Experience and the Evolution of Wind Power Project Costs in the United States

John W. Anderson*

November 2012

Abstract

I investigate empirically the relationship between accumulated experience completing wind power projects and the installed costs of wind generating capacity in the United States for the period 2001-2009. I develop a modeling framework that: (i) disentangles accumulated experience from other determinants of cost — in particular, input prices, scale economies, and technical change; and (ii) allows for alternative measures of experience and multiple channels through which experience can accumulate. For a variety of model specifications, I find evidence of small or modest firm-specific learning-by-doing but no evidence of knowledge spillovers across firms. I also find evidence consistent with firms being able to share experience with and purchase experience from competitors via joint ventures and acquisitions, respectively. Finally, I find evidence that firms' experience depreciates rapidly over time. These findings suggest that the cost-reducing benefits of experience in wind capacity installations are fully captured by the entity that undertakes these installations (and, subsequently, any acquiring entity thereof), rather than by other participants in the industry.

1 Introduction

Productivity change due to accumulated experience with a production process or technology — the phenomenon now known as *learning-by-doing* — has long been of interest to academics, managers, and policymakers.[1] In recent years, amid growing concern about climate

*Department of Economics, Stanford University. I am grateful for financial support from the Stanford Institute for Economic Policy Research (SIEPR) and the Kapnick Foundation.

[1] Alchian (1963), Hirsch (1956), and Wright (1936) were among the first to empirically investigate this type of productivity change, while Arrow (1962) was first to propose a comprehensive theoretical framework. The Boston Consulting Group (1968) later encouraged its clients to leverage such productivity change to their competitive advantage.

change and energy security, there has emerged a literature investigating whether learning-by-doing is characteristic of renewable energy technologies in general, and wind power in particular. Anecdotal evidence suggests that learning-by-doing on the part of wind power developers — the firms that design and build wind power projects — contributed to the dramatic fall in average wind power project costs in the United States from the early 1980s to the early 2000s.[2] Indeed, owing to this and other episodes of costs falling with accumulated experience, learning-by-doing has become a basis for a number of policies to promote wind and other renewables in the United States, including production and investment tax credits at the federal level and renewable portfolio standards at the state level.[3] For much of the 2000s, however, average wind power project costs in the United States increased, despite unprecedented investment in new wind generating capacity. The purpose of this paper is to investigate whether anecdotal evidence of learning-by-doing is corroborated by econometric evidence based on a much larger sample of wind power projects; in other words, does rhetoric match reality on the subject of learning-by-doing in U.S. wind?

Econometric estimation of learning-by-doing in this or any other setting is challenging for two principal reasons. First, it is necessary to define experience and explain how and to whom it accumulates. Most existing research defines experience in terms of cumulative output. I consider two alternative measures of output for U.S. wind power developers: megawatts of installed capacity and number of installed projects. Because the U.S. wind energy industry has witnessed significant technological change and has endured several boom-bust cycles, I allow for the possibility that output from the distant past counts less towards experience than does output from the recent past — i.e. I allow for the possibility that experience depreciates, as is the case in Argote et al. (1990), Benkard (2000), Kellogg (2011), Nemet (2012), and Thompson (2007).[4] Moreover, because the U.S. wind energy industry consists of many competing developers, and because I have assembled a project-level dataset, I quantify separately the accumulated experience of each individual developer. This makes it possible to distinguish between firm-specific learning-by-doing and inter-firm knowledge spillovers, which is important because these two learning effects have very different policy implications.[5] Finally, because there is a history of joint ventures and acquisitions in the U.S.

[2]According to Wiser and Bolinger (2010), average U.S. wind power project costs declined in real terms from about $4,800/kW in 1984 to about $1,300/kW in 2001.

[3]In his August 12, 2008 column, Thomas L. Friedman of the New York Times writes: "Tax credits [...] stimulate investments by many players in solar and wind so these technologies can quickly move down the learning curve and become competitive with coal and oil." In a February, 2012 interview, Minh Le of the U.S. Department of Energy states: "Renewable portfolio standards help drive down the learning curve and reduce solar energy cost in the long run."

[4]Baloff (1970) and Hirsch (1952) discuss how interruptions to production might adversely affect future productivity; Barradale (2010) discusses how unpredictability concerning the federal renewable electricity production tax credit (PTC) — the most important government incentive available to U.S. wind power projects — has caused such interruptions in the U.S. wind energy industry.

[5]If one firm's installation of wind generating capacity yields cost-reducing knowledge that is expropriable even by idle competitors, then the firm has a disincentive to invest in capacity. In this case, subsidies can

wind development business, I allow for the possibilities that developers can share experience with and purchase experience from one another.

The second challenge arises because accumulated experience is but one of many possible factors that determine costs. For instance, the increase in average U.S. wind power project costs during the 2000s is in large part attributable to higher prices for primary inputs like steel as well as technological changes like the advent of larger wind turbines (Bolinger and Wiser, 2011). Indeed, failure to account for other likely determinants of cost besides experience is a major shortcoming of much existing empirical work on learning-by-doing in wind and other renewable energy technologies (Nordhaus, 2009). It is therefore necessary to have a modeling framework that can disentangle learning from other contemporaneous determinants of cost. In this paper, I estimate minimum cost functions for installed wind generating capacity derived from an economic model of firm behavior in the U.S. wind energy industry. This approach allows me to estimate firm-specific learning-by-doing, inter-firm knowledge spillovers, the rate at which experience depreciates, and the degrees to which experience is shareable and transferable while controlling for the effects on costs of scale economies, changing input prices, and technical progress exogenous to wind power developers.

For a variety of model specifications, I find evidence consistent with small or modest firm-specific learning-by-doing: the estimation results suggest that doubling a firm's experience decreases its per-megawatt costs of installed wind generating capacity by 0.9-2.4 percent, all other things being equal. For no specification, however, do I find evidence consistent with inter-firm knowledge spillovers. Altogether, these findings suggest that the cost-reducing benefits of experience in wind capacity installations — slight as they may be — are fully captured by the entity that undertakes these installations, rather than by other industry participants. This calls into question the need for government subsidies to stimulate cost reductions in the U.S. wind energy industry.

I also find evidence consistent with the hypothesis that experience depreciates over time: the estimation results suggest that at most 62 percent of a firm's accumulated experience persists after one full year of inactivity — a finding in line with estimates produced elsewhere in the literature. Moreover, the results imply that one full year of inactivity increases a firm's per-megawatt costs of installed wind generating capacity by 1.6-2.1 percent, all other things being equal. This finding could in part explain why the largest U.S. wind power developers undertake new projects at fairly regular intervals: they may seek to prevent or at least slow the erosion of competitive advantages stemming from their comparatively large experience

serve to compensate the firm for the positive externality it bestows on its competitors. If, on the other hand, cost-reducing knowledge that results from the installation of capacity remains entirely within the firm, then there is no market failure; subsidies are not justified by a non-expropriable learning effect.

bases. At the same time, however, the finding that experience depreciates rather quickly could explain why fringe developers are able to compete for business.

Finally, evidence regarding the degrees to which firm-specific experience can be shared and transferred is inconclusive but nonetheless informative. For example, the data cannot reject the hypothesis that experience resulting from projects undertaken as joint ventures is equally as valuable as experience resulting from projects undertaken by just one firm. Likewise, the data cannot reject the hypothesis that acquired experience is a perfect substitute for organic experience — a result borne out by the fact that most acquisitions in the U.S. wind development business involve the purchase of an experienced incumbent by an inexperienced entrant.

The remainder of the paper is organized as follows: section 2 discusses anecdotal evidence of learning-by-doing in the design and construction of U.S. wind power projects, the growth of the U.S. wind energy industry, and the policies in place to support wind and other renewables. Section 3 introduces notation and discusses the unique dataset assembled for this paper. In section 4, I derive minimum cost functions for installed wind generating capacity from an optimizing model of firm behavior in the U.S. wind energy industry. In section 5, I discuss my estimation strategy and estimation results. I conclude in section 6.

2 Motivation

Figure 1 illustrates that average wind power project costs fell substantially in the United States from the early 1980s to the early 2000s, and there is much anecdotal evidence that this was due in part to learning-by-doing by wind power developers. As they accumulated design and construction experience, developers became adept at identifying sites well-suited for wind power projects — not just in terms of wind resource quality, but also proximity to transmission lines and other infrastructure.[6] Likewise, developers learned to navigate the myriad federal, state, and local regulations that govern the siting and construction of wind power projects.[7] Developers learned to optimize the logistics of transporting literally thousands of cargo loads to remote project sites and the logistics of managing complex construction operations: for instance, how best to build foundations in different types of terrain, how to optimize large networks of access roads and electrical wiring, and even

[6]Construction of new transmission is an extremely expensive undertaking — especially in the case of wind power projects, which are generally smaller than conventional power plants (e.g. gas and coal) and farther from major electricity demand centers.

[7]At just the federal level, a developer may need to obtain project permits from each of the Environmental Protection Agency, Federal Aviation Administration, Federal Communications Commission, Fish & Wildlife Service, and Army Corps of Engineers.

Figure 1: Average U.S. wind power project costs

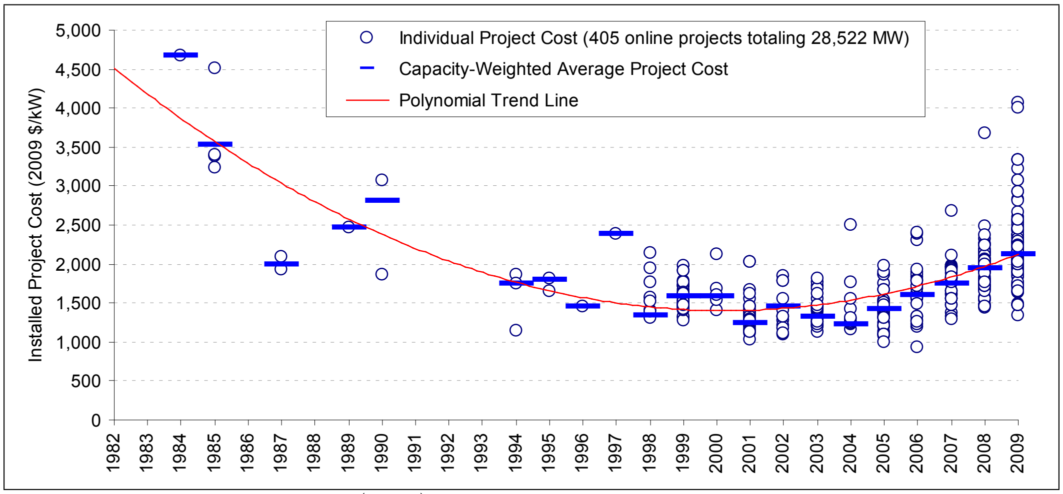

Source: Wiser and Bolinger (2010).

how best to move equipment around a project site. Developers' experience designing and building wind power projects also facilitated cost-reducing innovations upstream in the manufacturing of wind turbines: one example is the advent of modular tower sections, which are not only cheaper to manufacture but also to transport and install. Further anecdotal evidence of learning-by-doing in the design and construction of U.S. wind power projects is provided in appendix A1.

Figure 1, however, also shows that for much of the 2000s average wind power project costs actually increased in the United States; moreover, this increase in costs occurred during a period of unprecedented investment in new wind generating capacity, as can be seen in figure 2. An important question, then, is did learning-by-doing occur during the 2000s? Federal and state policies to promote wind and other renewable energy technologies are premised in part on the notion that costs will fall as experience accumulates. The federal renewable electricity production tax credit (PTC) pays generators $22 per megawatt-hour of electricity generated from eligible renewable resources; the U.S. Treasury Department estimates the PTC will cost $1.5 billion per year for each of the next ten years. Renewable portfolio standards (RPSs), state-level laws that require retailers of electricity to end users (e.g. public and investor-owned utilities) to procure from eligible renewable resources a certain percentage of their annual electricity sales, effectively guarantee wind generators higher-than-market prices for their power.[8,9] Finally, the Section 1603 grant program, part

[8]In practice, for the PTC and RPSs, "eligible renewable resources" more often than not means wind: as table A2 in appendix A2 makes clear, wind accounted for the vast majority of additions to U.S. renewable generating capacity in each of the years 2001-2011.

