 largest 50 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 -- 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.

Julian di Giovanni
International Monetary Fund
700 19th Street NW
Washington, DC 20431
jdigiovanni@imf.org

Andrei A. Levchenko
Department of Economics
University of Michigan
611 Tappan Street
Ann Arbor, MI 48109
and NBER
alev@umich.edu

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. 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.¹ 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, Kortum and Kramarz (2008). 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 extensive 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 firms 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

¹This has been documented by Axtell (2001) for the census of U.S. firms, and by di Giovanni, Levchenko and Rancière (2010) for the census of French firms. Similar findings obtain for several European countries (Fujiwara, Aoyama, Di Guilmi, Souma and Gallegati 2004) and Japan (Okuyama, Takayasu and Takayasu 1999). Other phenomena known to follow power laws include city size, income and wealth, and CEO compensation (Gabaix 2009).

welfare impact of the fall in the iceberg costs.² 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 or not, 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.³ 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, La Porta, Lopez-de-Silanes and Shleifer (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.⁴ 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

²The disappearing domestic varieties (the domestic extensive margin) have a correspondingly negative welfare impact.

³A 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.

⁴To give one example, the official cost of following all the procedures to set up a business ranges from 0.5% of per capita GDP in the U.S. to 4.6 times per capita GDP in the Dominican Republic.

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, Jensen, Redding and Schott 2007). The very same model we analyze in this paper has been used in dozens of studies to examine the firm’s decision whether or not to export (e.g., Chaney 2008), or how much to export (e.g., Arkolakis 2008). 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, Demidova, Klenow and Rodríguez-Clare (2008) and Arkolakis, Costinot and Rodríguez-Clare (2010) 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.⁵ Feenstra (2010) shows that in a Melitz model with free entry, the positive welfare impact of newly imported varieties is exactly cancelled out by the negative welfare impact of disappearing domestic varieties, resulting in gains from variety that are precisely nil.⁶

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

⁵In a dynamic two-country model of trade and innovation, Atkeson and Burstein (2010) argue that the impact of the extensive margin on the rate of innovation is likely to be small as well.

⁶This result does not appear to be general. As we show below, though the welfare impact of the domestic extensive margin can indeed be negative, the two extensive margins do not cancel perfectly in any of the analytical or quantitative models we consider in this paper, including models with a fixed mass of entrepreneurs; and models with free entry.

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 Feenstra's 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 cancelling out perfectly. Instead we show that they are both vanishingly small in absolute value. In our view, this is a more robust reason why the extensive margin is not important for welfare. 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 firms 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 not important, 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 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, Khandelwal, Pavcnik and Topalova 2009, 2010). 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 overwhelmingly used this type of model for the study of the extensive margin. In other words, some other mechanisms need to be specified for the extensive margin to have a discernible welfare impact.

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 main quantitative results. Section 4 develops a model featuring an alternative assumption on the entry of firms. Section 5 concludes.

2 Theoretical Framework

The world is comprised 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

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

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 number of varieties consumed in country i , coming from all countries. It is well known that demand for variety k is equal to

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

in country i , 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 Eaton et al. (2008), Chaney (2008), and Arkolakis (2008). 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. Section 4 presents an alternative model, in which the mass of potential entrepreneurs can adjust, as in Krugman (1980) and Melitz (2003). The main quantitative results are remarkably similar.

Each potential entrepreneur can produce a unique CES variety, and thus has some market power: it faces the demand for its variety given by (1). There are both fixed and variable costs of production and trade. Each entrepreneur's type is given by the marginal cost $a(k)$. On the basis of this cost, each entrepreneur in country j decides whether or not to pay the fixed cost of production f_{jj} , and which, if any, export markets to serve. To start exporting from country j to country i , a firm must pay the fixed cost f_{ij} , and an iceberg per-unit cost of $\tau_{ij} > 1$.⁷

If entry is important – be it into production, or the export markets – it is becoming clear that one of the ways it matters is through the varieties available as intermediate inputs in production. Jones (2007, 2008) shows that the use of intermediate inputs creates a TFP multiplier that goes some

⁷That is, the firm in country j must ship $\tau_{ij} > 1$ units to country i in order for one unit of the good to arrive there. We normalize the iceberg cost of domestic sales to one: $\tau_{jj} = 1$.

way to explaining observed income differences across countries. Acemoglu, Antràs and Helpman (2007), Cowan and Neut (2007), and Costinot (2009) argue that in countries with worse institutions, production will use fewer intermediates, adversely affecting productivity. On the trade side, it has also been argued that imported intermediates play an important role in domestic productivity.⁸

Our model features foreign and domestic intermediate inputs and the associated multiplier. 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 marginal cost $a(k)$ must use this many input bundles to produce one unit of output. An input bundle has a cost $c_j = w_j^\beta P_j^{1-\beta}$, 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 selling to country i faces a demand curve given by (1), and has a marginal cost $\tau_{ij} c_j a(k)$ of serving this market. 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 quantity supplied is equal to $\frac{X_i}{P_i^{1-\varepsilon}} \left(\frac{\varepsilon}{\varepsilon-1} \tau_{ij} c_j a(k) \right)^{-\varepsilon}$, and the total ex-post variable profits are:

$$\rho_{ij}^V(k) = \frac{X_i}{\varepsilon P_i^{1-\varepsilon}} \left(\frac{\varepsilon}{\varepsilon-1} \tau_{ij} c_j a(k) \right)^{1-\varepsilon}. \quad (3)$$

Note that these are variable profits of a firm in country j from selling its good to country i only. These expressions are valid for each country pair i, j , including domestic sales: $j = i$.

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 marginal cost a_{ij} , above which firms in country j do not serve market i . We assume (and later verify in the calibration exercise), that all firms that decide to export abroad are sufficiently productive to also serve their domestic markets. On the other hand, there is a range of productivities for which firms serve their domestic markets, but choose not to export. In this case, firms with marginal cost above a_{jj} in country j do not operate at all. The cutoff a_{ij} characterizes the entrepreneur in j who earns zero profits from shipping to country i :

$$a_{ij} = \frac{\varepsilon-1}{\varepsilon} \frac{P_i}{\tau_{ij} c_j} \left(\frac{X_i}{\varepsilon c_j f_{ij}} \right)^{\frac{1}{\varepsilon-1}}. \quad (4)$$

Closing the model involves finding expressions for a_{ij} , P_i , and w_i for all $i, j = 1, \dots, N$. The price level for country i can be expressed as follows:

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

⁸For instance, Amiti and Konings (2007), Kasahara and Rodrigue (2008), Goldberg et al. (2009, 2010), Luong (2008), and Halpern, Koren and Szeidl (2009) provide empirical evidence that newly available foreign intermediate inputs increased the TFP of individual firms.

where J_{ij} is the set of varieties exported from country j to country i . In order to solve the model, 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 dispersion.⁹ The cdf of productivity is given by:

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

It is then straightforward to show that the marginal cost, a , has a distribution function $G(a) = (ba)^\theta$. The price level then becomes, after plugging in the expression for a_{ij} in (4):

$$P_i = \left(\sum_{j=1}^N n_j \int_0^{a_{ij}} \left[\frac{\varepsilon}{\varepsilon-1} \tau_{ij} c_j a \right]^{1-\varepsilon} dG(a) \right)^{\frac{1}{1-\varepsilon}} \quad (5)$$

$$= \frac{1}{b} \left[\frac{\theta}{\theta - (\varepsilon - 1)} \right]^{-\frac{1}{\theta}} \frac{\varepsilon}{\varepsilon - 1} \left(\frac{X_i}{\varepsilon} \right)^{-\frac{\theta - (\varepsilon - 1)}{\theta(\varepsilon - 1)}} \left(\sum_{j=1}^N n_j \left(\frac{1}{\tau_{ij} c_j} \right)^\theta \left(\frac{1}{c_j f_{ij}} \right)^{\frac{\theta - (\varepsilon - 1)}{\varepsilon - 1}} \right)^{-\frac{1}{\theta}}. \quad (6)$$

Having expressed P_i and a_{ij} in terms of X_i and c_i , for all $i, j = 1, \dots, N$, it remains to close the model by solving for the X_i 's and w_i 's. To do this, 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 .

