HARVARD Kennedy School
MOSSAVAR-RAHMANI CENTER
for Business and Government

The Effects of Experience and Technological Innovation in the Offshore Wind Industry

Honorable Mention, 2021 Dunlop Undergraduate Thesis Prize

Keshav Rastogi

June 2021

M-RCBG Associate Working Paper Series | No. 171

Mossavar-Rahmani Center for Business & Government
Weil Hall | Harvard Kennedy School | www.hks.harvard.edu/mrcbg

The Effects of Experience and Technological Innovation in the Offshore Wind Industry

Keshav Rastogi

Presented to the Department of Economics
in partial fulfillment of the requirements
for a Bachelor of Arts degree with Honors

Harvard College
Cambridge, Massachusetts
March 11, 2021

Abstract

Using data on completed offshore wind farms, I seek to identify the primary drivers behind falling capital expenditure (CAPEX) in the global offshore wind sector. I test two hypotheses. One is that offshore wind developers and turbine manufacturers have experienced learning-by-doing. The other is that the most notable technological innovation in the industry, which has been the shift to larger turbines, has driven down CAPEX. After controlling for market selection of low-cost firms, I find evidence of no statistically significant returns to experience among either developers or turbine makers. In contrast, the empirical analysis indicates that a doubling in average turbine capacity is associated with a 19 percent decrease in CAPEX per watt of installed capacity, suggesting that technological innovation may be a significant part of the story of contemporaneous and future cost reductions in offshore wind. Many of the demand-pull policies that seek to take advantage of the learning curve and are currently in place in offshore powerhouses like the United Kingdom and the European Union, such as future capacity targets, may thus require alternative justifications. Government research and development spending as related to turbines, on the other hand, may be prescient.

Acknowledgements

This thesis would not be what it is without the support and advice of many individuals. I would like to thank:

My advisor Professor Jim Stock for his insights and encouragement over the last 1.5 years. In addition to his feedback, I immensely appreciated his unshakable faith in the pertinence of this research during an unusual and difficult year.

My seminar leader Andrés Maggi for providing advice, feedback, and positivity on a regular basis.

Robin Lee, Myrto Kalouptsidi, Robert Stavins, Joe Aldy, Ben Clinton, Sarah Armitage, Kristen McCormack, Michael Aziz, Dan Schrag, Frank Wolak, Richard Sweeney, and Brian Snyder for discussing this research and providing comments.

Representatives from Ørsted, Northland Power, Siemens Gamesa, Bloomberg New Energy Finance, Wood Mackenzie, and Sustainable Energy Advantage for their invaluable industry-related insights.

My friends and my seminar peers for their support and suggestions at various stages of the research process.

Most importantly, my family for their unwavering and inspiring love, encouragement, and work ethic. None of this would have happened without them.

Contents

1 Introduction

Decarbonizing the electricity supply is a critical component of the international fight against climate change. As a result, countries and subnational entities around the world are supporting the deployment of energy facilities that harness renewable resources, such as solar and wind energy. These governments seek to minimize the cost of these initiatives, not only because they have other social issues to tend to but also because doing so can maintain the support of the ratepaying public.

There are two forms of wind energy—onshore, which involves installing wind turbines on land, and offshore, which includes those placed in water bodies. Land-based wind is much more prevalent globally, with total installed capacity approximately twenty times that of the offshore sector (Taylor et al. 2020). Onshore wind is also significantly cheaper, both in terms of capital expenditure associated with installation (CAPEX) and the levelized cost of electricity (LCOE), which is a measure of the average present-value cost of each unit of electricity produced by a power plant over its lifetime.[1] However, offshore wind energy is an attractive option for countries with limited or densely populated land area but extensive coastline, and its costs are falling.[2] Offshore wind farms also tend to produce more energy with less intermittency because of stronger and more consistent winds over water bodies. For these reasons, certain governments, particularly those in Northern Europe and China, have subsidized the growth of the offshore sector. As these countries and others, such as the United States, seek to minimize the impact of renewable subsidies on ratepayers, choosing the right policies to maximally drive down costs in the industry is of the utmost importance.

[1] CAPEX includes spending on wind turbines, the balance of system (foundations and electrical infrastructure), and installation costs (permitting, construction equipment, labor, and cost of capital).

[2] See Figure 2.

In this study, I explore the drivers of CAPEX reductions in the offshore wind energy sector. Specifically, I investigate two hypotheses. One is whether there has been significant learning-by-doing among companies developing (planning and assembling) offshore wind farms and those manufacturing turbines. The other is whether technological innovations among turbine original equipment manufacturers (OEMs) have been primarily responsible for recent reductions in CAPEX. The primary technological advancement has been in the size of turbines, as offshore turbines have become continually larger than their onshore counterparts since the early 2010s. Learning-by-doing is the idea that productivity increases due to accumulated experience. In capital-intensive industries such as the offshore wind sector, these improvements theoretically stem from engineers and managers making slight modifications to the production process (Benkard 2000).

To distinguish innovation from learning-by-doing, Stein (1997) posits that innovations, which arrive in waves, lead to discrete cost reductions, while learning-by-doing allows firms to decrease production costs given a fixed technology. While Stein (1997) approaches this topic from a firm-specific perspective, this theory can also be applied to the offshore industry as a whole. Within this theoretical framework, discrete jumps in turbine size constitute the innovations, and turbine OEMs and developers experience learning-by-doing in manufacturing and installing turbines of a given size. The shift to larger turbines could itself be construed as a result of accumulated experience among OEMs and developers, but this interpretation likely requires stretching the definition of learning-by-doing.

The literature suggests the presence of learning-by-doing in the supply chains for other renewable energy sources. For example, Anderson et al. (2019) discover that there are cost-reducing benefits to developer-specific accumulated experience in America's onshore wind industry. Nemet (2019) finds evidence of significant learning-by-doing within the solar photovoltaic industry, as the actual technology itself has remained remarkably similar since the 1950s. I describe more of the relevant literature in Section 3.

In order to understand whether there exists learning-by-doing in the offshore wind sector, I build off of Anderson et al. (2019) and model the typical developer's project design problem, which includes, among others, the variables of interest measuring developer experience, OEM experience, and turbine size. By exploiting the fact that developers frequently have limited agency over deciding the size of their projects, I solve for the developer's profit maximizing decision and manipulate the resulting equation such that it can be empirically estimated by an OLS regression with fixed effects. I conduct the analysis using data from offshore energy consultancy 4C Offshore covering all fully commissioned, or completed, offshore wind farms. The preferred set of specifications, which control for market selection of low-cost firms, show no statistically significant evidence of learning-by-doing among either developers or turbine OEMs.

In contrast, the empirical results suggest that a doubling in average turbine rating, or size, is robustly associated with a 19 percent decrease in total costs, which corroborates the broad academic and industry-based consensus that larger turbines have been a key driver of CAPEX reductions. I then discuss the meaning of the primary ordinary least squares (OLS) and robustness check results. I first note that the selection robustness check demonstrates that any observed learning-by-doing effect in the offshore wind industry is likely due to low-cost firms capturing market share rather than these firms seeing their costs fall with increased experience. I also demonstrate how the cost-reducing effects of turbine rating nest within the developer's decision regarding turbine size at a particular site. I then use these estimates to make projections for the costs of offshore wind farms currently under development globally. I conclude by connecting the results of my analysis to the policy debate around encouraging the growth of offshore wind energy and renewables more broadly.

Overall, understanding cost drivers in the sector is important because of how quickly the international and domestic offshore wind industries are projected to grow. By the end of 2018, nearly 23 gigawatts (GW) of offshore wind capacity were installed globally, with 154 to 193 GW projected by 2030. Moreover, the United States currently has only two fully commissioned offshore wind farms, but the operational pipeline stood at more than 25 GW at the end of 2018 (Musial et al. 2019). For context, as of 2019, the United States has around 1,100 GW of installed electricity generation capacity (Energy Information Administration 2020). Discovering whether learning-by-doing specifically has been leading to significant CAPEX decreases has important climate modeling and policy implications. Large-scale energy-economic models, such as the EPA-MARKAL model and REGEN model, are increasingly incorporating learning curves in order to measure technological change endogenously (Rubin et al. 2015). These models are often used to identify efficient and cost-effective climate policies. Broadly speaking, if learning-by-doing has been a significant cost reducing factor in the offshore wind industry, then demand-pull policies that incentivize rapid deployment and, accordingly, cost reductions, such as future capacity targets, would be well advised. Such measures have historically dominated the offshore wind policy dialogue. The United Kingdom and Germany, the countries with the most installed offshore wind capacity, have set targets for 40 GW and 20 GW of capacity by 2030, respectively (Durakovic 2020a, 2020b). In contrast, if technological innovation, such as turbine upscaling, is a more significant cost lever, then technology-push policies, like increased public investment into research and development on turbines and other components of an offshore wind farm, may constitute a wiser path forward.

The rest of the paper is organized as follows. Section 2 provides relevant background on offshore wind farms and the industry as a whole, and Section 3 describes where this paper fits within the existing literature. I describe the data and the theoretical and empirical models in Sections 4 and 5, respectively. Section 6 displays the results of the empirical estimations, and Section 7 contextualizes the results. The paper finishes with Section 8, which uses the empirical model to make future CAPEX projections, and Section 9, which recapitulates the study and considers its policy implications.

2 Offshore Wind Industry Overview

This section describes the structure of an offshore wind farm and the history of the offshore wind industry in a manner that is germane to the cost reduction analysis conducted in this study.

2.1 Offshore Wind Farm Structure

Offshore wind facilities are complex projects that consist of several distinct components. The electricity is produced by turbines that are typically fixed to the seabed by foundations. The most common type of foundation is a monopile, which is simply a steel pile that is driven into the seabed and then connected to the turbine. The technology used to drill the first monopiles in the industry was not very different from 1980s offshore oil and gas technology. Generated electricity is carried by array cables to an offshore substation that sends the electricity along an export cable to an onshore substation, which is directly connected to the grid. The construction process involves specialized “jack-up” vessels, which install the turbines, in addition to cable-laying vessels and substation installation vessels (BVG Associates 2019).

