

International Trade: Linking Micro and Macro¹

Jonathan Eaton, Samuel Kortum, and Sebastian Sotelo

¹An earlier draft of this paper was presented at the Econometric Society World Congress, Paired Invited Session on Trade and Firm Dynamics, in Shanghai in August 2010. It was published in the volume edited by Daron Acemoglu, Manuel Arellano, and Eddie Dekel, *Advances in Economics and Econometrics, Tenth World Congress, Volume II, Applied Economics*, Cambridge University Press, 2013, 329-370. We have benefited from the valuable comments of Daron Acemoglu, Costas Arkolakis, Thomas Chaney, Peter Egger, Elhanan Helpman, Stephen Redding, Joao Santos Silva, Silvana Tenreyro, and Alain Trognon. Kelsey Moser provided excellent research assistance. Alvin Chung and Michelle Tong carefully revised this draft to be consistent with the published version. We gratefully acknowledge the support of the National Science Foundation under grant numbers SES-0339085 and SES-0820338.

Abstract

A recent literature has introduced heterogeneous firms into models of international trade. This literature has adopted the convention of treating individual firms as points on a continuum. While the continuum offers many advantages this convenience comes at some cost: (1) Shocks to individual firms can never have an aggregate effect. (2) It is hard to reconcile the small (sometimes zero) number of firms engaged in selling from one country to another with a continuum. (3) For such models to deliver finite solutions for aggregates, such as the price index, requires restrictions on parameter values that may not hold in the data. We show how a standard heterogeneous-firm trade model can be amended to allow for only an integer number of firms. The model overcomes the deficiencies of the continuum model enumerated above. Taking the model to aggregate data on bilateral trade in manufactures among 92 countries and to firm-level export data for a much narrower sample shows that it accounts for both the large share of a small number of firms in sales around the world and for zeros in bilateral trade data while maintaining the good fit of the standard gravity equation among country pairs with thick trade volumes. Randomness at the firm level adds substantially to aggregate variability.

1 Introduction

The field of international trade has advanced in the past decade through a healthy exchange between new observations on firms in export markets and new theories that have introduced producer heterogeneity into trade models. As a result, we now have general equilibrium theories of trade that are consistent with various dimensions of both the aggregate and the firm-level data. Furthermore, we have a much better sense of the magnitudes of key parameters underlying these theories.¹

This flurry of activity at the firm level has left the core aggregate relationships among trade, factor costs, and welfare largely untouched, however. Although we now have much better microfoundations for aggregate trade models, their predictions are much like those of the Armington model – for years a workhorse of quantitative international trade. Arkolakis, Costinot, and Rodríguez-Clare (2012) emphasize this (lack of) implication of the recent literature for aggregate trade.

We argue that a primary reason why models of heterogeneous producers deliver so little in the way of modification of how we think about aggregates is the device – initiated in the trade literature by Dornbusch, Fischer, and Samuelson (1977) – of treating the set of products as a continuum. The heterogeneous-firm literature embraced this approach, applying it to individual producers.

Treating individual producers as points on a continuum has a number of extremely convenient implications that the field (including work by two of the authors of this chapter) has exploited relentlessly. With a continuum, each producer has measure zero, so it has no effect

¹Bernard, Jensen, Redding, and Schott (2007) and, more recently, Redding (2011) provide surveys.

on aggregates. Invoking the law of large numbers, we are free to model what goes on at the aggregate level as driven by the parameters (which may be small in number) governing the distributions of the outcomes affecting individual units but not on the realizations of those outcomes themselves. With the right distributional assumptions about the processes underlying the outcomes of individual firms, one can readily integrate over them to get simple and familiar aggregate relationships. The continuum thus provides a wall of separation between smooth aggregate relationships and potentially jagged heterogeneity underneath, allowing us to deal with each realm in isolation.

Of course, the number of producers or products is not literally a continuum. Yet, are they so numerous that treating them as a continuum is an innocuous simplification? For many purposes, it undoubtedly is, but there are questions for which the continuum can lead us astray. First, some individual producers may indeed loom so large that their own individual fates have implications for the economy as a whole. Second, under the continuum assumption, anything that can happen (within the support of what is modeled as possible), will happen. An implication, for example, is that if we observe no exports from one country to another, as we often do, then exporting was impossible – not that it just so happened that no firm found exporting worthwhile. A third limitation is that obtaining well-behaved integrals across the continuum requires restrictions on distributional parameters that prevent the size distribution of firms from becoming too skewed. The skewness observed in the data is uncomfortably close to the limits imposed by these parameter restrictions.

Here, we explore the implications of having only a finite number (sometimes zero) of firms exporting. We develop a variant of a standard model of firm participation in exporting in

which the number of firms is an integer. The model confronts each of the issues raised in the previous paragraph: (1) Under parameter values consistent with the data, randomness in the situations of individual firms translates into substantial randomness in aggregates such as the price index; (2) the model predicts zero trade flows, with a frequency similar to what we see in the data, simply because no firm happened to be efficient enough – not because it was impossible for any firm whatever its luck of the draw; and (3) Our finite-firm model can easily deal with parameter values consistent with any degree of skewness in the firm-size distribution.

We use our model to perform a series of quantitative exercises. We first derive the model's implication for the specification of a gravity equation. Estimating this equation with aggregate bilateral trade and production data delivers estimates of the parameters governing the probability of a firm from each source entering each destination. We then take the model to firm-level data to learn about the cost of entry in different markets (and the other remaining parameters). With a fully parameterized version of the model in hand, we conduct two experiments. The first addresses the zeros issue. A simulation of 10 percent lower trade barriers worldwide introduces 206 new bilateral trading relationships (although the amount of trade involved is miniscule). The second gets to the heart of the law-of-large-numbers question. We find that resampling repeatedly the efficiencies of individual firms around the world generates a variance (of percentage deviations) in the manufacturing price index for the United States of 14 and for Denmark of 24 – far from the zero implied by a continuum model.

This chapter addresses a particular situation in which an aggregate relationship (here, a bilateral trade flow) is the outcome of decisions by heterogeneous individual agents (here, of

firms on whether and how much to export to a destination). However, the issues we raise apply to any aggregate variable (e.g., consumption or investment) whose magnitude is the summation of what a diverse set of individuals choose to do, which may include nothing.

The chapter proceeds as follows. We begin with a review of related literature followed by an overview of key features of the trade data. Next, we introduce our finite-firm model, which underlies the estimation approach that follows. We then confront the model with the data introduced in the previous section.

2 Related Literature

This chapter relates closely to another literature that emphasized the importance of individual firms in aggregate models. Gabaix (2011) used such a structure to explain aggregate fluctuations due to shocks to very large firms in the economy. This analysis is extended to international trade by Canals, Gabaix, Vilarrubia, and Weinstein (2007) and di Giovanni and Levchenko (2009), again highlighting the role of very large firms in generating aggregate fluctuations.

The literature on zeros in the bilateral-trade data includes Eaton and Tamura (1994); Santos Silva and Tenreyro (2006); Armenter and Koren (2008), Helpman, Melitz, and Rubinstein (2008), Martin and Pham (2008); and Baldwin and Harrigan (2011). Our underlying model of trade relates closely to Helpman, Melitz, and Rubinstein (2008), but rather than obtaining zeros by truncating a continuous Pareto distribution of efficiencies from above, zeros arise in our model because – as in reality – the number of firms is finite. Like us, Armenter and Koren (2008) assume a finite number of firms, stressing, as we do, the importance of the sparsity of

the trade data in explaining zeros. Theirs, however, was a purely probabilistic rather than an economic model.²

Our work also touches on Balistreri, Hillberry, and Rutherford (2011). Their paper discussed both estimation and general equilibrium simulation of a heterogeneous-firm model similar to the one we consider here. However, it does not draw out the implications of a finite number of firms, which is our main contribution.

Finally, a recent paper by Armenter and Koren (2012) took an approach complementary to ours in this chapter. In their framework individual buyers, rather than sellers, are finite in number, generating the possibility of zero sales.

3 The Data

We use macrodata and microdata on bilateral trade in manufactures among 92 countries. The macrodata are aggregate bilateral trade flows (in U.S. dollars) of manufactures X_{ni} from source country i to destination country n in 1992 (Feenstra, Lipsey, and Bowen 1997) and aggregate manufacturing production (described in Eaton, Kortum, and Kramarz 2011). The microdata are firm-level exports to destination n for four exporting countries i . The efforts of many researchers, who exploited customs records, are making such data more widely available. We were generously provided microdata for exports from Brazil, France, Denmark, and Uruguay.³

²Mariscal (2010) showed that the Armenter and Koren (2008) approach also goes a long way in explaining multinational expansion patterns.

³The French data for manufacturing firms in 1992 are from Eaton, Kortum, and Kramarz (2011). The Danish data for all exporting firms in 1993 are from Pedersen (2009). The Brazilian data for manufactured exports in 1992 are from Arkolakis and Muendler (2010). The Uruguayan data for 1992 and 1993 were compiled

The microdata allow us to measure the number K_{ni} of firms from i selling in n as well as mean sales per firm \bar{X}_{ni} when K_{ni} is reported as positive.⁴ In merging the data, we chose our 92 countries for the macrolevel analysis in order to have observations at the firm level from at least two of our four sources.

Figure 1. Micro and macro bilateral trade

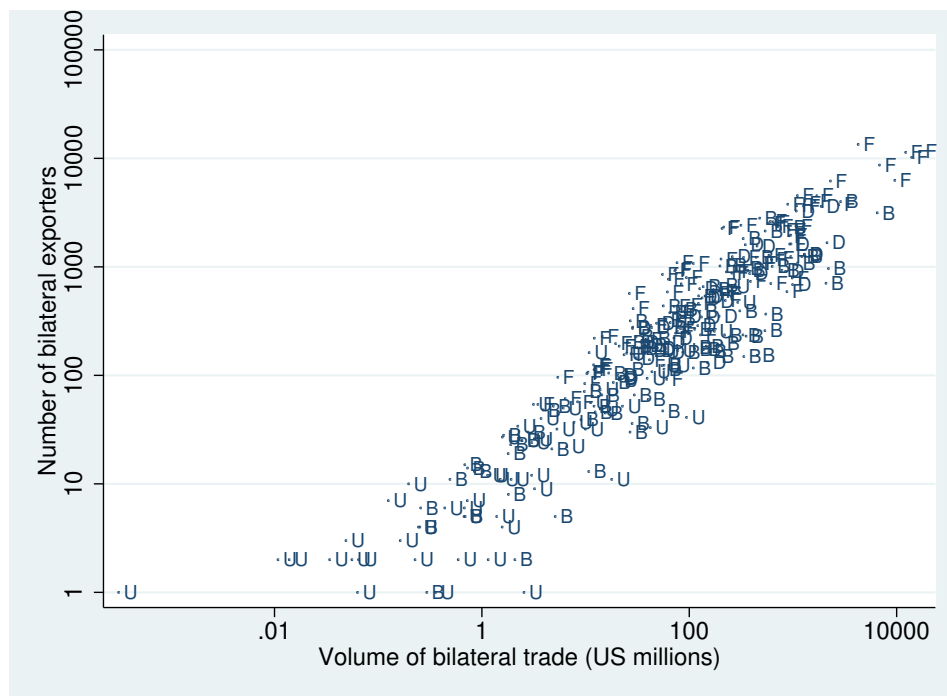


Table 1 lists our 92 countries and each country’s total exports to and imports from the other 91. The last two columns display the number of destinations for each country’s exports and the number of sources for its imports (each out of a maximum of 91). It is not surprising that a country trades with a greater number of other countries when it trades more in total.
by Sampognaro and used previously in Eaton Kortum, and Kramarz (2011). (Figure 1 includes only the 1992 data for Uruguay).

