Winds of Change: Technical Progress and Learning by Doing in the Wind Industry *

Thomas R. Covert †
Richard L. Sweeney ‡

March 2026

Abstract

We measure technical progress and cost declines in the wind turbine industry, accounting for large quality improvements that have occurred over time. As in many industrial settings, we do not observe costs. Instead, we infer them by modeling the turbine procurement process, and estimate the model on the near-universe of global commercial wind farms built from 2000 to 2019. To account for dynamic pricing incentives, we introduce a new method for controlling for dynamics in learning-by-doing settings without computing the full equilibrium of a dynamic game. We find that while commonly reported measures of price per unit of capacity fell by 30 percent during this time, quality-adjusted costs fell by over 75 percent. A doubling of cumulative manufacturing experience reduces manufacturing costs by 24 to 33 percent. Though inter-firm spillovers are relatively small, in aggregate, they are responsible for significant cost reductions over time.

Keywords: Innovation, Renewable Energy, Learning

*Both authors declare they have no interests, financial or otherwise, that relate to the research described in this paper, nor do they have any current ties, directly or indirectly to the energy industry. Financial support for data acquisition from Boston College, the University of Chicago's Becker Friedman Institute, and the NBER-Sloan conference on the Economics of Innovation in the Energy Sector is gratefully acknowledged. We also thank Jackson Dorsey, Gautam Gowrisankaran, Ashley Langer and Jing Li for helpful comments. Devin McNulty, Grace Park and Miriam Goldgeil provided excellent research assistance.

†Airbnb, thom.covert@gmail.com.
‡Boston College, sweeneri@bc.edu.

1 Introduction

In order to avoid severe damages from climate change, society needs to transition away from fossil fuels towards clean energy as fast as possible. The reason that the current share of clean energy is so low is that, historically, these technologies have been (privately) expensive. This situation can improve over time, through two channels. First, firms can learn to make existing clean technologies more cheaply. Second, firms can introduce better (more productive) clean technologies. Measuring the extent to which these channels have driven cost reductions is important for understanding the pace of technological progress in clean energy to date, evaluating the impact of public policy, and making forecasts about the energy transition.

Given the marked increase in adoption over the past two decades, it is clear that clean energy has become cheaper during this time. However, measuring the extent to which this reflects genuine improvements in manufacturing technology and reductions costs is challenging for two reasons. First, commonly reported measures of clean energy improvement are typically based on the observed *prices* of these technologies. But price can differ from costs due to firms' strategic incentives. Second, even if costs were observed, measures of capital cost alone would not reflect the dramatic improvements in "quality" that have occurred over time. Simply comparing, for example, the cost to make electric vehicle now to the cost of making an electric vehicle in 2010 would neglect the fact that today's EVs are dramatically better cars than earlier models.

In this paper, we address both these challenges and measure the extent to which manufacturing progress has occurred in the global wind turbine industry. A key feature of this setting is that the quality of a wind turbine and the cost to produce it are governed by relatively simple physics. The mapping from turbine characteristics to electricity production (the ultimate source of adopter value) is disciplined by observable engineering relationships, with output increasing approximately quadratically in turbine size. At the same time, the physical material necessary to make a turbine (the core driver of manufacturing costs) is approximately cubic in size. We use this "square-cube" structure to relate the costs of producing turbines of different vintages and quality levels, and thereby to estimate latent technology manufacturing cost states for the industry over time.

Our primary data consist of turbine choices made by the near universe of wind farms constructed globally between 2000 and 2019. For each project, we observe the full set of available turbine models, engineering estimates of each model's expected output at that location, and the turbine actually installed. Unlike conventional discrete choice settings where prices are posted, wind turbines are heavy industrial goods acquired through informal

procurement processes that generate no public transaction records. We therefore model each wind project's turbine selection as a procurement scoring auction, and show how data on turbine characteristics, site wind resources, and observed choices—together with standard discrete choice modelling tools—identify the otherwise latent "bids" that manufacturers submit. Unlike a conventional hedonic analysis, this framework separately identifies markups from underlying costs, allowing us to trace the evolution of manufacturing costs rather than prices.

Our bid estimates reveal substantial technical progress in the industry between 2000 and 2019 (figure 7). Average bid per unit of turbine capacity ended the sample roughly 30% lower than in 2000. Using this metric of renewable energy costs, which is standard in popular reporting, cost declines appear modest. However, this ignores the fact that turbine output per MW increased by roughly 50% during that time. Combining these two measures, the average bid per unit of output (or levelized cost of energy) was more than 50% lower at the end of the sample. Given the "sqaure-cube" nature of this device, generating this decline in cost per unit of output required an even larger improvement in manufacturing efficiency. Our preferred measure of wind turbine manufacturing prowess, average cost per unit of turbine volume, declined by over 75%. This simple decomposition underscores the importance of accounting for quality improvements when measuring technical progress in an industry where quality is improving.

Having recovered latent cost states, we then ask how much of this technical progress can be attributed to learning by doing. We relate changes in turbine costs to cumulative manufacturing experience at the turbine, firm, and industry levels. To allow for knowledge spillovers, we adopt the logic of Irwin and Klenow (1994) and estimate how experience accumulated on one turbine affects the costs of other turbines produced by the same firm, as well as turbines produced by rival firms. To address the endogeneity of experience, we exploit variation in renewable-energy policies across markets, interacted with the strong home-market patterns in global wind turbine demand. This provides variation in the pace and location of accumulated manufacturing experience that is informative about learning but plausibly orthogonal to contemporaneous cost shocks.

Using this framework, we document considerable learning by doing. A doubling of manufacturing experience reduces manufacturing costs by 24 to 33 percent, an effect which is similar to the results on aircraft manufacturing in Benkard (2000). Additionally, spillovers appear to be important, though small on the margin. Approximately 2 to 10 percent of experience spills over to other turbine models produced by the same firm, and spillovers from turbines produced by other firms are on the order 0.1 to 0.2 percent. Though the marginal effects of spillovers are small, aggregate experience within a firm is often 2 orders of magni-

tude or more larger than turbine-specific experience, and aggregate experience outside the firm is another order of magnitude larger. Thus, spillovers have generated significant cost reductions over time. As an example, we show that when the Chinese wind industry began in the late 2000's, Chinese manufacturers' entered at a cost structure much more commensurate with established western firms than their limited own manufacturing experience would suggest.

While our primary specification assumes that manufacturers set prices to maximize current period profits, this static assumption may be violated if firms anticipate that current sales reduce future costs and shade prices down accordingly—the dynamic markdown problem emphasized by Besanko et al. (2014) and Benkard (2004). If present, such incentives would cause us to overstate manufacturing costs early in the sample, when firms are on steeper portions of their learning curves, potentially biasing estimated learning rates downward. We therefore also conduct sensitivity analyses which attempt to account for these dynamic incentives.

To do so, we adopt a procedure proposed in Berry and Pakes (2000), which characterizes a first-order condition for dynamic oligopoly problems with continuous controls without requiring explicit computation of the dynamic equilibrium. The key insight is that when firms have optimally set a continuous control—such as a bid in a procurement auction—the dynamic component of the problem can be expressed as a rational expectations term plus a shock. This term is a function of the state transition distribution and subsequent realized observables, which are far easier to compute than the full equilibrium of a dynamic game, and scale naturally to settings with many firms and products while avoiding the curse of dimensionality. We use this approach to separately identify static manufacturing costs and dynamic pricing incentives.

In addition to measuring learning in an economically large and increasingly policy relevant industry, this paper thus also provides a novel and generic tool for measuring the effects of learning by doing, as well as any other dynamic component of a first-order optimality condition, in settings where only demand side data are available. The paper closest to ours in this dimension, and indeed a key inspiration for our functional form assumptions, is Irwin and Klenow (1994), which studied learning by doing and spillovers in the global semiconductor industry in the 1970s-1980s. Like this paper, Irwin and Klenow (1994) was only able to observe standard demand estimation data, like prices and quantities. However, to handle the dynamic incentives that semiconductor manufacturers may have faced, Irwin and Klenow (1994) assumed that Cournot quantity choices were first-order optimal with respect to a standard Euler condition, an approach that is not necessarily consistent with most modern models of dynamic oligopoly behavior. Our method of using the Berry and

Pakes (2000) approach to a learning setting thus complements the original idea in Irwin and Klenow (1994), bringing its insights to the modern structural IO research paradigm.

Our paper also contributes to the broader learning-by-doing and innovation literatures with a specific focus on energy and environmental economics (Acemoglu et al., 2019). Newell et al. (1999) demonstrate that energy price shocks cause manufacturers to adopt more energy saving technology, consistent with an "induced innovation" hypothesis. Using data similar to ours, Knittel (2011) estimates considerable improvements in the production capabilities of car manufacturers, despite being seemingly constrained by fuel economy standards. In a recent paper, Barwick et al. (2025) also use IO tools to infer costs and estimate learning by doing in the global EV battery industry. While largely complimentary, two differences between that paper and ours are that wind turbines have a natural structural cost function which we directly estimate, and the attempt to control for dynamic incentives rather than assuming firms maximize static profits.

The paper proceeds as follows. Section 2 describes the physical and industry features of wind turbine production that motivate our empirical approach. Section 3 describes the data. The next sections recover latent costs and use them to measure technical progress and learning by doing. The final sections discuss dynamics, policy implications, and conclude.

2 Background

In this section, we provide additional background on some key physical concepts and industry features that we leverage in estimation. To summarize, wind turbine production is quadratic in the size of the device, while the costs manufacturing a turbine are cubic in size. The industry is highly concentrated, and, while all the major players are active globally, "home" region preferences generate even more concentration at the market level. All manufacturers produce multiple turbine sizes at the same point in time, a fact we will later leverage in our estimation strategy.

2.1 Wind Power Basics

A wind turbine consists of a rotor with three long blades connected to a gearbox and generator atop a large tower. As wind passes through the blades, the rotor spins a drive shaft connected through a series of gears to a generator that converts this kinetic energy to electrical energy. The amount of energy such a device can capture is given by "Betz" law:

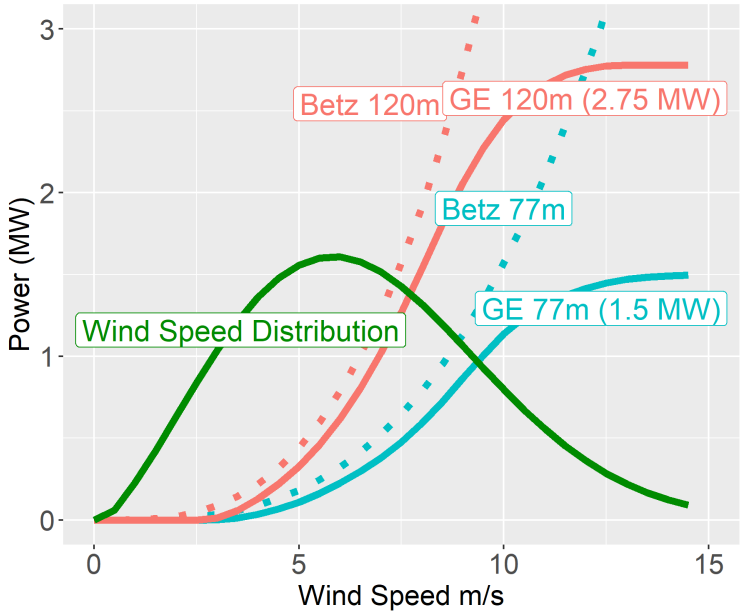
$$\text{Power} \quad Q = C_p \left[\frac{1}{2}\pi r^2\right] \left[dv^3\right] \tag{1}$$

where C_p is the "power coefficient", or the ratio of the power flowing through the device that is captured, and d is the density of the air the turbine is exposed to. Betz (1926) demonstrated that the theoretical limit on C_p is $C_p^{\max} = \frac{16}{27} \approx 0.593$.

In practice, commercial turbines produce less than this optimum for two reasons. First, real world turbine blade designs never quite achieve the theoretical maximum at any speed (although some are remarkably close for a wide range of v's). Second, and more importantly, generators, which convert captured power into electricity, have a maximum capacity which typically chokes off the devices power at high wind speeds. This is partially to avoid extreme stress on the device, but also largely for economic reasons: if generator costs are increasing in size, and high wind speeds occur infrequently, then it doesn't make sense to pay for generating capacity that will be rarely used.

Figure 1 presents the Betz frontier and power curves for two different sized General Electric wind turbines. Power curves are functions, provided by the manufacturer to prospective buyers, which map wind speeds into the device's output. They are the empirical analogue to Betz' law. The figure also includes the probability density function of wind speed for a typical location. Wind speeds are well approximated by a Weibull distribution, with means between 6 and 8 meters per second (m/s). In this example, both turbines closely match the Betz frontier, until they achieve their "rated power", denoted in megawatts (MW).

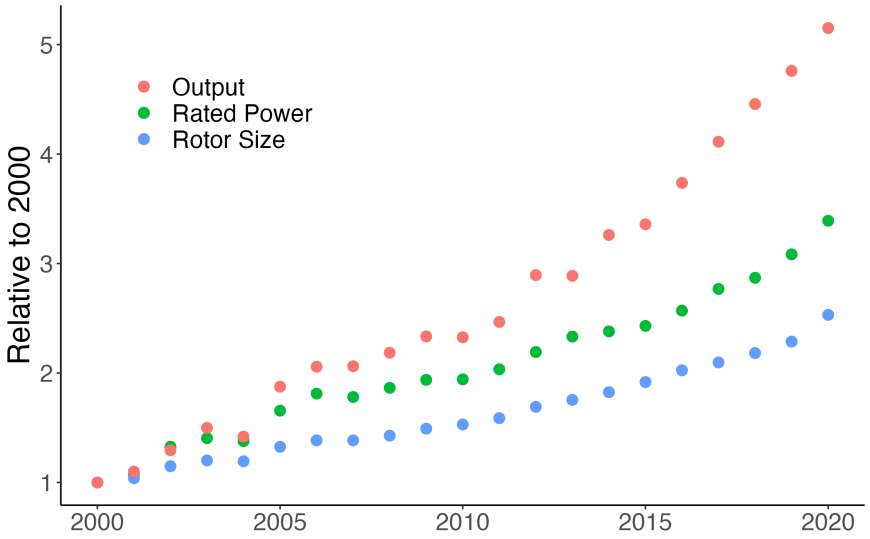
Figure 1: Example Power Curves



One immediate implication of Betz law is that the wind turbine production function is characterized by increasing returns to size. This means that, all else equal, *better* wind tur-

bines, are generally *bigger* wind turbines. This physical relationship has underpinned much of the rapid growth the industry has experienced in recent years. Since 2000, wind turbine rotors have more than doubled in size, while expected output per turbine has increased nearly fourfold (Figure 2).

