NBER WORKING PAPER SERIES

FIRM HETEROGENEITY AND THE STRUCTURE OF U.S. MULTINATIONAL ACTIVITY: AN EMPIRICAL ANALYSIS

Stephen Yeaple

Working Paper 14072
http://www.nber.org/papers/w14072

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
June 2008

The statistical analysis of firm-level data on U.S. multinational corporations reported in this study was conducted at the International Investment Division, U.S. Bureau of Economic Analysis, under arrangements that maintained legal confidentiality requirements. Views expressed are those of the author and do not necessarily reflect those of the Bureau of Economic Analysis. I thank Peter Egger, Richard Kneller, Bill Zeile, and Jonathan Eaton for their comments. The views expressed herein are those of the author(s) and do not necessarily reflect the views of the National Bureau of Economic Research.

Firm Heterogeneity and the Structure of U.S. Multinational Activity: An Empirical Analysis
Stephen Yeaple
NBER Working Paper No. 14072
June 2008
JEL No. F1,F23

ABSTRACT

We use firm-level data for U.S. multinational enterprises (MNE) and the model of firm heterogeneity first presented in Helpman, Melitz, and Yeaple (2004) to make four empirical contributions. First, we show that the most productive U.S. firms invest in a larger number of foreign countries and sell more in each country in which they operate. Second, we assess the importance of firm heterogeneity in the structure of MNE activity. Third, we use the model to identify the mechanisms through which country characteristics affect the structure of MNE activity. Finally, we provide a systematic assessment of the model's shortcomings in order to inform the development of new theory.

Stephen Yeaple
Department of Economics
The Pennsylvania State University
610 Kern Building
University Park, PA 16802-3306
and NBER
stephen.yeaple@colorado.edu

1 Introduction

A tiny minority of firms engage in international trade, and a still smaller fraction of firms own production facilities in more than one country. These internationally engaged firms are systematically different than their domestically oriented peers. Firms that export are larger and more productive than firms that do not (Bernard and Jensen, 1999), while firms that open foreign affiliates are still larger and more productive than firms that only export (Tomiura, 2007). A well-developed body of theory explains these phenomena as the sorting of heterogeneous firms into modes of foreign market access.[1]

In this paper, we investigate how well a model of firm heterogeneity adapted from Helpman, Melitz, and Yeaple (2004) can explain the cross-country structure of U.S. multinational activity. The model is built on two key assumptions. First, firms face a trade-off in serving foreign markets: By opening a local affiliate, firms avoid per unit transport costs associated with trade but must instead incur fixed costs associated with managing a foreign affiliate. Second, firms differ in their productivity. These assumptions imply that for each country there is a productivity cutoff, which is determined by the country's characteristics, such that only those firms whose productivity exceeds this cutoff will open an affiliate in that country. Hence, the model predicts a "pecking order" such that the most productive firms should open an affiliate in even the least attractive countries, while progressively less productive firms enter progressively more attractive countries.

In the model, country characteristics affect the aggregate volume of multinational activity, measured as the sales of affiliates to host customers, through two channels. First, country characteristics determine the productivity cutoff, and so affect the productivity composition of the firms that invest there. The key feature of the model is that a change in a country characteristic that encourages a greater number of foreign firms to open a local affiliate must be inducing progressively less productive firms to enter. Second, country characteristics determine the optimal level of sales, holding fixed the set of firms that own an affiliate there.

We use the model and firm-level data collected by the Bureau of Economic Analysis for all U.S. multinational enterprises for the year 1994 to make four empirical contributions. First, we show that more productive U.S. firms own affiliates in a larger number of countries and these affiliates generate greater revenue on sales in their host countries. Previous studies, such as Girma, Kneller, and Pisu (2005), Head and Ries (2003), and Tomiura (2007) show only that firms that become multinational are systematically different than firms that export. Our analyses demonstrate that this sorting extends to the scale and scope of multinational enterprises: more productive firms own affiliates in a larger set of countries and their affiliates are larger than those of less productive firms. This sorting has quantitatively important implications for the aggregate structure of U.S. multinational activity.

The second contribution of our analysis is to assess how important firm heterogene-

[1] See, for instance, Melitz (2003), Bernard et al (2003), and Helpman, Melitz, and Yeaple (2004).

ity is in determining the structure of U.S. multinational activity. We show that as a country becomes more attractive to U.S. multinationals, it attracts progressively smaller and less productive firms. For instance, our estimates suggest that a 10% increase in a country's GDP per capita leads to a 7.6% increase in the number of U.S. firms that enter that country, but because new entrants are less productive than old entrants, the average productivity of all entrants falls by 2.0%. Thus, the contribution of the extensive margin, adjusted for the productivity composition of entrants, is 5.6%. Although there are several papers exploring the importance of firm heterogeneity models in the structure of international trade (e.g. Eaton, Kortum, and Kramarz, 2005), little has been done to investigate whether this body of theory improves our understanding of the aggregate structure of multinational activity.

Our third contribution is to use the structure of the model to disentangle the mechanisms through which individual country characteristics affect the structure of U.S. multinational activity. For example, while it has long been known that horizontal FDI is primarily attracted to developed countries, previous analyses do not shed light as to exactly why this is the case. We show that multinational activity is increasing in a host country's GDP per capita because individual entrants face greater effective demand in richer countries not because these countries have relatively lower entry costs. The analysis generates similarly surprising conclusions concerning the mechanisms through which physical distance and a shared language affect the structure of U.S. multinational activity.[2]

The fourth contribution of this paper is to assess the manner in which the model fails. We document systematic deviations from the pecking order. Although multinational activity is highly concentrated in the most productive firms, the model predicts that an even greater concentration of affiliate sales in the largest firms than is actually observed. In particular, larger firms underinvest in the least attractive countries. These observations should prove useful for the future development of models of firm heterogeneity.[3]

The remainder of this paper is divided into five sections. In section 2, we use a version of Helpman, Melitz, and Yeaple (2004) to derive predictions over the investment behavior of individual firms and to specify a structural econometric model of aggregate multinational activity. In section 3, we describe the firm-level data and the set of firm and country characteristics used to estimate this model. In section 4, we present the main results of our empirical analyses. The results confirm that it is important to account for firm heterogeneity in order to understand the structure of aggregate multinational affiliate sales, and they illuminate the channels through which country characteristics influence

[2]In this respect, our analysis is similar in flavor to Helpman, Melitz, and Rubinstein (2008) who use a model of firm heterogeneity and trade to disentangle the effect of distance on trade patterns. There is also some similarity to the work of Head and Ries (2008). They propose a mechanism that probabilistically assigns firms to countries. They too use their structural model to interpret the data.

[3]In this sense, our paper is similar to Baldwin and Harrigan (2007) who show that the Melitz (2003) model of firm heterogeneity fits certain facts well while systematically failing along other dimensions. Unlike Baldwin and Harrigan (2007) we do not pursue specific adjustments to the model to address these short-comings.

multinational activity. In section 5, we calculate the structure of multinational affiliate sales that would be observed if the pecking order were strictly observed and compare this counterfactual measure to the actual structure in order to demonstrate precisely how the model falls short. The final section concludes and presents suggestions for future research.

2 The Analytical Framework

We use a framework based on Helpman, Melitz and Yeaple (2004) to organize an econometric analysis of the structure of U.S. multinational activity across a range of host countries.[4] We first specify the model and generate firm level predictions. We then aggregate over individual firms to form country-wide predictions. Finally, we develop a series of equilibrium conditions that can be taken to the data.

2.1 The Model

The preferences of the representative consumer are the same everywhere and are given by

$$U = \frac{\beta}{\alpha} \ln \left(\int_{\omega \in \Omega} x(\omega)^{\alpha} d\omega \right) + (1-\beta) \ln Y$$

where $\alpha = (\sigma - 1)/\sigma$, $\sigma > 1$ is the elasticity of substitution across differentiated goods, and Y is a freely-traded homogenous good that is produced in every country. These preferences imply the following demand curve in country j,

$$x_j(\omega) = \beta E_j \left(P_j\right)^{\sigma-1} p_j(\omega)^{-\sigma}, \tag{1}$$

where E_j is gross national expenditure in country j, P_j is the price index in country j, and $p_j(\omega)$ is the price of variety ω in country j.

