1 Introduction

2 Industry Background and Data

2.1 Data

2.2 Preliminary Analysis of the Border Effect

3 Model

3.1 Project Bidding Game

3.2 Entry Game

4 Estimation

4.1 Parameter Estimates

4.2 Fixed Cost Bounds

5 Border Frictions, Market Segmentation, and Welfare

5.1 Market Shares and Segmentation

6 Conclusion

Appendices

FOR ONLINE PUBLICATION (Cosar,Grieco,Tintelnot 2014)

Appendix A Gravity in Wind Turbine Trade

In order to get a rough comparison of the relevance of trade costs within the wind turbine industry versus common benchmarks in the literature, we estimate a gravity equation using the 6-digit HS 2007 product category that is associated with the industry. The precise goal is to compare distance and contiguity coefficients to the values obtained in the literature using aggregate data.

Gravity variables come from the CEPII dataset (`http://www.cepii.fr/CEPII/en/bdd_modele/bdd.asp`) made available by Head and Mayer (2013). USITC (2009) helps us to identify the product code associated with wind turbines: "wind-powered generating sets" with the HS 2007 code 8502.31. We obtain bilateral trade data on this product from WITS database (`http://wits.worldbank.org/wits/`). Data is available for the period 2002-2010. The estimation equation takes the form

$$\ln X_{sd} \ = \ \psi_s + \psi_d + \alpha \cdot Contig_{sd} + \beta \cdot \ln(distance_{sd}) + \Gamma \cdot Z_{sd} + \epsilon_{sd}, \tag{1}$$

where the dependent variable is the natural logarithm of trade volume X_{sd} between source country s and destination country d averaged over 2002-2010. (ψ_s, ψ_d) are importer-exporter fixed effects. The variable $Contig_{sd}$ equals one if the two countries are contiguous. Z_{sd} includes a set of standard controls such as common language, common currency, bilateral tariffs, regional or bilateral free trade agreement, and colonial links. We estimate this equation with OLS using data on country pairs with positive trade flows $X_{sd} > 0$. Table 1 reports the results.

Table 1: Gravity of Wind Turbine Trade

Contiguity	0.585*
	(0.306)
Distance	-1.027***
	(0.199)
Country fixed effects	Yes
Observations	1366
R^2	0.594

Notes: Standard errors in parenthesis
* significance at 10 percent level
*** significance at 1 percent level.

The results indicate that the industry is remarkably representative in terms of distance and contiguity. The elasticity of trade flows with respect to distance is -1.027, which is consistent with the typical unit elasticity reported by the literature for aggregate trade flows. The contiguity coefficient is 0.585. In a survey of 159 papers from the gravity literature, Head and Mayer (2013) report summary statistics on the coefficients of most frequently used variables. The mean distance elasticity and contiguity coefficient across structural gravity estimates are -1.1 and 0.66, respectively (Table 4 in their paper). This gives us some assurance that the results of the paper are not specific to an industry that is itself an outlier in terms the effects of distance and contiguity.

Appendix B Data

B.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/`) and spans the period 1982-2005. Before 1987, however, both countries have low levels of annual installations: in Germany, there are only 48 wind farms in operation as of 1987, whereas after this year, there are at least 50 new projects annually.

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, 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.

B.2 Project Characteristics

Table 2, and Figures 1-3 provide some summary statistics on project characteristics in the two countries. Differences in distance to producers reflect heterogeneity in country size. Evidently, key observable characteristics such as electricity generating capacity, tower height and rotor diameter are remarkably similar in the two markets, ruling out product differentiation as a source of market segmentation. Slightly higher tower heights in Germany are due to lower wind speeds in southern regions. In such an environment, larger turbines are needed to attain the same capacity. What matters for this paper is that wind conditions do not change at the border. The European wind atlas available at the following link verifies that this is the case. (`http://www.wind-energy-the-facts.org/en/appendix/appendix-a.html`).

Table 2: 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	1977-1981	76	0
	1982-1987	362	48
	1988-1994	1030	1452
	1995-1996	296	929
	1997-2005	1373	4148

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

B.3 List Prices

The survey of the German wind turbine market published by Interessenverband Windkraft Binnenland (various years) provides information on list prices for various turbine models as advertised by producers. These prices, however, are only suggestive and do not reflect project-specific final transaction prices. We use this information to verify the validity of our constant-returns-to-scale assumption. Figure 4 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 3 shows, a majority of the turbines installed in this period were in the 400-600 KW range.

