BORDERS, GEOGRAPHY, AND OLIGOPOLY: EVIDENCE FROM THE WIND TURBINE INDUSTRY

A. Kerem Coşar, Paul L. E. Grieco, and Felix Tintelnot*

Abstract—Using a microlevel data set 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. In the data, firms' market shares drop precipitously at the border. We find that 40% to 50% of the gap can be attributed to national border costs. Counterfactual analysis indicates that eliminating national border frictions would increase total welfare in the wind turbine industry by 4% in Denmark and 6% in Germany.

I. Introduction

DISTANCE and political borders lead to geographic 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 structure.[2] To avoid these potentially confounding effects, we use spatial microdata from wind turbine installations in Denmark and Germany to estimate a structural model of oligopolistic competition with border frictions. Our main findings are that (a) border frictions are large within the wind turbine industry, (b) fixed and variable costs of exporting are both important in explaining overall border frictions, and (c) these frictions have a substantial impact on welfare.

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 the energy portfolios of many countries, 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. National subsidies are directed only toward the generation of renewable electricity. By the Single European Act, they do not discriminate against other European producers of turbines. The border costs in this setting are therefore due to factors other than formal barriers to trade.

Despite substantial formal integration, the data indicate significant market segmentation between Denmark and Germany. Examining the sales of turbines in 1995 and 1996, we find that domestic manufacturers have a substantially higher market share than foreign manufacturers. For example, the top five German manufacturers possess a market share of 60% in Germany and only 2% in Denmark. There appear to be border frictions on both the extensive and intensive margins: in the extensive margin, only one of the five large German firms exports to Denmark. In contrast, all five large Danish firms have sales in Germany. In the intensive margin, however, their market share is substantially lower in the foreign market and drops discontinuously at the border.

We propose a model to explain these patterns and study the welfare implications of the border. Firms are heterogeneous in their qualities, costs, and primary manufacturing location. The model has two stages. In the first stage, turbine producers decide whether to export. Exporting firms must pay a fixed entry cost specific to them. 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. From the sourcing decision of project managers, we can identify internal and international variable border costs by exploiting variation across firms and projects in their distance to federal and national borders. Since borders are responsible for a discrete jump in costs, they can be separated from the steady increase in costs related to distance. The model thus delivers endogenous variation in prices, markups, and market

Received for publication April 16, 2013. Revision accepted for publication May 19, 2014. Editor: Gordon Hanson.

* Coşar: University of Chicago; Grieco: Pennsylvania State University; Tintelnot: University of Chicago and Princeton University.

For their insightful comments and constructive suggestions, we thank Costas Arkolakis, Lorenzo Caliendo, Allan Collard-Wexler, Steve Davis, Jonathan Eaton, Charles Engel, Oleg Itskhoki, Samuel Kortum, Kala Krishna, Jan De Loecker, Matthias Lux, Ferdinando Monte, Eduardo Morales, Brent Neiman, Dennis Novy, Ralph Ossa, Ayşe Pehlivan, Joris Pinkse, Mark Roberts, Andrés Rodríguez-Clare, Alan Spearot, Elie Tamer, Jonathan Vogel, Stephen Yeaple, Nicholas Ziebarth, the editor, two anonymous referees, and numerous seminar participants. For their valuable discussions with us about the industry, we thank Wolfgang Conrad, Robert Juettner, Alex Kaulen, and Jochen Twele.

A supplemental appendix is available online at http://www.mitpressjournals.org/doi/suppl/10.1162/REST_a_00485.

[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. The effectiveness of domestic regulation in some industries may hinge on the extent of trade exposure, as shown by Fowlie, Reguant, and Ryan (2012) for the U.S. Portland cement industry.

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

The Review of Economics and Statistics, July 2015, 97(3): 623–637

doi:10.1162/REST_a_00485

shares across points in space, so we are able to analyze the impact of the border on trade flows, as well as producer and consumer surplus.

Our results indicate substantial variable and fixed costs to sell wind turbines across the border between Denmark and Germany.[3] Whereas German firms face nonnegligible variable costs when competing outside their home state, the variable border costs associated with the national border are roughly 85% higher. While fixed foreign market entry costs are not point identified by our model, we are able to gauge their significance through counterfactual analysis. We conduct two counterfactual experiments in which we first eliminate fixed entry costs and then all international border frictions. Market segmentation declines as we remove frictions to both the extensive and intensive margins. Overall, we find that the elimination of international border frictions raises consumer surplus by 8.6% and 8.8% in Denmark and Germany, respectively. Total surplus increases by 4.3% in Denmark and 5.6% 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. Early contributions by McCallum (1995) and Anderson and van Wincoop (2003) use data on interstate, interprovincial, and international trade flows 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. Alternatively, Engel and Rogers (1996), Gopinath et al. (2011), and Goldberg and Verboven (2001, 2005) have documented market segmentation by studying internal versus cross-border price dispersion.

Rather than inferring a "border effect" or "width of the border" based on differences between intra- and international trade flows or price differentials, we estimate a structural model of market segmentation using spatial microdata. By doing so, we address several critiques raised by the literature. Hillberry (2002), Hillberry and Hummels (2008), and Broda and Weinstein (2008) have argued that sectoral, geographic, and product-level aggregation may lead to upward bias in the estimation of the border effect in studies that use trade flows. Holmes and Stevens (2012) emphasize the importance of controlling for internal distances. Our study addresses these critiques since our data enable us to precisely 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. In addition, we can use our model to quantify the producer and consumer surplus implications of cross-border trade barriers.

[3] We assume that all border frictions are related to costs rather than home bias on the part of project managers. In consumer goods industries, preferences may discontinuously change at the border if consumers act on a home bias toward domestic producers. In our setting, where consumers are profit maximizers purchasing an investment good, we expect that demand-driven home bias is small. Alternatively, we can interpret our cost estimates as incorporating the additional costs that exporting firms must incur to overcome any home bias in preferences.

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.

II. 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. Large-scale production and installation of electricity generating wind-turbines became popular after the introduction of feed-in-tariff subsidies for wind energy generation in 1984 in Denmark and in 1991 in Germany. 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 remuneration for renewable energy is not conditional on purchasing turbines from domestic 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 small independent investors.[4] The turbine manufacturing industry, however, is dominated by a small number of firms that manufacture, construct, and maintain turbines on the project owner's land. Manufacturers usually have a portfolio of turbine designs available with various generating capacities. Overall, their portfolios are relatively homogeneous in terms of observable characteristics.[5]

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

[4] Small purchasers were encouraged by financial incentive schemes that gave larger remuneration to small producers such as cooperative investment groups 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%. The Danish support scheme provided 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.