There are J countries indexed by j. In country j the mass of firms is N_j. Each firm is capable of producing a single variety of the differentiated good using a single input called labor. The price of labor in country j is w_j. The wage is determined in the homogenous-good industry Y. Firms are heterogeneous in terms of their productivity φ.[5] The empirical distribution of φ in each country G is assumed to be Pareto, i.e.

$$G(\varphi) = 1 - \left(\frac{b}{\varphi} \right)^{k},$$

[4] The model presented in this section differs from that presented in Helpman, Melitz, and Yeaple (2004) in two respects. First, the model is not closed via a free entry condition. Second, because we do not observe firm-level exports in our dataset, we do not include a fixed cost of exporting.

[5] Variation in productivity across firms can be thought of more generally as variation in firm characteristics that lead to a higher value of output per unit input. As pointed out by Melitz (2003), variation in productivity across firms is isomorphic to variation in quality across firms when preferences are CES.

where $k > \sigma - 1$.

Each firm sells its variety in each country. In serving foreign markets, a firm faces a trade-off in the way that it organizes its production. If a firm from country h chooses to export its variety from country h to a foreign country j, it faces wage w_h and incurs iceberg transport cost $\tau_{hj} > 1$. Alternatively, the firm may avoid this transport cost by opening an affiliate in country j. Setting up an affiliate in country j requires the firm to incur the fixed cost F_j. The firm must also pay the local labor cost w_j. We assume that technology transfer is perfect so that the firm's productivity is the same in every country. To focus on the pure proximity-concentration trade-off, we abstract away from the possibility that a firm engages in export platform FDI and vertical FDI by assuming that transport costs are relatively high compared to wage differences between countries so that $w_h \tau_{hj} > w_j$ for all h and j.[6]

In the remainder of this section, we derive equilibrium conditions to organize the empirical analyses, but we do not solve for the full general equilibrium of the model. In particular, the equilibrium price index will not be derived, because it is a complex function of a country's geography. As such, we will need to take a reduced form approach in relating the certain country variables, such as the price index, to country characteristics when implementing our empirical work.

2.2 Firm Mode Choice

We begin our analysis by considering the foreign direct investment decisions of individual firms with headquarters located in country h. Let the marginal cost of serving country j for a firm with productivity φ from country h be $c_{hj}(\varphi)$. As is well-known, a firm facing demand curve (1) will optimally charge a price in country j of $p_j(\varphi) = c_{hj}(\varphi)/\alpha$, generating revenues of

$$R_{hj}(\varphi) = A_j c_{hj}(\varphi)^{1-\sigma},$$

where $A_j \equiv \alpha^{\sigma-1}\beta E_j P_j^{\sigma-1}$ is the mark-up adjusted demand level in country j and P_j is the price index in country j. A firm from country h with productivity φ that sells its product in country j will face marginal costs of

$$c_{hj}(\varphi) = \begin{cases} w_h \tau_{hj}/\varphi & \text{if it exports to country } j, \\ w_j/\varphi & \text{if it produces in country } j. \end{cases}$$

and so will generate revenues of

$$R_{hj}(\varphi) = \begin{cases} \left[A_j (w_h \tau_{hj})^{1-\sigma}\right] \varphi^{\sigma-1} & \text{if it exports to country } j, \\ \left[A_j (w_j)^{1-\sigma}\right] \varphi^{\sigma-1} & \text{if it produces in country } j. \end{cases} \tag{2}$$

[6]For an analysis of a very different geography were affiliate exports are prominently featured in a model of firm heterogeneity, see Yeaple (2008).

Because $w_h\tau_{hj} > w_j$ and $\sigma > 1$ the revenue generated by an affiliate will exceed the revenue generated via exports. Finally, note that this firm's sales in its home market are

$$R_{hh}(\varphi) = \left[A_h(w_h)^{1-\sigma}\right]\varphi^{\sigma-1}.$$

A firm's sales revenue in each market (including its home market) is proportional to its productivity index $\varphi^{\sigma-1}$. Because each firm serves its home country h from a local plant, home sales are a conceptually valid measure of a firm's productivity. These observations will be important in the empirical work that follows.

A firm of productivity φ generates gross profits on sales in country j that are proportional to its revenues in that market: $\pi(c_{hj}(\varphi)) = R_{hj}(\varphi)/\sigma$. Using equation (2) it follows that the net profit of a firm of productivity φ that exports from country h to country j is

$$\pi_X(\varphi) = \left[\frac{A_j(w_h\tau_{hj})^{1-\sigma}}{\sigma}\right]\varphi^{\sigma-1}, \tag{3}$$

while the net profit of the same firm that opens an affiliate in country j is

$$\pi_M(\varphi) = \left[\frac{A_j w_j^{1-\sigma}}{\sigma}\right]\varphi^{\sigma-1} - F_j. \tag{4}$$

Both $\pi_X(\varphi)$ and $\pi_M(\varphi)$ are linear and increasing in a firm's productivity index $\varphi^{\sigma-1}$, but $\pi_M(\varphi)$ increases faster than $\pi_X(\varphi)$ because we have assumed that $w_h\tau_{hj} > w_j$. Because $\pi_M(0) < \pi_X(0)$, it follows that there exists a country-pair specific cutoff productivity index $\widehat{\varphi}_{hj}^{\sigma-1}$ such that all firms from source country h with productivity index $\varphi^{\sigma-1} \geq \widehat{\varphi}_{hj}^{\sigma-1}$ engage in FDI in country j and all firms with productivity $\varphi^{\sigma-1} < \widehat{\varphi}_{hj}^{\sigma-1}$ export instead. Setting equation (3) equal to (4) and solving for the cutoff productivity index yields

$$\widehat{\varphi}_{hj}^{\sigma-1} = \sigma\left(\frac{1}{A_j w_j^{1-\sigma}}\right)\left(\frac{F_j}{1-(w_h\tau_{hj}/w_j)^{1-\sigma}}\right). \tag{5}$$

Equation (5) shows that the determinants of the cutoff productivity index $\widehat{\varphi}_{hj}^{\sigma-1}$ can be decomposed into the two components shown in parentheses. Shown in the first parentheses is the inverse of a measure of the country-specific scale of operations $A_j w_j^{1-\sigma}$. Everything else equal, an increase in the mark-up adjusted demand level A_j or a reduction in the local cost of production w_j will lower the cutoff productivity index, inducing relatively less productive firms to invest. Shown in the second parentheses is a measure of the fixed cost of opening an affiliate in country j relative to a measure of the marginal-cost savings made possible by opening a local affiliate in country j. An increase in the fixed cost F_j or a reduction in the marginal-cost savings $[1-(w_h\tau_{hj}/w_j)^{1-\sigma}]$ increases the cutoff productivity index.

Earlier, we showed that firms with larger home market sales must have a higher productivity index $\varphi^{\sigma-1}$ and so earn larger revenues in each country in which they own an affiliate. Now note that firms with large home market sales also invest in a larger number of countries because their productivity index will exceed the cutoff productivity index for a larger number of countries. Let Θ_{hj} be the set of firms that invest in country j. Suppose that the number of firms that originate in h and invest in country j is increasing in j so that country J is the most attractive. Then, it must be the case that

$$\Theta_{h1} \subset \Theta_{h2} \subset \Theta_{h3} \subset ... \subset \Theta_{hJ}.$$

This is known as the "pecking order" because every firm from h that owns an affiliate in country j also owns an affiliate in every country $k > j$.

