

KOÇ UNIVERSITY-TÜSİAD ECONOMIC RESEARCH FORUM
WORKING PAPER SERIES

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

A. Kerem Cosar
Paul L. E. Grieco
Felix Tintelnot

Working Paper 1228
October 2012

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

A. Kerem Cosar
Chicago Booth

Paul L. E. Grieco
Penn State

Felix Tintelnot
Penn State

September 13, 2012

Abstract

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.

JEL Codes: F14, L11, L20, L60, R12

Keywords: trade costs, oligopoly, spatial competition, constrained MLE

*For their insightful comments, we thank Costas Arkolakis, Lorenzo Caliendo, Allan Collard-Wexler, Steve Davis, Jonathan Eaton, Charles Engel, Chang-Tai Hsieh, 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, and workshop participants at Penn State, Chicago Booth, Boston University, Boston FED, UCLA, U.S. Census Bureau, Ohio State, University of Minnesota, University of Munich, Koç University, Massachusetts Institute of Technology, Columbia University, Princeton University, EIT in Purdue University, Midwest Trade Conference, Cornell-PSU macro conference, NBER IO Summer Institute, LACEA-TIGN, Conference on Globalization in Denmark, and UECE Lisbon Meetings in Game Theory and Applications. For their valuable discussions with us about the industry, we thank Wolfgang Conrad, Robert Juettner, Alex Kaulen, and Jochen Twele.

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 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.

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 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.³ 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 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

¹Policy 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.

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

³All 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.

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.⁴ 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 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.⁵ 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, Hsieh, and Li (2011) 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

⁴Our 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.

⁵See Anderson and van Wincoop (2004) for a survey of the empirical literature.

that the markets are becoming more integrated over time.⁶

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 address 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. Holmes and Stevens (2012) 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.⁷ 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 6 presents an experiment based on our estimated model in which we regress imputed price differentials on distances and a border dummy to calculate the implied width of the border. This estimated width is substantially larger than the “true” width used to generate the data. A comparison of structural and reduced form equations illustrates the sources of bias.

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 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 4. In Section 5, we perform

⁶The 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.

⁷It 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.

a counterfactual analysis of market shares and welfare by re-solving the model without fixed and variable border costs. Section 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 7 with a discussion of policy implications.

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 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.⁸ 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.⁹ There 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 1). 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.

⁸Small 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.

⁹Main 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 8 in Appendix A.

Figure 1: TRANSPORTATION OF WIND TURBINE BLADES



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

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.¹⁰ Appendix A provides a detailed description of the data.

In this paper, we concentrate on the period from 1995 to 1996.¹¹ 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.¹²

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). Learning-by-doing would provide firms with an incentive to lower prices below a static profit maximizing level in return for anticipated dynamic gains.¹³ 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 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. 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

¹⁰As 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.

¹¹Appendix A.4 shows that the evidence on market shares and the border effect is stable in subsequent time periods.

¹²Moreover, the Danish onshore market saturates after the late 1990s, leaving us with little variation at that side of the border.

¹³In 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.

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

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

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 2, we present the project locations using separate markers for German and Danish produced projects. Figure 3 provides the location of the primary production facility for each turbine manufacturer.

2.2 Preliminary Analysis of the Border Effect

Table 1 and Figure 2 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 (930 projects were installed in Germany in this period, versus 296 in Denmark—see the map of projects in Figure 2), 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 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.

