The Pennsylvania State University

The Graduate School

College of the Liberal Arts

ESSAYS ON INTERNATIONAL TRADE AND MULTINATIONAL FIRMS

A Dissertation in
Economics
by
Felix Tintelnot

Submitted in Partial Fulfillment
of the Requirements
for the Degree of

Doctor of Philosophy

May 2013

The dissertation of Felix Tintelnot was reviewed and approved* by the following:

Jonathan Eaton
Professor of Economics
Dissertation Co-Adviser, Co-Chair of Committee

Stephen R. Yeaple
Associate Professor of Economics
Dissertation Co-Adviser, Co-Chair of Committee

James R. Tybout
Professor of Economics

Martin Bojowald
Professor of Physics

Vijay Krishna
Professor of Economics

*Signatures are on file in the Graduate School.

Abstract

This dissertation consists of two chapters.

Most international commerce is carried out by multinational firms, which use their foreign affiliates for the majority of their foreign sales. In **Chapter 1**, I examine the determinants of multinational firms' location and production decisions and the welfare implications of multinational production. The few existing quantitative general equilibrium models that incorporate multinational firms achieve tractability by assuming away export platforms – i.e. they do not allow foreign affiliates of multinationals to export – or by ignoring fixed costs associated with foreign investment. I develop a quantifiable multi-country general equilibrium model, which tractably handles multinational firms that engage in export platform sales and that face fixed costs of foreign investment. I first estimate the model using German firm-level data to uncover the size and nature of costs of multinational enterprise and show that fixed costs of foreign investment are large. Second, I calibrate the model to data on trade and multinational production for twelve European and North American countries. Counterfactual results reveal that multinationals play an important role in transmitting technological improvements to foreign countries as they can jump the barriers to international trade; I find that a twenty percent increase in the productivity of US firms leads to welfare gains in foreign countries an order of magnitude larger than in a world in which multinational production is disallowed. I demonstrate the usefulness of the model for current policy analysis by studying the pending Canada-EU trade and investment agreement; I find that a twenty percent drop in the barriers to foreign produc-

tion between the signatories would divert about seven percent of the production of EU multinationals from the US to Canada.

Chapter 2, which is joint work with Kerem Cosar and Paul Grieco, studies the implications of national borders on economic activity. Using a micro-level dataset of wind turbine installations in Denmark and Germany, we estimate a structural oligopoly model with cross-border trade and heterogeneous firms. Our approach separately identifies border-related from distance-related variable costs and bounds the fixed cost of exporting for each firm. Variable border costs are large: equivalent to roughly 400 kilometers (250 miles) in distance costs, which represents 40 to 50 percent of the average exporter's total delivery costs. Fixed costs are also important; removing them would increase German firms' market share in Denmark by 10 percentage points. Counterfactual analysis indicates that completely eliminating border frictions would increase total welfare in the wind turbine industry by 5 percent in Denmark and 10 percent in Germany. Finally, an experiment using our structural model shows that commonly used price difference regressions produce misleadingly high estimates of the impact of national boundaries.

Contents

List of Figures

List of Tables

Dedication

To my mother, Kathrin Tintelnot, and my father, Eduard Hemmerling.

Acknowledgments

I am deeply indebted to my advisors Jonathan Eaton and Stephen Yeaple for their excellent guidance, encouragement, and invaluable support. They have very generously shared their time and great expertise with me. I am also extremely grateful to Andrés Rodríguez-Clare, Paul Grieco, and Kerem Cosar for many great discussions which have been very inspring and influential on my work.

I wish to thank Anca Cristea, Anna Gumpert, Christian Gourieroux, David Jinkins, Corinne Jones, Sung Jae Jun, Konstantin Kucheryavyy, Kala Krishna, Matthias Lux, Joris Pinkse, Mark Roberts, and Jim Tybout for helpful comments and suggestions on my dissertation. I thank Vijaj Krishna, Kalyan Chatterjee, Ruilin Zhou, Neil Wallace, Mark Roberts, Jonathan Eaton, Stephen Yeaple, and Kala Krishna for their inspiring graduate teaching in Economics. I thank Venky Venkateswaran and Neil Wallace for teaching me how to give a good economic presentation. I thank the German Bundesbank for the hospitality and access to its Microdatabase Direct investment (MiDi). I am thankful to Barry Ickes for his encouragement and support with a research assistantship from the Center for Research on International Financial and Energy Security.

I was fortunate to have a great group of friends and peers in my graduate program. In particular, I thank Vikram Kumar, Pradeep Kumar, Bernardo Quiroga Gomez, and Yao Luo for collaboration and discussions during the core courses in our graduate studies.

I also want to thank Wolfram Schrettl for his inspiring teaching and Stephan Schilling for many exciting discussions about economics during my undergraduate studies.

I would not have been as productive if I did not have had these great friends who I spent time with outside my graduate studies. I thank Emily Crawford and Rachel Tanos for being much more than running partners. I thank my climbing buddies Colin Shaw, Sean Devlin, and Tom Karnezos for great times inside and outside the rock climbing gym. And I thank my friends Sarah Polborn, Frank Erickson, Yao Luo, Matthias Lux, Michael Coccia, Katie Gates, Jessica Heckert, Eleanora Rossi, Maria Cruz Martin Garcia, Joel Wirth, and my brother Jonas Tintelnot for many great times over the last 5 years.

Chapter 1

Global Production with Export Platforms

1.1 Introduction

Most international commerce is carried out by multinational firms, which use foreign affiliates for the majority of their foreign sales.[1] In structuring their global operations, these firms confront various costs of multinational production and trade. For instance, whether a firm should pursue a strategy of maintaining many plants to avoid shipping costs or a strategy of consolidating production in a few locations turns on the size of fixed costs of establishing foreign plants relative to the costs of shipping goods. Further, given a set of production locations, the choice of which product to produce where depends on the interaction of comparative advantage and the cost of shipping goods. Taken as a whole, the structure of multinationals' operations must reflect the nature of costs of international commerce. It is therefore interesting that multinationals' global operations reveal a strong home bias: despite

---

[1] A multinational firm is a company with enterprises in more than one country. I define its home country as the country in which the parent company of the enterprises is registered. Usually, this coincides with the country of the multinational firm's headquarters. According to Bernard, Jensen, and Schott (2009), in the year 2000 multinational firms accounted for nearly 80 percent of US imports and exports, respectively, and employed 18 percent of the entire U.S. civilian workforce. Publicly available BEA data shows that, in the manufacturing sector, the sales by U.S. MNEs' majority-owned foreign affiliates are more than twice as large as the aggregate U.S. exports.

the opportunity to move production anywhere, they keep most of their production in the domestic country.

In this paper, I develop a framework that is designed to answer several key questions. First, what are the costs associated with multinational production? How important are the fixed costs of establishing foreign operations relative to the efficiency losses due to remote management? Second, how does the process of globalization, measured as a fall in these costs, affect the structure of global production? Will globalization result in firms' consolidating production in a few favored locations, or will firms expand their global production networks? Third, how does allowing for multinational production affect our understanding of the welfare effects in a general equilibrium trade model?

Existing quantitative models of trade and multinational production have proven tractable only after excluding many of the strategies that firms actually use or by shutting down mechanisms that are almost universally thought to be important. The framework that I develop to answer these questions is suitable both for structural estimation of global production costs using firm-level data and for aggregate quantitative analysis in general equilibrium. My model incorporates, and so allows me to quantify, a wide range of mechanisms that appear in the theoretical literature. In this model, firms choose from a rich array of production strategies in a multi-country setting in which variable trade and multinational production costs interact with increasing returns at the plant level.

An example of the rich production strategies that can be addressed in my model is the case of export platforms. Export platform sales are exports from a foreign plant to other countries. For US multinationals' affiliates in Europe, Figure 1 documents the proportion of output exported to other countries from the host country. Export platform sales account for on average around 40 percent of multinationals' foreign output, a share which is systematically higher for smaller countries. It is implausible to assume that the anticipation of these sales does not affect location or production decisions. My model also incorporates fixed costs of establishing foreign plants, a component of my study which is suggested by many firms' concentrating their production in only a few locations and which

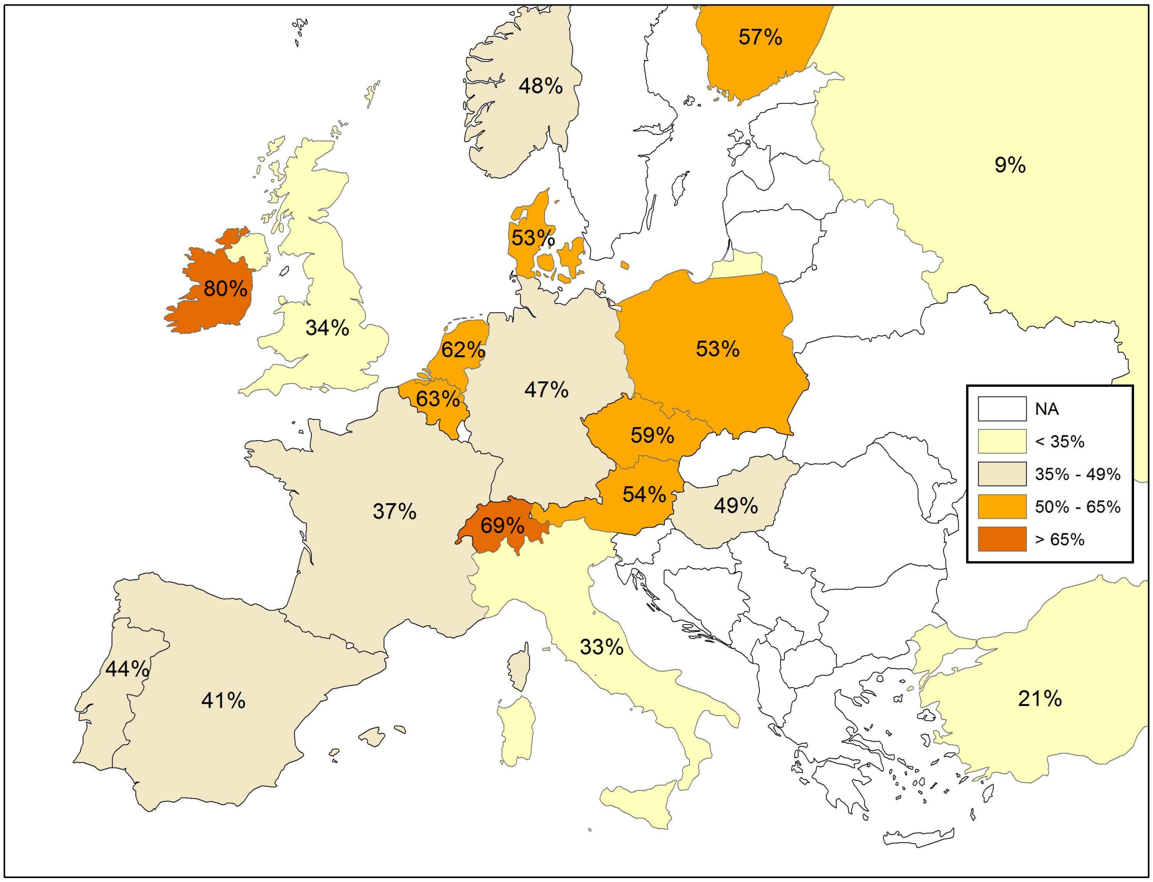

Figure 1: Export platform shares for US multinationals in Europe (year: 2004, source: BEA)

constitutes a key feature of much of the existing literature on multinational production. Nevertheless, the few empirical papers that incorporate export platforms ignore fixed costs of establishing foreign plants. I also incorporate multiple products per firm into the framework, an element for which a growing empirical literature provides much descriptive evidence.[2]

The possibility of export platform sales, together with the presence of fixed cost of establishing foreign plants, makes the decision as to which market to serve from which location interdependent across markets. For example, if the firm decides to serve France for a particular product from a local plant, then this decision affects the choice regarding which country to serve the Netherlands from,

---

[2] Evidence on the pervasiveness of multi-product firms is provided by Bernard, Jensen, Redding, and Schott (2007), Bernard, Redding, and Schott (2010), and Arkolakis and Muendler (2010).

because the fixed cost of establishing a plant in France has already been incurred. I solve the firm's problem in two stages. In the first stage, each firm chooses a set of countries in which to produce and incurs the fixed costs of establishing foreign plants. In the second stage, the firm decides for each product which market to serve from which location. In the countries in which a firm has established a plant, I treat its product-location-specific productivities as random variables, similarly to how Eaton and Kortum (2002) treat a country's productivities. By envisioning each firm as consisting of a continuum of products, I obtain intuitive, closed-form expressions for the output at each of the firm's plants. The firm's output is a function of the locations of the firm's plants, the productivity of each plant, the input costs in the plants' host countries, and the local and foreign market potential of the plants' host countries. Furthermore, the model delivers a probability with which a firm chooses a set of plants, as the fixed cost to establish a plant in a foreign country is stochastic and firm-country-specific.

With this framework, I conduct a two-tier empirical analysis. Using German firm-level data on output at the parent and affiliate levels, I estimate both the variable production costs in foreign countries as well as the distribution of fixed costs to establish a foreign plant. I find that German multinational firms face between 7 percent (Austria) and 42 percent (United States) larger variable production costs abroad than at home. In the data and estimated model, the share of foreign production of multinational firms is on average around 30 percent. If the variable production costs were the same in foreign countries as in Germany, the foreign output share would rise to 68 percent (taking into account firms' re-optimizing their locations). If, instead, variable production costs were at their estimated level and fixed costs to setting up foreign plants were zero (so each firm had a plant everywhere), the foreign output share would become 72 percent. Hence, fixed costs and larger variable production costs abroad are similarly important barriers to foreign production. If both variable production cost differences and fixed costs were eliminated, the foreign output share would rise to 88 percent (which is roughly equal to the share of foreign countries' GDPs in the set of countries considered).

In the second tier of my empirical inquiry, I turn my attention to general equilibrium welfare analysis. I calibrate the general equilibrium outcomes of my model to match data on bilateral trade flows, bilateral shares of foreign production, and the country-specific production cost estimates from German multinational firms. The cost estimates of German multinationals enable me to include both variable foreign production frictions and fixed costs in the analysis that otherwise includes only aggregate data. I solve for the endogenous relative wages and price indices in every country. With the calibrated model, I explore how globalization changes the structure of global production. For example, currently, Canada and the EU are negotiating a trade and investment agreement: CETA. If one supposes that the agreement is signed and yields a twenty percent reduction of variable and fixed production costs between the signatories, then – according to my calibrated model – EU multinationals would divert around seven percent of their production from the US to Canada. These findings hinge on the possibility of export platform sales from Canada to the US. Without this possibility, the location and output decisions of European firms are independent between Canada and the United States. Instead, I find that a Canada-EU trade and investment agreement could induce a strong third-party effect on the United States.

A more complete model of multinational production and trade can revise answers to classic questions in the trade literature. First, I evaluate the welfare gains from trade both in my global production model and in a classical trade model without multinational production offered by Anderson and van Wincoop (2003), which is a special case of my model when multinational production is shut down. Contrary to what one may expect, I find that the gains from trade estimates from this standard trade model without multinational production are very similar to the gains from trade estimates in my global production model. However, multinational production is instrumental for the analysis of gains from foreign technology improvements, a question studied by Eaton and Kortum (2002), among others. Suppose all US firms improve their technology by 20 percent. I find that the welfare gains in foreign countries from such a technology improvement are an order of magnitude larger when multinational production is taken into account.

The model presented in this paper contains elements of Helpman, Melitz, and Yeaple (2004) and Eaton and Kortum (2002). As in Helpman, Melitz, and Yeaple (2004), firms produce differentiated goods and can establish foreign plants at the expense of fixed costs.[3] I extend their framework by incorporating export-platform sales and multi-product firms. As in Eaton and Kortum (2002), countries differ in their comparative advantage in production. In my model, however, each product can be produced only by a single firm, which can also produce in foreign countries, while Eaton and Kortum (2002) instead assume that each firm operates only domestically and that firms from different countries can produce the same product. If multinational production is prohibitively costly, my model collapses with respect to its aggregate predictions to Anderson and van Wincoop (2003), and the product-location-specific productivity draws have no impact.

A vibrant area of ongoing research centers on the gains from multinational production and trade. Ramondo and Rodriguez-Clare (2012) investigate the gains from trade, multinational production, and openness.[4] They find that the gains from trade can be twice as large if multinational production is taken into account than without.[5] Arkolakis, Ramondo, Rodriguez-Clare, and Yeaple (2012) endogenize the allocation between production and innovation in a model of global production with monopolistic competition. A key difference between these papers and my work is that they assume away fixed costs of foreign plants. Their calibrated models' fit of the data on export platform sales is only successful in special cases.[6] While my model fits the export platform sales of US multinationals well (without having aimed to fit those in the calibration), a restricted version of my model without fixed costs does not. Both fixed and variable costs discourage foreign production, but it is

---

[3] Helpman, Melitz, and Yeaple (2004) combine key elements that appeared in Melitz (2003) and Horstmann and Markusen (1992).

[4] Their paper extends the Ricardian trade model by Eaton and Kortum (2002) insofar as it allows the technologies that originated in a country to be used for production abroad.

[5] One reason that our results differ is that in their paper, a complementarity between trade and MP is directly built into the input bundle of a multinational firm abroad, which is a function of intermediate shipments from the home country.

[6] In Ramondo and Rodriguez-Clare (2012), only when the productivity draws for ideas that originated in one country are uncorrelated across countries can the calibrated model come close to matching the data on export platform sales for US multinationals. The calibrated model in Arkolakis, Ramondo, Rodriguez-Clare, and Yeaple (2012) generates much lower export platform sales for US firms than in the data.

the fixed costs that induce firms to concentrate their production in a few locations.[7]

My findings that multinational firms face significantly larger variable production costs abroad and significant fixed costs of establishing foreign plants are in line with the findings of Irarrazabal, Moxnes, and Opromolla (2009). They use data from Norwegian firms and develop a structural model that extends Helpman, Melitz, and Yeaple (2004) by incorporating intra-firm trade, and they find that a very large share of intra-firm trade is necessary to rationalize the observed output data.[8] Their paper ignores export platform sales, however, which makes the set of production strategies across which a firm can choose much smaller. Without the possibility of export platform sales, the decision to set up an affiliate in Belgium is independent of the decision to set up an affiliate in the Netherlands.[9]

Since in my model firms choose a set of production locations instead of making independent decisions about whether to establish a plant for each country, this paper also joins a literature that studies large discrete choice problems at the firm level.[10] Morales, Sheu, and Zahler (2011) estimate a dynamic trade model in which the costs of serving a foreign market depend on the set of foreign markets the firm had served in the past. This creates an interdependency of the destination markets. Interdependent location choices within the firm also arise in Holmes (2011), who estimates the determinants of the expansion of Walmart stores within the United States. Both papers use moment

---

[7] Fixed costs and export platforms have been analyzed together only in very restrictive settings. Neary (2002) shows in a theoretical analysis that with export platform sales and fixed costs of establishing foreign plants, the European single-market policy increases foreign direct investment into the EU from outside countries. Ekholm, Forslid, and Markusen (2007) develop a three-country model that incorporates both fixed costs and export platform sales. Other three-country models with fixed costs and complex relationships between domestic and foreign plants have been developed by Yeaple (2003) and Grossman, Helpman, and Szeidl (2006). However, it is impractical to apply their model to the data of many countries. Head and Mayer (2004) apply a model with multiple countries, fixed costs, and sales to surrounding markets to data on Japanese affiliates under the restriction that each firm can only have a single production location. The interdependence between firms' location and production decisions has been reflected in empirical work by Baltagi, Egger, and Pfaffermayr (2008) and Blonigen, Davies, Waddell, and Naughton (2007), who apply spatial econometric methods to data on bilateral FDI and multinational firms' sales and point out significant third-country effects in their estimation results.

[8] Instead of assuming intra-firm trade, I allow the production efficiency of foreign affiliates to differ from the production efficiency at home (e.g., through communication costs with headquarters).

[9] Existing work on structural estimation with data on multinational firms is sparse. Exceptions are Feinberg and Keane (2006) who structurally estimate U.S. multinationals' decisions to invest and produce in Canada, and Rodrigue (2010) who structurally estimates a model of trade and FDI with data on Indonesian manufacturing plants.

[10] The decision as to where to establish facilities and which market to serve from which facility is known as the 'Facility Location Problem' in operations research. See Klose and Drexl (2005) for a survey of the literature on the 'Facility Location Problem,' which is primarily concerned with developing solution algorithms to the single firm's problem.

inequalities to conduct their estimations. By contrast, the parameters in my model are point-identified, which enables me to conduct general equilibrium analysis.

The following section outlines the model. Section 1.3 estimates country-specific fixed and variable production costs for German multinational firms via constrained maximum likelihood. Section 1.4 calibrates the general equilibrium, and Section 1.5 conducts the counterfactual exercises described above. Section 1.6 concludes.

1.2 A model of global production with export platforms

I develop a model that explains in which countries firms locate their plants, how much they produce in each country, and how much they ship from one country to another. Geography is reflected in three kinds of barriers between countries: variable iceberg trade costs, variable efficiency losses in foreign production, and fixed costs to establish foreign plants. Countries differ in endowments of labor and the mass and distribution of firms. While the technology of local firms is part of the endowments, the set of firms that produce in a country is determined endogenously. I assume a market structure characterized by monopolistic competition. For simplicity, I assume there are no fixed costs to exporting.[11] Consequently, every product is sold to every market. I start with the description of demand and then turn to the problem of the firm.

1.2.1 Demand

I assume standard CES preferences, with the distinction that here each firm has a continuum of products instead of a single product. A good is indexed by a firm ω and a variety υ. I assume a measure 1 of varieties per firm and a fixed measure of firms. If the representative consumer of country j consumes $q_j(\omega, \upsilon)$ units of each variety υ of each firm $\omega \in \Omega$, she gets the following utility:

---

[11] Fixed costs of exporting (at the firm level) could be incorporated, as in Eaton, Kortum, and Kramarz (2011), but they are omitted for simplicity and would require additional data to be identified.

$$U^j \equiv \left(\int_{\Omega} \int_{0}^{1} q_j(\omega, \upsilon)^{(\sigma-1)/\sigma} d\upsilon d\omega \right)^{\sigma/(\sigma-1)} . \tag{1}$$

The elasticity of substitution $\sigma > 1$ is identical between varieties inside and outside the firm. Assuming the same elasticity of substitution between varieties within the firm and between varieties from different firms simplifies the pricing decision by the firm. Consumers maximize their utility by choosing their consumption of goods subject to their budget constraint. I denote the aggregate income in country j by Y_j. Utility maximization implies that the quantity demanded in country j of variety υ supplied by firm ω at price $p_j(\omega, \upsilon)$ is

$$q_j(\omega, \upsilon) = p_j(\omega, \upsilon)^{-\sigma} \frac{Y_j}{P_j^{1-\sigma}}, \tag{2}$$

where P_j is the ideal price index in country j:

$$P_j \equiv \left[\int_{\Omega_j} p_j(\omega)^{(1-\sigma)} d\omega \right]^{1/(1-\sigma)} , \tag{3}$$

which is simply the standard CES price index over the firm-level price indices. The price index of firm ω to country j is

$$p_j(\omega) \equiv \left(\int_{0}^{1} p_j(\omega, \upsilon)^{1-\sigma} d\upsilon \right)^{1/(1-\sigma)} , \tag{4}$$

and the expenditure on goods produced by firm ω in country j is

$$s_j(\omega) = p_j(\omega)^{1-\sigma} \frac{Y_j}{P_j^{1-\sigma}}. \tag{5}$$

Next, I proceed to describe the problem of a single firm.

1.2.2 The firm's problem

Each firm behaves like a monopolist and faces a CES demand function for each of its products. Every firm is infinitesimal and takes aggregate price indices, income, and wages as given. The problem of the firm consists of two stages: first, the firm selects the set of countries in which to establish a plant in order to maximize expected profits; it then learns about the exact quality of each plant, and decides which market to serve from which location for each product. Note that the timing assumption – the firm learns about the quality of each plant after the set of production locations is selected – is not essential, but it simplifies the analysis of firm-level data for reasons that I will discuss in Section 3.

A firm is characterized by its country of origin, i, its core productivity parameter, ϕ, a vector of fixed cost levels in every country, η, and a vector of location-specific productivity shifters, ϵ. All these variables are firm-specific. There are N countries.

Production decisions after the plants are selected

Denote by Z the set of locations the firm has selected for production plants. I assume that a firm always has a plant in its home country. In those countries in which the firm has established a plant, the firm draws a location-specific productivity for each of its products from a Fréchet distribution.[12] Let ν_j be a random variable that denotes the productivity level in country j for a particular product. The cumulative distribution function of a product's productivity in country j is:

$$\Pr(\nu_j \leq x) = \exp\left(- \left(\phi \epsilon_j\right)^{\theta} \left(\gamma_{ij} x\right)^{-\theta} \right).$$

The product of the core productivity level, ϕ, and the plant-specific productivity shifter, ϵ_j, determines the level of the productivity draws in the plant in country j. Larger values of $\phi \epsilon_j$ imply

---

[12] See Kotz and Nadarajah (2000), Chapter 1, for a description of the Fréchet and other extreme value distributions.

better productivity distributions.[13] The dispersion of the productivity draws is decreasing in θ. All firms from country i may have lower productivity in country l, which is captured by an iceberg loss in production, γ_{il}. These losses may for example occur because of higher costs due to communication challenges, information frictions, or shipments of intermediate products. For technical reasons I impose $\theta > \max(\sigma - 1, 1)$.

At each location, the firm transforms units of labor into goods at a constant marginal cost inversely proportional to productivity. The wage in country j is denoted by w_j. Trade costs to ship goods from country l to m are of the iceberg type and are denoted by τ_{lm}. Given these assumptions about production and shipping technology, it is easy to derive that the costs to serve market m from country $l \in Z$ are distributed as

$$\Pr\left(\frac{w_l \tau_{lm}}{\nu_l} \leq c\right) = 1 - \exp\left(-\left(\frac{\gamma_{il} w_l \tau_{lm}}{\phi \epsilon_l}\right)^{-\theta} c^{\theta}\right).$$

Having its production plants in place, the firm selects, for each product and market, the production location that can supply that market at the minimum cost. Using the known properties of the Fréchet distribution, one can derive that the product-level costs with which the firm will serve market m are distributed according to

$$G_m(c|i, \phi, Z, \epsilon) = 1 - \exp\left(-\sum_{k \in Z}\left(\frac{\gamma_{ik} w_k \tau_{km}}{\phi \epsilon_k}\right)^{-\theta} c^{\theta}\right). \tag{6}$$

With CES preferences and monopolistic competition, the firm charges a constant mark-up, $\frac{\sigma}{\sigma-1}$, for each good over the unit cost of delivering the good to each market. Using the optimal pricing rule, and the distribution of product-level costs, (6), we can write the firm-level price index – defined in (4) – which aggregates the product-level prices that the firm (i, ϕ, Z, ϵ) charges in market m, as

---

[13] The reader familiar with Eaton and Kortum (2002) may recognize the similarity between the country-specific parameter T_j in their paper and the firm-country-specific parameter $\phi \epsilon_j$ in this paper.

$$p_m(i,\phi,Z,\epsilon) = \kappa^{\frac{1}{1-\sigma}} \phi^{-1} \left(\sum_{k \in Z} (\gamma_{ik} w_k \tau_{km})^{-\theta} \epsilon_k^{\theta} \right)^{-1/\theta}, \tag{7}$$

where $\kappa = \Gamma\left(\frac{\theta+1-\sigma}{\theta}\right)\left(\frac{\sigma}{\sigma-1}\right)^{1-\sigma}$ is a constant.[14] The total sales of firm (i,ϕ,Z,ϵ) in market m are

$$s_m(i,\phi,Z,\epsilon) = p_m(i,\phi,Z,\epsilon)^{1-\sigma} \frac{Y_m}{P_m^{1-\sigma}}. \tag{8}$$

The expressions for the firm's price index, (7), and total sales, (8), in market m have intuitive properties: the sales rise in the core productivity level of the firm; furthermore, the firm benefits particularly from having a plant in a country k in which the variable costs to supply market m are low (low $\gamma_{ik} w_k \tau_{km}$), and in which the firm has a large plant-wide productivity shifter (large ϵ_k).

Due to constant returns to scale in the variable production costs, the firm will simply choose for each variety the location with the lowest unit cost to serve a market. We can write the share of products for which the plant in country l is selected to serve country m as

$$\mu_{lm}(i,\phi,Z,\epsilon) = \Pr\left[\underset{j \in Z}{\operatorname{argmin}} \frac{\gamma_{ij} w_j \tau_{jm}}{\nu_j} = l \right] = \begin{cases} \frac{(\gamma_{il} w_l \tau_{lm})^{-\theta} \epsilon_l^{\theta}}{\sum_{k \in Z} (\gamma_{ik} w_k \tau_{km})^{-\theta} \epsilon_k^{\theta}} & \text{if } l \in Z \\ 0 & \text{otherwise} \end{cases}. \tag{9}$$

The share of goods that a firm ships from country l to country m is large if the plant in country l has low costs to serve market m relative to the firm's other plants. If the firm has a plant in country l ($l \in Z$), the product-level cost at which a firm actually supplies market m from location l also has the distribution $G_m(c|i,\phi,Z,\epsilon)$. Consequently, $\mu_{lm}(i,\phi,Z,\epsilon)$ equals not only the share of products that a firm with location set Z ships from location l to market m, but also the corresponding value share. Therefore, the sales from location l to market m for such a firm are

---

[14] This step is analogous to the calculation of the *overall* price index in Eaton and Kortum (2002) and uses the moment generating function for Fréchet distributed random variables. The calculation is valid under the restriction made earlier that $\theta > \sigma - 1$.

$$s_{lm}(i, \phi, Z, \epsilon) = s_m(i, \phi, Z, \epsilon)\mu_{lm}(i, \phi, Z, \epsilon). \tag{10}$$

The relationship between a firm's plants is described in Proposition 1, whose proof is in the appendix.

Proposition 1. *The firm-level sales to each market increase as additional production locations are added to the set of existing locations. However, there is a cannibalization effect across production locations. That is, a firm that adds a production location decreases the sales from the other locations.*

Next, I proceed to examine the optimal choice of the set of locations, Z.

Choice of production locations

There are various motivations for setting up foreign plants: a foreign plant yields proximity to the local and surrounding markets, may have lower factor costs, and, finally, has a comparative advantage in the production of some of the firm's products. On the other hand, the firm incurs a fixed cost for establishing a foreign plant, which motivates the firm to concentrate its production in as few locations as possible. The firm selects a set of production locations based on its core productivity level, ϕ, its fixed cost draws, η, and its country of origin, i. As it is assumed that a firm always has a plant in its home country, in total, there are 2^{N-1} feasible combinations of locations. I denote the set that contains all sets of locations for a firm from country i by $\mathcal{Z}^i$. Fixed costs have to be paid in units of labor from the host country. If the firm chooses the set of locations $Z \in \mathcal{Z}^i$, the firm incurs fixed costs equal to $\sum_{l \in Z} \eta_l w_l$.