B.4 Regression Discontinuity Design

We estimate the following linear probability model in Subsection 2.2:

$$y_i = \alpha_0 + \sum_{k=1}^{k=3} \alpha_k \cdot \text{distance}_i^k + \gamma \cdot \text{Germany}_i + \sum_{k=1}^{k=3} \eta_k \cdot \text{distance}_i^k \cdot \text{Germany}_i + \epsilon_i. \quad (2)$$

The dependent variable is $y_i = 1$ if the producer of project i is one of the five large Danish firms, and zero otherwise. The variable distance_i is the distance to the border. The effect of the border is picked up by the dummy variable Germany_i that equals one if the project is in Germany, and zero otherwise. The parameter of interest is γ. Table 3 reports the results for various specifications estimated with robust standard errors. The first column is the baseline featuring a cubic polynomial and interaction terms which allow distance to have a different effect on the two sides of the border. The border coefficient γ is significantly negative and of comparable magnitude in all four regressions.

Table 3: RDD Results for the 1995-1996 Period

	Baseline Specification	Cubic No interactions	Linear	Linear No interaction
Germany (γ)	-0.305*	-0.338***	-0.411***	-0.423***
	(0.126)	(0.07)	(0.066)	(0.047)
Constant (α_0)	0.925***	0.807***	0.851***	0.862***
	(0.112)	(0.049)	(0.059)	(0.027)
Distance				
α_1	0.0014	-6.7e-04***	-4.77e-04**	-3.91e-04***
	(0.0026)	(1.84e-04)	(3.41e-04)	(8.32e-05)
α_2	1.17e-05	3.37e-07		
	(1.8e-04)	(5.24e-07)		
α_3	2.04e-08	2.55e-10		
	(3.61e-08)	(7.17e-10)		
Interactions				
η_1	-0.004		-8.92e-05	
	(0.0027)		(3.52e-04)	
η_2	-4.94e-06			
	(1.81e-05)			
η_3	-2.59e-08			
	(3.61e-08)			
Observations	1226	1226	1226	1226
Adjusted R^2	0.284	0.279	0.278	0.2718

Notes: Standard errors in parentheses. *,**,***: significance at 10, 5, 1 percent levels.

Appendix C Additonal Results and Robustness Checks

C.1 Firm Profits

Table 4 presents the level of operating profits under the baseline and two counterfactual scenarios, calculated according to (5). While the scale of these profit figures is arbitrary (similar to f_j in Table 4, 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 48.77 in profits in Denmark, and 45.66 in Germany. If the national border were reduced to a state border, Bonus's profits in Denmark would fall to 37.66, while their profits in Germany would rise to 61.47. On overall, Bonus would

Table 4: Baseline and Counterfactual Profit Estimates

	Denmark			Germany	
	Estimates	No Fixed Costs	No National Border Costs	Estimates	No National Border Costs
Bonus (DK)	48.77	41.61	37.66	45.66	61.47
	(5.23)	(5.25)	(4.52)	(5.65)	(9.96)
Nordtank (DK)	42.70	36.41	32.94	43.56	58.72
	(4.49)	(4.54)	(3.84)	(5.28)	(9.74)
Micon (DK)	82.87	71.29	64.81	77.88	104.75
	(7.32)	(7.71)	(6.39)	(8.08)	(16.29)
Vestas (DK)	160.77	140.50	128.86	156.12	208.22
	(11.40)	(12.46)	(10.74)	(13.84)	(27.77)
WindWorld (DK)	21.57	18.53	16.84	16.74	22.60
	(3.58)	(3.26)	(2.95)	(3.04)	(4.74)
Enercon (DE)		24.04	36.46	474.18	398.38
		(7.85)	(6.05)	(33.46)	(48.68)
Fuhrlaender (DE)		0.77	1.18	15.42	12.43
		(0.44)	(0.55)	(5.10)	(4.20)
Nordex (DE)	7.34	6.14	9.32	78.47	63.08
	(3.13)	(2.14)	(1.79)	(9.25)	(9.68)
Suedwind (DE)		1.44	2.19	21.99	17.61
		(0.62)	(0.58)	(4.76)	(4.53)
Tacke (DE)		7.73	11.78	153.84	124.94
		(2.60)	(2.15)	(13.57)	(17.23)

Notes: Scale is normalized by variance of ϵ (see Footnote 8). Standard errors in parentheses.

see its total profits increase as a result of the elimination of national border frictions, as gains in Germany would more than offset loses 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—loose from the entire elimination of national border frictions. Underlying this result is the significant asymmetry in size and productivity between Germany and Denmark. The losses German firms face due to increased competition in the larger German market overwhelm all gains they receive from better access to the Danish market. Our model estimates Danish firms to be highly productive, so eliminating the national border is quite costly to German incumbents. Even a small Danish exporter like WindWorld gains from the reduction of national border frictions since increased profits in the larger German 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.