Figure 2: Global Turbine Size and Output Trends



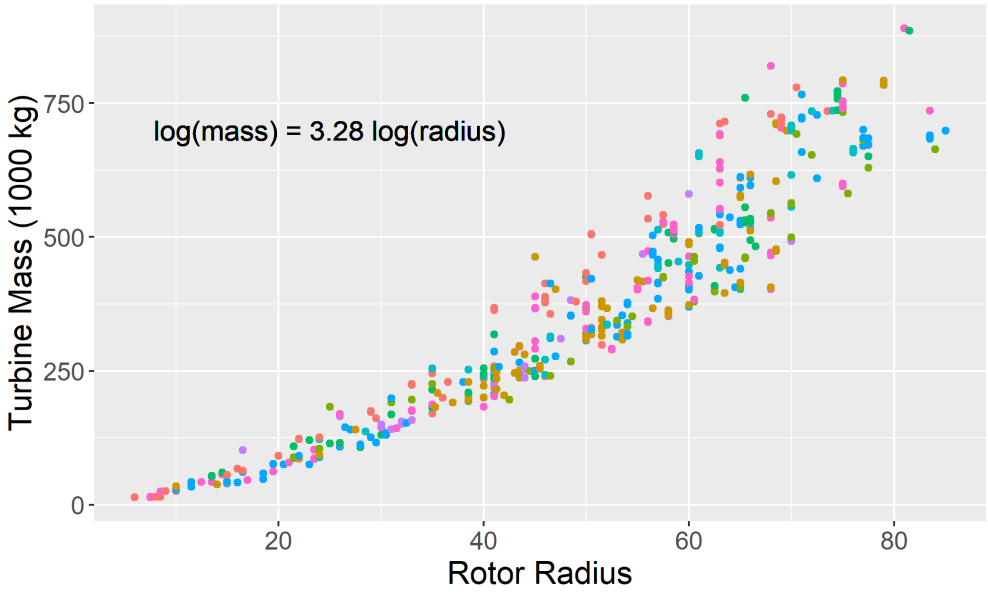
Global average installed rotor size, rated power and *predicted* output (author's calculations), relative to their year 2000 values. Source: BNEF.

Given that the underlying technology demonstrates increasing returns to scale (rotor size), it is natural to question why wind turbines weren't initially larger, and why they aren't even bigger now. It turns out that the constraints on rotor size are also physical in nature. Galileo's "square-cube" law, adapted to this setting, states that the volume of materials necessary to produce a rotor of radius r should be proportional to r^3. Thus, while large turbines produce more power than smaller turbines do, their costs of manufacturing are more than proportionately larger (for the same material and design). To profitably make bigger turbines, firms must either develop new turbine designs or they must employ newer manufacturing materials in existing designs (or both). Newer designs or materials allow firms to increase rotor size r with a quadratic benefit in power generation at a less than cubic increase in costs. Innovation in the wind turbine industry, as reflected in the arrival of ever bigger turbines, is grounded in these engineering investments.

Beyond a simple physical observation, this cubic mass relationship has also been demonstrated in real world turbines. In the 2000s, the DOE conducted a series of studies on the limits to manufacturing large turbines. As part of this, they produced a software tool which estimates the mass of materials necessary to produce a turbine given user-entered character-

istics. Figure 3 plots the log mass vs rotor size for all turbines in our sample, showing that estimated turbine mass indeed grows faster than quadratically in rotor size.

Figure 3: Turbine Mass vs Size (NREL)



2.2 Industry Economics

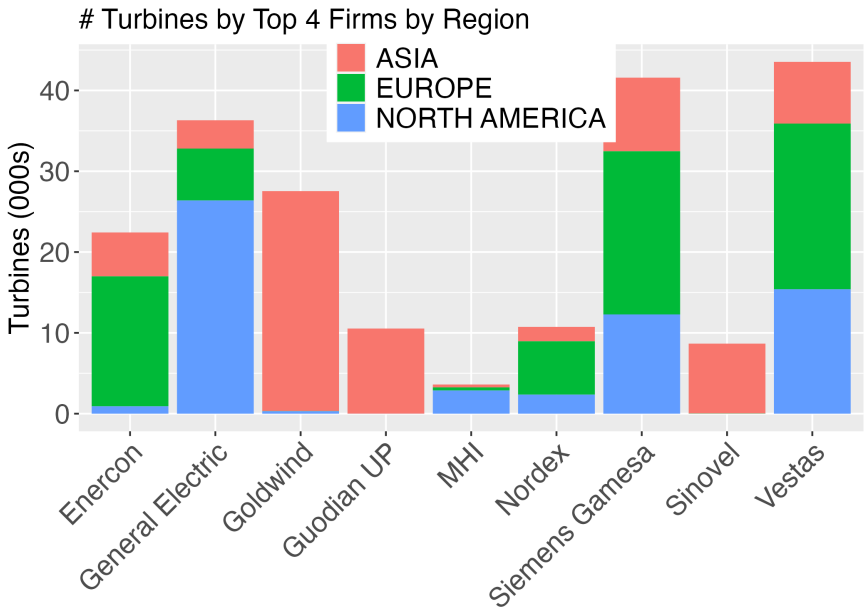
The wind turbine market is highly concentrated. Table 1 presents sales for the ten largest firms. The four largest firms have over half of global turbine sales, and the top ten have nearly 80 percent. Although the industry is global, in that the top firms supply every market, sales are regionally concentrated. This is at least partially due to home market bias (Coşar et al., 2015). Figure 4 presents the distribution of sales by region, for the top four firms in each region, confirming that General Electric does most of its sales in the US market, Enercon, Siemens Gamesa, and Vestas do most of their sales the European market, and the two major Chinese manufacturers, Goldwind and Guodian UP, sell primarily in Asia.

Table 1: Sales by Firm

Firm	Country	Models	Capacity	Projects	Turbines
Vestas	Germany	41	100958	4264	49911
General Electric	United States	28	68336	1433	37111
Gamesa	Spain	26	48932	1560	31212
Goldwind	China	22	47449	896	29286
Enercon	Germany	28	42967	3946	24054
Siemens	Germany	15	28722	644	13629
Nordex	Germany	21	22985	1215	9460
Guodian UP	China	11	18728	377	11256
Senvion	Germany	13	14702	1079	6919
Suzlon	India	13	14678	1040	8472

This table presents the total number of turbines sold, total capacity and number of wind farms for the ten largest firms (by turbine sales) in the BNEF data.

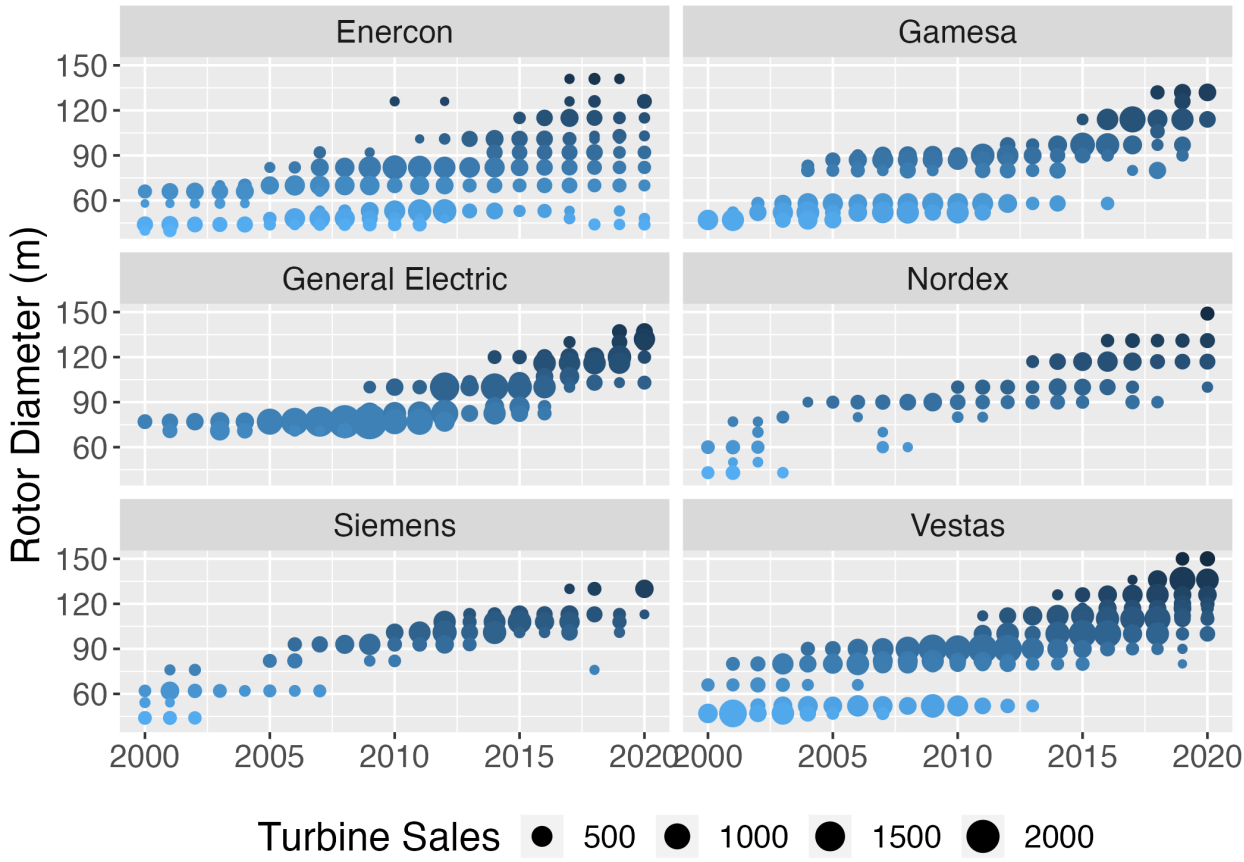
Figure 4: Sales by Region



At any given point in time, manufacturers produce multiple turbines. As new turbine models are introduced, old models are discontinued. Figure 5 plots the annual mix of turbine sizes offered, scaled by sales in that year, for the six largest western manufacturers. For all firms, the mix of rotor sizes increases over time, with the largest rotors in 2018 being 50%-100% bigger than the largest rotors in 2000, when our data begins. Additionally, Figure 5 also shows that when new, larger turbines are introduced, they do not immediately take

over all of the market share from existing smaller turbines, despite being more productive. Instead, these larger turbines gradually gain market share over time, which provides suggestive evidence that the relative cost difference must be declining over time, consistent with learning.

Figure 5: Sales by Turbine Size, Top Western Firms



3 Data

The primary data come from a proprietary list of global wind farms maintained by Bloomberg New Energy Finance (BNEF). For almost all projects, this list includes the exact location of the wind farm (geocoded), its capacity in turbines and output, the date proposed and the date commissioned (if ever). Importantly, for most commissioned wind farms, BNEF also records the exact turbine model installed. We match these turbines to a detailed database of turbine characteristics from The Wind Power (TWP), an energy marketing consultancy. TWP's turbine database includes information about the rotor diameter and rated power of

each turbine, along with dozens of other technical specifications. TWP also maintains a nearly comprehensive list of turbine power curves (as presented in Figure 1).

We supplement this database with project-specific information on wind speeds and power prices. We purchased information on the distribution of wind speeds from Vaisala, a commercial vendor widely used in the industry for siting purposes. For each project, we obtain the parameters of a Weibull distribution of wind speeds. Integrating a turbine's power curve over the distribution of wind speeds allows us to estimate the expected output from any given location-turbine combination, including those not selected in the data.

To estimate the value of this output at a given site, we employ data from several sources, of varying degrees of specificity depending on the region and time period. For all non-regulated US projects, EIA Form 860 provides average annual revenue from resale sales. Many countries outside the US support wind farms with feed-in-tariffs, which we obtained by year from the OECD and BNEF. In recent years, turbine contracts have become increasingly awarded via auction. BNEF maintains a database of all wind farm auctions, as well as the winning projects, their bids and the award price ($/MWh). Where none of these prices are available, we use the average wholesale price in a country-year, or country-state-year. Figure 6 plots average wind-specific power prices over time for selected large markets.

For analysis, we make a number of sample restrictions. Starting with the full BNEF database, which contains projects going back as far as the 1980s, and many not yet built, we exclude projects built prior to 2000 or not yet commissioned. We exclude small projects, of less than a megawatt. Of this sample, we exclude some projects where BNEF either does not have turbine information or we were unable to locate the turbine specified in the TWP database. Finally, we exclude offshore projects.

Table 2: Sample Construction

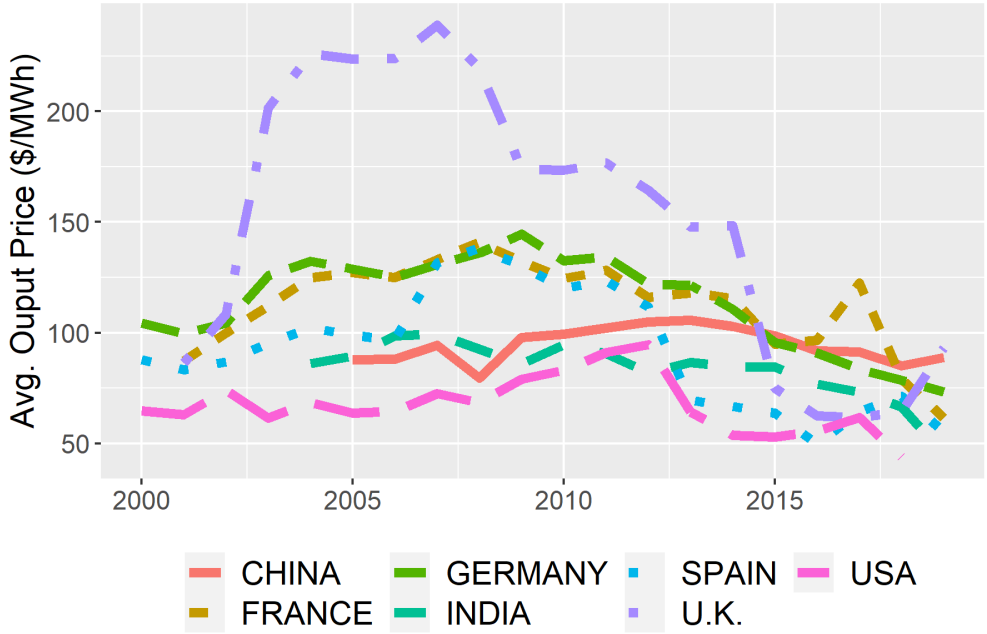
Group	N	mean
All Projects	22,186	1.00
Completed, Post 2000	21,316	0.96
Capacity .5 MW	21,057	0.95
Turbine match	20,543	0.93
Onshore	20,313	0.92

Source: BNEF Project Database

Table 3: Sample Summary

Variable	mean	sd	max	median	min
Capacity (MW)	26.43	39.72	600.00	12.00	0.50
# Turbines	14.97	22.56	426.00	6.00	1.00
Rated Power (MW)	1.95	0.81	12.00	2.00	0.06
Wind speed (m/s)	6.61	0.99	15.86	6.52	2.36
Rotor Radius (M)	42.71	11.40	110.00	41.25	10.00
Est. Turbine Output (kW/hour)	664.78	387.44	4,419.43	613.14	5.80

Figure 6: Average Wind Output Price by Country



4 Recovering Turbine Bids

In the typical dataset used for estimating marginal costs of producing consumer products, researchers have access to both aggregate quantities and posted (assumed common) prices. For large capital goods like wind turbines, there are no posted prices. Turbines are procured in a confidential, developer-specific process. Once completed, the terms agreed upon are rarely disclosed. Thus, before estimating manufacturing costs using the inversion of a demand system, we must recover the prices developers face when choosing which turbine to install. To

do this, we develop a model of turbine procurement auctions, and show how the parameters of this model represent the bids manufacturers submit.