[5] 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 appendix B in the online appendix.

delivery route. According to industry sources, transportation costs range between 6% and 20% of total costs (Franken & Weber, 2008). Plant-to-project distances also affect the cost of postsale services (such as maintenance), installing remote controllers to monitor wind farm operations, gathering information about sites farther away from the manufacturer's location, and maintaining relationships with local contractors who construct turbine towers.[6]

Intra- and international political boundaries impose other variable costs on firms. Industry experts highlight several sources of friction. In the case of state borders, these costs are related to administrative hurdles in coordinating transportation across different agencies, acquiring building permits, and interacting with regional operators to connect projects to the grid. The banking sector, which is critical for obtaining project financing, is also typically organized at the state or local level. Moreover, firms that are local employers benefit from greater visibility than their out-of-state competitors do. In addition to the cost of selling across an intranational border, the international border imposes even higher transaction costs. Additional channels include the cost of writing and enforcing international contracts and dealing with a different currency, language, and culture.

In contrast to distance and variable border costs, fixed market entry costs are incurred only on entering a foreign market. Differences in the electricity grid in Denmark and Germany require the development and installation of a country-specific software that regulates generation. Similarly, each turbine model undergoes a separate certification process in each country before it can be marketed. In order to overcome differences in language and business practices, firms may establish country-specific sales teams. These fixed entry costs may prevent a firm from competing for projects in the foreign market at all. Accounting for these costs will be important, as they may substantially change the market structure (i.e., the number of competitors) on either side of the border.

A. *Data*

We have collected data on every installation of a wind turbine in Denmark and Germany dating to the birth of the wind turbine industry. 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. Using the location of each manufacturer's primary production facility, we calculate road distances from each manufacturer to each project. This provides us with a spatial source of variation in manufacturer costs that aids in identifying the sources of market segmentation. While our data are rich in the spatial dimension, we do not observe transaction prices due to the business-to-business nature of the industry. (Appendix B provides a detailed description of the data.)

In this paper, we concentrate on the years 1995 to 1996. 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. This wave includes a cross-border acquisition, which would blur the distinction between a foreign and domestic firm and complicate our analysis of the border effect.[7] 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.

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 & Tybout, 1997; Das, Roberts, & 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. We must also abstract away from the possibility of collusion that could result from repeated interaction (Salvo, 2010), although we have no reason to expect that collusion occurred in this industry. 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.

Table 1 displays the market shares of the largest five Danish and German firms in both countries. We take these firms to be the set of manufacturers in our study. In the left panel of figure 1, we present the project locations using separate markers for German and Danish produced projects. The right panel provides the location of the primary production facility for each turbine manufacturer.

[6] For a rough comparison of the effect that distance has in this industry against common benchmarks in the literature, appendix A estimates a gravity equation on international trade in the six-digit HS 2007 product category associated with wind turbines. The results indicate that the industry is remarkably representative in terms of distance and contiguity. We take this as evidence that while distance is an important driver of costs in the industry, its effect is not inordinately large relative to other tradable manufactured goods.

[7] On the other hand, the specter of a merger wave presents the possibility of anticipatory effects. For example, if a firm was seen as a likely merger target, this might affect its reputation given that servicing a turbine in the future is typically the responsibility of the manufacturing firm. We control for these effects through firm fixed effects to allow the reputation of firms to be heterogeneous.

Table 1.—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

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

B. *Preliminary Analysis of the Border Effect*

Table 1 and figure 1 clearly suggest some degree of market segmentation between Germany and Denmark. Four of the five large German firms, including the German market leader, Enercon, have no 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 1), these fixed costs can be amortized over a larger number of projects in Germany.

For firms that do export, the lower market share in the foreign market 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 distance 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 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 in the left-hand panel of figure 2. This figure documents Vestas's declining market share as the distance from its main manufacturing location increases. While this figure suggests that costs may be increasing in 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 relation to all firms.

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 1 if it is supplied by a large Danish firm and 0 otherwise, to a cubic polynomial of distance from the project to the border, a Germany dummy (to capture the border effect), and interaction terms. The first column in table 2 reports the border dummy, which is a statistically significant -0.305. This market share drop of the largest five Danish firms is reflected in the right-hand panel of figure 2, which plots the fit of this regression (see appendix B.4 for details).

If the market share discontinuity at the border captures trade frictions, one may expect a declining effect throughout the 1990s, a period during which European integration deepened. To check this pattern and assess the representativeness of our 1995–1996 data, we estimate the RDD by pooling the data between 1982 and 2005 and allowing the border dummy to have a time trend. The second column of table 2 reports the results. There is a gradual but statistically significant declining trend in the market share discontinuity at the border. The market share discontinuity captured by the border dummy is 0.423 at 1982 and shrinks by about 1 percentage point annually. Enforcement of the European Single Market program, general reductions in trade costs due to globalization, and subsequent cross-border acquisitions and investments may have indeed reduced the frictions faced by foreign producers over time. Our period of study falls in the middle of this trend.

These results give us reason to believe that the border matters in the wind turbine industry; however, they tell 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. Motivated by this, the following section proposes a structural model that accounts for the change in market structure at the border.

III. Model

Denmark and Germany are indexed by $\ell \in \{D, G\}$. Each country has a discrete 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), both of which are exogenous. 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 to pay a fixed cost to 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.

A. *Project Bidding Game*

In this stage, active firms offer a separate price to each project manager, and project managers choose the offer that