⁴We cannot always determine in the export microdata if the lack of any reported exporter to a particular destination means zero exports there or that the particular destination was not in the dataset. Hence, our approach, which exploits the micro data only when $K_{ni} > 0$, leaves open the interpretation.

Table 1. Trade in manufactures

Country	Value of trade (Million USD)		Trade partners in sample (out of a total of 91)	
	Total exports	Total imports	No. destinations	No. sources
1 Algeria	262.02	6230.41	34	47
2 Angola	48.04	2149.29	20	38
3 Argentina	7111.71	12284.37	83	64
4 Australia	15566.94	30132.72	86	72
5 Austria	22085.23	21720.69	91	85
6 Bangladesh	1446.20	1188.85	72	48
7 Benin	15.96	448.10	17	36
8 Bolivia	305.03	1111.53	41	54
9 Brazil	27212.22	13626.56	91	70
10 Bulgaria	1341.33	1283.07	60	53
11 Burkina Faso	26.11	232.03	21	34
12 Burundi	5.08	88.01	21	35
13 Cameroon	390.73	877.53	38	45
14 Canada	106421.63	106100.68	91	84
15 Central African Republic	17.02	87.79	17	31
16 Chad	2.69	110.86	19	27
17 Chile	7067.69	7613.92	75	68
18 China	31071.30	39042.04	91	74
19 Colombia	2557.45	6204.99	70	69
20 Costa Rica	639.36	2363.57	47	55
21 Côte d'Ivoire	675.01	1457.22	45	47
22 Denmark	23624.13	19651.31	91	83
23 Dominican Republic	2294.14	2882.82	42	49
24 Ecuador	876.57	2565.07	43	55
25 Egypt	995.60	6324.02	76	65
26 El Salvador	326.56	1291.13	42	52
27 Ethiopia	31.62	535.79	18	49
28 Finland	17197.93	11243.78	91	71
29 France	141492.66	130104.82	91	91
30 Ghana	723.87	1184.87	49	67
31 Greece	4535.57	13795.85	85	81
32 Guatemala	514.37	2201.65	40	53
33 Honduras	122.73	910.98	27	52
34 Hungary	4567.63	5024.21	88	67
35 India	12955.11	8470.82	91	73
36 Indonesia	16126.92	18685.77	84	72
37 Iran	640.27	12368.96	51	48
38 Ireland	21663.64	17493.05	91	77
39 Israel	9252.63	11270.82	64	59
40 Italy	117066.40	93372.11	91	90
41 Jamaica	1071.58	1172.92	45	46
42 Japan	273219.72	121513.38	91	90
43 Jordan	353.57	1974.08	52	51
44 Kenya	327.22	1031.39	56	69
45 Korea	59662.13	47027.97	91	75
46 Kuwait	274.11	4757.93	44	51

continued next page

Country	Value of trade (Million USD)		Trade partners in sample (out of a total of 91)	
	Total exports	Total imports	No. destinations	No. sources
47 Madagascar	74.45	289.07	28	47
48 Malawi	33.71	448.13	28	43
49 Malaysia	21881.53	25116.63	86	72
50 Mali	28.84	270.31	21	38
51 Mauritania	215.04	363.36	23	36
52 Mauritius	749.66	1122.83	55	59
53 Mexico	36481.61	56450.13	77	69
54 Morocco	2723.01	4864.38	73	67
55 Mozambique	129.24	702.29	33	38
56 Nepal	124.93	290.90	26	36
57 Netherlands	63075.79	63236.59	91	91
58 New Zealand	7167.16	6989.50	77	60
59 Nigeria	261.50	5915.16	43	56
60 Norway	14116.79	18442.85	91	71
61 Oman	440.42	2292.31	45	52
62 Pakistan	4808.01	5441.02	86	63
63 Panama	320.01	7850.87	43	56
64 Paraguay	295.52	1532.92	43	47
65 Peru	2422.71	2731.93	63	57
66 Philippines	4675.29	8433.17	69	60
67 Portugal	12726.92	19680.55	90	86
68 Romania	2182.08	2094.73	83	55
69 Rwanda	5.51	114.88	17	33
70 Saudi Arabia	3088.77	27632.93	55	61
71 Senegal	373.17	804.17	32	39
72 South Africa	6671.92	10369.34	88	82
73 Spain	46963.64	63036.14	91	90
74 Sri Lanka	1476.41	2182.93	59	54
75 Sweden	40954.33	29656.78	91	83
76 Switzerland	44029.96	36146.51	91	87
77 Syrian Arab Republic	141.13	2141.40	41	48
78 Taiwan	65581.95	50130.16	64	58
79 Tanzania, United Rep. of	72.00	842.68	40	46
80 Thailand	21645.97	27416.26	91	80
81 Togo	20.69	489.79	28	43
82 Trinidad and Tobago	481.03	1068.05	46	52
83 Tunisia	2230.96	4130.15	56	54
84 Turkey	6824.79	12386.31	88	67
85 Uganda	23.50	266.95	31	41
86 United Kingdom	128688.75	137566.47	91	91
87 United States of America	359292.84	395010.78	91	91
88 Uruguay	1324.24	1672.66	56	56
89 Venezuela	2819.75	11546.50	57	60
90 Viet Nam	833.21	1695.58	53	37
91 Zambia	912.95	768.91	36	43
92 Zimbabwe	555.31	1286.70	52	56
Average			59.6	59.6
Variance			652.5	283.6

Nonetheless, the number of zero trade links is large, comprising more than one third of the 8,372 bilateral observations.

The average number of positive bilateral trade flows per country, either as an exporter or an importer, is 59.6. The variance of the number of export destinations, however, is 652.5 while the variance of the number of import sources is only 283.6. As we discuss in Section 5.2, our analysis provides an explanation for the large deviation between the variances.

For country pairs for which $K_{ni} > 0$, Figure 1 plots K_{ni} against X_{ni} on log scales, with source countries labeled by the first letter of the country name. The data cluster around a positively-sloped line, with no apparent differences across the four source countries.⁵

Where exporting does occur, how important are very large firms? Using data on individual French firms in 1986, we order exporters according to their total exports.⁶ Table 2 reports the contribution to total French exports of the top 10, 100, 1000, and 10,000 largest French exporters. The 100 largest exporters account for nearly half of total exports – nearly half of which is due to only the top 10 exporters. These same firms are the main contributors to French exports to individual destinations as well.

⁵The regression slope is 0.71 (standard error 0.17), slightly higher than the 0.65 Eaton, Kortum, and Kramarz (hereafter EKK, 2011) found for French firms in 1986. Allowing for source-specific intercepts (which differ significantly from a common intercept), we cannot reject the hypothesis of a common slope of 0.62 (with a standard error of 0.03).

⁶The sample consists of 34,035 French exporters, described in EKK (2011). We consider exports to the 112 destinations reported in that paper. Canals, Gabaix, Vilarrubia, and Weinstein (2007) reported results for Japanese exporters similar to those for France in Table 2. Freund and Pierola (2012) find, among 33 developing countries, that the top 1 percent of exporting firms typically account for more than half of export revenues. The number of such export "superstars" averages 54, and for 8 countries there are fewer than 10.

Table 2. Share of largest French exporters

	French exports to:			Standard Deviation of Shares across destinations
	Everywhere	United States	Denmark	
Top 10	23.6	22.4	22.2	18.9
Top 100	47.9	54.6	52.2	16.8
Top 1,000	80.5	84.8	83.5	12.4
Top 10,000	98.9	99.3	99.2	1.2

For example, Table 2 shows that the top 100 French exporters account for more than half of French exports to both the United States and Denmark. Although the United States and Denmark are typical, the last column of the table shows that these statistics vary considerably among countries.

4 A Finite-Firm Model of Trade

Our framework relates closely to work on trade with heterogeneous firms in Melitz (2003), Chaney (2008), and EKK (2011). The key difference is that we treat the range of potential technologies for these firms not as a continuum but rather as an integer. Some results from the existing work survive, others do not. We show the difficulties introduced by dropping the continuum and an approach to overcoming them. To highlight the similarities and differences, we report established results from the continuum case in parallel with our finite-firm variant.

4.1 Technology

As in the recent literature (but also close to the basic Ricardian model of international trade), our basic unit of analysis is a technology for producing a unique good. We represent technology

by the quantity Z of output produced by a unit of labor.⁷ We refer to Z as the efficiency of the technology. We call the owner of this technology a firm, even though – in equilibrium – many of these “firms” will be inactive.

A standard building block in modeling firm heterogeneity is the Pareto distribution of firm efficiency. We follow this tradition in assuming that for any given firm, conditional on its efficiency exceeding some threshold $\underline{z} > 0$, its efficiency is the realization of a random variable Z drawn from a Pareto distribution with parameter $\theta > 0$, so that:

$$\Pr[Z > z | Z \geq \underline{z}] = (z/\underline{z})^{-\theta}. \quad (1)$$

The Pareto distribution has a number of properties that make it analytically very tractable.⁸ Moreover, for reasons that were discussed by Simon and Bonini (1958), Gabaix (1999), and Luttmer (2011), the relevant data (e.g., firm-size distributions) often exhibit Pareto properties, at least in the upper tail.

⁷A higher Z can mean: (1) more of a product, (2) the same amount of a better product, or (3) any combination of the first two that renders the output of the good produced by a worker more valuable. For the results here the different interpretations have isomorphic implications. Here, “labor” can be interpreted to mean an arbitrary bundle of factors and the “wage” to mean the price of that bundle, common across all goods j . Eaton and Kortum (2002) and EKK (2011) made the input bundle a Cobb-Douglas combination of labor and intermediates, an extension we do not pursue in this chapter.

⁸To list a few: (1) Integrating across functions weighted by the Pareto density often yields simple closed-form solutions. Hence, for example, if a continuum of firms are charging prices that are distributed Pareto, under standard assumptions about preferences, a closed-form solution for the price index emerges. (2) Truncating the Pareto distribution from below yields a Pareto distribution with the same shape parameter θ . Hence, as we assume subsequently, if entry is subject to a cutoff, the distribution of the technologies that make the cut remains Pareto. (3) A Pareto random variable taken to a power is also Pareto. Hence, if individual prices have a Pareto distribution, with a constant elasticity of demand, so also do sales.

4.1.1 Continuum Case

With a continuum of firms, the measure of firms with efficiency greater than z is thus proportional to $z^{-\theta}$. Hence, we can write the measure of firms in country i with efficiency $Z \geq z$ as:

$$\mu_i^Z(z) = T_i z^{-\theta} \tag{2}$$

where $T_i > 0$ is a parameter reflecting the overall measure of firms in country i .