C.2 Alternative Cost Specifications

We implement several alternative specifications as robustness checks and extensions to our baseline cost specification. First, we estimate the cost function of the firm without the state border. In our second alternative, we allow distance costs to vary by manufacturing firm:

$$c_{ij} = \phi_j + \beta_{dj} \cdot \log(\text{distance}_{ij}) + \beta_b \cdot \text{border}_{ij} + \beta_s \cdot \text{state}_{ij}. \tag{3}$$

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.[1] 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 third alternative specification, we allow the per-megawatt cost of a project and the impact of national boundaries to vary by project size,

$$c_{ij} = \phi_j + \beta_d \cdot \log(\text{distance}_{ij}) + \beta_b \cdot \text{border}_{ij} + \beta_s \cdot \text{state}_{ij} + \gamma_1 \cdot S_i + \gamma_2 \cdot \text{border}_{ij} \cdot S_i. \qquad (4)$$

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 γ_2 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 be if costs are connected to delivery or legal liability associated with the value of cross-border contracts, then γ_2 will be small in magnitude and border costs will remain important even for large projects. The size coefficient, γ_1, 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.

The left-hand panel of Table 5 contains the estimates of the heterogeneous distance cost specification presented in (3). The border coefficients remain strongly significant, indicating that they are not an artifact of heterogeneity in distance costs. Turning to the distance costs themselves, small firms do not have systematically higher distance costs. Two small firms in our data, Suedwind and Nordex, are estimated to be distance loving, as they built several turbines in locations further away from their plants. While a formal likelihood ratio test rejects the null hypothesis of homogeneous distance costs, the estimation results indicate that heterogeneous distance costs are not driving cross-border differences in this industry. Therefore, we use our homogeneous distance cost specification for the counterfactuals in the following section.

The last column of Table 5 contains estimates from the size-varying per-megawatt cost specification, (4). The coefficient of interest is the interaction term, γ_2, which is negative, but neither economically nor statistically significant. (The average project size is 1 megawatt.) This is evidence that the variable national 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, γ_1, 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 effects 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 national border variable cost scales with project size.

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

Table 5: Alternative Specifications

	Heterogeneous Distance Costs		Economies of Scale
National Border Variable Cost, β_b	0.938		1.246
	(0.285)		(0.253)
State Border Variable Cost, β_s	0.683		0.650
	(0.240)		(0.224)
Log Distance Cost, β_d			0.535
			(0.092)
Project Size, γ_1			-0.723
			(0.108)
Project Size × Border, γ_2			-0.075
			(0.054)
Firm specific coefficients			
	Fixed Effects, ξ_j	Distance Costs, β_{dj}	Fixed Effects, ξ_j
Bonus (DK)	2.305	0.479	1.951
	(0.280)	(0.220)	(0.226)
Nordtank (DK)	3.051	1.040	1.998
	(0.342)	(0.272)	(0.233)
Micon (DK)	3.138	0.680	2.553
	(0.263)	(0.188)	(0.219)
Vestas (DK)	4.477	1.189	3.243
	(0.278)	(0.196)	(0.216)
WindWorld (DK)	1.215	0.271	1.111
	(0.348)	(0.188)	(0.262)
Enercon (DE)	3.823	0.490	3.340
	(0.243)	(0.177)	(0.223)
Fuhrlaender (DE)	0.963	1.863	0.099
	(0.403)	(0.339)	(0.329)
Nordex (DE)	0.988	-0.437	1.684
	(0.355)	(0.232)	(0.245)
Suedwind (DE)	0.226	-0.149	0.519
	(0.519)	(0.305)	(0.310)
Tacke (DE)	2.401	0.131	2.238
	(0.259)	(0.180)	(0.228)
Log-Likelihood	-2290.86		-2324.05
N	1225		1225

Notes: Standard errors in parentheses.