4.1 A model of turbine procurement

We begin by describing a procurement model with purely static incentives (dynamic incentives will be discussed in section 6). At time t, each firm f in the set of manufacturers F has a portfolio of turbines $j \in K_{ft}$. During t, each project i in a set of wind farm projects W_t requests bids for each turbine offered by each firm. If project i selects turbine j, i earns discounted expected revenues $R_{ijt} + \epsilon_{ijt}$, where ϵ is an iid shock to developer profits, known only by the developer.

The revenues that turbine j generates at site i are determined by the known wind speed distribution $F_i(v)$ for site i, the known power curve $P_j(v)$ which maps wind speeds v into megawatts of output for turbine j, and p_i^{output}, the output price site i receives for its production. Projects in the data have markedly different sizes, and we assume that project i will buy the same total capacity, regardless of which turbine j it selects. Because turbines have different rated capacities (MW_J), the number of turbines i must buy to meet its capacity target will vary with j. We therefore model discounted lifetime revenue per unit of capacity as:

$$R_{ij} = \underbrace{p_i^{\text{output}}}_{\text{price per megawatt hour}} \times \underbrace{\int P_j(v) dF_i(v) dv / MW_i}_{\text{expected capacity factor}} \times \underbrace{8760}_{\text{hours in a year}} \times \underbrace{\sum_{t=1}^{20} \delta^{t-1}}_{\text{discounted value for 20 years}}$$

Because the capacity factor (expected output per hour divided by turbine capacity) varies with j, expected revenue R_{ijt} varies with j as well even though total plant capacity is the same.

We assume that i conducts a second price procurement auction to determine which turbine it buys, and what price it pays. In this mechanism, each turbine manufacturer submits a (potentially) project-specific bid per MW b_{ijt} for each turbine j in its portfolio. Project i selects the turbine that delivers the highest net surplus, which we define as $R_{ijt} - b_{ijt} + \epsilon_{ijt}$. If turbine j from firm f wins the auction, the payment $\widetilde{b}_{ijkt}$ from i to f is designed to make i indifferent between choosing j and paying this price, and choosing the next highest net surplus turbine, k, and paying the actual bid offered by k's manufacturer. That is, $\widetilde{b}_{ijkt}$ satisfies

$$R_{ijt} - \widetilde{b}_{ijkt} + \epsilon_{ijt} = R_{ikt} - b_{ikt} + \epsilon_{ikt}$$

so that the payment is

$$\widetilde{b}_{ijkt} = \underbrace{b_{ikt}}_{\substack{\text{bid of second best} \\ \text{turbine proposal}}} + \underbrace{(R_{ijt} + \epsilon_{ijt})}_{\substack{\text{revenues from best} \\ \text{turbine proposal}}} - \underbrace{(R_{ikt} + \epsilon_{ikt})}_{\substack{\text{revenues from second} \\ \text{best turbine proposal}}}$$

Thus, though the probability that a firm wins a sale does depend on its bid, its payment, conditional on winning, does not, as a result of these second-price rules.

Asker and Cantillon (2008) show that the unique dominant strategy equilibrium for this game is for each firm to bid exactly its opportunity cost of delivering each turbine, and, as a result, site developers pick the turbine which maximizes social surplus, the sum of developer and manufacturer profits. This implies that, conditional on the opportunity cost of supply, bids will not be project-specific in equilibrium. Based on this result, we assume that manufacturing costs are constant within a given year, so that if two projects i and i' are built in the same year, $b_{ijt} = b_{i'jt}$. In a slight abuse of notation, we define the (common) bid for turbine j in year t as b_{jt}.

For tractability, we make three additional assumptions. First, we assume that the project-by-turbine revenue shocks are distributed as type-1 extreme value. Second, we assume that every project which solicits bids from manufacturers chooses a turbine (e.g., there is no outside option). Finally, we assume that developer's consider bids on all turbines offered by all active manufacturers at a given point in time (e.g., that choice sets are common). These additional assumptions imply that the probability firm f wins the site i auction with turbine j at time t is:

$$s_{ijt} = \frac{\exp(R_{ijt} - b_{jt})}{\sum_{f' \in F} \sum_{l \in J_{f'}} \exp(R_{ilt} - b_{lt})}$$

With this notation, we can define the static expected profits for firm f at time t. Let c_{jt} be the marginal cost (per MW) of delivering turbine j at time t. Given our previous definition for f's revenues when it sells turbine j, its static profits per MW from a successful sale to site i are:

$$\pi_{ijt} = R_{ijt} - b_{jt} + \epsilon_{ijt} - \max_{k \in \prod_{\phi \neq f} K_{\phi,t}} (R_{ikt} - b_{kt} + \epsilon_{ikt}) + (b_{jt} - c_{jt})$$

Because developers always choose the turbine which generates the highest net surplus, we can write f's total realized static profits from offering its turbines to site i as:

$$\pi_{ift} = \max_{j \in \prod_{\phi} K_{\phi,t}} (R_{ijt} - b_{jt} + \epsilon_{ijt}) - \max_{k \in \prod_{\phi \neq f} K_{\phi,t}} (R_{ikt} - b_{kt} + \epsilon_{ikt}) + \sum_{j \in K_{ft}} \mathbb{I}\left[i \text{ chooses } j\right] (b_{jt} - c_{jt})$$

Note that this expression includes the possibility that firm f does not make a sale: when i chooses a turbine *outside* of K_{ft}, the first expression above is identical to the second, so the entire expression is zero. Thanks to the type-1 extreme value assumption on the ϵ's, the firm's expected profits at site i are:

$$\mathbb{E}\pi_{ift} = \underbrace{\log \sum_{j \in \prod_\phi K_{\phi,t}} \exp\left(R_{ijt} - b_{jt}\right)}_{S_{it}} - \underbrace{\log \sum_{k \in \prod_{\phi \neq f} K_{\phi,t}} \exp\left(R_{ikt} - b_{kt}\right)}_{S_{it}^{-f}} + \sum_{j \in K_{ft}} s_{ijt}(b_{jt} - c_{jt})$$

In words, this is the logit inclusive value of the wind developer's net surplus among all turbines (S_{it}), minus the logit inclusive value of the wind developer's net surplus among all turbines excluding those from firm f (which we call S_{it}^{-f}), plus firm f's expected markup over its marginal costs.

4.2 Empirical implementation

We use this structure to recover the unobserved procurement auction bids by modeling the developer's discrete choice problem, allowing for a slightly richer notion of developer preferences. In equation 4.1, developer revenues per MW simply the product of observed (typically statutory) output prices and the discounted expected capacity factor. We allow developers to have preferences per unit of output that differ from this price along two dimensions. First, developers may value output at more or less than the statutory price, for a couple of reasons. First there may be measurement error, uncertainty or other wedges, such as curtailment, that reduce the realized payment. Second, although we assume that all turbines are available to all projects, in practice there are physical constraints which limit the of some turbines at especially windy sites. Because there are no documented hard and fast rules which indicate which turbines are allowed at which sites, we include dummy variables that indicate whether a site's measured wind speed distribution lies outside of the turbine's recommended range. Our modified term is

$$R_{ijt} = \exp(\sum_c \mathbb{I}[c(i) = c]\left(\alpha_0^c + \alpha_1^c \ln(p_{ijt})\right) + X_{ij}^W \beta^W) q_{ijt} \tag{2}$$

where p_{ijt} is the statutory output price for site i and turbine j at time t, and q_{ijt} is the discounted expected capacity factor. The vector of parameters $\{\alpha_0^c, \alpha_1^c\}$ allows for country-specific preferences over output prices, while the vector β^W captures the importance of turbine-site compatibility.

Project developer i in period t chooses a utility-maximizing turbine j given:

$$u_{ijt} = \alpha_0 R_{ijt} + \alpha_H R^H_{ijt} + \sum_c \alpha_c \mathbb{I}[c(i) = c] R_{ijt} + X^D_{ijt}\beta^D - b_{jt} + \epsilon_{ij} \tag{3}$$

The vector X^D_{ijt} captures additional sources of preference heterogeneity that vary with the turbine selected, independent of the output produced. Specifically, we control for well-documented home-market bias in the wind turbine industry (Coşar et al., 2015). To do this, we include dummy variables which are equal to 1 if the project's country is the same as the country where the turbine manufacturer's headquarters lie. Finally, we included a dummy for whether a firm has a manufacturing facility in the country.

Our assumption that manufacturer turbine bids are constant within each year means that we can estimate those bids as turbine-year fixed effects. Because our demand system does not have an outside option, what our estimates recover is actually the difference in bids between turbine j and the bid for a reference turbine, $\hat{b}_{jt} = b_{jt} - b_{0t}$. We define the reference turbine as the best selling turbine manufactured by Vestas in each year. With data on site-by-turbine expected revenues, other turbine characteristics, and each site's turbine choice, we estimate the b's, the α's and the vector β^D using maximum likelihood.

4.3 Bid estimation results

Table 4 presents the estimated preferences from four separate demand specifications. The top section presents the estimated value of output from equation 2 by country (standard deviations, not standard errors are presented). Notable lower values include China, where it has been documented that early large wind farms sent significantly less output to the grid than their plants implied production (Chan 2017), and Italy, where subsidies implemented via an uncertain and highly volative green certificate scheme rather than a fixed tariff.

Table 4: Wind Plant Turbine Preferences

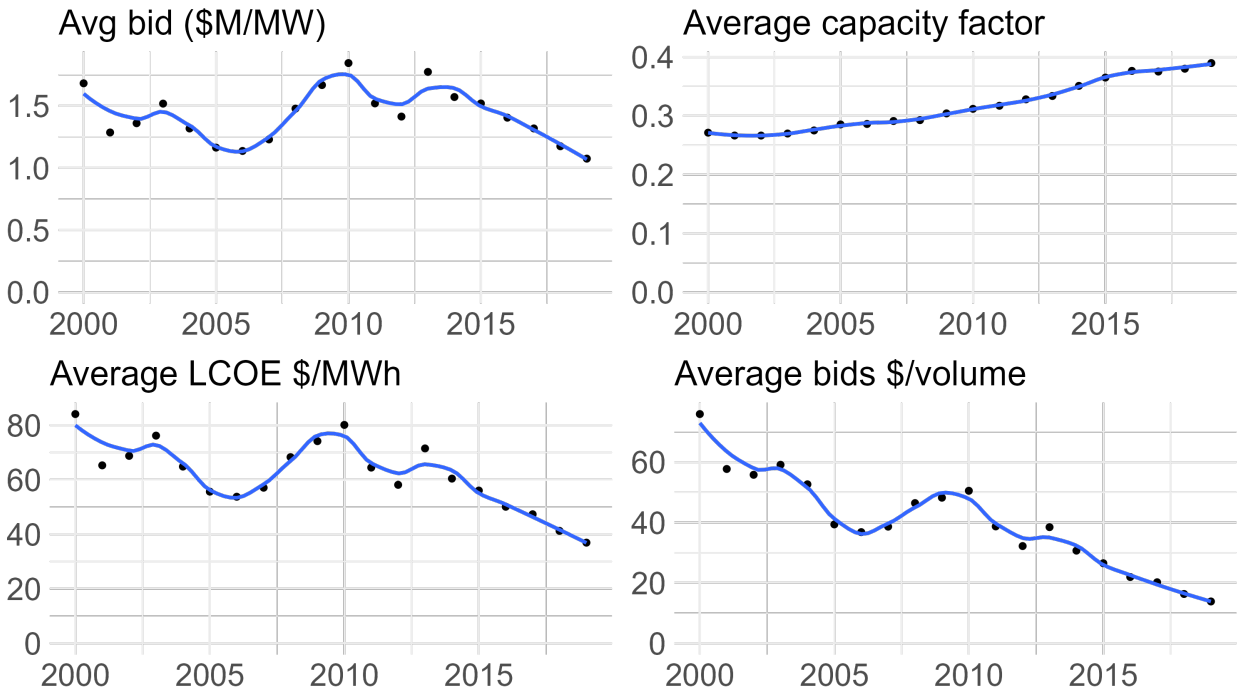
	Model 1		Model 2		Model 3		Model 4	
Parameter	Point Est.	Std. Err.	Point Est.	Std. Err.	Point Est.	Std. Err.	Point Est.	Std. Err.
Value of Output ($/MW)								
AUSTRIA	110.1	65.3	78.1	47.5	128.7	68.6	87.3	47.9
BELGIUM	81.4	79.3	56.0	50.3	97.0	84.0	66.0	54.2
CHINA	16.2	16.3	12.7	13.9	16.9	16.8	13.4	13.9
DENMARK	54.5	11.4	41.1	8.1	54.1	8.4	38.0	6.5
FRANCE	100.2	54.7	70.2	36.7	117.2	65.0	79.8	43.0
GERMANY	215.8	41.4	149.9	24.3	229.1	40.0	156.6	25.4
ITALY	39.0	60.7	28.5	40.3	43.4	63.1	29.3	40.9
PORTUGAL	187.6	45.5	138.4	34.5	223.5	39.1	156.3	29.6
SPAIN	127.5	111.3	94.0	80.5	130.3	120.9	92.0	85.3
SWEDEN	160.8	11.2	120.6	7.8	206.8	9.2	144.9	6.3
UNITED STATES	172.5	5.9	132.0	7.6	174.5	7.1	125.6	7.9
Site/Turbine Class Compatibility								
(Site: 1, IEC: II)	0.015	0.017	0.013	0.017	0.009	0.016	0.008	0.016
(Site: 1, IEC: III)	-0.045	0.024	-0.059	0.025	-0.043	0.023	-0.051	0.023
(Site: 2, IEC: III)	-0.081	0.016	-0.085	0.016	-0.077	0.015	-0.079	0.015
Home Bias/Manufacturing								
Foreign	-1.879	0.046	-1.446	0.198	-1.511	0.050	-1.138	0.147
China x Foreign	-4.547	0.320	-3.732	0.612	-4.455	0.321	-3.250	0.510
Import					-0.860	0.044	-0.577	0.080
China x Import					-0.324	0.246	-0.164	0.162
Heteroskedasticity Over Time								
2001			0.170	0.185			0.028	0.168
2002			-0.324	0.172			-0.323	0.163
2003			-0.127	0.171			-0.268	0.158
2004			-0.232	0.167			-0.321	0.156
2005			-0.134	0.171			-0.179	0.160
2006			-0.486	0.161			-0.478	0.152
2007			-0.565	0.158			-0.502	0.149
2008			-0.461	0.158			-0.440	0.148
2009			-0.517	0.155			-0.499	0.145
2010			-0.378	0.169			-0.359	0.156
2011			-0.573	0.163			-0.514	0.154
2012			-0.036	0.172			-0.112	0.161
2013			-0.343	0.191			-0.421	0.174
2014			-0.205	0.160			-0.295	0.152
2015			-0.166	0.158			-0.324	0.150
2016			-0.081	0.157			-0.219	0.148
2017			-0.171	0.153			-0.269	0.145
2018			-0.541	0.170			-0.577	0.159
2019			-0.440	0.167			-0.470	0.158