Proposition 1 *Total profits of firms based in country i are a constant multiple of total expenditure:*
 $\Pi_i = \frac{\varepsilon-1}{\theta\varepsilon} X_i$.

Proof: See Appendix A. ■

Since by definition 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, each country's GDP is a constant multiple its total labor income:

$$X_i = \frac{1}{\beta \left(1 - \frac{\varepsilon-1}{\theta\varepsilon}\right)} w_i L_i. \quad (7)$$

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

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

⁹The Pareto assumption is by far the most common distributional assumption made in the heterogeneous firms 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.

Using the expression for a_{ji} in (4), and P_j in (6), total exports from i to j become:

$$X_{ji} = \frac{n_i \left(\frac{1}{\tau_{ji} c_i} \right)^\theta \left(\frac{1}{c_i f_{ji}} \right)^{\frac{\theta - (\varepsilon - 1)}{\varepsilon - 1}}}{\sum_{l=1}^N n_l \left(\frac{1}{\tau_{jl} c_l} \right)^\theta \left(\frac{1}{c_l f_{jl}} \right)^{\frac{\theta - (\varepsilon - 1)}{\varepsilon - 1}}} X_j. \quad (8)$$

Using the trade balance conditions, $X_i = \sum_{j=1}^N X_{ji}$ for each $i = 1, \dots, N$, the expression for total GDP, X_i , in (7), and the definition of c_i leads to the following system of equations in w_i :

$$w_i L_i = \sum_{j=1}^N \frac{n_i \left(\frac{1}{\tau_{ji} w_i^\beta P_i^{1-\beta}} \right)^\theta \left(\frac{1}{w_i^\beta P_i^{1-\beta} f_{ji}} \right)^{\frac{\theta - (\varepsilon - 1)}{\varepsilon - 1}}}{\sum_{l=1}^N n_l \left(\frac{1}{\tau_{jl} w_l^\beta P_l^{1-\beta}} \right)^\theta \left(\frac{1}{w_l^\beta P_l^{1-\beta} f_{jl}} \right)^{\frac{\theta - (\varepsilon - 1)}{\varepsilon - 1}}} w_j L_j, \quad (9)$$

$i = 1, \dots, N$. 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 equations (9) for wages and (6) 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

It has been argued that in the data, the distribution of firm size follows a power law, with an exponent close to 1 in absolute value. In this section, we first build a bridge between the model and the data by showing that the distribution of firm sales in the workhorse model outlined above does indeed follow a power law. Consequently, we argue that the distribution of firm size in the data places a key restriction on the important parameter values in the model. Finally, we review the available empirical evidence on the firm size distribution.

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

$$\Pr(x > s) = cs^{-\zeta}. \quad (10)$$

It turns out that the baseline Melitz-Pareto model delivers a power law in firm size. In our model, the sales of a firm as a function of its marginal cost are: $x(a) = Ca^{1-\varepsilon}$, where the constant C reflects the size of overall demand, and we drop the country subscripts. Under the assumption that $1/a \sim \text{Pareto}(b, \theta)$, the power law follows:

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

satisfying (10) for $c = (b^{\varepsilon-1}C)^{\frac{\theta}{\varepsilon-1}}$ and $\zeta = \frac{\theta}{\varepsilon-1}$. In addition, this calculation shows that $x \sim \text{Pareto}\left(b^{\varepsilon-1}C, \frac{\theta}{\varepsilon-1}\right)$.

The key point for connecting the model to the data is that in the model, the slope of the power law is given by $\frac{\theta}{\varepsilon-1}$. 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 ζ ? Available estimates put it very close to 1, suggesting that the distribution of firm size follows Zipf's Law. Figure 1 reproduces the now famous power law for firm size in the U.S. estimated by Axtell (2001). The fit of this relationship is typically very close: it is common to observe R-squareds in excess of 0.99. Using a variety of estimation techniques, Axtell 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. It will become important below that the coefficient estimates are never significantly different from 1, and indeed never very far from 1 in absolute terms as well.¹⁰

The question remains whether Zipf's Law obtains in the firm size distributions for many countries. Currently, no comprehensive set of results exists. di Giovanni et al. (2010) show that Zipf's Law holds among French firms. 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 C, 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.¹¹

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

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

¹⁰Strictly speaking, when not all firms export, selection into exporting implies that the power law exponent estimated on total sales – domestic plus exporting – is lower than $\theta/(\varepsilon - 1)$. di Giovanni et al. (2010) explore this bias in detail using the census of French firms, and suggest several corrections to the estimating procedure that can be used to estimate $\theta/(\varepsilon - 1)$ in an internally consistent way. Their analysis shows that the bias introduced by selection into exporting is not large. Corrected estimates obtained by di Giovanni et al. (2010) show that $\theta/(\varepsilon - 1)$ is about 1.05, roughly the same as the value used in this paper.

¹¹Other 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, Reshef and Schwenkenberg 2009) also follow power laws with an exponent close to -1 .

to w_i/P_i , which is also a measure of welfare.¹² 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 j . Using equation (8), setting $i = j$ and rearranging yields the following relationship:

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

Plugging this expression into the price level (6) and rearranging, welfare under trade in this economy can be written as:

$$\frac{w_i}{P_i} = \left\{ \frac{1}{b} \left[\frac{\theta}{\theta - (\varepsilon - 1)} \right]^{-\frac{1}{\theta}} \frac{\varepsilon}{\varepsilon - 1} n_i^{-\frac{1}{\theta}} \left(\frac{L_i}{f_{ii} \varepsilon \beta \left(1 - \frac{\varepsilon - 1}{\varepsilon \theta} \right)} \right)^{-\frac{\theta - (\varepsilon - 1)}{\theta(\varepsilon - 1)}} \pi_{ii}^{\frac{1}{\theta}} \right\}^{-\frac{1}{\beta - (1 - \beta) \frac{\theta - (\varepsilon - 1)}{\theta(\varepsilon - 1)}}}. \quad (11)$$

This allows us to represent real income per capita in each country relative to the U.S. as a product of several components:

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

A special case of this expression is obtained if we adopt the assumption in Alvarez and Lucas (2007), Chaney (2008), and di Giovanni and Levchenko (2009) 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:

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

This expression is similar in spirit to Waugh (2009), with some key differences. The similarity is in the contribution of trade to income differences, which is summarized simply by the relative openness $\left(\frac{\pi_{ii}}{\pi_{US,US}} \right)$. The difference is that in our model entry costs also matter (the $\frac{f_{ii}}{f_{US,US}}$ term), and there is a “home market 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:

¹²Welfare is proportional to the real wage even though in this economy there are profits. From Proposition 1, 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. (7).

$\beta = 0.5$ from Jones (2008), $\varepsilon = 6$ (Anderson and van Wincoop 2004), and $\theta = 5.3$, designed to match the power law exponent on firm size to U.S. data, $\frac{\theta}{\varepsilon-1} = 1.06$ (Axtell 2001). 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 $\frac{f_{ii}}{f_{US,US}}$ 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 $\frac{\pi_{ii}}{\pi_{US,US}}$ 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 $\frac{\pi_{ii}}{\pi_{US,US}}$ 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 $\frac{f_{ii}}{f_{US,US}}$ term, the exponent goes up from 0.02 to 0.22 in absolute value, a tenfold increase. By contrast, the exponent on the $\frac{\pi_{ii}}{\pi_{US,US}}$ 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 to:¹³

$$\pi_{ii}^{-\frac{1}{\theta}} \frac{1}{\beta - (1-\beta) \frac{\theta - (\varepsilon - 1)}{\theta(\varepsilon - 1)}}. \quad (12)$$

A few things are notable about this expression. First, at a given π_{ii} , gains from trade are decreasing in $\frac{\theta}{\varepsilon - 1}$ as long as $\beta\varepsilon > 1$.¹⁴ 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%).