Figure 1: CAPEX Breakdown

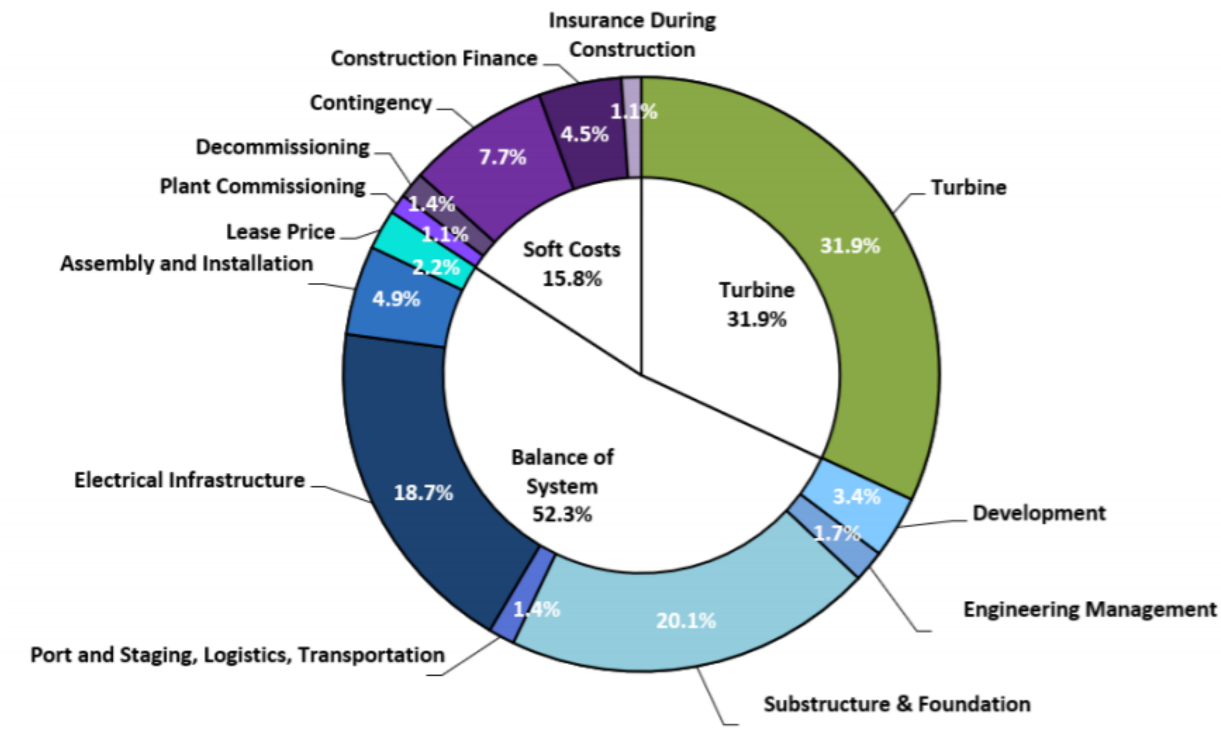

Notes: Source is Stehly et al. (2020).

Developers such as Ørsted and Vattenfall are responsible for planning the offshore wind farm and bringing all of these components together. They frequently put out requests for proposals (RFPs) to select the turbine OEM, cable manufacturer, and other firms that have specialized roles in the installation process. Developers must also handle the permitting and consent processes, which include feasibility studies, environmental surveys, and geophysical surveys (BVG Associates 2019). Figure 1 provides a breakdown of the CAPEX of a typical, fixed-bottom offshore wind facility.

It is worth noting that, under ideal circumstances, this analysis would be conducted with the LCOE instead of CAPEX as a measure of costs. What ultimately matters is the cost of each unit of generated electricity, not necessarily the upfront cost of the power plant. CAPEX data, however, is much more available than LCOE statistics, which also require analysts to make assumptions about an offshore wind farm's output over its lifetime. CAPEX, as opposed to operational expenditures (OPEX) and other costs, typically accounts for around 70 percent of LCOE (Crabtree et al. 2015). The remainder primarily stems from OPEX related to operating and maintaining the turbines and associated electrical infrastructure over the lifetime of the offshore wind farm. Because the offshore wind industry is so capital-intensive, it is reasonable to use CAPEX data as a proxy for overall electricity costs.

2.2 Offshore Wind Industry History

The industry is relatively young, as the first offshore wind farm was installed in Danish waters in 1991. The Vindeby project produced 5 megawatts (MW) of power in total, consisting of 11 0.45 MW turbines installed in water depths less than 7 meters (Lehn-Christiansen 2017). For context, Vindeby could power 2,200 Danish homes on average. Like Vindeby, early offshore wind facilities were predominantly Danish and small. They were also cheap because they utilized close derivatives of onshore wind turbines based on concrete foundations in shallow waters no deeper than 10 meters. These projects were primarily driven by direct government orders (Gottlieb et al. 2019).

Figure 2: CAPEX Over Time

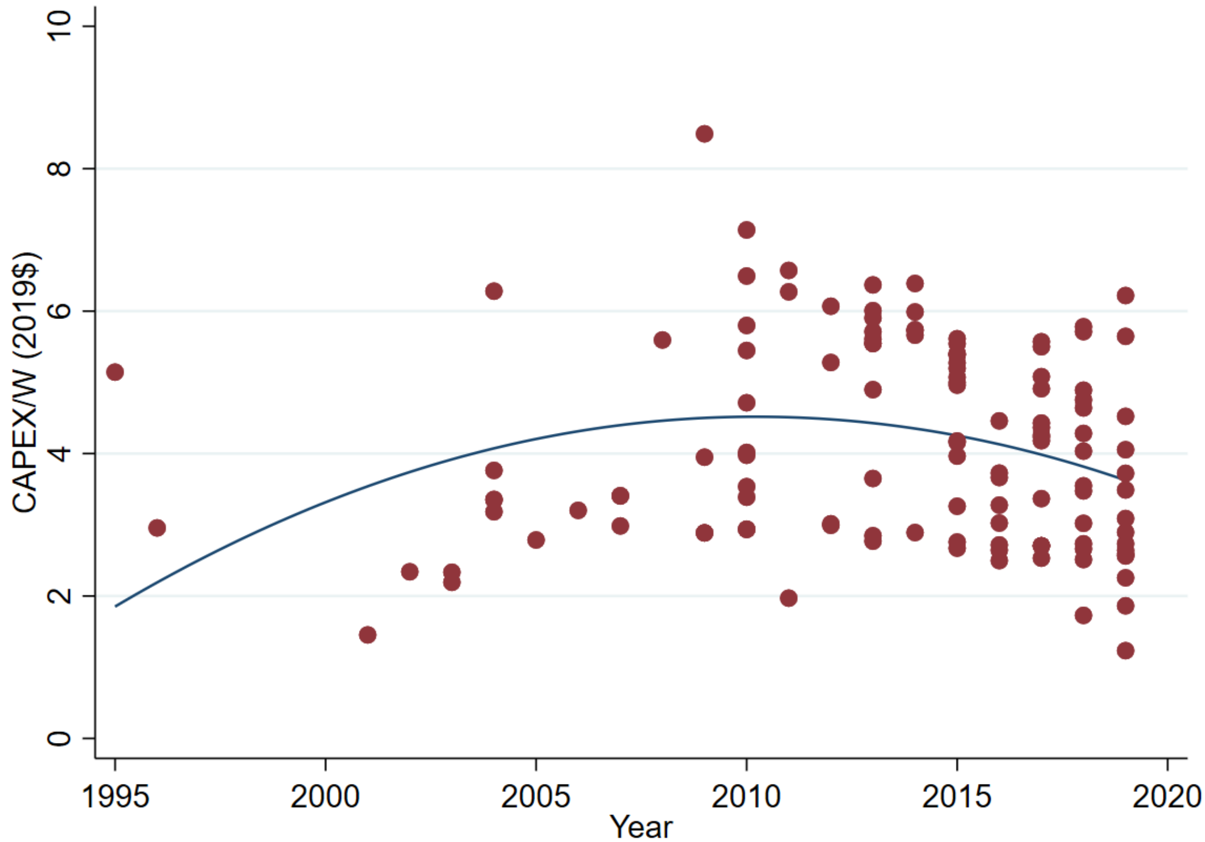

Notes: Each point corresponds to a unique wind farm, plotting its unitized CAPEX (total project CAPEX divided by project capacity) against the year when it was fully commissioned.

In the 2000s, Danish governmental plans and offshore leasing rounds held by the Crown Estate in the United Kingdom led to the commissioning of the first utility-scale offshore wind farms, which each had a total capacity in the hundreds of MW. These supportive policies, combined with generous feed-in tariffs and demand-pull national buildout plans in Denmark, the United Kingdom, and other Northwestern European countries, increased market volume and encouraged turbine OEMs to begin producing dedicated offshore wind turbines (Gottlieb et al. 2019). Feed-in-tariffs are long-term contracts that guarantee a fixed, administratively set offtake price for each megawatt hour of electricity produced by a facility over its entire lifetime. Feed-in tariffs differed slightly from country to country, but they were all government-set subsidies that were often higher for offshore wind energy than for other renewable and nonrenewable energy sources.

This rapid rise in demand, along with an underdeveloped supply chain, rising commodity prices, and a trend towards installing offshore wind farms in deeper waters, led to a rise in CAPEX in the late 2000s and early 2010s (Van der Zwaan et al. 2012). Figure 2 plots all of the CAPEX observations for individual offshore wind farms in my sample against time, displaying a parabolic trajectory for CAPEX over the course of the offshore wind sector's history. The decrease in CAPEX since it peaked around 2013 is one of the motivations for this paper.

By the mid-2010s, governments interested in growing their domestic offshore wind sectors began prioritizing cost reductions. Rather than providing fixed, expensive subsidies, the United Kingdom, the Netherlands, Germany, and Denmark began running competitive auctions that selected the projects with the lowest-price bids. These bids are based on the LCOE of output rather than the upfront CAPEX, but as I explain above, CAPEX is a reasonable proxy for LCOE. Simultaneously, turbine OEMs began selling larger turbines with longer blades and higher output (Gottlieb et al. 2019). Due to the long lead times associated with offshore wind development, most of the projects in my sample, which are all fully commissioned, were not procured under these competitive tenders such that this shift in policy design cannot explain the post-2013 decreasing trend in CAPEX. It is also worth noting that, throughout the 2010s, China was the only non-European country to scale its offshore wind sector.

Overall, the trends of increasing water depth, average turbine rating, and project size have led to a dramatic evolution of the sector. Hornsea Project One, which was fully commissioned in 2019, has 7 MW turbines installed in water depths ranging from 20 to 40 meters for a total capacity of 1200 MW. Figure 3 provides a visualization of the scale of these turbines in order to show the results of technological innovation on the part of the turbine OEMs.