The next panel in the table documents the importance of turbine-site compatibility. High wind sites (class 1) appear to strongly dislike turbines designed for low wind sites (classes II and III), and class 2 sites dislike class III turbines. Next we also find considerable heterogeneity in firm preferences across countries, independent of output. Turbines are large, heavy pieces of equipment, and having a production facility inside the country makes developers significantly more likely to select a firm's turbine. Even conditional on that though, we find that developers have a large willingness to pay for home market turbines, consistent with (Coşar et al., 2015). Finally, in looking at the raw data, it is clear that Chinese developer preferences along these dimensions appear different from the rest of the

sample. We rationalize this with the inclusion of a foreign producer dummy for the country, and estimate this effect to be twice as large as the other home bias effects combined.

Next we plot the average estimated bid, normalized by different notions of quality change over time (figure 7). The first panel plots the average bid per MW, which is the metric most commonly reported in press and industry analyses of this sector. Average bids hovered between \$1.5 and \$1.2 million per MW until 2012, at which point they steadily declined, ending the sample roughly 30% lower than in 2000. Using this measure, cost declines appear modest. However, this ignores the fact that turbines in 2019 are much higher quality than in 2000. The second panel plots the average capacity factor (output per MW) in the sample. This shows that productivity per MW increased by roughly 50%. Combining these two measures, the bottom left panel plots the average levelized cost of energy (cost per unit of output over the life of the turbine). This shows that the average bid per unit of output, which is what developers care about, is more than 50% lower at the end of the sample than 2000. The final panel plots the average cost per unit of volume over time. This is the closest proxy to actual improvements in manufacturing efficiency. Here we see that bids per unit of volume have declined by over 75%. This simple decomposition underscores the importance of accounting for quality improvements when measuring technical progress in an industry where quality is improving.

Figure 7: Average bid trends

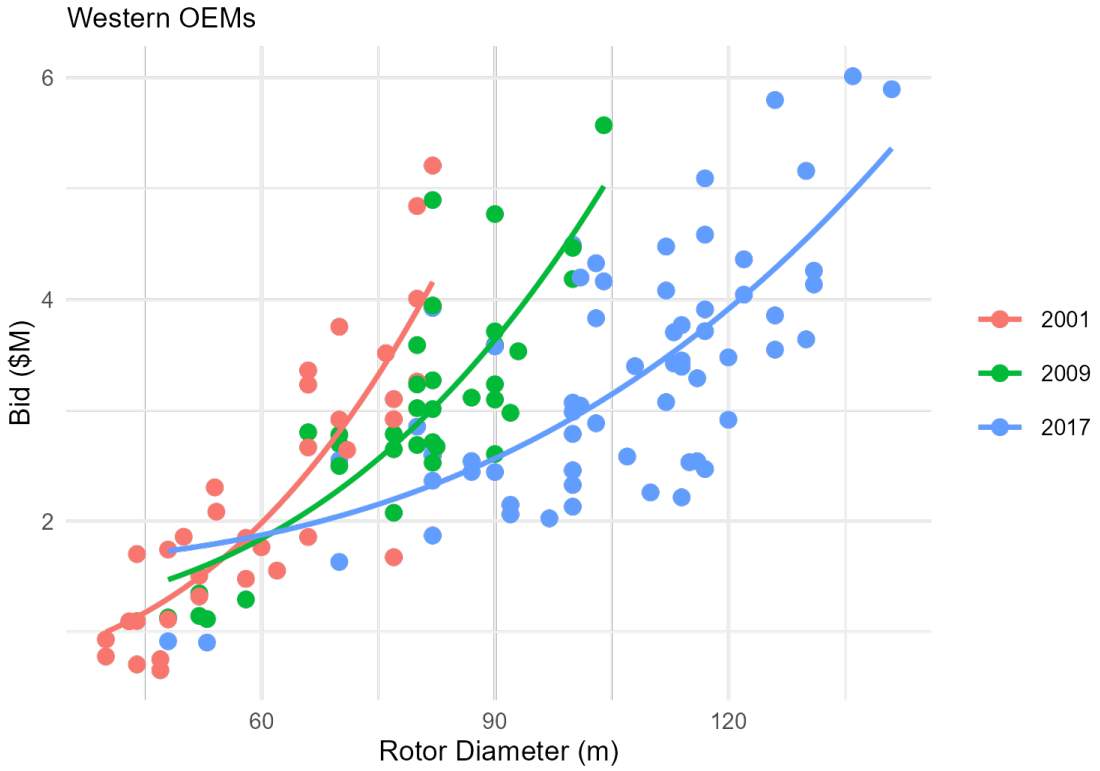


This figure plots annual capacity weighted averages. In the bottom left panel, LCOE stands for levelized cost of energy, which is the average bid per unit of lifetime output.

Figure 8 plots the estimated turbine bids against turbine size for three years in our sample. Looking across these three years, two patterns stand out. First, many turbine sizes are supplied over a long horizon. There are plenty of turbines in the 80m range that are sold in all three years (spanning nearly two decades), and there are also many turbines in the 100m range that are sold in both of the later years we plot. Moreover, across all sizes, the price of a turbine at a given size comes down considerably. Second, although the price of a turbine increases in turbine size, consistent with the discussion in section 2, the price gradient with respect to size flattens out dramatically over time. This change in cost structure is what underlies the remarkable shift towards bigger, therefore more productive, turbines over the past twenty years.

An alternative way to visualize the estimated bids is to look at the evolution over time within the same device. Figure 9 plots estimated bids for the most popular turbines sold by the two largest European manufacturers, Enercon and Vestas. Although the estimates are inherently a bit noisy, it appears that bids come down on average over time. However, it is important to remember that the direction of prices over time in an oligopolistic market with learning is theoretically ambiguous (Besanko et al., 2010, 2014). On the other hand, some of the *increases* in prices that occur leading up to 2010 are potentially driven by raw materials

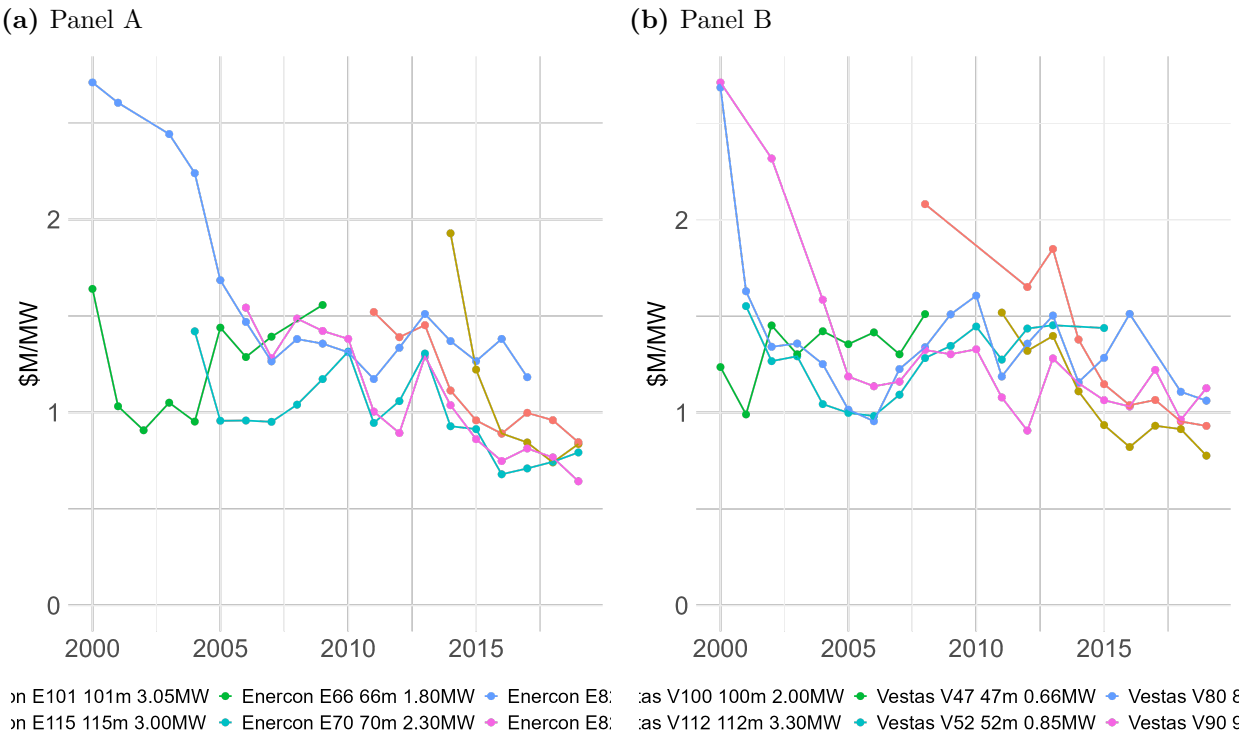
Figure 8: Estimated bids vs rotor size over time



Estimated bids are recovered relative to a reference turbine, which we pick to be the most common Vestas turbine sold that year. To plot levels, we impute the implied cost of this turbine from Vestas financials.

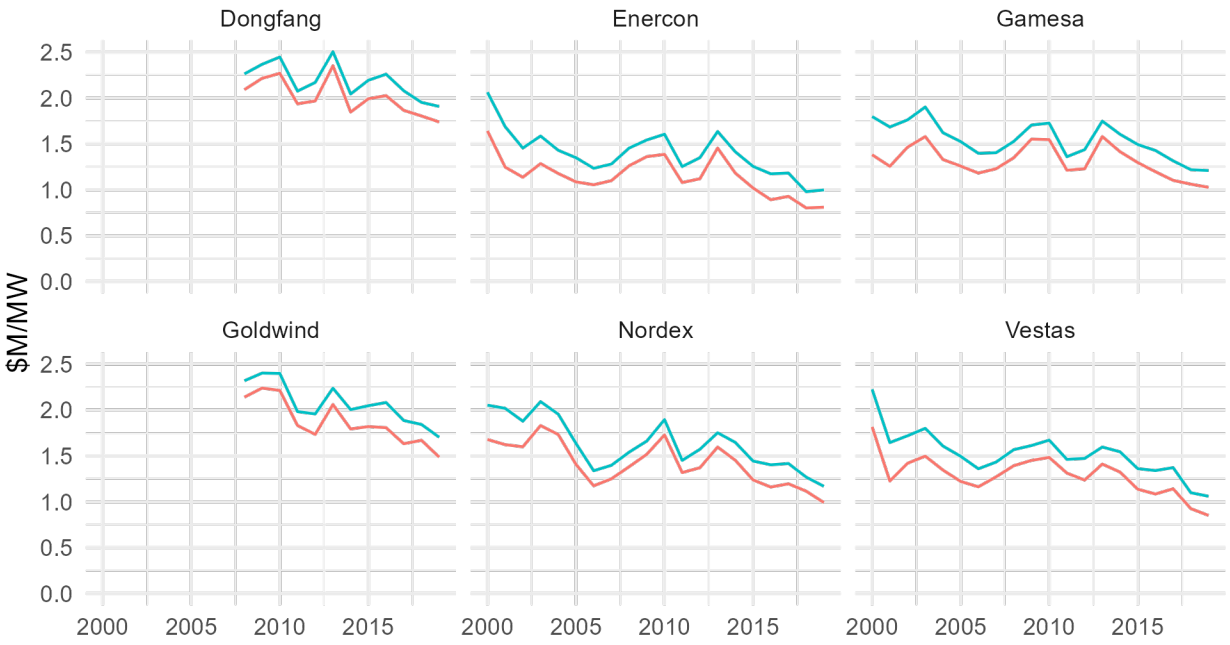
prices, including steel, which increased nearly 50% between 2005 and 2009.

Figure 9: Estimated bids for popular Enercon and Vestas turbines



Finally, while we are primarily interested in bids, our framework does allow us to estimate the implied "price" of a turbine, which is the expected payment from the developer to the manufacturer for that turbine. Figure 10 plots the average price and bid per MW for the six prominent firms. The difference between these two lines is the average markup, which appears to be decreasing over time for the major western manufacturers. This could be due to increased competition. But it is also consistent with the idea that learning-by-doing creates dynamic incentives for firms to price below marginal cost early on in order to walk down the learning curve.