[9]Strictly speaking, the PTC and RPSs incentivize production of wind-generated electricity; because,

Figure 2: Annual and cumulative growth in U.S. wind generating capacity

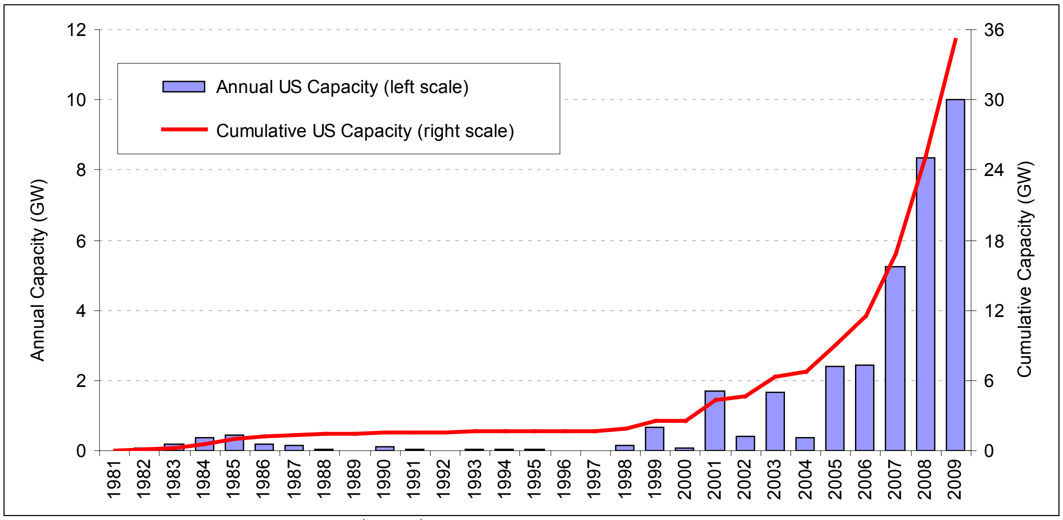

Source: Wiser and Bolinger (2010).

of the American Recovery and Reinvestment Act (ARRA) passed in the wake of the 2008 financial crisis, had awarded almost $10 billion in cash grants to wind power projects as of September, 2012. Insofar as federal and state policies are intended to address knowledge spillovers in the renewable energy sector, the results in this paper should provide some evidence as to whether this rationale has any basis in actual market outcomes.

3 Data

According to the American Wind Energy Association (AWEA), 470 utility-scale wind power projects had been completed in the United States by the end of the year 2009. For each of these projects, I obtained data from AWEA on: the project's nameplate generating capacity, q; the state, s, in which the project is situated; the developer(s), d, that designed and built the project; the nationality, n, of the manufacturer of the project's wind turbines; and the year, T, in which the project was completed. Approximately ten percent of the projects completed through 2009 were undertaken as joint ventures between two or more developers, such that d, strictly speaking, is a set. For example, $d = \{\text{BP,Clipper}\}$ for the 60 MW Silver Star wind farm in Texas, whereas $d = \{\text{Iberdrola}\}$ for the 160 MW Barton wind farm in Iowa. I also used the U.S. Energy Information Administration (EIA) Form EIA-860 database

however, there is generally no excess wind generating capacity in the United States from which to squeeze additional output, these policies strongly incentivize investment in new wind generating capacity.

Table 1: Project-level data variables and definitions

Variable	Definition
q	Nameplate generating capacity (MW)
s	State
d	Developer(s)
n	Nationality of turbine manufacturer
T	Year of completion
t	Quarter of completion
C	Total completion cost ($M)

to identify the year-quarter, t, in which each project was completed (e.g. t = 2008:Q3 for Silver Star), and to verify the accuracy of much of the AWEA data.

Project cost estimates were identified for 225 of the 339 projects completed between 2001 and 2009, and are from a variety of sources: Bloomberg New Energy Finance, business publications (in particular, Project Finance, Power Finance & Risk, and Global Power Report), state public utilities commissions' filings and testimony, corporate press releases, national and regional newspapers, and personal correspondence with wind power developers. A project's total completion cost, C, is the sum of its development, equipment purchase, and construction costs. Development costs include the costs of measuring and assessing the wind resource at a candidate project site, acquiring land usage rights, and completing environmental impact assessments. Equipment purchase costs are the costs of procuring the materials necessary to construct the wind power project, such as turbines, towers, and wires. Construction costs are the costs of erecting the wind turbines and connecting them and their attendant equipment to the electrical grid.

Table 1 summarizes the key project-level variables used throughout this paper. Appendix A2 presents annual summary statistics, while appendix A3 examines heterogeneity across developers in terms of number of projects completed, frequency with which projects are undertaken, market shares, and costs. Because reliable cost estimates could not be identified for all 339 projects completed from 2001 to 2009, appendix A4 makes a case for missing instances of cost data that occur at random.

4 Model

Econometric estimation of learning-by-doing is challenging for two main reasons: first, it is necessary to define experience and explain how and to whom it accumulates, and second, it is necessary to account for other determinants of cost besides accumulated experience. In section 4.1, I develop a framework for quantifying experience in the U.S. wind development business that: (i) allows for alternative definitions of experience; (ii) differentiates between experience internal and external to firms; (iii) allows experience to depreciate over time; and (iv) allows experience to accumulate through non-conventional channels — i.e. joint ventures and acquisitions. In section 4.2, I derive a minimum cost function for installed wind generating capacity that integrates my experience measures into a coherent econometric model from which I can estimate firm-specific learning-by-doing, inter-firm knowledge spillovers, the rate at which experience depreciates, and the degrees to which experience is shareable and transferable while controlling for the effects on costs of scale economies, changing input prices, and exogenous technical progress.

4.1 Quantifying experience

This section constructs variables Q_{d_i,t_i} and Q_{-d_i,t_i} that quantify two distinct stocks of accumulated experience available to the developer(s) of project i at the time of the project's undertaking. The former quantifies experience internal to firm(s) d_i and will be used to estimate firm-specific learning-by-doing; the latter quantifies experience external to d_i and will be used to estimate inter-firm knowledge spillovers. Experience is typically measured in terms of cumulative output, and here I consider two different measures of output for U.S. wind power developers: megawatts of installed wind generating capacity and number of installed wind power projects.[10] If learning is thought to be proportional to project size, then megawatts of installed capacity is arguably the better measure of output: a 100 MW project counts twice as much as a 50 MW project. On the other hand, if learning is thought to be invariant to project size, then number of installed projects is arguably the better measure of output: two 50 MW projects count twice as much as one 100 MW project. The remainder of this section assumes megawatts of installed capacity is the measure of output; the exposition is analogous, however, for the case where number of installed projects is the measure of output (each occurrence of q is replaced with 1).

As a first step, define firm d's *organic* experience at time t:

$$Q^O_{d,t} = \sum_{j \in J} (1-\delta)^{t-t_j-1} \cdot \lambda_{|d_j|} \cdot q_j \cdot \mathbf{1}\{d \in d_j\} \cdot \mathbf{1}\{t_j < t\} \tag{1}$$

where J is the set of all U.S. wind power projects completed through 2009, $|d_j|$ is the cardinality of the set d_j (i.e. the number of firms that developed project j — in most cases

[10] I also considered a third measure of output: number of installed wind turbines; however, this produced results very similar to those for the first measure of output (megawatts of installed capacity). Estimation results for this third measure of output/experience are thus relegated to appendix A6.

just one), and $\mathbf{1}\{\cdot\}$ is the indicator function. In keeping with recent work on organizational forgetting — the hypothesis that production experience depreciates over time — in settings as diverse as aircraft manufacturing (Benkard, 2000), oil drilling (Kellogg, 2011), shipbuilding (Argote et al., 1990; Thompson, 2007), and wind power production (Nemet, 2012), the parameter δ measures the quarterly rate of depreciation of experience, such that all other things being equal, capacity installed in the distant past counts less towards experience than does capacity installed in the recent past. That experience accumulated by U.S. wind power developers should depreciate seems plausible for at least two reasons. First, wind turbine technology has evolved considerably (see figures A10 and A11 in appendix A5) and experience with antiquated technology may not be as useful as experience with state-of-the-art technology. Second, the U.S. wind energy industry has endured several boom-bust cycles on account of the pattern of repeated expiration and short-term renewal of the PTC (Barradale, 2010). Periods of actual or anticipated unavailability of the PTC tend to result in significant labor force turnover — one of the most recognized explanations in the literature for organizational forgetting.[11]

Approximately ten percent of all wind power projects completed in the United States through 2009 were undertaken as joint ventures between two or more firms; accordingly, the λ parameters in equation (1) allow a project's relative contribution to developer d's organic experience base to depend on the number of co-developers. In my dataset, no project has more than three co-developers — i.e. $|d_j| \in \{1, 2, 3\}$ for all $j \in J$. I normalize $\lambda_1 = 1$ such that capacity completed by a single firm is the numeraire against which I measure capacity completed by joint ventures between two or three firms. I can then test a variety of hypotheses about the manner in which firms share experience. For instance, if $\lambda_2 = \lambda_3 = 1$, each partner in a joint venture is credited with having installed the total capacity of the project; alternatively, if $\lambda_2 = 1/2$ and $\lambda_3 = 1/3$, each partner is credited with having installed an equal proportion of the project's total capacity.

In addition to growing their experience bases organically as described by equation (1), it seems plausible that firms can accumulate experience by purchasing competitors. Table 2 reports ten major acquisitions in the U.S. wind development business through 2009; notably, eight of these acquisitions involved the purchase of an experienced incumbent by an inexperienced entrant.[12] We might therefore define firm d's *acquired* experience at time t

[11] Incidentally, the PTC is once again set to expire at the end of this year. Recent headlines in the New York Times include "An Expiring Tax Credit Threatens the Wind Power Industry" (Sept. 13, 2012), and "Tax Credit in Doubt, Wind Power Industry Is Withering" (Sept. 20, 2012).

[12] According to a Nov. 1, 2008 article in Windpower Monthly magazine, new entrants to the U.S. wind development business may need six or more months to get their bearings; acquiring an incumbent could be a means of short-circuiting this process. Indeed, an executive at one of the acquiring firms listed in table 2 explained to me that the target firm's experience in the U.S. wind development business was an important motivation behind the acquisition.

Table 2: Major acquisitions in the U.S. wind development business

Date	Acquired Firm	Acquiring Firm	Acquisition Marks Entry
1997:Q1	Zond	Enron	Yes
2002:Q2	Enron	GE	Yes
2003:Q1	Navitas	Gamesa	No
2005:Q1	Atlantic	PPM	No
2005:Q1	SeaWest	AES	Yes
2006:Q1	PacifiCorp	MidAmerican	Yes
2006:Q3	Padoma	NRG	Yes
2006:Q4	Orion	BP	Yes
2007:Q2	PPM	Iberdrola	Yes
2008:Q3	Catamount	Duke	Yes

as follows:

$$Q^A_{d,t} = \mu \cdot \sum_{d' \in a(d,t)} Q^O_{d',t} \tag{2}$$

where $a(d,t)$ is the set of all firms acquired by d as of time t. Organic experience transfers from d' to d — that is, from first to second owner — at rate μ. If $\mu = 1$, for instance, then acquired experience is a perfect substitute for a firm's own organic experience. Table 2, however, shows two instances in which an acquiring firm later found itself the target of an acquisition (Enron in 2002 and PPM in 2007). Accordingly, I generalize equation (2) to allow for the possibility that experience can change owners twice:

$$Q^A_{d,t} = \mu \cdot \sum_{d' \in a(d,t)} \left(Q^O_{d',t} + \mu \cdot \sum_{d'' \in a(d',t)} Q^O_{d'',t} \right) \tag{3}$$

In equation (3), organic experience transfers from d'' to d — that is, from first to third owner — at rate μ^2.

Firm d's total accumulated experience at time t is just the sum of its organic experience and its acquired experience:

$$Q_{d,t} = Q^O_{d,t} + Q^A_{d,t} \tag{4}$$

Then, for a given wind power project i, the stock of accumulated experience that is *internal* to developer(s) d_i at the time of the project's undertaking, t_i, is:

$$Q_{d_i,t_i}(\delta, \lambda_2, \lambda_3, \mu) = \lambda_{|d_i|} \cdot \sum_{d \in d_i} Q_{d,t_i} \tag{5}$$

where I have made explicit the dependency of Q_{d_i,t_i} on the parameters δ, λ_2, λ_3, and μ.

Notice that if project i has just one developer (i.e. $|d_i| = 1$) then (5) reduces to (4). If, on the other hand, project i is a joint venture between two or three developers (i.e. $|d_i| > 1$) then the interpretation of (5) hinges on the λ parameters. If $\lambda_{|d_i|} = 1$ then Q_{d_i,t_i} is the sum of the joint venture partners' individual experience bases, as given by (4); alternatively, if $\lambda_{|d_i|} = 1/|d_i|$ then Q_{d_i,t_i} is the mean of the partners' individual experience bases.