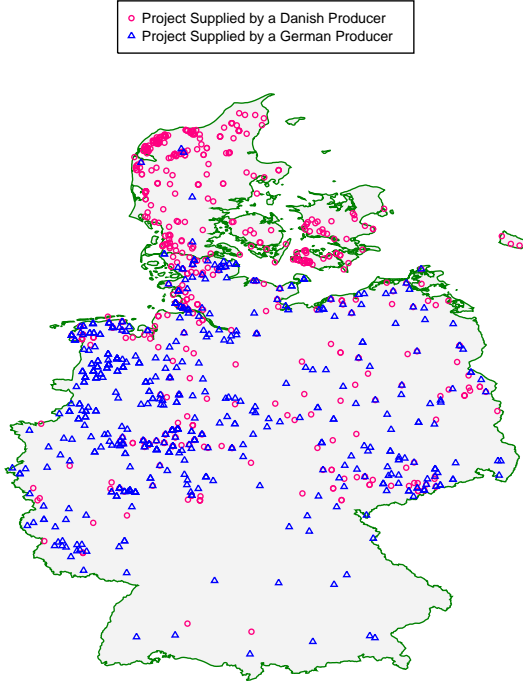


Figure 2: PROJECT LOCATIONS

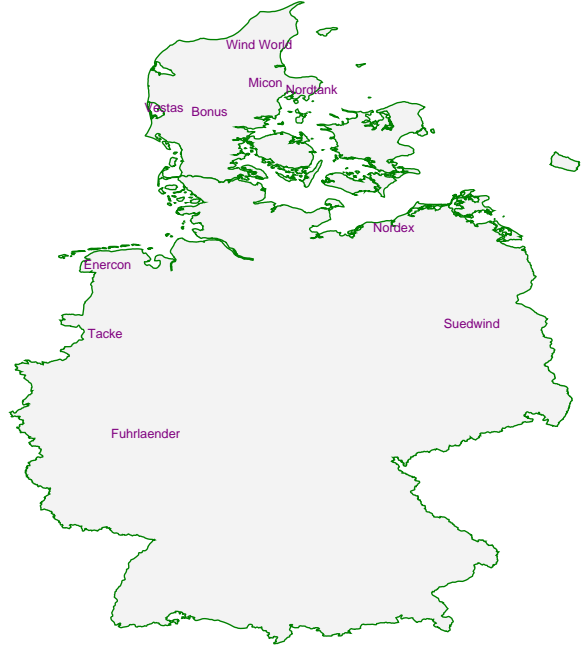



Figure 3: PRODUCER LOCATIONS

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 4. This figure documents Vestas’s declining market share as the distance from its main manufacturing location increases. Whereas Figure 4 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 160 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 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 firm 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 A.4 for details). Figure 5 plots the fit of this regression. The border dummy is a statistically significant -0.284 , which is reflected in the sharp drop in the market share of the largest five Danish firms

Figure 4: MARKET SHARE OF VESTAS BY PROXIMITY TO PRIMARY PRODUCTION FACILITY



Notes: Proportion of projects won by Vestas projected on a cubic polynomial of distance to Vestas's production facility. Dots are aggregated market shares in bins of 75 km width.

Figure 5: MARKET SHARE OF DANISH FIRMS BY PROXIMITY TO THE BORDER



Notes: Regression discontinuity fit of projects supplied by five large Danish firms using a cubic polynomial of distance to border, a Germany dummy and interaction terms. Regression details are in Appendix A.4. Dots are aggregated market shares in bins of 75 km width.

from around 90 to 60 percent at the border.¹⁴

These results give us reason to believe that the border matters in the wind turbine industry, 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.¹⁵ 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.

3 Model

We begin by describing the environment. 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}_ℓ 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.

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

¹⁴These results are robust to considering projects within various bandwidths of the border, as is standard within the RDD framework. For expository purposes, Figure 5 includes projects in the [-300,600] band.

¹⁵ 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 B 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.

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,¹⁶ and an idiosyncratic choice-specific shock ϵ_{ij} .¹⁷ 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.¹⁸ The ϵ_i vector is private information to project managers who collect project-specific price bids from producers. The assumption that ϵ_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.¹⁹ 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,

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

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

¹⁶Since 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.

¹⁷ We assume away project-level economies of scale by making price bids per-megawatt. 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.

¹⁸Project 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.

¹⁹For 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 B, we address the robustness of our estimate to local unobservables of this type.

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:²⁰

$$c_{ij} = \phi_j + \beta_d \cdot \text{distance}_{ij} + \beta_b \cdot \text{border}_{ij}, \quad (2)$$

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 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.

Firms engage in Bertrand competition by submitting price bids for projects in the markets in which they are active.²¹ They observe the identities and all characteristics of their competitors (i.e., their quality and marginal cost for each project) except the valuation vector ϵ_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:²²

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

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

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

²⁰Here, 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.

²¹Industry 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.

²²We 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 address the robustness of our model to the presence of economies of density in Appendix B.

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

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

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 (3) into (1), we get a fixed-point problem with $|\mathcal{J}|$ unknowns and $|\mathcal{J}|$ equations for each project i :

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

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 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.

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:

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

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, (6) 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:

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

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.²³ We now turn to the estimation of the model.

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 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.²⁴

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 border_{ij}^ℓ equals zero if project i and firm j are both located in country ℓ and one otherwise. Substituting the cost function (2) into the winning probability (4), we get

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

From this equation, one can see that firms' production costs ϕ_j and quality level d_j are not separately identified given our data.²⁵ 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 $\theta = (\beta_b, \beta_d, \xi_1, \dots, \xi_{|\mathcal{M}_D| + |\mathcal{M}_G|})$. We estimate the model via constrained maximum likelihood, where the likelihood of the data is maximized subject to the

²³Several papers (e.g., Pakes, Porter, Ho, and Ishii, 2006; Ciliberto and Tamer, 2009) proposed using bounds to construct moment inequalities for use in estimating structural parameters. Holmes (2011) 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.

²⁴Note 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.