4.1.2 Finite-Firm Case

We propose an alternative in which, instead, each country i has access to an integer number of technologies, with the number having efficiency $Z \geq z$ the realization of a Poisson random variable with parameter $\mu_i^Z(z) = T_i z^{-\theta}$ (instead of a measure $\mu_i^Z(z)$).⁹ It is useful to rank these technologies according to their efficiency; that is, $Z_i^{(1)} > Z_i^{(2)} > Z_i^{(3)} > \dots > Z_i^{(k)} > \dots$ ¹⁰

⁹In either specification, the level of T_i may reflect a history of innovation, as discussed in Eaton and Kortum (2010, Chapter 4). Furthermore, \underline{z} can be set arbitrarily close to zero. For the finite-firm case, for example, we can consider taking D draws from the Pareto distribution (1), where D is distributed Poisson with parameter $T\underline{z}^{-\theta}$. The number of draws for which $Z > z$ is distributed Poisson with parameter $Tz^{-\theta}$, as we previously assume. The largest Z , called $Z^{(1)}$, is distributed:

$$\Pr[Z^{(1)} \leq z] = \exp(-Tz^{-\theta}),$$

the type II extreme value (Fréchet) distribution. Letting \underline{z} approach zero, as we do throughout this paper, this distribution is defined over all positive values of z .

¹⁰In the Ricardian model of Eaton and Kortum (2002), all technologies in this sequence would be used to produce the same good j so that if country i produces good j , it uses $Z_i^{(1)}$, with all the rest irrelevant. The same is true of production in Bernard, Eaton, Jensen, and Kortum (2003), although $Z_i^{(2)}$ can be relevant in determining the price of good j . In each model, as in Dornbusch, Fischer, and Samuelson (1977), the space

4.2 Costs

We introduce impediments to trade in a standard way: Selling a unit of a good to market n from source i requires exporting $d_{ni} \geq 1$ units, where we set $d_{ii} = 1$ for all i . It also requires hiring a fixed number F_n workers in market n , which we allow to vary by n but, for simplicity, keep independent of i .¹¹

We denote the wage in country i as w_i . Then, a potential producer from country i with efficiency Z_i can sell in country n at a unit cost:

$$C_{ni} = \frac{w_i d_{ni}}{Z_i}.$$

4.2.1 Continuum Case

Under the continuum specification (2), the measure of firms from country i that can sell in country n at unit cost $C_{ni} \leq c$ is:

$$\mu_{ni}^C(c) = \Phi_{ni} c^\theta,$$

where:

$$\Phi_{ni} = T_i (w_i d_{ni})^{-\theta}. \tag{3}$$

of goods is $j \in [0, 1]$. In our finite-firm model here, as in models of monopolistic competition, each technology $Z_i^{(1)}, Z_i^{(2)}, Z_i^{(3)}, \dots$ produces a unique good. How far up the list to go is determined (endogenously) by entry costs.

¹¹As we discuss subsequently, our data handle a cost that is common across sources with relative equanimity but balk at the imposition of an entry cost that is common across destinations. Because assuming a cost that is the same for all entrants in a market yields some simplification, we take that route. Chaney (2008) and EKK (2011) showed how to relax it.

Summing across sources $i = 1, \dots, N$, the measure of firms from anywhere that can sell in n at unit cost c or less is:

$$\mu_n^C(c) = \sum_{i=1}^N \mu_{ni}^C(c) = \Phi_n c^\theta$$

where:

$$\Phi_n = \sum_{i=1}^N \Phi_{ni}. \quad (4)$$

Among firms with unit cost $C_n \leq c$, the fraction from country i is:

$$\pi_{ni} = \frac{\Phi_{ni}}{\Phi_n} \quad (5)$$

regardless of c .

4.2.2 Finite-Firm Case

With only an integer number of firms we can associate each technology $Z_i^{(k)}$ in market i with a unit cost of delivering in market n of

$$C_{ni}^{(k)} = w_i d_{ni} / Z_i^{(k)},$$

so that $C_{ni}^{(1)} < C_{ni}^{(2)} < C_{ni}^{(3)} < \dots$. An implication of this connection between costs and efficiency is that the number of firms from country i that can sell in country n at a cost $C \leq c$ is the realization of a Poisson random variable with parameter $\mu_{ni}^C(c) = \Phi_{ni} c^\theta$ (instead of a measure $\mu_{ni}^C(c)$). Furthermore, the total number of firms that could sell in n at a cost $C \leq c$ is the realization of a Poisson random variable with parameter $\mu_n^C(c) = \Phi_n c^\theta$ (instead of a measure $\mu_n^C(c)$) where Φ_{ni} and Φ_n are still given by (3) and (4).

We can rank these firms according to their unit costs in n irrespective of their source $C_n^{(1)} < C_n^{(2)} < C_n^{(3)} < \dots$. To keep track of the source, we can define an indicator $I_{ni}^{(k)}$ to

equal 1 if the k th lowest-cost firm in n is from i (and 0 otherwise). Properties of the Poisson distribution imply:

$$\Pr \left[I_{ni}^{(k)} = 1 \right] = \pi_{ni},$$

where π_{ni} is defined in (5). The probability that the firm is from i is independent of its rank k in country n or its unit cost there, $C_n^{(k)}$. Unlike the continuum model, π_{ni} is now the expected fraction of firms from i selling in n , rather than the realized fraction.

4.3 Entry

To close the model, we specify total spending in a market as X_n , and make the standard assumption that demand derives from an aggregator with a constant elasticity of substitution $\sigma > 1$.

Under these assumptions, a firm selling in market n with a unit cost C charging a price p makes a profit in that market, gross of the entry cost $E_n = w_n F_n$, of:

$$\Pi_n(p, C) = \left(1 - \frac{C}{p}\right) \left(\frac{p}{P_n}\right)^{-(\sigma-1)} X_n, \quad (6)$$

where P_n is the price index in country n .

4.3.1 Continuum Case

In the case of a continuum of firms, each firm – no matter how efficient – has no effect on the overall price index P_n . It therefore sets a price $p_n(C)$ to maximize (6) taking P_n as given, thereby choosing the standard Dixit-Stiglitz markup:

$$p_n(C) = \bar{m}C,$$

where:

$$\bar{m} = \frac{\sigma}{\sigma - 1}.$$

Variable profit is decreasing in unit cost C . Hence, firms enter market n up to the point at which their unit cost implies zero profit, net of the fixed entry cost, at a cost threshold \bar{c}_n satisfying:

$$\left(\frac{\bar{m}\bar{c}_n}{P_n}\right)^{-(\sigma-1)} \frac{X_n}{\sigma} = E_n. \quad (7)$$

The resulting price index is:

$$P_n = \left[\int_0^{\bar{c}_n} (\bar{m}c)^{-(\sigma-1)} d\mu_n^C(c) \right]^{-1/(\sigma-1)}. \quad (8)$$

Under the restriction that

$$\theta > \sigma - 1 \quad (9)$$

Equations (7) and (8) together deliver nice analytic expressions for the price index:

$$P_n = \bar{m} \left[\frac{\theta}{\theta - (\sigma - 1)} \right]^{-1/\theta} \left(\frac{X_n}{\sigma E_n} \right)^{-[\theta - (\sigma - 1)]/[\theta(\sigma - 1)]} \Phi_n^{-1/\theta}$$

and cutoff:

$$\bar{c}_n = \left[\frac{\theta}{\theta - (\sigma - 1)} \right]^{-1/\theta} \left(\frac{X_n}{\sigma E_n} \right)^{1/\theta} \Phi_n^{-1/\theta}.$$

Without the restriction (9), however, the price index and cutoff are undefined. The reason is that technological heterogeneity and the elasticity of substitution are so large that buyers achieve zero cost by squeezing all of their spending into the lower tail of the cost distribution. Hence, it is standard in models with a continuum of goods to impose a restriction like (9).

In the continuum model, the firm sales distribution is Pareto with parameter $\theta/(\sigma - 1)$:¹²

$$\Pr [X_n(C) \geq x | C \leq \bar{c}_n] = \left(\frac{x}{\sigma E_n} \right)^{-\theta/(\sigma-1)}. \quad (10)$$

Restriction (9) prevents this parameter from falling to 1 or below. Hence, the model cannot predict a highly skewed sales distribution without wandering into “forbidden” territory.

4.3.2 Finite-Firm Case

With only an integer number of firms, the restriction (9) is not needed because at no point do we integrate over the distribution of prices. Solving for the equilibrium is less straightforward, however. In the continuum model, an individual firm (of measure zero) naturally takes aggregate spending X_n , the wage w_n , and the price index P_n as given in deciding what price to charge and whether to enter.

So as not to introduce too many complications into our finite-firm case at once, we continue to assume that firms take expenditure X_n and the wage w_n as given but incorporate the effect

¹²The sales of a firm with cost $C \leq \bar{c}_n$ are:

$$X_n(C) = \left(\frac{\bar{m}C}{P_n} \right)^{-(\sigma-1)} X_n,$$

whereas the distribution of costs for such firms is

$$\Pr [C \leq c | C \leq \bar{c}_n] = \frac{\mu_n^C(c)}{\mu_n^C(\bar{c}_n)} = \left(\frac{c}{\bar{c}_n} \right)^\theta.$$

Combining these results using (7) yields (10). The expected sales of a firm in market n are:

$$\begin{aligned} \bar{X}_n &= \int_{\sigma E_n}^{\infty} x \frac{\theta}{\sigma - 1} x^{\theta/(\sigma-1)-1} (\sigma E_n)^{\theta/(\sigma-1)} dx \\ &= \frac{\theta \sigma E_n}{\theta - (\sigma - 1)}, \end{aligned}$$

which is finite only under the parameter restriction (9).

of their decisions on the price index P_n .¹³

We treat equilibrium in any market as determined in two stages. In stage two, the number of firms K_n entering into each market is given. The firms present in each market engage in Bertrand competition. This competition establishes a price associated with each unit cost C , denoted by $p_n(C)$.¹⁴ Gross profit in market n of a firm with unit cost C is:

$$\Pi_n(C) = \left[1 - \frac{C}{p_n(C)}\right] \left(\frac{p_n(C)}{P_n}\right)^{-(\sigma-1)} X_n, \quad (11)$$

where now the price index is:

$$P_n = \left(\sum_{k=1}^{\infty} [p_n(C_n^{(k)})]^{-(\sigma-1)} I_n^{(k)}\right)^{-1/(\sigma-1)}$$

where $I_n^{(k)} = 1$ if the firm with the k th lowest unit cost enters and $I_n^{(k)} = 0$ otherwise. Hence,

$$\sum_{k=1}^{\infty} I_n^{(k)} = K_n.$$

¹³A huge benefit of treating X_n and w_n as unaffected by firms' entry and price decisions is that we can separately analyze equilibrium in each market. Otherwise, a pricing or entry decision in any market would affect sales and wages in every country of the world, making the equilibrium difficult to compute. Various assumptions can justify our treatment of X_n and w_n as exogenous to the entry decision, all of them inelegant. One assumption is that wages are determined by trade in a different sector; a second is that profits are all spent on output from a different sector. We then can set $X_n = \alpha w_n L_n$ where L_n is the labor force in country n and α is the Cobb-Douglas share of manufactures in workers' spending. The first assumption has a tradition in this literature, being a case considered in Eaton and Kortum (2002); it also appears in Chaney (2008) and Melitz and Ottaviano (2008). The second assumption resembles that of the "absentee landlord" sometimes posited in the economic geography literature. We hope that future research explores the implications of more attractive assumptions, but we do not expect results to differ much from those we report.