Finally, for a given project i, the stock of accumulated experience that is *external* to developer(s) d_i at time t_i is:

$$Q_{-d_i,t_i}(\delta) = \sum_{j \in J} (1-\delta)^{t-t_j-1} \cdot q_j \cdot \mathbf{1}\{d_j \cap d_i = \emptyset\} \cdot \mathbf{1}\{t_j < t_i\} \cdot \mathbf{1}\{d_j \cap a(d,t_i) = \emptyset \ \forall d \in d_i\} \tag{6}$$

where I have made explicit the dependency of Q_{-d_i,t_i} on the parameter δ. Equation (6) is simply the depreciated sum of all U.S. wind generating capacity installed prior to time t_i by developers other than d_i (or any acquirees thereof).

4.2 Technology and behavior

The production function for installed wind generating capacity is:

$$q_i = f^D\left(K_i^D, L_i^D, E_i^D, M_i^D, W_i; A_{d_i,t_i}\right) \tag{7}$$

where K_i^D, L_i^D, E_i^D, and M_i^D are, respectively, the quantities of capital, labor, energy, and materials used "downstream" in installing project i, W_i is the quantity of wind turbines installed at project i, and A_{d_i,t_i} is total factor productivity of the developer(s) of project i at the time of the project's undertaking. Wind turbines, however, are likewise produced by using capital, labor, energy, and materials, albeit "upstream" such that the production function for the wind turbines installed at project i can be written:

$$W_i = f^U\left(K_i^U, L_i^U, E_i^U, M_i^U\right) \tag{8}$$

Hence, the effective production function for the completed wind power project i is:

$$q_i = f\left(K_i^U, L_i^U, E_i^U, M_i^U, K_i^D, L_i^D, E_i^D, M_i^D; A_{d_i,t_i}\right) \tag{9}$$

Assume the production function $f(\cdot)$ is Cobb-Douglas, such that:

$$q_i = A_{d_i,t_i} \cdot \left(K_i^U\right)^{\alpha_{K^U}} \cdot \left(L_i^U\right)^{\alpha_{L^U}} \cdot \ldots \cdot \left(M_i^D\right)^{\alpha_{M^D}} \tag{10}$$

It follows that $\gamma = \alpha_{K^U} + \alpha_{L^U} + \alpha_{E^U} + \alpha_{M^U} + \alpha_{K^D} + \alpha_{L^D} + \alpha_{E^D} + \alpha_{M^D}$ measures returns to scale in the design and construction of wind power projects. Assume as well the following functional form for total factor productivity:[13]

$$A_{d_i,t_i} = \left[Q_{d_i,t_i}\left(\delta, \lambda_2, \lambda_3, \mu\right)\right]^{\beta} \left[Q_{-d_i,t_i}\left(\delta\right)\right]^{\theta} \exp\left(\phi_{T_i}^{TFP} + \psi_{s_i}^{TFP} + \epsilon_i^{TFP}\right) \tag{11}$$

In equation (11), the parameter β measures the extent to which productivity is enhanced by the stock of accumulated experience that is internal to developer(s) d_i (i.e. learning-by-doing), whereas the parameter θ measures the extent to which productivity is enhanced by the stock of accumulated experience that is external to d_i (i.e. knowledge spillovers). A fixed effect for the year in which project i was completed provides a means of controlling for technological advancements that, while exogenous to U.S. wind power developers, might nonetheless affect the costs of designing and building wind power projects. Likewise, a fixed effect for the state in which project i is situated provides a means of controlling for different policy environments that, all other things being equal, make the designing and building of wind power projects more costly in some states than in others.[14] Finally, total factor productivity depends on a mean-zero, project-specific productivity shock, ϵ_i^{TFP}, the realization of which is observed by developer(s) d_i once work on project i is underway, but unobserved by the econometrician.

Profit-maximizing wind power developers are assumed to minimize the costs of completing wind power projects of predetermined capacities given prevailing input prices. Many firms in the U.S. wind development business are publicly traded and, as such, have fiduciary obligations to maximize profits (and therefore minimize costs) for their shareholders; it seems probable that other firms in the business will likewise minimize costs in order to compete with the publicly-traded firms. Moreover, in the United States, developers generally build wind power projects to the specifications of other entities, such as independent power producers (IPPs) or electric utilities; consequently, the sizes of U.S. wind power projects can be thought of as predetermined to the developers that build them. Finally, the prices of the upstream and downstream inputs to the production function (10) are set in large markets in which wind power developers are relatively small actors — as such, these prices can be taken as exogenous to individual developers. Altogether, these assumptions lead to

[13] Equation (11) is based on Irwin and Klenow (1994), who use a similar specification in their study of learning-by-doing and knowledge spillovers in the semiconductor industry. The key differences are: (i) the experience variables in equation (11) are functions of unknown parameters (δ, λ_2, λ_3, and μ); and (ii) equation (11) includes deterministic terms (year and state fixed effects) in addition to a stochastic term.

[14] Wiser and Bolinger (2012) present evidence that average wind power project costs in the United States vary by region. In particular, Texas is regularly the lowest-cost region, whereas California and New England are regularly the highest-cost regions.

Table 3: Assumed temporal and geographical variation in input prices

Input	Price	Description
Upstream		
Capital	$P_{K_i^U} = \phi_{T_i}^{K^U}$	Completion-year FE
Labor	$P_{L_i^U} = P_{L,T_i,n_i}$	Avg. manufacturing wage in year T_i in country n_i
Energy	$P_{E_i^U} = P_{E,T_i,n_i}$	Avg. crude oil spot price in year T_i in country n_i
Materials	$P_{M_i^U} = \phi_{T_i}^{M^U} + \pi_{n_i}^{M^U}$	Completion-year and turbine country-of-origin FE
Downstream		
Capital	$P_{K_i^D} = \phi_{T_i}^{K^D}$	Completion-year FE
Labor	$P_{L_i^D} = P_{L,t_i,s_i}$	Avg. construction wage in quarter t_i in state s_i
Energy	$P_{E_i^D} = P_{E,t_i,s_i}$	Avg. gasoline price in quarter t_i in state s_i
Materials	$P_{M_i^D} = \phi_{T_i}^{M^D} + \psi_{s_i}^{M^D}$	Completion-year and state FE

the following cost minimization problem for each wind power project i:

$$\min_{K_i^U, L_i^U, \ldots, M_i^D} \quad P_{K_i^U} \cdot K_i^U + P_{L_i^U} \cdot L_i^U + \ldots + P_{M_i^D} \cdot M_i^D \quad \text{s.t.} \quad q_i \leq f\left(K_i^U, L_i^U, \ldots, M_i^D; A_{d_i,t_i}\right) \tag{12}$$

Table 3 summarizes the assumptions I make about temporal and geographical variation in input prices for purposes of solving the cost minimization problem (12). The prices of upstream and downstream capital are assumed to vary only over time — I capture this variation with completion-year fixed effects. Likewise, materials prices vary over time on account of changing supply and demand conditions, but different prices at different locations should merely reflect different (time-invariant) costs of transportation — otherwise there would exist opportunities for arbitrage. Accordingly, I capture variation in upstream materials prices with completion-year and turbine country-of-origin fixed effects, and variation in downstream materials prices with completion-year and state fixed effects. Upstream labor and energy prices are assumed to vary by year and by country: I use manufacturing wage data from the U.S. Bureau of Labor Statistics (BLS) to measure the former and crude oil spot price data from the U.S. Energy Information Administration (EIA) to measure the latter. Downstream labor and energy prices are assumed to vary by quarter and by state:

I measure the former using construction wage data from BLS and the latter using gasoline price data from EIA.

Under the stated assumptions concerning variation in input prices, the solution to (12) yields the following cost function:

$$
\begin{aligned}
\log C_i = & \frac{\alpha_{L^U}}{\gamma} \log P_{L,T_i,n_i} + \frac{\alpha_{E^U}}{\gamma} \log P_{E,T_i,n_i} + \frac{\alpha_{L^D}}{\gamma} \log P_{L,t_i,s_i} + \frac{\alpha_{E^D}}{\gamma} \log P_{E,t_i,s_i} \\
& + \frac{1}{\gamma} \log q_i - \frac{\beta}{\gamma} \log Q_{d_i,t_i}(\delta, \lambda_2, \lambda_3, \mu) - \frac{\theta}{\gamma} \log Q_{-d_i,t_i}(\delta) \\
& + \phi_{T_i} + \psi_{s_i} + \pi_{n_i} + \epsilon_i
\end{aligned}
\tag{13}
$$

The fixed effects ϕ_{T_i} and ψ_{s_i} in equation (13) now reflect both variation in input prices and (exogenous) variation in total factor productivity — consequently, it is not possible to separately identify the effects on cost of certain input prices, exogenous technical progress, and time-invariant state characteristics. More importantly for purposes of this paper, it is possible to identify from equation (13) firm-specific learning-by-doing, inter-firm knowledge spillovers, the rate at which experience depreciates, and the degrees to which experience is shareable and transferable while *controlling* for the effects on cost of changing input prices and technical progress exogenous to wind power developers. The goal of the next section is to estimate the parameters of equation (13).

5 Estimation

5.1 Estimation strategy

The assumption that capacity q_i is predetermined when developer(s) d_i undertakes project i means q_i and ϵ_i are uncorrelated in the cost function (13), such that the parameters in (13) are consistently estimated by a least squares estimation procedure. I believe this assumption is in keeping with the manner in which most wind power projects are completed in the United States. Before construction of a wind power project begins, the project's owner (an IPP, for instance) typically negotiates a long-term, fixed-price power purchase agreement (PPA) with an electricity retailer; the revenue stream guaranteed by this PPA allows the owner to secure financing for the project from a commercial or investment bank.[15] The owner then hires a wind power developer to design and construct the project with sufficient generating capacity for the owner to meet its contractual obligations to the retailer. In preparation

[15] Barradale (2010), for instance, shows that long-term PPAs were the dominant offtake arrangement for U.S. wind power projects completed in the 2000s.

for construction of the project, orders are placed for the necessary wind turbines and their attendant equipment. The revelation at this point that the project will be either more or less costly to complete than anticipated has no bearing on the quantity of capacity the developer must install: the owner still requires the previously decided-upon quantity of capacity to fulfill its PPA obligations, and it may be costly to the developer to cancel or alter an outstanding order for wind turbines. So, the productivity shock ϵ_i affects the completion costs of project i, and hence the profits of developer(s) d_i, but does not affect the capacity q_i of project i — i.e. q_i and ϵ_i are uncorrelated in equation (13).

I use nonlinear least squares (NLS) to estimate the parameters of the cost function (13). Write equation (13) compactly as follows:

$$\log C_i = h\left(\mathbf{x}_i, \boldsymbol{\xi}\right) + \epsilon_i \tag{14}$$

where $\boldsymbol{\xi} = \left(\alpha_{L^U}, \alpha_{E^U}, \alpha_{L^D}, \alpha_{E^D}, \gamma, \beta, \theta, \delta, \lambda_2, \lambda_3, \mu, \boldsymbol{\phi}', \boldsymbol{\psi}', \boldsymbol{\pi}'\right)'$ is the vector of parameters — including the year, state, and country fixed effects — to be jointly estimated and $\mathbf{x}_i$ is the data used to construct the i^{th} observation. I wish to find the estimate $\widehat{\boldsymbol{\xi}}$ of $\boldsymbol{\xi}$ that minimizes the sum of the squared residuals:

$$\text{SSR}\left(\boldsymbol{\xi}\right) = \sum_{i=1}^{N} \left[\log C_i - h\left(\mathbf{x}_i, \boldsymbol{\xi}\right)\right]^2 \tag{15}$$

Because I can compute analytically the gradient $\partial \text{SSR}\left(\boldsymbol{\xi}\right)/\partial \boldsymbol{\xi}'$, I use a quasi-Newton algorithm to search for a solution to the above minimization problem, subject to fixed bounds on the parameters (such that the final solution is interior).

Let $\widehat{\mathbf{X}}$ be the matrix with i^{th} row $\partial h\left(\mathbf{x}_i, \widehat{\boldsymbol{\xi}}\right)/\partial \boldsymbol{\xi}'$. Once a solution has been identified, I compute a heteroskedasticity-consistent estimate of the covariance matrix of $\widehat{\boldsymbol{\xi}}$ as follows:[16]

$$\text{Var}\left(\widehat{\boldsymbol{\xi}}\right) = \left(\widehat{\mathbf{X}}'\widehat{\mathbf{X}}\right)^{-1} \widehat{\mathbf{X}}' \,\widehat{\boldsymbol{\Omega}}\, \widehat{\mathbf{X}} \left(\widehat{\mathbf{X}}'\widehat{\mathbf{X}}\right)^{-1} \tag{16}$$

where

$$\widehat{\boldsymbol{\Omega}} = \text{diag}\left(\widehat{\omega}_1, \ldots, \widehat{\omega}_N\right) \tag{17}$$

and

$$\widehat{\omega}_i = \frac{N}{N-k} \left[\log C_i - h\left(\mathbf{x}_i, \widehat{\boldsymbol{\xi}}\right)\right]^2 \tag{18}$$

N is the number of observations on (14), and k is the number of parameters estimated (i.e. $k = \dim\left(\boldsymbol{\xi}\right)$).