Figure 10: Average price (blue) and bid (red) for top firms



5 Learning by Doing

5.1 Cost parameterization

Our goal is to measure the extent to which manufacturing costs correlate with various notions of manufacturing experience, at the turbine, firm, and industry-level. To do this, we need to impose more structure on costs c_{jt}. We assume turbine manufacturing costs factor into two terms: *resources* and *experience*. By "resources," we mean that bigger turbines require more material inputs to produce, conditional on firm experience. As we discussed in section 2, material inputs for manufacturing a wind turbine grow approximately cubically in rotor size r. However, firms with more experience may be able to economize on unobservable inputs, like labor, wasted materials, etc. We assume that this learning reduces firms' cost per unit turbine volume, such that the marginal cost of manufacturing turbine j at time t is

$$c_{jt} = \omega_{jt} r_j^3 = \omega_0 \left(\widetilde{E}_{jt}\right)^{\delta} r_j^3 \tag{4}$$

where ω_0 reflects the baseline cost of volume in the industry. The actual cost faced by turbine j at time t declines with "effective" experience $\widetilde{E}_{jt}$. Following Irwin and Klenow (1994), we

parameterize this as

$$\widetilde{E}_{jt} = \beta_J E_{jt} + (E_{ft} - E_{jt}) + \beta_W (E_t - E_{f(j),t}) \tag{5}$$

Effective experience is a linear combination of a firm's experience manufacturing a specific turbine (E_{jt}), the firm's experience manufacturing *other* turbines in its portfolio ($E_{f(j),t} - E_{jt}$), and the experience that all other firms in the industry have thus far accumulated ($E_t - E_{f(j),t}$).[1] We measure experience as the total cumulate number of turbines that firm $f(j)$ has previously shipped.[2] The parameter β_J thus measures the extent to which experience manufacturing turbine j is more or less useful than the firm's experience producing other turbines. Similarly, the parameter β_W measures the extent to which other firms' aggregate experience is as useful as a firm's own experience making other turbines. That is, a unit of turbine-specific experience is β_J times as valuable as firm experience, and a unit of industry experience is β_W times as valuable as firm experience. Finally, the parameter δ represents learning economies. Holding everything else equal, a 1% increase in effective experience increases costs by δ%. When $\delta < 0$, the learning-by-doing literature often reports $1 - 2^{\delta}$, the so-called "Spence coefficient" which measures the proportional effect of doubling of effective experience on costs.

Finally, recall that we do not actually observe or estimate the level of bid b_{jt}, as our discrete choice model does not have an outside option. Instead, our discrete choice model estimates deliver $\widehat{b}_{jt} = b_{jt} - b_{0t}$ for a pre-specified base turbine 0. We capture the effects of the base turbine bid b_{0t} using year fixed effects. Accounting for all of this, our estimating equation becomes:

$$\widehat{b}_{jt} = \omega_0 \left(\beta_J E_{jt} + (E_{ft} - E_{jt}) + \beta_W (E_t - E_{f(j),t})\right)^{\delta} r_j^3 - b_{0t} - \nu_{jt} \tag{6}$$

We estimate the model using generalized method of moments. A potential threat to identification comes from the possibility that there are other unobservable determinants of bids which may be correlated with observable determinants of costs, like firm- or turbine-specific experience. This would be the case if, for example, firms had serially correlated productivity shocks. If this were true, high productivity firms in previous years would have lots of experience and also lower than average costs in the present year. To deal with this potential source of endogeneity, we construct predictions of turbine sales that are driven exclusively by home-market bias and changes to country level demands. We then use these predictions to construct predictions of the experience firms accumulate from these forces alone, which

[1] Formally, firm experience $E_{f(j),t} = \sum_{l \in K_{f(j),t}} E_{lt}$, and industry experience $E_t = \sum_f E_{ft}$.

[2] We also consider cumulative megawatts and volume shipped measures in the appendix.

we use as instrumental variables for our observed experience measures.[3] The validity of this "Bartik" style instrument relies on the assumption that year-to-year changes in total demand for wind turbines across countries are uncorrelated with unobservable determinants of manufacturer costs.

5.2 Static LBD results

Table 5 presents the results. We restrict the estimating sample to include turbines sold by large firms, with at least two sales in the demand estimation sample in year t. In column 1, we assume that the cost of volume from equation 5 only depends on firm experience. The estimated learning parameter δ is -0.2. The corresponding Spence coefficient is reported at the bottom of the table. A 1% increase in firm experience decreases marginal costs by about 0.13%. In column 2 we allow a turbine's effective experience to include the experience of other firms. The estimate in the row labelled "Industry Expc," which represents β_W in equation 5, implies that a one unit increase in global experience at other firms generates 4 percent of the learning benefits as a unit increase in a firms own experience. In this model, the learning parameter is now much larger, implying that a 1% increase in effective experience decreases costs by 24%. In column 3, we allow effective experience to evolve differently across turbines within the same firm. The row labelled "Turbine Expc," which represents β_J in equation 5, implies that a one unit increase in experience for turbine j provides nearly 6 times the learning benefits generated by the firm's other turbines. In this model, where turbine j's experience has been separated out, global experience is now worth 5 percent as much as firm's own experience selling other turbines. The estimated learning parameter is relatively similar, implying that a doubling of effective experience decreases costs by approximately 27%.

[3]Specifically, we compute choice probabilities for each turbine-project pair using a demand system that includes turbine-project revenues and country-by-firm dummies. These choice probabilities reflect changes in each country's overall demand and the share of demand in a country that each turbine would typically expect to receive, but do not depend on year-to-year changes in turbine pricing.

Table 5: Learning Regressions - Turbine Experience

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
δ	-0.201	-0.401	-0.443	-0.501	-0.049	-0.404	-0.539	-0.500
	(0.012)	(0.012)	(0.012)	(0.012)	(0.014)	(0.016)	(0.018)	(0.018)
Turbine Expc			5.721	4.843			60.895	66.386
			(1.186)	(1.146)			(10.260)	(10.263)
Industry Expc		0.043	0.050			0.329	0.239	
		(0.003)	(0.004)			(0.034)	(0.021)	
Firm Expc - Post				0.131				0.000
				(0.124)				(0.441)
Industry Expc - Pre				0.039				0.179
				(0.003)				(0.017)
Industry Expc - Post				0.108				0.063
				(0.014)				(0.049)
Spence coef $(1-2^{\delta})$	0.130	0.243	0.265	0.293	0.033	0.244	0.312	0.293
	(0.007)	(0.006)	(0.006)	(0.006)	(0.010)	(0.008)	(0.009)	(0.009)
Year FE	X	X	X	X	X	X	X	X
OEM FE					X	X	X	X

All models estimated with GMM. Experience is measured by cumulative turbines sold. To account for endogeneity in turbine and firm experience, all models use predicted experience terms as instruments. Robust standard errors presented in parentheses.

Thinking about the way in which spillovers might occur, one can imagine two broad channels. First, firms may use knowledge gleaned from experience to improve the design of new turbines. Second, firms may use knowledge gleaned from experience to improve the manufacturing processes involved in producing a turbine once it is in production. Given the timing of these two channels, a natural further decomposition is to separate out experience gained before and after a turbine is introduced. This model is implemented in column 4. The reference category in this model is a firms own experience prior to the introduction of turbine j. "Firm Expc - Post" reflects experience gained after turbine j is introduced, and analogous categories are defined for industry experience. In this model, a firm's experience producing turbine j is worth 5 times as much as the experience it gained producing other turbines prior to j being introduced. The value of experience producing other turbines after introduction is only 13%, and this estimate is imprecise. Experience observed at other firms after j introduction is worth a similar amount, although other firms experience producing turbines prior to j introduction is worth only 4 percent as a firms own experience. The estimated learning parameter is now larger, implying that a doubling of effective experience

decreases costs by approximately 29%.

Columns 4 through 6 repeat these models but include manufacturer (OEM) fixed costs per turbine, to allow for other time invariant unobservable cost difference across firms that might be correlated with sales. Comparing these estimates to the first set, the learning parameters (δ) are similar. However, the importance of turbine specific experience relative to other firm experience is much larger, suggesting that when a firm produces turbine j, it learns 60 times more about producing the next unit of j than it does from experience producing other turbines. Since firm experience producing other turbines is the reference category, this in turn implies that the experience of other firms (producing other turbines) is relatively larger than the first four columns. These models are quite saturated, however, and the estimated spillover parameters in the most flexible model, column 8, are very imprecise.

6 Dynamic Bidding Incentives

Although it is optimal for firms to bid their true *opportunity* costs in the turbine procurement mechanism we assume, the dynamic incentives implied by learning by doing mean that manufacturing costs and opportunity costs are not identical in this setting. As noted by Irwin and Klenow (1994) and Benkard (2004), the fact that sales today may generate cost reductions tomorrow means that rational and forward-looking manufacturers have an incentive to *underprice* recently introduced turbines. Thus, opportunity costs are the sum of true manufacturing costs, which may reflect past production experience, and anticipated learning benefits, or future cost reductions from marginal sales today ("dynamic markdowns"). To learn about the effects of past experience on manufacturing costs, we need to separately control for these dynamic markdowns.

Previous theoretical and applied research on dynamic markdowns explicitly computes them using dynamic programming techniques (Besanko et al., 2014; Benkard, 2004). While this approach is feasible for markets with mature production technologies and a small number of firms and products, it is not helpful in our setting. The wind turbine manufacturing industry has always had several active firms, and in any given year, 30 or more distinct wind turbine models are available for sale. It is not currently computationally possible to compute the equilibria of a dynamic game with this many firms and/or products, especially in a non-stationary setting like this.

Instead of attempting to solve such a complicated dynamic game in order to *exactly* characterize the dynamic markdown, we use a rational expectations method developed in Berry and Pakes (2000). The primary advantage of this approach is that rational expectations techniques provide a way to construct a noisy measure of the dynamic markdown without

explicitly solving for the equilibrium of the underlying game firms may be playing.

In this section, we derive the static component of opportunity costs, using the properties of our procurement auction assumptions (Section 6.1). Next, we derive the dynamic markdown term and show how to use the Berry-Pakes method to construct a noisy measure of it (Section 6.2). We then derive the state transition process for this industry, a necessary ingredient to the Berry-Pakes method, in Section 6.3. In Section 6.4 we describe two ways of implementing the rational expectations cash flows needed for the Berry-Pakes method. Finally, we derive our estimating equations (Section 6.5) and present the results (Section 6.6)

6.1 Optimal bidding behavior and static opportunity costs

To relate our observed bids to costs, we will assume that firms choose bids that are dynamically first-order optimal, in a sense we will formalize below. The dynamic component of this problem arises from the fact that past manufacturing experience is a key determinant of today's manufacturing costs, so sales today, which increase manufacturing experience, affect costs tomorrow. To allow for this, we define a state variable x as the cumulative vector of sales experience for all turbines: x_{jt} is the count of all sales of turbine j before period t. This state variable is common to all firms in the industry, which will allow us to account for spillovers in experience across firms, to the extent that they are empirically relevant.

During period t, the industry at state x_t supplies turbines to a set of wind farm projects $i \in W_t$, using the procurement auction mechanism described in the previous section. The aggregate sales from this process are a vector q_t, which, for each $j \in \cup_f K_{ft}$, is defined by:

$$q_{jt} = \sum_{i \in W_t} n_i \mathbb{I}[i \text{ chooses } j \text{ in } t]$$

These sales in turn update the cumulative sales vector for the next period, $x_{t+1} = x_t + q_t$. The probability distribution over q_t, which depends on the set of available projects W_t and the bids firms set b_t, is $dF(q_t \mid W_t, b_t)$, which we derive in Section 6.3.

We assume that firms choose turbine-specific bids at the start of the year, before any sales occur, and submit these fixed bids to each individual procurement auction. In a world with no dynamic incentives, this assumption is innocuous, due the Asker and Cantillon (2008) efficiency result: bids should always equal costs, and so they are common across projects by definition. However, if dynamic pricing incentives are indeed important, then as a firm accumulates experience over the course of a time period, its costs change. Moreover, the learning benefits from selling heterogeneous projects need not be homogenous. For example,

bigger projects induce larger changes in experience than smaller projects do. Thus, our assumption that firms set bids once at the start of a period, and submit the same vector of bids to every project's procurement auction is not without loss of generality.

However, we impose this assumption out of necessity. Because we do not directly observe bids, we must infer them from the demand system, in the form of turbine-by-time fixed effects, and there are limits to how aggressively we can divide the data up in order to generate finer turbine-time fixed effects. Accordingly, in our derivation of optimal bidding that follows, we will assume that firms receive all bidding-relevant information before they set bids. In this sense, our assumption that bids are constant over a time period is similar to a timing assumption in the productivity literature.

Firms choose bids in order to maximize the sum of expected profits from selling to available projects that period (the set $i \in W_t$), and the expected value of a firm-time specific continuation value function of x_{t+1}, where expectations are taken with respect to the distribution $dF(q_t \mid W_t, b_t)$. Firm f's objective is:

$$\max_{b_{ft}} \sum_{i \in W_t} n_i \mathbb{E}\pi_{ift}(b_t, x_t) + \int V_{ft}(x_t + q_t) dF(q_t \mid W_t, b_t)$$

Here, we are making explicit the idea that expected profits depend on the entire vector of bids (b_t). We also allow the expected profit function to depend on x because a firm's present manufacturing costs may depend on its past manufacturing experiences, as well as other firm's manufacturing experiences if spillovers are present.

Optimal bidding for turbine l will satisfy a first-order condition:

$$0 = \underbrace{\sum_{i \in W_t} n_i \frac{\partial}{\partial b_{lt}} \mathbb{E}\pi_{ift}(b_t, x_t)}_{\nabla\pi\text{=marginal static profits}} + \underbrace{\frac{\partial}{\partial b_{lt}} \int V_{ft}(x_t + q_t) dF(q_t \mid W_t, b_t)}_{\nabla V\text{=marginal dynamic benefits}}$$

The first component of this expression shows how marginal pricing changes affect the firm's marginal static profits. In particular, a marginal change in the bid for turbine l induces a

change in static profits $(\nabla\pi)_{lt}$:

$$
\begin{aligned}
(\nabla\pi)_{lt} &= \sum_{i\in W_t} n_i \frac{\partial}{\partial b_{lt}} \mathbb{E}\pi_{ift}(b_t, x_t) \\
&= \sum_{i\in W_t} n_i \frac{\partial}{\partial b_{lt}} \left(\log \sum_{j\in\prod_\phi K_{\phi,t}} \exp\left(R_{ijt} - b_{jt}\right) - \log \sum_{k\in\prod_{\phi\neq f} K_{\phi,t}} \exp\left(R_{ikt} - b_{kt}\right) + \sum_{j\in K_{ft}} s_{ijt}(b_{jt} - c_{jt}) \right) \\
&= \sum_{i\in W_t} n_i \left(-s_{ilt} - 0 + \sum_{j\in K_{ft}} \frac{\partial s_{ijt}}{\partial b_{lt}} (b_{jt} - c_{jt}) + s_{ilt} \right) \\
&= \sum_{i\in W_t} n_i \sum_{j\in K_{ft}} \frac{\partial s_{ijt}}{\partial b_{lt}} (b_{jt} - c_{jt})
\end{aligned}
$$

The first equality comes from our assumption that firms set the same turbine bids for all projects in the same time period. The second equality is a result of our procurement auction assumptions: expected (per-turbine) profits on plant i are a "net surplus" term, insensitive to which turbine is actually chosen, plus a share-weighted average bid-cost markup term. The third equality comes from differentiating the expected profit expression with respect to the bid on turbine l. Thanks to the logit errors, the derivative of the net surplus term is simply the negative of the probability that plant i picks turbine l. Similarly, the derivative of the bid-cost markup term is the "standard" two terms common in all discrete choice demand models: the elasticity-weighted sum of changes in margins, plus the gains from infra-marginal buyers. The final expression above includes the first of these standard terms, but not the second, because second price payment rules mean that the firm does not capture infra-marginal benefits from higher prices. As a result, marginal static expected profits are equal to the elasticity weighted average bid-cost markup, summed over all plants in the market.