¹⁴In fact, the price chosen by each firm in the Bertrand equilibrium in market n depends on the unit costs of all firms present. We list only the firm's own unit cost C as an argument of p_n because two firms that happened to have the same C , a probability zero event, would charge the same price $p_n(C)$.

We consider a situation in which $I_n^{(k)} = 1$ for $k \leq K_n$ and 0 otherwise, giving the price index:

$$P_n^{K_n} = \left(\sum_{k=1}^{K_n} [p_n^{K_n}(C_n^{(k)})]^{-(\sigma-1)} \right)^{-1/(\sigma-1)}$$

where the K_n superscript denotes the dependence of prices on the number of entrants. The corresponding gross profit of the k th lowest-cost firm, with $k \leq K_n$, is:

$$\Pi_n^{K_n}(C_n^{(k)}) = \left[1 - \frac{C_n^{(k)}}{p_n^{K_n}(C_n^{(k)})} \right] \left(\frac{p_n^{K_n}(C_n^{(k)})}{P_n^{K_n}} \right)^{-(\sigma-1)} X_n$$

The following (unsurprising) result, which we call *profit monotonicity*, is useful in determining entry.¹⁵

$$\Pi_n^{K+1}(C_n^{(K+1)}) \leq \Pi_n^K(C_n^{(K)}) \tag{12}$$

¹⁵An outline of the proof is as follows. First, we note that in any Bertrand equilibrium:

$$\Pi_n^{K+1}(C_n^{(K+1)}) \leq \Pi_n^{K+1}(C_n^{(K)})$$

The reason is that the firm with unit cost $C_n^{(K)}$ could always earn a higher profit than the firm with cost $C_n^{(K+1)}$ simply by charging the same price as that firm (hence, selling the same quantity at a lower cost). Second, we note that removing the firm with unit cost $C_n^{(K+1)}$ raises the profit of all remaining firms. We can consider the profit of the k 'th firm as a function of the price of each entrant (not necessarily its equilibrium price), which we denote by $\Pi_n^{(k)}(p_n^{(1)}, p_n^{(2)}, \dots, p_n^{(K+1)})$. Removing firm $K+1$ is a special case of raising its price arbitrarily. That the profit of all remaining firms rises follows from the fact that both:

$$\frac{\partial \Pi_n^{(k)}}{\partial p_n^{(k')}} \geq 0$$

and

$$\frac{\partial^2 \Pi_n^{(k)}}{\partial p_n^{(k)} \partial p_n^{(k')}} \geq 0$$

for $k' \neq k$. The first inequality implies that a higher price on the part of a rival raises profit and the second implies that a higher price by a rival raises the price charged by any other firm. Hence, a higher price from the firm with unit cost $C_n^{(K+1)}$ causes every other firm to raise its price, which raises the profit of all remaining

In the first stage, firms decide whether to enter each market. To avoid uninteresting multiple equilibria, we assume that they make their entry decisions sequentially, starting with the firm with the lowest unit cost $C_n^{(1)}$, followed by the firm with unit cost $C_n^{(2)}$, and so on.¹⁶ When making its decision to enter, each firm anticipates perfectly what its profit would be in the subsequent second-stage Bertrand equilibrium.

An immediate implication of profit monotonicity (12) is that the two conditions:

$$\Pi_n^{K_n}(C_n^{(K_n)}) \geq E_n$$

and:

$$\Pi_n^{K_n+1}(C_n^{(K_n+1)}) < E_n.$$

determine K_n . Firms will enter up to the point at which the firm with unit cost $C_n^{(K_n+1)}$ would not be able to cover its entry cost.¹⁷

We have now completed the statement of the finite-firm model. With a finite number of firms, the full set of parameters $\theta, \sigma, T_i, d_{ni}, X_n$, and w_n are not enough to determine the firms, including that of the firm with unit cost $C_n^{(K)}$. Letting $C_n^{(K+1)}$ rise without bound:

$$\Pi_n^{K+1}(C_n^{(K)}) \leq \Pi_n^K(C_n^{(K)}).$$

Combining these two profit inequalities delivers the profit-monotonicity result (12).

¹⁶With a discrete number of firms, a possible outcome is entry by one or more less-efficient firms blocking a more-efficient firm from entering. With a continuum of firms, this possibility does not arise because no firm has any effect on the aggregate outcome.

¹⁷To avoid the outcome $K_n = 0$ (in which case we could not have taken X_n as given), we assume $X_n \geq E_n$. A possible outcome is $K_n = 1$, in which case the monopolist charges a price approaching infinity, supplies a quantity approaching 0, and obtains gross profit of $\Pi_n^1(C_n^{(1)}) = X_n$. In subsequent sections, we fit the model to data and simulate entry. With realistic parameter values, the monopoly outcome never comes close to happening.

equilibrium of the model. We also need the realizations of the technologies $Z_i^{(k)}$ in each source i determining ordered costs $C_n^{(k)}$ in each destination n . The equilibrium in each destination determines overall entry K_n and the price level P_n , as well as entry by individual firms, as indicated by $I_{ni}^{(k)}$, and their sales there:¹⁸

$$X_n^{(k)} = \left(\frac{p_n^{K_n}(C_n^{(k)})}{P_n^{K_n}} \right)^{-(\sigma-1)} X_n. \quad (13)$$

The total number of firms from i entering n is thus:

$$K_{ni} = \sum_{k=1}^{K_n} I_{ni}^{(k)} \quad (14)$$

and their total sales:

$$X_{ni} = \sum_{k=1}^{K_n} I_{ni}^{(k)} X_n^{(k)}. \quad (15)$$

We conclude with three implications of the discrete model important for the quantitative analysis that follows:

1. The probability π_{ni} that a firm selling in country n is from country i is independent of its rank k or its unit cost $C_n^{(k)}$ in market n , and hence, of its sales there, $X_n^{(k)}$.
2. Because the number of firms K_{ni} from i selling in n is determined by a finite number of Bernoulli trials, zero is a possibility.
3. Unlike the continuum model, we need no restrictions on θ other than $\theta > 0$.

¹⁸The definition of the price index ensures that the sales of each entrant sum to total spending:

$$\sum_{k=1}^{K_n} X_n^{(k)} = X_n.$$

5 Quantification

Our goal is to determine whether our finite-firm model can capture patterns of trade at both the aggregate and firm levels. We proceed in the following five steps, which culminate in a fully parameterized version of the finite-firm model:

1. We specify a gravity equation consistent with our firm-level model, which we estimate using data on bilateral trade in manufactures. This step provides estimates of the market-entry probability π_{ni} given in (5).
2. We use firm-level data to extract an estimate of mean sales per firm \bar{X}_n in each market n , from which we can estimate total entry $K_n = X_n/\bar{X}_n$. Our estimates of π_{ni} and K_n allow us to calculate the probability of zero exports from each source i to each destination n .
3. We construct cost draws that allow us to simulate an entire matrix of entry by firms from each source i in each destination n .
4. We use these cost draws to calculate the Bertrand equilibrium in each destination. This calculation yields the sales distribution of firms across markets.
5. We infer entry costs E_n , which completes the parameterization.

At the completion of the fifth step, we are prepared to perform the counterfactuals in the subsequent section.

5.1 The Gravity Equation

The gravity equation has a long and successful history of capturing empirically how much one country sells to another. A standard formulation is:

$$X_{ni} = \frac{Y_i X_n}{k_{ni}}$$

where Y_i is production in the exporting country i and k_{ni} is the distance from i to n . Although there has been much progress in deriving such an equation (suitably modified) from theories of trade, an important remaining issue in taking this equation to data is the specification of the error term.

Our finite-firm model implies that randomness can emerge from two sources. First, given the provenance of firms that have entered a market, firms from some source might have drawn particularly low C s and thus sell more. Second, given the expectation π_{ni} that a firm in n is from i , firms from i might have had particularly lucky rolls of the die and therefore have a larger than expected presence in n .

To capture these two sources of error, we divide each side of (15) by total expenditure on manufactures (i.e., absorption) in market n and take the expectations of each side to obtain:

$$E \left[\frac{X_{ni}}{X_n} \right] = E \left[\sum_{k=1}^{K_n} I_{ni}^{(k)} \frac{X_n^{(k)}}{X_n} \right] = E \left[\sum_{k=1}^{K_n} E \left[I_{ni}^{(k)} | X_n^{(k)} \right] \frac{X_n^{(k)}}{X_n} \right].$$

The first implication of our model enumerated at the end of Section 4 – that the probability of a firm being from i is independent of k and $X_n^{(k)}$ – allows us to write this expression as:

$$E \left[\frac{X_{ni}}{X_n} \right] = \pi_{ni} E \left[\sum_{k=1}^{K_n} \frac{X_n^{(k)}}{X_n} \right].$$

Because the remaining summation is over all firms selling in n , it is identically 1 (and, hence,

its expectation is also). We have simply:

$$E \left[\frac{X_{ni}}{X_n} \right] = \pi_{ni}, \quad (16)$$

The expectation of country i 's market share in n is the probability that any particular firm in n is from i .

We can use equations (3), (4), and (5) to connect π_{ni} to our model and write it as a multinomial logit function:

$$E \left[\frac{X_{ni}}{X_n} \right] = \pi_{ni} = \frac{\exp(\ln T_i - \theta \ln w_i - \theta \ln d_{ni})}{\sum_{l=1}^N \exp(\ln T_l - \theta \ln w_l - \theta \ln d_{nl})}. \quad (17)$$

Equation (17) is the basis of our gravity estimation, with X_{ni}/X_n measured by actual trade shares. We parameterize the right-hand side of (17) as follows. We set:

$$S_i = \ln T_i - \theta \ln w_i,$$

capturing source-specific determinants of trade as a fixed effect. We use geographical measures to capture the costs of exporting from i to n . Specifically, for $i \neq n$, we set:

$$-\theta \ln d_{ni} = m_n + g'_{ni} \alpha + \ln \nu_{ni}.$$

Here, m_n is a destination fixed effect capturing differences in openness to imports and g_{ni} is a vector of observables potentially raising trade costs (in our case, the log of distance and indicators for lack of a common border, lack of a common language, lack of a common legal origin, lack of a common colonizer, and lack of colonial ties).¹⁹ Because these indicators are unlikely to reflect all aspects of trade costs, we also introduce an unobservable component of trade costs ν_{ni} (with $\nu_{nn} = 1$ because $d_{nn} = 1$).

¹⁹These variables are from Head, Mayer, and Ries (2010), available on the CEPII web site.