FIGURE 1.—PROJECT AND PRODUCER LOCATIONS

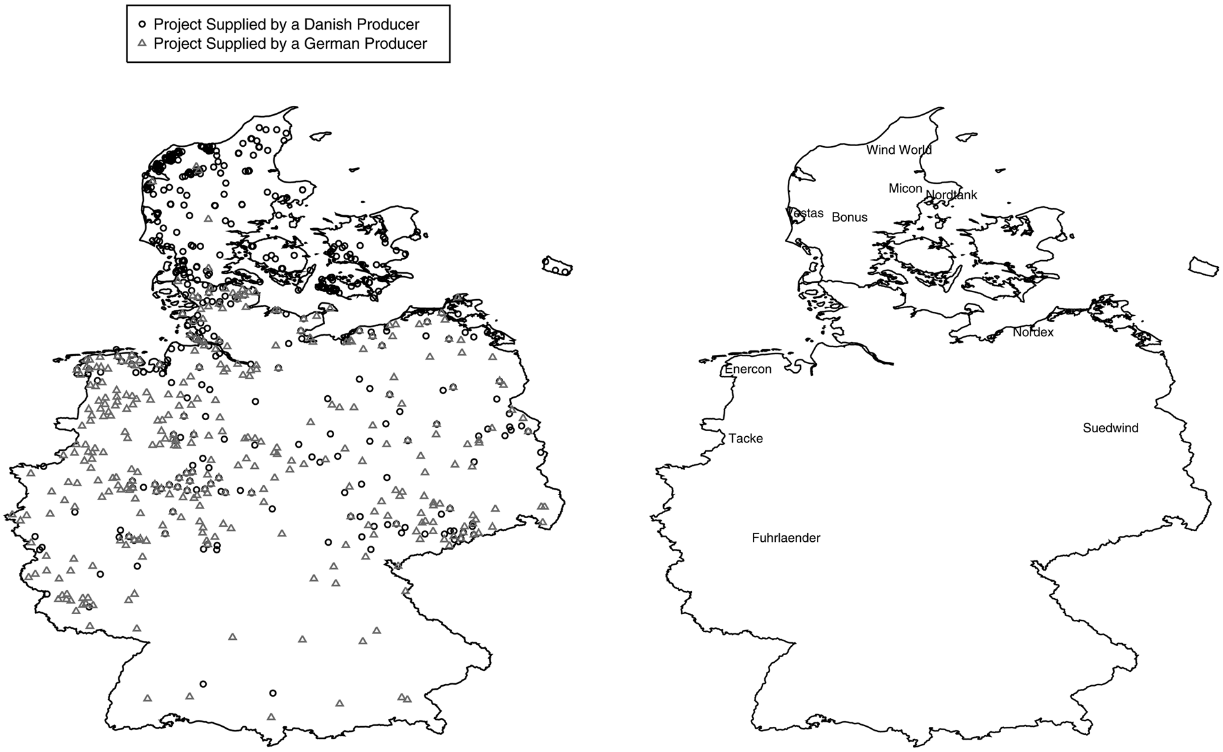

FIGURE 2.—MARKET SHARES BY DISTANCE AND ACROSS THE NATIONAL BORDER

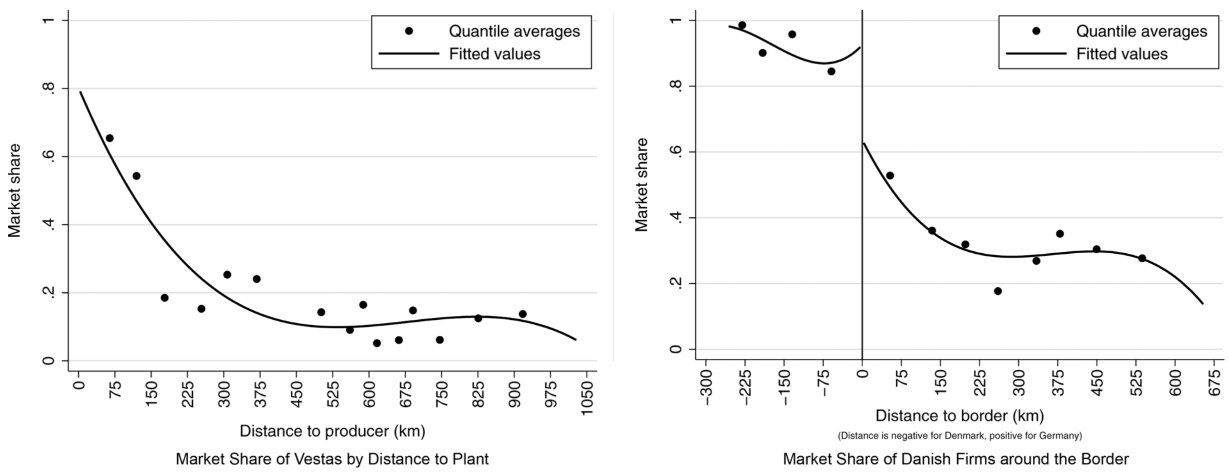

(Left) The line follows from the linear regression of a dummy variable that takes the value 1 if a project is supplied by Vestas and 0 otherwise, on a cubic polynomial of projects' road distances to Vestas's headquarters. Dots are local market shares (i.e., proportion of projects supplied by Vestas) within fifteen distance quantiles. (Right) The line follows from the linear regression of a dummy variable that takes the value 1 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 B.4. Dots are local market shares (i.e., proportion of projects supplied by Danish firms) within twelve distance quantiles.

TABLE 2.—RDD RESULTS

Germany	−0.305	−0.423
	(0.126)	(0.057)
Time trend		0.011
		(0.003)
Data period	1995–1996	1982–2005
Observations	1,226	9,622
R^2	0.284	0.274

Standard errors in parentheses.

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,[8] and an idiosyncratic choice-specific shock ϵ_{ij}.[9] It is well known that discrete choice models identify only relative differences in valuations. We thus model a nonstrategic 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 1 extreme value distribution. The $\boldsymbol{\epsilon}_i$ vector is private information to project managers who collect project-specific 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.[10] 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

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

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 it is a foreign or domestic out-of-state producer:

$$c_{ij} = \phi_j + \beta_d \times \log(\text{distance}_{ij}) + \beta_b \times \text{border}_{ij} + \beta_s \times \text{state}_{ij}, \tag{2}$$

where the dummy variable border_{ij} equals 1 if i and j are located in different countries and 0 otherwise. Similarly, state_{ij} equals 1 if both i and j are located in Germany, but in different states, and 0 otherwise.[11] Due to data limitations, this cost function is meant to capture and distinguish in a reduced form those costs that are related simply to distance (e.g., shipping costs, communication difficulties) from those that directly relate to political boundaries (differences in laws and regulations) and those specifically related to national boundaries (cultural and language differences and international contracting). While we are unable to directly understand why distance and political boundaries both impart costs on trade, we believe our study takes a step in the direction of understanding the role these costs play in segmenting national and international markets.

Firms engage in Bertrand competition by submitting price bids for projects in the markets in which they are active.[12] 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:[13]

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

where S_i is the size of the project in megawatts. Firm i's first-order condition is

$$0 = \frac{\partial \rho_{ij}(p_{ij}, \mathbf{p}_{i,-j})}{\partial p_{ij}} (p_{ij} - c_{ij}) + \rho_{ij}(p_{ij}, \mathbf{p}_{i,-j}),$$

$$p_{ij} = c_{ij} - \frac{\rho_{ij}(p_{ij}, \mathbf{p}_{i,-j})}{\partial \rho_{ij}(p_{ij}, \mathbf{p}_{i,-j}) / \partial p_{ij}}.$$

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

$$p_{ij} = c_{ij} + \frac{1}{1 - \rho_{ij}(p_{ij}, \mathbf{p}_{i,-j})}. \tag{3}$$

The markup 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 equation (3) into equation (1), we arrive at a fixed-point

[8] Since we do not directly observe prices, we will use the manufacturer's first-order condition to derive prices in units of the project owner's 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.