The firm chooses the set of locations that maximizes its expected profits. The expected variable profits from Z are simply the sum of the expected sales to all markets multiplied by the proportion of sales that represents variable profits:

$$E_\epsilon(\pi(i,\phi,Z,\epsilon)) = \frac{1}{\sigma}\sum_m E_\epsilon(s_m(i,\phi,Z,\epsilon)). \tag{11}$$

The total expected profits of set Z are the expected variable profits minus the fixed cost payments associated with the locations contained in the set. I assume that no fixed costs have to be paid for the domestic plant (or that they have been paid in the firm's entry stage that I do not include in this model). The expected total profits from choosing a set of locations Z are thus:

$$E_\epsilon(\Pi(i,\phi,Z,\epsilon,\eta)) = E_\epsilon(\pi(i,\phi,Z,\epsilon)) - \sum_{k\in Z, k\neq i}\eta_k w_k. \tag{12}$$

I write the set of locations that maximizes the expected profits as

$$Z(i,\phi,\eta) \in \arg\max_{Z\in\mathcal{Z}^i} E_\epsilon(\Pi(i,\phi,Z,\epsilon,\eta)). \tag{13}$$

While, in general, multiple sets of locations could be optimal for the firm, as long as the fixed cost vector η is drawn from a continuous distribution (where the draws are independent across countries), the set of fixed cost shock vectors for which the firm is indifferent across two or more location sets has measure zero.

In the following subsection, I turn to describing the endowments of each country, the aggregation of the firms' choices, and the global production equilibrium.

1.2.3 Equilibrium

Country j is endowed with a population L_j and a continuum of heterogeneous firms of mass M_j. I assume that the elements of the fixed cost vector, η, are drawn independently across countries from a distribution denoted by $F^i(\eta)$ that can differ by the country of origin, i, is continuous, and has the positive orthant as its support.[15] The core productivity level, ϕ, and the vector of location-specific productivity shifters, ϵ, can be realizations of arbitrary (potentially degenerate) distributions, which are denoted by $G(\phi)$ and $H(\epsilon)$, respectively.

Now I proceed to aggregate over the individual firms' choices to establish expressions that I use in the definition of the global production equilibrium below. The share of firms from country i with core productivity ϕ that choose location set Z is

$$\rho_Z^{i,\phi} = \int\limits_{\eta} \mathbb{1}\left[Z(i,\phi,\eta) = Z\right] dF^i(\eta). \tag{14}$$

This formulation is used in the derivation of the total sales of firms that originated in country i from country l to country m, X_{ilm}. We can simply integrate over the core productivity levels of the firms from country i, and write their sales as the weighted sum of the sales a firm would make from country l to country m conditional on a location set, where the weights are the probabilities with which the firm actually chooses this location set:

$$X_{ilm} = M_i \int\limits_{\phi} \sum_{Z' \in \mathcal{Z}^i} \rho_{Z'}^{i,\phi} E_\epsilon(s_{lm}(i,\phi,Z',\epsilon)) dG(\phi). \tag{15}$$

Aggregate trade flows from country l to m are then simply the sum of the term X_{ilm} across all countries of origin:

$$X_{lm} = \sum_i X_{ilm}. \tag{16}$$

Following (3), the consumer price index in market m, P_m, consists of the firm-level price indices for market m of firms from all countries. Again, the expression is the integral over the core productivity

---

[15] For instance, the fixed costs to produce domestically are assumed to be zero, which generates differences among the fixed cost contributions across countries.

levels of the firms and a weighted sum of the firms' price indices conditional on their location choice:

$$P_m = \left[\sum_i M_i \int_\phi \sum_{Z' \in \mathcal{Z}^i} \rho_{Z'}^{i,\phi} E_\epsilon (p_m(i, \phi, Z', \epsilon)^{1-\sigma}) dG(\phi) \right]^{1/(1-\sigma)} . \tag{17}$$

In order to establish the labor market clearing condition for country k, I define the set of feasible location sets for firms from country i that include a location in country k as $\Delta_k^i = \{Z \in \mathcal{Z}^i \mid k \in Z\}$. Total labor income in country k is equal to the sum of the wages paid in production in country k by firms from all countries and of the wages paid in plant construction by foreign companies:

$$w_k L_k = \frac{\sigma - 1}{\sigma} \sum_m X_{km} + \sum_{i \neq k} M_i \int_\phi \int_\eta \sum_{Z \in \Delta_k^i} \mathbb{1}\left[Z(i, \phi, \eta) = Z\right] \eta_k w_k dF^i(\eta) dG(\phi). \tag{18}$$

I assume that a representative household owns the domestic firms.[16] The aggregate income in country m is then the sum of the labor payments and the profits by firms that originated in country m.

$$Y_m = w_m L_m + M_m \int_\phi \int_\eta \sum_{Z \in \mathcal{Z}^m} \mathbb{1}\left[Z(i, \phi, \eta) = Z\right] E_\epsilon (\Pi(i, \phi, Z, \epsilon, \eta) dF^i(\eta) dG(\phi) \tag{19}$$

Now that I have defined the expressions above, I can define the global production equilibrium.

Definition 1. *Given* $\tau_{ij}, \gamma_{ij}, F^i(\eta), G(\phi), H(\epsilon), M_i, \mathcal{Z}^i, \forall i, j = 1, ..., N$, *a* **global production equilibrium** *is a set of wages,* w_i*, price indices,* P_i*, incomes,* Y_i*, allocations for the representative consumer,* $q(\omega, \upsilon)$*, prices,* $p_m(i, \phi, Z, \epsilon)$*, and location choices,* $Z(i, \phi, \eta)$*, for the firm, such that*

---

[16] This seems to be a reasonable assumption: according to Cummings, Manyika, Mendonca, Greenberg, Aronowitz, Chopra, Elkin, Ramaswamy, Soni, and Watson (2010), in 2007, U.S. residents held 86 percent of the total market value of all U.S. companies' equities either directly as individual investors or indirectly through pension funds and retirement and insurance accounts.

(i) *equation* (2) *is the solution of the consumer's optimization problem.*

(ii) $p_m(i, \phi, Z, \epsilon)$ *and* $Z(i, \phi, \eta)$ *solve the firm's profit maximization problem.*

(iii) P_i *satisfies equation* (17).

(iv) *The labor market clearing condition,* (18)*, holds.*

(v) Y_m *satisfies equation* (19).

Since the model is static, utility maximization implies current account balance. However, it is possible that a country runs a trade deficit, which is financed by the profits that this country's multinational firms make abroad.

In the following section I apply this model to data from German multinational firms to identify the determinants of firms' production and location choices. In this first tier of my empirical analysis, I take wages, aggregate income, and price indices in countries as given.

1.3 Estimation of fixed and variable production costs

This section estimates the barriers to foreign production faced by German multinationals. Subsection 1.3.1 documents that German firms tend to concentrate their production in only a few countries, and – conditional on being active in a foreign country – produce less in that foreign country than the relative size of the foreign economy (measured in GDP or gross production) would suggest if multinationals were completely footloose. Subsection 1.3.2 describes the estimation of fixed and variable costs of foreign production with constrained maximum likelihood, whose parameter estimates are presented in Subsection 1.3.3. Finally, Subsection 1.3.4 conducts counterfactual analysis to document the quantitative importance of each of these barriers.

1.3.1 Data description and preliminary evidence on barriers to foreign production

My analysis in this section is based on firm-level data on German multinational firms in the manufacturing sector. By law, German resident investors are required to report on the activities of foreign affiliates if the affiliate has a balance sheet total above 3 million Euro and the investor has a share of voting rights of 10 percent or more. The information about the foreign affiliates is contained in the Microdatabase Direct Investment (MiDi) which is maintained by the German Bundesbank.[17] I use data for the year 2005 for affiliates that belong to the manufacturing sector and that are majority-owned by a parent firm in the manufacturing sector. I focus on German multinationals' activities in twelve Western European and North American countries.[18] I take the set of countries in which a multinational owns an affiliate (including the home country) as the corresponding data analogue to the set of production plants in the model. I observe the total sales for each affiliate as well as the total sales for the parent company.[19]

The data for manufacturing firms in these host countries contains 1,711 positive firm-country output observations from 665 firms. The United States and France are the most popular destination countries for German multinational firms. Table 17 in Appendix B describes the activities at the country level. Most multinationals concentrate their production in very few countries: the average number of production locations (including the home country) is 2.57. Further, the fraction of multinationals' production that occurs abroad is small relative to the fraction of aggregate production (by all – including foreign – firms) that occurs abroad. I call this phenomenon 'home bias in production.' On average, across all German multinationals, the share of foreign production in total output is 0.29. Table 18 in Appendix B shows that the share of foreign production in total output is rising in the number of foreign affiliates. However, even for firms with more than six production locations, the

---

[17] Other research uses of the database include Muendler and Becker (2010), who study the margins of multinational labor substitution for multinational firms, and Buch, Kleinert, Lipponer, and Toubal (2005), who characterize the patterns of German firms' multinational activities.

[18] These countries are Austria, Belgium, Canada, Switzerland, Germany, Spain, France, United Kingdom, Ireland, Italy, Netherlands, and the United States.

[19] I consolidate multiple affiliates in the same country by the same parent company into one entity.

average share of total output that is produced abroad is only around 50 percent. Suppose a firm's output in country k were proportional to the value of gross production in country k. This would result in an average share of foreign output to global output of 0.44, controlling for the set of locations in which the firm is active.[20] As this figure is larger than the actual foreign output share of firms, 0.29, this finding suggests that, beyond fixed costs, differences in variable production costs drive home bias in production.[21]

Additionally, I use data on gross production and bilateral trade flows from the OECD STAN database to calculate country-specific manufacturing absorption (described in Appendix B), and I use estimates from a standard gravity pure trade model as proxies for bilateral trade costs and price indices.

1.3.2 Estimation

Next, I complete the empirical specification of the model, and then I show how fixed and variable production costs can be estimated from location set and output data from German multinationals via constrained maximum likelihood.

Parameterization

Let $\tilde{\eta}_{t,k} = \eta_{t,k} w_k$ denote the value of the fixed costs that firm t must pay to erect a production facility in country k. Let $\tilde{w}_k = w_k \gamma_{ik}$ denote the unit input costs in country k of German firms (firms from country i). I add a subscript t to the variables that are firm-specific. I assume that the fixed cost that a firm has to pay to start production in country k, $\tilde{\eta}_{t,k}$, is drawn independently across countries and firms from a log-normal distribution with mean $\mu_{\tilde{\eta}}$ and standard deviation $\sigma_{\tilde{\eta}}$. I set

---

[20] Specifically, I calculate for each firm with location set Z, $\frac{\sum_{k \neq i, k \in Z} y_k}{\sum_{k \in Z} y_k}$, where y_k denotes gross production in manufacturing in country k and i denotes the country of origin of the firm (here Germany). The average of this measure across firms is 0.44 as opposed to 0.29 for the average foreign output share of the firms.

[21] This pattern is robust across various sub-sectors of the manufacturing sector (see Table 19 in Appendix B), with the exception being '*other non-metallic mineral products*' in which the mean share of foreign host countries' production exceeds the mean share of foreign production by German firms from this sector.

the fixed costs in Germany to zero and normalize the unit input costs in Germany to one. Further, I assume that the location-specific productivity shifter ϵ is drawn from a log-normal distribution, $\log\mathcal{N}(0, \sigma_\epsilon)$, independently across countries and firms, and that the core productivity levels of the German multinationals are drawn from a Pareto distribution with scale parameter μ_ϕ and shape parameter σ_ϕ.

I set the value of the elasticity of substitution between products, σ, to six. This implies a reasonable mark-up of 20 percent above marginal costs. The estimates are robust to various parameters for the dispersion parameter, θ, of the distribution of the country-firm specific productivity shifters. I use a benchmark value of seven for the dispersion parameter ($\theta = 6$ and $\theta = 9$ give very similar results).[22]

Constrained Maximum Likelihood Estimation

Under the new notation with firm subscripts, $\tilde{\eta}_{t,k} = \eta_{t,k} w_k$, $\tilde{w}_k = w_k \gamma_{ik}$, and equations (11) and (12) from the model, the expected profits from selecting location set Z for firm t with core productivity ϕ_t and fixed cost draws $\tilde{\eta}_t$ are:

$$E_\epsilon(\Pi(\phi_t, Z, \epsilon, \tilde{\eta}_t; \sigma_\epsilon, \tilde{w})) = \frac{1}{\sigma}\kappa\phi_t^{\sigma-1}\sum_m \int_\epsilon \frac{Y_m}{P_m^{1-\sigma}}\left(\sum_{k\in Z}(\tilde{w}_k\tau_{km})^{-\theta}\epsilon_k^\theta\right)^{(\sigma-1)/\theta} dH(\epsilon;\sigma_\epsilon) - \sum_{k\in Z, k\neq i}\tilde{\eta}_{t,k}. \tag{20}$$

The first term represents expected variable profits from having production facilities in the countries contained in the location set, and the second term represents the fixed costs that the firm would have to pay. Recall that the level of fixed costs is known at the time the firm makes its decision, but the firm only learns how productive these facilities are after selecting its plants. Following equation

---

[22] Ideally I would estimate $\frac{\theta}{\sigma-1}$ from product-level bilateral export data or sales data in a particular country. The distribution of costs to serve market m in (6), together with the optimal pricing rule and the demand function implies that the product-level sales of a particular firm are distributed Fréchet with dispersion parameter $\frac{\theta}{\sigma-1}$. Data for the entire manufacturing sector would be most appropriate to use, as this is my selection criteria for the multinationals and trade data. When using car model sales data in five European countries available from Goldberg and Verboven (2001), I find an estimate of $\widehat{\frac{\theta}{\sigma-1}} = 1.02$.

(14) from the model, we can write the probability that a firm with core productivity level ϕ_t selects location set Z_t as

$$\Pr(Z = Z_t \mid \phi_t; \tilde{w}, \sigma_\epsilon, \mu_{\tilde{\eta}}, \sigma_{\tilde{\eta}}) = \int\limits_{\tilde{\eta}} \{E_\epsilon(\Pi(\phi_t, Z, \epsilon, \tilde{\eta}; \sigma_\epsilon, \tilde{w})) \geq E_\epsilon(\Pi(\phi_t, Z', \epsilon, \tilde{\eta}; \sigma_\epsilon, \tilde{w})) \quad \forall Z' \in \mathcal{Z}^i\} dF(\tilde{\eta}; \mu_{\tilde{\eta}}, \sigma_{\tilde{\eta}}). \tag{21}$$

This is a good place to discuss the timing assumption made in the model. The model attends to the possibility that, after the plants are established, the operations in every country are hit by productivity shocks whose realizations were not known to the firm when the production locations were established. The timing assumption simplifies the computation: the firm chooses its optimal location only conditional on its core productivity level, ϕ_t, the vector of fixed cost draws, $\tilde{\eta}_t$, and other parameters that are common across firms, $(\tilde{w}, \sigma_\epsilon)$, but not also conditional on the firm-country-specific productivity levels.

Since the firm-level data contains only the observations for German multinationals, but not for those firms that decided to operate only domestically, I also specify the probability that firm chooses location set Z_t conditional on choosing to become a multinational (which is the selection criteria of the data):

$$\Pr^*(Z = Z_t \mid \phi_t; \tilde{w}, \sigma_\epsilon, \mu_{\tilde{\eta}}, \sigma_{\tilde{\eta}}) = \frac{\Pr(Z = Z_t \mid \phi; \tilde{w}, \sigma_\epsilon, \mu_{\tilde{\eta}}, \sigma_{\tilde{\eta}})}{1 - \Pr(Z = Z_{\text{domestic}} \mid \phi; \tilde{w}, \sigma_\epsilon, \mu_{\tilde{\eta}}, \sigma_{\tilde{\eta}})}. \tag{22}$$

Aside from the information about the multinational's chosen locations, we observe its total output in each country in which it is active. Given a parameter guess of the unit input costs across countries, we can learn about the country-specific productivities of the multinational from the country-specific output levels. The productivity of firm t in country l is the product of the core productivity level, ϕ_t, and the firm-country specific productivity shifter, $\epsilon_{t,l}$. I denote this expression by $\psi_{t,l} = \phi_t \epsilon_{t,l}$. Let $r_{t,l}(\tilde{w}, Z_t, \psi_t) = \sum\limits_m s_{lm}(i_t, \phi_t, Z_t, \epsilon_t)$ denote the total revenue from sales to all countries of firm t in country l. Plugging in equations (9), (7), (8), and (10), we get the following equation for the output of firm t in country l:

$$r_{t,l}(\tilde{w}, Z_t, \psi_t) = \kappa \sum_m \frac{Y_m}{P_m^{1-\sigma}} \frac{(\tilde{w}_l \tau_{lm})^{-\theta} \psi_{t,l}^{\theta}}{\left(\sum\limits_{k \in Z_t} (\tilde{w}_k \tau_{km})^{-\theta} \psi_{t,k}^{\theta} \right)^{\left(\frac{\theta+1-\sigma}{\theta}\right)}}. \tag{23}$$

We have such an equation for every location in which firm t has a production location. Let r_t denote the vector of outputs of firm t in its production locations. Knowing the output of a firm in each of its locations and all other parameters allows us to pin down exactly its productivity level, $\psi_{t,l} = \phi_t \epsilon_{t,l}$, in each of its locations l. Proposition 2 states that given all other parameters, the solution to this system of equations is unique (the proof is in the appendix).

Proposition 2. *Let* $r : {}^K_{++} \times \mathcal{Z}^i \times \Psi \to {}^K_{++}$ *be the stacked vector of revenues as defined in equation* (23)*, where* K *denotes the number of countries in which firm* t *has a plant and* $\Psi = [\psi_{min}, \psi_{max}]^K$ *with* $\psi_{min} > 0$ *and* $\psi_{min} < \psi_{max} < \infty$*. Then for any triple* $\{r_t, \tilde{w}, Z\}$*, the vector* ψ *that solves* $r_t - r(\tilde{w}, Z, \psi) = 0$ *is unique.*

The likelihood function for each firm consists of the probability of its chosen location set and the density of the plant-specific revenues of the firm conditional on its location set and its core productivity level. I integrate out the core productivity level of each firm, which is observed by the firm but unobserved by the researcher. The likelihood function of the parameters $\Theta = \{\tilde{w}, \sigma_\epsilon, \mu_{\tilde{\eta}}, \sigma_{\tilde{\eta}}, \mu_\phi, \sigma_\phi\}$ given the observed data on location choice and revenues $\{Z_t, r_t\}_{t=1}^T$ can be written as:

$$L(\Theta; \{Z_t, r_t\}_{t=1}^T) = \prod_{t=1}^T \int \Pr{}^*(Z = Z_t \mid \phi; \tilde{w}, \sigma_\epsilon, \mu_{\tilde{\eta}}, \sigma_{\tilde{\eta}}) g(r_t \mid Z_t, \phi; \tilde{w}) dG(\phi; \mu_\phi, \sigma_\phi), \tag{24}$$

The first factor under the integral – the probability of the location choice – is specified directly in (22). The second factor – the density of the revenues – can be expressed in terms of the density of the plant-specific productivity shifters, $\epsilon_{t,l} = \frac{\psi_{t,l}}{\phi_t}$. It follows from Proposition 2 that the vector of revenues, r_t, can be inverted to get the vector of plant-specific productivity levels, ψ_t. The firm-location-specific productivity shifter $\epsilon_{t,l}$ is i.i.d. across firms and locations. I rewrite the likelihood function in (24) as

$$L(\Theta; \{Z_t, \psi_t\}_{t=1}^{T}) = \prod_{t=1}^{T} \int \Pr{}^*(Z = Z_t \mid \phi; \tilde{w}, \sigma_\epsilon, \mu_{\tilde{\eta}}, \sigma_{\tilde{\eta}}) \, |J_t(\phi, \tilde{w})| \prod_{l \in Z_t} h\left(\frac{\psi_{t,l}(\tilde{w})}{\phi} \mid \sigma_\epsilon \right) dG(\phi; \mu_\phi, \sigma_\phi), \tag{25}$$

where $h(\cdot \mid \sigma_\epsilon)$ denotes the univariate density of the firm-location-specific productivity shifter. The term $|J_t(\phi; \tilde{w})|$ is the determinant of the Jacobian which is included in the likelihood function because of the change of variables from the firm's revenues to the firm's productivity shifters.

Note that the firm-specific productivity shifter is not directly observed; we learn about the firm's productivity level in country k – given the current parameter guess and the observed country-specific output levels of the firm – from a system of equations that contains the output of the firm in each of its locations specified in (23). Therefore, I solve the following constrained optimization problem to estimate the parameters in which the objective function is the logarithm of the likelihood function specified in (25):

$$\begin{aligned} \max_{\Theta, \psi} \quad & \log L(\Theta; \{Z_t, \psi_t\}_{t=1}^{T}) \\ \text{subject to:} \quad & r_{t,l}(\tilde{w}, Z_t, \psi_t) = \kappa \sum_m \frac{Y_m}{P_m^{1-\sigma}} \frac{(\tilde{w}_l \tau_{lm})^{-\theta} \psi_{t,l}^{\theta}}{\left(\sum_{k \in Z_t} (\tilde{w}_k \tau_{km})^{-\theta} \psi_{t,k}^{\theta} \right)^{\left(\frac{\theta+1-\sigma}{\theta}\right)}} \end{aligned} \tag{26}$$

$$\forall \; t \in \{1, ...T\}, \; l \in \{1, ...N\} \text{ such that } l \in Z_t.$$

In summary, I use data on the chosen set of countries, Z_t, for each firm t – the probability of the location choice is the first term of the likelihood function – and the observed output in every country $r_{t,l}$ in which firm t is active – which is the left hand side of the constraints – to estimate the following parameters: the vector of unit input costs, $\tilde{w}$, the vectors that characterize the destination-country-specific distributions of fixed costs, $\mu_{\tilde{\eta}}$ and $\sigma_{\tilde{\eta}}$, the parameters for the core productivity distribution, μ_{ϕ} and σ_{ϕ}, and the parameter that characterizes the dispersion of the firm-country level productivity shocks, σ_{ϵ}. Given the structural parameters and the vector of location-specific outputs, the vector of the firm-country-specific productivity levels, ψ, solves the system of constraints. I control for unobserved heterogeneity in the core productivity levels of the firms and in the country-specific fixed cost draws.

The estimation is an implementation of the Mathematical Programming with Equilibrium Constraints (MPEC) procedure proposed by Su and Judd (forthcoming). They show that the estimator is equivalent to a nested fixed-point estimator in which the inner loop solves for the firm-country specific productivity levels, and the outer loop searches over parameters to maximize the likelihood. The estimator therefore inherits all the statistical properties of a nested fixed-point estimator. It is consistent and asymptotically normal as the number of firms tends to infinity and the number of simulation points used to evaluate the integrals rises proportionally to the number of firms.[23] As there are 1,711 positive firm-country output observations, the constrained optimization problem described in (39) has 1,711 equality constraints. In total, the data on the firm-output observations and the firms' location set choices is used to estimate 26 structural parameters. I compute standard errors via bootstrapping and use a logit-smoothed accept-reject simulator to evaluate the probability of location choice described in (21).[24]

---

[23] As the integrals are evaluated numerically in a finite sample with finite simulation draws, the Simulated Maximum Likelihood Estimator is necessarily biased (after taking logarithms of the Likelihood function). I find in a Monte-Carlo study of my estimation procedure that the bias is very small in practice for this problem.

[24] See Train (2009), Chapter 5, for a description of this and other methods of simulation.

1.3.3 Parameter Estimates

Table 1 displays the parameter estimates. I find that for German multinationals the variable costs of production (unit input costs) are systematically smaller in Germany than in foreign countries, which is not surprising given the low foreign output share abroad discussed in Section 1.3.1. The unit input costs in Germany are normalized to one. The smallest difference in unit input costs is found in Austria, in which German multinationals face only around seven percent larger variable production costs than at home. Within Western European countries, the production costs for German multinationals are largest in Italy and the United Kingdom (33-34 percent higher than in Germany). The production costs in the United States are around 42 percent higher than at home. The differences in production costs reflect both wage-level differences and efficiency losses that occur by producing outside the home country.

We can give the fixed costs a value interpretation as we observe the firms' output in Euro and, with CES preferences and monopolistic competition, we can easily determine that variable profits are proportional to output. Fixed costs are identified by observing the actual choice of production locations and variable profits together with the counterfactual scenarios of how variable profits would change if the firm altered its set of production locations. Note that my model does not distinguish between fixed costs to maintain a plant and sunk costs to establish a foreign plant. I use the estimates in Table 1 together with the structure of the model to calculate the mean fixed costs paid by firms that set up a production location in the respective countries. The calculation of the mean fixed cost conditional on having established a plant in the country is described in Appendix C and the results are displayed in Table 2. For most countries the estimated mean fixed cost of plants that were actually established is 6-8 million Euro. The paid fixed cost is estimated to be larger in Canada (12 million) and Belgium (18 million). The larger fixed cost estimates for these countries are in accordance with the data in Table 17 in Appendix B. Belgium has almost the same geographic location as the Netherlands and a similar local and surrounding market potential. While the number of German firms that have production locations in these countries is about the same, the output of plants in Belgium is much larger. This is reflected in the estimation of a lower variable production cost in Belgium and a larger fixed cost to keep the number of entrants at the same level with the Netherlands. Similarly, only a small number of firms have a plant in Canada, but they tend to have very large outputs.

1.3.4 Decomposing the sources of home bias in production

While the copious literature on the proximity-concentration trade-off has provided evidence for the presence of fixed costs, little is known about their quantitative importance. The parameter estimates above demonstrate both significant fixed costs to starting production in a foreign country and higher variable production costs abroad. In this section, I let firms re-optimize their location decisions as well as their decisions about which market to serve from which location, under different levels of fixed and variable costs.

Table 3 contains the results. The model effectively fits the average share of foreign output across firms. Both in the data and in the estimated model the average foreign output share is only around 0.30. If the unit input costs in the foreign countries were the same as in Germany, and there were no fixed costs for setting up foreign plants, then every firm would have a plant in each country and the average foreign output share across firms would be 0.88. The question arises as to whether fixed costs or larger variable production costs abroad are the more important barrier to foreign production. If unit input costs were equalized across countries, and fixed costs were kept at their estimated level, then firms would re-optimize their production locations and output decisions such that the foreign output share would be 0.68. If, instead, fixed costs were eliminated (and unit input costs held at their estimated level), the average foreign output share would rise even further to 0.72. Overall, I find that both fixed costs and differences in unit input costs significantly contribute to home bias in production.

While both factors have a large quantitative effect, fixed costs are slightly more important.

Table 1: MAXIMUM LIKELIHOOD ESTIMATES

		Unit input costs $\tilde{w}$	Fixed costs $\mu_{\tilde{\eta}}$
Country			
Austria		1.076	4.659
		(0.021)	(0.423)
Belgium		1.144	5.609
		(0.038)	(0.500)
Canada		1.324	5.067
		(0.080)	(0.571)
Switzerland		1.264	4.468
		(0.055)	(0.472)
Spain		1.223	3.912
		(0.018)	(0.335)
France		1.229	3.683
		(0.023)	(0.243)
United Kingdom		1.341	3.906
		(0.021)	(0.321)
Ireland		1.127	6.149
		(0.052)	(0.671)
Italy		1.334	3.978
		(0.039)	(0.309)
Netherlands		1.194	5.303
		(0.029)	(0.513)
United States		1.420	3.847
		(0.016)	(0.250)
S.d. log fixed cost, $\sigma_{\tilde{\eta}}$	2.1902		
	(0.320)		
Scale parameter productivity, μ_{ϕ}	1.1329		
	(0.017)		
Shape parameter productivity, σ_{ϕ}	5.1026		
	(0.620)		
S.d. log productivity shock, σ_{ϵ}	0.1844		
	(0.009)		
Log-Likelihood	-1.21E+004		
Number of firms, T	665		

Notes: Unit input costs in Germany are normalized to one. Standard errors in parentheses.

Table 2: Fixed cost by country

Country	Mean fixed cost of firms who set up a plant in the respective country in million Euro
Austria	7.107
	(1.338)
Belgium	18.063
	(7.515)
Canada	11.718
	(6.497)
Switzerland	5.814
	(2.715)
Spain	7.370
	(2.474)
France	7.037
	(1.423)
United Kingdom	6.653
	(1.966)
Ireland	6.069
	(1.665)
Italy	6.103
	(1.041)
Netherlands	7.499
	(2.332)
United States	6.799
	(1.257)

Notes: Standard errors in parentheses.

Table 3: Average share of foreign production in the output of German multinationals

Data	Model	No fixed costs	Same unit input costs as in Germany	No fixed costs and same unit input costs as in Germany
0.288	0.317	0.716	0.676	0.883
	(0.013)	(0.009)	(0.021)	(0.001)

Notes: Trade costs and price indices are held fixed. Standard errors in parentheses.

1.4 Calibration

In the second tier of my empirical inquiry, I focus on general equilibrium welfare analysis. In this section, I calibrate the key parameters to the general equilibrium outcomes of the model using data for many countries. Specifically, I calibrate trade costs, variable foreign production costs, and fixed costs of setting up foreign affiliates, to data on bilateral trade flows, the values of output of firms from country i in country l, and the estimates of the country-specific variable production costs of German multinationals from the previous section. The estimates of fixed and variable production costs for German multinationals from the previous section enable me to include both variable foreign production frictions and fixed costs in the analysis. I solve for the endogenous relative wages and price indices in every country.

1.4.1 Data

The analysis incorporates the same 12 Western European and North American countries as the previous section. Data on multinational production comes from Ramondo, Rodriguez-Clare, and Tintelnot (in process).[25] Gross manufacturing production and bilateral trade data comes from OECD STAN. Figures on labor endowments are drawn from the Penn World Tables, and statistics on educational attainment levels by country are from Barro and Lee (2010). Data on trade and multinational production (MP) are averages across the years 1996 to 2001, and the figures on population and educational attainment are for the year 2000.

---

[25] Unlike bilateral trade flow data, data on production activities of multinationals in foreign countries is documented only sporadically. They use available data from UNCTAD, BEA, Bundesbank and other sources on non-financial affiliate sales together with information on M&A from Thomson & Reuters to predict the aggregate total sales of (non-financial) affiliates from country i in country l.