We now have an additional source of randomness. The connection between observables and π_{ni} is itself random.²⁰

To obtain an expression suitable for estimation, we define:

$$\varphi_{ni} = \begin{cases} \exp(S_i + m_n + g'_{ni}\alpha) & i \neq n \\ \exp(S_n) & i = n \end{cases},$$

and:

$$\Lambda_{ni} = \frac{\varphi_{ni}}{\sum_l \varphi_{nl}}.$$

We then can write:

$$\pi_{ni} = \frac{\varphi_{ni}\nu_{ni}}{\sum_l \varphi_{nl}\nu_{nl}}. \quad (18)$$

To apply a standard estimation procedure, ν_{ni} must have the property that:

$$E[\pi_{ni}|\Lambda] = E\left[\frac{\varphi_{ni}\nu_{ni}}{\sum_l \varphi_{nl}\nu_{nl}} \middle| \Lambda\right] = \Lambda_{ni}, \quad (19)$$

where conditioning on the observables Λ means that the ν_{ni} are treated as random variables.

Thus, constructing Λ_{ni} from the true parameters and observables delivers an unbiased predictor of π_{ni} .²¹

²⁰An analogy is the likelihood of a 3 from a roll of a die with unknown bias. There is randomness due not only to multinomial sampling but also to the uncertainty of the bias. Our error term ν_{ni} introduces such bias.

We assume that the distribution of ν_{ni} is such that there is no ex-ante bias.

²¹One way to construct ν_{ni} that satisfy (19) is based on the gamma distribution. A random variable X is distributed Gamma(a, b) (with mean ab and variance ab^2) if its distribution is:

$$\Pr[X \leq x] = \frac{1}{\Gamma(a)} \int_0^{x/b} t^{a-1} \exp(-t) dt$$

where

$$\Gamma(a) = \int_0^\infty t^{a-1} \exp(-t) dt$$

Putting together these sources of error, the moment conditions we use for our estimation

are:

$$E \left[\frac{X_{ni}}{X_n} \middle| \Lambda \right] = \Lambda_{ni} = \begin{cases} \frac{\exp(S_i + m_n + g'_{ni}\alpha)}{\exp(S_n) + \sum_{l \neq n} \exp(S_l + m_n + g'_{nl}\alpha)} & i \neq n \\ \frac{\exp(S_n)}{\exp(S_n) + \sum_{l \neq n} \exp(S_l + m_n + g'_{nl}\alpha)} & i = n \end{cases}.$$

These conditions are nonlinear in the parameters that we need to estimate. However, because the Λ_{ni} sum to 1 across all sources i (for any n), the parameters can be estimated quite easily by multinomial pseudo-maximum likelihood (PML), as described in Gouriéroux, Monfort, and Trognon (1984). We apply this estimator to our data on bilateral trade X_{ni} between 92

is the complete gamma function. We let $\nu_{ni} = (V_{ni}/V_{nn})$ so that

$$\pi_{ni} = \frac{\varphi_{ni}\nu_{ni}}{\sum_l \varphi_{nl}\nu_{nl}} = \frac{\Lambda_{ni}V_{ni}}{\sum_l \Lambda_{nl}V_{nl}}$$

and assume that V_{ni} is distributed $\text{Gamma}\left(\frac{\Lambda_{ni}}{\eta^2}, \frac{\eta^2}{\Lambda_{ni}}\right)$. From the properties of the gamma distribution, we have:

$$\Lambda_{ni}V_{ni} \approx \text{Gamma}\left(\frac{\Lambda_{ni}}{\eta^2}, \eta^2\right),$$

and

$$\sum_l \Lambda_{nl}V_{nl} \approx \text{Gamma}\left(\frac{1}{\eta^2}, \eta^2\right).$$

The vector of π_{ni} 's is therefore distributed:

$$(\pi_{n1}, \pi_{n2}, \pi_{n3}, \dots, \pi_{nN}) \approx \text{Dirichlet}\left(\frac{\Lambda_{n1}}{\eta^2}, \frac{\Lambda_{n2}}{\eta^2}, \frac{\Lambda_{n3}}{\eta^2}, \dots, \frac{\Lambda_{nN}}{\eta^2}\right)$$

with mean:

$$E[\pi_{ni}] = \frac{\Lambda_{ni}/\eta^2}{\sum_l \Lambda_{nl}/\eta^2} = \Lambda_{ni}$$

so that (19) is satisfied. The variance is given by:

$$\text{Var}[\pi_{ni}] = \frac{\eta^2}{\eta^2 + 1} \Lambda_{ni}(1 - \Lambda_{ni})$$

The derivation follows the derivation of the random effects negative binomial model in Hausman, Hall, and Griliches (1984).

countries, where we include home sales, X_{nn} , for $i = n$.²²

The results appear in the last column of Table 3 showing the coefficients $\hat{\alpha}$ on the gravity variables. In line with most gravity equations specified in a more conventional form, the coefficient on the log of distance is estimated to be near -1.²³ All of the other geography variables have the expected negative effect on trade as well. For comparison, the first two columns of the table show estimates of the same parameters obtained by approaches that have been used in earlier work.²⁴ Focusing on the coefficient of log distance, our results are in between what is delivered by running a regression in logs, dropping observations with zero

²²In the continuum model of Eaton and Kortum (2002), instead of (16), we would have:

$$\frac{X_{ni}}{X_n} = \pi_{ni}$$

(without the expectation). In that case, we can write (18) as:

$$\frac{X_{ni}}{X_n} = \frac{\varphi_{ni}\nu_{ni}}{\sum_l \varphi_{nl}\nu_{nl}}.$$

Eaton and Kortum could normalize by X_{nn}/X_n without violating Jensen's inequality to obtain the specification:

$$\frac{X_{ni}}{X_{nn}} = \frac{\varphi_{ni}\nu_{ni}}{\varphi_{nn}}.$$

If $X_{ni} > 0$ for all country pairs (as it was in Eaton and Kortum's (2002) OECD sample), this equation could be estimated by ordinary least squares (OLS) after taking logarithms of both sides. Closer to our approach here, and also tackling the zeros problem, is the Poisson PML approach taken by Santos Silva and Tenreyro (2006).

²³Chaney (2011) provided a theoretical explanation for this regularity.

²⁴The first column follows Eaton and Kortum (2002), in which the dependent variable was $\ln(X_{ni}/X_{nn})$. Although this approach can be estimated using OLS, it requires dropping observations with zero trade. The middle column follows Santos Silva and Tenreyro (2006), applying Poisson PML, with the dependent variable X_{ni} . Neither of these other approaches is fully consistent with our finite-firm model.

trade flows, and the one delivered by Poisson PML.²⁵

Table 3. Bilateral trade regressions

	OLS	Poisson	Multinomial
Distance	-1.418*** (0.0379)	-0.699*** (0.0444)	-1.072*** (0.0511)
Lack of Contiguity	-0.442** (0.156)	-0.694*** (0.181)	-0.370** (0.136)
Lack of Common Language	-0.686*** (0.0808)	0.121 (0.131)	-0.511*** (0.106)
Lack of Common Legal Origin	-0.184** (0.0593)	-0.281*** (0.0778)	-0.133 (0.0721)
Lack of Common Colonizer	-0.212 (0.146)	0.222 (0.199)	-0.306 (0.204)
Lack of Colonial Ties	-0.684*** (0.126)	0.226 (0.122)	-0.953*** (0.139)
Adjusted R sq.	0.968		
Pseudo R sq.		0.993	0.563
Number of observations	5483	8464	8464

Standard errors in parentheses

* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

²⁵One explanation for the different results is that the three estimation approaches apply different penalties to deviations between model and data for large and small trade flows. By taking logs, the approach in the first column treats proportional deviations as equally likely across all observations. At the other extreme, Poisson PML applies a much greater penalty to a given proportional deviation in a large trade flow than in a small one (because a proportional deviation from the mean becomes less likely for a Poisson distributed random variable as the mean is increased). Our current approach is between the two. Our dependent variable normalizes bilateral trade flows by the importers' total absorption, thus eliminating different penalties for proportional deviations across large and small trade flows, to the extent that they vary with the size of the destination. Yet, our approach is more like Poisson PML in that that proportional deviations from a large exporting country are much less likely than for a small exporter.

We use our estimates of the gravity equation to calculate

$$\hat{\pi}_{ni} = \frac{\hat{\varphi}_{ni}}{\sum_{k=1}^N \hat{\varphi}_{nk}}$$

using the estimated coefficients \hat{S}_i , \hat{m}_n , and $\hat{\alpha}$ and data on source and geography. These estimated entry probabilities $\hat{\pi}_{ni}$ have the desirable properties of lying *strictly* between 0 and 1, even though they are based on trade shares that are frequently zero in the data. Thus, they predict a positive trade flow even when none is observed in the data.

5.2 Mean Sales per Firm

Because we assume that the cost of entry E_n does not vary with the source country i , our model implies that in expectation, mean sales in a destination should be the same for all i :

$$E[\bar{X}_{ni}] = E[\bar{X}_n].$$

We can exploit this restriction to estimate mean sales \bar{X}_n of firms in a market using our limited firm-level data. To do so, we pool our data on sales across the four source countries (i.e., Brazil, Denmark, France, and Uruguay). As described previously, we restrict ourselves to destinations for which we have data from at least two sources.²⁶ Our estimate of mean sales is simply:

$$\hat{\bar{X}}_n = \frac{\sum_{i \in \Omega_n} K_{ni} \bar{X}_{ni}}{\sum_{i' \in \Omega_n} K_{ni'}} \tag{20}$$

²⁶We drop the home-country observations (when available) because the universe of firms selling in the home market typically is measured very differently. The customs data indicate the number of exporters and their sales in a foreign market. The total number of active firms in a country is more difficult to determine because many may not be counted. Because there are so few exporters from Uruguay, we merge the data for that country in 1992 and 1993.

where Ω_n is the set of source countries for which we have firm-level data on exports to destination n . The results are shown in Table 4.²⁷ Mean sales range from \$47,000 in the Central African Republic to \$1.6 million in the United States, with an elasticity with respect to total expenditure X_n of 0.33.²⁸

Because we treat expenditures on manufactures in a market X_n as fixed at its actual value, our estimate \widehat{X}_n provides a way to infer the number of firms $\widehat{K}_n = X_n/\widehat{X}_n$ that sell there. Simulating the model to generate $K_n = \widehat{K}_n$, our simulation will automatically match our estimate of mean sales per firm.

With $\widehat{\pi}_{ni}$ and \widehat{K}_n in hand, we can calculate the likelihood of a zero bilateral export as follows. Without the trade cost shocks ν_{ni} , we can use the binomial distribution to obtain an expression for the probability of a zero:

$$\Pr[K_{ni} = 0] = (1 - \pi_{ni})^{K_n} \quad (21)$$

²⁷To gauge the plausibility of our restriction that $E_{ni} = E_n$, we examine whether mean sales \overline{X}_{ni} of our four source countries (which are diverse in economic size and development) differ among each other in a systematic way. We run a weighted OLS regression of the unbalanced panel \overline{X}_{ni} on a full set of destination-country effects and source-country effects. (The weights, $K_{ni}/(\widehat{X}_n)^2$, correct for the fact that the observations are averages over different numbers of firms, and destination countries differ in mean sales.) Table 5 reports the source-specific intercepts relative to France, which is normalized to zero. The estimates imply modest variation across sources, with Brazil's mean sales about \$70,000 higher than in France, whereas Denmark's and Uruguay's mean sales are about \$25,000 lower. Although we can easily reject the joint hypothesis of equal mean sales by source, the small magnitude of the deviations suggests that we will not do great violence to the data by simply ignoring them.