²⁵The 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.

equilibrium constraints (8). The likelihood function of the project data has the following form:

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

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:

$$\begin{aligned} & \max_{\theta, \rho} && L(\rho) \\ \text{subject to:} &&& \rho_{ij}^\ell = \frac{\exp\left(\xi_j - \beta_d \cdot \text{distance}_{ij} - \beta_b \cdot \text{border}_{ij}^\ell - \frac{1}{1-\rho_{ij}^\ell}\right)}{1 + \sum_{k=1}^{|\mathcal{J}_\ell|} \exp\left(\xi_k - \beta_d \cdot \text{distance}_{ik} - \beta_b \cdot \text{border}_{ik}^\ell - \frac{1}{1-\rho_{ik}^\ell}\right)} \\ &&& \sum_{k=1}^{|\mathcal{J}_\ell|} \rho_{ik}^\ell + \rho_{i0}^\ell = 1 \end{aligned} \quad (10)$$

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

Our estimation is an implementation of the Mathematical Programming with Equilibrium Constraints (MPEC) procedure proposed by Judd and Su (2011). 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 (10) in order to simplify its Jacobian. In our baseline specification, this is a problem with 12,314 variables (12 structural parameters and 12,302 equilibrium win probabilities for all firms competing for each project) and 12,302 equality constraints. We describe the details of the computational procedure in Appendix C.

Once the structural parameters are recovered, one can calculate bounds on the fixed costs of entry for each firm, f_j , using the 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. Finally, a parametric bootstrap procedure helps to calculate the standard errors for these bounds.²⁶

We also implement several alternative specifications as robustness checks and extensions to our

²⁶To 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.

baseline specification. In our first alternative, we allow distance costs to vary by manufacturing firm:

$$c_{ij} = \phi_j + \beta_{dj} \cdot \text{distance}_{ij} + \beta_b \cdot \text{border}_{ij}. \quad (11)$$

Note that the difference between this and the baseline specification (2) is that distance cost coefficients are heterogeneous (β_{dj} vs. β_d). This cost function is consistent with Holmes and Stevens (2012), who document that in U.S. data large firms tend to ship further away, even when done domestically.²⁷ 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,

$$c_{ij} = \phi_j + \beta_d \cdot \text{distance}_{ij} + \beta_b \cdot \text{border}_{ij} + \beta_s \cdot S_i + \beta_{sb} \cdot \text{border}_{ij} \cdot S_i. \quad (12)$$

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 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.

²⁷They rationalize this observation in a model where heterogeneous firms invest in their distribution networks. Productive firms endogenously face a lower “iceberg transportation cost.”

4.1 Parameter Estimates

Estimation results are presented in Table 2, with the baseline specification reported in the first column. Both variable costs are economically and statistically significant. Based on our estimate, the cost of supplying a foreign project is equivalent to an additional 432 kilometers of distance between the manufacturing location and the project site ($\beta_b/\beta_d = 0.432$). The mean distance from Danish firms to German projects in our data is 623 kilometers. As a consequence, border frictions represent 41 percent of Danish exporters' total delivery costs ($\beta_b/(\beta_b + \beta_d \cdot 6.23) = 41$ 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 ($\beta_b/(\beta_b + \beta_d \cdot 4.2) = 51$ percent).²⁸

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.²⁹ For exporters, the median distance elasticity ranges from .95 to 1.40. 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 1.40 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.

As discussed above, the firm fixed effects reflect the combination of differences in quality and productivity across firms. We find significant differences between them. 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 6 presents this comparison. 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 5, 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

²⁸The magnitudes are very similar if we condition distances on winners of projects (43 and 53 percent, respectively).


²⁹The 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); Eaton and Kortum (2002); Anderson and van Wincoop (2003))

Table 2: MAXIMUM LIKELIHOOD ESTIMATES

	Baseline	Heterogeneous Distance Costs	Project Size Scaling
Border Variable Cost, β_b	0.869 (0.219)	0.867 (0.239)	0.948 (0.228)
Distance Cost (100km), β_d	0.201 (0.032)		0.194 (0.032)
<i>Bonus (DK)</i>		0.169 (0.066)	
<i>Nordtank (DK)</i>		0.277 (0.073)	
<i>Micon (DK)</i>		0.134 (0.051)	
<i>Vestas (DK)</i>		0.287 (0.049)	
<i>WindWorld (DK)</i>		0.016 (0.068)	
<i>Enercon (DE)</i>		0.296 (0.063)	
<i>Fuhrslaender (DE)</i>		1.794 (0.236)	
<i>Nordex (DE)</i>		-0.071 (0.087)	
<i>Suedwind (DE)</i>		-0.231 (0.104)	
<i>Tacke (DE)</i>		0.103 (0.071)	
Project Size, β_s			-0.686 (0.108)
Project Size x Border, β_{sb}			-0.062 (0.053)
Firm Fixed Effects, ξ_j			
<i>Bonus (DK)</i>	2.473 (0.223)	2.332 (0.297)	1.956 (0.231)
<i>Nordtank (DK)</i>	2.526 (0.229)	2.811 (0.326)	2.004 (0.238)
<i>Micon (DK)</i>	3.097 (0.221)	2.786 (0.268)	2.575 (0.230)
<i>Vestas (DK)</i>	3.805 (0.215)	4.180 (0.265)	3.288 (0.224)
<i>WindWorld (DK)</i>	1.735 (0.273)	0.818 (0.418)	1.212 (0.281)
<i>Enercon (DE)</i>	3.533 (0.175)	3.859 (0.270)	3.015 (0.194)
<i>Fuhrslaender (DE)</i>	0.330 (0.263)	3.305 (0.506)	-0.174 (0.269)
<i>Nordex (DE)</i>	1.782 (0.203)	0.683 (0.400)	1.265 (0.216)
<i>Suedwind (DE)</i>	0.537 (0.270)	-1.188 (0.510)	0.024 (0.279)
<i>Tacke (DE)</i>	2.389 (0.177)	2.104 (0.263)	1.878 (0.197)
Log-Likelihood	-2363.00	-2315.82	-2354.50
N	1226	1226	1226

Notes: Standard errors in parentheses.