2.3 The Aggregate Structure of Multinational Activity

In this section, we compute the value of aggregate local sales in country j of the affiliates owned by firms from country h. In the empirical analysis below, country h will be the United States. These sales can be written

$$S_{hj}^{j} = N_h \int_0^\infty R_{hj}(\varphi) dG(\varphi).$$

The superscript on S_{hj}^{j} indicates the country of sales while the subscripts indicate the sales are determined by variation in the number of firms from country h that own affiliates in country j. Using equations (2) and (5), the aggregate sales of affiliates from country h located in country j can be written

$$S_{hj}^{j} = \left(A_j w_j^{1-\sigma}\right) \left(N_h \int_{\widehat{\varphi}_{hj}}^\infty \varphi^{\sigma-1} dG(\varphi)\right). \tag{6}$$

The first term in parentheses is the country-specific scale of operations of foreign affiliates, holding fixed the identity of the firms investing in country j, and the second term in parentheses is a measure of the number and the productivity composition of firms from country h that open an affiliate in country j. Equation (6) illustrates an important feature of this model of firm heterogeneity. Holding fixed the sale of operations $(A_j w_j^{1-\sigma})$, an increase in multinational activity is driven by expansion of the number of firms investing in the country, or a decrease in $\widehat{\varphi}_{hj}$. A decrease in $\widehat{\varphi}_{hj}$ implies, however, that the additional firms being attracted are less productive than the incumbent firms. Hence, as the number of firms that invests in the country expands, the average productivity of the entrants falls.

To analyze the role of sorting in the structure of aggregate multinational activity, we will use the following decomposition of aggregate affiliate sales of firms from country h in country j:

$$S_{hj}^{j} = \left(\frac{S_{hj}^{j}}{S_{hj}^{h}}\right) \times N_{hj} \times \left(\frac{S_{hj}^{h}}{N_{hj}}\right), \tag{7}$$

where S^h_{hj} is the sales of the set of firms from country h that own an affiliate in country j in their home country h, and N_{hj} is the size of that set in terms of the number of firms from country h that invest in country j. Note that because S^h_{hj} and S^j_{hj} are calculated using using the same set of firms, the first term in parentheses is a measure of how countries differ in the *scale* of firm operations:

$$\frac{S^j_{hj}}{S^h_{hj}} = \frac{A_j w_j^{1-\sigma}}{A_h w_h^{1-\sigma}}. \tag{8}$$

The second term in (7) is the *number* of firms from country h that invest in country j and is equal to

$$\begin{aligned} N_{hj} &= N_h \left(1 - G(\widehat{\varphi}_{hj})\right) \\ &= N_h b^k \widehat{\varphi}_{hj}^{-k}. \end{aligned} \tag{9}$$

The second line of (9) is obtained by substituting the definition of $G(\varphi) = 1 - b^k \varphi^{-k}$. The third term in (7) measures the *average productivity* of the set of firms from country h that open an affiliate in country j, as measured by their parent firms' sales in the home country h, and it is equal to

$$\begin{aligned} \frac{S^h_{hj}}{N_{hj}} &= \frac{A_h w_h^{1-\sigma}}{1 - G(\widehat{\varphi}_{hj})} \int_{\widehat{\varphi}_{hj}}^{\infty} \varphi^{\sigma-1} dG(\varphi) \\ &= \frac{A_h w_h^{1-\sigma} k}{k - \sigma + 1} \widehat{\varphi}_{hj}^{\sigma-1} \end{aligned} \tag{10}$$

Note that the home country scale $(A_h w_h^{1-\sigma})$ does not vary across destination countries j, so that only the average productivity index varies across destination countries.

Combining equations (7) through (10) with equation (5) allows us to link the relevant country characteristics identified by theory (those associated with $A_j w_j^{1-\sigma}$ and $F_j(1 - (w_h \tau_{hj}/w_j)^{1-\sigma})^{-1}$) to the aggregate volume of multinational activity and to the individual components (*scale*, *number*, and *average productivity*). A key feature of the model can be obtained by comparing equations (9) and (10). The number of affiliates in country j is decreasing in the cutoff productivity $\widehat{\varphi}_{hj}$ while the average productivity is increasing in the cutoff productivity. Combining these two equations,we find that

$$\log \frac{S^h_{hj}}{N_{hj}} = \log \left(\frac{k b^{-k}}{N_h (k - \sigma + 1)} \right) - \left(\frac{k}{\sigma - 1} \right)^{-1} \log N_{hj}. \tag{11}$$

This equation shows that the relationship between the number of entrants and the average productivity of entrants is governed by $\frac{k}{\sigma-1}$, which is a measure of the degree of size dispersion across firms. Everything else equal, the greater the size dispersion across firms (the smaller is $\frac{k}{\sigma-1}$), the greater the degree of heterogeneity in the productivity index, and the stronger is the negative relationship between the *number* and *average productivity* of entrants across countries.

2.4 Country Characteristics and Mechanisms

In this section, we show how components of the decomposition presented in (7) can be reorganized to isolate meaningful sets of model parameters. These sets of parameters are the scale of foreign affiliates, holding fixed the composition of the entrants, or $(A_j w_j^{1-\sigma})/(A_h w_h^{1-\sigma})$, and the ratio of the fixed cost of FDI to the marginal-cost benefit of FDI, or $F_j / [1 - (w_h \tau_{hj}/w_j)^{1-\sigma}]$. We do not attempt to find an empirical analog for every variable in the model. For instance, the mark-up adjusted demand level A_j is a function of the price index P_j, which is an endogenous and complicated function of a country's location. Further, differences in marginal cost and fixed costs across countries are difficult to specify. Instead, we take a reduced-form approach, relating components of the model to observable country characteristics. We then use the model to identify which of two mechanisms country characteristics affect the structure of multinational sales.

As noted above, there are two mechanisms through which variation in country characteristics affect the structure of multinational activity across countries. The first is directly related to *scale* as illustrated in equation (8). We assume that the country-specific scale, which is directly measurable using data on the sales of affiliates in country j to the sales of their parent firms in country h, is related to a vector of observable country characteristics Z_j through the following equation:

$$\log\left(\frac{S_{hj}^j}{S_{hj}^h}\right) = \frac{A_j w_j^{1-\sigma}}{A_h w_h^{1-\sigma}} \equiv \Phi \log Z_j + \varepsilon_j. \tag{12}$$

The second mechanism through which country characteristics affect the cross-country structure of multinational activity is through their effect on the relative magnitude of fixed costs F_j to the measure of unit cost savings of multinational activity $[1 - (w_h \tau_{hj}/w_j)^{1-\sigma}]$. We assume that this ratio is related to the same set of country-characteristics Z_j through the following equation:

$$\log\left(\frac{F_j}{1 - (w_h \tau_{hj}/w_j)^{1-\sigma}}\right) = \Gamma \log Z_j + \eta_j. \tag{13}$$

Using equations (6), (9), (5), and (13), it is straightforward to show that

$$S_{hj}^j / N_{hj} = \frac{k\sigma}{k - \sigma + 1}\left(\frac{F_j}{1 - (w_h \tau_{hj}/w_j)^{1-\sigma}}\right),$$

which motivates the following econometric specification:

$$\log(S_{hj}^j / N_{hj}) = \log\left(\frac{k}{k - \sigma + 1}\right) + \Gamma \log Z_j + \eta_j. \tag{14}$$

By estimating equations (12) and (14), we can identify the mechanisms through which individual country characteristics influence the aggregate structure of multinational activity. The vector of coefficients Φ summarize the effect of country characteristics on affiliate scale, while the coefficients Γ summarize the effect of country characteristics on the relative size of concentration costs versus proximity benefits of FDI.

3 Data

This study relies on the confidential enterprise-level data collected by the Bureau of Economic Analysis (BEA). The BEA conducts annual surveys of U.S. Direct Investment Abroad where U.S. direct investment is defined as the direct or indirect ownership or control by a single U.S. legal entity of at least 10 percent of the voting securities of an incorporated foreign business enterprise or an equivalent interest in an unicorporated business enterprise. A U.S. multinational entity (MNE) is the combination of a single U.S. legal entity that has made the direct investment, called the U.S. parent, and at least one foreign business enterprise, called the foreign affiliate. The International Investment and Trade in Services Survey Act requires that all firms larger than a certain size file detailed financial and operating items for the parent firm and each affiliate.

The dataset includes every manufacturing affiliate that is majority owned by a U.S. parent company whose main-line-of-business is a manufacturing industry for the benchmark survey year 1994. For each affiliate, we observe the country in which the affiliate is located, the main-line-of-business of the affiliate and the volume of sales to customers based in that country.[7] For every parent firm, we observe the firm's main-line-of-business, its sales to customers located in the United States, and which affiliates it owned. To calculate each parent firm's total factor productivity, we also extracted employment and capital stock data from the BEA dataset. When a firm owned more than one manufacturing affiliate in a given foreign country, we aggregated the volume of sales across affiliates to form a single firm-country observation. When we disaggregate the sample by industry, we use the industry of the parent firm.