[16] See, for instance, chapter 16 of Davidson and MacKinnon (1993).

5.2 Estimation results

Table 4 presents the results of NLS estimation of the cost function (13) for the case where megawatts of installed wind generating capacity is the measure of cumulative output (model 1); table 5 does likewise for the case where number of installed wind power projects is the measure of cumulative output (model 2). For each of models 1 and 2, I estimate seven model variants. The first model variant is the most general: it uses all available data and imposes no restrictions on equation (13). The second model variant, for reasons I will make clear later, excludes from the analysis all wind power projects with three developers. The third through seventh model variants impose on the parameters λ_2 and μ restrictions of economic interest that cannot be rejected by the estimation results for the second model variant. In all cases, reported standard errors are heteroskedasticity-consistent.

Point estimates of α_{L^U}, α_{E^U}, α_{L^D}, and α_{E^D} — i.e. the parameters in the Cobb-Douglas production function (10) associated with those inputs whose prices are explicitly modeled in (13) — have the expected positive signs in all cases except downstream energy in model 1. The point estimates suggest that of the four input prices explicitly modeled in (13), the prices of upstream energy and downstream labor have the most bearing on wind power projects' total completion costs. While α_{L^D} is the only one of the four parameters that is precisely estimated (in model 1 or 2), the hypothesis that α_{L^U}, α_{E^U}, α_{L^D}, and α_{E^D} are jointly zero is strongly rejected for all model variants.

Point estimates of the parameter γ range from 1.018 to 1.022, suggesting there are small economies of scale in the installation of wind generating capacity in the United States. These estimates imply that, all other things being equal, doubling a wind power project's nameplate generating capacity reduces the project's per-megawatt cost by 1.2-1.5 percent.[17] For all model variants, however, the hypothesis $\gamma = 1$ — i.e. constant returns to scale — cannot be rejected. Similarly, Wiser and Bolinger (2012) present evidence of weak returns to scale among small U.S. wind power projects (i.e. less than 20 MW) and constant returns to scale among larger projects.

Point estimates of β are positive and for the most part precisely estimated (no p-value is larger than 0.05 in the case of model 1 or 0.13 in the case of model 2). Point estimates of θ, on the other hand, are very imprecisely estimated — especially in the case of model 2, where the estimates actually show the wrong sign. Thus, there is evidence of firm-specific learning-by-doing, but no evidence of inter-firm knowledge spillovers.[18] Notice from equation (13)

[17] The percentage change in per-megawatt cost from doubling a project's capacity is $100 \times \left(2^{(1-\gamma)/\gamma} - 1\right)$.

[18] It is instructive perhaps to compare this qualitative finding to other studies of knowledge spillovers in electricity generation technologies: Joskow and Rose (1985) do not find evidence of spillovers in the construction of coal power plants, while Nemet (2012) and Zimmerman (1982) do find evidence of spillovers in the operation of wind power plants and the construction of nuclear power plants, respectively.

Table 4: NLS estimation results for cost function (13) (output measure 1: megawatts of installed capacity)

Model variant	Parameter										
	α_{L^U}	α_{E^U}	α_{L^D}	α_{E^D}	γ	β	θ	δ	λ_2	λ_3	μ
1-1	0.0266	0.0973	0.7256	-0.0014	1.0197	0.0153	0.0666	0.2911	1.4545	0.0002	46.873
	(0.2112)	(0.8592)	(0.1924)	(0.0817)	(0.0188)	(0.0073)	(0.0663)	(0.1606)	(3.2477)	(0.0012)	(201.08)
1-2	0.0594	0.1019	0.7206	-0.0010	1.0180	0.0168	0.0739	0.2517	1.3710		23.817
	(0.2132)	(0.8647)	(0.1935)	(0.0809)	(0.0192)	(0.0072)	(0.0737)	(0.1190)	(2.8037)		(93.312)
1-3	0.0604	0.1191	0.7201	-0.0014	1.0183	0.0170	0.0734	0.2487	1		25.298
	(0.2114)	(0.8574)	(0.1924)	(0.0803)	(0.0199)	(0.0068)	(0.0743)	(0.1138)			(96.457)
1-4	0.0628	0.1561	0.7172	-0.0015	1.0191	0.0169	0.0700	0.2501	0.5		31.587
	(0.2128)	(0.8602)	(0.1917)	(0.0807)	(0.0203)	(0.0071)	(0.0736)	(0.1181)			(123.17)
1-5	0.0602	0.0764	0.7198	-0.0052	1.0181	0.0173	0.0782	0.2417	1.4858		1
	(0.2117)	(0.8645)	(0.1933)	(0.0806)	(0.0193)	(0.0072)	(0.0760)	(0.1135)	(2.9655)		
1-6	0.0552	0.0591	0.7023	-0.0048	1.0198	0.0137	0.0606	0.3086	1.3074		0
	(0.2120)	(0.8678)	(0.1873)	(0.0808)	(0.0194)	(0.0071)	(0.0625)	(0.1559)	(3.2142)		
1-7	0.0619	0.0980	0.7183	-0.0059	1.0184	0.0177	0.0780	0.2373	1		1
	(0.2099)	(0.8572)	(0.1917)	(0.0802)	(0.0199)	(0.0068)	(0.0770)	(0.1077)			

Heteroskedasticity-consistent standard errors in parentheses.
$N = 225$ for model 1-1; $N = 223$ for models 1-2 through 1-7.

Table 5: NLS estimation results for cost function (13) (output measure 2: number of installed projects)

Model variant	Parameter										
	α_{L^U}	α_{E^U}	α_{L^D}	α_{E^D}	γ	β	θ	δ	λ_2	λ_3	μ
2-1	0.0100	0.1047	0.7200	0.0281	1.0215	0.0316	-0.0016	0.1420	1.1972	0.0275	5.2596
	(0.2134)	(0.8769)	(0.1931)	(0.0872)	(0.0195)	(0.0210)	(0.1825)	(0.1176)	(1.3616)	(0.0807)	(11.693)
2-2	0.0444	0.0804	0.6965	0.0366	1.0199	0.0320	-0.0223	0.1283	1.1566		4.0045
	(0.2145)	(0.8781)	(0.1941)	(0.0843)	(0.0197)	(0.0201)	(0.1779)	(0.0936)	(1.2982)		(8.3648)
2-3	0.0452	0.0932	0.6953	0.0362	1.0202	0.0324	-0.0234	0.1270	1		4.0409
	(0.2131)	(0.8752)	(0.1932)	(0.0835)	(0.0202)	(0.0192)	(0.1785)	(0.0899)			(8.2822)
2-4	0.0483	0.1514	0.6906	0.0355	1.0221	0.0307	-0.0346	0.1345	0.5		5.3079
	(0.2149)	(0.8781)	(0.1911)	(0.0849)	(0.0206)	(0.0201)	(0.1750)	(0.0991)			(12.245)
2-5	0.0499	0.0773	0.6880	0.0337	1.0197	0.0347	-0.0285	0.1126	1.1854		1
	(0.2135)	(0.8790)	(0.1959)	(0.0832)	(0.0197)	(0.0197)	(0.1872)	(0.0779)	(1.2375)		
2-6	0.0570	0.0919	0.6880	0.0336	1.0213	0.0282	-0.0490	0.1479	1.0539		0
	(0.2157)	(0.8791)	(0.1892)	(0.0859)	(0.0198)	(0.0196)	(0.1674)	(0.1102)	(1.3596)		
2-7	0.0512	0.0933	0.6861	0.0332	1.0199	0.0353	-0.0295	0.1110	1		1
	(0.2121)	(0.8757)	(0.1952)	(0.0825)	(0.0202)	(0.0188)	(0.1881)	(0.0749)			

Heteroskedasticity-consistent standard errors in parentheses.
$N = 225$ for model 2-1; $N = 223$ for models 2-2 through 2-7.

that the elasticity of cost with respect to firm-specific experience is $-\beta/\gamma$. Estimates of this elasticity range from -0.013 to -0.017 in the case of model 1 and from -0.028 to -0.035 in the case of model 2. All other things equal, then, doubling a firm's experience base decreases its per-megawatt costs of installed wind generating capacity by 0.9-1.2 percent in the case of model 1 and by 1.9-2.4 percent in the case of model 2.[19] Ultimately, the results suggest that cost-reducing knowledge arising from the design and construction of wind power projects — slight as it may be — remains entirely within the firm. There is no market failure, it seems, due to non-appropriability of knowledge, which calls into question the need for government subsidies to stimulate cost reductions in the design and construction of U.S. wind power projects.

Point estimates of δ are in general quite large, and although they are more precisely estimated in the case of model 1 (p-values range from 0.03 to 0.07) than in the case of model 2 (p-values range from 0.14 to 0.23), they are consistent with the hypothesis that experience depreciates over time. In the case of model 1, just 23-34 percent of a firm's accumulated experience persists after one full year of inactivity, whereas in the case of model 2 the corresponding range is 53-62 percent. For comparison, estimates elsewhere in the literature of the percentage of experience that persists after one year include: 51-61 percent in aircraft manufacturing (Benkard, 2000), 40 percent in oil drilling (Kellogg, 2011), 5-65 percent in shipbuilding (Argote et al., 1990; Thompson, 2007), and 3 percent in wind power production (Nemet, 2012). Figures 3 and 4 plot the percentage increase in per-megawatt costs of installed wind generating capacity due to prolonged inactivity for the cases of models 1 and 2, respectively. After one year of inactivity, for instance, costs are about 2 percent higher in the case of model 1 and about 1.8 percent higher in the case of model 2, all other things being equal.[20] This finding could in part explain why the largest U.S. wind power developers undertake new projects at fairly regular intervals. Table A3 in appendix A3 shows that from 2005 to 2009 the average spell of inactivity among large developers lasted just two quarters; it is possible these developers seek to prevent or at least slow the erosion of competitive advantages stemming from their comparatively large experience bases. At the same time, however, the finding that experience depreciates rather quickly could explain why fringe developers are able to compete for business — see, for instance, the market share figures A7 and A8 in appendix A3.

For each of models 1 and 2, the point estimate of λ_3 in the first model variant is very nearly zero and is imprecisely estimated. In my dataset, just six of the wind power projects completed through 2009 were undertaken as joint ventures between three firms. These six projects, however, constitute one super-project of nearly 600 MW total capacity that was

[19]The percentage change in per-megawatt cost from doubling a firm's experience base is $100 \times \left(2^{-\beta/\gamma} - 1\right)$.

[20]The percentage change in per-megawatt cost due to R quarters of inactivity is $100 \times \left((1-\delta)^{-R\beta/\gamma} - 1\right)$.