²⁸The slope in Figure 1 implies that exports per firm rise with a country's total exports with an elasticity of 0.29. As is known from the gravity literature, total exports increase with destination expenditures with an elasticity close to 1, so the two results are very much in line.

Table 4. Mean sales per firm

Destination country	No. of source countries	Mean sales per firm
Algeria	2	0.426
Angola	2	0.272
Argentina	4	0.638
Australia	4	0.324
Austria	4	0.334
Bangladesh	2	0.391
Benin	2	0.079
Bolivia	3	0.174
Brazil	3	0.493
Bulgaria	4	0.211
Burkina Faso	2	0.065
Burundi	2	0.065
Cameroon	2	0.096
Canada	4	0.301
Central African Republic	2	0.047
Chad	2	0.070
Chile	4	0.345
China	3	1.811
Colombia	3	0.351
Costa Rica	3	0.190
Côte d'Ivoire	2	0.134
Denmark	3	0.323
Dominican Republic	3	0.258
Ecuador	3	0.229
Egypt	4	0.486
El Salvador	3	0.118
Ethiopia	2	0.099
Finland	4	0.223
France	3	0.904
Ghana	2	0.194
Greece	4	0.354
Guatemala	3	0.151
Honduras	3	0.090
Hungary	4	0.226
India	4	0.452
Indonesia	3	1.162
Iran	4	1.121
Ireland	4	0.301
Israel	3	0.235
Italy	4	1.375
Jamaica	3	0.132
Japan	4	1.124
Jordan	3	0.171
Kenya	3	0.230
Korea	4	0.715
Kuwait	4	0.256

continued next page

Destination country	No. of source countries	Mean sales per firm
Madagascar	2	0.079
Malawi	2	0.126
Malaysia	3	0.435
Mali	2	0.082
Mauritania	2	0.107
Mauritius	2	0.101
Mexico	4	0.835
Morocco	3	0.258
Mozambique	2	0.519
Nepal	3	0.173
Netherlands	4	0.884
New Zealand	4	0.108
Nigeria	3	0.618
Norway	4	0.290
Oman	2	0.422
Pakistan	3	0.414
Panama	3	0.195
Paraguay	3	0.229
Peru	3	0.199
Philippines	4	0.502
Portugal	4	0.346
Romania	4	0.292
Rwanda	2	0.055
Saudi Arabia	4	0.536
Senegal	2	0.093
South Africa	3	0.238
Spain	4	0.992
Sri Lanka	3	0.291
Sweden	4	0.446
Switzerland	4	0.314
Syrian Arab Republic	2	0.341
Taiwan	4	0.607
Tanzania, United Rep. of	2	0.130
Thailand	4	0.692
Togo	3	0.077
Trinidad and Tobago	3	0.170
Tunisia	3	0.240
Turkey	4	0.497
Uganda	2	0.061
United Kingdom	4	1.311
United States of America	4	1.603
Uruguay	2	0.176
Venezuela	3	0.330
Viet Nam	3	0.548
Zambia	2	0.110
Zimbabwe	2	0.195

Table 5. Source-country coefficients

	Mean sales*
Denmark	-0.0279 (0.0216)
Brazil	0.0724** (0.0221)
Uruguay	-0.0265 (0.0680)
p-value for F test of joint significance	0.0050
Number of observations	282

Standard errors in parentheses

*OLS Regression also includes all destination country effects as independent variables

* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

which we evaluate at $K_n = \widehat{K}_n$ and $\pi_{ni} = \widehat{\pi}_{ni}$.²⁹ Figure 2a is a histogram with the predicted probability of zero trade along the horizontal axis and the frequency of observations with that predicted probability on the vertical axis, for all country pairs in which $X_{ni} = 0$. (Figure 2b reports the corresponding histogram for pairs in which $X_{ni} > 0$.) Although we predict a low probability of zero trade when there in fact is trade (Figure 2b), we sometimes also predict a low probability of zero trade even when there is no trade (Figure 2a). Including trade-cost shocks helps to reduce the errors in Figure 2a.

²⁹We can also incorporate ν_{ni} s along the lines proposed in Footnote 16 as follows. Our assumptions there imply that the vector of π_{ni} s for each destination n are distributed Dirichlet($\Lambda_{n1}/\eta^2, \Lambda_{n2}/\eta^2, \dots, \Lambda_{nN}/\eta^2$). The corresponding marginal distribution for any source i is distributed Beta($\Lambda_{ni}/\eta^2, (1 - \Lambda_{ni})/\eta^2$). The probability of a zero is then:

$$\begin{aligned} \Pr[K_{ni} = 0] &= \frac{\Gamma(1/\eta^2)}{\Gamma(\Lambda_{ni}/\eta^2)\Gamma((1 - \Lambda_{ni})/\eta^2)} \int_0^1 (1-x)^{K_n} x^{(\Lambda_{ni}/\eta^2)-1} (1-x)^{(1-\Lambda_{ni})/\eta^2-1} dx \\ &= \frac{\Gamma(1/\eta^2)\Gamma(K_n + (1 - \Lambda_{ni})/\eta^2)}{\Gamma((1 - \Lambda_{ni})/\eta^2)\Gamma(K_n + 1/\eta^2)}, \end{aligned}$$

which we evaluate at $K_n = \widehat{K}_n$ and $\Lambda_{ni} = \widehat{\pi}_{ni}$ and $\eta^2 = 0.0001$ (the results deteriorate with larger values of η^2).

Figure 2a. Probabilities of observing zero trade, given no trade





Figure 2b. Probabilities of observing zero trade, given positive trade



We can also use equation (21) to simulate the number of export destinations and import sources for each country (the actual values are shown in the last two columns of Table 1).³⁰ The simulated average number of unidirectional trade links per country is 70.5 (from a maximum of 91), somewhat overpredicting the actual number. The simulations also fit the fact that the variance is higher across export destinations, 1,077, than across import sources, 48.6, although this difference is substantially magnified relative to the data. Figures 3a and 3b provide a more detailed comparison of the simulations and the data, plotting each against a country's expenditure on manufactures, a convenient measure of country size. Whereas the model captures the basic pattern that trade links rise with country size, for small countries it typically undershoots the number of export destinations and overshoots the number of import sources.

Why is our model able to predict that zeros are so much more variable across export destinations than across importer sources? A reason is that a country's success in penetrating a market as an exporter depends on the efficiency of its most efficient firm, generating enormous correlation

³⁰A simulation proceeds as follows, letting q_{ni} denote the estimated value of (21). We draw W_i independently for $i = 1, 2, \dots, N$ from the uniform distribution on $[0,1]$. We construct the indicator $\delta_{ni} = 1$ if $W_i > q_{ni}$ and $\delta_{ni} = 0$ otherwise. We count i 's export destinations as:

$$N^E(i) = \sum_{n \neq i} \delta_{ni}$$

and count n 's import sources as

$$N^I(n) = \sum_{i \neq n} \delta_{ni}.$$

The results presented are based on averages from carrying out this simulation 1,000 times.

Figure 3a. Actual and Simulated Number of Destinations


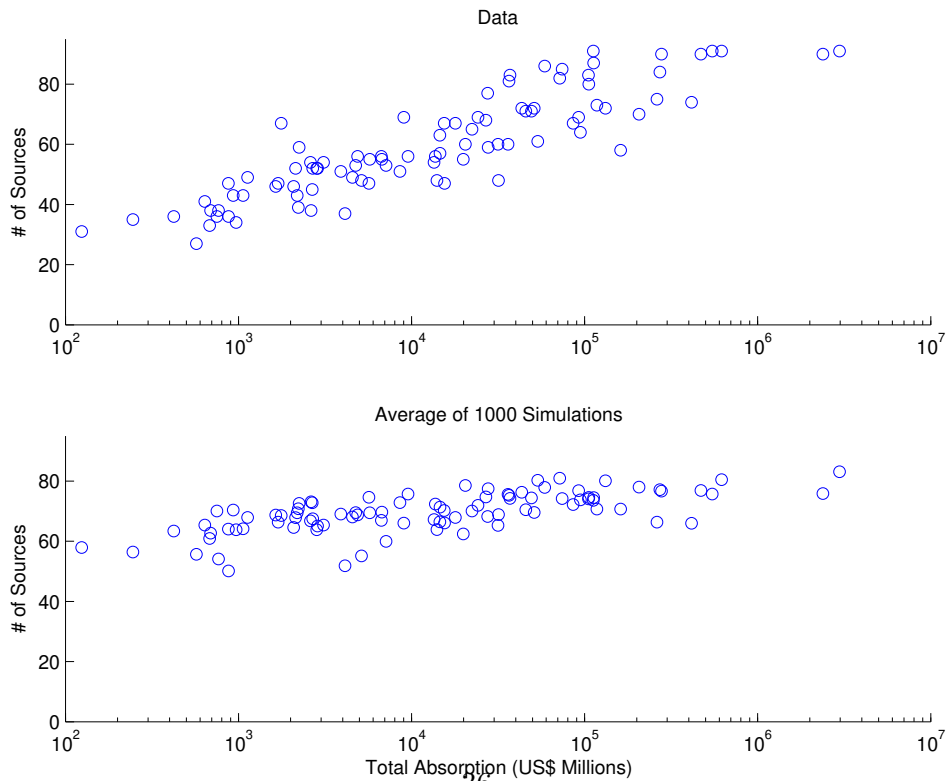


Figure 3b. Actual and Simulated Number of Sources



across foreign markets in entry. Thus, two countries with the same geography and size likely would be very similar in terms of their ability to attract entry from other countries. However, the two countries would differ enormously in their ability to penetrate foreign markets if the lead firm in one was much more efficient than the lead firm in the other. Thus, the much greater variance in exporter zeros is consistent with the finite-firm model.

5.3 Simulating Unit Costs

An advantage of formulating the model in terms of an ordering of efficiencies and unit costs is that we can exploit properties of order statistics to simulate these objects. In particular, our model implies that the most efficient firm from each source i has an efficiency $Z_i^{(1)}$ drawn from the Type II extreme-value (Fréchet) distribution:

$$\Pr[Z_i^{(1)} \leq z] = e^{-T_i z^{-\theta}}.$$

It follows that $U_i^{(1)} = T_i \left(Z_i^{(1)} \right)^{-\theta}$ is distributed exponential, free of any parameters:

$$\Pr \left[U_i^{(1)} \leq u \right] = 1 - e^{-u}.$$

We can proceed up the ordered efficiencies, defining:

$$U_i^{(k)} = T_i \left(Z_i^{(k)} \right)^{-\theta} \tag{22}$$

for any $k > 1$. In Eaton and Kortum (2010), we show that the spacings in this sequence also have an exponential distribution:

$$\Pr \left[U_i^{(k+1)} - U_i^{(k)} \leq u \right] = 1 - e^{-u}.$$

For each source i we construct $U_i^{(k)}$ for k up to \bar{K} 3.2 million, much more than we ever need.