To construct our measure of local affiliate sales in country j, S_j^j (note the source country subscript has been dropped because the United States is the sole source country the superscript indicates that these are the sales of the set of firms investing in country j in country j), we aggregate the sales of all U.S. affiliates in country j. Our measure of the number of entrants N_j corresponds to the total number of U.S. manufacturing firms that own manufacturing affiliates in country j. Our measure of the productivity composition of U.S. firms that invest in country j, S_j^{US}, is the sum of all the U.S. sales of all U.S. parents that own a manufacturing affiliate in country j. The dependent variables used in the analysis are then the logarithms of the appropriate combinations of these variables, i.e. S_j^j, N_j, S_j^{US}/N_j, and S_j^j/S_j^{US}.

Following the recent literature we interpret the model as applying to aggregate manufacturing activity rather than as at the industry level.[8] Each firm is then interpreted as producing a unique good, and the fixed, transport, and production costs are assumed

[7] Because our model is in the proximity-concentration tradition of Brainard (1997), we focus on the local sales of the affiliates. These sales account for roughly sixty percent of the sales of majority-owned foreign affiliates of U.S. multinational enterprises.

[8] We have in mind, Eaton, Kortum, and Kramarz (2005), Baldwin and Harrigan (2007), Head and Ries (2008), and Helpman, Melitz, and Rubinstein (2008).

to be common across manufacturing. As a robustness check, however, we also consider a sample in which firms are disaggregated by industry. Hence, these measures were computed at two different levels of industrial aggregation: at the level of total manufacturing and at the level of the three-digit BEA industry. As a robustness check, we also measure the country-specific cutoff productivity index $\widehat{\varphi}_j^{\sigma-1}$ directly using the local affiliate sales of the smallest U.S. affiliates in country j.

We now turn to the country characteristics that are included in the set of independent variables Z_j. We do not attempt to define proxies for every variable in the model. Proxies for variables such as country-level fixed costs tend to be highly correlated with many other country characteristics. Instead, we choose to use standard gravity variables, which have been shown to explain most of the country variation in multinational activity. The key innovation of our analysis is to show how these variables affect the various margins of foreign affiliate activity and to use the structure of the model to identify the mechanism through which these variables affect foreign affiliate activity.

The gravity variables included are GDP, which is the logarithm of a country's gross domestic product in 1994; $GDPPC$, which is the logarithm of a country's GDP per capita in 1994; $DIST$, which is the logarithm of the distance between Washington D.C. and the capitals of each country in the dataset; and $ENGLISH$, which is an indicator for whether English is a major language in the destination country. The variables GDP and $GDPPC$ are taken from the *Penn World Tables*. The variable $DIST$ is taken from Andy Rose's website.[9]

4 Results

In this section, we report the results of several related analyses. We begin with a descriptive analysis of the foreign investment behavior of individual U.S. firms to establish that a key prediction of the model, that more productive firms have larger multinational networks, is true in the data. We then present the main results of the paper: those associated with estimating the effect of country characteristics on manufacturing affiliate sales and the three components identified in the decomposition shown in equation (7). We also use the model to interpret the mechanisms through which country characteristics affect the structure of U.S. multinational activity. Finally, we conduct several robustness checks.

[9]The 49 countries included in the dataset are Argentina, Australia, Austria, Belgium, Canada, Chile, China, Colombia, Costa Rica, Czech Republic, Denmark, Dominican Republic, Ecuador, Egypt, Finland, France, Germany, Greece, Guatemala, Hong Kong, Hungary, India, Indonesia, Ireland, Israel, Italy, Japan, Korea, Malaysia, Mexico, Netherlands, New Zealand, Norway, Panama, Peru, Philippines, Poland, Portugal, Singapore, South Africa, Spain, Sweden, Switzerland, Taiwan, Thailand, United Kingdom, Russia, Uruguay, and Venezuela.

4.1 The Investment Behavior of Individual Firms

In this section, we provide a simple descriptive analysis of the foreign production activity of individual U.S. multinational firms. We first explain the propensity of firms to invest in any given foreign market as a function of a parent firm's measured productivity. We use two measures of a parent firm's productivity: the logarithm of the value of the parent firm's sales in the United States and the logarithm of the parent firm's TFP.[10] We estimate a simple linear probability model that allows us to control for host country and industry fixed effects. Second, we regress the logarithm of individual firms' local sales in foreign countries on the logarithm of their parent firms' U.S. sales and their TFP. Descriptive statistics can be found in appendix table 1.

The results are shown in Table 1. The first two columns correspond to the coefficient estimates obtained from the linear probability model. The dependent variable is equal to one if a firm owns an affiliate in that country and is equal to zero otherwise. The first column corresponds to the specification using the logarithm of parent firm size as measure of firm efficiency, while the second corresponds to the specification using the logarithm of parent firm TFP. In both specifications, the measure of firm efficiency is positive and statistically significant, indicating that more efficient firms are more likely to own an affiliate in any given host country. The results are very similar if industry fixed effects are not controlled.

The third and fourth columns of Table 1 show the results of regressing the logarithm of a foreign affiliate's sales on the logarithm of its parent firm's productivity and country and industry fixed effects. The third column corresponds to the specification in which firm productivity is measured using parent firm sales in the United States, while the fourth column corresponds to the specification in which firm productivity is measured using parent firm TFP. In both specifications, affiliate size is increasing in the parent firm's productivity. The results are very similar if industry fixed effects are not controlled.[11]

Earlier studies (e.g. Tomiura 2007) demonstrate that larger, more productive firms are more likely to become multinationals. The results shown in Table 1 extend the body of knowledge, showing that the scope (number of affiliates) and scale (size of its affiliates) of a firm's multinational activities, are also increasing in the firm's productivity.

4.2 Firm Heterogeneity and the Aggregate Structure of Multinational Activity

In this section, we report the effects of regressing the logarithm of aggregate affiliate sales and the logarithm of the components of these sales as expressed in equation (7) on the

[10] Given that the model features constant mark-ups, parent sales in the home market are the preferred measure of parent firm productivity. We include TFP for completeness.

[11] The fact that the coefficient on the logarithm of parent firm sales in the United States is less than one is inconsistent with the predictions of the model. We discuss this fact in greater detail later in the paper.

logarithm of the gravity variables. We first show that country characteristics that predict a larger number of domestic entrants also predict a lower average productivity of the parent firms. We then use these coefficient estimates to interpret the mechanisms through which country characteristics influence the cross-country structure of U.S. multinational activity. Finally, we provide an initial assessment of some of the strict predictions of the model. The results reported in this section are generated using data aggregated to the level of U.S. manufacturing. In the next section, we compare the results obtained when cross-industry heterogeneity is controlled.

We begin by reporting the estimate of $(\sigma - 1)/k$ obtained by estimating equation (11), which relates the logarithm of the average parent firm sales in the United States of the set of firms that own an affiliate in a given country to the logarithm of the number of U.S. firms that own an affiliate in that country. This coefficient summarizes the importance of firm heterogeneity in the structure of multinational activity. The larger the coefficient the greater the degree of firm heterogeneity implied in the sample, and hence, the more responsive is the productivity composition is to changes in the number of entrants.

The coefficient estimate is found to be -0.258 with a standard error of 0.04 and an R-squared of 0.36. This implies a value of $k/(\sigma - 1)$ approximately equal to 4. This result provides direct evidence of the type of sorting predicted by the model: A 10% increase in the number of entrants into a given country is associated with a 2.5% decrease in the average productivity index of entrants. In their study of the exporting behavior of French firms, Eaton, Kortum, and Kramarz (2005) estimate a value of $k/(\sigma - 1)$ or approximately 1.5, which implies a much stronger composition effect.[12]

We now turn to the results of regressing the various measures of multinational activity on country characteristics. The main results are shown in Table 2.[13][14] Column 1 reports the coefficient estimates obtained by regressing aggregate multinational sales onto the set of gravity variables. Consistent with other studies, we find that the four variables account for much of the variation in affiliate sales across countries as indicated by the R-squared of 0.738. The local affiliate sales of U.S. multinationals are increasing in a country's GDP, in its per capita GDP, and are higher in English-speaking countries. Affiliate sales are also strongly decreasing in distance: a 10% increase in the distance between the U.S. and a given host country is associated with a 7.6% decrease in these sales.