Second, at a given level of trade openness (π_{ii}), gains from trade are increasing in the share of intermediate goods in the input bundle, $(1 - \beta)$. This is an intermediate goods multiplier effect akin to Jones (2008): the more foreign varieties are used as intermediate goods in production, the more the country reaps a double benefit from trade: first as an increase in labor productivity due to foreign intermediates, and second as consumers of those foreign varieties.

Finally, in order to get a sense of the gains from trade, it is sufficient to simply look at the share of spending on domestic goods. This feature has been noted about Ricardian models (Eaton and Kortum 2002), as well as monopolistic competition models such as the one in this paper (Arkolakis et al. 2008). Arkolakis et al. (2010) provide a unified treatment of the conditions under which this result holds.

At the same time, in the presence of intermediate goods it is no longer the case that in order to calculate the overall gains from trade, one needs to know only π_{ii} and the estimated elasticity

¹³To find an expression for gains from trade, it is useful to write out the autarky price level and welfare. Setting $N = 1$ in equation (6) and dropping country subscripts, the autarky price level becomes:

$$P_A = \frac{1}{b} \left[\frac{\theta}{\theta - (\varepsilon - 1)} \right]^{-\frac{1}{\theta}} \frac{\varepsilon}{\varepsilon - 1} n^{-\frac{1}{\theta}} c \left(\frac{X}{\varepsilon c f} \right)^{-\frac{\theta - (\varepsilon - 1)}{\theta(\varepsilon - 1)}}.$$

Using the expression for X in (7) and c , we can write welfare in autarky as follows:

$$\frac{w_A}{P_A} = \left\{ \frac{1}{b} \left[\frac{\theta}{\theta - (\varepsilon - 1)} \right]^{-\frac{1}{\theta}} \frac{\varepsilon}{\varepsilon - 1} n^{-\frac{1}{\theta}} \left(\frac{L}{f\varepsilon\beta \left(1 - \frac{\varepsilon - 1}{\varepsilon\theta}\right)} \right)^{-\frac{\theta - (\varepsilon - 1)}{\theta(\varepsilon - 1)}} \right\}^{-\frac{1}{\beta - (1-\beta) \frac{\theta - (\varepsilon - 1)}{\theta(\varepsilon - 1)}}}.$$

It is immediate from comparing this expression for autarky welfare to (11) that the two differ only by the term $\pi_{ii}^{-\frac{1}{\theta} \frac{1}{\beta - (1-\beta) \frac{\theta - (\varepsilon - 1)}{\theta(\varepsilon - 1)}}}$, yielding equation (12).

¹⁴This 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 2008).

of trade with respect to (variable) trade costs, as argued by Arkolakis et al. (2008). As can be gleaned from equation (8), the elasticity of bilateral trade with respect to τ_{ji} is $-\theta$. It is true that, without intermediate goods, the gains from trade are given by $\pi_{ii}^{-\frac{1}{\theta}}$, so that the exponent on π_{ii} is exactly the inverse of that elasticity. However, in the presence of intermediate inputs, that is not the case: the exponent on π_{ii} is now $-\frac{1}{\theta} \frac{1}{\beta - (1-\beta) \frac{\theta - (\varepsilon - 1)}{\theta(\varepsilon - 1)}}$. In other words, in order to assess the gains from trade, we can no longer rely on one potentially observable object – the elasticity of bilateral trade with respect to τ_{ji} – and instead need to take a stand on other parameters of the model, namely β and ε .

In this context, it is worth noting that the firm size distribution provides an alternative source of information regarding the model parameters. While at first blush, estimating the elasticity of trade volumes with respect to trade costs may seem straightforward, in fact one must typically make a series of parametric assumptions about how the true trade costs τ_{ji} are related to observables such as distance or tariff barriers. Things are further complicated by the fact that in the Melitz model both τ_{ji} and f_{ji} affect trade volumes, but with different elasticities. Since a typical gravity regression cannot distinguish between the two, yet more assumptions on the nature of fixed and variable costs are needed to back out θ . The firm size distribution, by providing an estimate of $\theta/(\varepsilon - 1)$, is an arguably cleaner way to calibrate the parameters of the model. In fact, as we show in this paper, some key results actually depend on this combination of parameters, rather than θ and ε individually.

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 in exports (or indeed domestic production) goes to zero.

The price level, (5), can be rewritten as a function of the extensive margin as follows:

$$P_i = \left(\frac{\varepsilon}{\varepsilon - 1} b^{\varepsilon - 1} \frac{\theta}{\theta - (\varepsilon - 1)} \sum_{j=1}^N n_j (\tau_{ij} c_j)^{1 - \varepsilon} G(a_{ij})^{\frac{\theta - (\varepsilon - 1)}{\theta}} \right)^{\frac{1}{1 - \varepsilon}}. \quad (13)$$

Here, the price level is expressed in terms of the share of firms from country j supplying country i , $G(a_{ij})$, precisely because it is the extensive margin: in any policy experiment, the 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:

$$P = \left\{ \frac{\varepsilon}{\varepsilon - 1} b^{\varepsilon-1} \frac{\theta}{\theta - (\varepsilon - 1)} n \left(G(a_D)^{\frac{\theta - (\varepsilon - 1)}{\theta}} + (N - 1) \tau^{1-\varepsilon} G(a_X)^{\frac{\theta - (\varepsilon - 1)}{\theta}} \right) \right\}^{\frac{1}{\beta(1-\varepsilon)}}, \quad (14)$$

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, $G(a_X)$. From the expression for the price level (14), 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 $\theta \rightarrow (\varepsilon - 1)$.

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 the many additional inferior brands of television sets become available.

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 extensive margin matters less under Zipf's Law, and trade

openness matters more, as it allows the country to access the extremely productive varieties from abroad.

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 an additional remark regarding our modeling approach to fixed costs. Arkolakis (2008) 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 B presents a 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. For these reasons, we choose to adopt the standard formulation of fixed costs of entry in our analysis.

3 Quantitative Evidence

In order to fully solve the model numerically, we must find the wages and price levels for each country, w_i and P_i , using the system of equations given by (6) and (9). To solve this system, we must 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. We now discuss how we calibrate each parameter value.

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 firm sales follow a power law with the exponent equal to $\frac{\theta}{\varepsilon-1}$. In the data, firm sales follow a power law with the exponent close to 1. Axtell (2001) reports 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 (2008).

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 extremely 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, Melitz and Rubinstein (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.¹⁵ 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 absolute level of $f_{US,US}$ is set to ensure an interior solution for the domestic production cutoff.¹⁶ Then, we use the information from the Doing Business Indicators database (The World Bank 2007a) to set f_{ii} for every other country relative to the U.S.. In this application, 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

¹⁵In di Giovanni and Levchenko (2009), as a robustness check, we also computed τ_{ij} using the estimates of Eaton and Kortum (2002). The advantage of the Helpman et al. (2008) estimates is that they are obtained in an empirical model that accounts explicitly for both fixed and variable costs of exporting, and thus corresponds most closely to the theoretical structure in our paper.

¹⁶That is, we set $f_{US,US}$ to a level just high enough that $a_{ji} < 1/b$ 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 $a(k)$.

inputs rather than value. To be precise, if according to the Doing Business Indicators database, in country i it takes 10 times longer to register a business than in the U.S., then $f_{ii} = 10 \times f_{US,US}$.

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 .¹⁷ 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.¹⁸ 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 1.

3.1 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. However, since the object of the paper is to examine the role of trade openness in welfare, it is more important that the model matches well the bilateral and overall trade volumes observed in the data. Comparing bilateral trade patterns generated by the model to the actual data is a good test of the model's success in describing the world economy, since the calibration procedure does not use any information on actual trade patterns, only country GDPs and estimated bilateral trade costs.