Figure 6: MODEL FIT: EXPECTED DANISH MARKET SHARE BY DISTANCE TO THE BORDER



Notes: Data line is the same as in Figure 5. The model line is the regression discontinuity fit of probability of winning each project by Danish firms on a cubic polynomial of distance to border, Germany dummy, and interaction terms.

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 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.

The second column of Table 2 contains the estimates of the heterogeneous distance cost specification presented in (11). 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.

Finally, the third column of Table 2 contains estimates from the size-varying per-megawatt cost

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

specification, (12). The coefficient of interest is the interaction term, β_{sb} , which is negative, but 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.

4.2 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.³¹ Since our model naturally allows for heterogeneity in firm operating profits, this section considers whether heterogeneity in firms’ fixed costs of exporting are 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 3.2. Based on the parameter estimates in Table 2, 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 (6), and a lower bound on fixed costs for firms that stay out of the foreign market using (7). 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.³²

Table 3 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

³¹The 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.

³²As mentioned above, since transaction prices are not observed we cannot express profits in currency units (see Footnote 17).

Table 3: EXPORT FIXED COST BOUNDS (f_j)

	Lower	Upper
Bonus (DK)		47.55 (19.52)
Nordtank (DK)		43.29 (8.91)
Micon (DK)		80.13 (13.62)
Vestas (DK)		164.32 (23.60)
WindWorld (DK)		17.35 (3.93)
Enercon (DE)	22.32 (4.87)	
Fuhrlaender (DE)	0.66 (0.32)	
Nordex (DE)		6.33 (1.82)
Suedwind (DE)	1.26 (0.45)	
Tacke (DE)	7.24 (1.71)	

Notes: Scale is normalized by the variance of ϵ (see Footnote 17). Standard errors in parentheses.

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, which would need to form contacts with Danish customers from scratch.³³

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 degree of heterogeneity in fixed costs that is necessary

³³Because 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.

to explain entry patterns.³⁴

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.³⁵ 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.³⁶ 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.

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 4 presents the market share of the major firms of Denmark and Germany in 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 more than 10 percentage points.³⁷ 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

³⁴It is important to note that the variable cost estimates presented in Table 2, 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).

³⁵We implicitly assume that the change in market structure does not induce domestic firms to exit the industry, or new firms to be created.

³⁶We 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.

³⁷For space and clarity, we do not report standard errors of changes in market shares in Table 4. All of the (non-zero) changes in market shares across counterfactuals are statistically significant.

Table 4: COUNTERFACTUAL MARKET SHARES OF LARGE FIRMS (%)

		Data	Baseline Estimates	No Fixed Costs	No Border
Denmark	Danish Firms	92.57	92.65 (1.52)	83.95 (2.26)	74.26 (3.64)
	German Firms	1.69	2.18 (0.60)	11.56 (2.05)	21.94 (3.88)
Germany	Danish Firms	32.37	32.42 (5.42)	32.42 (5.42)	49.32 (7.55)
	German Firms	59.57	59.24 (3.93)	59.24 (3.93)	44.90 (5.80)

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

on the German side of the border when fixed costs of exporting are removed.³⁸ 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 why we see more Danish firms entering Germany than vice versa.³⁹ 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 4 displays the model prediction of national market shares if the border were entirely eliminated. In addition to setting f_j equal for all firms, we also eliminate variable border costs by setting β_b equal to zero.⁴⁰ This results in a large increase in imports on both sides of the border. The domestic market share of Danish firms falls from 92.6 percent to 74.3 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

³⁸Because of this duplication, we simply omit the column which removes fixed cost of entry in Germany in the tables below.

³⁹This argument assumes fixed costs of exporting are of the same order of magnitude for both countries, which appears to be the case.

⁴⁰Because 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.

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 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 7 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 6. 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.⁴¹


5.2 Firm Profits

We now turn to an analysis of winners and losers from border frictions, starting with individual firms. Table 5 presents the level of operating profits under the baseline and two counterfactual scenarios, calculated according to equation (5). While the scale of these profit figures is arbitrary (similar to f_j in Table 3, 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 47.06 in profits in Denmark, and 47.55 in Germany. If the border were removed entirely, Bonus’s profits in Denmark would fall to 34.83, while their profits in Germany would rise to 75.46. 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 losses from 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—lose from the entire elimination of border frictions. Underlying this result is the significant

⁴¹The 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 7: COUNTERFACTUALS: EXPECTED DANISH MARKET SHARE BY DISTANCE TO THE BORDER



Notes: Regression discontinuity fit of projects won by Danish firms on a cubic polynomial of distance to border, Germany dummy, and interaction terms.