Figure 3: Growth in Turbine Rating

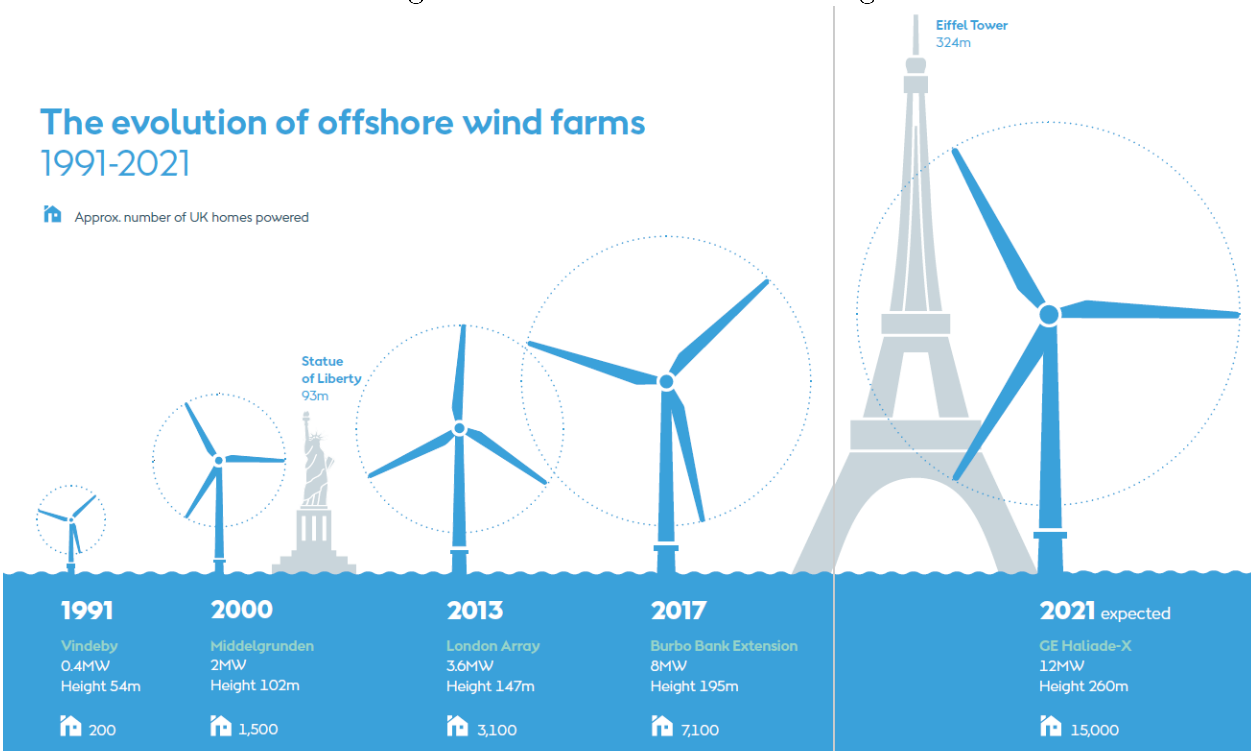

Notes: Source is Gottlieb et al. (2019).

3 Related Literature

This section reviews relevant industry reports and economic literature related to the hypotheses about learning-by-doing and technological innovation.

3.1 Cost Reductions in Offshore Wind

The continual increase in offshore wind turbine capacity is emphasized as a major CAPEX reduction driver in the majority of the literature (Van Hoof and Velthuijsen 2018; Musial et al. 2019; New York Power Authority 2019; Jennings et al. 2020; Taylor et al. 2020). Larger turbines are associated with lower CAPEX per unit of capacity and economies of scale during the installation process. During interviews, employees at Ørsted and at Siemens Gamesa, which are the world's most prominent developer and turbine manufacturer, respectively, also agreed with these assessments. Analysts at the energy research outfits Bloomberg New Energy Finance (BNEF) and Wood Mackenzie endorsed this interpretation as well. Separately, certain exogenous factors, such as commodity prices (especially that of steel) and the average water depth, are almost ubiquitously mentioned as important determinants of CAPEX by these sources as well.

In contrast, it is unclear whether learning-by-doing is a major cost reduction driver in the offshore wind industry, and the author is unaware of any econometric studies that answer this question at the time of writing. The literature examining learning in the offshore wind sector tends to be limited to industry reports. Some literature points to anecdotal evidence about learning-by-doing among both developers and manufacturers. Snyder and Kaiser (2008) cite evidence that, for individual wind farms, firms install the last turbines faster than they do the first ones in order to assert that firms with greater offshore installation experience can boast shorter installation times. Sources from Ørsted corroborated this claim. Jennings et al. (2020) compile interviews from a variety of stakeholders in the U.K. offshore wind sector and find that learning-by-doing is frequently associated with successful deployment of each generation of wind turbines, encouraging the introduction of the next, typically larger generation. Taylor et al. (2020) argue that both accumulated developer and turbine manufacturer experience have been important in driving down CAPEX. Van Hoof and Velthuijsen (2018) and New York Power Authority (2019), however, do not discuss cost reductions due to learning-by-doing. Furthermore, during an interview, an employee at Siemens Gamesa did not immediately point to learning as an important factor.

There are other drivers, similar to learning-by-doing, that are stressed only by a subset of articles, stakeholders, and analysts, and these sources sometimes disagree on the direction of these drivers' impacts on CAPEX. New York Power Authority (2019), Jennings et al. (2020), Taylor et al. (2020), and employees at Ørsted assert that competition among developers and turbine manufacturers, respectively, has been important in bringing down costs. In contrast, Jennings et al. (2020) note that it is possible that increased competition has hampered knowledge exchange between firms, and Ibenholt (2002) states that greater competition can create volatile markets for wind developers. Van Der Zwaan et al. (2012), Vieira et al. (2019), and Jennings et al. (2020) point to the maturation of the supply chain, especially the growth in manufacturing capacity and the development of "jack-up" vessels specialized for offshore wind installation, as an important factor pushing down CAPEX. Analysts at Wood Mackenzie, Musial et al. (2019), and Jennings et al. (2020) note that technological innovation among array and export cable manufacturers has reduced costs, while Van Hoof and Velthuijsen (2018) and Taylor et al. (2020) fail to mention this CAPEX driver.

Overall, there is disagreement in the industry about the importance of certain CAPEX reduction drivers. Of those that are not cited ubiquitously in the literature, increasingly competitive procurement regimes and supply chain maturation are highly endogenous determinants of CAPEX whose effects would be difficult to separate from those of other drivers. For example, it is unclear how one would distinguish the impact of slack in the supply chain from technological innovation in the turbine manufacturing sector. In contrast, the phenomenon of learning-by-doing has been well studied by the field of industrial organization as described in the following section.

3.2 Learning Curve Analysis in Renewables

Studies of learning-by-doing have been conducted for various industries, including renewable energy sectors. The concept of learning-by-doing is not new and was first formally introduced by Wright (1936) in the context of the aircraft industry. Many analyses of learning-by-doing simply regress some measure of unitized costs on a variable representing cumulative experience, such as the number of airplanes a manufacturer has produced or the MW of electricity generation capacity a project developer has installed.

Some of the studies that have most convincingly estimated learning-by-doing, such as Benkard (2000), use data on an individual manufacturing company's marginal cost for each unit of a particular product. For the offshore wind industry, performing a similar analysis on production costs of a specific type of turbine or other wind farm component would be ideal. However, the industry is relatively nascent and certain markets in its supply chain have high levels of concentration, making it difficult to obtain or recover data on individual wind farm component costs. Total wind farm CAPEX data is more readily available and provides a more holistic picture of offshore wind costs regardless.

Learning curve analysis is a popular approach in cost studies of the energy industry, and particularly in research on renewable sectors. The prospect of cost reductions due to learning-by-doing is often used to justify policies supporting the initial deployment of early-stage renewable technologies. Partially due to its simplicity, the one-factor learning curve, which typically measures the log-log relationship between unit costs and cumulative installed capacity and does not account for any other explanatory variables, dominates the renewables-focused literature (Rubin et al. 2015).

There are significant concerns with interpreting simple learning curve models as causal. The omission of a measure of exogenous technical change can lead to significant upward bias in estimates of the learning parameter, as empirically demonstrated by Nordhaus (2014). Nordhaus (2014) notes that exogenous technical change includes most sources of cost declines other than the learning curve, such as the returns to research and development and spillover inventions from other economic sectors. Abernathy and Wayne (1974) maintain that most empirical research on learning curves incorrectly assumes that cost reductions due to learning continue in perpetuity. In short, Ibenholt (2002) and Anderson et al. (2019) note that cost reductions in renewable energy industries are driven by four main factors: learning-by-doing, input price changes, exogenous technical change, and economies of scale. If an empirical strategy can separate learning-by-doing from the other factors, it should constitute a valid modeling approach.

The literature has evolved to try to address some of these concerns. One approach is to estimate two-factor learning curves, which include some measure of research and development spending in addition to a cumulative capacity variable (Jamasb 2007). Furthermore, to account for heterogeneity among different firms, Van Benthem et al. (2008) and Bollinger and Gillingham (2019) estimate residential photovoltaic solar cost reductions due to learning by measuring each solar installer's accumulated experience, rather than using an industry-wide measure. This approach also allows one to separate appropriable learning from non-appropriable learning. Theoretically, governments using learning as a justification to subsidize a renewable technology should only do so when there exists significant non-appropriable learning because only non-appropriable learning constitutes a positive externality. Anderson et al. (2019) build on this work by developing a structural model to describe cost minimization on the part of onshore wind developers in the United States and empirically estimating the model to find if there are appropriable and non-appropriable returns to experience in the sector. I use this paper's model as a theoretical basis and extend it for my analysis.

4 Data

This section surveys the data used in this study, commenting on its particularities and limitations when relevant. The primary dataset I use is from the U.K.-based marine consultancy 4C Offshore. The dataset is global and includes information on every fully commissioned, decommissioned, and future offshore wind farm. Each observation pertains to a single offshore wind facility and, inter alia, has data on the facility's home country, the location's geographic attributes (i.e. water depth, distance from shore, wind speed), the wind farm's development timeline, its developers, its turbines, its revenue mechanism, and its overall CAPEX. 4C Offshore collects these data through close correspondence with the companies developing these wind farms. I verify the data and fill in missing observations for variables that are important in the main analyses, primarily through online research.

Appendix A.1 includes more information on manipulations made to the original dataset.

I obtain data on commodity prices and other controls from a variety of sources. I get data on average annual prices of refined copper (cents/lb), Brent Crude ($/barrel), and steel ($/tonne) from the Economist Intelligence Unit. I also use data on the annual hydraulic cement manufacturing producer price index, indexed to June 1989 and sourced from Federal Reserve Economic Data. These commodities, especially steel, are important in the offshore wind construction process. To account for country-level differences in labor costs, I utilize monthly wage data averaged at the annual level for various countries from the Economist Intelligence Unit. Finally, as a proxy for exogenous technological change, I obtain global weighted-average unitized CAPEX data for the onshore wind sector from IRENA's *Renewable Power Generation Costs in 2019* report.

There are some limitations with this dataset. There is a relatively small number of observations, which is naturally a drawback of studying a nascent sector. I use data from offshore wind farms that had been fully commissioned by 2019, and there were fewer than 130 completed facilities by that time. Even though I verify most of the relevant observations in the 4C Offshore database, there may still be some slight inaccuracies in the data. The 4C Offshore data stems from quotes provided by developers and are thus not adjusted for variations in policy that directly impact CAPEX. For example, transmission policy varies by country, and these regulations determine which stakeholder is responsible for securing investments for transmission assets, such as onshore and offshore substations and undersea cables. In the United Kingdom, for example, developers are initially responsible for making sure their offshore wind facilities can link up to the grid, while, in Germany, it is the transmission system operator's responsibility from start to finish (New York Power Authority

2019). Because 4C Offshore collects its CAPEX data from developer quotes, its CAPEX figures for offshore wind farms in the United Kingdom would include transmission investment, while the statistics for their German counterparts would not. This is a potential issue because transmission typically comprises around one-fifth of total CAPEX. I address how I deal with this potential measurement error in Section 5.1.