Figure 2 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

¹⁷An 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.

¹⁸We 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.

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. 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. Figure 3 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 2 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 0.72.¹⁹

Overall, though the model calibration does not use any information on trade volumes, it fits bilateral trade data quite well. We now turn to the analysis of welfare gains from reduction in entry costs and trade barriers implied by the model.

3.2 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.. 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 3 reports the associated welfare gains. The top panel presents the baseline calibration, in which firm-size distribution is set to match Zipf's Law. The bottom panel reports the alternative counterfactual calibration, in which $\theta/(\varepsilon - 1) = 2$. Since by construction f_{ij} affects entry, but not the variable costs of existing firms, we attribute all of the welfare gains to the extensive margin. The welfare gains are small. We can see that even a dramatic drop (6-fold on average) in the fixed

¹⁹We also experimented with increasing the number of countries in the simulation to 60. The model fit the data well, though it over-predicted the overall average trade openness of countries by slightly more than the 50-country model. In addition, there are more zeros in bilateral trade data in the 60-country sample compared to the 50-country one. (With 50 countries, among the 2500 possible unidirectional bilateral trade flows, only 18 are zeros.) For these reasons we confine our analysis to the largest 50 countries.

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%. Zipf’s Law matters a great deal for this conclusion. The bottom panel reports that the welfare gain from the same reduction in entry barriers is on average 40.87% in the non-Zipf world. This is 12 times higher than in the Zipf’s Law 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 that features 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.

3.3 Counterfactual II: Reduction in 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 equation (13), welfare can be expressed, up to a constant that is the same in all countries and trade regimes, as follows:

$$W_i = \left[\sum_{j=1}^N n_j \left(\tau_{ij} \frac{c_j}{c_i} \right)^{1-\varepsilon} G(a_{ij})^{\frac{\theta-(\varepsilon-1)}{\theta}} \right]^{\frac{1}{\beta(\varepsilon-1)}}. \quad (15)$$

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

²⁰An 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.

expansion, we can write the proportional increase in welfare as a function of the two margins:

$$\frac{\Delta W_i}{W_i} \approx \frac{1}{\beta} \sum_{j=1}^N \varphi_{ij} \left[\underbrace{-\frac{\Delta \left(\tau_{ij} \frac{c_j}{c_i} \right)}{\tau_{ij} \frac{c_j}{c_i}}}_{\text{Intensive Margin}} + \underbrace{\frac{\theta - (\varepsilon - 1)}{\theta (\varepsilon - 1)} \frac{\Delta G(a_{ij})}{G(a_{ij})}}_{\text{Extensive Margin}} \right], \quad (16)$$

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

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

It is immediate from (16) 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. As we saw above, the calibrated value of this ratio is about 0.01.

Table 3 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, which 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 one percent), 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 Figure 4. 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 variable costs matter *more* for welfare as the economy approaches Zipf’s Law, while changes in fixed costs matter *less*.

3.4 Impact of Varying Intermediate Goods Share in Final Production

How important is the intermediate goods multiplier in our counterfactual exercises?²¹ In order to examine this question, we vary β , which is set to 0.5 in the baseline (equal shares of labor and intermediate goods for final production). Table 4 presents the total welfare changes in the two counterfactual exercises for β equal to 0.33 (smaller share of labor/larger share of intermediate goods) and 0.67 (larger share of labor/smaller share of intermediate goods) for the Zipf simulations. We also include the baseline case for comparison. The welfare impact of the same firm entry cost and trade openness changes decreases monotonically as β increases (the intermediate goods share falls). Increasing the intermediate input share from $1/2$ to $2/3$ ($\beta = 0.33$) increases the welfare gain from a reduction in fixed costs by about 50%. The same absolute change in β in the other direction ($\beta = 0.67$) reduces the welfare gains by less than a percentage point relative to the baseline case. Finally, the last column of the table presents the case of no intermediate goods multiplier ($\beta = 1$, and therefore $c = w$). We see that the welfare gains from a reduction in entry costs are reduced roughly in half compared to the baseline case. To summarize, the variation in the share of intermediates in production has an important effect on welfare. Though we do not pursue this point further here, our comparative statics suggest that gains from trade may differ substantially across countries depending on their export specialization: countries that specialize in industries requiring lots of (foreign) intermediates will gain from trade substantially more than countries that

²¹The idea of the intermediate goods multiplier in the closed-economy setting is due to Jones (2008), who in addition assumes that intermediate inputs are complements in production to get an even larger effect, explaining potentially all the variation in per-capita incomes across countries. In the multi-country model of production with endogenous varieties, it would not be possible to incorporate complementarities of inputs, since producers of varieties are monopolistically competitive, and their profit maximization problem is not well defined when the elasticity of substitution is less than 1. Thus, we adopt the setup in which the elasticity of substitution in production and consumption is the same.

simply produce output using the domestic factors of production.

4 Free Entry

The preceding analysis was carried out under the assumption of a fixed mass of firms n_i . This assumption is common in the recent literature on trade with heterogeneous firms (see, among many others, Eaton et al. 2008, Chaney 2008, and Arkolakis 2008). However, because the total profits in the economy are strictly positive, this model should be seen as a metaphor for the (relatively) short run. At the opposite extreme, the assumption adopted by the literature building on Krugman (1980) and Melitz (2003) is that the economy is populated by an infinite mass of potential entrepreneurs with zero outside option, and at any point in time there is sufficient entry to drive the firms' aggregate profits to zero. This should be seen as a metaphor for the long run, as neither the number of entrepreneurs nor the number of ideas are truly infinite at a point in time, but we would indeed expect higher net entry in response to a positive expected profit given enough time.

How does allowing for free entry affect our results? In this section, we set up a model that allows n_i to adjust, and carry out the quantitative exercise under this alternative assumption. Following Melitz (2003), suppose that there is an infinite mass of potential entrepreneurs. In order to produce, each entrepreneur in country i must first pay a fixed “exploration” cost of f_{Ei} input bundles in order to find out its marginal cost $a(k)$. Once the marginal cost is revealed, each entrepreneur decides whether to produce domestically (paying a fixed production cost f_{ii}), and whether or not to enter each of the possible foreign markets j (paying a fixed cost f_{ji}). The equilibrium mass of entrants n_i is such that expected profits are zero in each country:

$$E \left[\sum_{j=1}^N (\rho_{ij}^V(k) - c_i f_{ji}) \right] = c_i f_{Ei} \quad (17)$$

for all $i = 1, \dots, N$, where $\rho_{ij}^V(k)$ is the variable profits from serving market j from country i , given by (3). This is a system of N equations in N unknowns, n_i . Since profits are zero, the expression for the total sales in the economy, instead of (7) becomes:

$$X_i = \frac{1}{\beta} w_i L_i. \quad (18)$$

Given this expression for total sales, the rest of the equilibrium conditions remain the same. In particular, in addition to the system of equations (17) defining n_i , the model solution satisfies the equations for wages (9) and prices (6). All in all, when $N = 50$ as in our quantitative exercise, this represents a system of 149 equations in 149 unknowns, with wages in one of the countries as the numéraire.