Column 2 reports the coefficients obtained by regressing the logarithm of the *number* of U.S. firms that own an affiliate in a given country on the set of country characteristics. The coefficient estimates reported in column 2 have the same signs as those in column 1 and are smaller in absolute value. A 10% percent increase in GDP is associated with a 5.4% increase in the number of U.S. firms present in that country and a 9.3% increase in the sales of those affiliates (see column 1), implying that the average affiliate size increases

[12] Using different data, Chaney (2007) reports a value closer to 2.

[13] The descriptive statistics for this sample can be found in Appendix Table 2.

[14] The nature of ordinary least squares guarantees that the sum of the coefficient estimates of each component (columns 2-4) is equal to the coefficient estimate for in the total sales equation (column 1).

by approximately 4.0%. As indicated by the fact that the coefficients in column 2 are in general more than half the size of the coefficients in column 1, variation in the extensive margin (the number of entrants) accounts for more than half of the variation in affiliate sales. This is particularly true in the case of the variable $ENGLISH$, where the extensive margin accounts for all of this variation.

Column 3 reports the effect of regressing the logarithm of the *average productivity* of the parent firms (as measured by sales in the U.S.) that own an affiliate in a given country on the same set of country variables. Theory predicts that these coefficients should be opposite in sign to those in column 3 and that the absolute value of the ratio of coefficients in column 3 to those in column 2 should be the same across country characteristics. The first prediction of theory is consistent with the data. All four coefficient estimates in column 3 have the opposite sign as the corresponding coefficient estimate in column 2. For instance, a 10% increase in a country's GDP is associated with a 5.4% increase in the number of U.S. firms owning an affiliate there, but a 1.4% decrease in the average productivity of the entrants. Of the four country characteristics, only $DIST$ is not statistically significant.

According to the theory, the coefficient on each variable in the *average productivity* equation should be equal to $-(\sigma - 1)/k$ multiplied by coefficient on the same variable in the *number* equation (column 2). This is an implication of a single productivity cutoff for each country. Hence, the absolute value of all of the coefficient estimates in column 3 should be 0.26 of the absolute value of the coefficient estimates in column 2. This is true of the coefficients on GDP (0.26) but is off for the coefficient on $DIST$ (0.50) and very far off for $ENGLISH$. From these estimates, we conclude that the structure of U.S. aggregate multinational activity suggests sorting of the type present in the model, but the prediction of a single threshold determined by one set of country characteristics is not strictly consistent with the data.

Columns 4 report the coefficient estimates from the specifications relating the logarithm of scale to the set of country characteristics. Holding fixed the composition of entrants, U.S. affiliates are larger in countries with larger market size (GDP), that are richer ($GDPPC$), that are more proximate to the U.S. market ($DIST$), and where the population speaks English ($ENGLISH$). The negative effect of distance and the positive effect of a shared language are particularly interesting. The negative coefficient on distance suggests that the marginal cost of production facing affiliates in countries far from the United States is higher than in more proximate countries. There is nothing in the standard proximity-concentration model that gives rise to this result, suggesting a need to augment the model. One way to adjust the model to replicate this fact is to include intermediate input purchases by the affiliate from their host country. The positive coefficient on $ENGLISH$ suggests that either the products of U.S. multinationals are somehow more popular in English-speaking countries (larger market size) or that the marginal cost of production is lower in these countries.

In the last column are the results obtained by regressing the logarithm of average affiliate sales to local customers on the same gravity variables. As noted above, these results can be interpreted as describing the effect of these country characteristics on the relative magnitude to country-specific fixed costs to country-specific proximity unit cost benefits of FDI.

The positive coefficients on GDP and $GDPPC$ are probably best interpreted as evidence that fixed costs are rising in market size and level of development. The negative coefficient on $DIST$ must be interpreted as either fixed costs that are decreasing in distance from the parent firm or as proximity cost benefits that are increasing in distance. The latter interpretation is more natural given that most models of horizontal FDI should have this feature. The negative effect of distance on aggregate multinational activity thus appears to be due to the effect that higher marginal costs have on the scale of multinational operations (and indirectly also through the cutoff productivity level $\widehat{\varphi}^{\sigma-1}$) rather than through the effect of fixed costs.[15]

Using the model to interpret the effect of a shared language on the structure of aggregate activity, there appears to be no effect of language on the relative size of fixed costs to unit proximity advantages as indicated by the statistical insignificance of the coefficient on $ENGLISH$. Rather, the positive effect of a shared language on aggregate volumes of multinational activity appears to work through its effect on *scale* (both directly and through its effect on the cutoff productivity level).

4.3 Robustness

We now estimate the same specifications on a sample in which the unit of observation is the country-industry. The benefit of considering a sample disaggregated by industry is (1) we can control for unobserved heterogeneity across industries, and (2) we can compare firms that are more likely to be competing in similar product markets. The disadvantage is that the number of observations in which there is no observed entry becomes substantial: more than half the country-industry pairs contain no U.S. multinational activity. There is no immediate way to accommodate these zero observations as any country characteristic that affects the likelihood of entry must, by the construction of the model, also affect the identity of the entering firms. That is, there cannot exist sensible exclusion restrictions.[16] Another problem that arises is that some industry categories contain very few U.S. parent firms.

The coefficients obtained by estimating the same econometric model on the disaggregated sample are shown in Table 3.[17] The main qualitative difference in these estimates compared to those of the aggregate analysis is that the absolute value of the coefficients in

[15] Note that according to the model, the logarithm of the productivity cutoff index $\widehat{\varphi}^{\sigma-1}$ is equal to the logarithm of average affiliate size minus the logarithm of affiliate scale. Hence, country characteristics that affect scale have both direct and indirect effects on the structure of multinational activity.

[16] Note also that the average productivity of entrants when there are no entrants is not defined when there are no entrants, so techniques, such as Poisson estimation, are not appropriate in this context.

[17] For descriptive statistics see Appendix Table 3.

the number and average productivity regressions is smaller, reflecting a mixture of composition effects that have been controlled by including industry fixed effects and attenuation bias stemming from the large number of zeros.

Focusing on the average productivity specification, we see that each coefficient estimate continues to have the opposite sign as that in the number regression but the pattern of statistical significance is considerably different. Now, the coefficient on $DIST$ is marginally statistically significant while the coefficient on GDP is not.

We have taken the model seriously in measuring the average productivity of a firm by the parent firm's sales in the United States. As a robustness check, we consider the composition of the productivities directly using the TFP of the parent firm. The first two columns of Table 4 correspond to a specification in which the logarithm of the average TFP of the parent firms that own an affiliate in a country are regressed on the same set of gravity variables. The first column corresponds to the aggregate sample and the second column corresponds to the sample disaggregated by industry.

The coefficient estimates for the aggregate sample shown in column 1 have the expected sign pattern. The average productivity of parent firm entrants is decreasing in GDP and $GDPPC$. Both coefficients are statistically significant. The coefficients on $DIST$ and $ENGLISH$ have their expected signs but are not statistically significant. The coefficient estimates in the disaggregate sample shown in column 2 again have the expected signs. None of the coefficients are individually statistically significant at standard levels, but an F-test of the joint significance of the four country variables indicate that they are statistically significant at the five-percent level (F-statistic of 2.69). The lack of individual statistical significance could reflect substantial multicollinearity across the independent variables.

The model has the strong implication that there should exist a"pecking order" across countries: the largest most productive firms should invest in all countries while the least productive firms should be located in only the most productive locations. To get at this prediction, we follow Manova (2007) in focusing on the cutoff directly. Our measure of the productivity cutoff is the logarithm of the U.S. sales of the smallest parent firm that invests in a country-industry pair. Given the Pareto assumption, the coefficients reported in these columns ought to be identical to those presented in the column 3 of Tables 2 and 3, respectively.