6.2 Dynamic opportunity costs

To characterize the marginal effect of higher prices on future benefits, which we call the dynamic markdown, first note that since we have assumed our continuation values are functions of the present state x_t and the realized aggregate sales vector q_t, turbine price changes only affect the probability distribution of q_t, not the function $V_{ft}(\cdot)$ itself. Thus, our first simplification is:

$$
\frac{\partial}{\partial b_{lt}} \int V_{ft}(x_t + q_t) dF(q_t \mid W_t, \mathbf{b}_t) = \int V_{ft}(x_t + q_t) \frac{\partial}{\partial b_{lt}} dF(q_t \mid W_t, \mathbf{b}_t)
$$

Next, note that for any realizable value of q_t, $dF(q_t \mid W_t, b_t) > 0$, so we can re-arrange terms to get:

$$\begin{aligned}\int V_{ft}(x_t + q_t)\frac{\partial}{\partial b_{lt}} dF(q_t \mid W_t, b_t) &= \int V_{ft}(x_t + q_t)\frac{\frac{\partial}{\partial b_{lt}} dF(q_t \mid W_t, b_t)}{dF(q_t \mid W_t, b_t)} dF(q_t \mid W_t, b_t) \\ &= \mathbb{E}\left[V_{ft}(x_t + q_t)\frac{\frac{\partial}{\partial b_{lt}} dF(q_t \mid W_t, b_t)}{dF(q_t \mid W_t, b_t)} \mid W_t, b_t\right]\end{aligned}$$

This re-arrangement makes clear that the dynamic markdown, which is the gradient of expected future benefits with respect to the price of turbine l, can be expressed as the expectation of the *product* of future benefits at a given realization of q and the relative change in the probability that this value of q is realized resulting from the price change. This is the first insight highlighted in Berry and Pakes (2000).

We also employ the second key idea from Berry and Pakes (2000), and assume that firms have rational expectations about this object, conditional on equilibrium bids. We'll assume that:

$$\mathbb{E}\left[V_{ft}(x_t + q_t)\frac{\frac{\partial}{\partial b_{lt}} dF(q_t \mid W_t, b_t)}{dF(q_t \mid W_t, b_t)} \mid W_t, b_t\right] = \underbrace{V_{ft}(x_t + q_t^*)}_{\text{Realized Discounted Cashflows}} \times \underbrace{\frac{\frac{\partial}{\partial b_{lt}} dF(q_t^* \mid W_t, b_t)}{dF(q_t^* \mid W_t, b_t)}}_{\text{Berry-Pakes factor}} + \nu_{lt}$$

where q_t^* is the realized turbine sales vector, $V_{ft}(x_t + q_t^*)$ is the realized discounted cashflows for firm f following time t, and $\mathbb{E}[\nu_{lt} \mid x_t, b_t] = 0$ is a rational expectations error. This assumption is useful because it allows us to decompose the otherwise-infeasible dynamic markdown term into the sum of a product of two feasible expressions, and a rational expectations error. The term $\frac{\frac{\partial}{\partial b_{lt}} dF(q_t^* \mid W_t, b_t)}{dF(q_t^* \mid W_t, b_t)}$ is feasible because it can be computed from knowledge of the realized sales vector q_t^*, the set of projects W_t, and the demand system parameters. We explore two feasible approaches to handling $V_{ft}(x_t + q_t^*)$ in section 6.4.

6.3 State transitions and their sensitivity to bids

To finish our derivation of the firm's pricing FOC, we define $dF(q_t \mid W_t, b_t)$ and compute its gradient with respect to the price of turbine l. There are many allocations of turbines to projects that can generate a given value of q, so the probability that q is realized is the sum of the probabilities of many "feasible" allocations. For example, if there are three projects, A with a demand for 3 turbines, B with a demand for 2 turbines, and C with a demand for 1 turbine, and 2 turbine models, 1 and 2, there are two ways to attain $q_1 = 3$ and $q_2 = 3$. First, turbine 1 could sell to project A, and turbine 2 could sell to projects B and C. The

reverse could also happen: 1 sells to B and C while 2 sells to A. This idea generalizes to more than two projects and/or turbines. Recognizing this, let $M(q, W)$ represent the set of allocations of turbines to the sites in W which generate an aggregate sales vector q. An allocation $m \in M(q, W)$ is a vector of turbine choices, so that $m_i = l$ for the chosen turbine l. Let p_m be the unconditional probability of allocation m:

$$p_m = \prod_{i \in W} s_{i,m_i}$$

Then the probability of outcome q with plants W and bids b, denoted by $dF(q \mid W, b)$, is $\sum_{m \in M(q,W)} p_m$. Because $dF(q \mid W, b)$ depends on individual choice probabilities, which, in turn, depend on the vector of turbine bids, the gradient of $dF(q \mid W, b)$ with respect to a given bid is nonzero:

$$\begin{aligned}
\frac{\partial}{\partial b_l} dF(q \mid W, b) &= \frac{\partial}{\partial b_l} \sum_{m \in M(q,W)} \prod_{i \in W} s_{i,m_i} \\
&= \sum_{m \in M(q,W)} \frac{\partial}{\partial b_l} \prod_{i \in W} s_{i,m_i} \\
&= \sum_{m \in M(q,W)} \left(\prod_{i \in W} s_{i,m_i} \right) \sum_{i \in W} \frac{\partial}{\partial b_l} \log s_{i,m_i} \\
&= \sum_{m \in M(q,W)} \left(\prod_{i \in W} s_{i,m_i} \right) \sum_{i \in W} (s_{il} - \mathbb{I}[m_i = l]) \\
&= dF(q \mid W, b) \sum_{i \in W} s_{il} - \sum_{m \in M(q,W)} p_m \sum_{i \in W} \mathbb{I}[m_i = l] \\
&= dF(q \mid W, b) \left(\sum_{i \in W} s_{il} - \frac{\sum_{m \in M(q,W)} p_m \sum_{i \in W} \mathbb{I}[m_i = l]}{dF(q \mid W, b)} \right) \\
&= dF(q \mid W, b) \left(\sum_{i \in W} s_{il} - \mathbb{E}\left[\# \text{ of projects pick } l \mid q, W, b \right] \right)
\end{aligned}$$

A marginal change in bid b_l *proportionally* increases $dF(q \mid W, b)$ by the difference between the unconditional expected number of *projects* which pick turbine l and the same expectation, conditional on the sales vector q. Note that in the special case where $n_i = n$ for all i (homogenous project sizes), in any feasible allocation, the number of projects that pick turbine l is, by definition, equal to q_l. When this happens, $\mathbb{E}\left[\# \text{ of projects pick } l \mid q, W, b \right] = q_l$. For notational brevity, we define $N_l(q, W, B) = \mathbb{E}\left[\# \text{ of projects pick } l \mid q, W, b \right]$.

With this derivation, we can now write the Berry-Pakes factor as:

$$\frac{\frac{\partial}{\partial b_{lt}} dF(q_t^* \mid W_t, b_t)}{dF(q_t^* \mid W_t, b_t)} = \sum_{i \in W_t} s_{ilt} - N_l(q_t, W_t, b_t)$$

6.4 Discounted Cash Flows

The Berry-Pakes approach requires a measure of *realized* discounted future cashflows for each firm and time period, $V_{ft}(x_t + q_t^*)$. In our analyses below, we construct two different types of discounted cash flow measures: those that are implied by our demand and cost models, as suggested in Berry and Pakes (2000), and those derived from public company accounting data, which is available for a subset of our firms and time periods.

In the Berry-Pakes "model" approach, we combine data on realized future turbine choice probabilities and project characteristics with our model for turbine costs. Specifically, we'll compute:

$$V_{ft}^M(x_t + q_t^* \mid \theta) = \sum_{\tau=1}^{T-t} \beta(\tau) \sum_{i \in W_{t+\tau}} n_i \left(R_{i,f,\tau+t} - \sum_{j \in K_{f,t+\tau}} s_{ij} c(E_{f,t+\tau,j}, \theta) \right)$$

where $\beta(t)$ is a discounting factor, $R_{i,f,\tau+t}$ is the expected revenue per turbine for firm f at project i during a procurement auction in period $\tau + t$, and $c(\widetilde{E}_{f,t+\tau,j}, \theta)$ is the cost to firm f of manufacturing a turbine with manufacturing "experience" $E_{f,t+\tau,j}$. For all but the final year T, we'll write $\beta(t) = \beta^t$. However, to capture future cash flows that have yet to be realized but which may affect firm's perceptions of the dynamic markdown, we'll annuitize the final year's cashflows and write $\beta(T) = \frac{\beta^T}{1-\beta}$. We assume $\beta = 0.9$.

In this "modeled" approach, future discounted revenues come from our demand system estimates and future discounted costs come from the shares implied by our demand system estimates, realized experience levels, and the functional form and parameters which relate experience to manufacturing costs. Because discounted cash flows are a function of the cost parameters, this future realized cash flow construction may directly affect our estimates of the cost function.

In the Berry-Pakes "accounting" approach, we instead rely on public financial reporting which is available for a subset of the industry. Vestas, Nordex and Gamesa, as well as some of the larger Chinese manufacturers, are publicly traded companies that are either wind turbine "pure plays" or have detailed wind segment reporting. These firms directly report the revenues earned and costs incurred selling wind turbines, and we can use this information to directly construct discounted cash flows, again annuitizing the final year to capture future

cash flows which have yet to be realized:

$$V_{ft}^{A}(x_t + q_t^*) = \sum_{\tau=1}^{T-t} \beta(\tau) \left(R_{f\tau}^{\text{Accounting}} - C_{f\tau}^{\text{Accounting}} \right)$$

6.5 Estimating Equations

To recap, the firm's first order condition for dynamically optimal bidding is:

$$0 = \sum_{i \in W_t} n_i \sum_{j \in K_{ft}} \frac{\partial s_{ijt}}{\partial b_{lt}} (b_{jt} - c_{jt}) + V_{ft}(x_t + q_t^*) \left(\sum_{i \in W_t} s_{ilt} - N_l(q_t, W_t, b_t) \right) + \nu_{lt} \tag{7}$$

We can collect these terms to describe a practical estimating equation. Let b_{ft} and c_{ft} represent vectors of the bids and marginal cost functions for firm f's turbines in period t. Let the matrix of project-size weighted demand elasticities among firm f's turbines be $\Delta_{ft} = \sum_{i \in W_t} n_i \sum_{j \in K_{ft}} \nabla s_{ift}$, where s_{ift} is the vector of firm f's choice probabilities at project i. Similarly, let $N_{ft}(q_t, W_t, b_t)$ be the vector of the expected number of projects that pick f's turbine's in market realization q_t for projects W_t with the full vector of bids b_t. Finally, let $s_{ft} = \sum_{i \in W_t} s_{if}$. Then our FOC in vector form is:

$$0 = \Delta_{ft}(b_{ft} - c_{ft}) + V_{ft}(x_t + q_t^*) \left(s_{ft} - N_{ft}(q_t, W_t, b_t) \right) + \nu_{ft}$$

We can re-arrange this to express the vector of firm f's bids in terms of its costs, its dynamic incentives, and a rational expectations shock in period t:

$$b_{ft} = c_{ft} - V_{ft}(x_t + q_t^*)(\Delta_{ft})^{-1} \left(s_{ft} - N_{ft}(q_t, W_t, b + t) \right) - (\Delta_{ft})^{-1} \nu_{ft} \tag{8}$$

To specify this for a single turbine j, let ξ_{jt} be the j-th entry of the vector $(\Delta_{ft})^{-1} \left(s_{ft} - \widehat{N}_{ft}(q_t, W_t, b + t) \right)$, and let the j-th entry of $(\Delta_{ft})^{-1} \nu_{ft}$ be $\widetilde{\nu}_{jt}$. Then our estimating equation becomes:

$$b_{jt} = c_{jt} - V_{ft}(x_t + q_t^*) \xi_{jt} - \widetilde{\nu}_{jt}$$

To the extent that our discounted cash flow calculations do not capture all the firm's dynamic bidding incentives, bids and $V_{ft}(x_t + q_t^*)\xi_{jt}$ will not necessary covary one-to-one. This could happen because our discount factor is wrong (we are assuming a nominal discount rate of 10% per year), because we do not account for future fixed cost expenditures, or because the price of turbine j may only affect a portion of future discounted cash flows. To allow for this, we'll estimate a coefficient μ on the dynamic markdown terms, so that our

estimating equation becomes:

$$\widehat{b}_{jt} = \omega_0 \left(\beta_J E_{jt} + (E_{ft} - E_{jt}) + \beta_W (E_t - E_{f(j),t})\right)^{\delta} r_j^3 + \mu V_{ft}(x_t + q_t^*)\xi_{jt} - b_{0t} - \widetilde{\nu}_{jt} \qquad (9)$$

If our model were perfectly specified, we would expect to find $\mu = -1$.

We again estimate the model using generalized method of moments. In order for either Berry-Pakes approach to deliver consistent estimates of the cost function, we must make assumptions about how the unobservable terms are correlated with observables or other variables we may use as instruments. A potential challenge comes from the fact that any specification that includes a version of our rational expectations dynamic markdown term will depend on q_t^* and other variables which are not known by firms when they set bids. As a result, these variables cannot serve as valid instruments, as they will be mechanically correlated with the rational expectations shock. If we had a fully-specified data generating process for the set of projects that come to the market in each year, the power prices they face, and evolution of the market for wind turbine materials, then functions of the current value of those state variables would be valid instruments. This is the approach in Berry and Pakes (2000) and other rational expectations settings. However, we have left these details unspecified, focusing only on various components of world/firm/turbine experience as state variables in the firms' dynamic problems. Thus, we need other variables which correlate with $V_{ft}(x_t + q_t^*)\left(\sum_{i \in W_t} s_{ilt} - N_l(q_t, W_t, b_t)\right)$ but which are uncorrelated with the rational expectations shock.