1.4.2 Calibration procedure

The model delivers predictions for MP and trade shares, which I use as moments to calibrate the parameters.[26] The share of expenditures by consumers from country m that is spent on goods produced in country l ('trade-share') is

$$\xi_{lm} = \frac{X_{lm}}{Y_m}, \tag{27}$$

and the share of output produced by firms from country i in country l ('MP-share') is

$$\kappa_{il} = \frac{\sum\limits_{m} X_{ilm}}{\sum\limits_{m} X_{lm}}. \tag{28}$$

As an additional set of moments, I include the relative unit production costs of German firms in various countries that were estimated in Section 3. These are driven both by the foreign efficiency losses, γ, and endogenous relative wages, w. Relative variable production costs for German (j) firms in country l are $\frac{\tilde{w}_l}{\tilde{w}_j} = \frac{w_l \gamma_{jl}}{w_j}$. Note that I do not impose a restriction that multinational firms from other countries have the same variable production costs as German firms in foreign countries. Rather, the variable production costs of German multinationals relative to their production costs at home, are targeted to fit the estimates that were obtained with firm-level data in the previous section.

I estimate the parameters that characterize the trade costs between countries l and m, τ_{lm}; efficiency losses of foreign production, γ_{il}; and the distribution of fixed costs to set up plants in foreign countries as a firm from country i, $F^i(\eta)$. I make the following restrictions on the functional form for trade and foreign production iceberg costs:

$$\begin{aligned} \tau_{lm} &= \beta^{\tau}_{const}(dist_{lm})^{\beta^{\tau}_{dist}}(\beta^{\tau}_{contig})^{contig_{lm}}(\beta^{\tau}_{lang})^{language_{lm}} && \text{for } l \neq m \\ \gamma_{il} &= \beta^{\gamma}_{const}(dist_{il})^{\beta^{\gamma}_{dist}}(\beta^{\gamma}_{contig})^{contig_{il}}(\beta^{\gamma}_{lang})^{language_{il}} && \text{for } i \neq l. \end{aligned}$$

---

[26] The construction of MP and bilateral trade shares from the data is described in detail in Appendix B.

Domestic production iceberg costs and trade costs are normalized to one, while fixed costs for the domestic production location are set to zero. For all $l \neq i$, the fixed costs to set up a plant in location l for a firm from i (in units of labor in the destination country) are drawn independently across firms and locations. Formally, $\eta_l \sim \log\mathcal{N}(\ln f_{il}, \beta^f_\sigma)$, where

$$f_{il} \;=\; \beta^f_{const}(dist_{il})^{\beta^f_{dist}}(\beta^f_{contig})^{contig_{il}}(\beta^f_{lang})^{language_{il}} \qquad \text{for } i \neq l.$$

The mass of firms in country i, M_i, is set proportional to the product of population size and average years of schooling in country i, while the size of the labor force, L_i, is set proportional to the population in country i. As in the previous section, the value for the dispersion parameter of the product level productivity shock distribution, θ, is set to seven, and the elasticity of substitution, σ, is fixed to six. Following Chaney (2008), the core productivity levels for all firms are drawn from a Pareto distribution. Axtell (2001) estimates that US firm sizes are Pareto distributed with shape parameter 1.098. This suggests a shape parameter, $\vartheta = 5.5$, of the Pareto distribution for the core productivity levels, ϕ.[27] As I do not aim to explain an individual firm's data in this section, I abstract from the distribution of firm-location specific productivity shifters from here onwards.[28] This enables me to avoid evaluating numerically large-dimensional integrals without any explanatory power, since only aggregated information is used in this section.

The three sets of moments are stacked into the following vector:

---

[27] In the restricted version of my model with only a single production location for each firm (which is true for most firms in practice), the firm size distribution inherits the distribution of the core productivity levels and therefore will be Pareto distributed with shape parameter $\frac{\vartheta}{\sigma-1}$. This value is within one standard error from the point estimate of the shape parameter in the previous section based on German multinationals.

[28] I set $\epsilon_l = 1 \quad \forall l$ from here onwards.

$$d(\beta, w, A) = \begin{bmatrix} \xi(\beta, w, A) - \xi \\ \kappa(\beta, w, A) - \kappa \\ \frac{\tilde{w}(\beta,w,A)}{w_j} - \frac{\tilde{w}}{\tilde{w}_j}. \end{bmatrix}$$

This vector $d(\beta)$ is a 300×1 vector in which each element characterizes the distance between the respective model outcome (given the parameter vector β) and the outcome in the data. The calibration's objective is to minimize the sum of the squared differences between the model outcomes and the data targets for these outcomes. As we vary the parameter vector β, the equilibrium values of wages, profits (income), and price indices change. Note that in order for firms to choose their optimal policy, only the equilibrium wages and the market potential $A_m = \frac{Y_m}{P_m^{1-\sigma}}$ need to be known. Let $A_m(\beta, A, w)$ denote the market potential in country m that comes out of the policy functions of the firms and equations (19) and (17). Searching for an equilibrium, we seek a vector of market potentials A and wages w such that $A_m(\beta, A, w) = A_m \quad \forall m = 1, ..., N$, and the labor market clearing condition, $L_l^{\mathrm{d}}(\beta, A, w) = L_l \quad \forall l = 1, ..., N-1$, which is specified in (18), holds. As in the previous section, this suggests a constrained optimization procedure to calibrate the parameters.

Formally, the calibration solves the following constrained optimization problem:

$$\begin{aligned} &\min_{\beta,w,A} d(\beta, w, A)' d(\beta, w, A) \\ &\text{subject to:} \\ &A_m(\beta, w, A) = A_m \quad \forall m = 1, .., N \\ &L_l^{\mathrm{d}}(\beta, w, A) = L_l \qquad \forall l = 1, .., N-1. \end{aligned} \tag{29}$$

As only relative wages matter, I normalize one country's wage and drop one labor market clearing condition. As I have 300 moments as targets and only 13 parameters, an obvious question is how to weight these moments. I decide to be agnostic and give each moment the same weight.

1.4.3 Calibration results

The parameter estimates are displayed in the second column of Table 4. The iceberg loss in foreign production, γ_{il}, is relatively invariant to the distance between the firm's country of origin, i, and the country of production, l. Instead, fixed costs rise with distance.

Identification of the variable MP cost parameters comes from the moments on variable production costs for German multinational firms in different countries. In Figure 12 in the appendix, I compare German firms' variable production costs in various destination countries implied by the calibrated model with the estimates from the firm-level data in the previous section; the numbers are closely matched. The identification of the fixed cost parameters comes from the moments on bilateral MP shares. Note that the calibration results for the fixed cost parameters imply that the fixed cost is rising in distance between the source and destination country, which was not a pattern of the German firm-level estimates in the previous section. However, data for many more country pairs is used for this section. The estimates reflect that bilateral MP declines with distance.

I also calibrate a restricted version of my model in which the fixed costs to set up foreign plants are set to infinity. As no multinational production arises under this restriction, I call this the 'pure trade model.' It is observationally equivalent to the model by Anderson and van Wincoop (2003) in terms of aggregate trade flows between countries.[29] Using only the trade shares as the targets, I calibrate the same gravity parameters of the trade cost function for this restricted model. The trade cost estimates are very similar across the global production and pure trade models (first column of Table 4). However, the distance coefficient is slightly larger and the constant slightly lower in the pure trade model.

The sum of the squared deviations from the MP and trade shares in the data and calibrated model are displayed under 'Norm MP fit' and 'Norm trade fit' in Table 4. The model of global

---

[29] If fixed costs of exporting were included, the restricted model with no multinational production would be equivalent to Chaney (2008).

Table 4: CALIBRATED PARAMETERS

	Pure trade model	Global production model
Trade cost		
constant, β^{τ}_{const}	0.722	0.789
distance, β^{τ}_{dist}	0.139	0.121
language, β^{τ}_{lang}	0.922	0.929
contiguity, β^{τ}_{contig}	0.934	0.925
Variable MP cost		
constant, β^{γ}_{const}		1.259
distance, β^{γ}_{dist}		0.006
language, β^{γ}_{lang}		0.962
contiguity, β^{γ}_{contig}		0.963
Fixed MP cost		
constant, β^{f}_{const}		0.089
distance, β^{f}_{dist}		0.073
language, β^{f}_{lang}		1.025
contiguity, β^{f}_{contig}		1.105
dispersion, β^{f}_{σ}		0.299
Norm trade fit	0.258	0.262
Norm MP fit		0.158

production in this paper fits trade flows similarly well to a pure trade model and additionally fits multinational production. I present scatter plots on the model's fit of trade and MP shares in Appendix D.

1.4.4 Fit of export platform shares

The calibration is targeted to fit bilateral trade and MP shares, as well as the relative variable production costs of German multinationals in various countries. How does the calibrated model perform with respect to moments it did not try to fit? I use data from the BEA on the export platform share of US multinational firms in all countries other than the US included in my estimation to compare the model's predictions with the actual data. The fit is good and displayed in Figure 2. Notice that existing work on multinational production by Ramondo and Rodriguez-Clare (2012) and Arkolakis, Ramondo, Rodriguez-Clare, and Yeaple (2012) get close to fitting the export platform share of US multinationals only under special cases [e.g., uncorrelated productivities across countries for the same idea in Ramondo and Rodriguez-Clare (2012)]. The key difference between the framework in this paper and in their work is that this paper incorporates fixed costs to establish foreign plants. Fixed costs cause firms to concentrate their production in fewer locations, which generates export platform sales. In Appendix E, I show for a symmetric world how additional production locations lead to lower shares of export platform sales. I also calibrate a restricted model with zero fixed costs to set up a foreign plant. Under this restriction, every firm operates a plant in every country. For this model I use as targets both trade and MP shares, but not the variable production cost estimates for German multinationals.[30] In F, I present the results for the restricted model with only variable MP costs but no fixed costs. Table 21 in F shows that the export platform shares are on average 1.4 times larger in the benchmark global production model than in the restricted model without fixed costs.

---

[30] These estimates would not be consistent with this kind of model, as they were derived from a model with positive fixed costs to set up foreign plants.

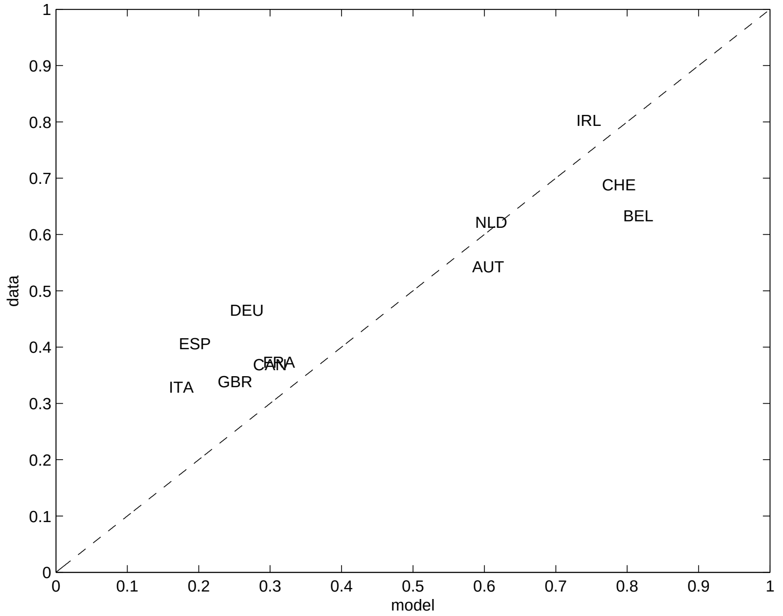

Figure 2: Export platform shares for US multinationals - data and model

1.5 General equilibrium counterfactuals

I proceed by conducting counterfactual analysis based on the calibrated global production model. In each counterfactual, the general equilibrium is resolved for the new parameter values. I begin with an analysis of an important current policy issue.

1.5.1 Potential effects from a Canada-EU trade and investment agreement

The EU and Canada are currently negotiating a trade and investment agreement: CETA.[31] What would be the effects of such an agreement – if it is signed – on the signatories and the U.S.? My setup is particularly suitable for addressing this question. As multinational firms tend to concentrate production in just a few locations and serve the rest of the world via export platform sales, investment liberalization agreements may have particularly strong third-country effects. Third-country effects

---

[31] Comprehensive Economic and Trade Agreement (CETA). Currently discussed measures to remove investment barriers between the EU and Canada include harmonization in taxes and regulation, opening of capital markets, the removal of barriers to executive labor mobility, and improvement in access to information for foreign investors. More information from the Canadian government can be found here: `http://www.international.gc.ca/trade-agreements-accords-commerciaux/agr-acc/eu-ue/negotiations-negoiations.aspx?lang=eng&view=d`. More information from the EU commission can be found here: `http://ec.europa.eu/enterprise/policies/international/cooperating-governments/canada/index_en.htm`.

arise in pure trade models due to the terms of trade effects. With multinational production – in addition to the terms of trade effects – an additional effect arises: the firm can directly move its production locations and volume between countries, so the effect on third countries may be stronger as firms respond to changes in the global bilateral investment cost structure. For instance, a European firm may want to have only one plant in North America. As investment barriers to Canada fall, this firm may move its plant from the United States to Canada. This outcome would be missed by models that do not take into account export platforms, since without export platforms the firm's decision to establish a plant in a country would be independent across countries.

Suppose a deep investment agreement can be reached that lowers both variable and fixed MP costs between the EU countries and Canada by 20 percent. Table 5 displays the difference in the MP-shares before and after the liberalization, as well as the relative change in welfare. The aggregate MP-share of EU countries in Canada would increase from 9 to 32 percent. US firms would react to higher relative Canadian wages and reduce their investment in Canada such that the share of US production in Canada would fall from 21 to 9 percent. Finally, the total foreign production in Canada would increase by a factor of 1.3. Simultaneously, part of the EU countries' increased investment in Canada would crowd out their previous production in the US. EU countries' production share in the US would fall from 5.58 to 5.11 percent. In relative terms, this is a decline by 7 percent. Canadian firms would react to higher relative wages in Canada and increase their activities in the US, but not to the same extent that EU firms would decrease their activities. The overall share of foreign production in the US would fall by 6 percent. Of all countries, Canada would experience the largest welfare gains. The welfare gains in EU countries would be positive but moderate in size and larger for smaller countries. The US and Switzerland would experience small welfare losses. The US economy is large enough that even though the diversion of EU investment from the US to Canada would be substantial, it would not be affected much in terms of aggregate welfare.

As a comparison to the potential effects from CETA, which is currently under negotiation, I

Table 5: COUNTERFACTUAL CHANGES OF LOWER EU-CANADA MP COSTS

	Difference in inward MP-shares		Rel. welfare
	Canada	United States	
Canada	-10.56	0.05	102.45
EU countries	23.23	-0.47	100.07 - 100.19
Switzerland	-0.19	0.01	99.90
United States	-12.48	0.41	99.95

Notes: Counterfactual: Reduction in variable and fixed MP costs between EU and Canada by 20 percent. First two columns: Differences in MP shares, κ_{il}, before and after the counterfactual change (column: destination l, row: source i).

also compute the potential effects from a hypothetical EU-US agreement that would lower variable and fixed foreign production costs between the signatories by the same proportion. As expected, the effects on the non-signatory partners from such an agreement would be even larger: the share of EU multinationals' production in Canada would fall from 9 to 6 percent, and the welfare in Canada would fall by almost half of a percent. Table 22 in the appendix contains the predicted outcomes for such an EU-US agreement.

I proceed with the analysis of a classic question in the trade literature, which has been studied by Eaton and Kortum (2002) among others: how large are the benefits of foreign technology?

1.5.2 The benefits of foreign technology

It has been widely documented that countries' technologies evolve over time.[32] How much do countries benefit from foreign technology improvement, and how do our estimates of these gains differ between models that allow for multinational production and those that do not? To answer these questions, I follow Eaton and Kortum (2002) and study what happens to welfare in foreign countries if all US firms improve their core productivity levels by 20 percent.[33] Similarly to their paper, I find that in a pure

---

[32] For example, by Bernard and Jones (1996) and Levchenko and Zhang (2011).

[33] The core productivity level of the firms is Pareto distributed; I multiply the draws for US firms by 1.2.

trade model the percentage gains decay dramatically with distance and size (see the results in Table 6). With multinational production, an additional source of gains for foreign countries arises: multinational firms use the better technology in their foreign plants and crowd out some of the production of less productive domestic firms. Hence, the average productivity in foreign countries rises, which in turn lowers those countries' consumer price indices. Interestingly, in the global production model, the welfare gain from a US technology improvement in foreign countries is about an order of magnitude larger than in the pure trade model. Still, Canada – as the neighboring country – benefits most from a US technology improvement, as the costs both to ship goods and to produce abroad rise with distance.

Table 6: Gains from US technology improvement

	Relative to benchmark		Relative to US gains	
	Pure trade model	Global production model	Pure trade model	Global production model
Austria	1.0009	1.0319	0.45	14.52
Belgium	1.0005	1.0205	0.26	9.34
Canada	1.007	1.0630	3.53	28.69
Switzerland	1.0007	1.0203	0.37	9.26
Germany	1.0003	1.0155	0.15	7.07
Spain	1.0005	1.0310	0.26	14.11
France	1.0003	1.0170	0.17	7.76
United Kingdom	1.0006	1.0299	0.32	13.60
Ireland	1.0022	1.0460	1.12	20.93
Italy	1.0004	1.0240	0.18	10.92
Netherlands	1.0006	1.0286	0.32	13.03
United States	1.1987	1.2195	100.00	100.00

Notes: Counterfactual: Productivity improvement of all firms that originated in the United States by 20 percent. Columns 3 and 4: Welfare gains by country in percent relative to welfare gains in the United States.

The United States also gains more from its firms' technology improvement when multinational production is allowed (US welfare increases by a factor of 1.219 instead of 1.198). Without multinational production, the change in real profits is proportional to the change in real wages. With multinational production, US multinational firms can benefit from the relatively cheaper labor in foreign countries and receive larger profits, which raises the US consumer's income. Note that the US welfare gains exceed 20 percent under a unilateral technology improvement for US firms; by contrast, if all countries had improved their firms' technology by 20 percent, every country's welfare would have simply risen by 20 percent (not displayed in the table). As a consequence, for different initial conditions, e.g. after the increase in productivity of US firms, the welfare gains for the US from a technology improvement in all countries other than the US could be negative as the US profits would decline.[34]

To summarize, in the global production model, since the level of technology used in a country's production is endogenous, an overall improvement to the US firms' technology improves the technology used in foreign countries' production. In other words, multinational firms enhance the spread of technology to foreign countries.[35] Thus, the gains from foreign technology improvements are dramatically underestimated if multinational production is omitted from the analysis.

1.5.3 The gains from trade, multinational production, and openness

I continue by studying the gains from trade, multinational production, and openness. I use the calibrated model as a starting point and address how welfare, wages, and profits change when the costs of trade, the costs of multinational production, or the costs of both trade and multinational production are set to prohibitively high levels. I start by comparing the benchmark calibrated model with a hypothetical world in which costs of international trade are infinite.

---

[34] In Appendix G I present levels of parameters for which my model can be solved analytically and show that whether the welfare gain of the country whose firms improved their technology exceeds the rate of technology improvement depends on the increase in global market share due to the technology improvement.

[35] The gains from foreign technology improvements may be even larger when spillovers are included. Poole (forthcoming) finds evidence for knowledge transfers from multinational to domestic firms in Brazilian matched worker-establishment data. Alfaro and Chen (2012) estimate that around 2/3 of the gains from multinational production arise through technology spillovers.

The gains from trade

I define as the gains from trade in the global production model the change in welfare from the model with the benchmark parameters to a model with infinite trade costs.[36] While the trade in goods is prohibited in this counterfactual world without trade, I allow for the flow of remittances between countries; if trade is prohibited, current account balance implies that for each country the total inflows equal the total outflows of profits.

I compare my model's predictions with those from a workhorse model in multi-country trade analysis, such as Anderson and van Wincoop (2003) and Eaton and Kortum (2002), which abstract away from multinational production. How do the gains from trade in these pure trade models differ from the gains from trade in this model, which includes trade and multinational production? Suppose a researcher would use the same starting point, that is, observe the trade flows that are implied by my global production model and then use a pure trade model to evaluate the gains from trade. Following Arkolakis, Costinot, and Rodriguez-Clare (2012), the gains from trade for country j in a pure trade model such as that of Anderson and van Wincoop (2003) can simply be calculated as $\xi_{jj}^{\frac{1}{1-\sigma}}$. I use $\sigma = 6$ for both models, which implies both a reasonable mark-up for firms and is in the range of estimates for the trade elasticity in pure trade models.[37]

The results in Table 7 display two striking patterns: first, they suggest that the gains from trade are slightly larger in a pure trade model than in a model with trade and MP (though the difference is small). As expected, smaller countries benefit more in both models from trade openness than do larger countries. Second, for each country the increase in real profits because of trade is about more than four times larger than the increase in real wages. Trade enables firms to exploit comparative advantage, concentrate their production in a few locations, and economize on fixed costs. Without trade, firms need to incur fixed costs for every market they want to serve, as this requires them to

---

[36] Welfare in country j is equal to real income, $\frac{Y_j}{P_j}$.

[37] The survey by Anderson and van Wincoop (2004) finds that most estimates of the trade elasticity in the literature range from five to ten.

Table 7: Gains from Trade

	Global Production model			Pure Trade model
	Welfare change	Real profit change	Real wage change	Welfare / Real wage change
Austria	1.193	1.585	1.154	1.208
Belgium	1.344	1.837	1.296	1.379
Canada	1.098	1.356	1.068	1.108
Switzerland	1.317	1.843	1.268	1.342
Germany	1.060	1.175	1.043	1.068
Spain	1.050	1.188	1.031	1.054
France	1.075	1.232	1.053	1.084
United Kingdom	1.059	1.201	1.040	1.066
Ireland	1.306	1.795	1.263	1.324
Italy	1.043	1.155	1.027	1.048
Netherlands	1.189	1.524	1.151	1.208
United States	1.012	1.035	1.008	1.013

Notes: A number in this table represents the outcome from the benchmark model divided by the outcome from the same model with no trade.

establish a local plant. Further, each product needs to be produced locally even though, if trade were available, a plant in another country would have comparative advantage in producing some of these products.

How do prohibitive trade costs affect the share of production in each country that is conducted by foreign firms? Generally, the outcome is unclear; on the one hand, infinite trade costs eliminate the surrounding market potential of foreign plants, which reduces the total market potential of the foreign plant, since the share of export platform sales is large. With fixed costs to establish foreign plants, a lower market potential – ceteris paribus – should lead to fewer foreign plants being established. On the other hand, without trade, multinational production is the only means of serving a foreign market. I find that the level of multinational production increases when trade is prohibited. Another feature of a world without trade is that not all varieties are sold to every market, just as with fixed costs of establishing foreign plants not every firm is a multinational. This explains why the gains from trade are still large in my global production model, and multinational production can only insufficiently substitute for trade. My findings differ from those Ramondo and Rodriguez-Clare (2012), whose calibrated model indicates that the gains from trade were systematically larger in a world with MP than in a world without MP. However, in their model, aside from the absence of increasing returns at the plant level, the input bundle of a multinational firm abroad is a CES aggregate of the local input bundle in the country of the affiliate and the input bundle in the home country. Therefore, the availability of trade improves the cost structure of the foreign affiliates in their model.[38]

Overall, this section suggests that if one wants to evaluate the gains from trade, the use of a pure trade model that ignores multinational production provides results that are close to those from a more general model with trade and multinational production. Furthermore, firms benefit from trade, as trade enables them to economize on fixed costs and to exploit comparative advantage in production. I continue by comparing the outcomes from the benchmark calibrated model with a hypothetical world in which costs of multinational production are prohibitive.

The gains from multinational production

I define as the gains from multinational production the change in real income, Y_j/P_j, one finds when going from a version of my model with infinite costs to multinational production to the model with the calibrated parameters. The relative changes in the outcomes for welfare, real wages, and real profits are displayed in Table 8.

I find that the welfare gains from multinational production are smaller than the welfare gains from trade. Note that real wages and real profits respond quite differently to the availability of multinational production. The changes in real wages from allowing multinational production are comparable to the changes in real wages from allowing trade, which reflects that multinational production substantially lowers the price index; compared to the changes in real wages from allowing trade, they are

---

[38] See the paper by Head and Ries (2004) for an overview of intra-firm trade; if firms produce single final products – as assumed in Ramondo and Rodriguez-Clare (2012) – it is sensible to interpret the intra-firm trade from the home country to the foreign affiliate as shipments of intermediates. If firms produce multiple final products – as assumed in my paper – these intra-firm sales can arise simply from the comparative advantage of the home country in producing some of the firms' products.

Table 8: Gains from Multinational Production

	Global Production model		
	Welfare change	Real profit change	Real wage change
Austria	1.017	0.740	1.073
Belgium	1.015	0.746	1.069
Canada	1.021	0.779	1.069
Switzerland	1.018	0.731	1.075
Germany	1.006	0.879	1.031
Spain	1.011	0.817	1.049
France	1.008	0.857	1.038
United Kingdom	1.011	0.832	1.047
Ireland	1.021	0.684	1.088
Italy	1.009	0.844	1.042
Netherlands	1.011	0.783	1.056
United States	1.002	0.956	1.012

Notes: A number in this table represents the outcome from the benchmark model divided by the outcome from the same model with no multinational production.

larger for some countries (e.g. US), and lower for others. However, the welfare gains are systematically lower as real profits fall. Note that the effect of multinational production on real profits is ambiguous. On the one hand, multinational production raises production efficiency and lowers a multinational's marginal cost curve. On the other hand, the aggregate price index falls, which lowers demand, and multinational firms bear the burden of fixed costs for multinational production. Note that if fixed costs were zero, real profits would rise unambiguously. Furthermore, in a world in which countries are asymmetric in the ratio of labor size to mass of firms, real profits may rise in the country with a particularly large ratio of firms to labor size. In the calibrated model, the mass of firms is roughly proportional to the labor force, as human capital differences between the selected countries are small. Real profits in smaller countries tend to fall more than in larger countries because of the availability of multinational production. Therefore, fixed costs are vital not only for explaining firm-level global production choices and matching the aggregate data on export platform sales, but also for understanding the overall gains and distributional effects from multinational production.

In the next section, I describe the gains from openness.

The gains from openness

Finally, I evaluate how welfare changes if both trade and multinational production are shut down, so that each country operates in autarky. Non-surprisingly, as neither trade nor multinational production can substitute for the absence of the other, the welfare losses of autarky are more substantial than if only trade or multinational production were shut down. The results are displayed in Table 9. For most countries, the change in real wages is similar to the change in welfare, as real profits are roughly unchanged. Real profits tend to increase through the availability of trade, and tend to fall through the availability of multinational production. As expected, small countries benefit substantially more from openness than do large countries.

Table 9: Gains from Openness

	Global Production model		
	Welfare change	Real profit change	Real wage change
Austria	1.262	0.918	1.331
Belgium	1.440	1.058	1.516
Canada	1.154	0.880	1.208
Switzerland	1.414	1.015	1.494
Germany	1.083	0.947	1.110
Spain	1.076	0.870	1.117
France	1.104	0.939	1.137
United Kingdom	1.089	0.896	1.127
Ireland	1.400	0.938	1.492
Italy	1.065	0.891	1.100
Netherlands	1.245	0.965	1.301
United States	1.018	0.970	1.027

Notes: A number in this table represents the outcome from the benchmark model divided by the outcome from the same model with no multinational production and no international trade.

Remarks

While the gains in real wages due to multinational production are similar to the gains in real wages due to trade, the welfare gains from multinational production are much smaller, since real profits fall substantially due to the fixed costs of establishing foreign plants. Note that the welfare gains from multinational production may increase considerably if free entry is taken into account. Lower real profits under multinational production would lead to less entry and henceforth less expenditures on entry costs, which can substantially change the calculation of the welfare gains from multinational production. In future research it would be interesting to incorporate free entry into my model and estimation. Conceptually, this extension is straightforward. For many potential applications free entry does not matter, but it is likely to affect the calculation of the overall gains from multinational production.

1.6 Conclusion

My paper contributes to the literature on international trade and multinational firms by developing a new framework that is less restrictive in its assumptions about where multinational firms can produce and sell, and what cost structure they face. In particular, my model accounts for both export platform sales and fixed costs of establishing foreign plants. Existing quantitative models of trade and multinational production have proven tractable only after excluding many of the strategies that firms actually use or shutting down mechanisms that are almost universally thought to be important. My framework can be used for a wide range of possible empirical applications with either firm-level or aggregate data.

Multinational firms have been characterized as footloose and free to reorganize their global operations as the global economic environment changes [Caves (1996)].[39] My estimates on the variable

---

[39] See also *The Economist* on March 25, 2004: "Footloose firms: Are global companies too mobile for workers' good?"

efficiency losses to foreign production and increasing returns at the plant-level suggest that the view of footloose multinationals is inaccurate. Differences in variable production costs across countries and fixed costs of establishing foreign plants turn out to be similarly important barriers to foreign production for German multinational firms.

General equilibrium analysis reveals that multinational firms play an instrumental role in the transmission of technology improvements to foreign countries. As multinationals have the ability to re-optimize their production locations and output decisions when the cost structure across countries changes, trade and investment agreements can have a significant third-country effect, which would be missed if multinational firms are excluded from the analysis or modeled in a more restrictive way. My findings accord with the common perception that countries compete for multinational firms and that small countries would be hurt disproportionately if a country nearby were to improve the conditions for multinationals from other countries, as this policy-change induces a firm-delocation effect. My framework can be used to quantitatively investigate the implications of such policies or other changes to the economic environment.

Chapter 2

Borders, Geography, and Oligopoly: Evidence from the Wind Turbine Industry

2.1 Introduction

Distance and political borders lead to geographical and national segmentation of markets. In turn, the size and structure of markets depend crucially on the size and nature of trade costs. A clear understanding of these costs is thus important for assessing the impact of many government policies.[1] Since the seminal work of McCallum (1995), an extensive literature has documented significant costs related to crossing national boundaries. Estimated magnitudes of border frictions are so large that some researchers have suggested they are due to spatial and industry-level aggregation bias, a failure to account for within-country heterogeneity and geography, and cross-border differences in market

---

[1] Policy relevance goes beyond trade policies. According to Obstfeld and Rogoff (2001), core empirical puzzles in international macroeconomics can be explained as a result of costs in the trade of goods. Romalis (2007) shows that the interaction of tax policies and falling trade costs was key to the rapid growth of Ireland in the 1990s. Effectiveness of domestic regulation in some industries may hinge on the extent of trade exposure, as shown by Fowlie, Reguant, and Ryan (2011) for the US Portland cement industry.

structure.[2] To avoid these potentially confounding effects, we use spatial micro-data from wind turbine installations in Denmark and Germany to estimate a structural model of oligopolistic competition with border frictions. Our main findings are: (1) border frictions are large within the wind turbine industry, (2) fixed and variable costs of exporting are both important in explaining overall border frictions, and (3) these frictions have a substantial impact on welfare. Finally, we show that commonly used price difference regressions are likely to produce misleadingly high estimates of the impact of national boundaries.[3]

Our ability to infer various components of trade costs is a result of our focus on a narrowly defined industry: wind turbine manufacturing. In addition to being an interesting case for study in its own right due to the growing importance of wind energy to Europe's overall energy portfolio, the wind turbine industry in the European Union (EU) offers an excellent opportunity to examine the effects of national boundaries on market segmentation. First, we have rich spatial information on the location of manufacturers and installations. The data are much finer than previously used aggregate state- or province-level data. The use of disaggregated data allows us to account for actual shipping distances, rather than rely on market-to-market distances, to estimate border costs. Second, the data contain observations of both domestic and international trade. We observe active manufacturers on either side of the Danish-German border, some of whom choose to export and some of whom do not, allowing us to separate fixed and variable border costs. Third, intra-EU trade is free from formal barriers and large exchange rate fluctuations. It is also subject to wide-ranging efforts to minimize informal barriers.[4] By the Single European Act, national subsidies are directed only toward the generation of renewable electricity and do not discriminate against other European producers of turbines. The border costs

---

[2] See Hillberry (2002), Hillberry and Hummels (2008), Broda and Weinstein (2008) and Gorodnichenko and Tesar (2009).