[9] We assume away project-level economies of scale by making price bids per megawatt. In appendix B, 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.

[10] 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 C.3, we address the robustness of our estimate to local unobservables of this type.

[11] Unlike federal Germany, Denmark has a unitary system of government, so we treat Denmark as a single entity.

[12] Industry experts we interviewed indicated 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.

[13] We assume that firms are maximizing expected profits on a project-by-project level. This abstracts away from economies of density in project locations—the possibility that by having several projects close together, they could be produced and maintained at a lower cost. We address the robustness of our model to the presence of economies of density in appendix C.3.

problem with $|\mathcal{J}|$ unknowns and $|\mathcal{J}|$ equations for each project i:

$$\rho_{ij} = \frac{\exp\left(d_j - c_{ij} - \frac{1}{1-\rho_{ij}}\right)}{1 + \sum_{k=1}^{|\mathcal{J}|} \exp\left(d_k - c_{ik} - \frac{1}{1-\rho_{ik}}\right)} \quad \text{for } j \in \mathcal{J}. \quad (4)$$

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 markup 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.

Our approach bears a strong resemblance to models of differentiated demand used in industrial organization (Berry, 1994; Berry, Levinsohn, & Pakes, 1995). There are two key differences. First, the traditional approach assumes that the econometrician observes the overall market share of a product with a fixed set of characteristics within the market. In our case, because the turbine location affects each firm's costs, the characteristics of products are different at every project location. Since we have precise data on which manufacturers constructed which projects, we are thus able to exploit observed manufacturer-consumer differences (i.e., distance to project location) to identify trade costs. Second, the traditional approach requires that prices are observed. We do not observe transaction prices due to the business-to-business nature of the industry. To surmount this challenge, we assume manufacturers choose prices (and hence markups) for each project on the basis that a profit maximization condition derived from our model. 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. With price data, the traditional approach is able to allow for a market-level unobserved quality component, whereas we control for unobserved turbine quality through a firm fixed effect.

B. *Entry Game*

Before bidding on projects, an entry stage is played in which all large firms simultaneously decide whether to be active in the foreign market by incurring a firm-specific fixed cost f_j. This fixed cost 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.[14]

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:

$$\Pi_j(\mathcal{J}_{-j} \cup j) = \sum_{i=1}^{N} E[\pi_{ij}(\mathcal{J}_{-j} \cup j)]. \quad (5)$$

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

$$\Pi_j(\mathcal{J}_{-j} \cup j) \geq f_j. \quad (6)$$

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, equation (6) must hold for that firm. If firm j is not observed in the foreign market, one can infer the following lower bound on fixed export cost:

$$\Pi_j(\mathcal{J}_{-j} \cup j) \leq f_j. \quad (7)$$

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. This approach is similar to several studies that have proposed the use of bounds to construct moment inequalities in estimating structural parameters (Pakes et al., 2015; Eizenberg, 2014). Holmes (2011) and Morales, Sheu, and Zahler (2014) applied this methodology to the context of spatial entry and trade. Of course, because our only contain projects data from 1995 to 1996, our bounds do not account for the possibility of future payoffs resulting from the decision to be active in the foreign market during the sample period, as might occur if there were substantial sunk costs to initiate exporting relative to per period fixed costs. Accurately estimating sunk entry and fixed continuation costs separately would require a longer time period and a fully dynamic model. Moreover, because we observe only a single observation of each firm's entry decision, a moment inequality approach is not applicable in our setting: instead, we simply report the single bound for fixed cost imputed from the first stage. We now turn to the estimation of the model.

IV. Estimation

Estimation proceeds in two steps: In the first step, we estimate the structural parameters of the project-bidding game.

[14] One could imagine the entry decision being regional rather than nationwide. This does not appear to be the case in our data, as exporting Danish firms supply projects in most German states. Therefore, we maintain the assumption that fixed costs are paid at the national level while testing for the presence of state-level fixed costs in section IVB.

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 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.

We now reintroduce the country index: ρ_{ij}^{ℓ} is firm j's probability of winning project i in country ℓ, in which the number of active firms is $|\mathcal{J}_\ell|$. Substituting the cost function, equation (2), into the winning probability, equation (4), we find

$$\rho_{ij}^{\ell} = \frac{\exp\left(d_j - \phi_j - \beta_d \times \log(\text{distance}_{ij}) - \beta_b \times \text{border}_{ij} - \beta_s \times \text{state}_{ij} - \frac{1}{1-\rho_{ij}^{\ell}}\right)}{1 + \sum_{k=1}^{|\mathcal{J}_\ell|} \exp\left(d_k - \phi_k - \beta_d \cdot \log(\text{distance}_{ik}) - \beta_b \cdot \text{border}_{ik} - \beta_s \cdot \text{state}_{ij} - \frac{1}{1-\rho_{ik}^{\ell}}\right)}. \tag{8}$$

From this equation, one can see that firms' production costs ϕ_j and quality level d_j are not separately identified given our data. 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 $\theta = (\beta_b, \beta_d, \beta_s, \xi_1, \ldots, \xi_{|\mathcal{M}_D|+|\mathcal{M}_G|})$. We estimate the model via constrained maximum likelihood, where the likelihood of the data is maximized subject to the equilibrium constraints, equation (8). The likelihood function of the project data has the following form:

$$L(\rho) = \prod_{\ell \in \{D,G\}} \prod_{i=1}^{N_\ell} \prod_{j=0}^{|\mathcal{J}_\ell|} (\rho_{ij}^{\ell})^{y_{ij}^{\ell}}, \tag{9}$$

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

$$\max_{\theta,\,\rho} \quad L(\rho)$$

subject to:

$$\rho_{ij}^{\ell} = \frac{\exp\left(\xi_j - \beta_d \times \log(\text{distance}_{ij}) - \beta_b \times \text{border}_{ij} - \beta_s \times \text{state}_{ij} - \frac{1}{1-\rho_{ij}^{\ell}}\right)}{1 + \sum_{k=1}^{|\mathcal{J}_\ell|} \exp\left(\xi_k - \beta_d \times \log(\text{distance}_{ik}) - \beta_b \times \text{border}_{ik} - \beta_s \times \text{state}_{ij} - \frac{1}{1-\rho_{ik}^{\ell}}\right)} \tag{10}$$

$$\sum_{k=1}^{|\mathcal{J}_\ell|} \rho_{ik}^{\ell} + \rho_{i0}^{\ell} = 1 \quad \text{for } \ell \in \{D, G\},\, i \in \{1, \ldots, N_\ell\},\, j \in \mathcal{J}.$$

Examining equation (10) provides straightforward intuition for identification of the model. The model implies a probability that each manufacturer builds each product. These are directly related to the individual firm's competitiveness, its cost to build each product, and its optimal markup—a function of its own and other firm's costs. As a project moves closer to or farther away from a firm, its costs will vary, allowing us to identify the impact of costs directly. Crossing an internal or international boundary results in a discontinuous change in the firm's predicted probability of winning, which can be separated from the smooth effects of distance. The proximity of a firm to other producers, while not affecting its cost, also affects its markup and probability of winning. The maximum likelihood estimator searches for the parameterization of the model that best matches the pattern of manufacturer choice observed in the empirical distribution. We describe the details of the computational procedure in appendix D.