Before solving numerically the fully-fledged calibrated model, we present a special case that can be solved in closed-form: symmetric countries. This will allow us to perform some comparative statics analytically, in order to build intuition and cross-check the numerical results. In particular, suppose that all of the countries are identical, as in section 2.3: $L_i = L$, $f_{Ei} = f_E$, $f_{ii} = f \forall i$, and $\tau_{ij} = \tau$, $f_{ij} = f_X \forall i, j, j \neq i$. Under these conditions, all the wages are the same in all countries, and we normalize them to 1. The price levels P are the same across countries as well, and there is a single cutoff $a_D = a_{ii}$ for domestic production, and a single cutoff $a_X = a_{ji}$ for exporting, to all destinations. Straightforward manipulation of equations (17), (4), and (5) yields the following solution for the equilibrium mass of firms:

$$n = \frac{\varepsilon - 1}{\theta} \left[\frac{L}{\beta \varepsilon f} \left(\frac{f}{f_E} \right)^{1 - \frac{\theta - (\varepsilon - 1)}{\theta(\varepsilon - 1)} \frac{1 - \beta}{\beta}} \right]^{\frac{1}{\beta(\beta - \frac{1 - \beta}{\varepsilon - 1})}} \times \left\{ b \frac{\varepsilon - 1}{\varepsilon} \left(\frac{\varepsilon - 1}{\theta - (\varepsilon - 1)} \right)^{\frac{1}{\theta}} \left[1 + (N - 1) \tau^{-\theta} \left(\frac{f}{f_X} \right)^{\frac{\theta - (\varepsilon - 1)}{\varepsilon - 1}} \right]^{\frac{1}{\theta}} \right\}^{\frac{1}{1 - \beta(\beta - \frac{1 - \beta}{\varepsilon - 1})}}.$$

Though it may appear complicated, there are still several conclusions that we can draw in the symmetric case. First, in the absence of input-output linkages (when $\beta = 1$ and the input bundle uses only labor), the equilibrium mass of firms simplifies to:

$$n = \frac{\varepsilon - 1}{\theta} \frac{L}{\varepsilon f_E}.$$

Without intermediate input linkages, the equilibrium mass of potential firms does not depend on any of the trade costs, be it variable (τ), or fixed (f_X), nor does it depend on domestic fixed costs f . As would be easy to verify by setting $N = 1$, this implies that the equilibrium n is unchanged under trade compared to complete autarky. This result is reminiscent of the original Krugman (1980) model, and has been found in a heterogeneous firms model by Arkolakis et al. (2008). Thus, the first striking feature of the endogenous entry model is that for the purposes of our comparative statics – changes in τ_{ij} and f_{ij} – intermediate input linkages are crucially important. Without them, keeping n exogenous and fixed as we do in the rest of the paper is without loss of generality, to a first order.

In the presence of intermediate input linkages, τ_{ij} and f_{ij} do affect the equilibrium mass of entrants. How does equilibrium n respond to changes in trade costs? First, the elasticity of n with respect to f is:

$$\frac{d \log n}{d \log f} = - \frac{\theta - (\varepsilon - 1)}{\theta(\varepsilon - 1)} \frac{1 - \beta}{\beta - \frac{1 - \beta}{\varepsilon - 1}} \frac{1}{1 + (N - 1) \tau^{-\theta} \left(\frac{f}{f_X} \right)^{\frac{\theta - (\varepsilon - 1)}{\varepsilon - 1}}}.$$

Thus, as the economy approaches Zipf’s Law ($\theta \rightarrow \varepsilon - 1$), the elasticity of the equilibrium mass of entrants with respect to the fixed cost of production goes to zero! It is easy to check that the same is true with respect to the fixed cost of exporting, f_X . What is the intuition for this result? A fall in the fixed cost of production (f) or export (f_X) will change the incentives to become a potential entrepreneur (i.e. to pay f_E for an $a(k)$ draw) insofar as it changes the expected profits from becoming one. Changes in f and f_X will affect the expected profits primarily on the extensive margin: with lower fixed costs, some additional marginal entrepreneurs can produce domestically or export. In Zipf’s world, as we had argued above, the marginal firms do not matter much: compared to the inframarginal firms, they are tiny, do not sell much, and have exceedingly small profits. The real payoff from becoming a potential entrepreneur comes from the possibility of “scoring big:” becoming one of the large firms at the top of the distribution. Changes in f and f_X do not affect that possibility at all. Thus, in Zipf’s world they have only a vanishingly small impact on the decision to become an entrepreneur. This is a striking result. Among other things, it says that as the distribution of firm size approaches Zipf’s Law, the assumption of a fixed mass of firms becomes a better and better approximation for the more general, endogenous entry model.

The elasticity of n with respect to variable trade costs is given by:

$$\frac{d \log n}{d \log \tau} = - \frac{1 - \beta}{\beta - \frac{1-\beta}{\varepsilon-1}} \frac{(N-1)\tau^{-\theta} \left(\frac{f}{f_X}\right)^{\frac{\theta-(\varepsilon-1)}{\varepsilon-1}}}{1 + (N-1)\tau^{-\theta} \left(\frac{f}{f_X}\right)^{\frac{\theta-(\varepsilon-1)}{\varepsilon-1}}}.$$

Unlike the elasticity with respect to f , it does not go to zero as we approach Zipf’s Law. Thus, the mass of potential entrants is more sensitive to variable costs than to fixed. This accords well with the intuition developed above: changes in n will happen in response to changes in expected profits. Since changes in τ affect expected profits throughout the distribution, including the very top, they have a much greater potential to induce net entry than changes in fixed costs. Above, we argued that in the presence of Zipf’s Law in firm size, variable costs have a much greater impact on welfare than fixed costs. The endogenous entry model provides yet another reason why this is the case: changes in variable costs will generate greater changes in the mass of firms than fixed costs. Thus, in this respect endogenizing n reinforces some of the main conclusions of the paper.

Having derived some analytical results using the symmetric model, we now present the solutions to the fully-fledged asymmetric model. In order to proceed, we must take a stand on the value of the “exploration” cost f_E . We cannot appeal to data in order to back out cross-country variation in f_E . This parameter measures the costs of finding out one’s productivity. Unlike starting an actual business, this process is unobservable, and/or takes place in the informal sector. Thus, we cannot find direct counterparts to it in databases such as Doing Business. Alternatively, f_E could

be calibrated to the number of firms present in each country, but currently that is not feasible as there are no reliable cross-country databases reporting the total numbers of firms in the economy. For these reasons we choose to set the same value of f_E across countries. We set f_E to be such that $n_{US} = 10,000,000$ in equilibrium (as argued above, this value is consistent with what is observed in the 2002 U.S. Economic Census). This is the value of n_{US} that is used in the baseline simulation, allowing for maximum comparability across the two sets of results. The resulting f_E is about 10 times larger than $f_{US,US}$, and about 1.7 times larger than the mean f_{ii} in our sample of countries.²²

The model is solved numerically as described above. Table 5 reports the results. Once again, the top panel presents results for the Zipf case ($\theta/(\varepsilon - 1) = 1.06$), and the bottom panel for the non-Zipf case ($\theta/(\varepsilon - 1) = 2$). We perform the same two comparative statics. The first is a complete harmonization of domestic and exporting fixed costs f_{ii} and f_{ij} to the U.S. level. Under Zipf's Law, the average change in welfare in our sample of countries is 4.03%, with a range from 0.04 to 9.10%. As expected, it is larger than in the model with exogenous n (which was 3.26%), but the difference is less than 25%. In the non-Zipf world, the mean change in welfare from the harmonization of fixed costs is 47.32%, compared to 40.87% with exogenous n . Once again, the two figures are close. Thus, endogenizing n preserves the essential quantitative results from this exercise. Though welfare gains are larger when the mass of potential entrepreneurs can adjust, the basic result that fixed costs have a small welfare impact in the Zipf's world, and an order of magnitude larger impact in the non-Zipf one is unchanged. Our analytical results using the symmetric case showed that as the distribution of firm size approaches Zipf's Law, the equilibrium n becomes insensitive to f . We can confirm this in the numerical exercise. A complete harmonization of fixed costs to the U.S. level leads to an average increase in n of less than 2%, and a maximum increase of 4.5% in this sample of 50 countries. Quantitatively, this is not an important force.