[3] Measures of the "border width" using observed price dispersion have been proposed by Engel and Rogers (1996), Parsley and Wei (2001), and Broda and Weinstein (2008) among others.

[4] All tariffs and quotas between former European Economic Community members were eliminated by 1968. The Single European Act came into force in 1987 with the objective of abolishing all remaining physical, technical and tax-related barriers to free movement of goods, services, and capital within the EU until 1992. Between 1986 and 1992, the EU adopted 280 pieces of legislation to achieve that goal.

in this setting are therefore due to factors other than formal barriers to trade and exchange rate fluctuations.

Despite substantial formal integration, the data indicate substantial market segmentation between Denmark and Germany. Examining the sales of turbines in 1995 and 1996, we find that domestic manufacturers had a substantially higher market share than did foreign manufacturers. For example, the top five German manufacturers possessed a market share of 60 percent in Germany and only 2 percent in Denmark. The market share of Danish producers drops by approximately 30 percent at the border.

What are the sources of cross-national market segmentation? On one hand, a cursory glance at our data suggests that national borders affect the decisions of firms to enter the foreign market. To be specific, only one of the five large German firms exports to Denmark. On the other hand, all five large Danish firms have sales in Germany, but their market share is substantially lower in the foreign market and drops discontinuously at the border.

The difference in participation patterns across the border reflects fixed costs faced by exporting firms. The change in market share at the border may be generated by differences in competition (e.g., differences in the set of competitors and their underlying characteristics) or by higher variable costs for foreign firms. To explain differences in market shares along extensive and intensive margins, we propose a model of cross-border oligopolistic competition that embeds costs for exporting as primitives and controls for other sources of cross-border differences. This allows us to infer the costs that exporter firms face and quantify their impact on market shares, profits, and consumer welfare through counterfactual analysis.

In our model, firms are heterogeneous in their production costs, foreign market entry costs, and distance to project sites. To become active in the foreign country, firms must pay a *fixed border cost* specific to them. Fixed border costs include maintaining a foreign sales force, developing technology to connect turbines to the foreign electricity grid, and gaining certification for turbine models in foreign countries. The model incorporates two types of variable costs for supplying a project: First, all firms face a *distance cost* that increases with the distance between the location where they produce the turbines and the location of the project. Second, exporters pay an additional *variable border cost* to supply projects in their foreign market. While the distance friction is analogous to the standard iceberg cost in trade models, the variable border cost captures additional hurdles that exporters face independent of shipping distance. These hurdles may arise due to language or cultural differences between purchasers and manufacturers, legal complications due to the use of cross-border contracts, or the need to interact with multiple national transportation authorities to authorize turbine delivery. One of our objectives is to gauge the importance of each type of cost in segmenting the markets.

The model has two stages: In the first stage, turbine producers decide whether or not to export. This depends on whether their expected profit in the foreign market exceeds their fixed border cost. As a result, the set of competing firms changes at the border. In the second stage, turbine producers observe the set of active producers in each market and engage in price competition for each project. This gives rise to a spatial model of demand for wind turbine installations. Project managers face a discrete choice problem: they observe price bids and pick the producer that maximizes their project's value. However, as in many business to business industries, transaction prices are not observable to the econometrician. To surmount this challenge, we assume manufacturers choose prices (and hence, markups) for each project on the basis a profit maximization condition derived from our model.[5] Each producer's costs depend on the location of the project through both distance and the presence of a border between the producer and the project, which are observable to all firms and the econometrician. In equilibrium, each firm takes into account the characteristics of its competitors when choosing its own price. The model thus delivers endogenous variation in prices, markups, and market shares across points in space. Our data informs us about the suppliers of all projects. We estimate the model by

---

[5] Our approach uses profit maximization to derive a structural connection between quantities and prices when only quantities are observed. As such, it can be seen as complementary to the work of Thomadsen (2005) and Feenstra and Levinsohn (1995), who use a profit maximization condition to derive a relationship between prices and quantities when only prices are observable.

maximizing the likelihood of correctly predicting these outcomes.

Our results indicate that there are substantial costs to sell wind turbines across the border between Denmark and Germany. We find that the variable border costs are roughly equivalent to moving a manufacturer 400 kilometers (250 miles) further away from a project site. Given that the largest possible distance from the northern tip of Denmark to the southern border of Germany is roughly 1,400 kilometers (870 miles), this is a significant cost for foreign firms. Removing fixed costs of foreign entry, such that all firms compete on both sides of the border, raises the market share of German firms in Denmark from 2 to 12 percent; further eliminating variable border costs raises that market share from 12 percent to 22 percent. Counterfactual analysis provides further insights into the welfare effects of borders. A hypothetical elimination of all border frictions raises consumer surplus by 10.4 and 15.3 percent in Denmark and Germany, respectively. Overall, total surplus increases by 5 percent in Denmark and 10 percent in Germany.

By estimating a structural oligopoly model that controls for internal geography and firm heterogeneity, this paper adds to the empirical literature on trade costs.[6] Early contributions by McCallum (1995) and Anderson and van Wincoop (2003) use data on interstate, interprovincial, and international trade between Canada and the United States to document a disproportionately high level of *intranational* trade between Canadian provinces and U.S. states after controlling for income levels of regions and the distances between them. Engel and Rogers (1996) find a high level of market segmentation between Canada and the United States using price data on consumer goods. Gopinath, Gourinchas, Hseih, and Li (2010) use data on retail prices to document large retail price gaps at the border using a regression discontinuity approach. Goldberg and Verboven (2001, 2005) find considerable price dispersion in the European car market and some evidence that the markets are becoming more integrated

---

[6] See Anderson and van Wincoop (2004) for a survey of the empirical literature.

over time.[7]

Rather than inferring a “border effect” or “width of the border” based on differences between intra- and international trade flows or price differentials, we use spatial micro-data on shipments to estimate trade costs which induce market segmentation. By doing this, we addresses several critiques raised by the literature. Hillberry (2002) and Hillberry and Hummels (2008) show that sectoral and geographical aggregation lead to upward bias in the estimation of the border effect in studies that use trade flows. **?** emphasize the importance of controlling for internal distances. In a similar fashion, Broda and Weinstein (2008) find that aggregation of individual goods’ prices amplifies measured impact of borders on prices. Our data enables us to calculate the distances between consumption and production locations for a narrowly defined product. That, in turn, enables us to separate the impact of distance from the impact of the border.

Our structural model of oligopolistic competition controls for differences in market structure and competitors’ costs across space. The estimates from our structural model can thus be directly interpreted as costs that exporters must pay to market their products abroad.[8] This approach addresses the concern of Gorodnichenko and Tesar (2009) that model-free, reduced-form estimates fail to identify the border effect. To highlight the importance of using disaggregated data and a structural model, Section 2.6 presents an experiment based on our estimated model in which we follow the common procedure in the literature to estimate the width of border using a reduced form model of price dispersion between markets. For all of the exporting firms in our data, this estimated width is over 60 percent higher than the true width used to generate the data. A comparison of structural and reduced form equations illustrates the sources of bias, which include measurement error in transport

---

[7] The interest in border frictions partially stemmed from the realization that prices of tradable goods do not immediately respond to exchange rate fluctuations, leading to substantial price differentials across countries. The exchange rate between Germany and Denmark was extremely stable during our sample period: the median month-to-month variation is 0.23 percent. So, this source of border frictions is absent from our environment.

[8] It may also be that preferences change at the border such that consumers act on a home bias for domestic turbines. In our setting, demand comes from profit maximizing energy producers buying an investment good, so we expect that demand driven home bias are less likely to occur than they would for a consumption good. Within our model, home bias in consumer preferences cannot be separately identified from border costs. Alternatively, we can interpret our results as incorporating the additional costs exporting firms must incur to overcome any home-bias in preferences.

costs and failure to accommodate for markup differences.[9]

In summary, our industry-specific focus has three major advantages: First, the use of precise data on locations in a structural model allows for a clean identification of costs related to distance and border. Second, the model controls for endogenous variation in markups across markets within and across countries based on changes in the competitive structure across space. Third, by distinguishing between fixed and variable border costs, we gain a deeper insight into the sources of border frictions than we do from studies that use aggregate data.

In the following section, we discuss our data and provide background information for the Danish-German wind turbine industry. We also present some preliminary analysis that is indicative of a border effect. Section 2.3 introduces our model of the industry. We show how to estimate the model using maximum likelihood with equilibrium constraints and present the results in Section 2.4. In Section 2.5, we perform a counterfactual analysis of market shares and welfare by re-solving the model without fixed and variable border costs. Section 2.6 uses market-to-market price differentials from our model in a reduced-form regression to relate our approach to studies that estimate border frictions based on the law of one price. We conclude in Section 2.7 with a discussion of policy implications.

2.2 Industry Background and Data

Encouraged by generous subsidies for wind energy, Germany and Denmark have been at the forefront of what has become a worldwide boom in the construction of wind turbines. Owners of wind farms are paid for the electricity they produce and provide to the electric grid. In both countries, national governments regulate the unit price paid by grid operators to site owners. These “feed-in-tariffs” are substantially higher than the market rate for other electricity sources. Important for our study is that public financial support for this industry is not conditional on purchasing turbines from domestic

---

[9] For example, a key source of measurement error is that the reduced form specification proxies for differences in supplier-to-market distance with market-to-market distance. Since the triangle inequality implies the latter is weakly larger, this generates upward bias on the border effect estimate. Section 2.6 fully develops this point.

Figure 3: TRANSPORTATION OF WIND TURBINE BLADES

Notes: A convoy of wind turbine blades passing through the village of Edenfield, UK. Photo Credit: Anderson (2007)

turbine manufacturers, which would be in violation of European single market policy. Therefore, it is in the best interest of the wind farm owner to purchase the turbine that maximizes his or her profits independent of the nationality of the manufacturer.

The project manager's choice of manufacturer is our primary focus. In the period we study, purchasers of wind turbines were primarily independent producers, most often farmers or other small investors.[10] The turbine manufacturing industry, on the other hand, is dominated by a small number of manufacturing firms that manufacture, construct, and maintain turbines on the project owner's land. Manufacturers usually have a portfolio of turbines available with various generating capacities. Overall, their portfolios are relatively homogeneous in terms of observable characteristics.[11] There

---

[10] Small purchasers were encouraged by the financial incentive scheme that gave larger remuneration to small, independent producers such as cooperative investment groups, farmers, and private owners. The German Electricity Feed Law of 1991 explicitly ruled out price support for installations in which the Federal Republic of Germany, a federal state, a public electricity utility or one of its subsidiaries held shares of more than 25 percent. The Danish support scheme provided an about 30% higher financial compensation for independent producers of renewable electricity (Sijm (2002)). A new law passed in Germany in 2000 eliminated the restrictions for public electricity companies to benefit from above-market pricing of renewable energy.

[11] Main observable product characteristics are generation capacity, tower height, and rotor diameter. Distribution of turbines in terms of these variables is very similar in both countries. Further details are displayed in Table 23 in Appendix I.

could be, however, differences in quality and reliability that we do not directly observe.

The proximity of the production location to the project site is an important driver of cost differences across projects. Due to the size and weight of turbine components, oversized cargo shipments typically necessitate road closures along the delivery route (see Figure 3). Transportation costs range between 6 to 20 percent of total costs (Franken and Weber, 2008). In addition, manufacturers usually include maintenance contracts as part of the turbine sale, so they must regularly revisit turbine sites after construction.

2.2.1 Data

We have constructed a unique dataset from several sources which contains information on every wind farm developed in Denmark and Germany dating to the birth of the wind turbine industry. Feed-in tariff subsidies for wind energy generation were introduced 1984 in Denmark and 1991 in Germany. Prior to the introduction of feed-in tariffs, there was little activity in the turbine manufacturing industry. Up until the early 1990s, the industry was characterized by a large number of small, experimental start-up firms. A small number of these firms grew to be the firms we study in this paper, while the majority ceased operations.

The data include the location of each project, the number of turbines, the total megawatt capacity, the date of grid-connection, manufacturer identity, and other turbine characteristics, such as rotor diameter and tower heights. We match the project data with the location of each manufacturer's primary production facility, which enables the calculation of road-distances from each manufacturer to each project. This provides us with a spatial source of variation in manufacturer costs which aids in identifying the sources of market segmentation. A key missing variable in our data set is transaction price, which necessitates the use of our model to derive price predictions from first order conditions on profit maximization.[12] Appendix I provides a detailed description of the data.

In this paper, we concentrate on the period from 1995 to 1996.[13] This has several advantages. First, the set of firms was stable during this time period. There are several medium-to-large firms competing in the market. In 1997, a merger and acquisition wave began, which lasted until 2005. The merger wave, including cross-border mergers, would complicate our analysis of the border effect. Second, site owners in this period were typically independent producers. This contrasts with later periods when utility companies became significant purchasers of wind turbines, leading to more concerns about repeated interaction between purchasers and manufacturers. Third, this period contains several well-established firms and the national price subsidies for wind electricity generation had been in place for several years. Prior to the mid-1990s, the market could be considered an "infant industry" with substantial uncertainty about the viability of firms and downstream subsidies. Fourth, starting in the late 1990s, a substantial fraction of wind turbine installations are offshore, so road-distance to the turbine location is no longer a useful source of variation in production costs.[14]

In focusing on a two-year period, we abstract away from some dynamic considerations. Although this greatly simplifies the analysis, it comes with some drawbacks. Most important is that one cannot distinguish sunk costs from fixed costs of entering the foreign export market (Roberts and Tybout, 1997; Das, Roberts, and Tybout, 2007). Because of the small number of firms and the lack of substantial entry and exit, it would not be possible to reliably estimate sunk costs and fixed costs separately in any case. Instead, we model the decision to enter a foreign market as a one-shot game. This decision does not affect the consistency of our variable cost estimates, whereas our counterfactuals removing fixed costs should be interpreted as removing both sunk and fixed costs. We also abstract away from dynamic effects on production technologies, such as learning-by-doing (see Benkard, 2004).

---

[12] As in most business-to-business industries, transaction prices are confidential. Some firms do publish list prices, which we have collected from industry publications. These prices, however, do not correspond to relevant final prices due to site-specific delivery and installation costs.

[13] Appendix I.4 shows that the evidence on market shares and the border effect is stable in subsequent time periods.

[14] Moreover, the Danish onshore market saturates after the late 1990s, leaving us with little variation at that side of the border.

Learning-by-doing would provide firms with an incentive to lower prices below a static profit maximizing level in return for anticipated dynamic gains.[15] Learning-by-doing is less of a concern for the mid-1990s than for earlier years. By 1995, the industry has matured to the extent that it is reasonable to assume that firms were setting prices to maximize expected profits from the sale.

Table 10: MAJOR DANISH AND GERMAN MANUFACTURERS

Manufacturer	Nationality	% Market share in Denmark	% Market share in Germany
Vestas	(DK)	45.45	12.04
Micon	(DK)	19.19	8.17
Bonus	(DK)	12.12	5.05
Nordtank	(DK)	11.45	4.73
WindWorld	(DK)	4.38	2.73
Total		92.59	32.72
Enercon	(DE)		32.58
Tacke	(DE)		14.95
Nordex	(DE)	1.68	7.53
Suedwind	(DE)		2.37
Fuhrlaender	(DE)		2.15
Total		94.27	92.3

Notes: Market shares in terms of number of projects installed in 1995-1996. Shares are very similar when projects are weighted by megawatt size.

Table 10 displays the market shares or the largest five Danish and German firms in both countries. We take these firms to be the set of manufacturers in our study. All other firms had domestic market shares below 2 percent, no long-term presence in their respective markets, and did not export. In our model, we treat these small turbine producers as a competitive fringe. The German and Danish wind turbine markets were relatively independent from the rest of the world. There was only one firm exporting from outside Germany and Denmark: A Dutch firm, Lagerwey, which sold to 21 projects in Germany (2.26 percent market share) and had a short presence in the German market. We include Lagerwey as part of the competitive fringe. In Figure 4, we present the project locations

---

[15] In some cases, this could even lead firms to sell products below cost. See Besanko, Doraszelski, Kryukov, and Satterthwaite (2010) for a fully dynamic computational model of price-setting under learning-by-doing.

using separate markers for German and Danish produced projects. Figure 5 provides the location of the primary production facility for each turbine manufacturer.

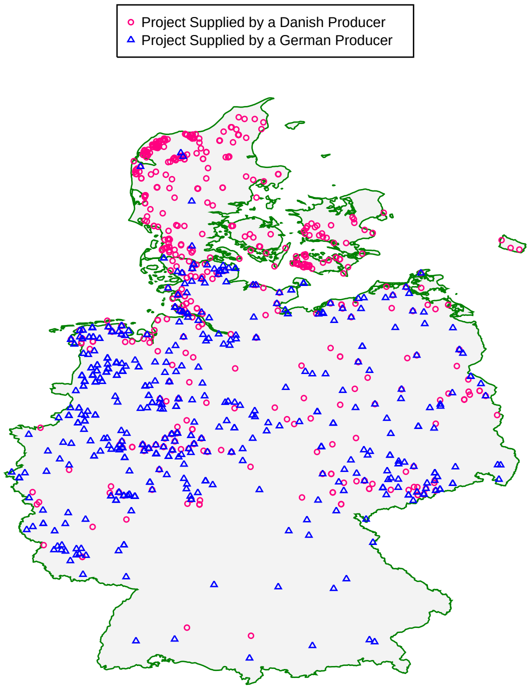

Figure 4: Project Locations

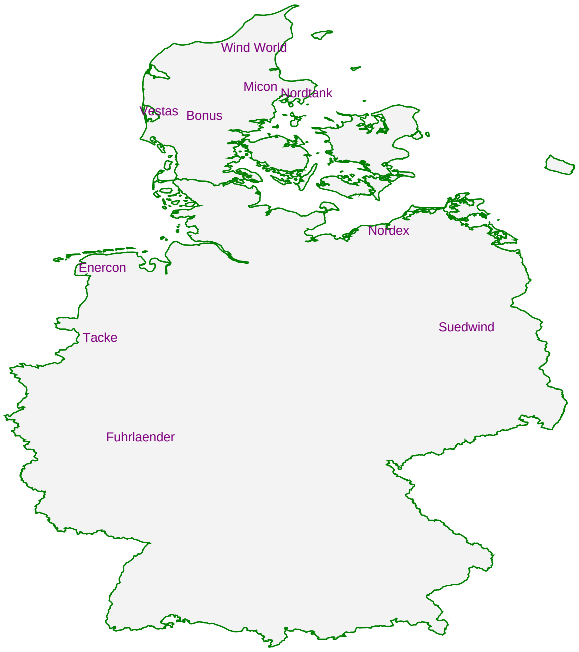

Figure 5: Producer Locations

2.2.2 Preliminary Analysis of the Border Effect

Table 10 and Figure 4 clearly suggest some degree of market segmentation between Germany and Denmark. Four out of five large German firms—including the German market leader, Enercon—do not have any presence in Denmark. That all Danish firms enter Germany whereas only one German firm competes in Denmark is consistent with the existence of fixed costs for exporting. Because the German market is much larger than the Danish market (929 projects were installed in Germany in this period, versus 296 in Denmark—see the map of projects in Figure 4), these fixed costs can be amortized over a larger number of projects in Germany.

For those firms that do export, the decline in market share by moving from foreign to domestic

Figure 6: Market Share of Vestas by Proximity to Primary Production Facility

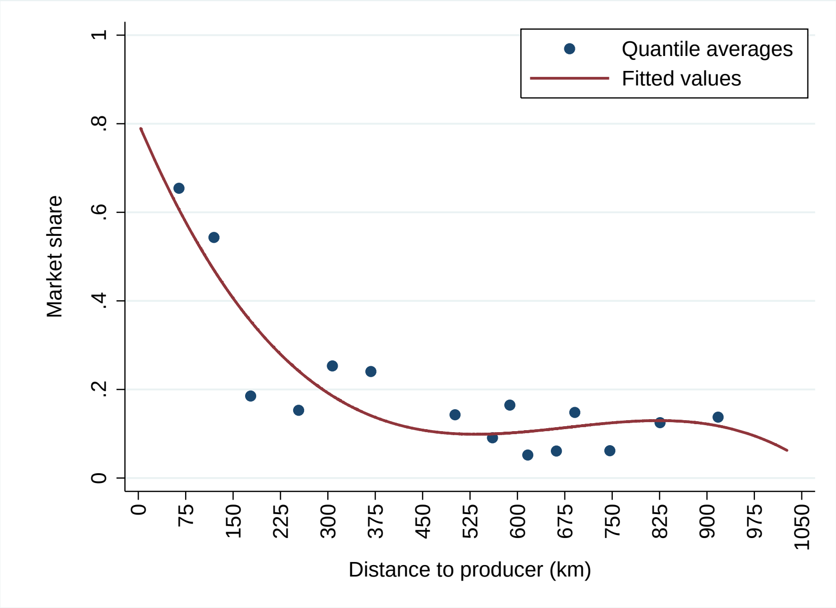

Notes: The fitted line follows from the linear regression of a dummy variable that takes the value one if a project is supplied by Vestas, and zero otherwise, on a cubic polynomial of projects' road distances to Vestas's headquarters. Dots are local market shares (i.e. proportion of projects supplied by Vestas) within 15 distance quantiles. We trim 12 out of 1225 projects at the top 1 distance percentile.

markets may have many different causes. First, market structure changes as the set of firms competing in Denmark is smaller than that in Germany. Second, due to transportation costs, foreign firms will have higher costs than domestic ones simply because projects are likely to be nearer to domestic manufacturing plants. Finally, there may be some variable border costs, which must be paid for each foreign project produced.

We start by exploring the effect of distance as a potential source of market segmentation. The impact of distance on firm costs is illustrated by Figure 6. This figure documents Vestas's declining market share as the distance from its main manufacturing location increases. Whereas Figure 6 suggests that costs increase with distance from the manufacturing base, it cannot easily be used to estimate distance costs. The impact of the border—roughly 150 kilometers from Vestas's manufacturing plant—confounds the relationship. Moreover, in an oligopolistic industry, Vestas's share is a function of not only its own costs but also those of competitor firms. Our model will jointly solve for the probability that each competing firm wins a project based on the project's location in

Figure 7: Market Share of Danish Firms by Proximity to the Border

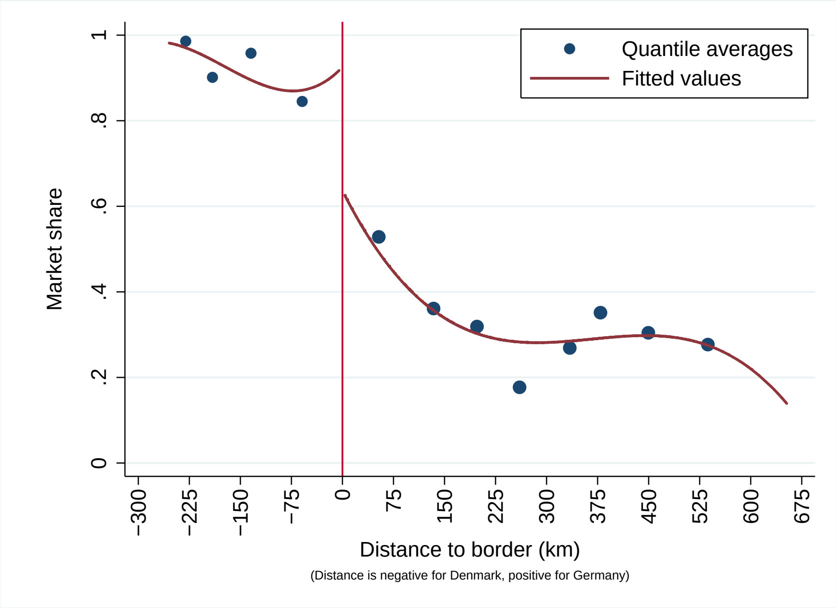

Notes: The fitted line follows from the linear regression of a dummy variable that takes the value one if a project is supplied by one of the five large Danish firms on a cubic polynomial of projects' great-circle distances to the border, a Germany dummy and interaction terms. Regression details are in Appendix I.4. Dots are local market shares (i.e. proportion of projects supplied by Danish firms) within 12 distance quantiles. We trim the sample at the bottom and top 1 distance percentiles which drops 24 out of 1225 projects.

relation to all firms. We are thus able to use the rich variation in projects across space to estimate the impact of distance on firm costs.

We next employ a regression discontinuity design (RDD) to quantify the effect of the border on large Danish firms' market share. Given that wind and demand conditions do not change abruptly, the RDD uncovers the impact of the border. To implement this, we regress a project-level binary variable that takes the value one if it is supplied by a large Danish firms and zero otherwise, to a cubic polynomial of distance from the project to the border, a Germany dummy (to capture the border effect), and interaction terms (see Appendix I.4 for details). Figure 7 plots the fit of this regression. The border dummy is a statistically significant -0.29, which is reflected in the sharp drop in the market share of the largest five Danish firms from around 90 to 60 percent at the border.[16]

These results give us reason to believe that the border matters in the wind turbine industry,

---

[16] These results are robust to considering projects within various bandwidths of the border, as is standard in the RDD framework. For expository purposes, Figure 7 includes projects in the [-300,675] band after trimming the data at the bottom and top 1 percentiles.

however it tells us little about how the discontinuity arises. For example, the discontinuity at the border does not separately identify the effect of changes in market structure between Germany and Denmark from the impact of variable border costs. Because variable border costs are incurred precisely at the point where market structure changes, we are unable to use the RDD approach to separate the two effects. This motivates our use of a structural model, which utilizes firm level data to separate these costs.[17] The following section proposes and estimates a model to account for the changes market structure at the border by treating the competition for projects as a Bertrand-Nash game.

2.3 Model

Denmark and Germany are indexed by $\ell \in \{D, G\}$. Each country has a discrete finite set of large domestic firms denoted by $\mathcal{M}_\ell$ and a local fringe. Large firms are heterogeneous in their location and productivity. There is a fixed number of N_ℓ projects in each country, and they are characterized by their location and size (total megawatt generation capacity). We assume the projects are exogenously located. The land suitable for building a wind turbine is mostly rural and diffuse, so it is unlikely that project location is affected by the presence or absence of a turbine manufacturer. Cross-border competition takes place in two stages: In the first stage, large firms decide whether or not to pay a fixed cost and enter the foreign market. In the second stage, firms bid for all projects in the markets they compete in (they do so in their domestic market by default). Project owners independently choose a turbine supplier among competing firms. We now present the two stages following backward induction, starting with the bidding game.

---

[17] While it is common to assume that border frictions arise due to costs associated with conducting business across national boundaries, they may have also other sources, such as home bias or even spatial collusion, as pointed out by Salvo (2010). We would expect home bias in preferences to be small in the setting of wind turbines since the consumers are profit maximizing businesses. We also do not believe spatial collusion to be a likely explanation for the discontinuity in our setting. Danish firms were active throughout Germany during this period, and our analysis in Appendix J does not reveal a strong degree of spatial clustering after controlling for borders and distance that we might expect if firms were cooperatively splitting the wind turbine market across space. Moreover, the industry receives a high degree of regulatory scrutiny due to its connection to electricity generation more broadly. No anti-trust cases have been filed with the European Commission against the firms studied in this paper.

2.3.1 Project Bidding Game

In this stage, active firms offer a separate price to each project manager, and project managers choose the offer that maximizes their valuation. The set of active firms is taken as given by all players, as it was realized in the entry stage. For ease of notation, we drop the country index ℓ for the moment and describe the project bidding game in one country. The set of active, large firms (denoted by $\mathcal{J}$) and the competitive fringe compete over N projects. $\mathcal{J}$ contains all domestic and foreign firms—if there are any—that entered the market in the first stage, so $\mathcal{M} \subseteq \mathcal{J}$.

The per-megawatt payoff function of a project owner i for choosing firm j is,

$$V_{ij} = d_j - p_{ij} + \epsilon_{ij}.$$

The return to the project owner depends on the quality of the wind turbine, d_j, the per-megawatt price p_{ij} charged by manufacturer j denominated in the units of the project owner's payoff,[18] and an idiosyncratic choice-specific shock ϵ_{ij}.[19] It is well known that discrete choice models only identify relative differences in valuations. We thus model a non-strategic fringe as an outside option. We denote it as firm 0 and normalize the return as

$$V_{i0} = \epsilon_{i0}.$$

We assume ϵ_{ij} is distributed i.i.d. across projects and firms according to the Type-I extreme value distribution.[20] The $\boldsymbol{\epsilon}_i$ vector is private information to project managers who collect project-specific

---

[18] Since we do not directly observe prices, we will use the manufacturer's first order condition to derive prices in units of the project owners payoff. As a result the "marginal utility of currency" coefficient on price is not identified and is simply normalized to 1. While this normalization does prevent us from presenting currency figures for consumer and producer surplus, it does not affect the ratio of consumer to producer surplus or the relative welfare implications of our counterfactual analyses.

[19] We assume away project-level economies of scale by making price bids per-megawatt. In Appendix I, we check whether foreign turbine manufacturers tend to specialize on larger projects abroad. We find that the average project size abroad is very similar to the average project size at home for each exporting firm.

[20] Project owners do not have any home bias in the sense that ϵ_{ij}'s are drawn from the same distribution for all producers in both countries.

price bids from producers. The assumption that $\boldsymbol{\epsilon}_i$ is i.i.d. and private knowledge abstracts away from the presence of unobservables that are known to the firms at the time they choose prices but are unknown to the econometrician.[21] After receiving all price bids, denoted by the vector $\mathbf{p}_i$, owners choose the firm that delivers them the highest payoff. Using the familiar logit formula, the probability that owner i chooses firm j is given by,

$$Pr[i \text{ chooses } j] \;\equiv\; \rho_{ij}(\mathbf{p}_i) \;=\; \frac{\exp(d_j - p_{ij})}{1 + \sum_{k=1}^{|\mathcal{J}|} \exp(d_k - p_{ik})} \qquad \text{for } j \in \mathcal{J}. \tag{30}$$

The probability of choosing the fringe is

$$Pr[i \text{ chooses the fringe}] \;\equiv\; \rho_{i0}(\mathbf{p}_i) \;=\; 1 - \sum_{j=1}^{|\mathcal{J}|} \rho_{ij}(\mathbf{p}_i).$$

Now we turn to the problem of the turbine producers. The per-megawatt cost for producer j to supply project i is a function of its heterogeneous production cost ϕ_j, its distance to the project, and whether or not it is a foreign producer:[22]

$$c_{ij} = \phi_j + \beta_d \cdot \text{distance}_{ij} + \beta_b \cdot \text{border}_{ij}, \tag{31}$$

where

$$\text{border}_{ij} = \begin{cases} 0 & \text{if both } i \text{ and } j \text{ are located in the same country,} \\ 1 & \text{otherwise.} \end{cases}$$

In other words, all firms pay the distance related cost ($\beta_d \cdot \text{distance}_{ij}$), but only foreign firms pay the variable border cost ($\beta_b \cdot \text{border}_{ij}$). The distance cost captures not only the cost of transportation

---

[21] For example, if local politics or geography favors one firm over another in a particular region, firms would account for this in their pricing strategies, but we are unable to account for this since this effect is unobserved to us. In Appendix J, we address the robustness of our estimate to local unobservables of this type.