Once the structural parameters are recovered, one can calculate bounds on the fixed costs of entry for each firm, f_j, using equations (6) and (7). 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.

A. *Parameter Estimates*

Estimation results are presented in table 3 starting in the first column with the baseline specification featuring national and state borders. The second column drops the state border, which is estimated to be significant in the baseline. In the third column, we bring back the state border but let the distance cost to be piecewise linear in three intervals in order to allow a more flexible specification in capturing the concavity of distance costs. Across all specifications, the national border coefficient is positive and statistically significant. Moreover, its magnitude is higher than the state border coefficient in the first and third columns.

While the state border reveals some regulatory hurdles that out-of-state producers face, such as the higher cost of obtaining local permits and coordinating transportation, the higher national border cost verifies the existence of additional frictions to exporting: dealing with foreign jurisdictions, visiting foreign locations for maintenance, and risks associated with long-term cross-border contracting and servicing. The baseline estimates indicate that the cost of crossing an international border is roughly 85% higher than that of crossing an internal boundary, a difference that is statistically significant.

Comparing the first and second columns, the elimination of the internal border causes the distance coefficient to fall and the national border coefficient to increase. This provides

TABLE 3.—MAXIMUM LIKELIHOOD ESTIMATES

	Baseline	National Border Only	Piecewise Linear Distance Costs
National Border Variable Cost, β_b	1.151	0.855	1.360
	(0.243)	(0.211)	(0.244)
State Border Variable Cost, β_s	0.622		0.799
	(0.223)		(0.212)
Log Distance Cost, β_d	0.551	0.679	
	(0.091)	(0.079)	
Distance, $[0, 50)$ km			6.096
			(1.475)
Distance, $[50, 100)$ km			0.442
			(0.616)
Distance, 100+ km			0.089
			(0.036)
Firm Fixed Effects, ξ_j			
Bonus (DK)	2.480	2.414	5.493
	(0.219)	(0.212)	(0.615)
Nordtank (DK)	2.531	2.492	5.487
	(0.225)	(0.221)	(0.625)
Micon (DK)	3.085	3.036	6.091
	(0.211)	(0.209)	(0.621)
Vestas (DK)	3.771	3.710	6.756
	(0.208)	(0.204)	(0.615)
WindWorld (DK)	1.641	1.594	4.570
	(0.256)	(0.255)	(0.623)
Enercon (DE)	3.859	3.526	6.850
	(0.208)	(0.166)	(0.605)
Fuhrlaender (DE)	0.598	0.199	3.465
	(0.324)	(0.302)	(0.566)
Nordex (DE)	2.198	1.806	5.198
	(0.235)	(0.188)	(0.609)
Suedwind (DE)	0.566	1.028	4.054
	(0.259)	(0.303)	(0.636)
Tacke (DE)	2.749	2.403	5.806
	(0.210)	(0.167)	(0.607)
Log likelihood	−2,333.76	−2,338.19	−2,328.99
N	1,225	1,225	1,225

Standard errors in parentheses. Distance is measured in units of 100 km.

some indication that controlling for internal borders is important to consistently recovering the impact of the national boundary. In particular, when the state border is not included, distance will act as an imperfect proxy for state borders, as higher distances will be correlated with crossing a state border, leading to an upward omitted variable bias. Similarly, since exporters do not face the internal border cost by construction, the state border dummy is negatively correlated with the national border dummy, leading to a downward bias when the state dummy is omitted.

The third column replaces the log distance specification with a piecewise linear specification. This confirms the concavity in distance costs implied by the use of log distance in the baseline. Distance costs are extremely steep very close to the production facility but decline substantially beyond 50 kilometers, and even farther beyond 100 kilometers. The magnitudes of the border cost variables are robust to this specification. While the magnitudes of the firm fixed-effect estimates rise, their relative magnitudes are very similar. The change in magnitude simply reflects the fact that the vast majority of projects are beyond 50 kilometers, so the higher marginal distance costs very near a manufacturer (sometimes referred to as first-mile costs) are captured by the fixed effect in the log specifications.[15]

Although ignoring internal border frictions in the second column leads to an overstatement of its effect, distance is also a significant driver of costs. To get a sense of its importance under the baseline estimates, we calculate the distance elasticity of the equilibrium probability of winning a project for every firm-project combination. The median distance elasticity ranges from 0.36 to 0.54. That is, the median effect of a 1 percent increase in the distance from a firm to a project (holding all other firms' distances constant) is a decline of 0.36% to 0.54% in the probability of winning the project. So distance has a sizable impact on costs and market shares for all firms.[16]

As discussed above, the firm fixed effects reflect the combination of differences in quality and productivity across firms. We find significant differences among 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 3 presents this comparison using the baseline results. The horizontal axis is the distance to the Danish-German border, where negative distance is inside Denmark. The solid line is the raw data fit. This is the same curve as that presented in the right-hand panel of figure 2, relating the probability of a Danish firm's winning a project to distance to the national border and a border dummy. In particular, this regression does not control for project-to-firm distances. The dashed 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 national border. The nonlinearity we see in winning probabilities captures not only the nonlinearity in distance costs but also the rich spatial competition patterns predicted by the model. Overall, the model fits the data well.

[15] One might be concerned that the concavity of distance costs is an artifact of an endogenous location decision on the part of firms. While firms may locate in areas where demand will be high, an endogeneity problem would arise if, rather than simply because of high demand for turbines (which would raise the profitability of all producers), Vestas located its assembly facility in a location where demand for Vestas-made turbine demand is high relative to other manufacturers' products. As we have discussed, turbines are largely homogeneous, and the most region-specific attribute, tower height, is easily customizable. Also, in appendix C.3, we check the robustness of our results to local unobservables that favor firms heterogeneously.

[16] The distance elasticities we report are a function of the characteristics of all firms at a particular project site in a single industry. It is difficult to directly compare them with gravity-based distance elasticities from the literature that rely on national or regional distance proxies (McCallum, 1995; Eaton & Kortum, 2002; Anderson & van Wincoop, 2003).