To deal with outliers, I remove demonstration projects with less than 10 MW of capacity and fewer than 4 turbines. These projects are usually built for research purposes and are thus not typically built with cost minimization in mind. I also remove the BARD Offshore 1 wind farm because it faced significant scheduling delays and cost overruns due to a series of unique engineering and construction setbacks (Karnitschnig 2014). I do the same for the Block Island wind farm because it was the first project built in U.S. waters and most of the installation infrastructure had to come from Europe, significantly increasing costs (McKenna 2017). The 4C Offshore dataset provides CAPEX in various currencies, so I standardize all CAPEX figures to 2019 USD. Finally, many developers form special purpose vehicles for each of their projects, so I manually separate these vehicles into their member companies to get a more accurate representation of the developers behind each wind farm.

The variables of interest are the measures of cumulative experience for developers and turbine OEMs, respectively, and the average turbine capacity. To compute these variables, I simply sum the MW of capacity that the developer has installed and that the turbine OEM has sold, respectively, prior to the observed wind farm's installation. If a project has multiple developers or multiple turbine manufacturers, I divide the project capacity by the number of developers and OEMs, respectively, and attribute the resulting fraction to each developer's and OEM's stocks of experience. I do not account for acquisitions in either portion of the supply chain and do not account for depreciation in experience across time and space.

Figure 4: CAPEX vs. Commonly Cited Explanatory Variables

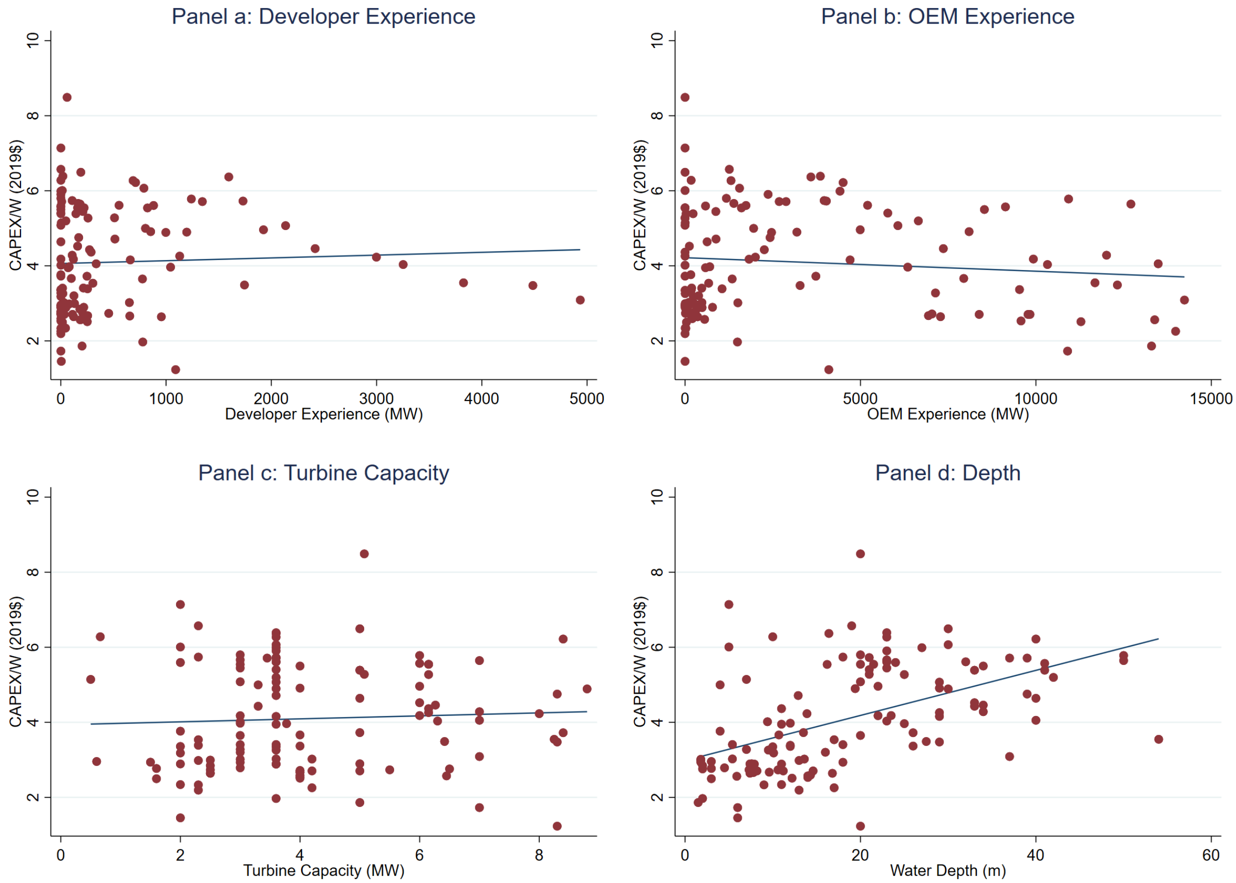

Notes: Within each panel, each point corresponds to a unique wind farm, plotting its unitized CAPEX against a potential explanatory variable. Developer Experience is measured by the cumulative MW installed by the developer(s) prior to the observed wind farm, and OEM Experience is the cumulative MW of turbines sold by the turbine OEM(s) prior to the observed wind farm.

Figure 4 includes various plots that suggest the direction of the relationships between the variables of interest and unitized CAPEX, which is the quotient of total project CAPEX and project capacity. At first glance, cumulative developer experience and turbine rating appear to have no correlation with CAPEX, while cumulative OEM experience appears to have a weakly negative relationship. Ceteris paribus, the theory of learning-by-doing would suggest that both experience variables should have negative relationships with CAPEX. Also, the positive slope of the line of best fit in Figure 4c seems to refute the aforementioned consensus that larger turbines have driven down capital expenditure.

It is worth noting that observations of developer experience are clustered below 1000 MW, indicating that most developers are relatively inexperienced, while observations of OEM experience are slightly more spread out. This is partially a result of the fact that the turbine OEM market is significantly more concentrated than the developer space. The only developer with more than 3000 MW of experience at any point in the dataset is Ørsted. In contrast, Siemens Gamesa, the market-leading turbine OEM, and Vestas have sold more than 15000 MW and nearly 5000 MW of turbines, respectively. As a result of this market concentration, there is more variation in observations of OEM experience than there is in observations of developer experience. Finally, of all the potential explanatory variables, water depth seems to have the strongest relationship with CAPEX.

5 Methodology

This section describes the theoretical model of the developer's project design problem, empirical estimations of the model, and associated robustness checks.

5.1 Model

The theoretical basis for my estimation strategy is similar to that employed by Anderson et al. (2019). I assume that the production function for developers in installing offshore wind capacity is Cobb-Douglas,

$$q_i = f(A_i, z_M, z_L, z_K) = A_i z_M^{\alpha_M} z_L^{\alpha_L} z_K^{\alpha_K}, \tag{1}$$

where q_i is the installed capacity of wind farm i and the choice variables z_M, z_L, and z_K represent factor inputs from raw materials, labor, and capital, respectively. Given the Cobb-Douglas functional form, α_j is a measure of how productive each input z_j is relative to the other inputs, and $\gamma = \alpha_M + \alpha_L + \alpha_K$. I assume the following functional form for total factor productivity:

$$A_i = [Exp_{Dev_i,t_i}]^{\beta}[Exp_{OEM_i,t_i}]^{\theta}[TurbineCap_i]^{\delta} \cdot e^{\phi_{C_i}^{TFP} + \psi_{T_i}^{TFP} + Depth_i + \epsilon_i}, \tag{2}$$

where the parameter β measures the returns to cumulative developer experience Exp_{Dev_i,t_i}, θ measures the returns to cumulative turbine manufacturer experience Exp_{OEM_i,t_i}, δ measures the effects of average turbine rating $TurbineCap_i$, and t_i is the date when the observed wind farm is fully commissioned. I include country fixed effects $\phi_{C_i}^{TFP}$ to account for different policy environments. These indicator variables should absorb the heterogeneity in transmission policy mentioned in Section 4, as transmission policy in each country does not fundamentally change over the time period of this data set. I also include year fixed effects $\psi_{T_i}^{TFP}$ to account for the exogenous technological trends that affect all wind farms in the same way, which accounts for the identification concern emphasized by Nordhaus (2014). To avoid overidentifying and reducing the power of the sample, which has 124 observations, I use fixed effects for five-year periods, rather than annual fixed effects. As noted in Section 3.1, water depth and turbine capacity (or rating) are almost ubiquitously listed as important determinants of CAPEX, so I include them in the model. The error term is ϵ_i.

While country fixed effects account for differences in transmission policy, subsidy policies vary between countries and over time. Most countries in the sample initially provided generous feed-in tariffs to stimulate their domestic offshore wind industries but have since shifted to more competitive procurement regimes, such as competitive tenders (Jansen et al. 2020).[3] Because of the long lead times associated with offshore wind development, which are often in the neighborhood of five years, the transition from feed-in tariffs to competitive tenders, which began after 2015 in most countries, does not significantly manifest itself in the sample. Only 5 of the 124 observed offshore wind farms were built under competitive auction schemes. Still, in all of the specifications, I include a dummy variable for whether an observed wind farm was built after participating in a competitive tender.

Even before the shift to competitive tenders, analysts at BNEF note in interviews that countries began reducing their feed-in tariffs as a response to falling costs in the industry. I do not account for feed-in tariff variation for two reasons. First, the variation in feed-in tariffs is endogenous, as governments' electricity regulation bodies often set the tariffs in response to cost trends in the industry. Second, under the cost minimizing model described below, it can be assumed that, given an administratively fixed price per unit of output, the developers will seek to minimize costs and accordingly maximize profits regardless of the magnitude of that price.

Given that most of the wind farms in the data set receive a feed-in tariff subsidy, which is based on projected output from a preset amount of installed capacity, I assume developers

[3]Feed-in tariff regimes are described in more detail in Section 2.

solve the following cost minimization problem:

$$argmin_{z_M,z_L,z_K} p_M z_M + p_L z_L + p_K z_K \; s.t. \; q_i \leq A_i z_M^{\alpha_M} z_L^{\alpha_L} z_K^{\alpha_K}, \tag{3}$$

where p_j is the price of input j. This problem yields the cost function

$$C(q_i, A_i) = \left[\frac{q_i}{A_i} p_M^{\alpha_M} p_L^{\alpha_L} p_K^{\alpha_K}\right]^{\frac{1}{\gamma}}. \tag{4}$$

5.2 Empirical Estimation

After a logarithmic transformation, the solution to the cost minimization problem results in the following equation (after flipping the necessary signs), which can be estimated econometrically:

$$\begin{aligned} log(C_i) = \frac{\alpha_M}{\gamma} logp_M + \frac{\alpha_L}{\gamma} logp_L + \frac{\alpha_K}{\gamma} logp_K + \frac{1}{\gamma} logq_i + \frac{\beta}{\gamma} Exp_{Dev_i,t_i} \\ + \frac{\theta}{\gamma} Exp_{OEM_i,t_i} + \frac{\delta}{\gamma} TurbineCap_i + \frac{1}{\gamma} Depth_i + \frac{1}{\gamma} \phi_{C_i}^{TFP} + \frac{1}{\gamma} \psi_{T_i}^{TFP} + \frac{1}{\gamma} \epsilon_i, \end{aligned} \tag{5}$$

where the parameters of interest are β, θ, and δ.