We obtain these variables by recognizing that the endogenous term we have constructed, $V_{ft}(x_t + q_t^*)\left(\sum_{i \in W_t} s_{ilt} - N_l(q_t, W_t, b_t)\right)$, is a noisy measure of the *ideal* term we'd like to construct: $\mathbb{E}\left[\frac{\partial}{\partial b_{lt}} V_{ft}(x_t + q_t) \mid W_t, b_t\right]$. In fact, the relationship between the observed quantity and the ideal quantity in this setting exactly satisfies the standard classical measurement error assumption, that the difference between the true and observed quantity is independent of the truth. This means that we can construct instruments using tools from the classical measurement error literature.

We follow the results from Lewbel (1997) and create instrumental variables from (centered) higher order moments of bids, Berry-Pakes factors, and components of the discounted cash flow terms. Lewbel (1997) shows that these variables are excludable (uncorrelated with ν) and relevant (correlated with $V_{ft}(x_t + q_t^*)\left(\sum_{i \in W_t} s_{ilt} - N_l(q_t, W_t, b_t)\right)$) whenever $\mathbb{E}\left[\frac{\partial}{\partial b_{lt}} V_{ft}(x_t + q_t) \mid W_t, b_t\right]$ has nonzero skew.

6.6 Dynamic results

Table 6 re-estimates the models estimated in section 5 with the two dynamic control strategies just discussed implemented. Column 1 repeats the model with both world and turbine experience (column 3 from table 5). Column 2 adds the model-based Berry-Pakes control. Compared to the first column, the implied Spence coefficient increases to 0.29 from 0.27. World experience and turbine experience are 70% more important, relative to firm experience, in this model compared to model 1. In columns 3 and 4 we implement the accounting based dynamic control term. A panel of reliable financials was only available for five firms (Vestas, Gamesa, Nordex, Goldwind and Guodian) whose turbines collectively cover about half of our overall sample. To have a baseline, no-dynamics comparison, we first estimate the static model (column 1) on this smaller sample. The learning coefficient in this sample is slightly larger, and the world and turbine experience terms are relatively smaller, compared to the model in column 1. In column 4, we include the accounting-based dynamic control. As in column 3, this has the effect of increasing the learning coefficient, and decreasing the relative importance of world and turbine experience by statistically significant amounts.

Table 6: Dynamic Markdown Controls - Turbine Experience

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
δ	-0.443	-0.498	-0.471	-0.541	-0.501	-0.536	-0.589	-0.579
	(0.012)	(0.028)	(0.016)	(0.034)	(0.012)	(0.027)	(0.014)	(0.030)
Turbine Expc	5.721	9.912	4.842	7.752	4.843	5.785	6.100	8.594
	(1.186)	(2.376)	(1.276)	(2.780)	(1.146)	(1.089)	(1.056)	(1.666)
Industry Expc	0.050	0.086	0.071	0.123				
	(0.004)	(0.014)	(0.006)	(0.018)				
Firm Expc - Post					0.131	0.000	0.090	0.097
					(0.124)	(0.076)	(0.098)	(0.110)
Industry Expc - Pre					0.039	0.054	0.068	0.081
					(0.003)	(0.009)	(0.005)	(0.012)
Industry Expc - Post					0.108	0.123	0.104	0.107
					(0.014)	(0.016)	(0.013)	(0.015)
Spence coef $(1-2^{\delta})$	0.265	0.292	0.279	0.313	0.293	0.311	0.335	0.331
	(0.006)	(0.014)	(0.008)	(0.016)	(0.006)	(0.013)	(0.007)	(0.014)
N	938	938	475	475	938	938	475	475
Dynamics	None	BP-Model	None	BP-Accnt	None	BP-Model	None	BP-Accnt
Sample	All	All	Accnt	Accnt	All	All	Accnt	Accnt

All models include year fixed effects. Experience is measured by cumulative turbine sales. Accounting ("Accnt") samples are restricted to five firms: Vestas, Gamesa, Nordex, Goldwind and Guodian. To account for endogeneity in turbine and firm experience, all models use predicted experience terms as instruments. Robust standard errors presented in parentheses.

The last four columns in table 6 repeat this exercise for models where we distinguish between experience accumulated before and after a turbine is introduced. The direction and magnitudes of the learning rate and world spillovers changes under dynamic controls are similar to the first four columns, although these estimates are less precise. Table 7 repeats the same exercise for models with fixed costs per turbine per firm (as in columns 6 and 8 of table 5). Compared to the previous table, importance of turbine experience and industry experience increase by similar proportions. However, the impact of the dynamic controls on the estimated learning rate is now much larger, increasing the Spence coefficient by roughly 20% in all specifications. These results suggest that firms bid relatively lower during initial, low experience periods, complicating our ability to estimate the learning rate, and echoing the results of Benkard (2004).

Table 7: Dynamic Markdown Controls - Turbine Experience - Firm Fixed Effects

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
δ	-0.539	-0.688	-0.525	-0.713	-0.500	-0.636	-0.613	-0.656
	(0.018)	(0.035)	(0.035)	(0.044)	(0.018)	(0.029)	(0.027)	(0.038)
Turbine Expc	60.895	196.885	56.816	158.301	66.389	86.759	68.532	30.586
	(10.260)	(119.745)	(16.493)	(129.781)	(10.263)	(26.417)	(12.548)	(18.170)
Industry Expc	0.239	0.565	0.225	0.693				
	(0.021)	(0.320)	(0.039)	(0.523)				
Firm Expc - Post					0.000	0.000	0.000	0.192
					(0.441)	(0.472)	(0.258)	(1.577)
Industry Expc - Pre					0.179	0.226	0.202	0.792
					(0.017)	(0.065)	(0.024)	(0.455)
Industry Expc - Post					0.063	0.090	0.000	1.188
					(0.049)	(0.064)	(0.039)	(0.765)
Spence coef $(1-2^{\delta})$	0.312	0.379	0.305	0.390	0.293	0.356	0.346	0.365
	(0.009)	(0.015)	(0.017)	(0.019)	(0.009)	(0.013)	(0.012)	(0.017)
N	938	938	475	475	938	938	475	475
Dynamics	None	BP-Model	None	BP-Accnt	None	BP-Model	None	BP-Accnt
Sample	All	All	Accnt	Accnt	All	All	Accnt	Accnt

All models include year fixed effects. Experience is measured by cumulative turbine sales. Accounting ("Accnt") samples are restricted to five firms: Vestas, Gamesa, Nordex, Goldwind and Guodian. To account for endogeneity in turbine and firm experience, all models use predicted experience terms as instruments. Robust standard errors presented in parentheses.

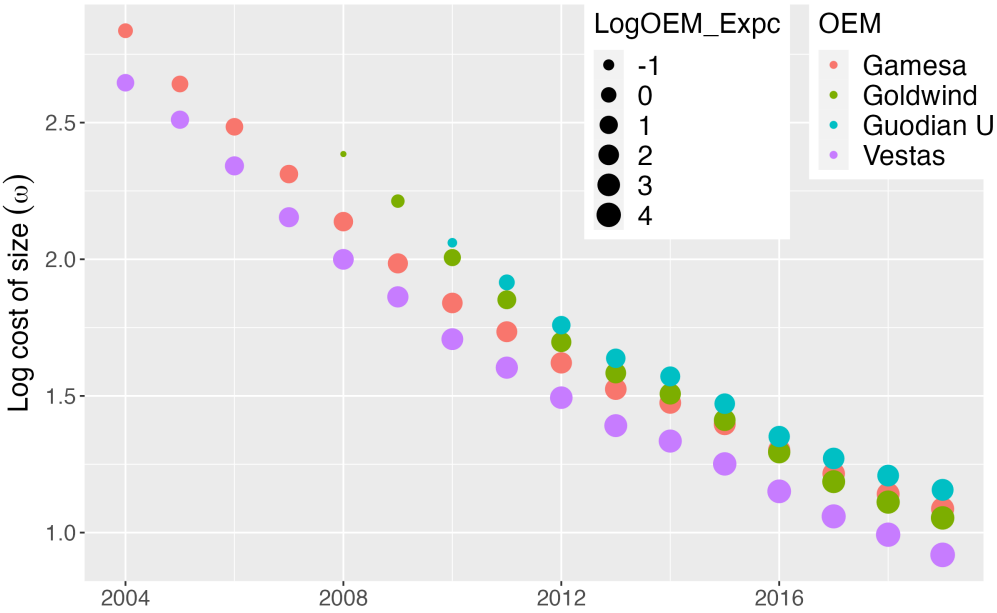
7 Discussion

We consistently find that experience today generates large cost reductions tomorrow, with a doubling of effective experience reducing costs by 24% to 33%. At the same time, while spillovers do exist, the effect of a turbine's own experience on its costs are much larger than the effect experience coming from the manufacturers other turbines. Taken together, these two facts present a puzzle: if learning effects are large, and within-turbine experience generates many times the experience benefits of other experience within the firm, why do manufactures *ever* introduce new turbines? Note that turbine lifespans are fairly short. For turbines introduced between 2000 and 2015, the median duration on the market was six years, and most sales occur in three or fewer years.

One reason why firms may introduce new turbines, even when existing turbines have

learning-driven cost advantages, is that future gains to learning for mature turbines are eventually small, due to the decreasing returns we estimate. While the *level* of a mature turbine's cost can be quite low after a few years of sales, the marginal gains to future sales are also quite small, while new turbines have an entire learning curve ahead of them. Moreover, the initial cost of a new turbine is declining over time due to spillovers within and across firms. Thus, firms may introduce new turbines in spite of temporary cost disadvantages, as a way of investing in new learning opportunities. Our cost function estimates support this idea. Figure 11 presents the initial cost of size, by year, for four large firms. On average, initial costs are falling by more than half a log point every four years. As the cost of size declines, the value proposition of introducing a new, larger turbine, with valuable future learning opportunities, becomes more attractive, even as costs of existing smaller turbines remain relatively small.

Figure 11: Example Estimated New Turbine Cost of Size

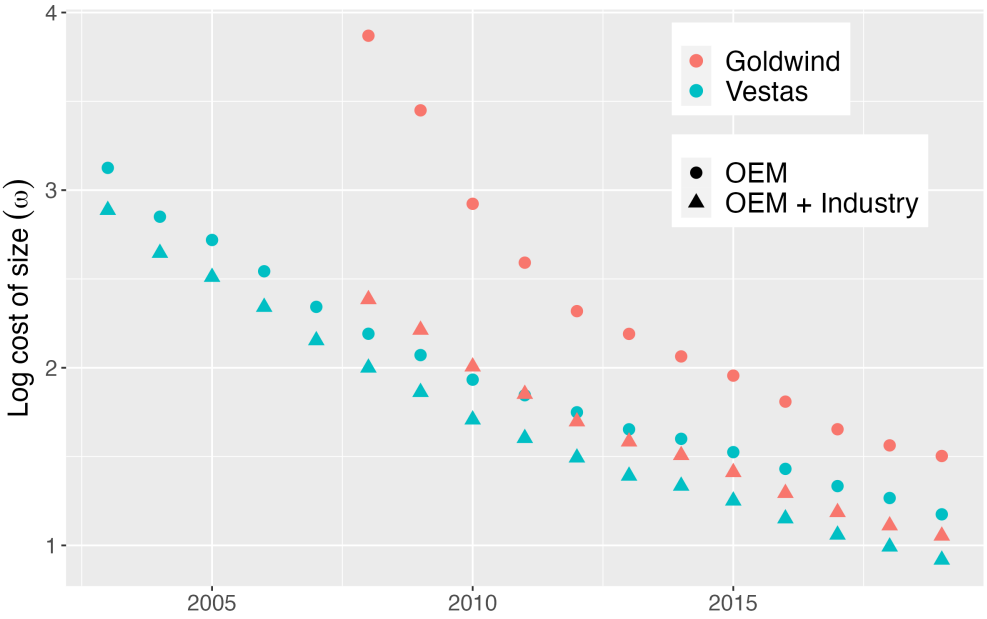


This figure presents the (log) estimated cost of size (ω_{ft}) for a newly introduced turbine ($E_{jt} = 0$), using the results from column 3 of table 5. Point sizes are proportional to cumulative manufacturer ("OEM") experience at the start of the year.

Given that spillovers are relatively small, why are new turbine costs declining so much? In figure 11, effective experience only depends on firm experience and global experience. And while other firm experience is worth only 5% as much as a firm's own experience, they are multiplied by a much larger set of turbines. So, while individual turbine spillovers are small, across all sales, they create a meaningful reduction over time. Figure 12 demonstrates

this for Vestas, the largest manufacturer at the start of the sample, and Goldwind, the largest Chinese manufacture, which not in the market until the mid 2000s. Despite having no experience when it enters in 2008, it's costs were much closer to Vesta's initially due to spillovers. By 2015, Goldwind's cost of size with world spillovers was already equivalent to Vestas' without it.

Figure 12: New Turbine Cost of Size with and without Spillovers



This figure presents the (log) estimated cost of size (ω_{ft}) for a newly introduced turbine ($E_{jt} = 0$), using the results from column 3 of table 5. In the series labeled "OEM", spillovers from other firms' experience are turned off ($\beta_W = 0$). In the models labeled "OEM + World", the estimated of β_W is taken from column 3 of table 5.

The other reason why firms introduce new turbines is that they have to in order to stay competitive. As noted in Stein (1997), firms who hope to move down the learning curve for a newly introduced product are often disappointed when their rivals quickly respond with a better product. At that point, the now-lagging firm faces a tradeoff: continue to produce its existing, and now inferior product, at ever lower costs, or abandon those learning opportunities and introduce its own newer product. Some combination of these two forces is presumably what drives the ever increasing sizes observed in Figure 5.

8 Conclusion

We estimate the extent to which learning by doing and spillovers have reduced costs in the wind turbine industry. Because neither costs nor inputs are recorded in public data,

we infer latent manufacturing costs from a structural model of turbine procurement which we estimate using the universe of wind turbine models procured by the near-universe of wind plants built in the last twenty years. To distinguish between manufacturing costs and the dynamic benefits inherent in any environment where learning by doing is present, we leverage insights from Berry and Pakes (2000) which allow us to control for dynamic pricing incentives without estimating or computing a dynamic game. We find that a doubling of manufacturing experience reduces manufacturing costs by 24 to 33 percent. Less than 10 percent of experience spills over to other turbine models produced by the same firm, and spillovers to turbines produced by other firms are on the order of 1 percent. Though relatively small, we show that, in aggregate, spillovers have generated significant cost reductions over time. These results are consistent with policymaker motivation for generously subsidizing the industry.