FIGURE 3.—MODEL FIT: EXPECTED DANISH MARKET SHARE

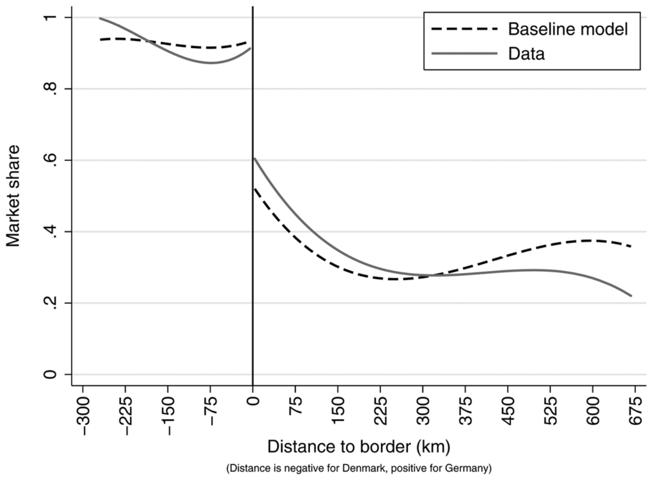

The data line is the same as the fitted values line in the right panel of figure 2. The model line is the linear fit of winning probability for each project by Danish firms on a cubic polynomial of projects' great circle distances to the border, a Germany dummy, and interaction terms.

Finally, our results relate to several studies that have attempted to get some sense of border magnitudes by reporting a border "width" (McCallum, 1995; Engel & Rogers, 1996) using market-to-market comparisons of prices and trade flows. We construct a similar statistic: the equivalent increase in distance that gives the same cost increase as crossing the national border as $\exp(\beta_b/\beta_d)$. Our baseline model implies an eight-fold increase in distance costs when crossing a border, while not controlling for internal boundaries causes this cost to fall to a 3.5-fold increase ($\exp(0.855/0.679) = 3.52$). Since the literature has typically not accounted for internal boundaries, the 3.5 figure is most appropriate for comparison. While large, both our figures are small relative to the Engel and Rogers's (1996) calculation. In a companion paper, we use a simulation exercise to illustrate how focusing on market-to-market price variation is susceptible to upward biases relative to source-to-market measures of border width due to specification error, measurement error, and omitted variable bias (Coşar, Grieco, & Tintelnot, 2015).

We include additional robustness checks and alternative specifications in the online appendix. In appendix C.2, we experiment with alternative specifications for the cost function of the firm, which allow for heterogeneity in distance cost coefficients (i.e., β_{dj} instead of common β_d), and scale economies in cross-border sales. In appendix C.3, we check the validity of the assumption on independent draws across projects, which may be violated due to the existence of spatial autocorrelation of unobservables across projects, economies of density, or spatial collusion among turbine manufacturers.[17] State and national border coefficients remain stable and significant across all these alternatives.

B. *Fixed-Cost Bounds*

Not all firms enter the foreign market; rather, firms optimally choose whether 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.[18] Since our model naturally allows for heterogeneity in firm operating profits, this section considers whether heterogeneity in firms' fixed costs of exporting is also needed to rationalize observed entry decisions.

Since we observe only 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 decisions based on their fixed market entry costs and on the operating profits they expect in the export market as described in section IIIB. Based on the parameter estimates in table 3, 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 equation (6): their fixed cost must have been lower than the expected value of entering the foreign market, for otherwise these firms would not have made any foreign sales. Similarly, equation (7) puts a lower bound on fixed costs for firms that stay out of the foreign market: their fixed costs must be at least as much as their expected profits from entering; otherwise, they would have bid on some foreign projects. 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.

Table 4 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 costs for entering Germany differ from those for entering Denmark. Since all Danish firms enter the Danish market, any fixed cost below 16.74 (the expected profits of WindWorld for entering Germany) would rationalize the observed entry pattern. In Germany, however, the lower and upper bounds of Enercon and Nordex have no intersection. Some background information about Nordex supports the implication of the model that Nordex may be

[17] Salvo (2010) offers a model of spatial competition in an oligopolistic industry where firms use geography to collude on higher prices. We do not believe spatial collusion is a likely explanation for the discontinuity in our setting. Danish firms were active throughout Germany during this period, and our analysis in appendix C.3 does not reveal a strong degree of spatial clustering that might be expected if firms were cooperatively splitting the wind turbine market across space. Moreover, the industry receives a high degree of regulatory scrutiny due to its importance in electricity generation. No antitrust cases have been filed with the European Commission against the firms studied in this paper.

[18] The canonical Melitz (2003) model assumes homogeneous 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.

TABLE 4.—EXPORT FIXED COST BOUNDS (f_j)

	Lower	Upper		Lower	Upper
Bonus (DK)	–	45.66	Enercon (DE)	25.22	–
		(5.65)		(8.72)	–
Nordtank (DK)	–	43.56	Fuhrlaender (DE)	0.91	–
		(5.28)		(0.59)	–
Micon (DK)	–	77.88	Nordex (DE)	–	7.34
		(8.08)			(3.13)
Vestas (DK)	–	156.12	Suedwind (DE)	1.70	–
		(13.84)		(0.83)	–
WindWorld (DK)	–	16.74	Tacke (DE)	8.77	–
		(3.04)		(3.38)	–

Scale is normalized by the variance of ϵ; see note 8. Standard errors in parentheses.

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, it could keep a foothold in the Danish market at a lower cost than other German firms, which would need to form contacts with Danish customers from scratch.[19]

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 et al. (2007). Nevertheless, our results suggest the degree of heterogeneity in fixed costs that is necessary to explain entry patterns.[20]

Our specification assumes that fixed entry costs are incurred at the national level. We think this is reasonable, as the biggest drivers in these fixed costs are associated with forming new sales and service teams to reduce transaction costs arising from lingual and cultural differences and dealing with foreign regulations and grid technology—factors that mostly vary by country rather than by state. To provide further reassurance, we use the model to test for the presence of state-level fixed entry costs. If these costs were a significant factor in firms' entry decisions, then our specification would incorrectly assume a firm competes in some region of Germany, say Bavaria, when in fact it does not. The model interprets zero wins in a given state as a firm simply losing all projects. But with state-level entry costs, the reality could be that it never competed at all. Therefore, a large number of "zeros" for a firm in terms of state-level number of projects supplied might be an indicator of state-level fixed costs.