The results for the aggregate sample are shown in column 3 of Table 4, while the results for the disaggregate sample are shown in column 4 of Table 4. As predicted by theory, the coefficient estimates in both columns are opposite in sign to those shown for the *number* specifications shown in Tables 2 and 3 respectively. The cutoff productivity level is lower in developed countries ($GDPPC$) countries with large economies (GDP) that are close to the United States ($DIST$) and where the main language is English ($ENGLISH$). Interestingly, the coefficient estimates are generally larger in absolute value in these specifications relative to the specifications in which the logarithm of average productivity is the dependent variable, and they tend to be much more precisely estimated. These results suggest that the "pecking order" can be observed in the sense that the smallest, and presumably least productive firms, invest in the most "attractive" countries.[18]

5 An Alternate Assessment of Sorting

An important, and very strong, prediction of the model is that the structure of multinational production across firms should follow a "pecking-order": the most productive firms invest in all foreign locations while the least productive firms invest in only the most attractive locations. While the results of the previous section suggest that on average this pattern is implied by the data, it is clear that there must be some deviation from the strictest form of the "pecking order." For instance, the composition regressions are stronger when the smallest parent firm size is used instead of average parent firm size.

In this section, we consider the following class of counterfactual exercises. First, we re-estimate the size distribution parameter $(k/(\sigma-1))$ that would be observed if the "pecking order" were strictly observed in the data. Second, we consider the number of countries in which individual firms would own an affiliate were the "pecking order" strictly observed. Finally, we consider the aggregate volume of local affiliate sales across countries that would be observed were the "pecking order" strictly observed, and we show how the structure of multinational activity would differ from the structure that is actually observed.

We begin by reporting the results obtained by regressing the logarithm of the average size of U.S. sales of parent firms that would be observed if the "pecking order" were strictly observed on the logarithm of the actual number of entrants by country. The coefficient estimate obtained is -0.746 with a standard error of 0.002 and an R-squared of 0.985. This implies a size-dispersion parameter of $(k/(\sigma-1))$ of approximately 1.5, which is more in line with that found by Eaton, Kortum, and Kramarz (2005). Compared to the number obtained from the actual structure of U.S. multinational operations, this implies a much greater degree of size dispersion than actually observed in the structure of multinational sales. Put another way, this result suggests that the observation of entry by U.S. firms is insufficiently concentrated in the largest and most productive firms than what the model would predict.

The actual and predicted propensity of firms to open foreign affiliates as a function of their productivity (as measured by sales in the United States) is illustrated in Figure 1. In this figure, firms are organized into five size categories depending on their share of the U.S. market. The largest firms appear in category 1 while the smallest appear in category 5. The lighter bar on the right for each grouping is the actual average number of countries in which firms in that size category own affiliates. Consistent with the results reported in Table 1, the more productive firms own more affiliates. The darker bar on the left for each grouping is the predicted number were the "pecking order" strictly followed. Clearly,

[18] Considering the second smallest parent firm generates very similar results.

the dispersion in the number of countries entered is too small given the dispersion in U.S. market shares: large firms invest in too few foreign locations and small firms invest in too many.

To get a sense of the types of locations that large firms underinvest and small firms overinvest, we created a counterfactual structure of multinational affiliate sales across locations as shown in Figure 2. In the figure, countries are sorted into five groups (10 countries per group) according to their "popularity." The lighter bar on the right corresponds to the actual average level (in logarithms) of affiliate sales of the U.S. firms investing in these countries. The darker bar on the left is the predicted average level (in logarithms) of affiliate sales if instead the "pecking order" were perfectly observed.[19] That is, holding fixed the country scale and the number of entering firms, the level of affiliate sales that would be observed if the productivities of those entrants perfectly followed the pecking order.

There are two key observations that can be made concerning the difference between actual and predicted levels of affiliate sales. First, in all countries there would be more U.S. multinational activity were the pecking order followed. Second, the percentage deviation is larger for the least attractive countries. This observation, combined with the observation that the largest firms invest in too few locations, suggests that the largest firms are underrepresented among the least attractive production locations. This conclusion is reinforced by the observation in the last section that the productivity composition of entrants as measured by the smallest entrant was better explained by country characteristics than the productivity composition of entrants as measured by the size of the average entrant.

6 Discussion and Suggestions for Future Research

This study has used a version of the model of Helpman, Melitz, and Yeaple (2004) to several ends. First, the paper demonstrates that firm heterogeneity does matter in understanding the structure of multinational activity: more productive firms enter a larger number of markets and sell more in each market that they enter than less productive firms. Second, firms tend to sort across countries in a manner that roughly approximates a "pecking order." On average, countries with characteristics that are more attractive to U.S. multinationals attract relatively less productive firms. This result was particularly strong when considering aggregate multinational activity, but was also present (albeit in a somewhat weaker form) when considering data disaggregated by industry. Third, the model was used to interpret the mechanisms through which country characteristics influence the structure of U.S. multinational activity. Among the conclusions reached from this analysis was that some variables typically hypothesized to work through fixed costs,

[19] In the decomposition given by (7), only the last term has been changed to compute this counterfactual level of local affiliate sales.

such as the effect of a common language or the effect of distance, actually appear to work through the market size facing the individual U.S. firm rather than through fixed costs.

An important contribution of the paper was also to demonstrate a number of short-comings of the simple model presented in Helpman, Melitz and Yeaple (2004). First, as shown in the individual firm analysis, the sales of a U.S. firm's affiliates are increasing in the parent firm's productivity as measured by the parent firm's U.S. sales, but at rate less than that predicted by the model. Second, the effect of distance on the structure of FDI suggests that standard models of the proximity-concentration trade-off are missing an ingredient that would explain why the unit cost of serving foreign markets appears to rise in distance. Finally, the counterfactuals considered in the last section demonstrate that the largest firms appear to invest in too few countries while the smaller firms tend to invest in too many. This underinvestment manifests itself in much less multinational activity in relatively less attractive countries than would be predicted by the naive model.

The manner in which the simple model falls short suggests potentially fruitful ways to design a model that is more consistent with the data. First, by including technology transfer costs and trade in intermediate goods the effect of distance on the structure of multinational activity can be better understood. Recent work by Keller and Yeaple (2008) shows that this mechanism can capture this particular short-coming. Second, the fact that the observed structure of multinational entry features too little entry by the largest firms suggests that diseconomies of scope across production locations may be at work. Combining this assumption with demand systems featuring variable mark-ups is a promising avenue for future work. As shown by Nocke (2006), when markups are decreasing in the number of entrants and firms can only invest in one location (extreme diseconomies of scope) more productive firms sort into larger markets and less productive firms into smaller markets. Allowing less extreme diseconomies of scope may be a useful avenue to extend Nocke (2006) to explain systematic deviations from the "pecking order" implied by standard models of firm heterogeneity.[20]

[20] Variable mark-ups could also weaken the link between the sales of parent firms in the United States (a large market) and the sales of their affiliates abroad. As pointed out in the discussion of Table 1, the coefficient on the logarithm of parent firm U.S. sales is less than one.

References

[1] Baldwin, Richard, and James Harrigan. 2007. "Zeros, Quality, and Space: Trade Theory and Trade Evidence", NBER Working Paper No. 13214.

[2] Bernard, Andrew, and J. Bradford Jensen. 1999. "Exceptional Exporter Performance: Cause, Effect, or Both?" *Journal of International Economics* 47(1): 1-25.

[3] Bernard, Andrew, J. Bradford Jensen, Stephen Redding, and Peter Schott. 2007. "Firms in International Trade." *Journal of Economic Perspectives* 21(3):105-130.

[4] Bernard, Andrew, Jonathan Eaton, Brad Jensen, and Samuel Kortum. 2003. "Plants and Productivity in International Trade." *American Economic Review* 93(4): 1268-1290.

[5] Brainard, S. Laul. 1997. "An Empirical Assessment of the Proximity-Concentration Trade-off Between Multinational Sales and Trade." *American Economic Review* 87(4): 520-544.

[6] Chaney, Thomas. 2008. "Distorted Gravity: Heterogeneous Firms, Market Structure, and the Geography of International Trade." mimeo, University of Chicago.