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 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.

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 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.⁴² The relative

⁴²Consumer 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 5: BASELINE AND COUNTERFACTUAL PROFIT ESTIMATES

	Denmark			Germany	
	Baseline Estimates	No Fixed Costs	No Border	Baseline Estimates	No Border
Bonus (DK)	47.06 (13.00)	41.02 (11.98)	34.83 (10.71)	47.55 (19.52)	75.46 (28.88)
Nordtank (DK)	44.70 (4.97)	38.98 (4.50)	33.11 (4.19)	43.29 (8.91)	68.72 (13.73)
Micon (DK)	82.76 (7.36)	72.63 (6.80)	62.07 (6.75)	80.13 (13.62)	126.74 (21.14)
Vestas (DK)	156.96 (14.60)	139.46 (12.46)	120.69 (11.83)	164.32 (23.60)	256.08 (37.23)
WindWorld (DK)	20.73 (3.19)	18.13 (2.76)	15.44 (2.49)	17.35 (3.93)	27.59 (6.32)
Enercon (DE)		21.46 (4.54)	42.56 (9.37)	428.91 (48.68)	305.06 (53.60)
Fuhrlaender (DE)		0.57 (0.26)	1.14 (0.56)	17.31 (4.20)	11.98 (3.28)
Nordex (DE)	6.33 (1.82)	5.43 (1.48)	10.79 (2.45)	75.69 (15.15)	51.24 (13.20)
Suedwind (DE)		1.09 (0.37)	2.16 (0.78)	21.74 (5.23)	14.85 (3.90)
Tacke (DE)		6.47 (1.42)	12.93 (3.19)	151.86 (16.60)	104.83 (17.33)

Notes: Scale is normalized by variance of ϵ (see Footnote 17). Standard errors in parentheses.

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.⁴³ 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.

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

⁴³Because 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.

Table 6: COUNTERFACTUAL WELFARE ANALYSIS BY COUNTRY

		Baseline (Levels)	No Fixed Costs (Levels) (% Chg)		No Border (Levels) (% Chg)	
Denmark	(A) Consumer Surplus	70.15 (4.94)	73.46 (4.97)	4.72 (1.03)	77.42 (5.38)	10.36 (2.19)
	(B) Danish Firm Profits	29.33 (0.54)	25.83 (0.74)	-11.92 (2.26)	22.16 (1.26)	-24.44 (4.47)
	(C) German Firm Profits	0.53 (0.15)	2.91 (0.55)	452.99 (122.97)	5.79 (1.13)	999.18 (297.29)
	Domestic Surplus (A+B)	99.47 (5.17)	99.29 (5.11)	-0.18 (0.07)	99.58 (5.09)	0.10 (0.25)
	Total Surplus (A+B+C)	100.00 (5.09)	102.21 (5.07)	2.21 (0.51)	105.37 (5.39)	5.37 (1.28)
Germany	(A) Consumer Surplus	68.99 (6.42)			79.57 (8.30)	15.34 (1.90)
	(B) Danish Firm Profits	10.43 (1.59)			16.41 (2.41)	57.27 (4.96)
	(C) German Firm Profits	20.58 (1.86)			14.44 (2.31)	-29.84 (5.62)
	Domestic Surplus (A+C)	89.57 (5.78)			94.01 (6.68)	4.96 (1.39)
	Total Surplus (A+B+C)	100.00 (6.72)			110.42 (8.59)	10.42 (1.77)

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.

owners. As a result, consumer surplus increases by 4.72 percent. Danish firms, facing harsher domestic competition, see profits decline by 11.92 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 losses 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 6 display the welfare effects of removing border frictions entirely. As we would expect, site owners see significant benefits, and consumer surplus rises by 10.36 percent in Denmark and 15.34 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 24.44 percent in Denmark and 29.84 percent in Germany.⁴⁴ In Denmark, the removal of border frictions results in a transfer of surplus from domestic firms to consumers, netting to essentially

⁴⁴Of course, these declines do not account for benefits realized in the export market. See Table 5 for an accounting of how each firm fairs as both a domestic producer and an exporter under our counterfactual scenarios.

no change in domestic surplus. When we include the benefits of exporters, however, total surplus increases by 5.37 percent. The story in Germany is a bit different. Consumer gains outweigh domestic firm losses in Germany and domestic surplus increases by 4.96 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 10.42 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.