There are 15 German states with at least one project. For the five Danish firms, this results in 75 state entry events. Of these, there are 28 instances where a Danish firm wins zero projects in a given state. On the other hand, in every German state, there is at least one Danish firm with positive sales. However, most of these zeros are for small states with very few projects, so it is reasonable to think that firms did compete but simply did not win any projects. To test this hypothesis against the alternative that the firm did not compete, we use the model to compute the implied probability a firm did compete in a state but did not win any projects, as assumed by our model. For the 28 cases where a Danish firm did not build a project in a German state, this probability is effectively a p-value of the null hypothesis above. In 25 of 28 cases, we fail to reject the null hypothesis that firms did compete and simply did not win; that is, in 25 of 28 cases, the p-value of the test is above 0.05. There are no instances where the model is rejected with 99% confidence: the p-value is never below 0.01. Likewise, we can test for the presence of state-level fixed costs among German firms. For German firms, there are 22 occasions when a firm fails to win a project in a German state. Running the identical test for each instance, we fail to reject the null hypothesis that the firm did compete but simply did not win any project (i.e., the p-values are always above 0.05).

While the above test by no means proves that state-level fixed costs do not exist, it do provide some comfort that the data do not strongly reject our assumption that fixed costs are incurred at the national level. The biggest worry relating to state-level fixed costs is that we are misspecifying the project managers' choice set of turbines. To be extra careful, we rerun the estimation eliminating the three states in which the model is rejected at the .05 level. This removes 272 projects from the data set. The coefficients for the national border, state border, and log distance remain significant and similar in magnitude.

V. 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 national and state borders.[21] Although 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 0, which implies that all firms enter the export market. The second step further reduces the variable cost of the national border by setting β_b equal to the state border coefficient, β_s.[22] In terms of

[19] Because of Nordex's connection to Denmark, we perform a robustness check by reestimating the model to allow 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.

[20] It is important to note that the variable cost estimates presented in table 3, 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 & Pakes, 1995).

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

[22] We first eliminate fixed costs and then change 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

TABLE 5.—COUNTERFACTUAL MARKET SHARES OF LARGE FIRMS (%)

		Data	Estimates	No Fixed Costs	No National Border Costs
Denmark	Danish firms	92.57	92.89	83.03	77.17
			(1.61)	(4.15)	(3.01)
	German firms	1.69	2.50	13.07	19.31
			(1.00)	(3.88)	(2.67)
Germany	Danish firms	32.29	32.12	32.12	42.10
			(1.49)	(1.49)	(4.60)
	German firms	59.63	59.40	59.40	51.07
			(1.57)	(1.57)	(4.03)

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

FIGURE 4.—COUNTERFACTUALS: EXPECTED DANISH MARKET SHARE

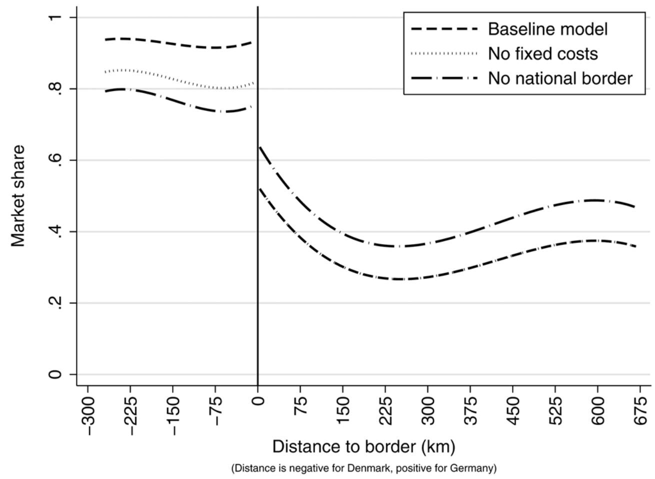

Regression discontinuity fit of projects won by large Danish firms under the baseline model and counterfactual scenarios. Since all Danish firms already compete in Germany, their market share does not change to the right of the border line when fixed costs are removed. See figure 3 for further details.

the model, this exercise makes Denmark simply another state of Germany.

A. Market Shares and Segmentation

We begin our analysis by considering how national market shares in each country react to the reduction 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 5 presents the market share of the major firms of Denmark and Germany in each country, with the fringe taking the remainder of the market. In a comparing of 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% 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 national border variable 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 by 10 percentage points. 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. 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. Because there are more projects in Germany than in Denmark, the payoff from entering Germany is much higher. This may be one reason that we see more Danish firms entering Germany than vice versa.[23] As a result, reducing fixed costs of exporting to Germany has no effect on market outcomes, whereas the impact of eliminating the fixed cost of exporting to Denmark is substantial.

The fourth and final column of table 5 displays the counterfactual market shares if the national border had the impact of only a state border. Here, in addition to setting f_j equal to 0 for all firms, we also reduce variable border costs by setting β_b equal to β_s. This results in a large increase in imports on both sides of the border. The domestic market share of Danish firms falls from 92.9% to 77.2%. 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 42% of the projects import Danish turbines once the national border is reduced to a state border, which reflects the strength of Danish firms (especially Vestas) in the industry.

Overall, our results indicate that national 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 partially responsible. However, when we remove national border frictions, our counterfactual analysis indicates that the gap shrinks to 35 percentage points. Almost half of the market share gap is thus attributable to national border frictions.

In addition to national market share averages, our model allows us to examine predicted market shares at a particular point in space. When the RDD approach describe above is used, figure 4 visualizes the impact of the counterfactual experiments. The dashed line represents expected market shares baseline model and is identical to that presented in figure 3. The dotted line displays counterfactual expected market shares when fixed border costs are removed. This reduces the domestic market share of Danish firms

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.

[23] This argument assumes that fixed costs of exporting are of the same order of magnitude for both countries.

TABLE 6.—COUNTERFACTUAL WELFARE ANALYSIS BY COUNTRY

		Baseline	No Fixed Costs		No National Border	
		(Levels)	(Levels)	(% Change)	(Levels)	(% Change)
Denmark	(A) Consumer surplus	70.63	74.35	5.26	76.69	8.58
		(3.68)	(3.03)	(1.89)	(3.21)	(1.52)
	(B) Danish firm profits	28.78	24.88	−13.55	22.68	−21.19
		(0.74)	(1.58)	(3.79)	(1.15)	(2.62)
	(C) German firm profits	0.59	3.24	446.37	4.92	729.51
		(0.25)	(1.04)	(76.90)	(0.74)	(231.52)
	Domestic surplus (A + B)	99.41	99.23	−0.18	99.37	−0.04
		(4.27)	(4.15)	(0.19)	(4.10)	(0.21)
	Total surplus (A + B + C)	100.00	102.46	2.46	104.28	4.28
		(4.10)	(3.49)	(1.08)	(3.58)	(0.92)
Germany	(A) Consumer surplus	68.04			73.99	8.75
		(2.62)			(3.39)	(4.33)
	(B) Danish firm profits	10.03			13.44	34.06
		(0.52)			(1.59)	(16.54)
	(C) German firm profits	21.94			18.18	−17.14
		(0.89)			(1.97)	(7.82)
	Domestic surplus (A + C)	89.97			92.17	2.44
		(3.00)			(2.86)	(1.37)
	Total surplus (A + B + C)	100.00			105.61	5.61
		(3.03)			(3.44)	(2.88)

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.

since more German firms enter, but it leaves market shares unchanged in Germany since all firms were already competing there. Finally, the dashed-dotted line shows the counterfactual estimates when the national border is turned into a state border. The discontinuity at the border remains due to the state border costs but is substantially reduced.