[22] Here, we assume costs scale with a project's production capacity. Below, we also consider an alternative specification where costs are allowed to vary with project size.

but also serves as a proxy for the cost of post-sale services (such as maintenance), installing remote controllers to monitor wind farm operations, gathering information about sites further away from the manufacturer's location, and maintaining relationships with local contractors who construct turbine towers. The border component captures additional variable costs faced by foreign manufacturers. This may include the cost of dealing with project approval procedures in the foreign market and coordinating transportation of bulky components with various national and local agencies. Note that delivery costs in (31) are per-unit rather than ad-valorem. While ad-valorem specification (i.e. iceberg-type trade costs that vary with the price of a good) is a common assumption in trade literature, it is inappropriate in the context of wind-turbine logistics. The presence of additive per-unit trade costs have been documented by **?** and more recently by **?**.

Firms engage in Bertrand competition by submitting price bids for projects in the markets in which they are active.[23] They observe the identities and all characteristics of their competitors (i.e., their quality and marginal cost for each project) except the valuation vector $\boldsymbol{\epsilon}_i$. The second stage is thus a static game with imperfect, but symmetric, information. In a pure-strategy Bayesian-Nash equilibrium, each firm chooses its price to maximize expected profits given the prices of other firms:[24]

$$E[\pi_{ij}] = \max_{p_{ij}} \; \rho_{ij}(p_{ij}, \mathbf{p}_{i,-j}) \cdot (p_{ij} - c_{ij}) \cdot S_i,$$

where S_i is the size of the project in megawatts. The first order condition reads as follows:

$$\begin{aligned} 0 &= \frac{\partial \rho_{ij}(p_{ij}, \mathbf{p}_{i,-j})}{\partial p_{ij}}(p_{ij} - c_{ij}) + \rho_{ij}(p_{ij}, \mathbf{p}_{i,-j}), \\ p_{ij} &= c_{ij} - \frac{\rho_{ij}(p_{ij}, \mathbf{p}_{i,-j})}{\partial \rho_{ij}(p_{ij}, \mathbf{p}_{i,-j})/\partial p_{ij}}. \end{aligned}$$

---

[23] Industry experts we interviewed indicated that there was an excess supply of production capacity in the market during these years. One indication of this is that many firms suffered from low profitability, sparking a merger wave. Therefore, it is not likely that capacity constraints were binding in this period.

[24] We assume that firms are maximizing expected profits on a project-by-project level. This abstracts away from economics of density in project locations–i.e., the possibility that by having several projects close together they could be produced and maintained at a lower cost. We are address the robustness of our model to the presence of economies of density in Appendix J.

Exploiting the properties of the logit form, this expression simplifies to an optimal mark-up pricing condition:

$$p_{ij} = c_{ij} + \frac{1}{1 - \rho_{ij}(p_{ij}, \mathbf{p}_{i,-j})}. \tag{32}$$

The mark-up is increasing in the (endogenous) probability of winning the project and is thus a function of the set of the firms active in the market and their characteristics. Substituting (32) into (30), we get a fixed-point problem with $|\mathcal{J}|$ unknowns and $|\mathcal{J}|$ equations for each project i:

$$\rho_{ij} = \frac{\exp\left(d_j - c_{ij} - \frac{1}{1-\rho_{ij}}\right)}{1 + \sum_{k=1}^{|\mathcal{J}|} \exp\left(d_k - c_{ik} - \frac{1}{1-\rho_{ik}}\right)} \qquad \text{for } j \in \mathcal{J}. \tag{33}$$

Our framework fits into the class of games for which Caplin and Nalebuff (1991) show the existence of a unique pure-strategy equilibrium. Using the optimal mark-up pricing condition, the expected profits of manufacturer j for project i can be calculated as,

$$E[\pi_{ij}] = \frac{\rho_{ij}}{1 - \rho_{ij}} S_i.$$

Potential exporters take expected profits into account in their entry decisions. We turn to the entry game in the next section.

2.3.2 Entry Game

Before bidding on projects, an entry stage is played in which all large firms simultaneously decide whether or not to be active in the foreign market by incurring a firm-specific fixed cost f_j. This captures expenses that can be amortized across all foreign projects, such as establishing a foreign sales office, gaining regulatory approvals, or developing the operating software satisfying the requirements set by national grids.

Let $\Pi_j(\mathcal{J}_{-j} \cup j)$ be the expected profit of manufacturer j in the foreign market where $\mathcal{J}_{-j}$ is the set of active bidders other than j. This is simply the sum of the expected profit of bidding for all foreign projects:

$$\Pi_j(\mathcal{J}_{-j} \cup j) = \sum_{i=1}^{N} E[\pi_{ij}(\mathcal{J}_{-j} \cup j)]. \tag{34}$$

Manufacturer j enters the foreign market if its expected return is higher than its fixed cost:

$$\Pi_j(\mathcal{J}_{-j} \cup j) \geq f_j. \tag{35}$$

Note that this entry game may have multiple equilibria. Following the literature initiated by Bresnahan and Reiss (1991), we assume that the observed decisions of firms are the outcome of a pure-strategy equilibrium; therefore, if a firm in our data is active in the foreign market, (35) must hold for that firm. On the other hand, if firm j is not observed in the foreign market, one we can infer the following lower bound on fixed export cost:

$$\Pi_j(\mathcal{J}_{-j} \cup j) \leq f_j. \tag{36}$$

We use these two necessary conditions to construct inequalities that bound f_j from above or from below by using the estimates from the bidding game to impute the expected payoff estimates of every firm for any set of active participants in the foreign market.[25] We now turn to the estimation of the model.

2.4 Estimation

Estimation proceeds in two steps: In the first step, we estimate the structural parameters of the project-bidding game. In the second step, we use these estimates to solve for equilibria in the project-bidding game with counterfactual sets of active firms to construct the fixed costs bounds. Before

---

[25] Several papers (e.g., Pakes, Porter, Ho, and Ishii, 2006; Ciliberto and Tamer, 2007) proposed using bounds to construct moment inequalities for use in estimating structural parameters. **?** and Morales, Sheu, and Zahler (2011) applied this methodology to the context of spatial entry and trade. Because we observe only a single observation of each firm's entry decision, a moment inequality approach is not applicable in our setting.

proceeding with the estimation, we must define the set of active firms in every country. Under our model, the set of firms that have positive sales in a country is a consistent estimate of the active set of firms; therefore, we define a firm as active in the foreign market if it has any positive sales there.[26]

We now reintroduce the country index: ρ_{ij}^{ℓ} is firm j's probability of winning project i in country ℓ. The number of active firms in market ℓ is $|\mathcal{J}_\ell|$, and $\text{border}_{ij}^{\ell}$ equals zero if project i and firm j are both located in country ℓ and one otherwise. Substituting the cost function (31) into the winning probability (33), we find,

$$\rho_{ij}^{\ell} = \frac{\exp\left(d_j - \phi_j - \beta_d \cdot \text{distance}_{ij} - \beta_b \cdot \text{border}_{ij}^{\ell} - \frac{1}{1-\rho_{ij}^{\ell}}\right)}{1 + \sum_{k=1}^{|\mathcal{J}_\ell|} \exp\left(d_k - \phi_k - \beta_d \cdot \text{distance}_{ik} - \beta_b \cdot \text{border}_{ik}^{\ell} - \frac{1}{1-\rho_{ik}^{\ell}}\right)}. \tag{37}$$

From this equation, one can see that firms' production costs ϕ_j and quality level d_j are not separately identified given our data.[27] We thus jointly capture these two effects by firm fixed-effects $\xi_j = d_j - \phi_j$.

We collect the parameters to estimate into the vector $\theta = (\beta_b, \beta_d, \xi_1, \ldots, \xi_{|\mathcal{M}_D|+|\mathcal{M}_G|})$. We estimate the model via constrained maximum likelihood, where the likelihood of the data is maximized subject to the equilibrium constraints (37). The likelihood function of the project data has the following form:

$$L(\rho) = \prod_{\ell \in \{D,G\}} \prod_{i=1}^{N_\ell} \prod_{j=0}^{|\mathcal{J}_\ell|} (\rho_{ij}^{\ell})^{y_{ij}^{\ell}}, \tag{38}$$

where $y_{ij}^{\ell} = 1$ if manufacturer j is chosen to supply project i in country ℓ and 0 otherwise. The constrained maximum likelihood estimator, $\hat{\theta}$, together with the vector of expected project win prob-

---

[26] Note that every active firm has a positive probability of winning every project. As the number of projects goes to infinity, every active firm wins at least one project. We thus consider firms with zero sales in a market as not having entered in the first stage and exclude them from the set of active firms there.

[27] The difference between productivity and quality would be identified if we had data on transaction prices. Intuitively, for two manufacturers with similar market shares, high prices would be indicative of higher quality products while low prices would be indicative of lower costs.

abilities, $\hat{\rho}$, solves the following problem:

$$\begin{aligned}
\max_{\theta,\,\rho} \quad & L(\rho) \\
\text{subject to:} \quad & \rho^{\ell}_{ij} = \frac{\exp\left(\xi_j - \beta_d \cdot \text{distance}_{ij} - \beta_b \cdot \text{border}^{\ell}_{ij} - \frac{1}{1-\rho^{\ell}_{ij}}\right)}{1 + \sum_{k=1}^{|\mathcal{J}_\ell|} \exp\left(\xi_k - \beta_d \cdot \text{distance}_{ik} - \ \beta_b \cdot \text{border}^{\ell}_{ik} - \frac{1}{1-\rho^{\ell}_{ik}}\right)} \\
& \sum_{k=1}^{|\mathcal{J}_\ell|} \rho^{\ell}_{ik} + \rho^{\ell}_{i0} = 1
\end{aligned} \tag{39}$$

$$\text{for } \ell \in \{D, G\},\ i \in \{1, ..., N_\ell\},\ j \in \mathcal{J}.$$

Our estimation is an implementation of the Mathematical Programming with Equilibrium Constraints (MPEC) procedure proposed by **?**. They show that the estimator is equivalent to a nested fixed-point estimator in which the inner loop solves for the equilibrium of all markets, and the outer loop searches over parameters to maximize the likelihood. The estimator therefore inherits all the statistical properties of a fixed-point estimator. It is consistent and asymptotically normal as the number of projects tends to infinity. For the empirical implementation, we reformulate the system of constraints in (39) in order to simplify its Jacobian. In our baseline specification, this is a problem with 12,303 variables (12 structural parameters and 12,291 equilibrium win probabilities for all firms competing for each project) and 12,291 equality constraints. We describe the details of the computational procedure in Appendix K.

Once the structural parameters are recovered, one can calculate bounds on the fixed costs of entry for each firm, f_j, using (35) and (36). This involves resolving the model with the appropriate set of firms while holding the structural parameters fixed at their estimated values. We use a parametric bootstrap procedure to calculate the standard errors for these bounds.[28]

---

[28] To be specific, we repeatedly draw θ_b from the asymptotic distribution of $\hat{\theta}$ and recalculate the bound each time. Under the assumptions of the model, the distribution of bound statistic generated by this procedure is a consistent estimate of the true distribution.

We also implement several alternative specifications as robustness checks and extensions to our baseline specification. In our first alternative, we allow distance costs to vary by manufacturing firm:

$$c_{ij} = \phi_j + \beta_{dj} \cdot \text{distance}_{ij} + \beta_b \cdot \text{border}_{ij}. \tag{40}$$

Note that the difference between this and the baseline specification (31) is that distance cost coefficients are heterogeneous (β_{dj} vs. β_d). This cost function is consistent with **?**, who document that in U.S. data large firms tend to ship further away, even when done domestically.[29] If heterogeneous shipping costs were present in the wind turbine industry, they might bias our baseline estimate of the border effect upward through a misspecification of distance costs, since smaller firms would not export due to higher transport costs instead of the border effect.

In a second alternative specification, we allow the per-megawatt cost of a project and the impact of national boundaries to vary by project size,

$$c_{ij} = \phi_j + \beta_d \cdot \text{distance}_{ij} + \beta_b \cdot \text{border}_{ij} + \beta_s \cdot S_i + \beta_{sb} \cdot \text{border}_{ij} \cdot S_i. \tag{41}$$

The primary purpose of this specification is to investigate economies of scale in the variable border cost. If variable border cost is primarily generated by a single per-project cost that does not vary with size, then β_{sb} will be negative and the border will matter relatively less for large projects than for small, since the cost is amortized across a more electric capacity. On the other hand, if the variable border costs are proportional to project size, as they would if costs are connected to delivery or legal liability associated with the value of cross-border contracts, then β_{sb} will be small in magnitude and border costs will remain important even for large projects. The size coefficient, β_s, affects all active producers equally and is meant to control for the fact that the competitive fringe is made up of small

---

[29] They rationalize this observation in a model where heterogeneous firms invest in their distribution networks. Productive firms endogenously face a lower "iceberg transportation cost."

firms and is less likely to have the resources to serve large projects.

Finally, we estimate a specification allowing for costs to be a quadratic function of distance. We find that the hypothesis that the cost function is in fact linear in distance cannot be rejected. We do not find this result surprising, as the "first mile" costs associated with loading equipment are subsumed within the firm-cost intercepts, ϕ_j whereas once trucks are loaded, linear costs seem like a reasonable assumption. Consequently, we focus on linear distance specifications in the results below.

2.4.1 Parameter Estimates

Estimation results are presented in Table 11, with the baseline specification reported in the first column. Both variable costs are economically and statistically significant. The border cost is equivalent to 436 kilometers in distance costs. (The next subsection provides a detailed analysis of trade cost estimates.) The second column contains the estimates of the heterogeneous distance cost specification presented in (40). The border variable cost coefficient, $\hat{\beta}_b$, is practically unchanged and remains strongly significant, indicating that the estimated border costs are not an artifact of distance cost heterogeneity. Turning to the distance costs themselves, most firms, particularly the larger ones, have distance costs that are close to our homogeneous distance cost estimate. It does not appear that small firms have systematically higher distance costs. The smallest firm in our data, Suedwind, is estimated to be distance loving; this firm is based in Berlin, but has built several turbines in the west of Germany.[30] While a formal likelihood ratio test rejects the null hypothesis of homogeneous distance costs, the estimation results indicate that heterogeneous distance costs are not driving cross-border differences in this industry. Therefore, we use our homogeneous distance cost specification for the counterfactuals in the following section.

The third column of Table 11 contains estimates from the size-varying per-megawatt cost specification, (41). The coefficient of interest is the interaction term, β_{sb}, which is negative, but

---

[30] Nordex, who is also located in the east of Germany, also has a negative coefficient, but it is statistically insignificant.

neither economically nor statistically significant. (The average project size is 1 megawatt.) This is evidence that the variable border cost does in fact scale with project size, and is not simply a per-project "hassle cost" that might be amortized away when a project is large. The coefficient on project size, β_s, is significant and reflects that the fringe firm has a more difficult time winning large projects (independent of the border). This is likely due to reputation and other practical difficulties which prevent small fringe firms from competing for large projects. Overall, these results provide support for our baseline assumption that the variable border cost scales with project size.

As discussed above, the firm fixed effects reflect the combination of differences in quality and productivity across firms. We find significant differences between firms. It is not surprising that the largest firms, Vestas and Enercon, have the highest fixed effects. Although there is significant within-country dispersion, Danish firms generally appear to be stronger than German ones. The results suggest that controlling for firm heterogeneity is important for correctly estimating border and distance costs.

Since our model delivers expected purchase probabilities for each firm at each project site, we can use the regression discontinuity approach to visualize how well our model fits the observed data. Figure 8 presents this comparison using the baseline specification results. The horizontal axis is the distance to the Danish-German border, where negative distance is inside Denmark. The red (solid) line is the raw data fit. This is the same curve as that presented in Figure 7, relating distance-to border and a border dummy to the probability of a Danish firm winning a project. In particular, this regression does not control for project-to-firm distances. The blue (dotted) curve is fitted using the expected win probabilities calculated from the structural model. These probabilities depend explicitly on our estimates of both firm heterogeneity and project-to-manufacturer distances but do not explicitly depend on distance to the border (as this indirectly affects costs for firms in our model). Note that predicted win probabilities are nonlinear despite the linearity of costs. This is due to the nonlinear nature of the model as well as the rich spatial variation of mark-ups predicted by the model. The

Table 11: MAXIMUM LIKELIHOOD ESTIMATES

	Baseline	Heterogeneous Distance Costs	Project Size Scaling
Border Variable Cost, β_b	0.873	0.880	0.953
	(0.219)	(0.239)	(0.228)
Distance Cost (100km), β_d	0.200		0.193
	(0.032)		(0.032)
Bonus (DK)		0.166	
		(0.066)	
Nordtank (DK)		0.275	
		(0.073)	
Micon (DK)		0.132	
		(0.051)	
Vestas (DK)		0.280	
		(0.049)	
WindWorld (DK)		0.013	
		(0.068)	
Enercon (DE)		0.297	
		(0.063)	
Fuhrlaender (DE)		1.826	
		(0.238)	
Nordex (DE)		-0.073	
		(0.087)	
Suedwind (DE)		-0.232	
		(0.104)	
Tacke (DE)		0.105	
		(0.071)	
Project Size, β_s			-0.687
			(0.108)
Project Size x Border, β_{sb}			-0.063
			(0.053)
Firm Fixed Effects, ξ_j			
Bonus (DK)	2.471	2.327	1.953
	(0.223)	(0.297)	(0.231)
Nordtank (DK)	2.524	2.806	2.001
	(0.229)	(0.326)	(0.238)
Micon (DK)	3.095	2.782	2.572
	(0.221)	(0.268)	(0.230)
Vestas (DK)	3.797	4.152	3.279
	(0.215)	(0.266)	(0.224)
WindWorld (DK)	1.733	0.810	1.208
	(0.273)	(0.418)	(0.281)
Enercon (DE)	3.531	3.868	3.012
	(0.175)	(0.270)	(0.194)
Fuhrlaender (DE)	0.328	3.359	-0.177
	(0.263)	(0.512)	(0.269)
Nordex (DE)	1.779	0.680	1.261
	(0.203)	(0.401)	(0.216)
Suedwind (DE)	0.534	-1.194	0.020
	(0.270)	(0.509)	(0.279)
Tacke (DE)	2.386	2.109	1.875
	(0.177)	(0.263)	(0.197)
Log-Likelihood	-2361.34	-2314.22	-2352.79
N	1225	1225	1225

Notes: Standard errors in parentheses.

Figure 8: Model Fit: Expected Danish Market Share by Distance to the Border

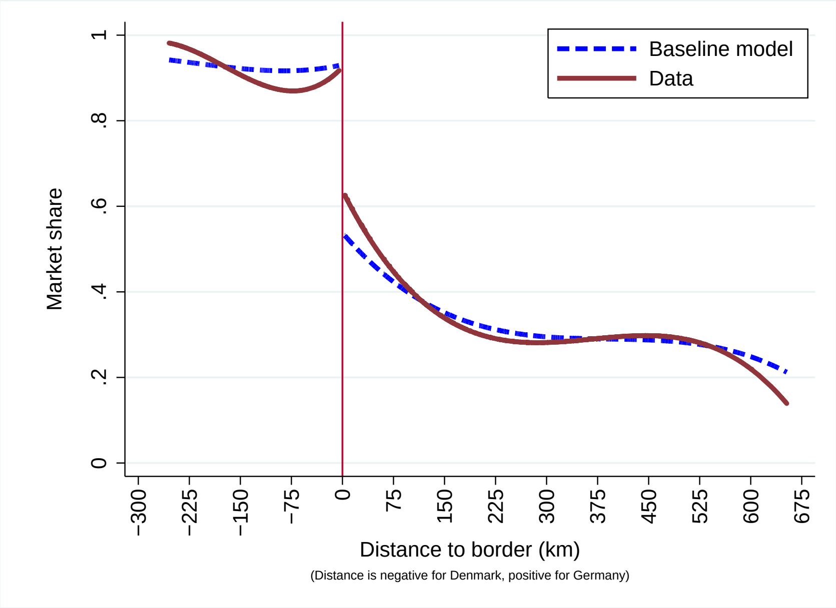

Notes: Data line is the same as in Figure 7. The model line is the linear fit of winning probability for each project by Danish firms on a cubic polynomial of projects' great-circle distances to the border, a Germany dummy, and interaction terms.

size of the discontinuity is somewhat larger using the structural model, although the qualitative result that the border effect is large is apparent using both methods. Overall, the model fits the data well. Next, we report various statistics on the magnitude and relative importance of estimated trade costs in order to highlight their economic content.

2.4.2 Analysis of Trade Cost Estimates

Our baseline estimates imply that the cost of supplying a foreign project is equivalent to an additional 436 kilometers of distance between the manufacturing location and the project site ($\beta_b/\beta_d = 0.436$). The mean distance from Danish firms to German projects in our data is 623 kilometers. As a consequence, border frictions represent 40 percent of Danish exporters' total delivery costs ($\beta_b/(\beta_b + \beta_d \cdot 6.23) = 40$ percent). Since German firms face a shorter shipping distance to Denmark (the mean distance is 420 kilometers), border frictions make up a larger share of total delivery costs for them ($\beta_b/(\beta_b + \beta_d \cdot 4.2) = 50$ percent).

To get a sense of the importance of distance-related costs on market outcomes, we calculate the distance elasticity of the equilibrium probability of winning a project for every project-firm combination.[31] For exporters, the median distance elasticity ranges from .95 to 1.38. That is, the median effect of a one percent increase in the distance from an exporting firm to a project abroad (holding all other firms' distances constant) is a decline of .95 to 1.38 percent in the probability of winning the project. For domestic firms, the median distance elasticities are lower, ranging from .17 to .83. The difference is due to both the smaller distances firms must typically travel to reach domestic projects and the impact of the border on equilibrium outcomes. It appears that distance costs have a significant impact on firm costs and market shares for both foreign and domestic firms.

Finally, we compute exporters' delivery costs relative to domestic delivery costs in order to gauge the combined effect of the border and geography. The average cost of supplying German projects for Danish firms is $\mathbf{C}_{DG} = \beta_b + \beta_d \cdot 6.23$. The same cost is $\mathbf{C}_{GG} = \beta_d \cdot 3.63$ for German firms who do not incur the border cost and are located 363 km away from their domestic projects on average. So, Danish firms face variable delivery costs that are on average 191% higher than their German competitors' costs in accessing the German market ($\mathbf{C}_{DG}/\mathbf{C}_{GG} - 1 = 1.91$).

2.4.3 Fixed Cost Bounds

Not all firms enter the foreign market; rather, firms optimally choose whether or not to export by weighing their fixed costs of entry against the expected profits from exporting. Hence, firm-level heterogeneity in operating profits, fixed costs, or both is necessary to rationalize the fact that different firms make different exporting decisions.[32] Since our model naturally allows for heterogeneity in firm operating profits, this section considers whether heterogeneity in firms' fixed costs of exporting are

---

[31] The distance elasticities we report are a function of the characteristics of all firms at a particular project site in a very specific industry. It is difficult to directly compare these distance elasticities with distance elasticities of aggregated trade volumes frequently reported in the trade literature that rely on national or regional capital distance proxies (e.g., McCallum (1995); **?**); Anderson and van Wincoop (2003))

[32] The canonical Melitz (2003) model assumes homogenous fixed costs and heterogeneity in operating profits. Eaton, Kortum, and Kramarz (2011) show that heterogeneity in fixed costs is also necessary to fit the export patterns in French firm-level data.

also needed to rationalize observed entry decisions.

Since we only observe a single export decision for each firm, fixed costs are not point identified. Nevertheless, the model helps to place a bound on them. Firms optimally make their export decision based on the level of fixed costs of foreign entry and on the operating profits they expect in the export market as described in Section 2.3.2. Based on the parameter estimates in Table 11, we can derive counterfactual estimates of expected operating profits for any set of active firms in the Danish and German markets. Therefore, we can construct an upper bound on fixed costs for firms entering the foreign market using (35), and a lower bound on fixed costs for firms that stay out of the foreign market using (36). While the scale of these bounds is normalized by the variance of the extreme-value error term, comparing them across firms gives us some idea of the degree of heterogeneity in fixed costs.[33]

Table 12 presents the estimates of fixed cost bounds for each firm. The intersection of the bounds across all firms is empty. For example, there is no single level of fixed costs that would simultaneously justify WindWorld entering Germany and Enercon not entering Denmark; hence, some heterogeneity in fixed costs is necessary to explain firm entry decisions.

One possibility is that fixed cost for entering Germany differ from those for entering Denmark. Since all Danish firms enter the Danish market, any fixed cost below 17.35 (the expected profits of WindWorld for entering Germany) would rationalize the observed entry pattern. In Germany however, the lower and upper bound of Enercon and Nordex have no intersection. Some background information about Nordex supports the implication of the model that Nordex may be subject to much lower costs than Enercon to enter into the Danish market. Nordex was launched as a Danish company in 1985 but shifted its center of business and production activity to Germany in the early 1990s. As a consequence, Nordex could keep a foothold in the Danish market at a lower cost than could the other German firms,

---

[33] As mentioned above, since transaction prices are not observed we cannot express profits in currency units (see Footnote 19).

Table 12: Export Fixed Cost Bounds (f_j)

	Lower	Upper
Bonus (DK)		47.53
		(19.52)
Nordtank (DK)		43.28
		(8.92)
Micon (DK)		80.12
		(13.64)
Vestas (DK)		163.36
		(23.47)
WindWorld (DK)		17.35
		(3.94)
Enercon (DE)	22.35	
	(4.89)	
Fuhrlaender (DE)	0.66	
	(0.32)	
Nordex (DE)		6.32
		(1.82)
Suedwind (DE)	1.26	
	(0.45)	
Tacke (DE)	7.25	
	(1.72)	

Notes: Scale is normalized by the variance of ϵ (see Footnote 19). Standard errors in parentheses.

which would need to form contacts with Danish customers from scratch.[34]

Of course, the Nordex anecdote also highlights some important caveats with regard to our bounds. By assuming a one-shot entry game, we are abstracting away from entry dynamics. If exporting is less costly to continue than to initiate, then the bounds we calculate—which consider only profits from operating in 1995 and 1996—will be biased downward. Data limitations, particularly the small number of firms, prevent us from extending the model to account for dynamic exporting decisions along the lines of Das, Roberts, and Tybout (2007). Nevertheless, our results illustrate the

---

[34] Because of Nordex's connection to Denmark, we perform a robustness check on our variable border cost estimate by re-estimating the model allowing Nordex to sell in Denmark without having to pay the border variable cost. The border cost estimate increases in this specification, but the difference is not statistically significant. Since Nordex is the only exporting German firm, this robustness check also serves as a check on our specification of symmetric border costs. See Balistreri and Hillberry (2007) for a discussion of asymmetric border frictions.

degree of heterogeneity in fixed costs that is necessary to explain entry patterns.[35]

2.5 Border Frictions, Market Segmentation, and Welfare

We now use the model to study the impact of border frictions on national market shares, firm profits, and consumer welfare. We perform a two-step counterfactual analysis. The first step eliminates fixed costs of exporting, keeping in place variable costs incurred at the border.[36] Even though we are unable to point identify firms' fixed costs of exporting, this counterfactual allows us to examine the implications of fixed border costs by setting them to zero, which implies that all firms enter the export market. The second step further removes the variable cost of the border by setting β_b equal to zero.[37] This eliminates all border frictions such that the only sources of differing market shares across national boundaries are plant-to-project distances and firm heterogeneity. While the results of this experiment constitute an estimate of what can be achieved if border frictions could be entirely eliminated, it is important to keep in mind that natural barriers, such as different languages, will be difficult to eliminate in practice.

2.5.1 Market Shares and Segmentation

We begin our analysis by considering how national market shares in each country react to the elimination of border frictions. Furthermore, because market shares are directly observed in the data, the baseline model's market share estimates can also be used to assess the fit of our model to national level aggregates. Table 13 presents the market share of the major firms of Denmark and Germany in

---

[35] It is important to note that the variable cost estimates presented in Table 11, as well as the counterfactual results below, are robust to dynamic entry as long as firm pricing decisions have no impact on future entry decisions. This assumption is quite common in the literature on structural oligopoly models, e.g., Ericson and Pakes (1995).

[36] We implicitly assume that the change in market structure does not induce domestic firms to exit the industry, or new firms to be created.

[37] We eliminate first fixed border costs and then variable costs because changes in variable border costs when fixed costs are still positive could induce changes in the set of firms that enter foreign markets. Because they are not point identified, we are unable to estimate fixed border costs. Even with reliable estimates, the entry stage with positive fixed costs is likely to result in multiple equilibria.

Table 13: Counterfactual Market Shares of Large Firms (%)

		Data	Baseline Estimates	No Fixed Costs	No Border
Denmark	Danish Firms	92.57	92.65	83.93	74.19
			(1.53)	(2.26)	(3.65)
	German Firms	1.69	2.17	11.57	22.02
			(0.60)	(2.06)	(3.90)
Germany	Danish Firms	32.29	32.36	32.36	49.34
			(5.43)	(5.43)	(7.56)
	German Firms	59.63	59.31	59.31	44.90
			(3.93)	(3.93)	(5.81)

Notes: Market share measured in projects won. Standard errors in parentheses.

each country, with the fringe taking the remainder of the market. Comparing the first two columns, the baseline predictions of the model closely correspond to the observed market shares. All of the market shares are within the 95 percent confidence interval of the baseline predictions, which suggests that the model has a good fit.

In the third column, we re-solve the model eliminating fixed costs of exporting and keeping the variable border cost in place. In response, the four German firms that previously competed only domestically start exporting to Denmark. As a result, the market share of German firms in Denmark rises almost 10 percentage points.[38] Danish firms, however, still maintain a substantial market share advantage in their home market. Since all five large Danish firms already compete in Germany, there is no change in market shares on the German side of the border when fixed costs of exporting are removed.[39] The difference in response to the elimination of fixed costs between the Danish and German markets is obvious, but instructive. The reduction or elimination of border frictions can have very different effects based on market characteristics. In our case, because there are more projects in Germany than in Denmark, the return to entry is much larger in Germany. This may be one reason

---

[38] For space and clarity, we do not report standard errors of changes in market shares in Table 13. All of the (non-zero) changes in market shares across counterfactuals are statistically significant.