In the second quantitative exercise, we simulate a 10% reduction in τ_{ij} across the board. In devising a breakdown of the welfare impact into extensive and intensive margins, we must now take into account changes in n . If a movement in τ_{ij} induces more entrepreneurs to draw from the productivity distribution – n to change – we will observe new producers and new exporters throughout the size distribution. Thus, changes in n are part of the extensive margin. To take this

²²There remains a tension regarding whether the data in the Doing Business indicators reflect f_{ii} or f_E . In the final analysis, it is clear that convincingly sorting this out is not possible. However, the Doing Business indicators reflect the costs of starting an actual business in the formal sector, which, taken at face value, corresponds to f_{ii} in the model. Though in the real world starting a business undoubtedly involves some uncertainty, and thus has an element of discovering one's type, we choose not to model variation in f_E explicitly. However, to the extent that f_E is truly a cost of "exploration" by which one discovers one's abilities, we should expect it not to vary as much across countries. In addition, under this interpretation, it would be much more difficult to affect f_E through policies.

into account, equation (16) must be modified as follows:

$$\frac{\Delta W_i}{W_i} \approx \frac{1}{\beta} \sum_{j=1}^N \varphi_{ij} \left[\underbrace{-\frac{\Delta \left(\tau_{ij} \frac{c_j}{c_i} \right)}{\tau_{ij} \frac{c_j}{c_i}}}_{\text{Intensive Margin}} + \underbrace{\frac{\theta - (\varepsilon - 1)}{\theta (\varepsilon - 1)} \frac{\Delta G(a_{ij})}{G(a_{ij})} + \frac{1}{\varepsilon - 1} \frac{\Delta n_j}{n_j}}_{\text{Extensive Margin}} \right] \quad (19)$$

The last term under “Extensive Margin” in equation (19) reflects movements in the mass of firms. The results are presented in Table 5. In the Zipf’s simulation, the average welfare gain from a reduction in variable trade costs is 4.97%, only 15% higher than in the baseline model with exogenous mass of potential entrepreneurs. Because movements in n are recorded as the extensive margin, its importance is now greater: some 11% of the total welfare gain is due to the extensive margin of imports, and an additional 10% is due to the increase in the mass of domestic entrepreneurs. The intensive margin still accounts for a large majority – 78.4% – of the total welfare gain, but the results are not as drastic as they were in the baseline model. This is not surprising. Changes in n have a first-order impact on welfare, as they increase the number of varieties available to consumers, as well as for intermediate input use, one-for-one. Even if the changes in n are not very large, they still contribute more to the extensive margin. As argued above, this model should be seen as a metaphor for the long run. It implies, for instance, that in response to a reduction in τ_{ij} , there is an increase in the number of very large firms/exporters at the top of the firm size distribution. Unsurprisingly, this leads to a greater role of the extensive margin, as the very large firms drive most of the welfare results. However, entry of the large, extremely productive new firms is unlikely to happen in the short- to medium-run, and thus this aspect of the extensive margin is likely to be realized only over a longer period of time.

The lower panel presents the results of the same exercise in the non-Zipf model. As was the case in the baseline model, the total welfare gains in the non-Zipf case are an order of magnitude lower, at 0.29%. Similarly, the extensive margin plays a bigger role, accounting for about 25% of the total gains. The majority of the welfare gains is still due to the intensive margin, even with endogenous n and a less fat-tailed distribution of firm size.

To summarize, the magnitudes of the overall welfare gains in these experiments are surprisingly similar to the model in which n is exogenous. The basic story about the relative importance of fixed versus variable costs for welfare, and how that depends on the parameters governing the distribution of firm size, is unchanged. Because endogenous entry leads to the appearance of additional giant firms at the top of the firm size distribution, the importance of the extensive margin for welfare is somewhat greater. However, that result should be interpreted as representing the long run, as it is unlikely that the new large firms will emerge instantaneously.

5 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. Our conclusions about the welfare impact of higher entry barriers and the extensive margin of trade are very sensitive to 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. By contrast, gains from reductions in trade costs are much higher than in a model that does not exhibit Zipf's Law in firm size. Finally, the extensive margin accounts for only 5% of the overall gains from trade.

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 Proof of Proposition 1

Proof: The total variable profits from selling to country j from country i are:

$$\Pi_{ji}^V = \int_{J_{ji}} \frac{X_j}{\varepsilon P_j^{1-\varepsilon}} \left(\frac{\varepsilon}{\varepsilon-1} \tau_{ji} c_i a(k) \right)^{1-\varepsilon} dk.$$

The total sales from i to j are:

$$X_{ji} = \int_{J_{ji}} \frac{X_j}{P_j^{1-\varepsilon}} \left(\frac{\varepsilon}{\varepsilon-1} \tau_{ji} c_i a(k) \right)^{1-\varepsilon} dk.$$

Therefore, $\Pi_{ji}^V = \frac{X_{ji}}{\varepsilon}$.

The total fixed costs paid by firms in country i to enter market j are equal to $f_{ji} c_i n_i (b a_{ji})^\theta$.

We need to show that this quantity is also a constant multiple of X_{ji} . To do so, write

$$\begin{aligned} X_{ji} &= \frac{X_j}{P_j^{1-\varepsilon}} \left(\frac{\varepsilon}{\varepsilon-1} \tau_{ji} c_i \right)^{1-\varepsilon} \int_{J_{ji}} (a(k))^{1-\varepsilon} dk \\ &= \frac{X_j}{P_j^{1-\varepsilon}} \left(\frac{\varepsilon}{\varepsilon-1} \tau_{ji} c_i \right)^{1-\varepsilon} n_i \frac{b^\theta \theta}{\theta - (\varepsilon - 1)} a_{ji}^{\theta - (\varepsilon - 1)} \\ &= \frac{X_j}{P_j^{1-\varepsilon}} \left(\frac{\varepsilon}{\varepsilon-1} \tau_{ji} c_i \right)^{1-\varepsilon} n_i \frac{b^\theta \theta}{\theta - (\varepsilon - 1)} a_{ji}^\theta c_i f_{ji} \frac{\varepsilon P_j^{1-\varepsilon}}{X_j} \left(\frac{\varepsilon - 1}{\varepsilon} \right) \\ &= \frac{\theta}{\theta - (\varepsilon - 1)} \varepsilon n_i (b a_{ji})^\theta c_i f_{ji}. \end{aligned}$$

Therefore, the total fixed costs paid by firms in i to export to j are a constant multiple of X_{ji} :

$$n_i (b a_{ji})^\theta c_i f_{ji} = \frac{\theta - (\varepsilon - 1)}{\theta} \frac{X_{ji}}{\varepsilon}.$$

Therefore, the total profits from selling to j from country i are:

$$\begin{aligned} \Pi_{ji} &= \Pi_{ji}^V - \frac{\theta - (\varepsilon - 1)}{\theta} \frac{X_{ji}}{\varepsilon} \\ &= \frac{X_{ji}}{\varepsilon} \left(1 - \frac{\theta - (\varepsilon - 1)}{\theta} \right) \\ &= X_{ji} \frac{(\varepsilon - 1)}{\varepsilon \theta}. \end{aligned}$$

This means that the total profits from selling to all countries equal:

$$\Pi_i = \sum_{j=1}^N \Pi_{ji} = \frac{(\varepsilon - 1)}{\varepsilon \theta} \sum_{j=1}^N X_{ji}.$$

Since in equilibrium total income equals total expenditure in each country, $X_i = \sum_{j=1}^N X_{ji}$, leading to the result that $\Pi_i = \frac{(\varepsilon - 1)}{\varepsilon \theta} X_i$. ■

Appendix B Model with Market Penetration Costs

A recent contribution by Arkolakis (2008) emphasizes that the model with simple fixed costs of accessing markets is too stark. Instead, Arkolakis (2008) 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 (2008) and Eaton et al. (2008) 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. (2008). 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 marginal cost $a(k)$ 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 marginal cost $a(k)$ is given by:

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

Finally, the firm will only enter market i if at zero market penetration, profits are increasing in s : $\frac{\partial \pi_i(k)}{\partial s} \Big|_{s=0} > 0$. 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 equation (4). That expression can be combined with equation (B.1) to write the sales of a firm with marginal cost $a(k)$ from country j to country i as:

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

As first observed by Arkolakis (2008), the baseline model with simple fixed costs provides the best approximation to the sales of the largest firms: as the marginal cost $a(k)$ 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:

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

Comparing equations (6) and (B.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 Proposition 1 still holds, and that the wages are still determined by equation (9). 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 (9) and (B.2). We now state the main result of this Appendix.