The resulting normalized ordered costs (inversely related to efficiency) for each source i are

simply a random walk of length \bar{K} with unit-exponential increments and an initial value drawn from a unit exponential.

We use these normalized ordered costs to construct ordered unit costs $C_{ni}^{(k)}$ of delivery to country n by firms from i , invoking (22), (5), and (3):

$$C_{ni}^{(k)} = \frac{w_i d_{ni}}{\left(U_i^{(k)}/T_i\right)^{-1/\theta}} = \left(\frac{U_i^{(k)}}{\Phi_n \pi_{ni}}\right)^{1/\theta},$$

for $k = 1, 2, 3, \dots, \bar{K}$. Estimates of π_{ni} and θ are needed in this step. We use $\hat{\pi}_{ni}$ for π_{ni} and set $\theta = 4.87$ from EKK (2011).³¹ (The term Φ_n cancels out of the relevant formulas.)

In any particular destination n , we can combine all the $C_{ni}^{(k)}$ from each source i and for all k and then order them once again (without regard to source) to form:

$$C_n^{(1)} < C_n^{(2)} < C_n^{(3)} < \dots < C_n^{(\hat{K}_n)}$$

(This ordering is invariant to Φ_n .) These ordered costs are the basis for calculating the Bertrand equilibrium in the next section. The source country i of any firm is irrelevant for calculating the Bertrand equilibrium. We nevertheless keep track of the source $I_{ni}^{(k)}$ for each firm in order to calculate who sells where.

5.4 Simulating Sales

We can focus on a particular destination n because the same routine applies to each market and our assumptions shut down any interactions among them. For a fixed K_n , all that is relevant for calculating the equilibrium in market n are the $C_n^{(k)}$ s and a value of σ . We start

³¹EKK's (2011) estimate is based on productivity and sales data from French exporters. Simonovska and Waugh (2011) found a similar value (4.12) using international price-comparisons data.

with $\sigma = 2.98$ from EKK (2011). In the continuum case, our values of θ and σ would imply that sales are distributed Pareto with parameter $\theta/(\sigma - 1) = 2.46$. We also try $\sigma = 5.64$ and $\sigma = 7.09$. In the continuum model, the implied parameters for the sales distribution would be 1.05 (with infinite variance) and 0.8 (with infinite mean and variance), respectively. For this last value, the continuum model would explode.

We solve for the Bertrand equilibrium prices $p_n^{K_n}(C_n^{(k)})$ in each country along with each firm's market share:³²

$$s_n^{(k)} = \frac{\left[p_n^{K_n}(C_n^{(k)}) \right]^{-(\sigma-1)}}{\sum_{k=1}^{K_n} \left[p_n^{K_n}(C_n^{(k)}) \right]^{-(\sigma-1)}}. \quad (23)$$

Our iterative numerical procedure exploits the condition for Bertrand-equilibrium markups given in Atkeson and Burstein (2008):

$$m_n^{(k)} = \frac{p_n^{K_n}(C_n^{(k)})}{C_n^{(k)}} = 1 + \frac{1}{(\sigma - 1)(1 - s_n^{(k)})}. \quad (24)$$

One issue of interest is how much these markups $m_n^{(k)}$ exceed the Dixit-Stiglitz markup $\bar{m} = \sigma/(\sigma - 1)$. Figure 4 shows the simulated distribution of markups among the $k = 1, \dots, 10$ largest firms across all markets (for $\sigma = 5.64$). Even at the top 95th percentile, the markup

³²As mentioned previously, we can simulate costs only up to an unknown constant $\Phi_n^{1/\theta} > 0$. Inspection of (23) and (24) shows that multiplying all costs $C_n^{(k)}$ by $\Phi_n^{1/\theta}$ leaves $s_n^{(k)}$ and $m_n^{(k)}$ unchanged, with


$$p_n^{K_n}(\Phi_n^{1/\theta} C_n^{(k)}) = \Phi_n^{1/\theta} p_n^{K_n}(C_n^{(k)}),$$

for all $k = 1, 2, \dots, K_n$. As a consequence, sales $X_n^{(k)} = s_n^{(k)} X_n$ and gross profits:

$$\Pi_n^{K_n}(C_n^{(k)}) = (1 - 1/m_n^{(k)}) X_n^{(k)}$$

are unchanged. Our solutions for what firms sell, their markups, and entry (discussed in Section 5.5) therefore are all invariant to Φ_n .

Figure 4. Markups of top 10 entrants (Bertrand competition)



of the largest firm in a market is around 1.29, as compared with a Dixit-Stiglitz markup of 1.22. For lower-ranked firms the markup is only negligibly higher than the Dixit-Stiglitz markup.³³

We also can calculate the sales of each firm where it has entered as:

$$X_{ni}^{(k)} = I_{ni}^{(k)} X_n^{(k)} = I_{ni}^{(k)} s_n^{(k)} X_n.$$

From these simulated sales, we identify the top 10, 100, 1000, and 10,000 French exporters across all foreign markets. We then calculate the contribution of each group both to total French exports and to French exports in each foreign market. Table 6 shows the results.

³³We also calculated markups under Cournot competition. At the 95th percentile the markup of the largest firm is 1.45 and markup of the second largest is 1.30.

Table 6. Share of largest French exporters

	Average			Standard deviation		
	(across 10 simulations)			(across 10 simulations)		
	$\sigma = 7.09$	$\sigma = 5.64$	$\sigma = 2.98$	$\sigma = 7.09$	$\sigma = 5.64$	$\sigma = 2.98$
Top 10	64.86	32.85	1.55	21.02	15.58	0.38
Top 100	82.99	51.27	6.13	10.12	11.34	0.39
Top 1,000	93.34	71.79	22.93	3.91	6.51	0.43
Top 10,000	98.73	92.05	65.86	0.76	1.86	0.24

We consider first the results with $\sigma = 2.98$, taken from EKK (2011). In that case, we do not come close to capturing the substantial contribution (shown in Table 2) of the largest French exporters. Conversely, increasing σ to 7.09 goes too far, with the top 100 firms accounting for more than 80 percent of French exports. The simulations with $\sigma = 5.64$ (and hence $\theta/(\sigma - 1) = 1.05$) match most closely the data in Table 2. The last three columns of Table 6 show that there is substantial variation in the contribution of the largest French exporters across simulation runs; however, in all cases, the middle value of σ delivers results that are closest to the data.³⁴

5.5 Entry Costs

From the results of the previous section, we can calculate the gross profits of the k th lowest cost firm in market n as:

$$\Pi_n^{\hat{K}_n}(C_n^{(k)}) = \left[1 - \frac{C_n^{(k)}}{p_n^{\hat{K}_n}(C_n^{(k)})} \right] \left(\frac{p_n^{\hat{K}_n}(C_n^{(k)})}{P_n^{\hat{K}_n}} \right)^{-(\sigma-1)} X_n$$

We calculate upper and lower bounds \bar{E}_n and \underline{E}_n on the entry costs as:

$$\bar{E}_n = \Pi_n^{\hat{K}_n}(C_n^{(\hat{K}_n)})$$

$$\underline{E}_n = \Pi_n^{\hat{K}_n+1}(C_n^{(\hat{K}_n+1)})$$

³⁴Our grid of parameter values is coarse. We leave it to future work to carry out a more formal estimation procedure.

In the simulations that follow, we set $\widehat{E}_n = \overline{E}_n$, although bounds are so tight that in the following figures the upper and lower would appear the same. The implied entry costs range from \$1000 to \$32,700.

We plot these values against mean sales across countries in Figures 5a. The relationship is tight, with an elasticity close to 1. In the continuum model, the elasticity would be exactly one. Figure 5b plots the entry costs against total absorption of manufactures. The relationship is noisier, with an elasticity of 0.29.³⁵

We are now fully equipped to look at the implications of various changes in the environment. In particular, knowing the entry costs, we can examine how such changes would affect the number of firms active in different markets.

6 Two Experiments

With values for E_n , we can combine our simulation of unit costs and simulation of sales from the previous section (which were conditional on the number of entrants K_n) to simulate a new global equilibrium in which entry into each market is endogenous. We continue to use the $\widehat{\pi}_{ni}$ (from our estimation of the gravity equation), $\theta = 4.87$, and $\sigma = 5.64$.

We do simulations of two types. One type examines the effect of changes in exogenous parameters, such as those governing trade costs d_{ni} and competitiveness S_i . We do these simulations using the same normalized cost draws $U^{(k)}$ as in the previous analysis. We hold the $U_i^{(k)}$'s fixed in conducting these counterfactuals in order to isolate the role of the parameters

³⁵Even though Cournot competition implied much larger markups for the biggest firms, the entry costs were very similar to the Bertrand case.


under consideration from changes introduced by a resampling of technology.

The second type focuses on the sensitivity of aggregate outcomes to the technology draws of the individual firms.³⁶ A major part of our analysis seeks to understand the sensitivity of the aggregate equilibrium to these draws. To assess their importance, we examine the implications of drawing different sets of $U_i^{(k)}$'s given the parameters of the model.

In either type of simulation, the number of firms K'_n entering a given market n is determined by the condition:

$$\Pi_n^{K'_n+1}(C_n^{(K'_n+1)}) < \widehat{E}_n \leq \Pi_n^{K'_n}(C_n^{(K'_n)}).$$


Figure 5a. Comparison of estimated \bar{E} and mean sales



$$\text{Regression: } \ln \widehat{E}_n = -4.1197 + 0.9377 \times \ln \bar{X}_n$$

³⁶In attempting to simulate the continuum model with a (computationally necessary) finite set of draws, sampling error would be a nuisance that researchers would want to minimize by sampling to the extent allowed by their hardware and patience. In our finite-firm model, sampling “error” is an integral part of the economic environment. Specifically, aggregate outcomes do indeed depend on the individual firms’ luck of the draw.

Figure 5b. Comparison of estimated \bar{E} and total absorption



$$\text{Regression: } \ln \hat{E}_n = -8.1375 + 0.2942 \times \ln X_n$$

6.1 Globalization

We begin by considering the consequence of a 10 percent decline in trade costs between countries. This experiment is similar to that carried out in EKK (2011) using a continuum model. In that experiment, some individual firms entered export markets and others were driven out of their home market. In our experiment here, with a finite-firm model, entire countries may enter new export markets. This counterfactual experiment illustrates the first type of simulation described previously, in which technology draws are held fixed as we consider the equilibrium implications of a change in the model parameters.