[39] Because of this duplication, we simply omit the column which removes fixed cost of entry in Germany in the tables below.

why we see more Danish firms entering Germany than vice versa.[40] As a result, reducing fixed costs of exporting to Germany has no effect on market outcomes, whereas the impact of eliminating fixed cost of exporting to Denmark is substantial.

The fourth and final column of Table 13 displays the model prediction of national market shares if the border were entirely eliminated. In addition to setting f_j equal to zero for all firms, we also eliminate variable border costs by setting β_b equal to zero.[41] This results in a large increase in imports on both sides of the border. The domestic market share of Danish firms falls from 92.7 percent to 74.2 percent. The domestic market share of large Danish firms remains high due to firm heterogeneity and the fact that they are closer to Danish projects. In Germany, roughly half of the projects import Danish turbines once the border is eliminated, which reflects the strength of Danish firms (especially Vestas) in the wind turbine industry. On both sides of the border, we see an approximate 20 percent increase in import share when the national boundary is eliminated.

Overall, our results indicate that border frictions generate significant market segmentation between Denmark and Germany. As a back of the envelope illustration, consider the difference between the market share of Danish firms in the two markets. The gap in the data and baseline model is roughly 60 percentage points. Not all of this gap can be attributed to border frictions since differences in transportation costs due to geography are also responsible for part of the gap. However, when we remove border frictions, our counterfactual analysis indicates that the gap shrinks to 25 percentage points. More than half of the market share gap is thus attributable to border frictions. When considering the sources of border frictions, we find that removing fixed costs of exporting alone accounts for one-third of the market share gap that is attributed to border frictions, while the remaining two-thirds are realized by removing both fixed and variable border frictions. Since fixed and variable costs interact, the overall impact of border frictions cannot be formally decomposed into fixed and variable

---

[40] This argument assumes fixed costs of exporting are of the same order of magnitude for both countries, which appears to be the case.

[41] Because adjustments to variable costs may result in a change in firms optimal entry decisions, we are unable to perform a counterfactual eliminating variable border costs alone.

cost components. We take these results as evidence that both fixed and variable border frictions are substantial sources of market segmentation.

In addition to national market share averages, our model allows us to examine predicted market shares at a particular point in space. Using the RDD approach describe above, Figure 9 visualizes the impact of the counterfactual experiments. The blue (dashed) line represents expected market shares baseline model, and is identical to that presented in Figure 8. The red (dotted) line displays counterfactual expected market shares when fixed border costs are removed. This reduces domestic market share of Danish firms since more German firms enter, but leaves market shares unchanged in Germany since all firms were already competing there. Finally, the green (dashed-dotted) line shows the counterfactual estimates when all border costs are eliminated. The discontinuity at the border is entirely eliminated, and only the impact of firm-to-manufacturer distances cause differences in market share on either side of the border.[42]

2.5.2 Firm Profits

We now turn to an analysis of winners and losers from border frictions, starting with individual firms. Table 14 presents the level of operating profits under the baseline and two counterfactual scenarios, calculated according to equation (34). While the scale of these profit figures is arbitrary (similar to f_j in Table 12, units are normalized by the variance of ϵ), they allow for comparison both across firms and across scenarios. The table separates profits accrued in Germany and Denmark for each firm. For example, in the baseline scenario, we see that Bonus made 47.14 in profits in Denmark, and 47.53 in Germany. If the border were removed entirely, Bonus's profits in Denmark would fall to 34.85, while their profits in Germany would rise to 75.61. On overall, Bonus would see its total profits increase as a result of the elimination of border frictions, as gains in Germany would more than offset loses from

---

[42] The kink at the boundary is an artifact of interaction terms in the RDD which implies that we estimate either side of the border as a separate cubic polynomial in distance to the border. The bottom line is that there is no discontinuity at the boundary when all border effects are removed.

Figure 9: COUNTERFACTUALS: EXPECTED DANISH MARKET SHARE BY DISTANCE TO THE BORDER

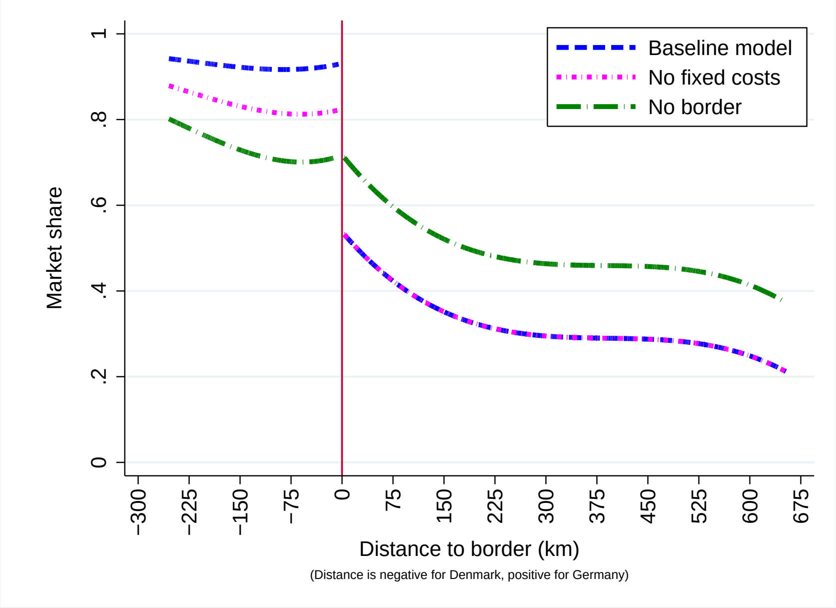

Notes: Regression discontinuity fit of projects won by large Danish firms under the baseline model and counterfactual scenarios. See notes to Figure 8 for details.

increased competition in Denmark.

The situation is different for German firms. When fixed costs are eliminated, the large German firms—Enercon and Tacke—take the lion's share of the gains. However, all German firms—even the largest firm, Enercon—loose from the entire elimination of border frictions. Underlying this result is the significant asymmetry in size and productivity between Germany and Denmark. The losses German firms face due to increased competition in the larger German market overwhelm all gains they receive from frictionless access to the Danish market. Our model estimates Danish firms to be highly productive, so eliminating the border is quite costly to German incumbents. In addition, variable border frictions are estimated to be so high that even a small Danish exporter like WindWorld becomes much more competitive in Germany when they are removed. Despite being a relatively small player, WindWorld gains from the elimination of border frictions since increased profits in the larger Germany market outweigh its losses at home. However, WindWorld's gains are insignificant when compared to the gains of the large Danish firms, such as Vestas. Overall, we find that because a

Table 14: BASELINE AND COUNTERFACTUAL PROFIT ESTIMATES

	Denmark			Germany	
	Baseline Estimates	No Fixed Costs	No Border	Baseline Estimates	No Border
Bonus (DK)	47.14	41.08	34.85	47.53	75.61
	(13.02)	(12.00)	(10.72)	(19.52)	(28.93)
Nordtank (DK)	44.77	39.04	33.12	43.28	68.87
	(4.98)	(4.51)	(4.19)	(8.92)	(13.77)
Micon (DK)	82.89	72.74	62.10	80.12	127.03
	(7.37)	(6.81)	(6.76)	(13.64)	(21.20)
Vestas (DK)	156.41	138.95	120.12	163.36	255.20
	(14.56)	(12.43)	(11.81)	(23.47)	(37.08)
WindWorld (DK)	20.76	18.15	15.44	17.35	27.65
	(3.19)	(2.76)	(2.49)	(3.94)	(6.34)
Enercon (DE)		21.48	42.74	429.44	304.95
		(4.55)	(9.41)	(48.71)	(53.63)
Fuhrlaender (DE)		0.57	1.14	17.34	11.98
		(0.26)	(0.57)	(4.21)	(3.28)
Nordex (DE)	6.32	5.42	10.81	75.75	51.19
	(1.82)	(1.48)	(2.45)	(15.16)	(13.19)
Suedwind (DE)		1.09	2.17	21.76	14.84
		(0.37)	(0.79)	(5.24)	(3.90)
Tacke (DE)		6.48	12.99	152.01	104.75
		(1.43)	(3.21)	(16.61)	(17.34)

Notes: Scale is normalized by variance of ϵ (see Footnote 19). Standard errors in parentheses.

German firm's domestic market is considerably larger than its export market, border frictions protect the profit of German firms over those of Danish firms.

2.5.3 Consumer Surplus and Welfare

We now analyze the overall impact of the border on welfare in the Danish and German wind turbine markets. For each country, Table 15 presents consumer surplus (i.e., surplus accruing to site owners) and firm profits (aggregated by producer's country) under the baseline and our two counterfactual scenarios.[43] The relative changes in consumer surplus across scenarios are invariant to the scale of

---

[43] Consumer surplus in country ℓ is equal to the sum of expected utility of all project owners:

$$CS^\ell = \sum_{i=1}^{N_\ell} S_i \log \sum_{j=1}^{|\mathcal{J}_\ell|} \exp\left(\xi_j - \beta_d \cdot \text{distance}_{ij} - \beta_b \cdot \text{border}_{ij}^\ell - \frac{1}{1-\rho_{ij}^\ell}\right).$$

Table 15: Counterfactual Welfare Analysis by Country

		Baseline (Levels)	No Fixed Costs (Levels)	No Fixed Costs (% Chg)	No Border (Levels)	No Border (% Chg)
Denmark	(A) Consumer Surplus	70.15	73.47	4.73	77.45	10.41
		(4.95)	(4.98)	(1.03)	(5.39)	(2.20)
	(B) Danish Firm Profits	29.32	25.82	-11.94	22.13	-24.53
		(0.53)	(0.74)	(2.27)	(1.26)	(4.48)
	(C) German Firm Profits	0.53	2.92	454.48	5.82	1005.55
		(0.15)	(0.55)	(123.38)	(1.14)	(299.03)
	Domestic Surplus (A+B)	99.47	99.29	-0.18	99.58	0.11
		(5.18)	(5.12)	(0.07)	(5.10)	(0.25)
	Total Surplus (A+B+C)	100.00	102.21	2.21	105.40	5.40
		(5.10)	(5.08)	(0.51)	(5.40)	(1.29)
Germany	(A) Consumer Surplus	68.98			79.62	15.42
		(6.43)			(8.32)	(1.90)
	(B) Danish Firm Profits	10.41			16.41	57.66
		(1.59)			(2.41)	(4.98)
	(C) German Firm Profits	20.61			14.44	-29.96
		(1.86)			(2.31)	(5.63)
	Domestic Surplus (A+B)	89.59			94.05	4.98
		(5.79)			(6.70)	(1.39)
	Total Surplus (A+B+C)	100.00			110.46	10.46
		(6.73)			(8.61)	(1.78)

Notes: Levels are scaled such that baseline total surplus from projects within a country is 100. "% Chg" is percent change from baseline level. Standard errors in parentheses.

ϵ, so we normalize the consumer surplus in the baseline scenario to 100 for expositional ease.[44] We define domestic surplus as the total surplus in the country that accrues to consumers and domestic firms.

The first column reports the breakdown of surplus under the baseline scenario, we see that in both Denmark and Germany, consumers receive roughly 70 percent of the total surplus. In Denmark, the bulk of the remaining 30 percent goes to Danish firms (recall that only one German firm is active in Denmark), while in Germany, approximately 10 percent goes to Danish firms and 20 percent to German firms.

---

[44] Because of its larger size, the total surplus in Germany is much larger than in Denmark, cross country comparisons of total surplus are available by request.

The next two columns present results from the counterfactual where only fixed costs of entry are removed. As discussed above, this counterfactual only affects the Danish market outcomes, since all Danish firms already sell in Germany in the baseline scenario. We report both surplus levels, and the percentage change from the baseline level. Note that, because of the correlation in the level estimates due to the uncertainty in firm fixed effects, the percent change estimates are much more precise than a naïve comparison of the level estimates would suggest. Removing fixed costs of exporting causes four German firms to enter the Danish market, which both increases price competition and provides additional variety to Danish site owners. As a result, consumer surplus increases by 4.73 percent. Danish firms, facing harsher domestic competition, see profits decline by 11.94 percent. Since the number of German firms increased from one to five, total German profits skyrocket in percentage terms, however this is due to a very small initial base. Even after removing fixed costs, German firms take less than three percent of the available surplus in Denmark in profits. The gains of Danish consumers from removing fixed export costs are almost perfectly offset by the loses from Danish firms. Domestic surplus actually declines by a statistically significant but economically negligible amount. When we account for the gains by German firms, total surplus increases by the statistically and economically significant 2.21 percent.

The final two columns of Table 15 display the welfare effects of removing border frictions entirely. As we would expect, site owners see significant benefits, and consumer surplus rises by 10.41 percent in Denmark and 15.42 percent in Germany. The increase in Denmark is more than twice as high as the increase realized from only removing fixed border costs. These increases come at the cost of domestic producers, who see home profits decline by 24.53 percent in Denmark and 29.96 percent in Germany.[45] In Denmark, the removal of border frictions results in a transfer of surplus from domestic firms to consumers, netting to essentially no change in domestic surplus. When we include

---

[45] Of course, these declines do not account for benefits realized in the export market. See Table 14 for an accounting of how each firm fairs as both an domestic producer and an exporter under our counterfactual scenarios.

the benefits of exporters, however, total surplus increases by 5.40 percent. The story in Germany is a bit different. Consumer gains outweigh domestic firm losses in Germany and domestic surplus increases by 4.98 percent. Essentially, removing border frictions improves German site owners access to high-productivity Danish firms and erodes Enercon's substantial market power in Germany. When we include the benefits to Danish exporters, elimination of the border raises surplus in the German market by a substantial 10.46 percent.

We conclude this section by repeating an important disclaimer. Our second counterfactual represents an elimination of all border frictions. In reality, these frictions are generated by a complex combination of political, administrative, and cultural differences between countries. It is unlikely that any policy initiative would succeed in eliminating these differences completely. Rather, our findings illustrate the magnitude of the border and its effect on firms and consumers in the wind turbine industry. Policy makers may view the results as an upper bound on what can be accomplished through economic and political integration.

2.6 Alternative Border Estimates

A large literature estimates the "width" of border using a reduced form model of price dispersion between markets. In a typical exercise, the researcher starts with price data p_k^j for identical, tradable goods indexed by j measured in locations indexed by k, and estimates the following relationship,[46]

$$|p_k^j - p_\ell^j| \;=\; \delta_d^j \cdot \text{distance}_{k\ell} + \delta_b^j \cdot \text{border}_{k\ell} + \delta_k^j + \delta_\ell^j + \epsilon_{k\ell}^j, \tag{42}$$

where $\text{distance}_{k\ell}$ is the distance between locations k and ℓ, and $\text{border}_{k\ell} = 1$ if k and ℓ are in the different countries. Location fixed effects $(\delta_k^j, \delta_\ell^j)$ are often included to control for market-specific

---

[46] Authors have estimated several variations of this specification. For example, one could use log distance instead of distance, use price covariances over time instead of absolute price differences (Engel and Rogers, 1996), or not include market fixed effects (Broda and Weinstein, 2008).

differences that impact price levels.

This equation is motivated by a no-arbitrage condition between markets, reflecting the possibility that consumers or middlemen can always travel to distant markets to purchase the good.[47] The border effect for good j is then interpreted in terms of the distance equivalent of the border dummy coefficient:

$$\textit{Border Effect}_j \ = \ \frac{\delta_b^j}{\delta_d^j}.$$

In contrast, we have explicitly modeled border costs within an oligopolistic framework, without appealing to a no-arbitrage condition across markets. Instead, the model we estimate reflects the costs that suppliers must incur to transport a good from its production location to the market, and prices in each location are endogenously chosen by firms. We believe our approach is a more realistic way to model many differentiated product markets where firms play a significant role in choosing prices. Moreover, it is likely to lead to a smaller "width" of the border in comparison with (42) for reasons we discuss below. In order to demonstrate how these approaches differ, this section uses our model-generated prices to estimate (42) and compares the implied border widths.

In our thought experiment, we parameterize our model using the estimates from Section 4.2 and calculate price bid differentials on the left hand side of (42). An econometrician with data on price bids of all manufacturers indexed by j tries to recover the width using (42) which regresses these differentials on the distances between locations ($\text{distance}_{k\ell}$)—but *not* the distances between locations and producers (distance_{kj})—a border dummy denoting whether or not the two markets are on opposite sides of a national boundary, and dummy variables representing market fixed effects. This is a close description of the information set used by researchers who estimate reduced-form regressions depicted above. These estimates are then used to construct the width of the border.

Table 16 reports the implied border width ($\hat{\delta}_b^j/\hat{\delta}_d^j$) from the OLS regresion of (42) for each

---

[47] As pointed out by Borraz, Cavallo, Rigobon, and Zipitria (2012), the majority of the literature implements the model as an equality although the no-arbitrage condition implies only an inequality.

Table 16: Border Effect Estimates from Reduced-Form Estimation

Firm	$\hat{\delta}_b/\hat{\delta}_d$ in km
Bonus	844.2
Nordtank	886.5
Micon	816.6
Vestas	709.1
WindWorld	888.9
Nordex	3253.8

manufacturer. Recall that our baseline model implies border width of 432 kilometers for all firms.[48] Evidently, this exercise overestimates the border effect in our model for all producers.[49] For Danish firms, estimates vary between 1.6 to 2 times the true value. The bias is much higher for the German firm, Nordex.

To gain intuition on the sources of this overestimation, contrast (42) with the price difference implied by our structural model using our estimates $(\hat{\beta}_d, \hat{\beta}_b, \hat{\rho}_{kj})$ in expressions (31) and (32):

$$|p_k^j - p_\ell^j| \; = \; \left|\hat{\beta}_d\big(\text{distance}_{kj} - \text{distance}_{\ell j}\big) + \hat{\beta}_b\big(\text{border}_{kj} - \text{border}_{\ell j}\big) + \Big(\frac{1}{1-\hat{\rho}_{kj}} - \frac{1}{1-\hat{\rho}_{\ell j}}\Big)\right| \qquad (43)$$

The three terms in this equation reflect the sources of producer-level spatial price differentials in our model: differences in project-to-producer distances are captured by the first term, differences in border frictions for each project are captured by the second term, and differences in project-specific mark-ups due to variation in competitive structure across space are captured by the last term. Note that the firm competitiveness parameter has canceled out through taking differences.

When we compare this data generating process with equation (42), it is apparent that the linear

---

[48] As discussed above, we calculate the model-based width as $\hat{\beta}_b/\hat{\beta}_d$, the width is based on the variable border cost because the fixed border cost is sunk when firms set prices.

[49] The overestimation is robust to whether or not we include location fixed effects, which are included in the reported results. In the underlying regressions, distance and border coefficients are statistically significant at .01 level for all producers. The detailed regression results are available from authors upon request.

reduced-form regression is misspecified. In the structural equation (43), price differentials are generated by the absolute value of several *differences* in project-to-producer distances, destination countries, and mark-ups, whereas (42) is a linear function of related, but different, variables. While trying to emulate this model-based expression, equation (42) suffers from two additional problems: First, using project-to-project distances ($\text{distance}_{k\ell}$) instead of the differences in project-to-producer distances differences ($\text{distance}_{kj} - \text{distance}_{\ell j}$) leads to (non-classical) measurement error. The triangle inequality implies that the actual difference of the project-to-producer distances is less than the project-to-project distances. This would tend to bias the estimate of δ_d towards zero relative to the distance parameter $\hat{\beta}_d$ in (43).[50] Second, (42) suffers from omitted variable bias due to the mark-up differentials being left out. Note that the vector of location fixed effects included in the regression cannot properly characterize the mark-up differences between project-pairs, since those dummies would capture information about levels instead of differences. Since markup differences are likely to be correlated with the border dummy, this would tend to bias δ_b upwards due to an endogeneity problem. The combined result is the border effect estimates $\hat{\delta}_b/\hat{\delta}_d$ in Table 16 are higher than their structural analogue, $\hat{\beta}_b/\hat{\beta}_d$.

While our thought experiment focuses on price deviations directly, it is easy to see that the linear specification error, measurement error and omitted variable bias would arise when using covariance measures are the dependent variable as well. The findings of this section resonate with Anderson and van Wincoop (2004) and Gorodnichenko and Tesar (2009) who argue that model-free border-effect regressions fail to identify border frictions based on price differentials. Moreover, we show the importance of using disaggregated data—in our case the knowledge of manufacturing locations—to properly control for variation in distance costs and markups. These issues apply to a large range of industries in which specific producers operate in only a few locations and the set of active firms differs

---

[50] The triangle inequality discrepancy explains why the measured effect is so much higher for Nordex in Table 16. Danish firms are all located at the north end of the set of projects. Hence, project-to-project distance is a better proxy for the distance differential, since the majority of projects are south of their manufacturing facility. Nordex, however, is more centrally located. As a result, two separate projects in Denmark and Germany that are equidistant to Nordex, and thus have a low firm-to-project distance differential will have a high project-to-project distance. Nordex's distinctively higher border effect estimate is thus in part due to a poorer distance proxy for many project pairs.

across national boundaries.

2.7 Conclusion

The large differences in national market shares in the wind turbine industry between Denmark and Germany arise not only through costs associated with distance, but also through barriers to foreign market entry and higher variable costs associated with crossing the border. This paper uses transaction-level data to document the impact of fixed and variable border costs while controlling for several sources of bias that plague analysis of aggregated trade flows. The model and the detailed geographical information on manufacturers and projects allow us to better control for distance costs and spatial differences in competition on either side of the border than the existing literature. In addition, the model enables us to conduct counterfactual analysis on the impact of border frictions on producer and consumer welfare. We find that border frictions are substantial; counterfactual analysis indicates that more than 50 percent of gap in cross-border market shares can be attributed to border costs. Our results also indicate that the welfare gains from a hypothetical removal of all border frictions between Germany and Denmark are large. These gains include removing barriers that are difficult if not impossible to remove, such as language, but represent an upper bound on what can be achieved through further integration.

Our study takes some strides towards identifying the underlying sources of border frictions. We separately document the role of a fixed cost to begin exporting and a variable border cost for each exported shipment. We also show that variable border costs in the wind turbine industry tend to scale with the size of the project (measured in megawatts) rather than being a simple per-project "hassle-cost" of contracting across borders. Of course, there is still more work to be done. We cannot, for example, separately identify the roles that bureaucratic, linguistic, or cultural differences play in generating border frictions. With data trade data from several countries, our model could easily be extended to investigate whether cultural or legal similarities appear to reduce the costs of crossing national boundaries. Moreover, while it is reasonable to attribute border frictions to costs in our setting of large capital goods trade in a business-to-business industry, in other industries cross-border differences in preferences—in particular home bias—may play a strong role. This is particularly true in consumer goods markets. Fortunately, in such markets prices transaction are much more readily available. With price data in hand, it would be possible to extend out model to untangle the difference between home-bias in preferences and border costs.

Nonetheless, the existence of large border frictions within the European wind turbine industry has important policy implications for the EU. Due to growing concerns about climate change, many governments, including EU members and the United States, subsidize renewable energy generation. The efficiency of subsidies in the wind electricity output market is closely related to the degree of competition in the upstream market for wind turbines themselves. If there are substantial frictions to international trade in turbines, a national subsidy to the downstream market may implicitly be a subsidy to domestic turbine manufacturers. This would be against the intensions of EU common market policy, which seeks to prevent distortions due to subsidies given by member states exclusively to domestic firms. In fact, Denmark, which has one of the most generous wind energy subsidies in Europe, is also home to the most successful European producers of wind turbines. Given our findings of large border frictions in the upstream market, EU members may wish to harmonize renewable energy tariffs to ensure equal treatment of European firms in accordance with the principles of the European single market project.

Bibliography

Alfaro, L., and M. X. Chen (2012): “Selection, Reallocation, and Spillover: Identifying the Sources of Gains from Multinational Production,” Working Paper 18207, National Bureau of Economic Research.

Anderson, J. E., and E. van Wincoop (2003): “Gravity with Gravitas: A Solution to the Border Puzzle,” *The American Economic Review*, 93(1), pp. 170–192.

——— (2004): “Trade Costs,” *Journal of Economic Literature*, 42(3), 691–751.

Anderson, P. (2007): “Scout Moor Wind Farm Under Construction,” `http://www.geograph.org.uk/gallery/scout_moor_wind_farm_under_construction_6568`.

Arkolakis, C., A. Costinot, and A. Rodriguez-Clare (2012): “New Trade Models, Same Old Gains?,” *American Economic Review*, 102(1), 94–130.

Arkolakis, C., and M.-A. Muendler (2010): “The Extensive Margin of Exporting Products: A Firm-level Analysis,” NBER Working Papers 16641, National Bureau of Economic Research, Inc.

Arkolakis, C., N. Ramondo, A. Rodriguez-Clare, and S. Yeaple (2012): “Innovation and Production in the Global Economy,” Pennsylvania State University.

Axtell, R. L. (2001): “Zipf Distribution of U.S. Firm Sizes,” *Science*, 293(5536), 1818–1820.

Balistreri, E. J., and R. H. Hillberry (2007): “Structural estimation and the border puzzle,” *Journal of International Economics*, 72(2), 451–463.

Baltagi, B. H., P. Egger, and M. Pfaffermayr (2008): “Estimating regional trade agreement effects on FDI in an interdependent world,” *Journal of Econometrics*, 145(1-2), 194–208.

Barro, R., and J.-W. Lee (2010): “A New Data Set of Educational Attainment in the World, 1950-2010.,” *NBER Working Paper*, (15902).

Benkard, C. L. (2004): “A Dynamic Analysis of the Market for Wide-Bodied Commericial Aircraft,” *Review of Economic Studies*, 71, 581–611.

Bernard, A. B., J. B. Jensen, S. J. Redding, and P. K. Schott (2007): “Firms in International Trade,” *Journal of Economic Perspectives*, 21(3), 105–130.

Bernard, A. B., J. B. Jensen, and P. K. Schott (2009): “Importers, Exporters and Multinationals: A Portrait of Firms in the U.S. that Trade Goods,” in *Producer Dynamics: New Evidence from Micro Data*, NBER Chapters, pp. 513–552. National Bureau of Economic Research, Inc.

Bernard, A. B., and C. I. Jones (1996): “Comparing Apples to Oranges: Productivity Convergence and Measurement across Industries and Countries,” *American Economic Review*, 86(5), 1216–38.

Bernard, A. B., S. J. Redding, and P. K. Schott (2010): "Multiple-Product Firms and Product Switching," *American Economic Review*, 100(1), 70–97.

Besanko, D., U. Doraszelski, Y. Kryukov, and M. Satterthwaite (2010): "Learning-by-Doing, Organizational Forgetting, and Industry Dynamics," *Econometrica*, 78(2), 453–508.

Blonigen, B. A., R. B. Davies, G. R. Waddell, and H. T. Naughton (2007): "FDI in space: Spatial autoregressive relationships in foreign direct investment," *European Economic Review*, 51(5), 1303–1325.

Borraz, F., A. Cavallo, R. Rigobon, and L. Zipitria (2012): "Distance and Political Boundaries: Estimating Border Effects under Inequality Constraints," NBER Working Paper 18122.

Bresnahan, T. F., and P. Reiss (1991): "Empirical Models of Discrete Games," *Journal of Econometrics*, 48(1), 57–81.

Broda, C., and D. E. Weinstein (2008): "Understanding International Price Differences Using Barcode Data," NBER Working Papers 14017, National Bureau of Economic Research, Inc.

Buch, C. M., J. Kleinert, A. Lipponer, and F. Toubal (2005): "Determinants and effects of foreign direct investment: evidence from German firm-level data," *Economic Policy*, 20(41), 52–110.

Caplin, A., and B. Nalebuff (1991): "Aggregation and Imperfect Competition: On the Existence of Equilibrium," *Econometrica*, 59(1), 25–59.

Caves, R. (1996): *Multinational Enterprise and Economic Analysis*. Cambridge University Press, Cambridge, 2nd edn.

Chaney, T. (2008): "Distorted Gravity: The Intensive and Extensive Margins of International Trade," *American Economic Review*, 98(4), 1707–21.

Ciliberto, F., and E. Tamer (2007): "Market Structure and Multiple Equilibria in Airline Markets," University of Virginia and Northwestern University.

Cummings, J., J. Manyika, L. Mendonca, E. Greenberg, S. Aronowitz, R. Chopra, K. Elkin, S. Ramaswamy, J. Soni, and A. Watson (2010): *Growth and competitiveness in the United States: The role of its multinational companies*. McKinsey Global Institute.

Das, S., M. J. Roberts, and J. R. Tybout (2007): "Market Entry Costs, Producer Heterogeneity and Export Dynamics," *Econometrica*, 75(3), 837–873.

Eaton, J., and S. Kortum (2002): "Technology, Geography, and Trade," *Econometrica*, 70(5), pp. 1741–1779.

Eaton, J., S. Kortum, and F. Kramarz (2011): "An Anatomy of International Trade: Evidence From French Firms," *Econometrica*, 79(5), 1453–1498.

Ekholm, K., R. Forslid, and J. R. Markusen (2007): "Export-Platform Foreign Direct Investment," *Journal of the European Economic Association*, 5(4), 776–795.

Engel, C., and J. H. Rogers (1996): "How Wide Is the Border?," *The American Economic Review*, 86(5), pp. 1112–1125.

Ericson, R., and A. Pakes (1995): "Markov-Perfect Industry Dynamics: A Framework for Empirical Work," *Review of Economic Studies*, 62(1), 53–82.

FEENSTRA, R. C., AND J. A. LEVINSOHN (1995): "Estimating Markups and Market Conduct with Multidimensional Product Attributes," *Review of Economic Studies*, 62(1), 19–52.

FEINBERG, S. E., AND M. P. KEANE (2006): "Accounting for the Growth of MNC-Based Trade Using a Structural Model of U.S. MNCs," *American Economic Review*, 96(5), 1515–1558.

FOWLIE, M., M. REGUANT, AND S. P. RYAN (2011): "Market-based Emissions Regulation and the Evolution of Market Structure," MIT Mimeo.

FRANKEN, M., AND T. WEBER (2008): "Heavy Duty," *New Energy*, 2008(5), 28–37.

GALE, D., AND H. NIKAIDO (1965): "The Jacobian Matrix and Global Univalence of Mappings," *Mathematische Annalen*, 159, 81–93.

GOLDBERG, P. K., AND F. VERBOVEN (2001): "The Evolution of Price Dispersion in the European Car Market," *Review of Economic Studies*, 68(4), 811–48.

GOLDBERG, P. K., AND F. VERBOVEN (2005): "Market integration and convergence to the Law of One Price: evidence from the European car market," *Journal of International Economics*, 65(1), 49 – 73.

GOPINATH, G., P.-O. GOURINCHAS, C.-T. HSEIH, AND N. LI (2010): "International Prices, Costs and Markup Differences," .