To enhance interpretability of the estimates and more closely align with standards set by prior literature, C_i is unitized and equal to the quotient of total project CAPEX and project capacity, or q_i (Rubin et al. 2015; Anderson et al. 2019).[4] I also take logs of the developer and OEM experience variables to improve interpretability, given that these

[4]The empirical estimations also keep $logq_i$ on the right-hand side to account for the potential cost-reducing effects of scale economies. It may be concerning that q_i is on both sides of the equation, but it is important to note that it is not a variable of interest such that the magnitude of the coefficient on q_i is not particularly relevant for this study. Regardless, to address these concerns, Appendix A.2 includes the results of alternative specifications in which C_i is not unitized.

variables frequently take on values in the hundreds or thousands of MW.

For several reasons, I only include two of the three factor inputs in the actual empirical estimation—labor and raw materials, which are denoted by the subscripts L and M, respectively. I exclude capital from the estimation strategy because it is slightly redundant with raw materials and is subject to data availability constraints. Broadly speaking, developers utilize two types of capital. One category includes the physical plant components of the offshore wind farm, such as the turbines, foundations, and cables. The other includes machinery necessary to link and install these components, such as jack-up vessels. Prices of components such as turbines and cables are highly endogenous and likely subject to markups that would be difficult to estimate. Controlling for commodity prices accounts for much of the potential exogenous variation in the prices of these components. Data on specialized installation capital, such as jack-up vessels, is even less available and reliable. In fact, at the time of writing, there are only 16 jack-up vessels in the world, making it difficult to find rental rates and to assume that these capital markets are exogenous to the developer's problem (Steinberg and Wallace 2021).

The subscripts on p_L and p_M indicate that the wage data is annual by country and that the commodity price data is annual and for the global market, respectively. The first specifications are OLS estimates of this equation.

Regarding country fixed effects, I only include dummy variables for the United Kingdom, Germany, China, Belgium, the Netherlands, and Denmark in an attempt to maximize the statistical power of the estimates. It is worth noting that these 6 countries account for over 98 percent of the total installed capacity in the cleaned sample.

While offshore wind farms use copper for array and export cabling, concrete for foundations, oil to fuel jack-up installation vessels, and steel for turbines and foundations, steel is the most heavily used and critical raw material input (BVG Associates 2019). I thus estimate a variety of specifications, controlling for the prices of all of these inputs in one of the specifications and only for the price of steel in the other two. Again, I do this to preserve as much statistical power for the estimates as possible.

In one of the specifications, I control for exogenous technological trends using annual weighted average cost data for the global onshore wind sector instead of utilizing five-year fixed effects. It is plausible that including onshore CAPEX could control for general exogenous trends in wind generation technology, such as certain turbine, cabling, and installation process improvements unrelated to cumulative offshore developer experience and cumulative offshore OEM experience. The onshore wind sector is substantially larger than the offshore sector, so this variable is likely exogenous.

5.3 Robustness Checks

While the model accounts for many cost reduction factors not included in prior learning-by-doing studies, the OLS specifications will only provide consistent estimates for β and θ if certain assumptions are met. It is critical that project size, or q_i, is independently assigned. If q_i is a choice variable rather than a parameter, then the observed coefficient on q_i may be biased. Insofar as there is a correlation between q_i and developer experience or turbine OEM experience (both measured using cumulative MW developed and sold, respectively), then there may be bias introduced into the estimates of β and θ, respectively. As I argue earlier in this section, most of the observed wind farms receive a feed-in tariff, which is allocated based on a preset amount of installed capacity. It is possible, however, that the developer has some control over deciding the amount of installed capacity and that its decision may be related to the size of the projects it has installed in the past. Project size would then be correlated with the experience variables measured in cumulative MW. During interviews, industry stakeholders at Ørsted and Siemens Gamesa and analysts at BNEF and Wood Mackenzie could not come to a consensus regarding the validity of the assumption that q_i is exogenous. In case the assumption does not hold, I run regressions identical to the previously described set of OLS specifications except the developer and turbine OEM experience variables are measured in cumulative projects rather than cumulative MW installed. For example, if the firm that developed wind farm i had built 2 offshore wind farms prior, the value for Exp_{Dev_i,t_i} would be 2 regardless of how large those prior projects were. Thus, even if q_i is correlated with ϵ_i, or the structural error term, it should not be correlated with these new measures of experience. If the parameter estimates do not radically change in statistical significance or direction, then I can conclude that the assumption that q_i is independently assigned is valid.

Separately, there is a potential selection concern in which the market selects for low-cost firms, allowing these companies to gain experience because of their low costs. This is the reverse direction of causality from that implied by the presence of learning-by-doing and thus constitutes a perennial concern in estimating learning curves. As done by Anderson et al. (2019) for the U.S. onshore wind industry, I carefully examine the history of the global offshore wind sector to understand whether this concern is legitimate in this specific industrial context. Evidence of high-cost firms exiting the market or being acquired would be a cause for concern.

Table 1: Offshore Wind Turbine Manufacturing Industrial History

Turbine OEM	Acquisition	Year of Last Installation
Nordtank	Vestas (2004)	1996
Bonus	Siemens Gamesa (2004)	2001
Senvion	Siemens Gamesa (2019)	2017
Siemens Gamesa	No	2019
Vestas	No	2019
GE Energy	No	2019
WinWind	No	2010
Adwen	Siemens Gamesa (2017)	2018
Sinovel	No	2012
Fuji Heavy Industries	No	2010
United Power	No	2010
MingYang	No	2019
SEwind	No	2019
Envision	No	2019
Sany	No	2010
Goldwind	No	2019
Haizhuang	No	2019
BaoNan	No	2010
Hitachi	No	2013
Dongfang Electric Corporation	No	2015
Doosan Heavy Industries	No	2017
XEMC Darwind	No	2016

Notes: In Column (2), company listed is the acquirer. Column (3) refers to the year of full commissioning of the most recent offshore wind farm to use turbines from the OEM.

The history of the offshore wind turbine manufacturer market does not definitively demonstrate whether there is a threat posed by selection. Table 1 includes information on whether turbine OEMs in the dataset were acquired by competitors and, in order to identify firms that have potentially exited, also notes the year in which each turbine OEM last sold turbines to an offshore wind farm. Table 1 shows that every non-acquired firm sold turbines into the 2010s. Most of the OEMs that were last active in the early 2010s, such as Fuji Heavy Industries, United Power, Sany, BaoNan, and Hitachi, sold turbines to only one project. This likely indicates that they were never legitimate players in the offshore wind space and were simply testing the waters. Corroborating this interpretation, a Siemens Gamesa employee mentioned in an interview that, because of high risks and liabilities, many turbine OEMs have refused to enter the offshore market or have done so hesitantly. These risks stem from the fact that technological or quality issues have severe financial consequences given the scale of most contemporary offshore wind farms. To say these companies ceded market share to low-cost firms due to high costs would be inaccurate given that they were never truly engaged in the market in the first place.

Still, there are some firms that seemed to have legitimately entered the market and yet have not been active for more than 5 years, including WinWind and Sinovel, so I do not easily dismiss the selection concern. Regarding mergers and acquisitions, while it appears that Vestas bought Nordtank and Siemens Gamesa bought Bonus in 2004 in order to initially gain footholds in the market, rather than acquire high-cost competitors, Senvion did declare bankruptcy before selling its assets to Siemens Gamesa in 2019 (Windpower Monthly 2004; Siemens Gamesa Renewable Energy 2017; Garcia Da Fonseca and Liu 2019). The aforementioned industry stakeholders at Ørsted and Siemens Gamesa and analysts at BNEF and Wood Mackenzie disagreed about whether the acquisitions listed in Table 1 occurred because high-cost OEMs were rendered not competitive, allowing low-cost OEMs to gain market share, or whether there were other market forces at play.

Thus, while it would be possible to argue that selection for low-cost firms is not a huge threat to my model, I try to account for it with alternative specifications that include turbine manufacturer fixed effects. To maintain the statistical power of the sample, I only include fixed effects for the 6 largest turbine manufacturers, which are Siemens Gamesa, Vestas, Adwen, Senvion, Goldwind, and Envision. They collectively account for approximately 95 percent of the offshore wind capacity installed in the dataset. The inclusion of these fixed effects accounts for inherent differences in the cost structures of these firms that may have allowed low-cost OEMs to capture market share when they entered the industry. These firms' initial rise in the industry would have then spurred cost reductions not due to learning-by-doing but rather because of inherent cost advantages. This robustness check deals with the selection concern, as any CAPEX reductions occurring after these firms sold turbines to their first few projects, controlling for the variables included in the primary OLS specifications, can be attributed to accumulated experience.

6 Results

This section presents and interprets the results from the OLS estimates of Equation (5) and the robustness checks related to the project size exogeneity and selection assumptions.

6.1 Primary OLS Results

Table 2 displays results from the OLS estimations of Equation (5). The first takeaway is that turbine OEM experience appears to affect project costs, while developer experience does not. The coefficients of interest are those on $log(Developer\,experience)$ and $log(OEM\,experience)$, which are estimations of β and θ, respectively.

The estimates for β, which measures the returns to cumulative developer experience, are not statistically significant at the five percent level and are rather closely bound around zero. As a result, at the five percent significance level, I cannot reject the null hypothesis that developer experience is not associated with offshore wind CAPEX variation. In contrast, the estimates for θ, which measures the cost reducing effects of cumulative turbine OEM experience, are significant in all three specifications. The point estimates and associated confidence intervals vary only slightly between the different regressions, demonstrating that they are robust to the inclusion of different measures of exogenous technological change and other controls. The point estimates, which are around -0.02, indicate that, holding developer experience and the control variables constant, a 1 percent increase in OEM experience, as measured by cumulative MW of turbines sold, is associated with a 0.02 percent decrease in unitized CAPEX.

The estimates for the coefficient on turbine rating, or δ, are directionally aligned with the general industry and academic consensus that larger turbines are associated with lower unitized CAPEX. The point estimates in all of the specifications are around -0.25, implying that, controlling for the other variables included in the specifications, a 1 percent increase in turbine size, is associated with 0.25 percent decrease in total CAPEX per MW. While the relative costs of increasing turbine size and gaining experience may be very different, it is clear that turbine rating has a significantly larger association with cost reductions than does developer experience or turbine OEM experience. It is important to note that turbine size is treated like a control in these specifications, so drawing causal implications from these estimates should be done with caution.