For each country pair $n \neq i$, we set the counterfactual trade cost to $d'_{ni} = d_{ni}/1.1$.³⁷

³⁷This change is equivalent to adding a constant $\theta \ln(1.1) = 0.464$ to each of the parameters m_n that appear

The decline in trade costs will alter the simulations by providing counterfactual values $\widehat{\pi}'_{ni}$ to replace $\widehat{\pi}_{ni}$ for constructing unit costs (for given $U_i^{(k)}$'s). These counterfactual values relate to the baseline values $\widehat{\pi}_{ni}$ according to:

$$\widehat{\pi}'_{ni} = \frac{\widehat{\pi}_{ni}(1.1)^\theta}{\widehat{\pi}_{nn} + \sum_{l \neq n}^N \widehat{\pi}_{nl}(1.1)^\theta}.$$

After computing equilibrium entry K'_n together with a Bertrand equilibrium in prices, we can evaluate the resulting counterfactual trade flows X'_{ni} .

World exports rise by 43 percent due to lower trade costs, in line with results in EKK (2011). Although nearly all of this increased trade occurs within pairs of countries that were already trading, 99.9984 percent, there are still perceptible changes along the extensive margin. Overall, 206 new trade flows emerge between country pairs for which one had not previously exported to the other. Among countries in the lowest-size quartile (measured by absorption of manufactures), the average number of export destinations increases by nearly 12 percent; however, sales in these new markets account for less than 0.2 percent of the increase in the value of their exports.

6.2 Granularity

Our second simulation evaluates the importance for aggregate outcomes of the luck of the technology draw at the level of individual firms. We look at variation in the manufacturing price level at the country level and at the role of the largest global firms. This experiment illustrates the second type of simulation in which parameters are held constant but technologies of all potentially active firms are redrawn repeatedly, with aggregate equilibrium outcomes in the gravity equation estimated in Section 5.1.

recalculated for each draw.

For each draw of technologies, generating a new set of costs $C_n^{(k)}$ in each market n , we calculate the equilibrium number of entrants K'_n and Bertrand equilibrium prices $p_n^{K'_n}(C_n^{(k)})$ to evaluate the log of the manufacturing price level:


$$\ln P'_n = \frac{-1}{\sigma - 1} \ln \left(\sum_{k=1}^{K'_n} \left[p_n^{K'_n}(C_n^{(k)}) \right]^{-(\sigma-1)} \right).$$

We present the $\ln P'_n$ for each simulation relative to its mean across all 200 simulations.

Figure 6 shows the results for two values of n , the United States and Denmark. Each point on the scatter represents the percentage deviation (calculated as the log difference) of the price level for the United States and Denmark for one of the 200 simulations. The variance for Denmark is 24.4 compared to 14.4 for the United States. Although it is notably smaller for the larger country, it remains substantial even in the largest destination. By plotting the results for both countries in the same figure, we capture the extent to which the model generates outcomes that move together across countries due to international trade. The results point to positive covariation, with a correlation of 0.48.

In some cases, the same firm is the top firm in both the United States and Denmark, in which case Figure 6 indicates its origin with a three-letter abbreviation. For example, along the northwest frontier of the scatterplot, the top firm in both the United States and Denmark often is a U.S. firm. It is not surprising that in these cases the United States has a low price level, whereas Denmark's is not so low. On the east-by-southeast frontier of the scatterplot, the largest firm in each country is often European, leading to a lower price level in Denmark than in the United States. In one case, the top firm is from Guatemala and in another case

Figure 6. Variation of P_n across simulations



Vietnam, demonstrating the possibility of extreme outcomes in this model. For the vast majority of the observations, the top firm in Denmark is different than the top firm in the United States. For the United States, a U.S. firm is at the top in 180 simulations, whereas Denmark’s top firm is Danish in 105 simulations.

7 Conclusion

In this chapter, we amend a standard heterogeneous-firm model of exporting by keeping the number of firms finite. Our quantification of the model suggests that it can fit well a number of features of the data.

Finiteness introduces both richness and complexity. To focus on its specific contribution, we keep our model simple in other dimensions. First, we introduce only one dimension of firm heterogeneity, underlying efficiency. Second, we do not incorporate endogenous entry costs, as in Arkolakis (2010).

As a consequence, the model makes some obviously false predictions. By stripping out additional dimensions of heterogeneity, firms from the same source will enter markets according to a strict hierarchy (i.e. a firm will always sell in an easy-to-enter market if it sells in a more-difficult market), and multiple firms from the same source selling in common destinations will always rank the same in terms of relative sales in each destination. By ignoring the endogeneity of entry costs, the model cannot account for systematic deviations from Zipf's law among small exporters.

EKK (2011) show how introducing heterogeneity to a firm's cost of entry and to its demand in each market, as well as adopting Arkolakis's formulation of endogenous entry costs, can break these rigid predictions. With these embellishments, the standard Melitz model can replicate multiple features of the data (although, of course, it fails to explain how zeros can arise in the trade data). Introducing additional sources of firm heterogeneity and endogenous entry costs into the model developed here should serve the same purpose in loosening this rigidity. In addition, we conjecture that introducing these features would improve the model's ability to predict zeros among very small source countries. These additions pose modeling challenges that we hope future research will overcome.

The domain of macroeconomics has been the study of aggregate relationships, whereas industrial organization focuses on the interaction of individual firms. Our exploratory analysis

in this chapter, in building a bridge (or perhaps only a tightrope) connecting their two domains, provides a new perspective on empirical relationships in international trade.

References

- Arkolakis, Costas (2010) “Market Penetration Costs and the New Consumers Margin in International Trade,” *Journal of Political Economy*, 118: 1151-1199.
- Arkolakis, Costas, Arnaud Costinot, and Andres Rodríguez-Clare (2012) “New Trade Model, Same Old Gains?,” *American Economic Review*, 102(1), 94-130.
- Arkolakis, Costas and Marc-Andreas Muendler (2010) “The Extensive Margin of Exporting Goods: A Firm-level Analysis,” *NBER Working Paper No. 16641*.
- Armenter, Roc and Miklos Koren (2008) “A Balls-and-Bins Model of Trade,” Philadelphia Federal Reserve Bank, unpublished working paper. .
- Armenter, Roc and Miklos Koren (2012) “Everything All the Time? Re-evaluating the Gains from Varieties,” Central European University, unpublished working paper.
- Atkeson, Andrew and Ariel Burstein (2008) “Pricing to Market, Trade Costs, and International Relative Prices,” *American Economic Review*, 98: 1998-2031.
- Baldwin, Richard and James Harrigan (2009) “Zeros, Quality and Space: Trade Theory and Trade Evidence,” *American Economic Journal: Microeconomics*, 3, 60-88
- Balistreri, Edward J., Russell H. Hillberry, and Thomas F. Rutherford (2011) “Structural Estimation and Solution of International Trade Models with Heterogeneous Firms,” *Journal of International Economics*, 83: 95-108.

- Bernard, Andrew B., Jonathan Eaton, J. Bradford Jensen, and Samuel Kortum, (2003) "Plants and Productivity in International Trade," *American Economic Review*, 93: 1268-1290.
- Bernard, Andrew B., J. Bradford Jensen, Stephen J. Redding, and Peter K. Schott (2007) "Firms in International Trade," *Journal of Economic Perspectives*, 21: 105-130.
- Canals, Claudia, Xavier Gabaix, Josep Vilarrubia, and David Weinstein (2007) "Trade Shocks, Trade Balances, and Idiosyncratic Shocks," Documentos de Trabajos No. 0721, Banco de España.
- Chaney, Thomas (2008) "Distorted Gravity: Heterogeneous Firms, Market Structure, and the Geography of International Trade," *American Economic Review*, 98: 1707-1721.
- Chaney, Thomas (2011) "The Gravity Equation in International Trade: an Explanation," unpublished working paper, University of Chicago.
- di Giovanni, Julian and Andrei Levchenko (2009) "International Trade and Aggregate Fluctuations in Granular Economies," forthcoming, *Journal of Political Economy*.
- Dornbusch, R., S. Fischer, and P.A. Samuelson (1977) "Comparative Advantage, Trade, and Payments in a Ricardian Model with a Continuum of Goods," *American Economic Review*, 67: 823-839.
- Eaton, Jonathan and Samuel Kortum (2002) "Technology, Geography, and Trade," *Econometrica* 70: 1741-1780.

- Eaton, Jonathan and Samuel Kortum (2010) *Technology in the Global Economy: A Framework for Quantitative Analysis*, unpublished manuscript, University of Chicago.
- Eaton, Jonathan, Samuel Kortum and Francis Kramarz (2011) “An Anatomy of International Trade: Evidence from French Firms,” *Econometrica*, 79: 1453-1498.
- Eaton, Jonathan and Akiko Tamura (1994) “Bilateralism and Regionalism in Japanese and U.S. Trade and Direct Foreign Investment Patterns,” *Journal of the Japanese and International Economies*, 8: 478-510.
- Feenstra, Robert C., Robert E. Lipsey, and Henry P. Bowen (1997) “World Trade Flows, 1970-1992, with Production and Tariff Data,” *NBER Working Paper No. 5910*.
- Gabaix, Xavier (1999) “Zipf’s Law for Cities: An Explanation,” *Quarterly Journal of Economics*, 94: 739-767.
- Gabaix, Xavier (2011) “The Granular Origins of Aggregate Fluctuations,” *Econometrica*, 79: 733-772.
- Gourieroux, C, A. Monfort, and A. Trognon (1984) “Pseudo Maximum Likelihood Methods: Theory,” *Econometrica*, 52: 681-700.
- Hausman, J. S., B. H. Hall, and Zvi Griliches, (1984) “Econometric Models Count Data with an Application to the Patents-R and D Relationship,” *Econometrica*, 52: 909-938.
- Head, Keith, Thierry Mayer, and John Ries (2010) “The Erosion of Colonial Trade Linkages After Independence,” *Journal of International Economics*, 81: 1-14.

- Helpman, Elhanan, Marc J. Melitz, and Yona Rubinstein (2008) “Estimating Trade Flows: Trading Partners and Trading Volumes,” *Quarterly Journal of Economics*, 123: 441-487.
- Luttmer, Erzo (2011) “On the Mechanics of Firm Growth,” *Review of Economic Studies*, 78: 1042-1068.
- Mariscal, Asier (2010) *Global Ownership Patterns*. Ph.D. dissertation, University of Chicago.
- Martin, Will and Cong S. Pham (2008) “Estimating the Gravity Model When Zero Trade Flows are Frequent,” unpublished working paper, World Bank.
- Melitz, Marc J. (2003) “The Impact of Trade on Intra-Industry Reallocations and Aggregate Industry Productivity,” *Econometrica*, 71: 1695-1725.
- Melitz, Marc J. and Gianmarco Ottaviano (2008) “Market Size, Trade, and Productivity,” *Review of Economic Studies*, 75: 295-316.
- Pedersen, Niels K. (2009) *Essays in International Trade*. Ph.D. dissertation, Northwestern University, Evanston, Illinois.
- Redding, Stephen J. (2011) “Theories of Heterogeneous Firms and Trade,” *Annual Review of Economics*, 3: 77-105.
- Santos Silva, J.M.C. and Silvana Tenreyro (2006) “The Log of Gravity,” *Review of Economics and Statistics*, 88: 641-658.
- Simon, H.A., and C.P. Bonini (1958) “The Size Distribution of Business Firms,” *American Economic Review*, 98: 607-617.

Simonovska, Ina and Michael E. Waugh (2011) "The Elasticity of Trade: Estimates and Evidence," unpublished, New York University.