[7] Eaton, Jonathan, Samuel Kortum, and Francis Kramarz. 2005. "An Anatomy of International Trade: Evidence from French Firms." mimeo, University of Chicago.

[8] Girma, Sourafel, Richard Kneller, and Mauro Pisu. 2005. "Exports versus FDI: An Empirical Test." *Journal of World Economics* 141(2): 193-218.

[9] Head, Keith, and John Ries. 2003. "Heterogeneity and the FDI versus Exports Decision of Japanese Manufacturers." *Journal of the Japanese and World Economies* 17(4): 448-467.

[10] Head, Keith, and John Ries. 2008. "FDI as an Outcome of the Market for Corporate Control: Theory and Evidence." *Journal of International Economics* 74(3): 2-20.

[11] Helpman, Elhanan, Marc Melitz, and Yona Rubinstein. 2008. "Trading Partners and Trading Volumes." *Quarterly Journal of Economics* 123(2): 441–487.

[12] Helpman, Elhanan, Marc Melitz, and Stephen Yeaple. 2004. "Exports versus FDI with Heterogeneous Firms." *American Economic Review* 94(1): 300-316.

[13] Keller, Wolfgang, and Stephen Yeaple. 2008. "Global Production and Trade in the Knowledge Economy." mimeo University of Colorado at Boulder.

[14] Manova, Kalina. 2006. "Credit Constraints, Heterogeneous Firms and International Trade." mimeo Stanford University.

[15] Melitz, Marc. 2003. "The Impact of Trade on Intra-Industry Reallocations and Aggregate Industry Productivity." *Econometrica* 71: 1695-1725.

[16] Nocke, Volker. 2006. "A Gap for Me: Entreprenuers and Entry." *Journal of the European Economics Association* 4(5): 929-956.

[17] Tomiura, Eiichi. 2007. "Foreign Outsourcing, Exporting, and FDI: A Productivity Comparison at the Firm Level." *Journal of International Economics* 72(1) 113-127.

[18] Yeaple, Stephen. 2008. "Firm Heterogeneity, Central Locations, and the Structure of Foreign Direct Investment." Forthcoming in Helpman, E., D. Marin, and T. Verdier, *The Organization of Firms in a Global Economy*, Harvard University Press.

Table 1: The Scope and Scale of Individual U.S. Multinationals' Foreign Activities as a Function of Parent Firms' Efficiency

	Scope: Linear Probability		**Scale: OLS Regression**	
	(1)	**(2)**	**(3)**	**(4)**
Parent Size	0.030 (0.001)		0.538 (0.013)	
Parent TFP		0.010 (0.003)		0.765 (0.067)
Country FE	YES	YES	YES	YES
Industry FE	YES	YES	YES	YES
N **R-Squared**	53,400 0.187	51,800 0.150	4,102 0.419	4,054 0.128

Notes: Standard errors shown in parentheses. Both independent variables are in logarithms. TFP is measured as the residual of the regression of the logarithm of output per worker on the logarithm of capital per worker and the logarithm of the number of workers. The coefficients in the TFP regression are 0.32 and 0.65 respectively. The number of parent firms in the sample is 1,036.

Table 2: The Structure of Aggregate Multinational Activity by Component

	(1) Aggregate Sales	(2) Number	(3) Average Productivity	(4) Scale	(5) Average Sales
GDP	0.931 (0.114)	0.536 (0.070)	-0.141 (0.035)	0.536 (0.065)	0.395 (0.054)
GDPPC	0.923 (0.207)	0.606 (0.126)	-0.278 (0.063)	0.599 (0.114)	0.317 (0.093)
DIST	-0.761 (0.199)	-0.536 (0.131)	0.050 (0.068)	-0.275 (0.107)	-0.224 (0.115)
ENGLISH	0.754 (0.249)	0.739 (0.154)	-0.374 (0.121)	0.389 (0.151)	0.015 (0.167)
Obs **R-Squared**	49 0.738	49 0.724	49 0.500	49 0.757	49 0.623

Notes: Heteroskedascity robust standard errors are shown in parentheses. Aggregate sales correspond to local affiliate sales of all manufacturing affiliates owned by U.S. parent firms in manufacturing industries. All independent variables, except ENGLISH, are in logarithms. By construction, the coefficient estimates in column (1) are equal to the sum of the coefficients in columns (2)-(4). The coefficients in columns (4) and (5) have structural interpretations as Scale and Relative Costs.

Table 3: The Structure of Multinational Activity by Component, Disaggregated by Three-Digit Industry

	(1) Affiliate Sales	(2) Number	(3) Average Productivity	(4) Scale	(5) Average Sales
GDP	0.969 (0.102)	0.390 (0.051)	-0.032 (0.022)	0.611 (0.063)	0.579 (0.061)
GDPPC	0.649 (0.255)	0.249 (0.110)	-0.125 (0.037)	0.524 (0.158)	0.400 (0.162)
DIST	-0.486 (0.116)	-0.313 (0.069)	0.038 (0.020)	-0.209 (0.057)	-0.172 (0.051)
ENGLISH	0.631 (0.182)	0.548 (0.131)	-0.237 (0.053)	0.320 (0.098)	0.083 (0.080)
Obs **R-Squared**	1239 0.624	1239 0.734	1239 0.769	1239 0.557	1239 0.624

Notes: All specifications include a full set of industry dummies. Coefficients estimated by weighted least squares where the weights reflect country and industry sizes. Standard errors are robust to heteroskedascity and clustering by country. Aggregate sales correspond to local affiliate sales of all manufacturing affiliates owned by U.S. parent firms in manufacturing industries. All independent variables, except ENGLISH, are in logarithms. The coefficients on the industry dummy variables are suppressed to ease exposition. By construction, the coefficient estimates in column (1) are equal to the sum of the coefficients in columns (2)-(4). The coefficients in columns (4) and (5) have structural interpretations as Scale and Relative Costs.

Table 4: Alternative Measures of Productivity Composition

	Average TFP		The Cutoff Productivity Level	
	(1) Aggregate Sample	**(2) Disaggregated Sample**	**(3) Aggregate Sample**	**(4) Disaggregated Sample**
GDP	-0.031 (0.008)	-0.007 (0.004)	-0.787 (0.142)	-0.585 (0.076)
GDPPC	-0.051 (0.013)	-0.008 (0.008)	-1.213 (0.284)	-0.573 (0.182)
DIST	0.014 (0.013)	0.005 (0.004)	0.527 (0.347)	0.466 (0.128)
ENGLISH	-0.024 (0.021)	0.001 (0.009)	-0.903 (0.437)	-0.955 (0.231)
N **R-Squared**	49 0.379	1237 0.61	49 0.500	1239 0.659

Notes: Standard errors shown in parentheses. The standard errors are robust to heteroskedascity and to clustering at the country level. Approximate TFP is calculated as in the Table 2. The cutoff productivity level is calculated as the logarithm of the U.S. sales of the smallest parent firm that owns an affiliate in a country-industry pair. Coefficients on the industry dummies in the disaggregated samples are suppressed to aid exposition.