Figure 3: Predicted percentage increase in cost due to inactivity (model 1)

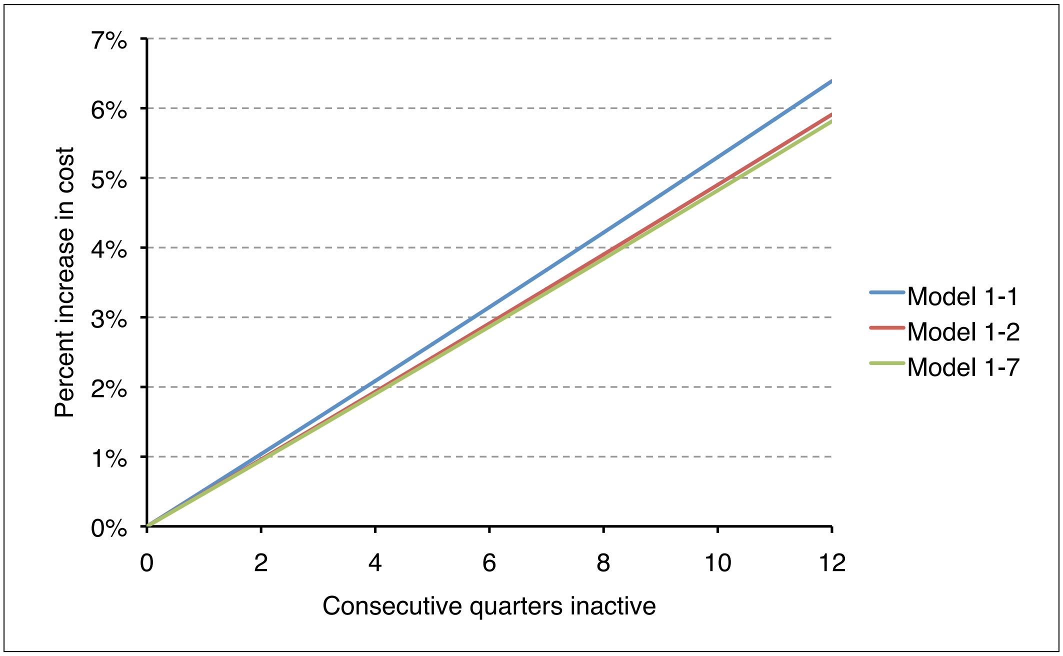

Figure 4: Predicted percentage increase in cost due to inactivity (model 2)

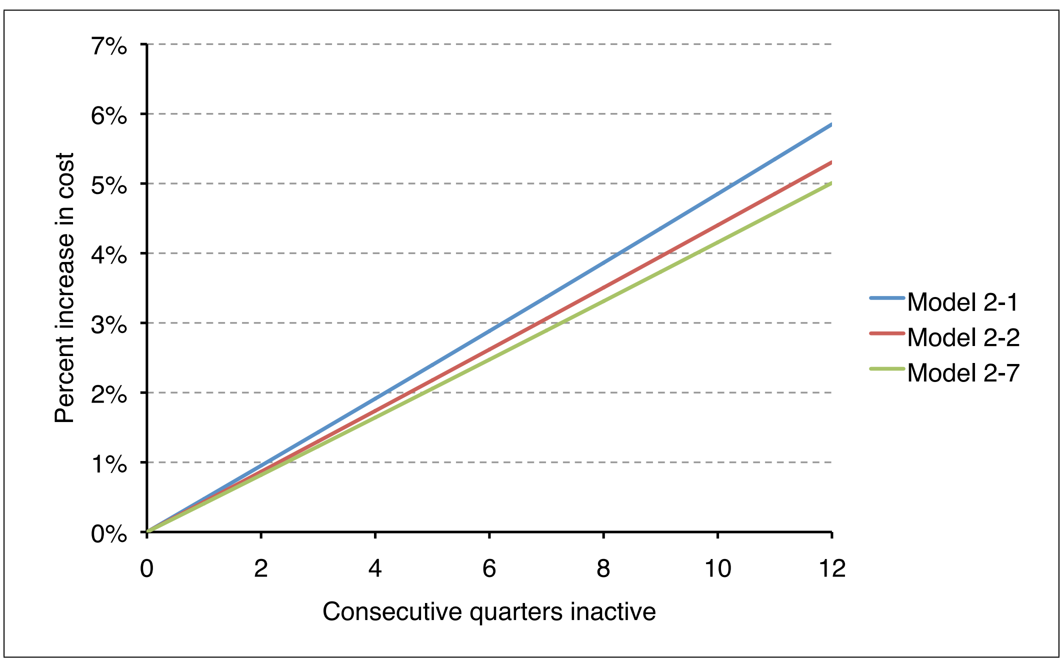

completed in phases between 2003 and 2007; moreover, it was the same three firms that partnered to develop all phases of this super-project. I therefore exclude projects with three co-developers from the remaining analysis. For each of models 1 and 2, the second model variant produces positive but imprecise estimates of λ_2 and μ: I cannot reject either of the hypotheses $\lambda_2 = 1$ or $\lambda_2 = 1/2$, nor can I reject either of the hypotheses $\mu = 1$ or $\mu = 0$. The third through sixth model variants impose restrictions on one of λ_2 or μ in an effort to improve the precision with which the other is estimated — however, this makes little difference. Ultimately, I estimate the seventh model variant, in which both λ_2 and μ are set to one. This is consistent with anecdotal evidence that one of the reasons developers undertake joint ventures is to take advantage of the experience that results from more or larger projects at a fraction of the cost. Likewise, this is consistent with the circumstantial evidence presented previously in table 2 that experience can be acquired. In any case, imposing restrictions on λ_2 and μ has no qualitative effect and little quantitative effect on the estimates of the other parameters in the model.

To conclude this section, consider a counterfactual scenario in which all wind power projects in the United States were completed by a single monopolistic developer. I use the parameter estimates in tables 4 and 5 to predict the cost to a large developer in my sample of completing a hypothetical wind power project given the developer's accumulated experience in each of the factual (competitive) and counterfactual (monopoly) cases.[21] I assume a project with the following attributes: 100 MW nameplate generating capacity, American-made wind turbines, completed in mid-2009, and situated in California. Based on the model 1 results, predicted per-megawatt costs are up to 3.6 percent lower in the monopoly case than in the competitive case; based on the model 2 results, predicted per-megawatt costs are up to 7.5 percent lower in the monopoly case than in the competitive case. Thus, granting one developer a monopoly in the U.S. wind development business — as extreme as it may sound — evidently is a means of reducing project completion costs through learning-by-doing. It is uncertain, however, whether lower project completion costs would in this case lead to lower prices to end users of wind-generated electricity.

6 Conclusion

If knowledge spillovers occur during the installation or operation of renewable generating capacity, then profit-maximizing firms will engage in these activities less than is socially desirable; public subsidies can overcome this market failure by compensating firms for the

[21]Because no point estimate of the spillover parameter, θ, is statistically significantly different from zero, I set θ equal to zero when computing predicted costs.

positive externalities their activities generate. For the particular case of the U.S. wind energy industry, however, I have found no empirical evidence of inter-firm knowledge spillovers in the design and construction of wind power projects — I have only found evidence of firm-specific learning-by-doing, which entails no externality. Thus, while federal and state policies like tax credits and renewable portfolio standards might accelerate reductions in wind power project costs, the empirical evidence presented in this paper suggests that cost reductions will occur even in the absence of government financial interventions.

I have presented evidence that experience accumulated by U.S. wind power developers depreciates over time. Ironically, the phasing out of the PTC could be beneficial to wind power developers insofar as this would reduce labor force turnover in the wind development business. The empirical evidence presented here also suggests learning-related cost reductions can be achieved through greater consolidation in the U.S. wind development business. Such consolidation could be either temporary, as in the case of joint ventures, or permanent, as in the case of acquisitions. In the former, firms reap the full experience benefits of undertaking large or numerous projects without having to bear the full costs. In the latter, not only is existing experience consolidated in a single firm, but socially-wasteful, duplicative learning is potentially avoided in the future. Owing to the number of firms active in the U.S. wind development business, it seems unlikely that greater consolidation poses any significant threat to competition.

Finally, I have argued that the assumptions that give rise to my econometric model of firm behavior in the U.S. wind energy industry are consistent with the manners in which this industry is organized and operates. Importantly, the key empirical results in this paper are qualitatively, if not always quantitatively, robust to minor changes in these assumptions. Alternative assumptions concerning functional forms and the nature of uncertainty in the model are potential areas for future research. Likewise, it would be interesting to see if similar models can be derived (if the assumptions are plausible) and estimated (if the data is available) for other technologies and countries. A better empirical understanding of the extent to which learning-by-doing is characteristic of renewable electricity generation technologies can help to ensure efficient use of public funds to support renewable energy.

Appendix

A1 Further anecdotal evidence of learning-by-doing in U.S. wind

This appendix elaborates on the anecdotal evidence of learning-by-doing in the design and construction of U.S. wind power projects presented in section 2 with regard to: (i) transportation logistics; (ii) construction logistics; and (iii) induced wind turbine innovations.

The developers with whom I have spoken have all made clear that experience plays an important role in keeping transportation costs down. Completion of a wind power project can entail hundreds or even thousands of cargo loads delivered to the project site. Delivery of just a single wind turbine, for instance, can require up to eight oversize loads: one for the nacelle, three for the blades, and four for the tower sections. Developers have learned to schedule and route deliveries to make best use of existing roads — in particular, those sufficiently flat, wide, and strong to accommodate oversize and overweight vehicles — without unduly disrupting local traffic patterns (due to road or bridge closures, for example). Moreover, they have learned to anticipate obstacles en route to a project site that could force the unloading and reloading of equipment or the complete rerouting of entire convoys of trucks. Consider the left-hand panel of figure A1: it was not left to chance that trucks hauling tower sections would ultimately fit across the bridge. Where unloading and reloading of equipment are unavoidable, however, as in the right-hand panel of figure A1, developers have learned how to do so quite effectively.

There is also anecdotal evidence that developer experience has lowered the construction costs of wind power projects. Wind turbine foundations, for instance, can require 20-40 tons of rebar and 250-450 cubic yards of concrete — see the left-hand panel of figure A2 — and can account for up to 16 percent of a project's capital costs (International Renewable Energy Agency (IRENA), 2012). Experienced developers have learned to adapt foundations to different turbine types and different ground and wind conditions so as to complete each foundation at low cost while (hopefully) avoiding the fate depicted in the right-hand panel of figure A2. Likewise, developers have learned how best to maneuver heavy equipment around a project site. For example, according to a contractor experienced in wind farm construction, the disassembling, transporting, and reassembling of a large crawler crane (e.g. the red cranes in figure A3) can take up to five days and cost as much as $70,000. Experienced developers therefore carefully sequence their construction activities so as to prevent or at least minimize such costly delays.

Finally, developers' experience designing and building wind power projects has also facilitated cost-reducing innovations upstream in the manufacturing of wind turbines. One example is the advent of modular tower sections, which as discussed in section 2 are cheaper not only to manufacture but also to transport and install. (According to IRENA (2012), towers make up about 17 percent of a wind power project's capital costs.) The left-hand panel of figure A1 shows a tower section in transit, while the left-hand panel of figure A3 shows tower sections being installed. A second example is rotors that can be assembled at ground level (left-hand panel of figure A3) and then lifted and installed in one piece (right-hand panel of figure A3).

Figure A1: Wind turbine transportation logistics

Figure A2: Wind turbine foundations (done right and done wrong)

Figure A3: Modular tower sections and ground-level rotor assembly

A2 Annual summary statistics

For the subset of 225 U.S. wind power projects completed from 2001 to 2009 and for which cost estimates are available, table A1 reports annual summary statistics for average cost of installed capacity (measured in millions of current-year dollars per megawatt) and nameplate generating capacity (measured in megawatts). As discussed elsewhere in this paper, average wind power project costs approximately doubled during the 2000s despite the completion of more and larger wind power projects than had ever previously been the case (i.e. despite potential for cost reductions due to learning-by-doing and economies of scale). Higher prices for primary inputs and the advent of larger wind turbines are two often-cited explanations for this period of rising costs (e.g. Bolinger and Wiser (2011)). Regarding the former, figure A4 plots four price indices for inputs important to the U.S. wind energy industry together with a GDP deflator; notably, all four price series increased at rates greater than the rate of overall inflation during the 2000s.[22] Regarding the latter, the hub height, rotor diameter, and capacity rating of the average wind turbine installed in the U.S. all increased significantly during the 2000s (see figure A11 in appendix A5); larger turbines are generally more costly because they require disproportionately more materials to support their greater weight and withstand severe wind forces. Table A1 is also indicative of the importance of government intervention to the growth of the U.S. wind energy industry: fewer and smaller projects were completed in 2002 and 2004 when the PTC was unavailable to new projects, whereas more and larger projects were completed during the later years of the sample when the PTC was consistently available and many more states adopted RPSs.[23,24]

Table A2 presents annual U.S. capacity additions for wind and four other renewable electricity generation technologies for the years 2001-2011. The table is presented simply to show that "wind energy" has become nearly synonymous with "renewable energy" in the United States. Not surprisingly, then, government subsidies to renewable electricity generation technologies in the U.S. accrue overwhelmingly to wind — see, for instance, EIA (2011).

A3 Developer heterogeneity

68 different wind development firms completed at least one wind power project in the United States between 2001 and 2009. Figure A5 shows the distribution of these firms by total

[22] The U.S. dollar-Euro exchange rate is included because many wind turbine components are imported from Europe.

[23] Of the wind power projects completed in 2002 and 2004, some were ineligible for the PTC (e.g. those owned by electric cooperatives or municipal utilities), while others received the credit retroactively.

[24] According to the Database of State Incentives for Renewables and Efficiency (DSIRE), the total number of states to have adopted mandatory RPSs was 5 in 2001, 11 in 2005, and 26 in 2009.