6 Alternative Border Estimates

Many authors (e.g., Engel and Rogers (1996), Parsley and Wei (2001), Broda and Weinstein (2008)) have estimated 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,⁴⁵

$$|p_k^j - p_\ell^j| = \delta_d^j \cdot \text{distance}_{k\ell} + \delta_b^j \cdot \text{border}_{k\ell} + \delta_k^j + \delta_\ell^j + \epsilon_{k\ell}^j, \quad (13)$$

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 included to control for market-specific 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.⁴⁶ The border effect for good j is then interpreted in terms of the distance equivalent of the border dummy coefficient:

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

⁴⁵Authors 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).

⁴⁶As pointed out by Borraz, Cavallo, Rigobon, and Zipitria (2012), the no arbitrage condition implies only an inequality, not an equality, but the majority of the literature implements the model as an equality.

Table 7: BORDER EFFECT ESTIMATES FROM REDUCED-FORM ESTIMATION

Firm	$\hat{\delta}_b/\hat{\delta}_d$ in km
Bonus	741.98
Nordtank	857.95
Micon	781.16
Vestas	516.61
WindWorld	1092.08
Nordex	3472.55

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 (13) for reasons we discuss below. In order to demonstrate how these approaches differ, this section uses our model-generated prices to estimate (13) and compares the implied border widths.

In our thought experiment, we parameterize our model using the estimates found in Section 4.2 to calculate price bid differentials on the left hand side of (13). An econometrician with data on price bids of all manufacturers indexed by j tries to recover the width using (13) by regressing 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 7 reports the implied border width ($\hat{\delta}_b^j/\hat{\delta}_d^j$) from the OLS regression of (13) for each manufacturer. Recall that our baseline model implies border width of 432 kilometers for all firms.⁴⁷ Evidently, this exercise overestimates the border effect in our model for all producers.⁴⁸ For Danish firms, estimates vary between 1.2 to 2.5 times the true value. The bias is much higher for the German firm, Nordex.

To gain intuition on the sources of this overestimation, contrast (13) with the price difference implied

⁴⁷As 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.

⁴⁸The 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.

by our structural model using our estimates $(\hat{\beta}_d, \hat{\beta}_b, \hat{\rho}_{kj})$ in expressions (2) and (3):

$$|p_k^j - p_\ell^j| = \left| \hat{\beta}_d(\text{distance}_{kj} - \text{distance}_{\ell j}) + \hat{\beta}_b(\text{border}_{kj} - \text{border}_{\ell j}) + \left(\frac{1}{1 - \hat{\rho}_{kj}} - \frac{1}{1 - \hat{\rho}_{\ell j}} \right) \right| \quad (14)$$

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 (13), it is apparent that the linear reduced-form regression is misspecified. In the structural equation (14), price differentials are generated by the absolute value of several *differences* in project-to-producer distances, destination countries, and mark-ups, whereas (13) is a linear function of related, but different, variables. While trying to emulate this model-based expression, equation (13) suffers from two additional problems: First, using project-to-project distances ($\text{distance}_{k\ell}$) instead of the differences in project-to-producer distances ($\text{distance}_{kj} - \text{distance}_{\ell j}$) 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 (14).⁴⁹ Second, (13) 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 7 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

⁴⁹The triangle inequality discrepancy explains why the measured effect is so much higher for Nordex in Table 7. 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.

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 across national boundaries.

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 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. 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.

References

- ANDERSON, J. E., AND E. VAN WINCOOP (2003): “Gravity with Gravititas: 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.
- BALISTRERI, E. J., AND R. H. HILLBERRY (2007): “Structural estimation and the border puzzle,” *Journal of International Economics*, 72(2), 451–463.
- BENKARD, C. L. (2004): “A Dynamic Analysis of the Market for Wide-Bodied Commercial Aircraft,” *Review of Economic Studies*, 71, 581–611.
- BESANKO, D., U. DORASZELSKI, Y. KRYUKOV, AND M. SATTERTHWAIT (2010): “Learning-by-Doing, Organizational Forgetting, and Industry Dynamics,” *Econometrica*, 78(2), 453–508.
- 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.
- CAPLIN, A., AND B. NALEBUFF (1991): “Aggregation and Imperfect Competition: On the Existence of Equilibrium,” *Econometrica*, 59(1), 25–59.
- CILIBERTO, F., AND E. TAMER (2009): “Market Structure and Multiple Equilibria in Airline Markets,” *Econometrica*, 77(6), 1791–1828.
- 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.
- 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 Multi-dimensional Product Attributes,” *Review of Economic Studies*, 62(1), 19–52.
- 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.
- GOLDBERG, P. K., AND F. VERBOVEN (2001): “The Evolution of Price Dispersion in the European Car Market,” *The Review of Economic Studies*, 68(4), 811–848.
- 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. HSIEH, AND N. LI (2011): “International Prices, Costs and Markup Differences,” *American Economic Review*, 101(6), 2450–86.
- 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.
- 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 the Economies of Density,” *Econometrica*, 79(1), 253–302.
- HOLMES, T. J., AND J. J. STEVENS (2012): “Exports, Borders, Distance, and Plant Size,” *Journal of International Economics*, *forthcoming*.
- INTERESSENVERBAND WINDKRAFT BINNENLAND (various years): “Windkraftanlagen Marktübersicht,” 1994 through 1996.
- JUDD, K. L., AND C.-L. SU (2011): “Constrained Optimization Approaches to Estimation of Structural Models,” *Econometrica*, *forthcoming*.
- 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.
- 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 variability, 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.
- 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.
- 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,” Energy Resource Center of the Netherlands, ECN-C-02-083.
- THOMADSEN, R. (2005): “The Effect of Ownership Structure on Prices in Geographically Differentiated Industries,” *The RAND Journal of Economics*, 36(4), pp. 908–929.