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, (6) and (9). Then, the vectors*

$$[w_1^{mp}, \dots, w_N^{mp}] = [w_1, \dots, w_N] \tag{B.3}$$

and

$$[P_1^{mp}, \dots, P_N^{mp}] = \delta [P_1, \dots, P_N] \tag{B.4}$$

are a solution to the system of equations (B.2) and (9) that define the equilibrium in the market penetration costs model.

Proof: It is immediate from examining (9) that the vector $[w_1, \dots, w_N]$ that solves (9) is the same under $[P_1, \dots, P_N]$ and $[P_1^{mp}, \dots, P_N^{mp}]$ when the latter is defined by (B.4), since δ cancels out from

the numerator and the denominator. We now show that as long as (B.3) is satisfied, (B.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 $\forall i$, (B.2) holds. We check directly whether the vector $\delta [P_1, \dots, P_N]$ satisfies that condition:

$$P_i^{mp} = \delta P_i = \frac{1}{b} \left[\frac{\theta}{\theta - (\varepsilon - 1)} - \frac{\theta}{\theta - (\varepsilon - 1)(1 - \lambda)} \right]^{-\frac{1}{\theta}} \frac{\varepsilon}{\varepsilon - 1} \left(\frac{w_i L_i}{\varepsilon \beta \left(1 - \frac{\varepsilon - 1}{\theta \varepsilon}\right)} \right)^{-\frac{\theta - (\varepsilon - 1)}{\theta(\varepsilon - 1)}} \times$$

$$\left(\sum_{j=1}^N n_j \left(\frac{1}{\tau_{ij} w_j^\beta (\delta P_j)^{1 - \beta}} \right)^\theta \left(\frac{1}{w_j^\beta (\delta P_j)^{1 - \beta} f_{ij}} \right)^{\frac{\theta - (\varepsilon - 1)}{\varepsilon - 1}} \right)^{-\frac{1}{\theta}}. \quad (\text{B.5})$$

After rearranging it becomes:

$$\delta^{\beta - (1 - \beta) \frac{\theta - (\varepsilon - 1)}{\theta(\varepsilon - 1)}} P_i = \frac{1}{b} \left[\frac{\theta}{\theta - (\varepsilon - 1)} - \frac{\theta}{\theta - (\varepsilon - 1)(1 - \lambda)} \right]^{-\frac{1}{\theta}} \frac{\varepsilon}{\varepsilon - 1} \left(\frac{w_i L_i}{\varepsilon \beta \left(1 - \frac{\varepsilon - 1}{\theta \varepsilon}\right)} \right)^{-\frac{\theta - (\varepsilon - 1)}{\theta(\varepsilon - 1)}} \times$$

$$\left(\sum_{j=1}^N n_j \left(\frac{1}{\tau_{ij} w_j^\beta P_j^{1 - \beta}} \right)^\theta \left(\frac{1}{w_j^\beta P_j^{1 - \beta} f_{ij}} \right)^{\frac{\theta - (\varepsilon - 1)}{\varepsilon - 1}} \right)^{-\frac{1}{\theta}},$$

which is the same as (6) 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 (6), we have shown that $\delta [P_1, \dots, P_N]$ satisfies (B.5), which completes the proof. ■

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.

²³This is the value of λ preferred by Arkolakis (2008). Using Simulated Method of Moments, Eaton et al. (2008) 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 $\lambda = \infty$.

Appendix C 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.

ORBIS is a multi-country database published by Bureau van Dijk that contains information on more than 50 million companies worldwide.²⁴ 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.²⁵ 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 (10), which in natural logs becomes:

$$\log(\Pr(x > s)) = \log(c) - \zeta \log(s). \tag{C.1}$$

For a grid of values of sales s , the estimated probability $\Pr(x > s)$ is simply the number of firms

²⁴The 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.

²⁵This 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.

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 (10) has a probability density function

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

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 $-(\zeta + 1)$.²⁶

Table A1 reports the results. The left panel reports estimates of equation (C.1), the right panel, equation (C.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

²⁶Finally, 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 (2009), 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.

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.

References

- Acemoglu, Daron, Pol Antràs, and Elhanan Helpman, “Contracts and Technology Adoption,” *American Economic Review*, June 2007, 97 (3), 916–943.
- Alvarez, Fernando and Robert E., Jr. Lucas, “General Equilibrium Analysis of the Eaton-Kortum Model of International Trade,” *Journal of Monetary Economics*, September 2007, 54 (6), 1726–1768.
- Amiti, Mary and Jozef Konings, “Trade Liberalization, Intermediate Inputs, and Productivity,” *American Economic Review*, December 2007, 97 (5), 1611–1638.
- Anderson, James and Eric van Wincoop, “Trade Costs,” *Journal of Economic Literature*, 2004, 42 (3), 691–751.
- Arkolakis, Costas, “Market Penetration Costs and the New Consumers Margin in International Trade,” 2008. Mimeo, Yale University.
- , Arnaud Costinot, and Andrés Rodríguez-Clare, “New Trade Models, Same Old Gains?,” January 2010. mimeo, Yale University, MIT, and Penn State University.
- , Svetlana Demidova, Peter J. Klenow, and Andrés Rodríguez-Clare, “Endogenous Variety and the Gains from Trade,” *American Economic Review Papers and Proceedings*, May 2008, 98 (2), 444–450.
- Atkeson, Andrew and Ariel Burstein, “Innovation, Firm Dynamics, and International Trade,” *Journal of Political Economy*, June 2010, 118 (3), 433–484.
- Axtell, Robert L., “Zipf Distribution of U.S. Firm Sizes,” *Science*, September 2001, 293 (5536), 1818–1820.
- Bekkers, Eddy and Joseph Francois, “Heterogeneous Firms, the Structure of Industry, and Trade under Oligopoly,” August 2008. Working Paper 0811, Department of Economics, Johannes Kepler University of Linz.
- Bernard, Andrew, J. Bradford Jensen, Stephen Redding, and Peter Schott, “Firms in International Trade,” *Journal of Economic Perspectives*, 2007, 21, 105–130.
- Broda, Christian and David Weinstein, “Globalization and the Gains from Variety,” *Quarterly Journal of Economics*, May 2006, 121 (2), 541–85.
- Chaney, Thomas, “Distorted Gravity: The Intensive and Extensive Margins of International Trade,” *American Economic Review*, September 2008, 98 (4), 1707–1721.
- Costinot, Arnaud, “On the Origins of Comparative Advantage,” *Journal of International Economics*, April 2009, 77 (2), 255–264.
- Cowan, Kevin and Alejandro Neut, “Intermediate Goods, Institutions, and Output Per Worker,” 2007. Central Bank of Chile Working Paper No. 420.
- di Giovanni, Julian and Andrei A. Levchenko, “International Trade and Aggregate Fluctuations in Granular Economies,” February 2009. RSIE Discussion Paper 585.
- , —, and Romain Rancière, “Power Laws in Firm Size and Openness to Trade: Measurement and Implications,” March 2010. CEPR Discussion Paper 7773.
- Djankov, Simeon, Rafael La Porta, Florencio Lopez-de-Silanes, and Andrei Shleifer, “The Regulation of Entry,” *Quarterly Journal of Economics*, 2002, 117 (1), 1–37.