GORODNICHENKO, Y., AND L. L. TESAR (2009): "Border Effect or Country Effect? Seattle May Not Be So Far from Vancouver After All," *American Economic Journal: Macroeconomics*, 1(1), 219–41.

GROSSMAN, G. M., E. HELPMAN, AND A. SZEIDL (2006): "Optimal integration strategies for the multinational firm," *Journal of International Economics*, 70(1), 216–238.

HEAD, K., AND T. MAYER (2004): "Market Potential and the Location of Japanese Investment in the European Union," *The Review of Economics and Statistics*, 86(4), 959–972.

HEAD, K., AND J. RIES (2004): "Exporting and FDI as Alternative Strategies," *Oxford Review of Economic Policy*, 20(3), 409–423.

HELPMAN, E., M. J. MELITZ, AND S. R. YEAPLE (2004): "Export Versus FDI with Heterogeneous Firms," *American Economic Review*, 94(1), 300–316.

HILLBERRY, R., AND D. HUMMELS (2008): "Trade responses to geographic frictions: A decomposition using micro-data," *European Economic Review*, 52(3), 527 – 550.

HILLBERRY, R. H. (2002): "Aggregation bias, compositional change, and the border effect," *Canadian Journal of Economics*, 35(3), 517–530.

HOLMES, T. J. (2011): "The Diffusion of Wal-Mart and Economies of Density," *Econometrica*, 79(1), 253–302.

HORSTMANN, I. J., AND J. R. MARKUSEN (1992): "Endogenous market structures in international trade (natura facit saltum)," *Journal of International Economics*, 32(1-2), 109–129.

INTERESSENVERBAND WINDKRAFT BINNENLAND (various years): "Windkraftanlagen Marktübersicht," 1994 through 1996.

Irarrazabal, A., A. Moxnes, and L. D. Opromolla (2009): "The Margins of Multinational Production and the Role of Intra-firm trade," CEPR Discussion Papers 7145, C.E.P.R. Discussion Papers.

Klose, A., and A. Drexl (2005): "Facility location models for distribution system design," *European Journal of Operational Research*, 162(1), 4 – 29.

Kotz, S., and S. Nadarajah (2000): *Extreme Value Distributions: Theory and Applications*. Imperial College Press.

Levchenko, A. A., and J. Zhang (2011): "The Evolution of Comparative Advantage: Measurement and Welfare Implications," NBER Working Papers 16806, National Bureau of Economic Research, Inc.

McCallum, J. (1995): "National Borders Matter: Canada-U.S. Regional Trade Patterns," *The American Economic Review*, 85(3), 615–623.

Melitz, M. J. (2003): "The Impact of Trade on Intra-Industry Reallocations and Aggregate Industry Productivity," *Econometrica*, 71(6), 1695–1725.

Morales, E., G. Sheu, and A. Zahler (2011): "Gravity and Extended Gravity: Estimating a Structural Model of Export Entry," Harvard University Mimeo.

Muendler, M.-A., and S. O. Becker (2010): "Margins of Multinational Labor Substitution," *American Economic Review*, 100(5).

Neary, J. P. (2002): "Foreign Direct Investment and the Single Market," *Manchester School*, 70(3), 291–314.

Obstfeld, M., and K. Rogoff (2001): "The Six Major Puzzles in International Macroeconomics: Is There a Common Cause?," in *NBER Macroeconomics Annual 2000, Volume 15*, NBER Chapters, pp. 339–412. National Bureau of Economic Research, Inc.

Pakes, A., J. Porter, K. Ho, and J. Ishii (2006): "Moment Inequalities and their Application," Harvard University and University of Wisconsin.

Parsley, D. C., and S.-J. Wei (2001): "Explaining the border effect: the role of exchange rate variablility, shipping costs, and geography," *Journal of International Economics*, 55, 87–105.

Pinkse, J., and M. E. Slade (1998): "Contracting in space: An application of spatial statistics to discrete-choice models," *Journal of Econometrics*, 85(1), 125 – 154.

Poole, J. (forthcoming): "Knowledge transfers from multinational to domestic firms: evidence from worker mobility," *The Review of Economics and Statistics*.

Ramondo, N., and A. Rodriguez-Clare (2012): "Trade, Multinational Production, and the Gains from Openness," Working paper, Arizona State University and University of California at Berkeley.

Ramondo, N., A. Rodriguez-Clare, and F. Tintelnot (in process): "Data Set and Facts on Bilateral Multinational Activity," Pennsylvania State University, Mimeo.

Revelt, D., and K. Train (2000): "Customer-Specific Taste Parameters and Mixed Logit: Households' Choice of Electricity Supplier," Economics Working Papers E00-274, University of California at Berkeley.

ROBERTS, M. J., AND J. R. TYBOUT (1997): "The Decision to Export in Colombia: An Empirical Model of Entry with Sunk Costs," *American Economic Review*, 87(4), 545–64.

RODRIGUE, J. (2010): "Foreign Direct Investment, Exports and Aggregate Productivity," Vanderbilt University.

ROMALIS, J. (2007): "Capital Taxes, Trade Costs, and The Irish Miracle," *Journal of the European Economic Association*, 5(2-3), 459–469.

SALVO, A. (2010): "Trade flows in a spatial oligopoly: gravity fits well, but what does it explain?," *Canadian Journal of Economics*, 43(1), 63–96.

SIJM, J. (2002): "The Performance of Feed-in Tariffs to Promote Renewable Electricity in European Countries," ECN-C–02-083.

SU, C.-L., AND K. L. JUDD (forthcoming): "Constrained Optimization Approaches to Estimation of Structural Models," *Econometrica*.

THOMADSEN, R. (2005): "The Effect of Ownership Structure on Prices in Geographically Differentiated Industries," *The RAND Journal of Economics*, 36(4), pp. 908–929.

TRAIN, K. (2009): *Discrete Choice Methods with Simulation*. Cambridge University Press, 2nd edn.

YEAPLE, S. R. (2003): "The complex integration strategies of multinationals and cross country dependencies in the structure of foreign direct investment," *Journal of International Economics*, 60(2), 293–314.

Appendix A

Propositions

Proposition 1. *The firm-level sales to each market increase as additional production locations are added to the set of existing locations. However, there is a cannibalization effect across production locations. That is, a firm that adds a production location decreases the sales from the other locations.*

Proof. Let $Z^1 \supset Z^2$. The proposition consists of two parts. Part (i) states, $s_m(i, \phi, Z^1, \epsilon) > s_m(i, \phi, Z^2, \epsilon) \quad \forall m$. Proof: Substituting equation (7) into (8) yields

$$s_m(i, \phi, Z, \epsilon) = \kappa \phi^{\sigma-1} \frac{Y_m}{P_m^{1-\sigma}} \left(\sum_{k \in Z} (\gamma_{ik} w_k \tau_{km})^{-\theta} \epsilon_k^{\theta} \right)^{(\sigma-1)/\theta},$$

which increases as additional locations are added to Z since $\sigma > 1$. Part (ii) states, $s_{lm}(i, \phi, Z^1, \epsilon) < s_{lm}(i, \phi, Z^2, \epsilon)$ if $l \in Z^2 \quad \forall m$. Proof: Substituting equations (9), (8) and (7) into (10) yields

$$s_{lm}(i, \phi, Z, \epsilon) = \begin{cases} \kappa \phi^{\sigma-1} \frac{Y_m}{P_m^{1-\sigma}} \frac{(\gamma_{il} w_l \tau_{lm})^{-\theta} \epsilon_l^{\theta}}{\left(\sum\limits_{k \in Z} (\gamma_{ik} w_k \tau_{km})^{-\theta} \epsilon_k^{\theta} \right)^{\left(\frac{\theta+1-\sigma}{\theta} \right)}} & \text{if } l \in Z \\ 0 & \text{otherwise.} \end{cases}$$

The denominator increases as additional locations are added to Z since $\theta > \sigma - 1$. ■

Proposition 2. *Let $r : {}^K_{++} \times \mathcal{Z}^i \times \Psi \rightarrow {}^K_{++}$ be the stacked vector of revenues as defined in equation (23), where K denotes the number of countries in which firm t has a plant and $\Psi = [\psi_{min}, \psi_{max}]^K$ with $\psi_{min} > 0$ and $\psi_{min} < \psi_{max} < \infty$. Then for any triple $\{r_t, \tilde{w}, Z\}$, the vector ψ that solves $r_t - r(\tilde{w}, Z, \psi) = 0$ is unique.*

Proof. The proof shows that the conditions for the univalence theorem of Gale and Nikaido (1965) are satisfied. Clearly $r(\tilde{w}, Z, \psi)$ is differentiable with respect to ψ and Ψ is a closed rectangular region. It is left to show that Jacobian matrix of the mapping r is a P-Matrix at all $\psi \in \Psi$.

I simplify the expression in equation (23) in the following way. I drop the constants and define $\alpha = \frac{\theta+1-\sigma}{\theta}$, and $\tilde{y}_m = \frac{Y_m}{P_m^{1-\sigma}}$. Given the assumptions made in the text, $0 < \alpha < 1$. I denote $c_{lm} = (\tilde{w}_l \tau_{lm})^{-\theta}$. Further, I drop the firm index t. Then $r_{t,l}$ becomes

$$r_l(c, Z, \psi) = \sum_m \tilde{y}_m \frac{c_{lm} \psi_l^{\theta}}{\left(\sum\limits_{k \in Z} c_{km} \psi_k^{\theta} \right)^{\alpha}}$$

Note that

$$\begin{aligned}
\frac{\partial r_l}{\partial \psi_l} &= \sum_m \tilde{y}_m \frac{-\alpha\theta c_{lm} c_{lm} \psi_l^{2\theta-1} + \theta c_{lm} \psi_l^{\theta-1}\left(\sum\limits_{k\in Z} c_{km}\psi_k^\theta\right)}{\left(\sum\limits_{k\in Z} c_{km}\psi_k^\theta\right)^{\alpha+1}} \\
&= \sum_m \tilde{y}_m \frac{(1-\alpha)\theta c_{lm} c_{lm} \psi_l^{2\theta-1} + \theta c_{lm} \psi_l^{\theta-1}\left(\sum\limits_{k\neq l, k\in Z} c_{km}\psi_k^\theta\right)}{\left(\sum\limits_{k\in Z} c_{km}\psi_k^\theta\right)^{\alpha+1}} \\
&> 0
\end{aligned}$$

and for $k \neq l$

$$\frac{\partial r_l}{\partial \psi_k} = \sum_m \tilde{y}_m \frac{-\alpha\theta c_{km}\psi_k^{\theta-1} c_{lm}\psi_l^\theta}{\left(\sum\limits_{k\in Z} c_{km}\psi_k^\theta\right)^{\alpha+1}} < 0$$

It is easy to see that at all $\psi \in \Psi$, $|\frac{\partial r_l}{\partial \psi_l}|\psi_l > \sum\limits_{k\neq l} |\frac{\partial r_l}{\partial \psi_k}|\psi_k \; \forall l$, hence the Jacobian matrix of r has a dominant diagonal in the sense of Gale and Nikaido (1965). This along with the fact that $\frac{\partial r_l}{\partial \psi_l} > 0$ $\forall l$ implies that the Jacobian matrix of r is a P-Matrix.

Then the univalence theorem of Gale and Nikaido (1965) implies, whenever $r(\tilde{w}, Z, a) = r(\tilde{w}, Z, b)$, where $a, b \in \Omega$, then $a = b$. ■

Appendix B

Data

B.1 German multinationals data

All parent companies and majority-owned affiliates are from the manufacturing sector. Table 17 provides descriptive statistics on the multinational firms' activities in the countries included in my dataset. Table 18 documents the average foreign production share by the number of production locations.

Table 17: Descriptive Statistics German multinationals activities

Country	Number	Mean output	Median output
Austria	91	76.3	34
Belgium	45	235.3	37
Canada	36	536.0	28.5
Switzerland	70	58.3	17
Germany	665	625.8	98
Spain	117	191.9	32
France	191	107.7	30
United Kingdom	121	119.4	23
Ireland	18	36.3	19.5
Italy	100	65.0	27.5
Netherlands	46	83.1	25
United States	211	569.0	26

Notes: Output in million Euro. Source: MiDi database.

Table 18: Foreign production shares by number of production locations

Number of production locations	Number of firms	Mean share of foreign production	Mean share of foreign gross production
2	474	0.26	0.37
		(0.20)	(0.24)
3	102	0.32	0.54
		(0.18)	(0.19)
4	40	0.35	0.65
		(0.19)	(0.13)
5	23	0.39	0.71
		(0.16)	(0.10)
6	14	0.46	0.75
		(0.15)	(0.08)
≥7	12	0.48	0.80
		(0.06)	(0.07)
all	665	0.29	0.44
		(0.20)	(0.25)

Notes: Statistics for German MNE activities in 12 Western European and North American countries. Standard deviations in parantheses. Source: MiDi database.

Table 19: Foreign production shares by sectors

Sector	Number of firms	Mean Number of production locations	Mean share of foreign production	Mean share of foreign host countries production
Manufacture of ...				
textiles	15	2.27	0.34	0.39
		(0.80)	(0.22)	(0.25)
Publishing, printing, and reproduction of recorded media	22	2.36	0.26	0.37
		(0.66)	(0.25)	(0.23)
chemicals and chemical products	85	3.05	0.33	0.45
		(1.79)	(0.22)	(0.26)
rubber and plastic products	67	2.73	0.32	0.45
		(1.21)	(0.21)	(0.25)
other non-metallic mineral products	23	2.65	0.39	0.34
		(1.19)	(0.24)	(0.21)
basic metals	31	2.35	0.22	0.40
		(0.66)	(0.14)	(0.24)
metal products	72	2.32	0.27	0.43
		(0.78)	(0.17)	(0.23)
machinery and equipment	138	2.49	0.25	0.46
		(1.16)	(0.17)	(0.26)
electrical machinery and apparatus	34	2.79	0.26	0.48
		(1.65)	(0.17)	(0.26)
radio, television, and communication equipment and apparatus	15	2.33	0.24	0.51
		(0.72)	(0.16)	(0.28)
medical, precision, and optical instruments, watches, and clocks	49	2.33	0.30	0.54
		(0.75)	(0.20)	(0.24)
motor vehicles, trailers and semi-trailers	57	2.82	0.30	0.48
		(1.28)	(0.21)	(0.25)
all	665	2.57	0.29	0.44
		(1.20)	(0.20)	(0.25)

Notes: Statistics for German MNE activities in 12 Western European and North American countries. Standard deviations in parantheses. Statistics are displayed for sectors with more than 10 German multinationals. Source: MiDi database.

B.2 Aggregate data

This appendix describes the construction of the trade and MP shares. All data comes from Ramondo, Rodriguez-Clare, and Tintelnot (in process). The trade data is for the manufacturing sector only, while the MP data covers the entire non-financial sector of the economy. I implicitly assume that the MP in the service sector is proportional to the MP in the trade sector. The same assumption is made by Arkolakis, Ramondo, Rodriguez-Clare, and Yeaple (2012). All data are averages over the years 1996-2001.

B.2.1 Trade shares

$$Absorption_m = GrossProduction_m + TotalWorldImports_m - TotalWorldlExports_m - TotOtherImports_m$$

where $TotOtherImports_m$ represents the total imports by country m from countries not included in the analysis.

$$TradeShare_{lm} = \frac{X_{lm}}{Absorption_m}$$

B.2.2 MP shares

Let Y_{il} denote the value of output produced in country l by firms originating from country i. The construction of the MP shares takes into account that the set of countries included in this study is only a subset of the entire global economy (though an important part of it, with good local coverage, e.g. Western Europe and North America). Further, the total production of firms at home is not directly observed. I therefore take data on gross non-financial production in the respective country, and subtract the MP from 50 other source countries contained in Ramondo, Rodriguez-Clare, and Tintelnot (in process), which has the same sectoral coverage. This gives me an estimate of the value of local production, Y_{ii}:

$$MPShare_{jl} = \frac{Y_{jl}}{\sum\limits_{i \in C} Y_{il}}$$

where C denotes the set of countries included in the analysis.

Appendix C

Calculation of individual level parameters

The estimation in Section 3 delivers a distribution of fixed costs faced by the observed multinational firms. With these estimates I derive the distribution of fixed costs for each multinational firm conditional on its observed location choice, Z_t, and the location-specific productivity vector, ψ_t. We can then calculate the mean value of fixed costs that were actually paid to set up a plant in the respective countries. To my knowledge, Revelt and Train (2000) were the first to use such a procedure to infer information about the tastes of each sampled customer from the estimates of the distribution of tastes in the population with a nonlinear - mixed logit - discrete choice model.

Let β denote the parameter vector of estimates in Section 3. The productivity vector across plants of firm t, ψ_t, can be calculated given r_t and β. The density of the fixed cost draws across countries conditional on having chosen a plant in country l can be written as

$$u(f \mid Z_t, \psi_t, \beta) = \frac{Pr(Z_t \mid \psi_t, f) z(f \mid \beta)}{\int_f Pr(Z_t \mid \psi_t, f) z(f \mid \beta) df},$$

where

$$Pr(Z_t \mid \psi_t, f) = \int_\phi Pr(Z_t \mid \phi, f) k(\phi \mid \psi) d\phi,$$

and

$$k(\phi \mid \psi_t) = \frac{g(\phi) \left| \frac{\mathrm{d}\psi_t/\phi}{\mathrm{d}\psi_t} \right| \prod_{l \in Z_t} h(\frac{\psi_{t,l}(\tilde{w},\sigma_\epsilon)}{\phi} \mid \beta)}{\int_{\phi'} g(\phi') \left| \frac{\mathrm{d}\psi_t/\phi'}{\mathrm{d}\psi_t} \right| \prod_{l \in Z_t} h(\frac{\psi_{t,l}(\tilde{w},\sigma_\epsilon)}{\phi'} \mid \beta) d\phi'},$$

and

$$Pr(Z_t \mid \phi, f) = \{E_\epsilon(\Pi|\phi, Z_t, \epsilon, f; \beta) \geq E_\epsilon(\Pi|\phi, Z', \epsilon, f; \beta) \quad \forall Z'\}.$$

The mean of fixed costs for firm t is

$$\bar{f}^t = \int f u(f \mid Z_{t,l}, \psi_t, \beta) df,$$

and the average fixed cost in country l of firms that actually have a plant there is

$$= \frac{\sum_{t=1}^{T} \bar{f}_l^t \{l \in Z_t\}}{\sum_{t=1}^{T} \{l \in Z_t\}}.$$

Appendix D

Fit of the calibrated global production model

D.1 Bilateral trade shares

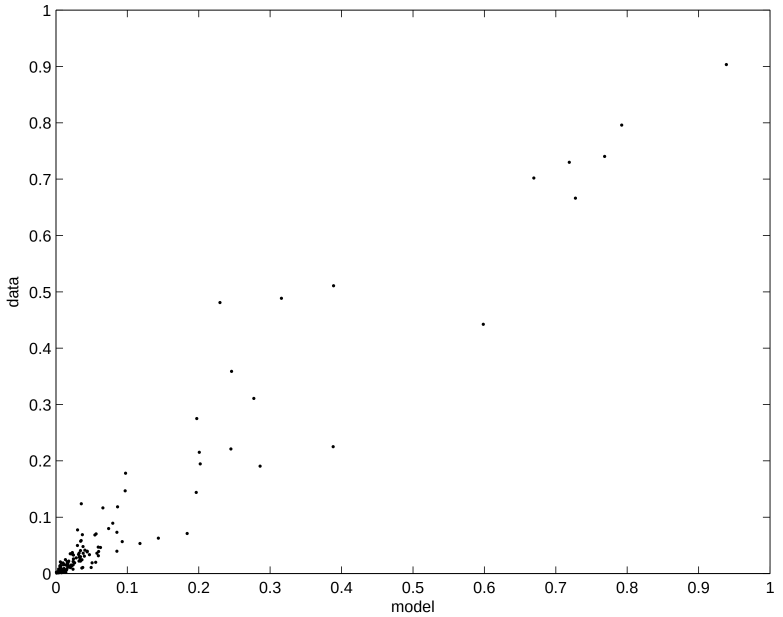

Figure 10: Bilateral trade shares - data and model

D.2 Bilateral MP shares

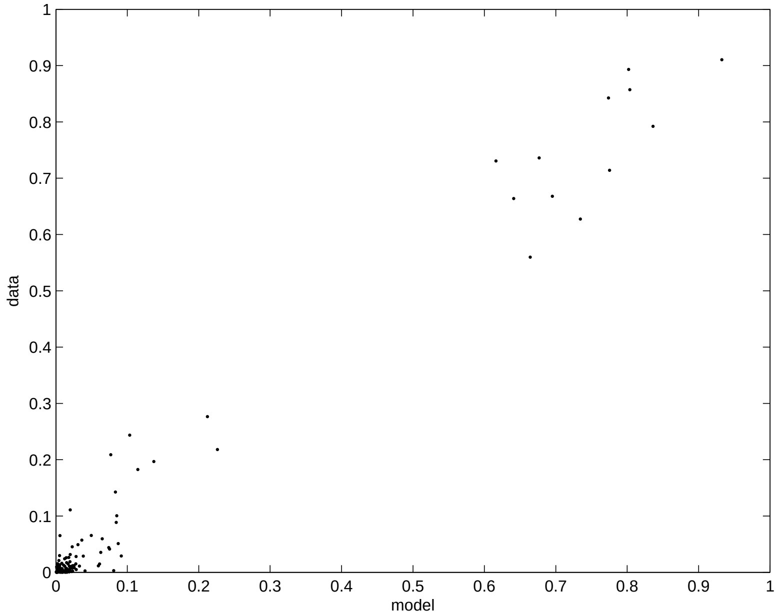

Figure 11: Bilateral MP shares - data and model

D.3 Variable production costs for German firms

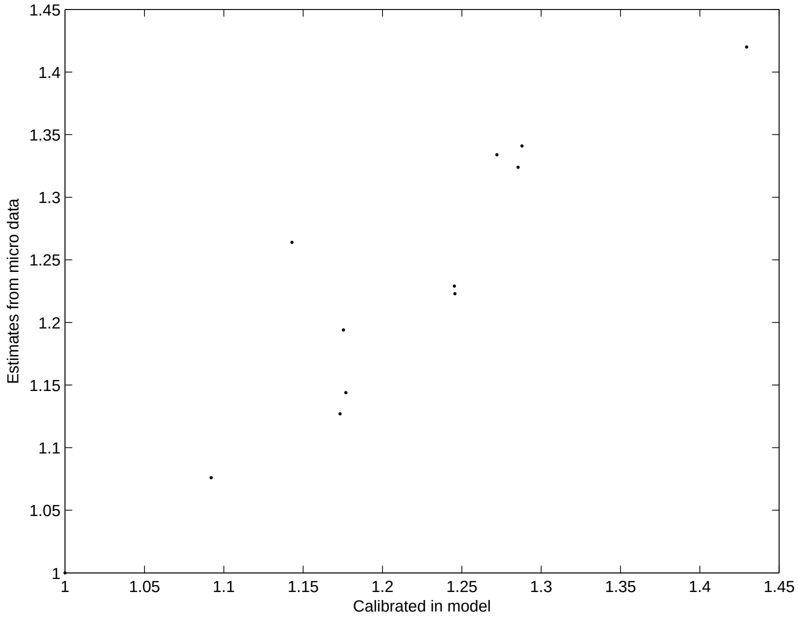

Figure 12: Variable production costs for German firms

Appendix E

Number of production locations and export platform shares

Commonly, the intuition is that additional production locations lower the average export platform shares of the firm. The export platform share of a plant is the plant's ratio of export to total sales. However, in general it is not true that any other additional plant decreases the export platform sales ratio of an existing plant: while it is true that the export platform sales decrease, it also matters by how much the local sales decrease. This section shows in a numerical example for a symmetric world that the export platform shares increase with fewer production locations. The numerical results are robust according to many different parameters. Nevertheless, it is crucial that trade costs between foreign countries are larger than domestic trade costs, which seems to be a plausible assumption.

I specify the following parameter values: $\sigma = 6$, $\theta = 7$, $\tau_{lm} = 1.6$, $\gamma_{il} = 1.2$. Figure 13 displays the export platform shares for plant $l \neq i$ as the number of plants increase. The export platform shares fall from 40 percent for a firm with just 2 plants to 29 percent for a firm with 12 plants.[1]

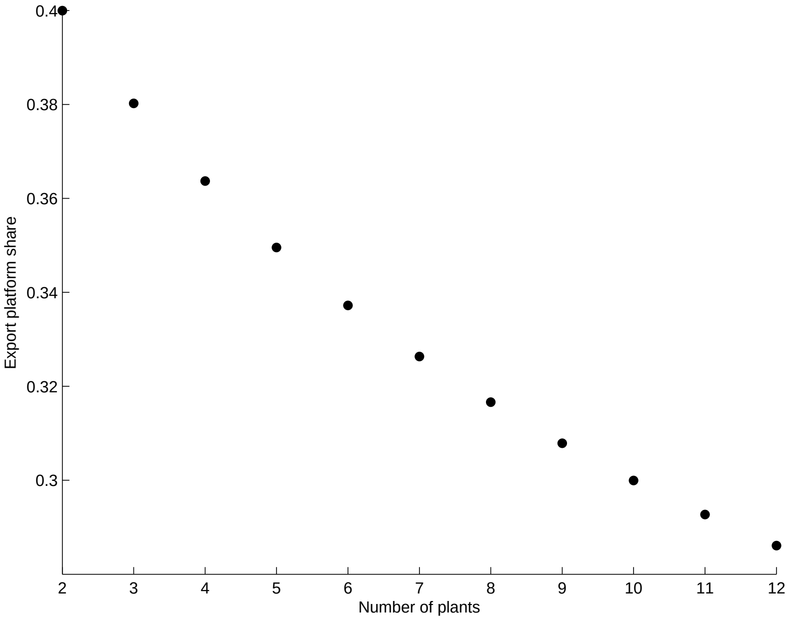

Figure 13: Export platform shares - Numerical example

---

[1] The export share of the domestic plant falls from 49 to 35 percent.

Appendix F

Results for a model with export platforms but without fixed costs

Here, I present the results for a calibrated model without fixed costs to establish foreign plants. Missing fixed costs imply that every firm establishes a plant in every country, which is obviously contrary to the firm-level evidence presented in Section 3. I calibrate the model to match aggregate trade and MP shares (the variable production cost estimates for German multinationals are not included as targets, because those were estimated from a model with both fixed and variable costs).

One can observe that this restricted model fits the MP data much worse; it does a slightly better job at fitting the bilateral trade data, but the sum of the two fitting norms is considerably higher. The share of export platform sales are on average 1.4 times lower than in the full model with fixed cost.

Table 20: CALIBRATED PARAMETERS - RESTRICTED GLOBAL PRODUCTION MODEL WITHOUT FIXED COST

	Parameters
Trade cost	
constant, β^{τ}_{const}	0.786
distance, β^{τ}_{dist}	0.118
language, β^{τ}_{lang}	0.925
contiguity, β^{τ}_{contig}	0.936
Variable MP cost	
constant, β^{γ}_{const}	1.870
distance, β^{γ}_{dist}	0.025
language, β^{γ}_{lang}	1.011
contiguity, β^{γ}_{contig}	0.847
Norm trade fit	0.221
Norm MP fit	0.339

Table 21: Export platform shares - Data and Models

	Data	Global Production Model	Restricted Global Production Model
Austria	0.54	0.58	0.45
Belgium	0.63	0.79	0.74
Canada	0.37	0.28	0.23
Switzerland	0.69	0.76	0.67
Germany	0.47	0.25	0.16
Spain	0.41	0.17	0.10
France	0.37	0.29	0.19
United Kingdom	0.34	0.23	0.14
Ireland	0.80	0.73	0.62
Italy	0.33	0.16	0.09
Netherlands	0.62	0.59	0.48

Notes: Global Production model: Full model characterized in the main text. Restricted Global Production Model: Model without fixed costs as described in the appendix.

Appendix G

Special case: gains from technology improvements

Section 1.5.2 on the benefits of foreign technology has two main results. The first result is that starting from the calibrated model, the magnitude of the gains in foreign countries is much larger if multinational production is taken into account. The second result is that with multinational production the gains from a technology improvement by factor $x > 1$ may yield welfare gains to that country by factor $y > x$. In order to demonstrate the economics behind the second results, I develop an analytic example in this section. In the example, I show that the size of the welfare gains of the country whose technology improved turns on how much the country's firms can increase their world market share in production.

Proposition 3. *Consider a symmetric world with identical size of the labor force in every country and* $\tau_{lm} = 1$, $\gamma_{il} = 1$, $\eta_{il} = 0$, $\forall i, l, m$. *Suppose* $\sigma = 6$, $M_i = L = 1$, $N = 3$, $x = 1.2$. *Then, an increase in productivity to one country by factor* x *raises its welfare by factor* $y > x$.

I only show the key expressions. Detailed derivations are available from the author upon request. I abstract away from firm heterogeneity (it does not matter for the results) and denote the productivity of all firms in country i by $\phi(i)$.

Welfare under the old scenario, $\phi(i) = \phi$ $\forall i$, is:

$$\frac{Y_1}{P} = \frac{\frac{\sigma}{\sigma-1}L}{N^{-1/\theta}\left(\sum_i M_i \kappa \phi(i)^{\sigma-1}\right)^{1/(1-\sigma)}}$$

Welfare under the new scenario, $\phi'(1) = x\phi$, $\phi'(j) = \phi$ $\forall j = 2, ..., N$, is:

$$\frac{Y_1'}{P'} = \frac{\frac{(\sigma-1+N\lambda_1')}{\sigma-1}L}{N^{-1/\theta}\left(\sum_i M_i \kappa \phi'(i)^{\sigma-1}\right)^{1/(1-\sigma)}}$$

where λ_i denotes the market share of firms from country i in the expenditures of each country:

$$\lambda_i = \frac{M_i \phi(i)^{\sigma-1}}{\sum_k M_k \phi(k)^{\sigma-1}}$$

Note that the equilibrium price index will always change at a rate less than the factor of technology improvement to country 1's firms, x. However, if the market share of country 1 goes up enough, which depends on the size of σ, the ratio of the two welfare expressions may exceed x. Plugging in the numbers, $\lambda_1' = 0.5544$ instead of the old $\lambda_1 = 1/3$. Relative price index: $\frac{P'}{P} = 0.9226$ and the welfare change in country 1 is 1.2036. For a lower value of σ, the welfare in country 1 would have increased less.

Appendix H

Potential effects from an EU-US trade and investment agreement

As a comparison to the potential effects from CETA, which is currently under negotiation, I also compute the potential effects from a hypothetical EU-US agreement that would lower variable and fixed foreign production costs between the signatories by the same proportion. As expected, the effects on the non-signatory partners from such an agreement would be even larger: the share of EU multinationals' production in Canada would fall from 9 to 6 percent, and the welfare in Canada would fall by almost half of a percent.