Table 2: OLS Results

	(1)	(2)	(3)
VARIABLES		log(CAPEX/W)	
log(Developer Experience)	0.001	0.000	0.007
	(0.007)	(0.007)	(0.008)
log(OEM Experience)	-0.016**	-0.018**	-0.016**
	(0.007)	(0.007)	(0.007)
log(Average Turbine Rating)	-0.222***	-0.263***	-0.214***
	(0.075)	(0.076)	(0.081)
log(Project Capacity)	-0.041	-0.049	-0.020
	(0.035)	(0.034)	(0.030)
log(Steel Price)	0.090	-0.023	0.241**
	(0.242)	(0.126)	(0.109)
log(Copper Price)	0.119		
	(0.374)		
log(Wage)	0.171***	0.181***	0.122***
	(0.045)	(0.044)	(0.042)
log(Oil Price)	-0.164		
	(0.208)		
log(Cement Price)	-0.758		
	(0.581)		
Water Depth	0.013***	0.012***	0.013***
	(0.003)	(0.003)	(0.003)
Competitive Procurement	-0.054	-0.053	-0.097
	(0.124)	(0.123)	(0.135)
Average Onshore CAPEX/W			0.129
			(0.208)
Observations	124	124	124
Adjusted R-squared	0.674	0.672	0.609
Country Fixed Effects	Y	Y	Y
Time Fixed Effects	Y	Y	N

Notes: Table reports regression coefficients with robust standard errors. Developer and OEM Experience are measured in total MW of capacity developed and total MW of turbines sold prior to the observed wind farm, respectively. Steel price, copper price, oil price, cement price, and weighted average onshore CAPEX are aggregated globally and annually. Wage is aggregated by country and annually. The "Competitive Procurement" dummy equals 1 if the project received revenue under a competitive procurement, 0 otherwise. All CAPEX and price variables are in 2019 USD. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$.

Most of the other coefficients make intuitive sense, which would indicate that the OLS regression models do not terribly skew or otherwise misinterpret the data. The coefficient on the price of steel is positive in all of the specifications. It is only significant at a five percent level in Column (3), which excludes five-year fixed effects. This would imply that the time fixed effects account for steel price variation, which makes sense because the steel data is aggregated at the annual level. The lack of statistical significance for the coefficients on copper, oil, and cement prices in Column (1) likely results from the fact that commodity prices tend to be correlated, especially in the sample period, as China's rapid development in the 2000s and early 2010s and subsequent slowdown helped drive volatility in commodity markets.

The direction of the coefficient on wages is positive and significant at a five percent level in all of the specifications, which also makes sense intuitively. The coefficient on water depth in all of the specifications is significantly positive, which aligns with the data as presented in Figure 4d in Section 4. The point estimates imply that, controlling for the other variables included in the specifications, a 1 meter increase in depth is associated with a 1 percent increase in CAPEX per MW. Instead of five-year fixed effects, I include a measure of global onshore wind CAPEX as a proxy for exogenous technological change in Column (3) and its positive, albeit insignificant, coefficient aligns with the notion that there is some relation between onshore wind and offshore wind CAPEX trends.

The greatest threat to the validity of these results, other than the issues that are accounted for in the following robustness checks, is the sheer size of the sample, which has only 124 observations. This small sample is unfortunately a natural result of doing research related to such a nascent industry. One other concern stems from how closely offshore wind developers and turbine OEMs have historically worked with governments in designing policy

(Gottlieb et al. 2019). It is difficult to account for this factor given the endogeneity of such interactions. For example, one can argue that developers successfully influencing policy is a result of their accumulated experience in the offshore wind industry.

6.2 Robustness Checks

This section describes the results of the robustness checks and their implications.

6.2.1 Project Size Exogeneity Assumption Test

Table 3 displays results from the alternative specifications in which developer and OEM experience are measured in projects, rather than in MW of capacity. This is meant to deal with concerns that project size, or q_i, is not exogenously assigned. As described in more

Table 3: Project Size Exogeneity Test Results

	(1)	(2)	(3)
VARIABLES	log(CAPEX/W)		
Developer Experience	-0.002	-0.002	0.001
	(0.005)	(0.005)	(0.006)
OEM Experience	-0.003***	-0.003***	-0.003**
	(0.001)	(0.001)	(0.001)
log(Average Turbine Rating)	-0.178**	-0.214***	-0.176**
	(0.073)	(0.073)	(0.078)
Observations	124	124	124
Adjusted R-squared	0.674	0.672	0.607
Country Fixed Effects	Y	Y	Y
Time Fixed Effects	Y	Y	N

Notes: Table reports regression coefficients with robust standard errors. The three specifications are identical to those in Table 2 except developer and OEM experience are measured in total projects developed and total projects to which turbines were sold prior to the observed wind farm, respectively. These experience variables are not logarithmically transformed. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$.

detail in Section 5.3, if project size is endogenous and correlated with developer and OEM experience measured in MW of capacity, then this may introduce bias to the estimates of β and θ in Table 2. A correlation between experience measured so and project size may result if developers or turbine OEMs associated with large projects in the past retain their propensity for sizable offshore wind farms going forward.

The results of this robustness check would constitute a cause for concern if the coefficient on turbine OEM experience flips its sign or is no longer statistically significant. I focus on direction and statistical significance because the magnitude of the estimates in Table 3 is not directly comparable to that in Table 2. The experience variables in Table 3 are not logarithmically transformed as the values they take on when measured in total projects are significantly smaller than when measured in total capacity. Regardless, the direction and statistical significance of the coefficient on turbine OEM experience remain even when the project size exogeneity assumption is relaxed. Thus, any potential bias introduced by this assumption does not remove the association between OEM experience and CAPEX.

6.2.2 Selection Assumption Test

Table 4 displays results from a robustness check that simply adds turbine manufacturer fixed effects to the primary OLS specifications in Table 2. In order to preserve statistical power, the regressions in Table 4 include dummies for the most prominent turbine manufacturers that account for 95 percent of the installed capacity in the data set, which are Siemens Gamesa, Vestas, Adwen, Senvion, Goldwind, and Envision.

Table 4: Selection Test Results

VARIABLES	(1)	(2) log(CAPEX/W)	(3)
log(Developer Experience)	0.001	0.000	0.008
	(0.008)	(0.008)	(0.009)
log(OEM Experience)	-0.010	-0.021	-0.003
	(0.023)	(0.022)	(0.022)
log(Average Turbine Rating)	-0.261***	-0.302***	-0.267***
	(0.082)	(0.078)	(0.082)
Observations	124	124	124
Adjusted R-squared	0.666	0.667	0.609
Country Fixed Effects	Y	Y	Y
Time Fixed Effects	Y	Y	N
Turbine Manufacturer Fixed Effects	Y	Y	Y

Notes: Table reports regression coefficients with robust standard errors. The three specifications are identical to those in Table 2 except with turbine manufacturer fixed effects added in. Dummies were included for the turbine manufacturers Siemens Gamesa, Vestas, Adwen, Senvion, Goldwind, and Envision, which account for 95 percent of capacity installed in the sample. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$.

The coefficients on OEM experience are no longer statistically significant. This would suggest that the estimates in Table 2 face bias because of historical selection for low-cost turbine OEMs. Following the logic laid out in Section 5.3, greater turbine OEM experience did not necessarily reduce costs, but rather OEMs with inherently low costs captured market share and gained experience as a result. Thus, the perennial selection concern with learning curve analysis appears to be legitimate in this case, or, at the very least, it is difficult to make a strong empirically founded argument that the offshore wind sector has experienced significant learning-by-doing during its history.

It is worth noting that, even though they are no longer statistically significant, the coefficients on OEM experience are all still negative. In Column (2), the coefficient is almost one standard error in magnitude less than zero. Also, the coefficients on turbine size remain negative and statistically significant. In fact, they are larger in magnitude than their counterparts in Table 2, suggesting that additional cost reducing effects of larger turbines were being falsely attributed to turbine OEM experience by the primary OLS specifications. This only further suggests that turbine rating is one of the most important CAPEX reduction drivers, or, at the very least, is more significant than learning-by-doing.[5]

7 Discussion

This section interprets the empirical estimates of θ in the context of previous learning curve analyses and separately uses turbine size decision models to contextualize the empirical estimates of δ, which represents the association between turbine size and CAPEX.

7.1 Learning Curve Interpretation

Even though the estimates of developer and turbine OEM learning-by-doing are not statistically significant once I account for the selection concern, it is still instructive to interpret the estimates of learning-by-doing by turbine OEMs from Table 2, especially so they can be compared to learning estimates for other renewables and the cost reducing impact of turbine capacity. In its simplest, single-factor form, as explained by Ibenholt (2002), the learning curve can be formulated as

$$log(C_t) = \alpha log(Q_t) + log(C_0) + \epsilon_t, \tag{6}$$

[5]Table 2, Table 3, and Table 4 are based on commodity prices and wages in the year of full commissioning of the observed wind farms. It could be argued that, given the lead time in building these projects, the commodity price and wage data should be based in the year that construction begins for the observed wind farms. Appendix A.3 presents versions of Table 2, Table 3, and Table 4 in which the data for these control variables is from the year that construction commences. The results do not change significantly.

where C_t and Q_t are the cost and cumulative quantity produced of a good at time t, and C_0 is the initial cost, or the cost of the first unit produced. Learning curve analysis allows researchers to identify simple correlations between industry or firm experience and production costs. Even though my estimations of the coefficients of interest β and θ take into account other factors related to reducing costs, they are analogous to α, which is often called the learning parameter. The learning rate (LR), which uses the learning parameter to estimate how much costs fall with each doubling in cumulative capacity produced, can be mathematically derived as

$$LR = 1 - 2^{\alpha}. \tag{7}$$

The estimates of θ in Table 2 thus translate to a learning rate of around 1.4 percent, indicating that a doubling in the cumulative MW sold by a turbine manufacturer is associated with a 1.4 percent decrease in the CAPEX per watt of an offshore wind farm to which it is selling its wares.

This estimate appears relatively low and certainly is when compared to the results of previous empirical studies of learning rates among renewable technologies, which vary heavily but tend to fall in the range of 5 to 20 percent (Rubin et al. 2015). It is worth noting that most of these studies employ simple one-factor learning curves or two-factor learning curves. These papers likely introduce some upward bias to their estimates of energy technology learning rates by failing to account for economies of scale and exogenous technological change, as argued by Nordhaus (2014). These studies also tend to analyze the effects of industry-wide experience, rather than accounting for developer-specific or OEM-specific experience as I do in my empirical strategy. Van Der Zwaan et al. (2012), who conduct perhaps the only prior study that focuses specifically on learning in offshore wind energy, find an industry-wide learning rate of approximately 3 percent using data on offshore wind facilities fully commissioned by 2008 and after controlling for commodity prices but not for turbine size, water depth, and other factors. This estimate is closer to the ones presented in Table 2.

As mentioned above, the 1.4 percent learning rate loses statistical significance when the selection concern is accounted for. In contrast, using the coefficients on turbine rating in Table 4 and a formula analogous to that of the learning rate, a doubling in turbine capacity is associated with a 19 percent decrease in unitized CAPEX. Even if the 1.4 percent learning rate were statistically significant, it is substantially smaller, by an order of magnitude, than the cost reductions associated with larger turbines.