C.3 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^{ℓ}_{ij}, 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 firms' projects to be more tightly clustered together than our model would predict, leading to spatial autocorrelation in firms' error terms across projects. To test for the presence of spatial autocorrelation, we consider the following parametric model for the error term:

$$\epsilon_j = \gamma + \psi W \epsilon_j + \nu_i. \tag{5}$$

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

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

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

where distance is the direct distance (as the crow flies) in kilometers between projects i and j.[2]

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

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

function,[3]

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

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 misspecified. Nonetheless, they are useful to conduct a hypothesis test for $\psi = 0$.

Table 6: Results from Auto-Correlation Tests

Manufacturer	$\hat{\psi}$	Std. Error	t-Stat.
Fringe	0.026	0.008	3.400
Bonus (DK)	0.028	0.006	4.932
Nordtank (DK)	0.024	0.004	6.177
Micon (DK)	0.030	0.005	6.544
Vestas (DK)	0.033	0.005	6.806
WindWorld (DK)	0.029	0.007	4.203
Enercon (DE)	0.048	0.007	6.651
Fuhrlaender (DE)	0.035	0.006	5.847
Nordex (DE)	0.045	0.010	4.393
Suedwind (DE)	0.042	0.014	2.898
Tacke (DE)	0.033	0.005	6.879

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 6.[4] 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 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

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

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

capturing differences in the installed bases of foreign and domestic firms in addition to the variable cost of exporting. Alternatively, if serial correlation is due to local unobserved characteristics then the location of previous installations, while not cost reducing in and of themselves, serve as proxies for unobservable local conditions. In this spirit, we propose the following specification to check the robustness of our results to mild spatial autocorrelation. We re-estimate the model with the augmented cost function,

$$c_{ij} = \phi_j + \beta_d \cdot \log(\text{distance}_{ij}) + \beta_b \cdot \text{border}_{ij} + \beta_s \cdot \text{state}_{ij} + \beta_c \cdot \text{installed}_{ij},$$

where,[5]

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

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 7. The coefficient on having a nearby installation has the expected negative sign (nearby installations are indicative of lower costs) and is of substantial magnitude. 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 actually increases to being equivalent to a 9.8-fold (exp(0.92/0.4)) increase in distance (from an 8-fold increase (exp(1.151/0.551)) in column 2 of Table 3). 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.

[5] We also experimented with including in the cost function the distance to the nearest installed project and using only projects installed in 1993-1994, and obtained qualitatively similar results.

Table 7: Robustness Check: Nearby Installed Turbines

	Coefficient	Std. Error
National Border Variable Cost, β_b	0.917	(0.251)
State Border Variable Cost, β_s	0.633	(0.228)
Log Distance Cost, β_d	0.401	(0.092)
Nearby Installation, β_c	-1.199	(0.107)
Firm Fixed Effects, ξ_j		
Bonus (DK)	1.357	(0.238)
Nordtank (DK)	1.562	(0.243)
Micon (DK)	2.118	(0.226)
Vestas (DK)	2.757	(0.227)
WindWorld (DK)	0.675	(0.266)
Enercon (DE)	3.142	(0.220)
Fuhrlaender (DE)	0.311	(0.335)
Nordex (DE)	1.389	(0.255)
Suedwind (DE)	0.398	(0.309)
Tacke (DE)	2.026	(0.225)
Log-Likelihood	-2263.27	
N	1225	

Appendix D Computational Method

D.1 Estimation of the Project Bidding Game

We formulate the estimation of the project bidding game as a constrained optimization problem.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

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

For the baseline estimation, there are 11 constraints for every German project, and 7 constraints for every Danish project ($|\mathcal{J}_G| = 10$ and $|\mathcal{J}_D| = 6$ plus one fringe firm in every market). Since we have 929 German and 296 Danish projects, this aggregates to 12,291 constraints. In our baseline specification we are choosing 12,304 variables (13 structural parameters and 12,291 equilibrium win probabilities for each firm in each market)