References

Acemoglu, Daron, David Hemous, Lint Barrage, Philippe Aghion et al., "Climate change, directed innovation, and energy transition: The long-run consequences of the shale gas revolution," in "2019 Meeting Papers" number 1302 Society for Economic Dynamics 2019.

Asker, John and Estelle Cantillon, "Properties of scoring auctions," *The RAND Journal of Economics*, 2008, *39* (1), 69–85.

Barwick, Panle Jia, Hyuk-Soo Kwon, Shanjun Li, and Nahim B Zahur, "Drive down the cost: Learning by doing and government policies in the global EV battery industry," Technical Report, National Bureau of Economic Research 2025.

Benkard, C. Lanier, "Learning and Forgetting: The Dynamics of Aircraft Production," *American Economic Review*, September 2000, *90* (4), 1034–1054.

Benkard, C Lanier, "A dynamic analysis of the market for wide-bodied commercial aircraft," *The Review of Economic Studies*, 2004, *71* (3), 581–611.

Berry, Steve and Ariel Pakes, "Estimation from the Optimality Conditions for Dynamic Controls," December 2000.

Besanko, David, Ulrich Doraszelski, and Yaroslav Kryukov, "The economics of predation: What drives pricing when there is learning-by-doing?," *American Economic Review*, 2014, *104* (3), 868–97.

_ , _ , _ , and Mark Satterthwaite, "Learning-by-doing, organizational forgetting, and industry dynamics," *Econometrica*, 2010, *78* (2), 453–508.

Betz, Albert, *Wind-Energie und ihre Ausnutzung Durch Windmühlen*, Vandenhoeck, 1926.

Coşar, A Kerem, Paul L E Grieco, and Felix Tintelnot, "Borders, Geography, and Oligopoly: Evidence from the Wind Turbine Industry," *The Review of Economics and Statistics*, July 2015, *97* (3), 623–637.

Irwin, Douglas A and Peter J Klenow, "Learning-by-Doing Spillovers in the Semiconductor Industry," *The Journal of Political Rconomy*, 1994, *102* (6), 1200–1227.

Knittel, Christopher R, "Automobiles on Steroids: Product Attribute Trade-Offs and Technological Progress in the Automobile Sector," *The American Economic Review*, December 2011, *101* (7), 3368–3399.

Lewbel, Arthur, "Constructing instruments for regressions with measurement error when no additional data are available, with an application to patents and R&D," *Econometrica: journal of the econometric society*, 1997, pp. 1201–1213.

Newell, Richard G, Adam B Jaffe, and Robert N Stavins, "The induced innovation hypothesis and energy-saving technological change," *The Quarterly Journal of Economics*, 1999, *114* (3), 941–975.

Stein, Jeremy C, "Waves of Creative Destruction: Firm-Specific Learning-by-Doing and the Dynamics of Innovation," *The Review of Economic Studies*, April 1997, *64* (2), 265–288.

Appendix A Computational details

A.0.1 Computation of N_l

How do we compute $N_l(q, W, b)$, the expected number of projects out of market W that pick turbine l, when the bids are the vector b and the realized turbine sales vector is q? As indicated above, in the case of homogenously sized projects, this object is simply the vector q. However, there is meaningful heterogeneity in size across projects which we must account for. Recall that the formal definition of this object is:

$$\begin{aligned} N_l(W_t, q_t, b_t) &= \mathbb{E}\left[\#\text{ of projects pick } l \mid q_t, b_t, W_t\right] \\ &= \frac{\sum_{m \in M(q_t, W_t)} \left(\prod_{i \in W_t} s_{i,m_i}\right) \sum_{i \in W_t} \mathbb{I}[m_i = l]}{\sum_{m \in M(q_t, W_t)} \prod_{i \in W_t} s_{i,m_i}} \end{aligned}$$

Exact computation of this requires a full enumeration of the set $M(q_t, W_t)$, which can be defined as the set of integer solutions to an under-determined (**I think??**) system of linear equations with integer coefficients. Let μ_{ij} for $i \in W_t$ and $j \in \cup_f K_{ft}$ be the entries of a matrix representing an allocation of turbines to plants. When plant i receives turbine j, $\mu_{ij} = 1$, and zero otherwise. In addition to the requirement that all entries of μ_{ij} are binary, feasible values of μ_{ij} satisfy two constraints. First, each plant chooses exactly one turbine, so for all i:

$$\sum_j \mu_{ij} = 1$$

Second, each turbine must sell in the quantities we ultimately observe in the aggregate sales vector q, so for all j:

$$\sum_i n_i \mu_{ij} = q_j$$

If the number of wind farms and turbines were sufficiently small, it would be possible to exhaustively enumerate all feasible solutions μ_{ij}, either using an integer linear programming solver, or specialized software for finding the vertices of the polyhedron defined by these equations.[4] However, with hundreds of projects per year and as many as 50 turbines in some years, complete enumeration is computationally impossible. Moreover, many feasible allocations may have vanishingly small probabilities of occurring, and as such may not contribute much to the exact value of $N_l(q, W, b)$.

In light of this, we approximate this object by using an integer linear programming solver in the "solution pool" mode.[5] We ask the solver to find the $L = 200$ "best" feasible solutions, where solution quality is the log-likelihood of the observed allocation, or $\sum_i \sum_j \mu_{ij} \log s_{ij}$.[6] Let $\widehat{M}(q, W)$ be our approximate set of solutions for realized sales q under market W. Then

[4] See, for example, http://cgm.cs.mcgill.ca/~avis/C/lrs.html

[5] Both CPLEX and Gurobi offer this option. We have used Gurobi here.

[6] This problem is NP-hard, and so instead of allowing the solver to run indefinitely, we collect the 200 best solutions available after 15 minutes of solution time per market.

we approximate $N_l(q_t, W_t, b_t)$ with:

$$\widehat{N}_l(q_t, W_t, b_t) = \frac{\sum_{m \in \widehat{M}(q_t, W_t)} \left(\prod_{i \in W_t} s_{i,m_i}\right) \sum_{i \in W_t} \mathbb{I}[m_i = l]}{\sum_{m \in \widehat{M}(q_t, W_t)} \prod_{i \in W_i} s_{i,m_i}}$$

for each l in $\cup_f K_{ft}$.

Appendix B Additional Tables and Figures

Table A.1: Demand Estimation Observations by Country

Year	AU-DK-PT	CHINA	DEU	FRA	ITA	SPAIN	SWE	U.K.	U.S.A.	Excluded	Share
2000	25	3	190	9	11	47	14	7	8	69	0.82
2001	12	0	308	3	19	44	11	7	38	89	0.83
2002	21	0	339	7	10	46	15	9	13	91	0.83
2003	37	6	281	9	31	59	30	7	39	119	0.81
2004	39	2	237	16	24	93	22	5	19	119	0.79
2005	71	16	201	40	17	62	15	16	29	164	0.74
2006	37	25	219	74	26	88	26	22	39	245	0.69
2007	31	47	187	79	34	96	43	27	58	218	0.73
2008	34	93	129	103	37	96	36	33	97	261	0.72
2009	39	187	213	98	55	86	46	35	107	291	0.75
2010	39	220	162	100	43	47	64	30	61	435	0.64
2011	19	305	184	79	31	35	48	32	87	549	0.60
2012	24	221	122	50	26	29	32	32	144	397	0.63
2013	39	200	138	49	14	11	19	67	12	207	0.73
2014	34	227	317	79	6	0	25	45	50	406	0.66
2015	34	324	344	65	11	2	18	30	61	452	0.66
2016	44	212	412	95	16	3	17	40	72	422	0.68
2017	24	164	511	123	11	2	8	58	62	365	0.73
2018	31	130	226	88	22	22	14	21	54	392	0.61
2019	9	115	107	95	18	71	15	14	40	298	0.62
2020	7	278	150	74	7	56	13	6	0	359	0.62

Table A.2: Demand Estimation Observations by Country, Total Capacity

Year	AU-DK-PT	CHINA	DEU	FRA	ITA	SPAIN	SWE	U.K.	U.S.A.	Excluded	Share
2000	120	42	1185	39	167	751	14	30	40	455	0.84
2001	51	0	2003	23	144	839	16	78	1427	755	0.86
2002	227	0	2204	43	53	937	20	79	643	692	0.86
2003	286	101	2092	56	493	1350	61	99	1562	1115	0.85
2004	480	46	1499	120	364	2649	42	110	376	1253	0.82
2005	768	359	1601	418	286	1591	15	372	2145	1773	0.81
2006	836	775	1769	794	474	2077	94	607	2534	3698	0.73
2007	551	2245	1255	847	798	2285	162	322	5278	3314	0.81
2008	659	4451	923	1237	767	2378	151	488	8275	4252	0.82
2009	762	10045	1696	1225	1227	2410	287	741	9596	6115	0.82
2010	599	13660	992	1374	1127	1053	516	560	4234	7104	0.77
2011	410	15798	1339	986	855	976	508	489	6397	10012	0.73
2012	404	10126	1354	700	775	824	645	1008	12712	8989	0.76
2013	722	10087	1603	649	393	331	578	1288	604	7404	0.69
2014	681	12074	3472	1031	135	0	950	820	4957	14183	0.63
2015	539	18731	3372	1084	193	14	723	537	8232	14435	0.70
2016	708	13456	4279	1378	320	13	491	748	8757	12716	0.70
2017	393	9747	4860	1782	250	47	160	2079	6008	12198	0.67
2018	502	7772	2491	1233	398	242	497	688	6785	11609	0.64
2019	191	6589	915	1320	490	1892	728	495	6148	13740	0.58
2020	193	17104	1296	1166	93	1568	1851	89	0	15658	0.60

B.1 Additional Results: Static Estimation

Table A.3: Learning Regressions - MW Experience

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
δ	-0.239	-0.348	-0.367	-0.335	-0.101	-0.311	-0.416	-0.391
	(0.009)	(0.009)	(0.009)	(0.009)	(0.010)	(0.012)	(0.014)	(0.014)
Turbine Expc			2.680	4.616			49.286	39.458
			(0.693)	(0.821)			(8.838)	(5.706)
Industry Expc		0.017	0.015			0.200	0.246	
		(0.002)	(0.002)			(0.028)	(0.031)	
Firm Expc - Post				0.290				0.000
				(0.103)				(0.322)
Industry Expc - Pre				0.007				0.112
				(0.002)				(0.015)
Industry Expc - Post				0.010				0.020
				(0.007)				(0.034)
Spence coef $(1-2^{\delta})$	0.153	0.214	0.225	0.207	0.067	0.194	0.250	0.237
	(0.006)	(0.005)	(0.005)	(0.005)	(0.006)	(0.007)	(0.007)	(0.007)
Year FE	X	X	X	X	X	X	X	X
OEM FE					X	X	X	X

All models estimated with GMM. To account for endogeneity in turbine and firm experience, all models use predicted experience terms as instruments. Robust standard errors presented in parentheses.

B.2 Additional Results: Dynamic Controls

Table A.4: Dynamic Markdown Controls - MW Experience

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
δ	-0.367	-0.349	-0.403	-0.439	-0.335	-0.362	-0.400	-0.450
	(0.009)	(0.024)	(0.012)	(0.027)	(0.009)	(0.024)	(0.011)	(0.024)
Turbine Expc	2.680	4.928	2.923	6.535	4.616	3.416	2.984	5.300
	(0.693)	(1.263)	(0.864)	(2.109)	(0.821)	(0.550)	(0.779)	(0.978)
Industry Expc	0.015	0.012	0.023	0.050				
	(0.002)	(0.005)	(0.004)	(0.012)				
Firm Expc - Post					0.290	0.029	0.300	0.260
					(0.103)	(0.053)	(0.105)	(0.081)
Industry Expc - Pre					0.007	0.008	0.019	0.033
					(0.002)	(0.003)	(0.004)	(0.007)
Industry Expc - Post					0.010	0.033	0.035	0.033
					(0.007)	(0.008)	(0.010)	(0.009)
Spence coef $(1-2^{\delta})$	0.225	0.215	0.243	0.262	0.207	0.222	0.242	0.268
	(0.005)	(0.013)	(0.006)	(0.014)	(0.005)	(0.013)	(0.006)	(0.012)
N	938	938	475	475	938	938	475	475
Dynamics	None	BP-Model	None	BP-Accnt	None	BP-Model	None	BP-Accnt
Sample	All	All	Accnt	Accnt	All	All	Accnt	Accnt

All models include year fixed effects. Experience is measured by cumulative MW sold. Accounting ("Accnt") samples are restricted to five firms: Vestas, Gamesa, Nordex, Goldwind and Guodian. To account for endogeneity in turbine and firm experience, all models use predicted experience terms as instruments. Robust standard errors presented in parentheses.

Table A.5: Dynamic Markdown Controls - MW Experience - Fixed Effects

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
δ	-0.416	-0.497	-0.387	-0.590	-0.391	-0.564	-0.453	-0.558
	(0.014)	(0.029)	(0.028)	(0.038)	(0.014)	(0.025)	(0.022)	(0.030)
Turbine Expc	49.286	146.532	47.741	93.526	39.457	63.171	45.224	24.685
	(8.838)	(115.138)	(15.858)	(65.858)	(5.706)	(64.610)	(9.714)	(12.817)
Industry Expc	0.246	0.554	0.158	0.346				
	(0.031)	(0.449)	(0.041)	(0.218)				
Firm Expc - Post					0.000	0.000	0.000	0.000
					(0.321)	(1.939)	(0.264)	(1.315)
Industry Expc - Pre					0.112	1.466	0.147	0.552
					(0.015)	(1.596)	(0.027)	(0.318)
Industry Expc - Post					0.020	1.414	0.000	0.934
					(0.034)	(1.504)	(0.043)	(0.609)
Spence coef $(1-2^{\delta})$	0.250	0.291	0.235	0.336	0.237	0.323	0.269	0.321
	(0.007)	(0.014)	(0.015)	(0.017)	(0.007)	(0.012)	(0.011)	(0.014)
N	938	938	475	475	938	938	475	475
Dynamics	None	BP-Model	None	BP-Accnt	None	BP-Model	None	BP-Accnt
Sample	All	All	Accnt	Accnt	All	All	Accnt	Accnt

All models include year fixed and fixed firm costs per turbine. Experience is measured by cumulative MW sold. Accounting ("Accnt") samples are restricted to five firms: Vestas, Gamesa, Nordex, Goldwind and Guodian. To account for endogeneity in turbine and firm experience, all models use predicted experience terms as instruments. Robust standard errors presented in parentheses.