Table A1: Summary statistics, U.S. wind power project data, 2001-2009

		Average cost ($M/MW)			Capacity (MW)		
Year	Projects	Min	Med	Max	Min	Med	Max
2001	16	0.63	1.09	1.67	1.3	65.0	299.6
2002	9	0.84	1.11	1.55	3.8	40.9	160.5
2003	19	0.93	1.09	1.39	2.6	50.4	204.0
2004	4	1.09	1.19	1.37	11.6	23.7	60.0
2005	15	0.84	1.19	1.60	10.5	114.0	213.0
2006	17	1.05	1.57	2.17	7.5	100.5	231.0
2007	29	1.33	1.59	2.11	14.7	100.7	400.5
2008	54	1.30	1.92	2.67	4.5	99.0	300.3
2009	62	0.95	2.00	3.27	2.0	100.5	400.3

Figure A4: Selected U.S. price indices of relevance to wind energy industry

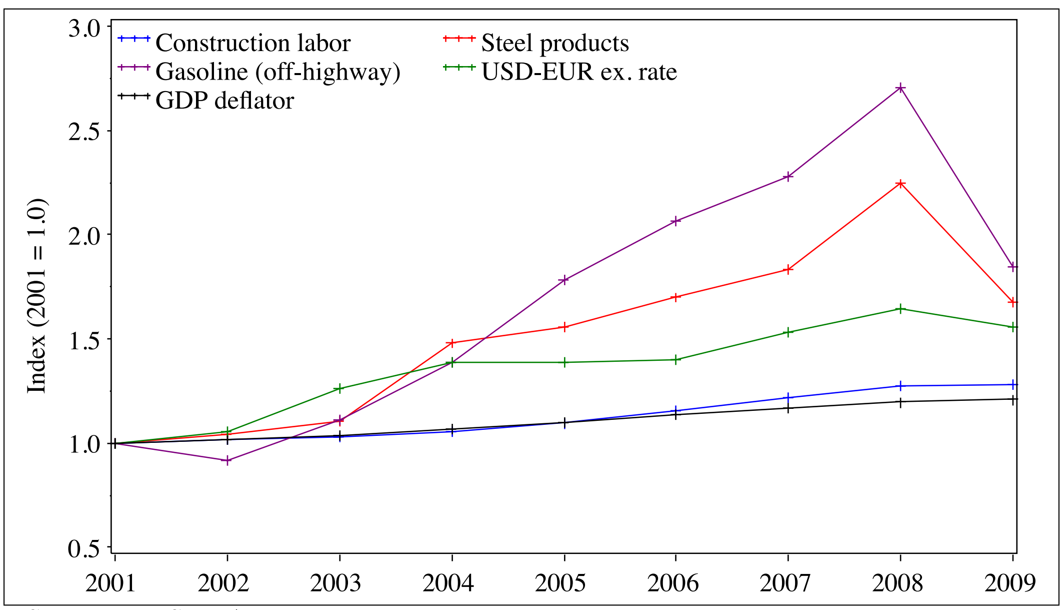

Sources: BLS; EIA.

Table A2: Annual additions to U.S. renewable generating capacity, 2001-2011 (MW)

	Technology				
Year	Biomass	Geothermal	Hydro	Solar	Wind
2001	97	0	132	5	1,403
2002	38	7	344	2	773
2003	105	0	83	0	1,609
2004	95	0	79	0	393
2005	44	30	30	2	2,156
2006	170	53	20	2	2,646
2007	114	39	20	107	5,148
2008	170	48	18	34	8,262
2009	279	213	26	89	9,766
2010	151	24	6	263	4,633
2011	222	7	151	579	6,189

Source: Form EIA-860 data; excludes CHP and non-grid-connected units.

number of projects completed from 2001 to 2009. Evidently, there are a number of large, experienced actors in this business; however, there are also many fringe competitors. Figure A6 plots average costs of installed capacity by year for eight of the largest developers in the sample. These firm-specific cost figures show the same upward trend over time as the industrywide figures presented in table A1. Admittedly, figure A6 disregards potentially important heterogeneity across projects (in terms of size and location, for instance) that might explain within-year variance in average costs across firms. Nevertheless, it is telling that the firms' per-megawatt cost rankings change from year to year — in particular, no firm is lowest-cost for a significant span of the 2001-2009 period. Perhaps for this reason, no firm has seen its market share grow to the significant detriment of other large competitors (figures A7 and A8). Thus, while the literature in industrial organization (e.g. Cabral and Riordan (1994) and Spence (1981)) recognizes that learning-by-doing can increase industry concentration through the emergence of a low-cost dominant firm, the evidence suggests this is not a concern in the present setting. Finally, it is noteworthy that among large developers, spells of inactivity are of relatively short duration. Table A3 shows that for the 2005-2009 period, rarely did more than two consecutive quarters pass without a large developer completing a new wind power project.

Figure A5: Distribution of wind developers by number of projects completed 2001-2009

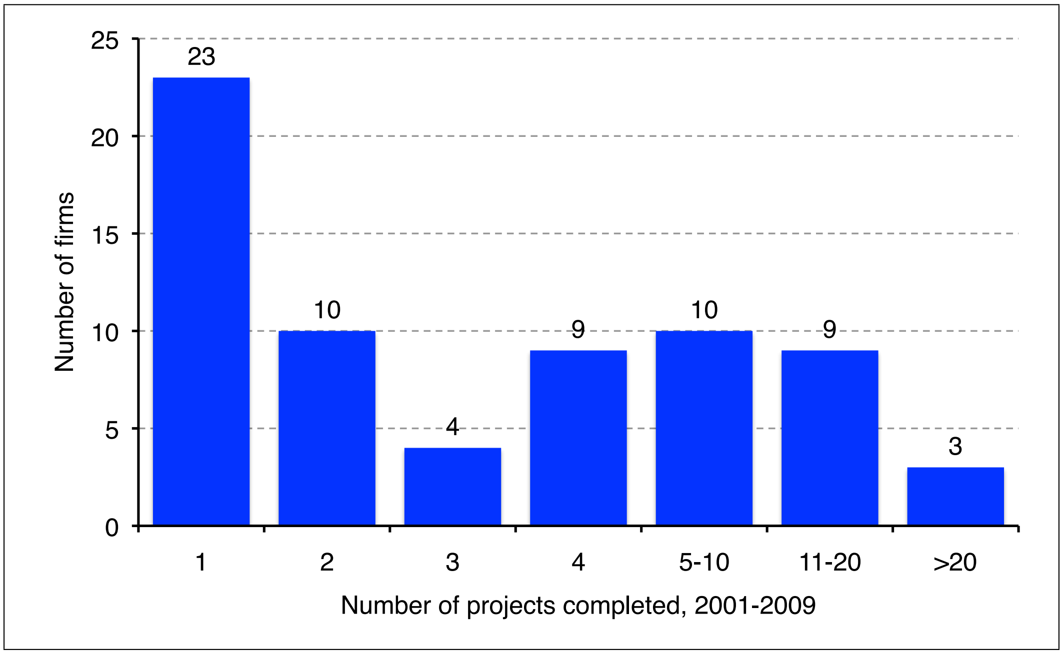

Figure A6: Average costs of installed capacity by year, selected developers

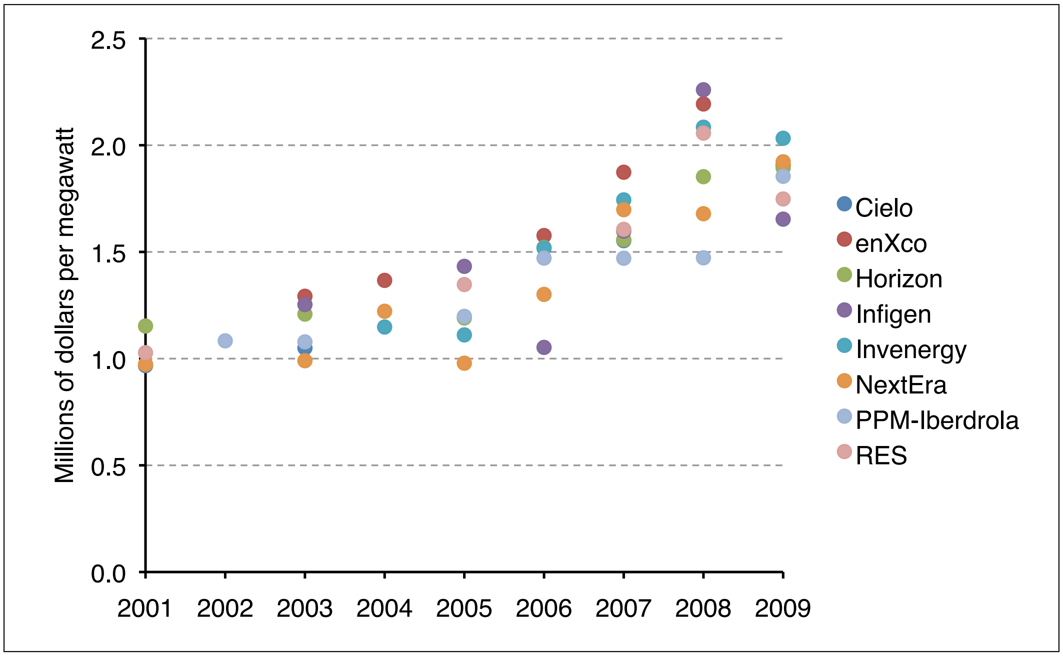

Figure A7: Market shares by year, selected developers (percent of installed MW)

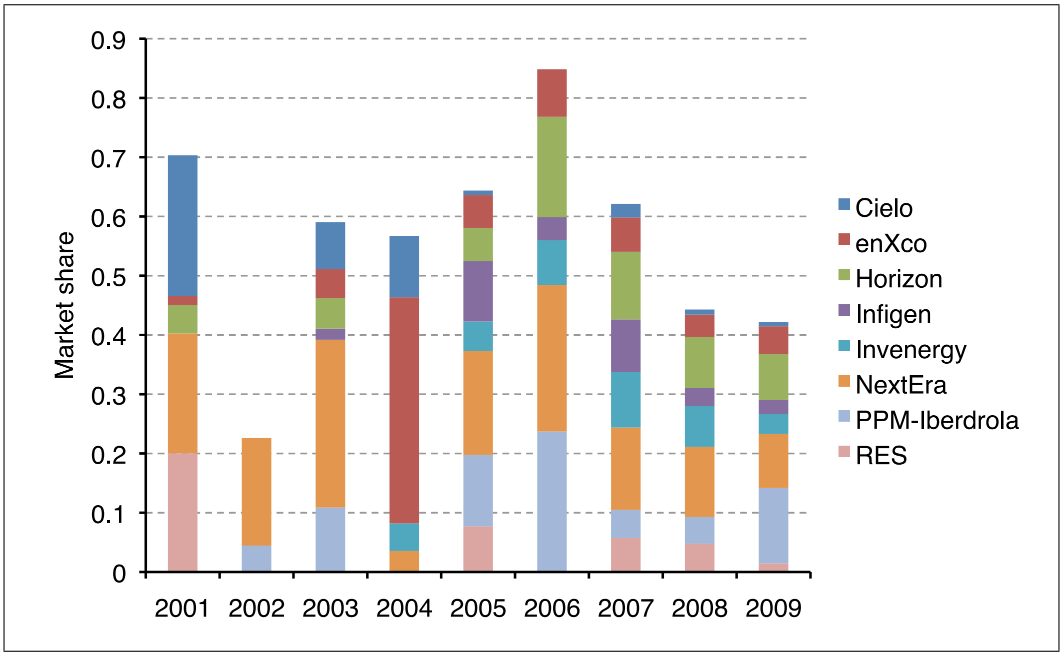

Figure A8: Market shares by year, selected developers (percent of completed projects)

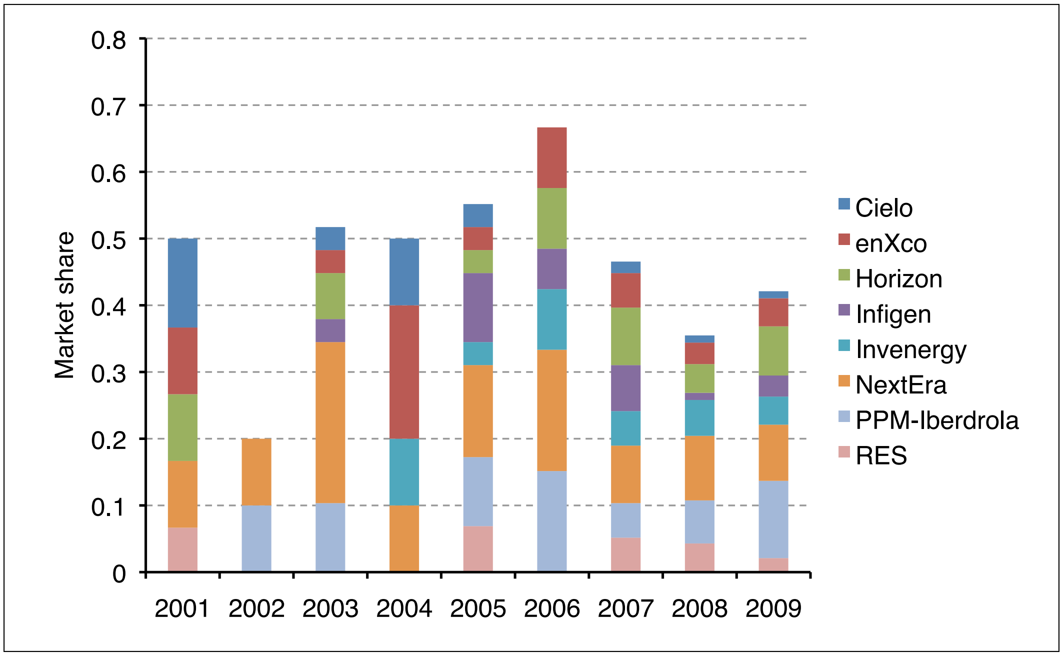

Table A3: Duration of spells of inactivity, selected developers, 2005-2009

Developer	N spells	Duration (quarters)			
		Mean	SD	Min	Max
Cielo	5	3.2	2.3	1	7
enXco	7	1.7	0.8	1	3
Horizon	5	1.8	0.8	1	3
Infigen	6	2.2	1.0	1	3
Invenergy	5	1.6	0.9	1	3
NextEra	4	1.3	0.5	1	2
PPM-Iberdrola	4	1.5	0.6	1	2
RES	4	3.0	2.2	1	6

A4 Randomness of missing cost data

Here, I consider the econometric problems that could arise on account of my having cost data for only 225 of the 339 U.S. wind power projects completed between 2001 and 2009. It is well established in the economics literature that estimation based on nonrandomly selected samples can result in biased estimates of parameters of economic interest. Although Heckman (1976, 1979) proposed a two-step estimation procedure to overcome this selection bias, implementing his procedure requires additional modeling assumptions (a second equation explaining entrance into the sample) and data (to estimate the second equation). Because neither requirement is necessarily straightforward, it behooves the researcher to weigh the evidence for and against the randomness of his sample before abandoning least squares in favor of a more complicated estimation procedure.