B. 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 6 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.[24] The relative changes in consumer surplus across scenarios are invariant to the scale of ϵ, so we normalize the consumer surplus in the baseline scenario to 100 for expositional ease.

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

The next two columns present results from the counterfactual where only fixed costs of entry are removed. We report both the levels, and percentage changes from baseline levels. 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 5%. Danish firms, facing harsher domestic competition, see profits decline by 14%. 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 3% of the available surplus in Denmark in profits. The gains of Danish consumers from removing fixed export costs are almost perfectly offset by the losses from Danish firms. Domestic surplus actually declines, but the decline is economically negligible and statistically insignificant. When we account for the gains by German firms, total surplus increases by a statistically and economically significant 2.46%.

The final two columns of table 6 display the welfare effects of reducing the national border frictions to the level of a state border.[25] As we would expect, site owners see significant benefits, and consumer surplus rises by 9% in Denmark and Germany. These increases come at the cost of domestic producers, who see home profits decline by 21% in Denmark and 17% in Germany.[26] In Denmark, the removal of national border frictions results in a transfer of surplus from domestic firms to consumers, netting to essentially no change in domestic surplus. When we include the benefits of exporters,

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

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

[25] For the welfare analysis, our assumption that the barriers to trade are driven by costs, not preferences, is important. We think this assumption is plausible for this industry.

[26] Of course, these declines do not account for benefits realized in the export market. See appendix C.1 for an accounting of how each firm fares as both an domestic producer and an exporter under our counterfactual scenarios.

however, total surplus increases by 4%. The story in Germany is a bit different. Consumer gains outweigh domestic firm losses in Germany, and domestic surplus increases by 2%. Essentially, removing border frictions improves the access of German site owners 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 6%.

We conclude this section with an important disclaimer. Our second counterfactual represents a reduction of all national border frictions to the level of only a state border. 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 national border and its effect on firms and consumers in the wind turbine industry. Policymakers may view the results as an upper bound on what can be accomplished through economic and political integration.

VI. Conclusion

This paper uses spatial microdata 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 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 almost 50% of the gap in cross-border market shares can be attributed to national border costs. Our study makes some strides toward 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.

Of course, there is 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 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 a large capital good traded in a business-to-business market, in other industries, cross-border differences in preferences—in particular, home bias—may play an important role. This is particularly true in consumer goods markets.

Finally, 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 intentions 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. Our findings of large border frictions in the upstream market imply that harmonizing renewable energy tariffs may be necessary to ensure equal treatment of European firms in accordance with the principles of the European single-market project.

REFERENCES

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

Balistreri, E. J., and R. H. Hillberry, "Structural Estimation and the Border Puzzle," *Journal of International Economics* 72 (2007), 451–463.

Berry, S., "Estimating Discrete-Choice Models of Product Differentiation," *RAND Journal of Economics* 25 (1994), 242–262.

Berry, S., J. Levinsohn, and A. Pakes, "Automobile Prices in Market Equilibrium," *Econometrica* 63 (1995), 841–890.

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

Broda, C., and D. E. Weinstein, "Understanding International Price Differences Using Barcode Data," NBER working paper 14017 (2008).

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

Coşar, A. K., P. L. E. Grieco, and F. Tintelnot, "Bias in Estimating Border- and Distance-Related Trade Costs: Insights from a Oligopoly Model," *Economics Letters* 126 (2015), 147–149.

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

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

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

Eizenberg, A., "Upstream Innovation and Product Variety in United States Home PC Market," *Review of Economic Studies* 81 (2014), 1003–1045.

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

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

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

Fowlie, M., M. Reguant, and S. P. Ryan, "Market-Based Emissions Regulation and Industry Dynamics," NBER working paper 18645 (2012).

Franken, M., and T. Weber, "Heavy Duty," *New Energy* 2008 (2008), 28–37.

Goldberg, P. K., and F. Verboven, "The Evolution of Price Dispersion in the European Car Market," *Review of Economic Studies* 68 (2001), 811–848.

——— "Market Integration and Convergence to the Law of One Price: Evidence from the European Car Market," *Journal of International Economics* 65 (2005), 49–73.

Gopinath, G., P.-O. Gourinchas, C.-T. Hsieh, and N. Li, "International Prices, Costs and Markup Differences," *American Economic Review* 101 (2011), 2450–2486.

Gorodnichenko, Y., and L. L. Tesar, "Border Effect or Country Effect? Seattle May Not Be So Far from Vancouver After All," *American Economic Journal: Macroeconomics* 1 (2009), 219–241.

Hillberry, R. H., "Aggregation Bias, Compositional Change, and the Border Effect," *Canadian Journal of Economics* 35 (2002), 517–530.

Hillberry, R., and D. Hummels, "Trade Responses to Geographic Frictions: A Decomposition Using Micro-Data," *European Economic Review* 52 (2008), 527–550.

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

Holmes, T. J., and J. J. Stevens, "Exports, Borders, Distance, and Plant Size," *Journal of International Economics* 88 (2012), 91–103.

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

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

Morales, E., G. Sheu, and A. Zahler, "Gravity and Extended Gravity: Using Moment Inequalities to Estimate a Model of Export Entry," Harvard University mimeograph (2014).

Obstfeld, M., and K. Rogoff, "The Six Major Puzzles in International Macroeconomics: Is There a Common Cause?" (pp. 339–412), in Ben S. Bernanke and Kenneth Rogoff, eds., *NBER Macroeconomics Annual 2000, 15* (Cambridge, MA: National Bureau of Economic Research, 2001).

Pakes, A., J. Porter, K. Ho, and J. Ishii, "Moment Inequalities and Their Application," *Econometrica* 83 (2015), 315–334.

Roberts, M. J., and J. R. Tybout, "The Decision to Export in Colombia: An Empirical Model of Entry with Sunk Costs," *American Economic Review* 87 (1997), 545–564.

Salvo, A., "Trade Flows in a Spatial Oligopoly: Gravity Fits Well, But What Does It Explain?" *Canadian Journal of Economics* 43 (2010), 63–96.

Sijm, J., "The Performance of Feed-in Tariffs to Promote Renewable Electricity in European Countries," Energy Resource Center of the Netherlands, ECN-C–02-083 (2002).

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