7.2 Turbine Rating Interpretation

The robustness of the estimates of the coefficient on turbine rating likely confirms the conventional wisdom that turbine size is an important CAPEX reduction driver, especially when compared to the impact of developer and OEM experience. Two-factor learning curve analyses of various renewables, which, ceteris paribus, display the relative returns to cumulative experience and research and development spending, typically show that research spending is the more important cost reduction driver (Rubin et al. 2015). A cursory glance at the results in Table 2 and Table 4 would suggest that the offshore wind sector is no different, with research-driven turbine innovation being more significant than turbine OEM and developer experience in reducing costs.

These are interesting results given the context of how developers and turbine manufacturers choose turbine size to maximize profits at different sites.[6] As part of an ongoing research project, Richard Sweeney and Thomas Covert built a simple model representing this decision (Sweeney 2020). For the sake of simplicity, the model only considers the turbine component of a wind farm and removes the distinction between the developer and turbine OEM by considering a fully vertically integrated wind power company's profit maximization problem.[7] For each turbine, the company seeks to maximize profit π,

$$\pi = pr^2v^3 - \omega r^3, \tag{8}$$

where p is the wholesale, power purchase agreement price, or feed-in tariff received per unit of electricity produced, r is the radius of the circular area swept by the turbine blades, v is the wind speed, and ω represents the manufacturing technology. The first term on the right-hand side represents revenue, and the second term represents the manufacturing cost. Up to a certain limit, power production from a wind turbine increases with the cube of the wind speed and the square of the radius (Danish Wind Industry Association 2003). By the square-cube law, while the area swept by the blades increases with the square of the radius, the volume increases with the cube of the radius, explaining why r is cubed in the

[6] Profit maximization is not always the developer's objective. A developer can only be profit maximizing if it knows the electricity price it will be offered for its output. 119 of the 124 offshore wind farms in the sample received government-set feed-in tariffs, so the profit-maximizing approach generally makes sense. However, the remaining projects in the sample and most of the facilities currently under construction and in earlier stages of development originated from competitive auctions. As described in Section 2.2, under competitive procurement regimes, the developer must submit the lowest LCOE bid in order to secure the project site in the first place. In turn, developers often hold RFPs directed at turbine OEMs, seeking to simply minimize the LCOE of electricity produced by the turbines. It is thus instructive to investigate the turbine size implications of a LCOE minimization model, which is done in Appendix A.4.

[7] Of course, developers and turbine manufacturers are separate entities in the real world. Still, Sweeney and Covert find that the optimal turbine size suggested by this model is also obtained from a second price procurement auction model that involves turbine OEMs bidding into an auction held by a distinct developer. It is worth noting that their research focuses on the onshore wind industry, however.

manufacturing cost term of the equation. It has been empirically observed that turbine OEMs face diseconomies of scale in making turbines larger than 1 MW (Samadi 2016).

I make one modification to this model such that it can better represent the particularities of the offshore wind sector. I include an installation cost term that is separate from the manufacturing cost term. This is necessary because of how significant installation CAPEX is for offshore wind, especially when compared to onshore wind and other renewables. I assume a constant installation cost I for each turbine regardless of its size, which is plausible because the marginal cost of creating a larger turbine foundation likely pales in comparison to the upfront cost of drilling to create the foundation in the first place. The basic Sweeney and Covert model then becomes

$$\pi = pr^2v^3 - \omega r^3 - I. \tag{9}$$

I take the first-order condition with respect to r, and, after solving for optimal r^*, I find that

$$r^* = \frac{2pv^3}{\omega}. \tag{10}$$

This solution indicates that, conditional on manufacturing technology, the size of the turbine should theoretically increase with the price of electricity and wind speed at a site. This helps explain why offshore turbines are larger than their onshore counterparts, as wind speeds are substantially higher at sea and, until recently, the feed-in tariffs offered to offshore wind facilities were much greater than those provided to other renewables, let alone wholesale market electricity rates.

It is less clear why offshore wind turbines have been getting larger over time. With offshore wind feed-in tariffs decreasing and the shift to more competitive procurement regimes, p has likely fallen on average. With regards to wind speed, or v, the locations of recent offshore wind farms are little different from those of their predecessors. This result for optimal r^* would thus imply that ω has decreased, indicating that the manufacturing technology has improved. This mechanism is plausible but would require a completely different strand of research, one focused on recovering the cost functions of turbine OEMs, to be confirmed.

The most important shortcoming of this model, however, is that I is not even present in the expression for r^*. The reason why is that Equation (10) treats I as a fixed cost, which is a result of the model's focus on maximizing profits for an individual turbine rather than for the whole offshore wind farm. The absence of I is problematic because it is widely agreed that, for a given total installed capacity at an offshore wind farm, larger turbines reduce total construction costs by requiring fewer installations in aggregate. Fewer installations means fewer turbine foundations, array cables, and other balance-of-plant components (Snyder and Kaiser 2009; Van Hoof and Velthuijsen 2018; Musial et al. 2019; New York Power Authority 2019; Jennings et al. 2020; Taylor et al. 2020). In essence, due to economies of scale, 6 MW turbines require half as many installations as 3 MW turbines, and the overall project can be installed more quickly.

To fully interpret my estimates of the coefficient on turbine rating, it will thus be necessary to use this model to visualize unitized CAPEX. To do this, I divide the sum of manufacturing and installation CAPEX by r^2, which is a good approximation of the turbine capacity in watts as this quantity usually scales with the surface area swept by the blades:

$$\frac{\omega r^3 + I}{r^2} = \omega r + \frac{I}{r^2}. \tag{11}$$

The result implies that, assuming the same manufacturing and installation technology across turbines of different sizes, it is unclear if larger turbines actually lower CAPEX. The model would suggest that the per-watt manufacturing CAPEX is higher and per-watt installation CAPEX lower for larger turbines, and the aggregate effect depends on the relative magnitude of the increase and decrease. My empirical estimates of the coefficient on turbine rating suggest that larger turbines are associated with lower CAPEX per watt, theoretically implying that the decrease in installation CAPEX has historically dominated the increase in manufacturing CAPEX in magnitude. As mentioned above, there are tremendous scale economies from installing larger turbines, so further research investigating the importance of this potential causal mechanism in explaining the relationship between turbine size and costs may be useful.

This result presents an interesting empirical puzzle. If larger turbines reduce total CAPEX, then developers should always select the largest available turbines when developing any given site. Developers would not face an increased unitized CAPEX and would actually pay even less upfront in order to take advantage of the long-term benefits provided by larger turbines, such as lower operations and maintenance costs and access to the higher wind speeds at greater altitudes (Scheu and Stegelmann 2019).[8] This conclusion seems too good to be true and is belied by the data. Figure 5, which plots the average turbine rating

[8] v could also thus be considered a function of r.

of each offshore wind farm against the year in which the facility was fully commissioned, shows that, even though there has been a clear historic trend towards larger turbines, many developers were still opting for relatively small (<4 MW) turbines into the late 2010s.

Figure 5: Turbine Rating Over Time

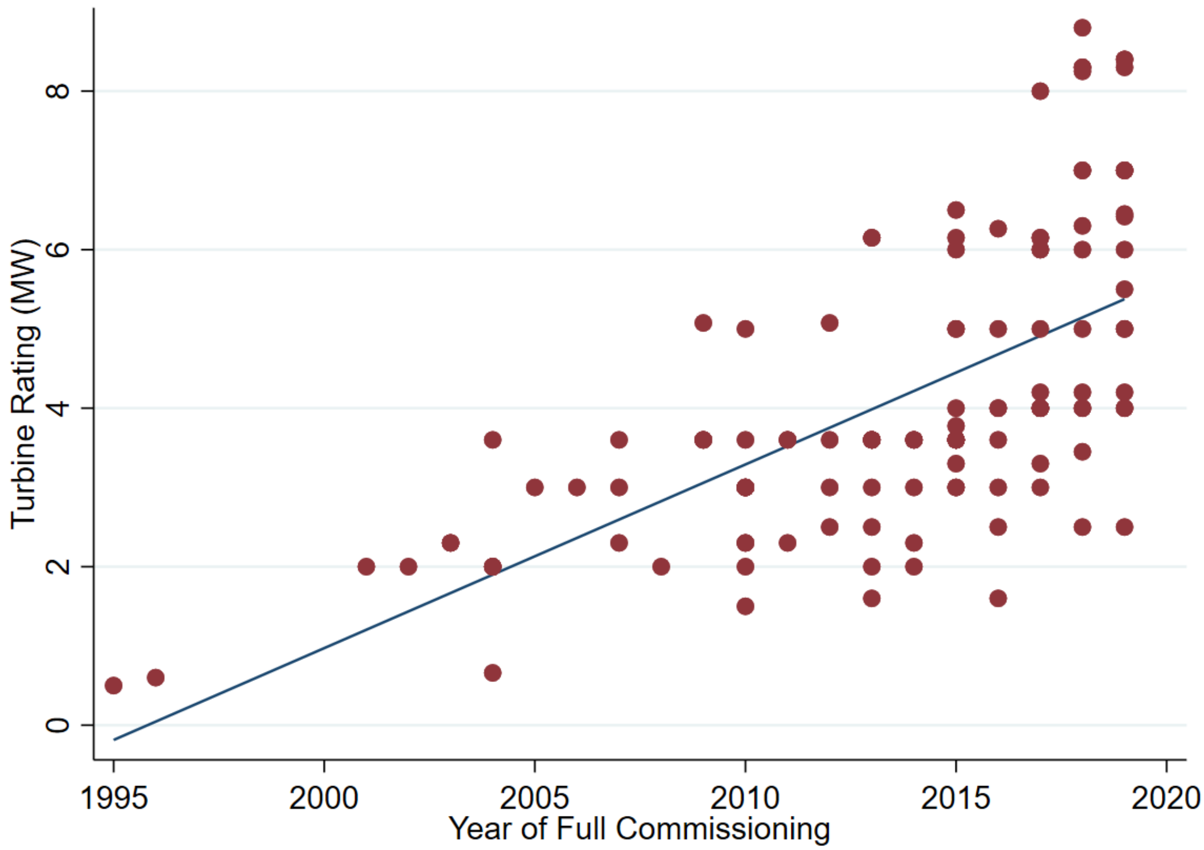

Notes: Each point corresponds to a unique wind farm, plotting its average turbine capacity in MW against the year when it was fully commissioned.

I turn to insights from industry stakeholders and researchers to understand this inconsistency. During interviews, representatives from Ørsted, Northland Power,[9] Siemens Gamesa, and BNEF note that one major reason some developers have chosen smaller turbines recently is because there is inherent risk in selecting newer, larger turbines that have less of a track

[9]Another developer.

record in real-world applications. Given that higher-capacity turbines often have price tags upwards of 10 million dollars, even a slightly risk-averse developer may not want to be an early adopter. Other explanations provided by these stakeholders and academics include governmental permitting constraints, often stemming from environmental considerations, and the fact that Chinese turbine OEMs, which take advantage of a strong home market bias, are lagging behind their European rivals in turbine development. This has resulted in fully commissioned Chinese offshore wind farms usually using turbines with ratings below 5 MW.