Table A4 compares the subsample of 225 U.S. wind power projects for which cost data is non-missing to the subsample of 114 projects for which cost data is missing. If the proportion of projects sharing a particular attribute *within* each subsample does not differ significantly *across* the two subsamples, then there is evidence that the 114 instances of missing cost data occur at random. From the table, it is apparent that the biggest differences between the two subsamples concern project size and geography. First, projects with non-missing cost data tend to be larger (in terms of total installed capacity) than projects with missing cost data — perhaps because large projects attract public scrutiny and require that developers disclose significant information. Second, among projects with non-missing cost data, greater proportions are located in the NPCC and RFC reliability regions, and smaller proportions are located in the MRO and TRE regions, than is the case among projects with missing cost data. (Figure A9 presents a map of NERC reliability regions in the U.S.) This could reflect attitudes or policies towards the disclosure of project information that differ across states or regions. For both capacity and NERC region, a multinomial test of equality of

Figure A9: North American Electric Reliability Corporation (NERC) U.S. region map

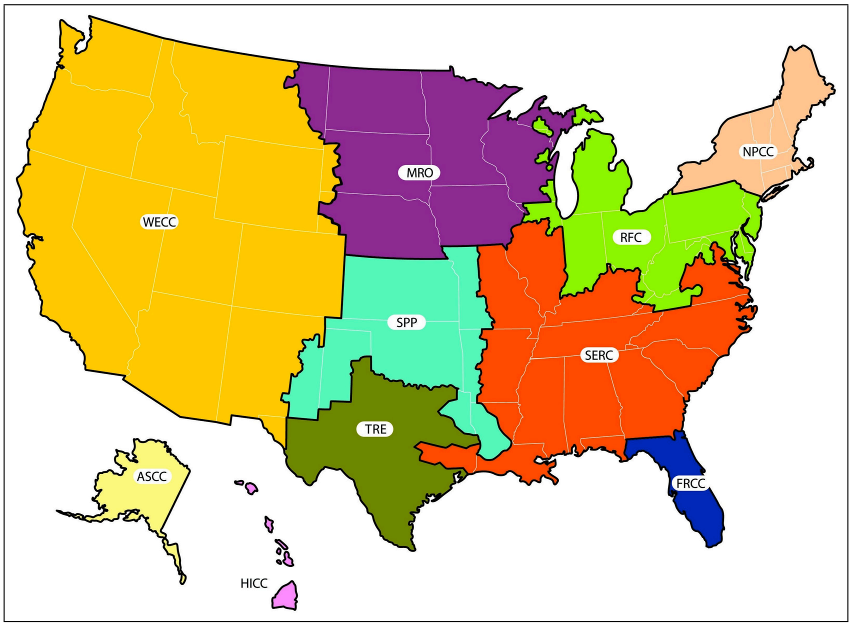

proportions rejects the null hypothesis that within-subsample proportions are jointly equal across the two subsamples, although binomial tests of equality of proportions in many cases fail to reject pairwise equality across the two subsamples.[25]

With respect to the remaining attributes in table A4, the differences between the two subsamples of wind power projects are much less pronounced. The proportion of projects completed in a given year or quarter does not appear to vary significantly across the two subsamples. Moreover, projects completed by multiple or foreign developers, and projects owned by independent power producers (IPPs), make up only slightly greater proportions of the subsample of projects with missing cost data than the subsample of projects with non-missing cost data — possibly because such projects are subject to less onerous disclosure requirements. No joint or pairwise test rejects the null hypothesis of within-subsample proportions that are equal across the two subsamples. Table A5 presents an additional comparison of project sizes across the two subsamples. Consistent with table A4, mean project capacity is in most years greater for projects with non-missing cost data than for projects with missing cost data; for all years except 2001, however, pairwise t-tests fail to reject the hypothesis that mean project capacity is equal across the two subsamples. Likewise, for all years except 2001, mean turbine rating is not statistically significantly different across the

[25]For joint hypothesis tests, the p-values in table A4 are based on the chi-squared distribution. For pairwise hypothesis tests, the p-values are based on the normal approximation to the binomial distribution.

Table A4: Attributes of wind power projects with and without cost data

Attribute	Non-missing cost data		Missing cost data		p-value*
	Count	Percent	Count	Percent	
Capacity (MW)					**0.0193**
$q \in (0, 25]$	34	15.1%	32	28.1%	0.0074
$q \in (25, 50]$	34	15.1%	14	12.3%	0.4671
$q \in (50, 100]$	54	24.0%	33	28.9%	0.3333
$q \in (100, 150]$	61	27.1%	20	17.5%	0.0390
$q \in (150, 500]$	42	18.7%	15	13.2%	0.1786
Completion year					**0.2256**
2001	16	7.1%	7	6.1%	0.7313
2002	9	4.0%	1	0.9%	0.0469
2003	19	8.4%	4	3.5%	0.0512
2004	4	1.8%	4	3.5%	0.3711
2005	15	6.7%	7	6.1%	0.8507
2006	17	7.6%	13	11.4%	0.2660
2007	29	12.9%	18	15.8%	0.4772
2008	54	24.0%	31	27.2%	0.5270
2009	62	27.6%	29	25.4%	0.6751
Completion quarter					**0.4254**
Q1	44	19.6%	17	14.9%	0.2754
Q2	37	16.4%	26	22.8%	0.1705
Q3	32	14.2%	18	15.8%	0.7046
Q4	112	49.8%	53	46.5%	0.5668
NERC region					**0.0001**
ASCC	1	0.4%	0	0.0%	0.3162
HICC	1	0.4%	2	1.8%	0.3163
MRO	51	22.7%	39	34.2%	0.0278
NPCC	15	6.7%	2	1.8%	0.0175
RFC	35	15.6%	4	3.5%	0.0001
SERC	6	2.7%	0	0.0%	0.0130
SPP	20	8.9%	9	7.9%	0.7530
TRE	32	14.2%	28	24.6%	0.0264
WECC	64	28.4%	30	26.3%	0.6767
Industry sector					**0.3746**
Electric Utility	42	18.7%	17	14.9%	—
IPP	183	81.3%	97	85.1%	—
Multiple developers					**0.4431**
No	195	86.7%	102	89.5%	—
Yes	30	13.3%	12	10.5%	—
Foreign developer(s)					**0.4349**
No	167	74.2%	80	70.2%	—
Yes	58	25.8%	34	29.8%	—

* Test of equality of proportions: boldface for joint test, lightface for pairwise test.

Table A5: Mean capacity and turbine rating of projects with and without cost data, by year

	Capacity (MW)			Turbine rating (MW)		
Year	Non-missing cost data	Missing cost data	p-value*	Non-missing cost data	Missing cost data	p-value*
2001	87.8	18.0	0.0097	1.06	0.68	0.0014
2002	48.9	97.7	0.3539	1.22	0.66	0.1997
2003	72.7	49.5	0.4745	1.41	1.44	0.9280
2004	29.8	59.7	0.4724	1.28	1.45	0.5983
2005	102.0	65.5	0.1668	1.50	1.29	0.1335
2006	95.4	89.3	0.8367	1.68	1.54	0.4099
2007	124.7	89.1	0.1240	1.76	1.86	0.4941
2008	105.4	88.1	0.2624	1.82	1.78	0.6594
2009	108.3	86.8	0.1437	1.85	1.84	0.9005

* Pairwise t-test of equality of means.

two subsamples. On balance, I believe the evidence tilts in the direction of cost data that is missing at random; as such, I use a least squares estimation procedure in section 5 rather than a more complicated two-step procedure.

A5 Installed cost of capacity vs. levelized cost of energy (LCOE)

In studying learning-by-doing and the evolution of costs in this paper, I have focused on installed costs of wind generating capacity; it could be argued, however, that what really matters is the cost of wind-generated electricity. For a given wind power project, define the levelized cost of energy, $LCOE$, to be the constant per-megawatt-hour cost that solves the following equation:

$$\sum_{t=1}^{T} \frac{LCOE \cdot E_t}{(1+r)^t} = \sum_{t=0}^{T} \frac{C_t}{(1+r)^t} \tag{A1}$$

where E_t is electricity generated by the project in period t, C_t is project costs incurred in period t, T is the useful life of the project, and r is the discount rate. The sum on the right-hand side of equation (A1) begins at $t = 0$ because installed costs, C_0, are incurred before the project produces any electricity. Operating and maintenance (O&M) costs, $C_1, \ldots, C_T$, are incurred in each production period, but pale in comparison to C_0. From equation (A1) it is apparent that $LCOE$ is the cost of electricity at which the project's owners will just break even on their investment.

Figure A10 shows that wind turbine technology has evolved considerably over the past several decades — notably, average hub height, rotor diameter, and turbine capacity rating have all increased (figure A11 depicts these increases for the period 2000-2009). On account

Figure A10: Antiquated vs. state-of-the-art wind turbine technology

Figure A11: Average U.S. wind turbine hub height, rotor diameter, and capacity rating

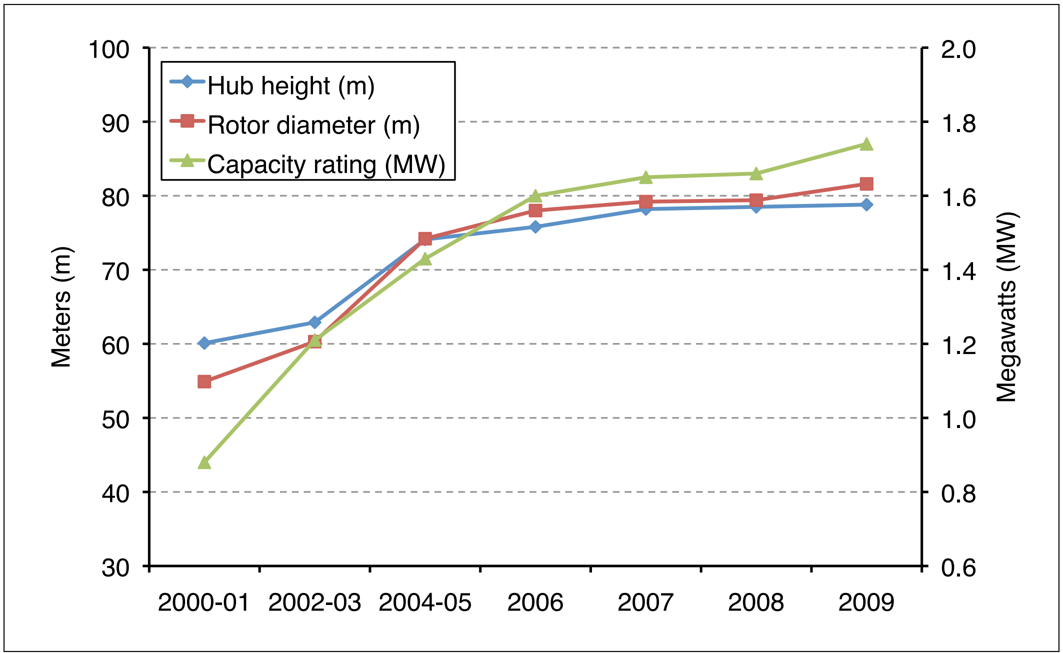

Source: Wiser and Bolinger (2010).

of such technological advances, newer wind power projects ought to be capable of generating more electricity than older wind power projects with the same amount of total installed capacity, all other things being equal. Accordingly, increasing costs of installed wind generating capacity, insofar as they result from deployment of larger wind turbines (see appendix A2), are not necessarily indicative of increasing costs of wind-generated electricity. In fact, equation (A1) shows that an increase in C_0, if accompanied by sufficiently large increases in $E_1, \ldots, E_T$, can actually decrease the levelized cost of wind-generated electricity.