Appendices

Appendix A Data

A.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.

A.2 Project Characteristics

Table 8, and Figures 8-10 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>).

A.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 CRTS assumption. Figure 11 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 10 shows, a majority of the turbines installed in this period were in the 400-600 KW range.


Table 8: 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	930
	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.

Figure 8: KW CAPACITY HISTOGRAMS BY MARKET



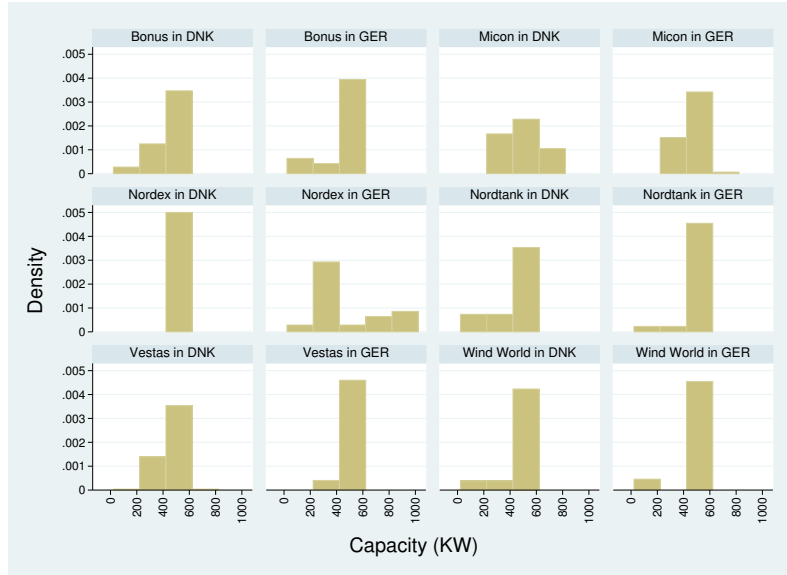
Notes: An observation is average kw capacity of turbines in a project. Years 1995 and 1996 only.

Figure 9: 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 10: 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 11: PER KW LIST PRICES OF VARIOUS TURBINES OFFERED IN 1995-1996



Notes: Pooled over all producers.

A.4 Regression Discontinuity Design

We estimate the following local linear probability model in Subsection 2.2 to implement the regression discontinuity design:

$$y_i = \alpha_0 + \sum_{k=1}^{k=3} \alpha_k \cdot \text{distance}_i^k + \gamma \cdot \text{Germany}_i + \sum_{k=1}^{k=3} \eta_k \cdot \text{distance}_i^k \cdot \text{Germany}_i + \epsilon_i. \quad (15)$$

The dependent variable is $y_i = 1$ if the producer of project i is one of the five large Danish firms in our model (Bonus, Micon, Nordtank, WindWorld or Vestas), and zero otherwise. 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 . The band we use for distance is [-300km, 600km] which preserves 1201 out of 1226 observations for the years 1995-1996. Table 9 reports the results for various specifications. 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. Figure 5 is the fit of this specification. The border coefficient γ is significantly negative and of comparable magnitude in all four regressions.

Table 9: RDD RESULTS FOR THE 1995-1996 PERIOD

	Baseline Specification	Cubic No interactions	Linear	Linear No interaction
Germany (γ)	-0.284* (-2.06)	-0.311*** (-4.27)	-0.406*** (-5.96)	-0.416*** (-8.49)
Constant (α_0)	0.922*** (7.52)	0.807*** (16.01)	0.848*** (13.88)	0.860*** (30.13)
Distance				
α_1	0.001 (0.44)	-8.08e-04 (-3.90)	4.83e-04 (-1.36)	4.093e-04*** (-4.56)
α_2	1.32e-06 (0.47)	-3.10e-08 (-0.05)		
α_3	2.41e-08 (0.36)	1.49e-09 (1.32)		
Interactions				
η_1	-0.004 (-1.31)		7.93e-05 (0.22)	
η_2	-3.91e-06 (-0.14)			
η_3	-3.27e-08 (-0.49)			
Observations	1201	1201	1201	1201
Adjusted R^2	0.282	0.278	0.276	0.276

Notes: t statistics in parentheses

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

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 10 reports the border dummy and the constant for all three periods. In all cases, Germany dummy is negative and statistically significant at the 5% level. Moreover, the border effect is very stable over time.