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

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

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

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

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

D.2 Counterfactuals

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

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

Figure 1: KW Capacity Histograms by Market

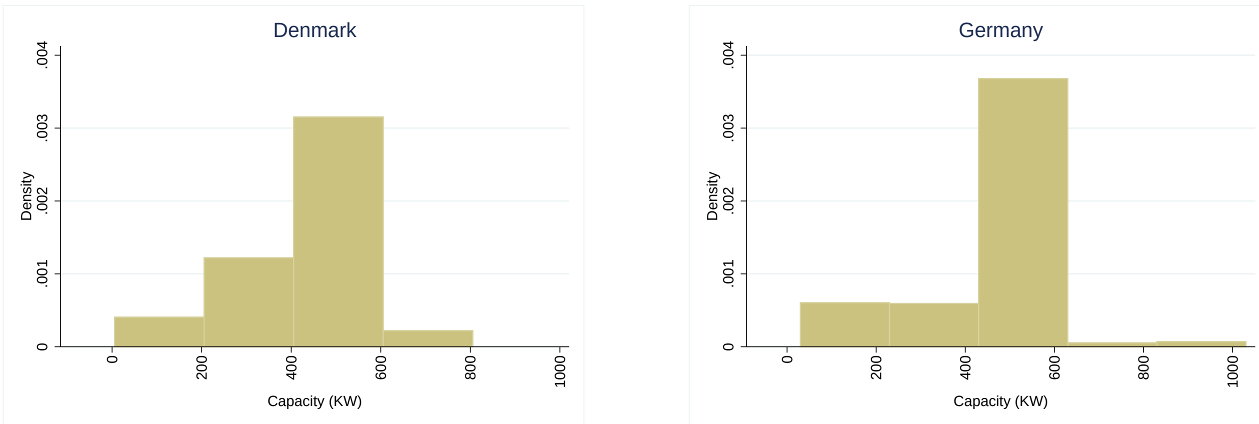

Notes: An observation is average kw capacity of turbines in a project. Years 1995 and 1996 only.

Figure 2: Tower Height Histograms by Producer and Market

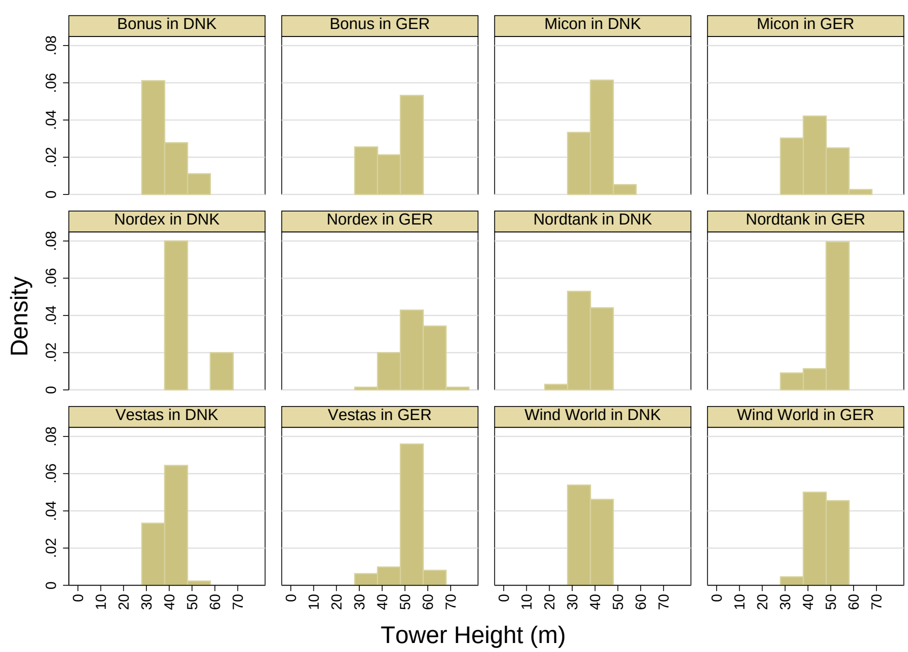

Notes: An observation is average tower height of turbines in a project. Years 1995 and 1996 only. "Bonus in DNK (GER)" indicates projects supplied by Bonus in Denmark (Germany).