As it happens, however, all other things were *not* equal over the course of the 2000s: as figure A12 shows, wind resource quality at projects completed in the late 2000s was, on average, poorer than wind resource quality at projects completed in the early 2000s.[26] Indeed, the turbine innovations discussed in the previous paragraph were in large part intended to improve electrical output of projects built at low wind speed sites (Wiser and Bolinger, 2012). The net effect of progressively larger wind turbines installed at progressively lower wind speed sites was project capacity factors that, on average, were stagnant for most of the 2000s, as shown in figure A13.[27] Because projects completed in the late 2000s do not appear to systematically outperform projects completed in the early 2000s, increasing costs of installed wind generating capacity over the course of the 2000s (as depicted in figure 1 in section 2 and table A1 in appendix A2) are likely indicative of increasing costs of wind-generated electricity. In other words, an empirical analysis of learning-by-doing based on wind power projects' levelized costs (were such data available) would likely arrive at the same conclusions found in this paper.

A6 Additional estimation results

Table A6 presents NLS estimation results for seven model variants of the cost function (13) in which number of installed turbines is the measure of output/experience; the results are very similar to those presented in table 4 for the case where megawatts of installed capacity is the measure of output.

[26] This is not to say that all high-wind-resource-quality sites in the United States have already been developed; rather, constraints such as inadequate access to transmission presently make development of many otherwise high-quality sites uneconomical.

[27] The capacity factor of a power plant is the ratio of the plant's actual electrical output over a period of time to its potential electrical output were it to operate at full capacity for the same period of time.

Figure A12: Index of average U.S. wind resource quality at 80 meters by completion period

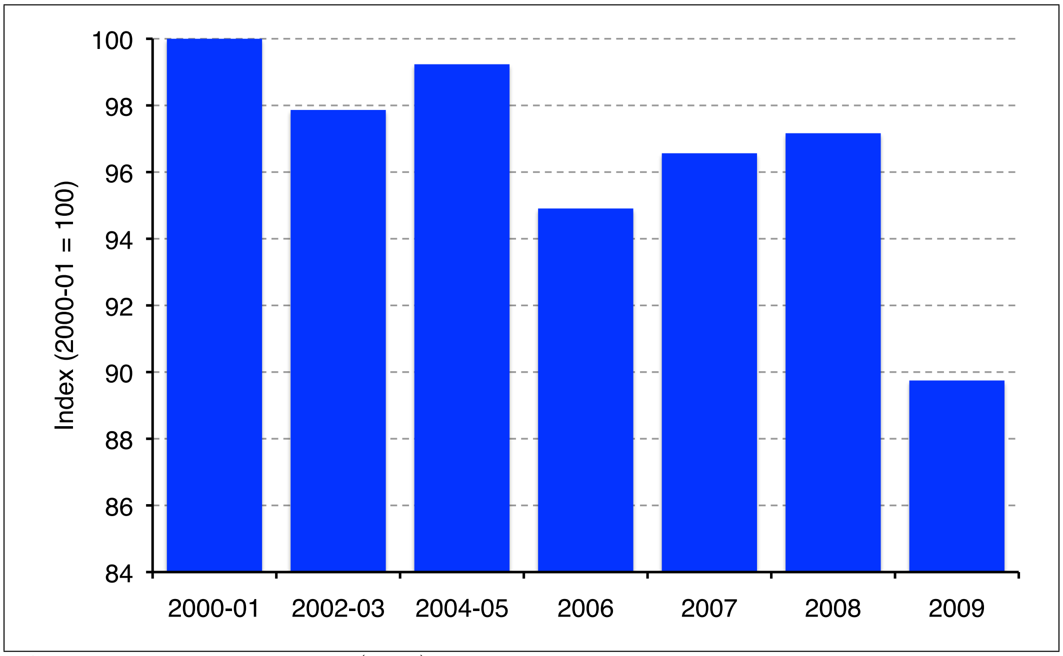

Source: Wiser and Bolinger (2012).

Figure A13: 2010 project capacity factors by commercial operation date

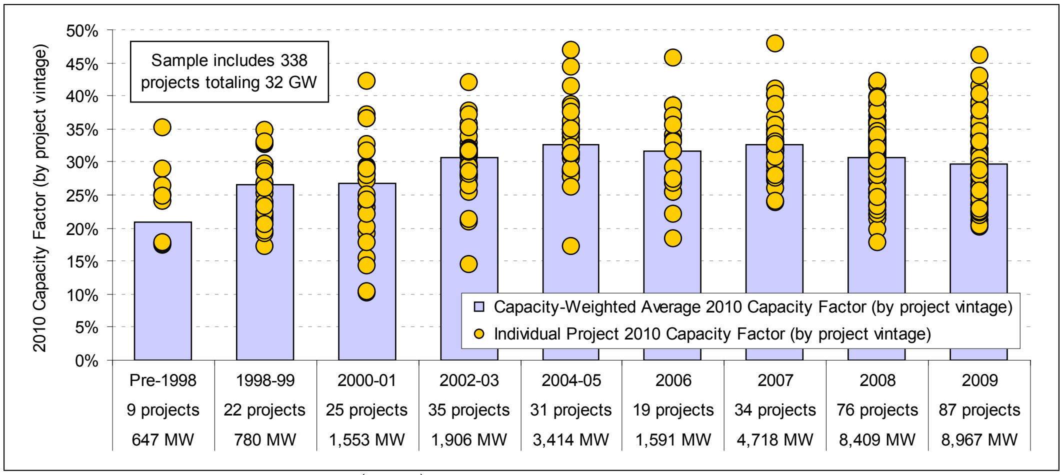

Source: Wiser and Bolinger (2011).

Table A6: NLS estimation results for cost function (13) (output measure 3: number of installed turbines)

Model variant	Parameter										
	α_{L^U}	α_{E^U}	α_{L^D}	α_{E^D}	γ	β	θ	δ	λ_2	λ_3	μ
3-1	0.0284	0.1455	0.7003	0.0002	1.0199	0.0157	0.0757	0.2928	1.5684	0.0003	47.612
	(0.2131)	(0.8634)	(0.1909)	(0.0818)	(0.0191)	(0.0080)	(0.0760)	(0.1610)	(3.4658)	(0.0016)	(201.68)
3-2	0.0591	0.1507	0.6949	0.0003	1.0183	0.0173	0.0839	0.2562	1.4634		26.021
	(0.2152)	(0.8682)	(0.1910)	(0.0813)	(0.0196)	(0.0079)	(0.0829)	(0.1185)	(2.9454)		(98.483)
3-3	0.0609	0.1718	0.6949	-0.0004	1.0186	0.0177	0.0840	0.2521	1		27.827
	(0.2128)	(0.8604)	(0.1902)	(0.0805)	(0.0203)	(0.0073)	(0.0840)	(0.1113)			(101.64)
3-4	0.0644	0.2073	0.6930	-0.0008	1.0196	0.0176	0.0806	0.2539	0.5		34.560
	(0.2144)	(0.8638)	(0.1899)	(0.0809)	(0.0208)	(0.0078)	(0.0836)	(0.1165)			(130.40)
3-5	0.0590	0.1198	0.6920	-0.0041	1.0185	0.0177	0.0887	0.2489	1.6842		1
	(0.2134)	(0.8665)	(0.1904)	(0.0811)	(0.0197)	(0.0079)	(0.0850)	(0.1148)	(3.3744)		
3-6	0.0560	0.0917	0.6785	-0.0036	1.0202	0.0140	0.0674	0.3145	1.4935		0
	(0.2137)	(0.8674)	(0.1864)	(0.0812)	(0.0197)	(0.0078)	(0.0703)	(0.1594)	(3.6916)		
3-7	0.0620	0.1484	0.6905	-0.0052	1.0188	0.0184	0.0898	0.2422	1		1
	(0.2113)	(0.8596)	(0.1890)	(0.0805)	(0.0203)	(0.0073)	(0.0870)	(0.1060)			

Heteroskedasticity-consistent standard errors in parentheses.
$N = 225$ for model 3-1; $N = 223$ for models 3-2 through 3-7.

References

Alchian, Armen, "Reliability of Progress Curves in Airframe Production," *Econometrica*, 1963, *31* (4), 679–693.

Argote, Linda, Sara L. Beckman, and Dennis Epple, "The Persistence and Transfer of Learning in Industrial Settings," *Management Science*, 1990, *36* (2), 140–154.

Arrow, Kenneth J., "The Economic Implications of Learning by Doing," *Review of Economic Studies*, 1962, *29* (3), 155–173.

Baloff, Nicholas, "Startup Management," *IEEE Transactions on Engineering Management*, 1970, *17*, 132–141.

Barradale, Merrill J., "Impact of public policy uncertainty on renewable energy investment: Wind power and the production tax credit," *Energy Policy*, 2010, *38*, 7698–7709.

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

Bolinger, Mark and Ryan Wiser, *Understanding Trends in Wind Turbine Prices Over the Past Decade*, United States Department of Energy, October 2011.

Boston Consulting Group, *Perspectives on Experience*, Boston: The Boston Consulting Group, Inc., 1968.

Cabral, Luis M.B. and Michael H. Riordan, "The Learning Curve, Market Dominance, and Predatory Pricing," *Econometrica*, 1994, *62* (5), 1115–1140.

Davidson, Russell and James G. MacKinnon, *Estimation and Inference in Econometrics*, New York: Oxford University Press, 1993.

Heckman, James J., "The Common Structure of Statistical Models of Truncation, Sample Selection and Limited Dependent Variables and a Simple Estimator for Such Models," *The Annals of Economic and Social Measurement*, 1976, *5*, 475–492.

_ , "Sample Selection Bias as a Specification Error," *Econometrica*, 1979, *47* (1), 153–161.

Hirsch, Werner Z., "Manufacturing Progress Functions," *The Review of Economics and Statistics*, 1952, *34* (2), 143–155.

_ , "Firm Progress Ratios," *Econometrica*, 1956, *24* (2), 136–143.

International Renewable Energy Agency, *Renewable Energy Technologies: Cost Analysis Series: Wind Power*, Bonn, Germany: International Renewable Energy Agency, June 2012.

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

Joskow, Paul L. and Nancy L. Rose, "The Effects of Technological Change, Experience, and Environmental Regulation on the Construction Cost of Coal-Burning Generating Units," *The RAND Journal of Economics*, 1985, *16* (1), 1–27.

Kellogg, Ryan, "Learning by Drilling: Interfirm Learning and Relationship Persistence in the Texas Oilpatch," *The Quarterly Journal of Economics*, 2011, *126* (4), 1961–2004.

Nemet, Gregory F., "Subsidies for New Technologies and Knowledge Spillovers from Learning by Doing," *Journal of Policy Analysis and Management*, 2012, *31* (3), 601–622.

Nordhaus, William D., "The Perils of the Learning Model for Modeling Endogenous Technological Change," January 2009. Cowles Foundation Discussion Paper No. 1685, Yale University.

Spence, A. Michael, "The learning curve and competition," *The Bell Journal of Economics*, 1981, *12* (1), 49–70.

Thompson, Peter, "How Much Did the Liberty Shipbuilders Forget?," *Management Science*, 2007, *53* (6), 908–918.

United States Energy Information Administration, *Direct Federal Financial Interventions and Subsidies in Energy in Fiscal Year 2010*, Washington, DC: United States Department of Energy, July 2011.

Wiser, Ryan and Mark Bolinger, *2009 Wind Technologies Market Report*, United States Department of Energy, August 2010.

_ **and** _ , *2010 Wind Technologies Market Report*, United States Department of Energy, June 2011.

_ **and** _ , *2011 Wind Technologies Market Report*, United States Department of Energy, August 2012.

Wright, Theodore P., "Factors Affecting the Cost of Airplanes," *Journal of Aeronautical Sciences*, 1936, *3* (4), 122–128.

Zimmerman, Martin B., "Learning Effects and the Commercialization of New Energy Technologies: The Case of Nuclear Power," *The Bell Journal of Economics*, 1982, *13* (2), 297–310.