Table 10: RDD RESULTS FOR SUBSEQUENT PERIODS

	(1)	(2)	(3)
	1995-1996	1997-1998	1999-2005
Germany	-0.284* (-2.06)	-0.253* (-2.24)	-0.318** (-2.67)
Constant	0.922*** (7.52)	0.911*** (11.45)	0.862*** (8.00)
Observations	1201	1237	3318
Adjusted R^2	0.282	0.380	0.192

Notes: t statistics in parentheses

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

Appendix B Robustness to Local Unobservables, Economies of Density, and Spatial Collusion

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,

$$\epsilon_j = \gamma + \psi W \epsilon_j + \nu_i. \quad (16)$$

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 through a road network), inclusion within the same region, or nearest neighbor adjacency. In practice we specify W as,

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

where distance is the direct distance (as the crow flies) in kilometers between projects i and j .⁵⁰

We are unable to directly test for spatial autocorrelation in ϵ_{ij}^ℓ because as with all discrete choice models, ϵ_{ij}^ℓ 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_{ij}^ℓ 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,⁵¹

$$\hat{\epsilon}_{ij}^\ell = \begin{cases} \gamma - \log \rho_{ij}^\ell & \text{if } y_{ij}^\ell = 1, \\ \gamma + \frac{\rho_{ij}^\ell}{1 - \rho_{ij}^\ell} \log \rho_{ij}^\ell & \text{if } y_{ij}^\ell = 0. \end{cases}$$

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_{ij}^ℓ must be in order for selection to occur. Hence, $\hat{\epsilon}_{ij}^\ell$ is decreasing in the ex-ante probability of firm j being selected. The fact that the distribution of $\hat{\epsilon}_{ij}^\ell$ 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 auto-correlation is violated, $\hat{\epsilon}_{ij}^\ell$ 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 11.⁵² While the magnitude of the estimated $\hat{\psi}$ is small, the test strongly rejects the null hypothesis for every firm, due in part to 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

⁵⁰Our results are robust to raising or lowering the distance cutoff and using a specification of W based on inverse distance.

⁵¹The derivation is available from the authors upon request.

⁵²It 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 11: RESULTS FROM AUTO-CORRELATION TESTS

Manufacturer	$\hat{\psi}$	Std. Error	t-Stat.
Fringe	0.024	0.008	3.096
Bonus (DK)	0.027	0.005	5.079
Nordtank (DK)	0.024	0.004	6.122
Micon (DK)	0.032	0.004	7.225
Vestas (DK)	0.034	0.005	7.124
WindWorld (DK)	0.031	0.007	4.635
Enercon (DE)	0.043	0.007	6.000
Fuhrlander (DE)	0.034	0.005	7.165
Nordex (DE)	0.048	0.006	8.194
Suedwind (DE)	0.038	0.014	2.757
Tacke (DE)	0.029	0.005	6.118

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,⁵³

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

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 12. The coefficient on having a nearby installation has the expected negative sign (nearby installations are indicative of lower 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.

⁵³We 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 12: ROBUSTNESS CHECK: NEARBY INSTALLED TURBINES

	Coefficient	Std. Error
Border Variable Cost, β_b	0.688	(0.178)
Distance Cost (100km), β_d	0.137	(0.031)
Nearby Installation, β_c	-1.249	(1.167)
Firm Fixed Effects, ξ_j		
<i>Bonus (DK)</i>	1.256	(0.189)
<i>Nordtank (DK)</i>	1.462	(0.183)
<i>Micon (DK)</i>	2.046	(0.160)
<i>Vestas (DK)</i>	2.689	(0.170)
<i>WindWorld (DK)</i>	0.640	(0.211)
<i>Enercon (DE)</i>	2.719	(0.147)
<i>Fuhrlaender (DE)</i>	-0.010	(0.266)
<i>Nordex (DE)</i>	0.858	(0.184)
<i>Suedwind (DE)</i>	-0.187	(0.206)
<i>Tacke (DE)</i>	1.578	(0.152)
Log-Likelihood	-2286.15	
N	1226	

Appendix C Computational Method

C.1 Estimation of the Project Bidding Game

We formulate the estimation of the project bidding game as a constrained optimization problem.⁵⁴ 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 (10) for the computational implementation. The reformulated constraints are mathematically equivalent to those in (10). 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

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

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 930 German and 296 Danish project this aggregates to 12,302 constraints. In our baseline specification we are

⁵⁴See Judd and Su (2011) for a seminal paper that explains why constrained optimization of structural models is often superior to estimation via nested fixed points.

choosing 12,314 variables (12 structural parameters and 12,302 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.⁵⁵ 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.

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 $\hat{\theta}$ 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:

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

C.2 Counterfactuals

The point estimate $\hat{\theta}$ 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 $\hat{\theta}$ 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 $\hat{S}(\hat{\theta})^{-1}$). 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.

⁵⁵Prior to the estimation we check via finite differences that our analytical gradients are correct.