Table 22 contains the predicted outcomes for an EU-US agreement that lowers both variable and fixed MP costs between the EU countries and Canada by 20 percent. The structure of this table is analogous to Table 5 in the main text.

Table 22: Counterfactual changes of lower EU-US MP costs

	Difference in inward MP-shares		Rel. welfare
	Canada	United States	
Canada	1.93	-0.23	99.54
EU countries	-2.84	7.06	100.82 - 101.78
Switzerland	0.08	-0.03	99.68
United States	0.83	-6.80	100.91

Notes: Counterfactual: Reduction in variable and fixed MP costs between EU and US by 20 percent. First two columns: Differences in MP shares, κ_{il}, before and after the counterfactual change (column: destination l, row: source i).

Appendix I

Wind turbine installations data

I.1 Description

The register of Danish wind turbines is publicly available from the Danish Energy Agency (`http://www.ens.dk/en-US/Info/FactsAndFigures/Energy_statistics_and_indicators/OverviewOfTheEnergySector/RegisterOfWindTurbines/Sider/Forside.aspx`). This dataset spans the entire universe of Danish turbine installations since 1977 until the most recent month. The data on German installations is purchased from the private consulting company Betreiber-Datenbasis (`http://www.btrdb.de/`). Typically, several turbines are part of one wind farm project. The German data comes with project identifiers. We aggregate Danish turbines into projects using the information on installation dates, cadastral and local authority numbers. Specifically, turbines installed in the same year, by the same manufacturer, under the same cadastral and local authority number are assigned to the same project. The fine level of disaggregation provided by cadastral and local authority numbers minimize the measurement error.

Data on manufacturer locations was hand-collected from firms' websites and contacts in the industry. As of 1995 and 1996, seven out of ten large firms we use for our analysis were operating a single plant. Bonus, Vestas and Nordex had secondary production facilities. For these firms, we use the headquarters. Our industry contacts verified that these headquarters were also primary production locations with the majority of value-added. Equipped with the coordinates of projects and production locations, we calculated road distances as of June 2011 using the Google Maps API (`http://code.google.com/apis/maps/`). Therefore, our road distances reflect the most recent road network. For developed countries such as Germany and Denmark, we believe the error introduced by the change in road networks over time is negligible. Using direct great-circle distances in estimation generated virtually the same results.

I.2 Project Characteristics

Table 23, and Figures 14-16 provide some summary statistics on project characteristics in the two countries. Differences in distance to producers reflect heterogeneity in country size. Evidently, key observable characteristics such as electricity generating capacity, tower height and rotor diameter are remarkably similar in the two markets, ruling out product differentiation as a source of market segmentation. Slightly higher tower heights in Germany are due to lower wind speeds in southern regions. In such an environment, larger turbines are needed to attain the same capacity. What matters for this paper is that wind conditions do not change at the border. The European wind atlas available at the following link verifies that this is the case. (`http://www.wind-energy-the-facts.org/en/appendix/appendix-a.html`).

Table 23: SUMMARY STATISTICS OF PROJECTS

		Denmark	Germany
Capacity (KW)	Mean	475.81	472.59
	St. Dev.	207.93	175.98
	Median	600	500
	$10th$ percentile	225	225
	$90th$ percentile	600	600
Tower height (m)	Mean	38.34	49
	St. Dev.	7.96	8.64
	Median	40	50
	$10th$ percentile	30	40
	$90th$ percentile	46	65
Rotor diameter (m)	Mean	37.43	38.51
	St. Dev.	9.13	7.02
	Median	42	40.3
	$10th$ percentile	29	29.5
	$90th$ percentile	44	44
Distance to the border (km)	Mean	159.38	296.88
	St. Dev.	72.33	162.23
	Median	169.45	295.12
	$10th$ percentile	51.59	90.68
	$90th$ percentile	242.58	509.20
Distance to producers* (km)	Mean	154.02	366.58
	St. Dev.	31.26	100.19
	Median	169.45	344
	$10th$ percentile	117.52	258.20
	$90th$ percentile	192.65	510.78
Number of turbines per project	Mean	1.94	1.95
	St. Dev.	2.07	2.52
Number of projects	1995-1996	296	929
	1997-2005	1373	4148

Notes: Summary statistics of product characteristics in the first six panels are from the sub-sample of projects installed in 1995-1996. Onshore projects only.
(*): Average distance to firms with positive sales in that market.

I.3 List Prices

The survey of the German wind turbine market published by Interessenverband Windkraft Binnenland (various years) provides information on list prices for various turbine models as advertised by producers. These prices, however, are only suggestive and do not reflect project-specific final transaction prices. We use this information to verify the validity of our constant-returns-to-scale assumption. Figure 17 plots the per kilowatt price of various models against their total kilowatt capacity. Evidently, there are increasing returns up to 200 KWs. Beyond that range, per unit prices are mostly flat. As Figure 16 shows, a majority of the turbines installed in this period were in the 400-600 KW range.

Figure 14: KW Capacity Histograms by Market

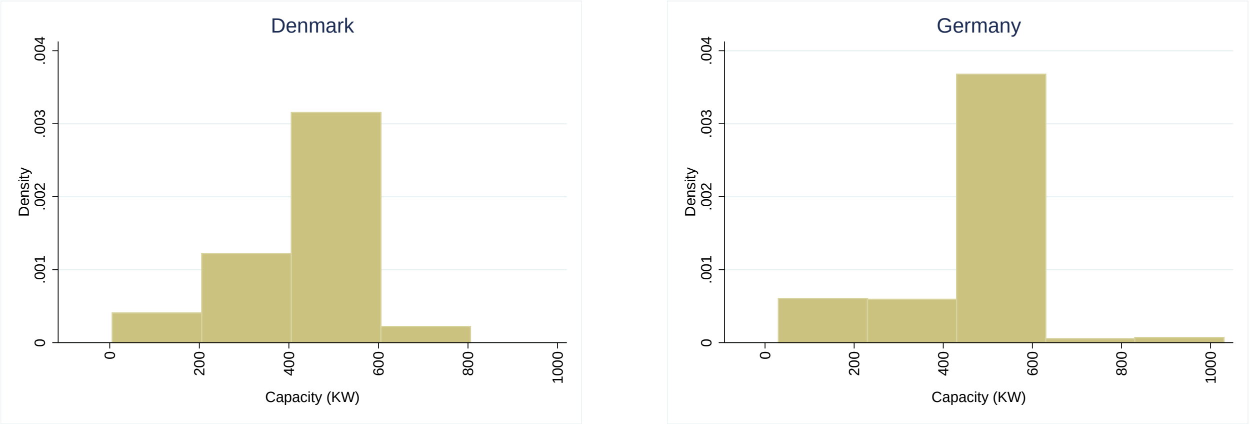

Notes: An observation is average kw capacity of turbines in a project. Years 1995 and 1996 only.

Figure 15: Tower Height Histograms by Producer and Market

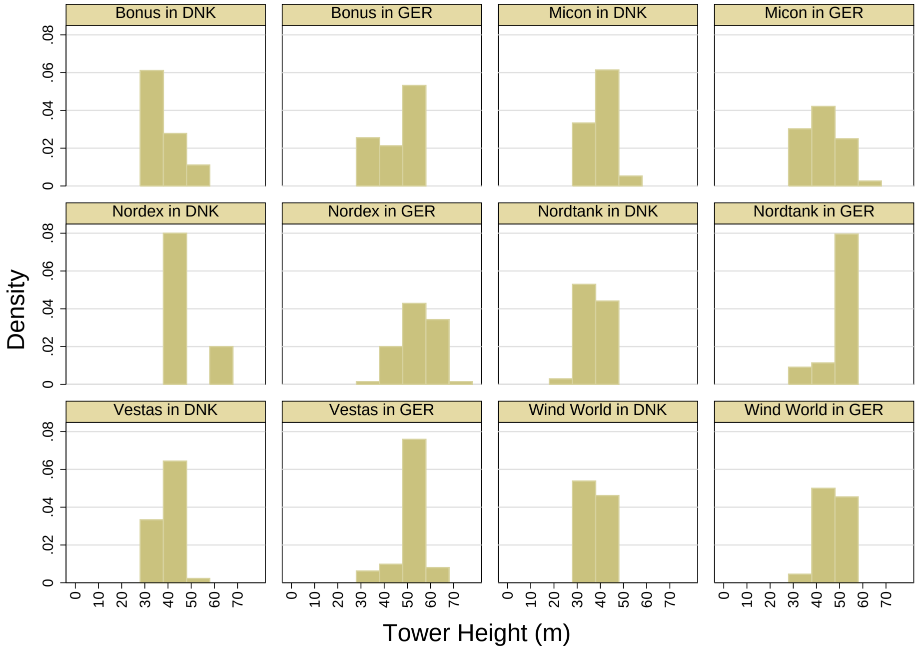

Notes: An observation is average tower height of turbines in a project. Years 1995 and 1996 only. "Bonus in DNK (GER)" indicates projects supplied by Bonus in Denmark (Germany).

Figure 16: KW Capacity Histograms by Producer and Market

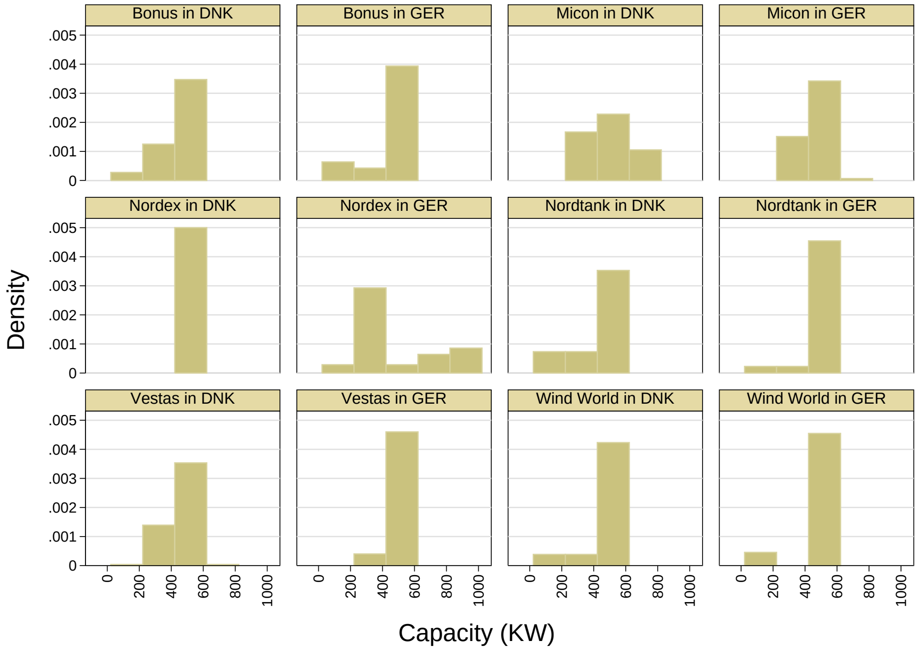

Notes: An observation is average kw capacity of turbines in a project. Years 1995 and 1996 only. "Bonus in DNK (GER)" indicates projects supplied by Bonus in Denmark (Germany).

Figure 17: Per KW List Prices of Various Turbines Offered in 1995-1996

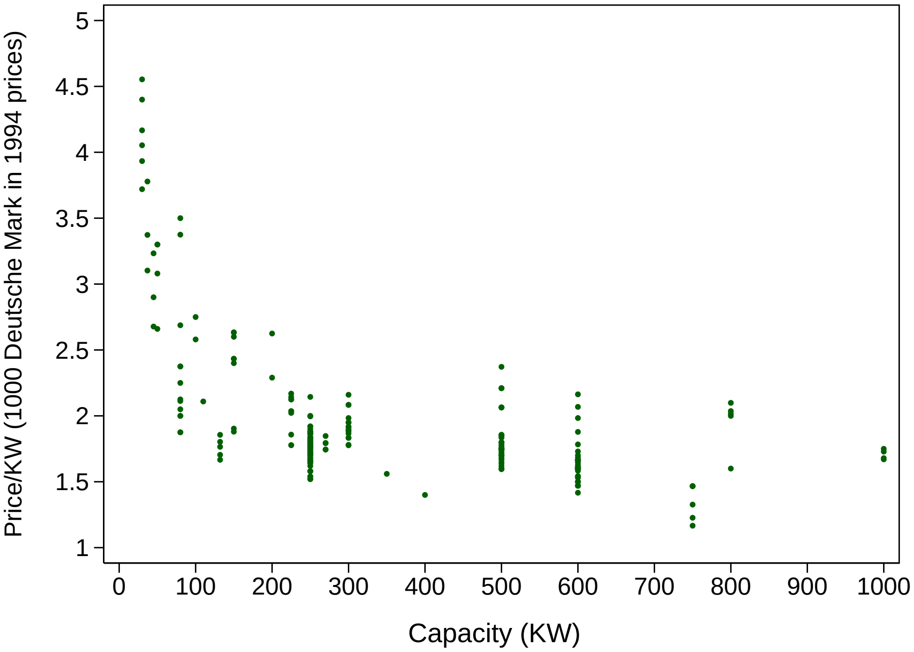

Notes: Pooled over all producers.

I.4 Regression Discontinuity Design

We estimate the following linear probability model in Subsection 2.2.2:

$$y_i = \alpha_0 + \sum_{k=1}^{k=3} \alpha_k \cdot \text{distance}_i^k + \gamma \cdot \text{Germany}_i + \sum_{k=1}^{k=3} \eta_k \cdot \text{distance}_i^k \cdot \text{Germany}_i + \epsilon_i. \qquad (44)$$

The dependent variable is $y_i = 1$ if the producer of project i is one of the five large Danish firms, and zero otherwise. The variable distance_i is the distance to the border. The effect of the border is picked up by the dummy variable Germany_i that equals one if the project is in Germany, and zero otherwise. The parameter of interest is γ. We trim the sample at the bottom and top 1 distance percentiles which keeps 1201 out of 1225 observations. We do this for expositional purposes related to the fit plots: while estimated regression coefficients are robust to using the entire sample, the fit on few observations in the tails are imprecise (see Figure 7). Table 24 reports the results for various specifications estimated with robust standard errors. The first column is the baseline featuring a cubic polynomial and interaction terms which allow distance to have a different effect on the two sides of the border. The border coefficient γ is significantly negative and of comparable magnitude in all four regressions.

To verify that we are not focusing on a peculiar period by using data from 1995-1996 in our structural estimation, we run the baseline RDD estimation for two subsequent subperiods, 1997-1998 and 1999-2005. The last subperiod pools data over seven years to ensure that there are enough observations in the neighborhood of the border at both sides. This becomes an issue because of the saturation of the Danish market after late 1990s. Table 25 reports the border dummy and the constant for all three periods. In all cases, Germany dummy is negative and statistically significant at 1% or 5% level. Moreover, the border effect is very stable over time.

Table 24: RDD RESULTS FOR THE 1995-1996 PERIOD

	Baseline Specification	Cubic No interactions	Linear	Linear No interaction
Germany (γ)	-0.289**	-0.319***	-0.413***	-0.423***
	(-2.52)	(-4.78)	(-7.38)	(-8.53)
Constant (α_0)	0.925***	0.807***	0.851***	0.862***
	(10.37)	(20.71)	(19.14)	(38.98)
Distance				
α_1	.0017	-7.7e-04***	-4.68e-04**	-3.91e-04***
	(0.59)	(-4.25)	(-2.01)	(-4.09)
α_2	1.46e-05	1.14e-07		
	(0.6)	(0.24)		
α_3	2.80e-08	1.08e-09		
	(0.47)	(1.11)		
Interactions				
η_1	-.0048		-8.19e-05	
	(-1.62)		(-0.32)	
η_2	-5.55e-06			
	(-0.22)			
η_3	-3.63e-08			
	(-0.61)			
Observations	1201	1201	1201	1201
Adjusted R^2	0.279	0.274	0.271	0.271

Notes: t statistics in parentheses

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

Table 25: RDD Results for Subsequent Periods

	(1) 1995-1996	(2) 1997-1998	(3) 1999-2005
Germany	-0.289**	-0.262**	-0.335***
	(-2.52)	(-2.44)	(-3.14)
Constant	0.925***	0.892***	0.822***
	(10.37)	(17.65)	(8.76)
Observations	1201	1254	3452
Adjusted R^2	0.28	0.38	0.18

Notes: t statistics in parentheses

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

Appendix J

Robustness to Local Unobservables and Economies of Densities

In order to derive the pricing equation, our model assumes that turbine manufacturers are independently maximizing project-level profits and that the unobservable shock to project owners' profits, ϵ_{ij}^{ℓ}, is unknown to firms, but drawn from a known distribution which is independent across projects and firms. Thus, we abstract away from the existence of spatial autocorrelation of unobservables across projects, economies of density in project location, and spatial collusion among turbine manufacturers. This section assesses whether this assumption has the potential to bias our estimate of the border effect.

There are several reasons for being concerned about the independence assumption which underlies the pricing equation. The assumption will be violated if firms directly observe sources of firm-project cost variation which are not explicitly controlled for by the model. While we feel that firms' productivity levels, firm-project distances, and the border dummy are the primary determinants of costs, other potential sources of variation could relate to unobservable local conditions being more amenable to a particular firm (e.g., local politics or geographic features of an area could result in lower cost for some firms). The independence assumption will also be violated if economies of density can be realized by a firm constructing several projects located geographically close together. Economies of density might be present if, for example, clustering projects together reduces travel costs for routine maintenance. Such economies of density might make the individual projects less expensive to maintain on a per-unit basis, leading firms with nearby installed projects to have a cost advantage over other firms that is not recognized in our model. Finally, if firms are colluding, then they are not maximizing prices, and the entire model is misspecified.

The existence of local unobservables would generate spatial autocorrelation in the error terms between projects which are geographically close. These could be due to unobserved characteristics of the terrain or local population which favor one manufacturer over another. Such an unobservable could also represent a spatially collusive agreement between firms to advantage a particular firm in a particular region. The existence of these unobservables would violate our assumption that the errors are independent across projects. Moreover, if firms are responding to economies of density of projects, firms pricing decisions become dynamic in nature. Since winning a project today lowers the firms' costs on other projects in the future, firms would not choose prices to maximize project-level profits, but rather the present discounted value of profits on this project and future projects. In short each of these forces, spatial unobservables, economics of density, and collusion, would lead firm's projects to be more tightly clustered together than our model would predict, leading to spatial autocorrelation in firms' error terms across projects. To test for the presence of spatial autocorrelation, we consider the following parametric model for the error term,

$$\epsilon_j = \gamma + \psi W \epsilon_j + \nu_i. \tag{45}$$

Here, ϵ_j is the vector of private shocks for firm j in all projects, γ is Euler's constant—the mean of the extreme value distribution, W is a known spatial weight matrix that determines the degree of influence one project has on another, and ν_i are independent and identically distributed mean-zero shocks. The scalar ψ determines the degree of spatial autocorrelation, we wish to test the null hypothesis that spatial autocorrelation is not present, i.e., that $\psi = 0$ and the ϵ_{ij} are in fact independent across projects.

In order to perform the test, we must specify the spatial weight matrix W. An element of the spatial weight matrix, W_{ik} provides an indication of how strongly project k is related to project i. Clearly many different specifications are possible, including inverse distance (measured either directly or though a road network), inclusion within the same region, or nearest neighbor adjacency. In practice we specify W as,

$$W_{ik} = \begin{cases} 1 & \text{if } dist(i,k) < 30\text{km}, \\ 0 & \text{otherwise,} \end{cases}$$

where distance is the direct distance (as the crow flies) in kilometers between projects i and j.[1]

We are unable to directly test for spatial autocorrelation in ϵ^ℓ_{ij} because as with all discrete choice models, ϵ^ℓ_{ij} is not directly recoverable. Instead, we follow Pinkse and Slade (1998) and test our results for spatial autocorrelation using the generalized errors. The generalized errors are the expectation of ϵ^ℓ_{ij} conditioned on the observed data and the model being correctly specified. Given the structure of our model, the generalized errors can be derived using the extreme-value density function,[2]

$$\hat{\epsilon}^\ell_{ij} = \begin{cases} \gamma - \log \rho^\ell_{ij} & \text{if } y^\ell_{ij} = 1, \\ \gamma + \frac{\rho^\ell_{ij}}{1-\rho^\ell_{ij}} \log \rho^\ell_{ij} & \text{if } y^\ell_{ij} = 0. \end{cases}$$

Again γ, represents Euler's constant—the unconditional expectation of the extreme value distribution. While the derivation of these expectations is algebraically tedious, the result is intuitive: the more likely a manufacturer j is to be selected by the project manager, the lower ϵ^ℓ_{ij} must be in order for selection to occur. Hence, $\hat{\epsilon}^\ell_{ij}$ is decreasing in the ex-ante probability of firm j being selected. The fact that the distribution of $\hat{\epsilon}^\ell_{ij}$ conditional on j not being chosen is independent of the actual choice observed in market i is a consequence of the well known independence of irrelevant alternatives (IIA) property of extreme-value discrete choice models. Note that, if the null hypothesis of no autocorrelation is violated, $\hat{\epsilon}^\ell_{ij}$ will be mis-specified. Nonetheless, they are useful to conduct a hypothesis test that $\psi = 0$.

We can use ordinary least squares to estimate ψ from the equation,

$$\hat{\epsilon}_j = \gamma + \psi W \hat{\epsilon}_j + \nu_i$$

and test whether $\psi = 0$. Note that, the estimate we generate, $\hat{\psi}$, is only consistent under the null hypothesis since the null is assumed in the construction of $\hat{\epsilon}_j$ and ordinary least squares is only consistent if $\psi = 0$.

The results are reported in Table 26.[3] While the magnitude of the estimated $\hat{\psi}$ is small, the

---

[1] Our results are robust to raising or lowering the distance cutoff and using a specification of W based on inverse distance.

[2] The derivation is available from the authors upon request.

[3] It is important that the test be conducted with heteroskedasticity-robust variance estimates, since there is little reason to believe that the generalized errors are homoscedastic.

Table 26: Results from Auto-Correlation Tests

Manufacturer	$\hat{\psi}$	Std. Error	t-Stat.
Fringe	0.023	0.007	3.375
Bonus (DK)	0.028	0.006	4.900
Nordtank (DK)	0.024	0.004	5.770
Micon (DK)	0.030	0.004	7.121
Vestas (DK)	0.035	0.005	7.083
WindWorld (DK)	0.030	0.007	4.594
Enercon (DE)	0.046	0.007	6.508
Fuhrlaender (DE)	0.035	0.005	6.394
Nordex (DE)	0.048	0.006	7.798
Suedwind (DE)	0.045	0.015	3.098
Tacke (DE)	0.033	0.005	6.908

test strongly rejects the null hypothesis for every firm, due in part to the the high precision of the estimates. We conclude that some degree of spatial autocorrelation is present, although it appears to be mild.

The presence of spatial autocorrelation has the potential to bias our estimate of the border effect. In particular, if spatial autocorrelation is due to cost or demand advantages in installing near already completed projects constructed by the same manufacturer, and if exporters have a smaller installed base within a country than do domestic firms, then the border effect may be capturing differences in the installed bases of foreign and domestic firms in addition to the variable cost of exporting. Alternatively, if serial correlation is due to local unobserved characteristics then the location of previous installations, while not cost reducing in and of themselves, serve as proxies for unobservable local conditions. In this spirit, we propose the following specification to check the robustness of our results to mild spatial autocorrelation. We re-estimate the model with the augmented cost function,

$$c_{ij} = \phi_j + \beta_d \cdot \text{distance}_{ij} + \beta_b \cdot \text{border}_{ij} + \beta_c \cdot \text{installed}_{ij},$$

where,[4]

$$\text{installed}_{ij} = \begin{cases} 1 & \text{if firm } j \text{ installed a turbine within 30km of project } i \text{ between 1991 and 1994,} \\ 0 & \text{otherwise.} \end{cases}$$

The new coefficient, β_c is able to capture the relationship between previously installed turbines and the costs of future projects. We are unable, however, to determine whether β_c is a causal effect, a proxy for local unobservables, or some combination of the two. Firms within our model continue to price according to static profit maximization. They do not take into account the possibility that building a turbine will make nearby projects less expensive in the future. This is consistent with the idea that the existence of local installations being merely a proxy variable and having no causal impact on future costs.

The results from this robustness specification are presented in Table 27. The coefficient on having a nearby installation has the expected negative sign (nearby installations are indicative of lower

---

[4] We have also experimented with a using distance to the nearest installed project in the cost function and using only projects installed between 1993 and 1994, and have found qualitatively similar results.

Table 27: ROBUSTNESS CHECK: NEARBY INSTALLED TURBINES

	Coefficient	Std. Error
Border Variable Cost, β_b	0.692	(0.178)
Distance Cost (100km), β_d	0.136	(0.031)
Nearby Installation, β_c	-1.249	(1.166)
Firm Fixed Effects, ξ_j		
Bonus (DK)	1.254	(0.189)
Nordtank (DK)	1.460	(0.183)
Micon (DK)	2.045	(0.160)
Vestas (DK)	2.682	(0.170)
WindWorld (DK)	0.638	(0.211)
Enercon (DE)	2.717	(0.147)
Fuhrlaender (DE)	-0.012	(0.266)
Nordex (DE)	0.855	(0.184)
Suedwind (DE)	-0.190	(0.206)
Tacke (DE)	1.576	(0.152)
Log-Likelihood	-2284.52	
N	1225	

costs) and is of substantial magnitude, but is statistically insignificant. The estimates of both distance costs, β_d and variable border costs, β_b both decrease slightly, but remain strongly significant. The estimated impact of the border relative to distance actually increases from 432 km to 502 km. Overall, these results appear to indicate that while unobservable local conditions of economies of density may induce some spatial autocorrelation between projects, the effect is mild and is not substantially impacting our primary results on the size of the border effect. In future work, we hope to investigate whether there is a causal effect of installations on the cost of future projects, but this question will require a fully dynamic pricing model which is outside the scope of our investigation of border costs.

Appendix K

Computational Method

K.1 Estimation of the Project Bidding Game

We formulate the estimation of the project bidding game as a constrained optimization problem.[1] The objective is to maximize the likelihood function subject to satisfying the firm-project specific winning probabilities expressions that come out of our model. We reformulate the problem defined in (39) for the computational implementation. The reformulated constraints are mathematically equivalent to those in (39). They come with two major advantages: First, when we reformulate the system maximizing the log-likelihood instead of the likelihood function, and rewrite the constraints, we are removing most of the nonlinearity. Second, winning probabilities only affect their respective equation and the adding-up constraint for the respective project. The sparse structure of the Jacobian of the constraints makes this large optimization problem feasible. The reformulated problem is

$$\begin{aligned}
\max_{\theta,\,\rho} \quad & \sum_{\ell\in\{D,G\}} \sum_{i=1}^{N_\ell} \sum_{j=0}^{|\mathcal{J}_\ell|} y_{ij}^{\ell} \log \rho_{ij}^{\ell} \\
\text{subject to:} \quad & \log \rho_{ij}^{\ell} - \log \rho_{i0}^{\ell} = \xi_j - \beta_d \cdot \text{distance}_{ij} - \beta_b \cdot \text{border}_{ij}^{\ell} - \frac{1}{1-\rho_{ij}^{\ell}} \\
& \sum_{k=1}^{|\mathcal{J}_\ell|} \rho_{ik}^{\ell} + \rho_{i0}^{\ell} = 1 \\
& \\
& \text{for } \ell \in \{D,G\},\ i \in \{1,...,N_\ell\},\ j \in \mathcal{J}.
\end{aligned}$$

For the baseline estimation, there are 11 constraints for every German project, and 7 constraints for every Danish project ($|\mathcal{J}_G| = 10$ and $|\mathcal{J}_D| = 6$ plus one fringe firm in every market). Since we have 929 German and 296 Danish project this aggregates to 12,291 constraints. In our baseline specification we are choosing 12,303 variables (12 structural parameters and 12,291 equilibrium win probabilities for each firm in each market)

We use the constrained optimization solver KNITRO to solve the problem. To improve speed and accuracy of the estimation, we hand-code the analytical derivatives of the object of function and the constraints and provide the sparsity structure of the Jacobian to the solver. In order to find a global maximum we pick 10 random starting values for the structural parameters. The estimation converges to the same solution for all attempted starting values.

---

[1] See **?** for a seminal paper that explains why constrained optimization of structural models is often superior to estimation via nested fixed points.

We calculate the covariance matrix of the parameter estimates using the outer product rule.

1. First, we calculate the score of each winning firm project pair, $\partial \log \rho_i^* / \partial \theta$, using numerical derivatives. This involves perturbing the $\widehat{\theta}$ vector. Note that the step size to perturb θ should be larger than the numerical tolerance level of the equilibrium constraints. Then the equilibrium constraints are resolved.

2. We then calculate the inverse of the covariance matrix:

$$\widehat{S}(\widehat{\theta}) = \sum_{i=1}^{N} \frac{\partial \log \rho_i^*(\widehat{\theta})}{\partial \theta} \frac{\partial \log \rho_i^*(\widehat{\theta})}{\partial \theta}'$$

K.2 Counterfactuals

The point estimate $\widehat{\theta}$ automatically satisfies the equilibrium constraints in the benchmark scenario with fixed entry and variable border costs. In the counterfactual "No fixed border costs" we use $\widehat{\theta}$ to then resolve the equilibrium constraints, with every firm being active in every market, $|\mathcal{J}_D| = |\mathcal{J}_G| = 10$. In the counterfactual "No border costs" we resolve the same system of equilibrium constraints with the variable border cost coefficient set to zero.

We use a parametric bootstrap procedure to calculate the standard errors for our counterfactuals. We draw 200 parameter vectors from the distribution of estimated parameters (multivariate normal distribution with mean θ and covariance matrix $\widehat{S}(\widehat{\theta})^{-1}$). First we resolve the baseline equilibrium constraints, then the constraints for the scenario with no fixed entry costs, and finally the constraints for the no border costs scenario (with each firm active in every market and the variable border costs coefficient set to zero). We store the equilibrium outcomes from each of these draws and use them to calculate the standard errors for our counterfactuals.

Vita

Felix Tintelnot

Education

Ph.D. in Economics, The Pennsylvania State University, May 2013.

Diplom in Economics (with distinction), Free University Berlin, July 2008.

Research Fields

Primary : International Trade.

Secondary : Industrial Organization.

Research Papers

"Global Production with Export Platforms", (December 2012).

"Borders, Geography, and Oligopoly: Evidence from the Wind Turbine Industry", with Kerem Cosar and Paul Grieco (December 2012).

"Exporting and the Environment: A New Look with Micro-Data", with Aoife Hanley and Sourafel Girma (June 2008).

Fellowships

German National Academic Foundation (*Studienstiftung des Deutschen Volkes*), 2005-2008.

Institute for Computational Economics Fellow, University of Chicago, 2010.

Personal

Date of Birth: 20/06/1983.

Place of Birth: Wuerzburg, Germany.

Gender: Male.

Languages: German (native), English (fluent).