Figure 1: Predicted Versus Actual Number of Entrants by Size of Parent

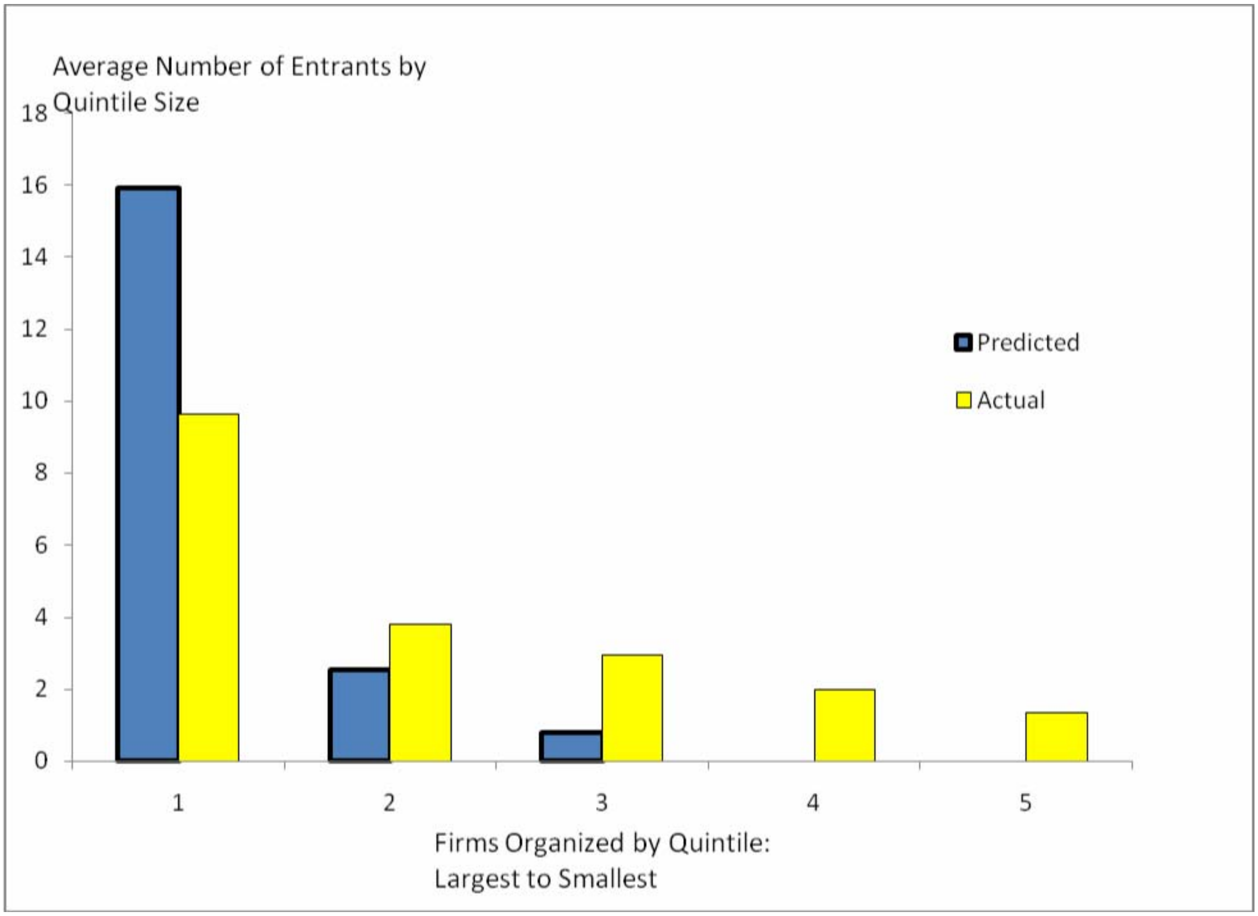

Notes: There are approximately 200 firms per size category. The measure of size is the parent firms' sales in the United States.

Figure 2: Predicted versus Actual Level of Foreign Affiliate Sales by Country

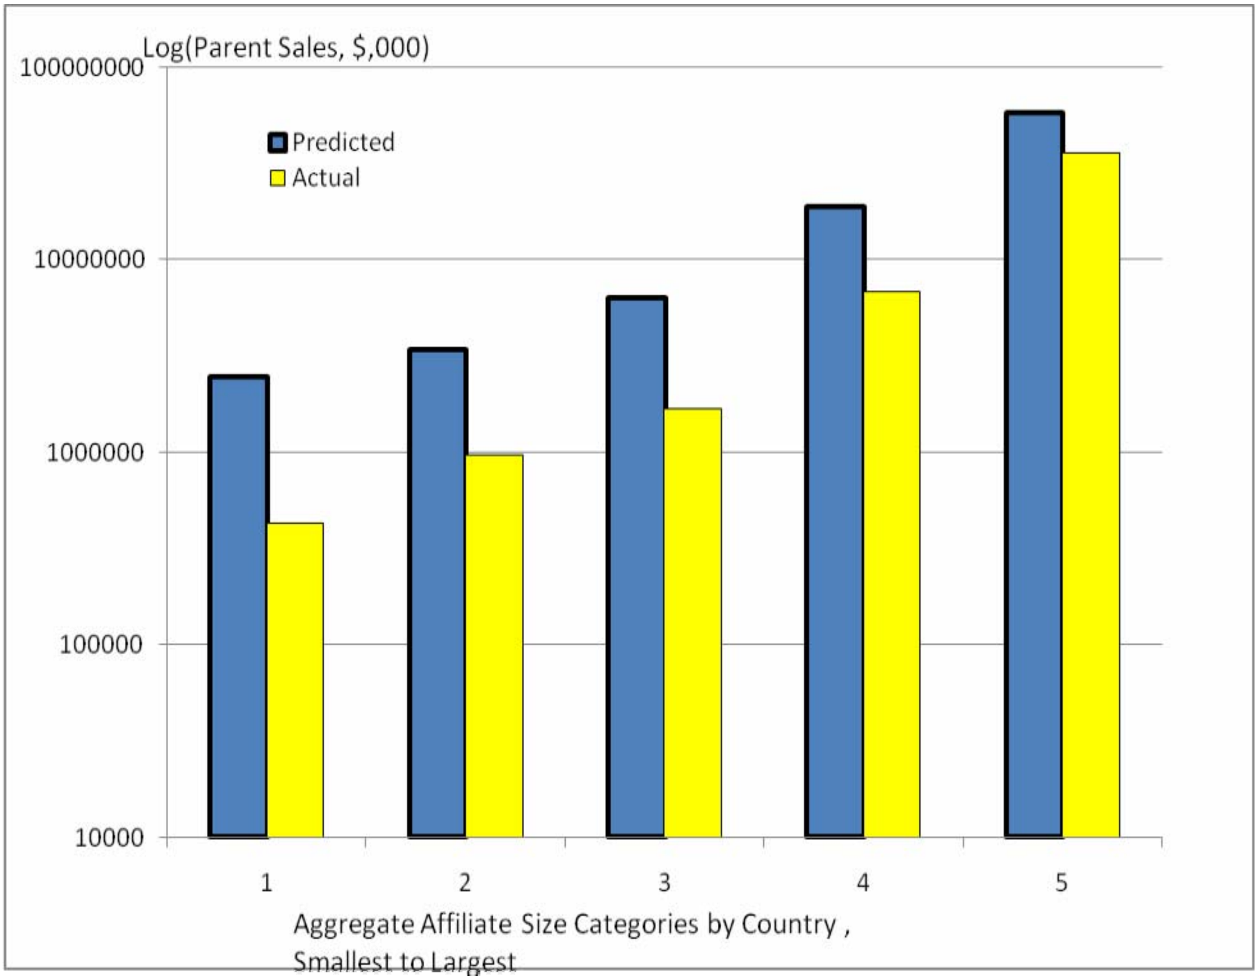

Appendix Table 1: Descriptive Statistics for Table 1

Variable	Mean	Standard Deviation
Entry	0.077	0.26
Log(affiliate sales)	9.97	1.60
Log(parent sales)	12.49	1.79
Log(parent TFP)	-0.073	0.45

Appendix Table 2: Descriptive Statistics for the Aggregate Sample (Tables 2 and 4)

Variable Name	Mean	Standard Deviation
Log(Aggregate Sales)	14.59	1.61
Log(Number)	3.84	1.02
Log(Average Productivity) measure: parent sales	15.25	0.44
Log(Average Productivity) measure: parent TFP	0.134	0.09
Log(Scale)	-4.50	0.93
GDP	19.25	1.26
GDPPC	9.25	0.71
DIST	8.89	0.57
ENGLISH	0.20	0.41

Appendix Table 3: Descriptive Statistics for the Disaggregated Sample (Tables 3 and 4)

Variable Name	Mean	Standard Deviation
Log(Aggregate Sales)	10.90	1.93
Log(Number)	0.76	0.85
Log(Average Productivity) Measure: parent sales	14.45	1.36
Log(Average Productivity) Measure: parent TFP	0.04	0.37
Log(Scale)	-4.31	1.49
GDP	19.63	1.18
GDPPC	9.45	0.64
DIST	8.82	0.64
ENGLISH	0.25	0.43