Figure 3: KW Capacity Histograms by Producer and Market

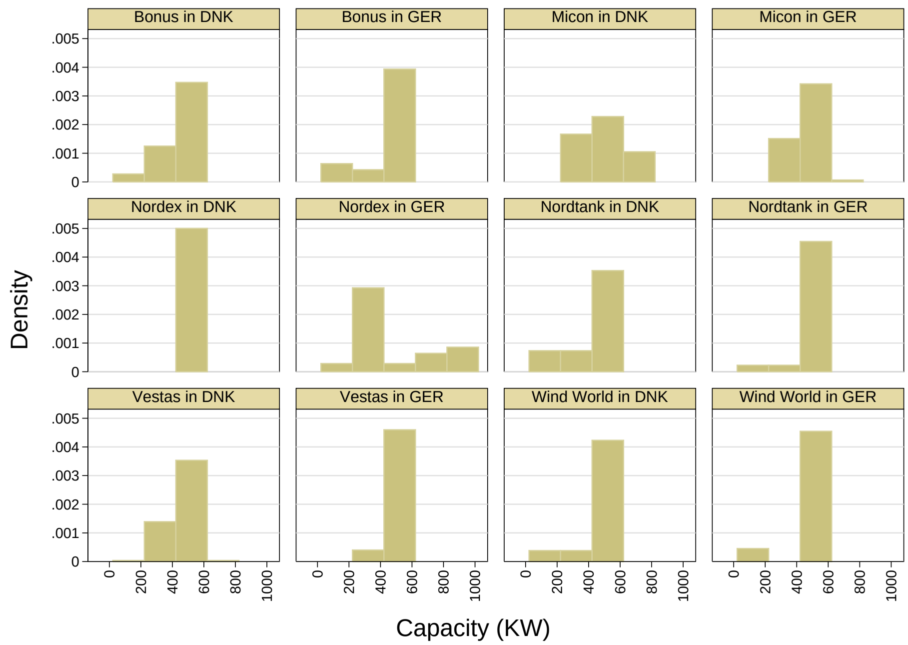

Notes: An observation is average kw capacity of turbines in a project. Years 1995 and 1996 only. "Bonus in DNK (GER)" indicates projects supplied by Bonus in Denmark (Germany).

Figure 4: Per KW List Prices of Various Turbines Offered in 1995-1996

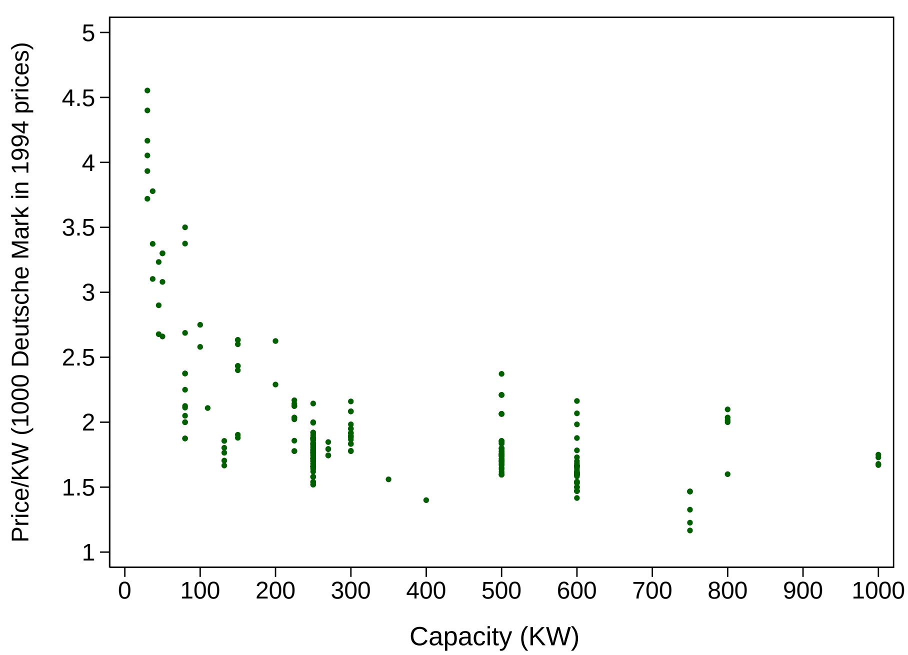

Notes: Pooled over all producers.

References

Head, K., and T. Mayer (2013): “Gravity Equations: Toolkit, Cookbook, Workhorse,” in *Handbook of International Economics*, vol. IV. Elsevier.

Holmes, T. J., and J. J. Stevens (2012): “Exports, borders, distance, and plant size,” *Journal of International Economics*, 88(1), 91–103.

Interessenverband Windkraft Binnenland (various years): “Windkraftanlagen Marktübersicht,” 1994 through 1996.

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.

USITC (2009): “Wind Turbines: Industry and Trade Summary,” United States International Trade Commission Publication ITS-02.