- Easterly, William, Ariell Reshef, and Julia Schwenkenberg, “The Power of Exports,” October 2009. World Bank Policy Research Working Paper 5081.
- Eaton, Jonathan and Samuel Kortum, “Technology, Geography, and Trade,” *Econometrica*, September 2002, 70 (5), 1741–1779.
- and —, “Technology in the Global Economy: A Framework for Quantitative Analysis,” 2005. Book manuscript, New York University and University of Chicago.
- , —, and Francis Kramarz, “An Anatomy of International Trade: Evidence From French Firms,” December 2008. NBER Working Paper No. 14610.
- Feenstra, Robert C., “Measuring the Gains From Trade Under Monopolistic Competition,” *Canadian Journal of Economics*, February 2010, 43 (1), 1–28.
- Fujiwara, Yoshi, Hideaki Aoyama, Corrado Di Guilmi, Wataru Souma, and Mauro Gallegati, “Gibrat and Pareto-Zipf Revisited with European firms,” *Physica A: Statistical Mechanics and its Applications*, 2004, 344 (1-2), 112–116.
- Gabaix, Xavier, “Power Laws in Economics and Finance,” *Annual Reviews of Economics*, 2009, 1, 255–93.
- and Rustam Ibragimov, “Rank-1/2: A Simple Way to Improve the OLS Estimation of Tail Exponents,” May 2009. Forthcoming, *Journal of Business Economics and Statistics*.
- Goldberg, Pinelopi, Amit Khandelwal, Nina Pavcnik, and Petia Topalova, “Trade Liberalization and New Imported Inputs,” *American Economic Review Papers and Proceedings*, May 2009, 99 (2), 494–500.
- , —, —, and —, “Imported Intermediate Inputs and Domestic Product Growth: Evidence from India,” January 2010. Forthcoming, *Quarterly Journal of Economics*.
- Halpern, László, Miklós Koren, and Adam Szeidl, “Imported Inputs and Productivity,” April 2009. mimeo, Hungarian Academy of Sciences, Central European University, and UC Berkeley.
- Helpman, Elhanan, Marc Melitz, and Yona Rubinstein, “Estimating Trade Flows: Trading Partners and Trading Volumes,” *Quarterly Journal of Economics*, May 2008, 123 (2), 441–487.
- Hinloopen, Jeroen and Charles van Marrewijk, “Comparative advantage, the rank-size rule, and Zipf’s law,” 2006. Tinbergen Institute Discussion Paper 06-100/1.
- International Monetary Fund, “Directions of Trade Statistics,” 2007.
- Jones, Charles, “Input-Output Multipliers, General Purpose Technologies, and Economic Development,” September 2007. Mimeo, University of California-Berkeley.
- , “Intermediate Goods and Weak Links: A Theory of Economic Development,” March 2008. NBER Working Paper No. 13834.
- Kasahara, Hiroyuki and Joel Rodrigue, “Does the use of imported intermediates increase productivity? Plant-level evidence,” *Journal of Development Economics*, 2008, 87, 106–118.
- Krugman, Paul, “Scale Economies, Product Differentiation, and the Pattern of Trade,” *American Economic Review*, 1980, 70 (5), 950–959.
- Luong, Tuan A., “The Impact of Input and Output Tariffs on Firms’ Productivity: Theory and Evidence,” 2008. mimeo, Princeton University.
- Luttmer, Erzo, “Selection, Growth, and the Size Distribution of Firms,” *Quarterly Journal of Economics*, 2007, 122, 1103–1144.

- Melitz, Marc J., “The Impact of Trade on Intra-Industry Reallocations and Aggregate Industry Productivity,” *Econometrica*, November 2003, *71* (6), 1695–1725.
- Neary, J. Peter, “Two and a Half Theories of Trade,” *The World Economy*, January 2010, *33* (1), 1–19.
- Okuyama, K., M. Takayasu, and H. Takayasu, “Zipf’s Law in Income Distribution of Companies,” *Physica A: Statistical Mechanics and its Applications*, 1999, *269* (1), 125–131.
- The World Bank, “Doing Business Indicators,” 2007.
- , “World Development Indicators,” 2007.
- Waugh, Michael, “International Trade and Income Differences,” 2009. Forthcoming, *American Economic Review*.

Table 1. Top 49 Countries and the Rest of the World in Terms of 2004 GDP

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. Source: The World Bank (2007b).

Table 2. Bilateral Trade Shares: Data and Model Predictions for the 50-Country Sample

	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. Source: International Monetary Fund (2007).

Table 3. Welfare Gains

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

Notes: This table reports the welfare increase, 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.

Table 4. Welfare Gains when Varying Share of Intermediate Goods in Final Production

Counterfactual	Total change in welfare			
	$\beta = 0.33$	$\beta = 0.5$ (baseline)	$\beta = 0.67$	$\beta = 1$
Complete harmonization of entry costs	5.03 (0.06, 11.37)	3.26 (0.03, 7.32)	2.41 (0.02, 5.39)	1.60 (0.00, 3.58)
10% reduction in τ	4.79 (0.22, 8.87)	4.33 (0.28, 8.26)	4.07 (0.32, 7.85)	3.80 (0.37, 7.37)


Notes: This table reports the welfare increase, in percentage points, due to each counterfactual experiment. The numbers in parentheses indicate the range across the 50 countries in the sample. $1 - \beta$ equals the share of intermediate goods in final production. These counterfactuals are done assuming $\frac{\theta}{\epsilon-1} = 1.06$ (Zipf's World).

Table 5. Welfare Gains Under Endogenous Mass of Potential Entrepreneurs

$\frac{\theta}{\epsilon-1} = 1.06$ (Zipf's World)				
Counterfactual	Total change in welfare	Intensive margin	Extensive margin	
Complete harmonization of entry costs	4.03 (0.04, 9.10)		4.03 (0.04, 9.10)	
10% reduction in τ	4.97 (0.54, 22.33)	3.83 (0.41, 13.03)	Foreign 0.48 (0.06, 1.42)	Domestic 0.66 (-0.91, 8.82)
		0.784	0.114	0.103
$\frac{\theta}{\epsilon-1} = 2$				
Counterfactual	Total change in welfare	Intensive margin	Extensive margin	
Complete harmonization of entry costs	47.32 (0.00, 123.93)		47.32 (0.00, 123.93)	
10% reduction in τ	0.29 (0.01, 1.73)	0.22 (0.01, 1.37)	Foreign 0.03 (0.00, 0.17)	Domestic 0.04 (0.00, 0.20)
		0.745	0.118	0.137


Notes: This table reports the welfare increase, in percentage points, due to each counterfactual experiment, using the model with endogenous mass of potential entrepreneurs described in section 4. 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.

Figure 1. Estimated Power Law in Firm Size in the U.S. (Axtell, 2001).




Notes: Reproduced from Axtell (2001). This figure depicts the power law in firm size in the U.S.: it plots the log frequency of the firms against log of firm size, measured by the number of employees. The solid line is the OLS regression fit through the data. The estimated slope coefficient is -2.059 (s.e. 0.054), which implies $\zeta = 1.059$. The adjusted R^2 is 0.992. Similar relationships are also reported for sales.

Figure 2. Bilateral Trade Shares: Data and Model Predictions




Notes: This figure reports the scatterplot of domestic output (π_{ii}) and bilateral trade (π_{ij}), both as a share of domestic absorption. The values implied by the model are on the horizontal axis. Actual values are on the vertical axis. Solid dots represent observations of π_{ii} , while hollow dots represent bilateral trade observations (π_{ij}). The line through the data is the 45-degree line.

Figure 3. Trade Openness: Data and Model Predictions



Notes: This figure reports total imports as a share of domestic absorption ($1 - \pi_{ii}$). The values implied by the model are on the horizontal axis. Actual values are on the vertical axis. The line through the data is the 45-degree line.

Figure 4. 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.

Table A1. 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 & 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: ** – significant at the 1% level. 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 R^2 , 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.