8 Extension: Future Cost Predictions

In this section, I test the predictive power of the CAPEX model from the selection assumption test, the results of which are detailed in Table 4. Specifically, I compare the model's CAPEX estimates for future offshore wind farms against estimates provided by project developers in the 4C Offshore dataset. I use the selection test specifications because they identify omitted variable bias in the regular OLS specifications' estimates of the coefficients on developer and OEM experience. I specifically use the model detailed in Column (2) because it does not include commodity prices other than steel, which appear to have little predictive power, and includes the year fixed effects instead of the onshore wind CAPEX proxy for exogenous technological change.

I test the model on all non-demonstration wind farms that are projected to be fully commissioned in the five years after the end of the sample, which is a period that runs from 2020 to 2024. I only include wind farms that have data for all of the relevant variables in the model, such as turbine rating, project size, and water depth. I also exclude floating wind farms, as they utilize a significantly different technology. I remove any facilities that do not have total CAPEX estimates in the 4Coffshore data such that I can compare CAPEX estimates from my model to those quoted by the developers. It is worth noting that developer quotes may not accurately estimate a wind farm's cost before it is fully commissioned, which is one of the main reasons why the sample used for the primary analyses in this paper only includes wind farms fully commissioned by the end of 2019. All of these sampling restrictions leave 66 offshore wind farms that are projected to be fully commissioned between 2020 and 2024.

Regarding the predictive model itself, I rerun the model displayed in Column (2) of Table 4 without the steel price, wage, and developer experience variables. I exclude steel price and wage because of a lack of reliable future data for these variables. I exclude developer experience because of its negligible effect on CAPEX and difficulties in measuring it for future offshore wind farms. Figure 6 displays the in-sample and out-of-sample CAPEX estimates from the model, the actual in-sample CAPEX values, and the developer quotes for the out-of-sample CAPEX values. Table 5 has summary statistics on the percent error of the out-of-sample predictions.

Figure 6 and Table 5 together suggest that the prediction model is not particularly accurate outside of the sample, with a mean percent error of 26 percent. The model appears to systematically underestimate the CAPEX for future wind farms. One explanation may be that larger turbines will not continue to be as strongly associated with reduced costs in the future as they were in the pre-2020 data that the prediction model is based on. Even

Figure 6: CAPEX Predictions vs. Developer Quotes

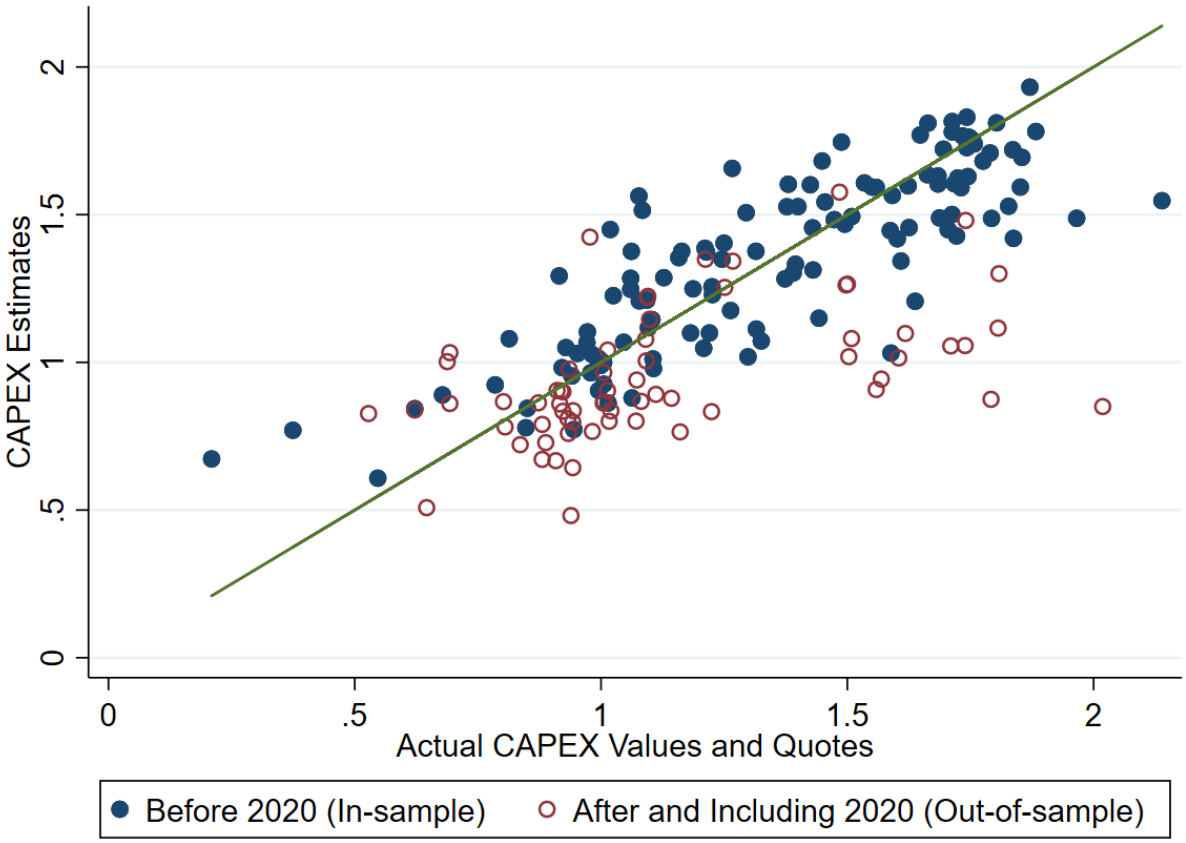

Notes: Each point corresponds to a unique wind farm, plotting the prediction model's CAPEX estimate against the actual CAPEX value (if in-sample) or the developer's quoted CAPEX estimate (if out-of-sample). Plotted values are logarithmic transformations of unitized CAPEX, which in turn is in terms of 2019 USD.

Table 5: Prediction Error Summary Statistics

Variable	Obs	Mean	Std. Dev.	Min	Max
Prediction Error	66	26.077	15.022	.394	62.132

Notes: Error measured as absolute value of the percent error.

with the prediction error, the clustering of the out-of-sample predictions and the developer quotes around a unitized CAPEX of around \$2.70 (after taking the antilogarithmic transformation) suggests that global offshore wind CAPEX is continuing its post-2013 downward trend. This potentially suggests that more massive turbines will still be an important CAPEX reduction force in the near future.

9 Conclusion

As far as I know, this is the first paper to test in an econometrically rigorous manner whether there is learning-by-doing in the offshore wind industry. It is also among the first papers to empirically document the relationship between offshore turbine rating and costs. Using a cost minimization model of the development process, I measure the returns to experience for both offshore wind developers and turbine manufacturers, finding evidence of no statistically significant learning-by-doing among either developers or turbine manufacturers once I account for the empirical threat of market selection for low-cost firms. In contrast, higher-capacity turbines are strongly associated with unitized CAPEX reductions across all of the specifications. A doubling in average turbine rating is associated with a 19 percent decrease in CAPEX per watt of installed capacity, corroborating the consensus among industry stakeholders and research analysts about the importance of larger turbines. Further research building off of the models presented in Section 7.2 will be important in establishing the potential causal mechanisms behind this relationship. Together, these results suggest that the common argument that aggressive policy-driven deployment of renewables can be justified because it takes advantage of a steep learning curve may hold less weight when it comes to offshore wind policy design.[10] As a result, many of the demand-pull policies

[10] As mentioned in Section 3.2, welfare-maximizing governments should only subsidize output from industries on the basis of learning-by-doing if there exists non-appropriable learning, which is a positive externality. The lack of firm-specific, or appropriable, learning in the offshore wind industry as indicated by the results of this study makes it very unlikely that there has been non-appropriable learning in the sector.

that are currently in place in offshore powerhouses like the United Kingdom and the European Union, such as future capacity targets, may require an alternative justification in these mature markets.

This is not to say that governments should play no role in the sector. In nascent markets like that in the United States, states and the federal government will play an important role in preparing grid transmission for the massive projects scheduled to be built in the next ten years. Demand-side national or subnational build-out plans may actually still be relevant in such countries without developed offshore wind supply chains, as aggressive capacity targets in the 2002-2011 period in Northwestern Europe helped the region's industry secure critical investment for upgrading its supply chain. For all countries interested in offshore wind, government research and development spending with regards to turbine innovation may be prescient given the cost reducing impacts of larger turbines. Governments could also facilitate greater knowledge sharing among prominent turbine OEMs, such as Siemens Gamesa, Vestas, and Goldwind, in order to accelerate the development of higher-rated turbines. Such policies would not be unprecedented, as European governments like Denmark's have previously supported common large-scale test facilities for various offshore wind components and turbines (Gottlieb et al. 2019).

It is possible that the learning rate, which is an inherently dynamic concept, may actually manifest as the industry continues to mature. Several competitively procured projects that have secured revenue offtake structures but are currently seeking permitting or are under construction (and were thus excluded from my sample) have won government-organized auctions with extremely low bids, suggesting that the post-2013 trend of falling CAPEX is only set to accelerate (Evans 2019). When they come into commercial operation as early as 2023, some of these projects are expected to produce cheaper electricity than existing natural-gas-fired power plants. The results of this paper would imply that these bids are most likely being driven down by contracts for larger turbines rather than by any learning effect. Regardless, policy will need to shift accordingly to maximize deployment in order to advance climate mitigation while minimizing costs faced by ratepayers.

A Appendix

A.1 Data

To create the sample, I first filter the 4C Offshore Wind Farm Online Database for non-demonstration projects fully commissioned by the end of 2024 and copy and paste the remaining rows into a new Excel spreadsheet called "4Coffshore Stata_v2.xlsx." The Stata .do file located here details the data cleaning steps, which include correcting inaccurate/empty observations and merging in external control variable datasets. In order to calculate the developer and OEM experience variables, the .do file outputs an intermediary Excel spreadsheet labelled "developer_experience_raw.xlsx," which should be fully copied and pasted into cell A1 of another spreadsheet labelled "developer_experience.xlsx." For the Project Size Exogeneity Test (Table 3), "developer_experience_raw.xlsx" should be fully copied and pasted into cell A1 of another spreadsheet labelled "developer_experience_v2.xlsx." The two "non-raw" spreadsheets are read in sequentially by the Stata .do file. Versions of these two spreadsheets with the formulae but without the data (because of proprietary concerns) are located here. The publicly available datasets, which are the currency converter, FRED cement manufacturing PPI dataset, and the IRENA Onshore Wind CAPEX dataset, are also located here.

A.2 Total (Non-Unitized) CAPEX Alternative Results

The purpose of this appendix is to quell concerns about q_i appearing on both sides of Equation (5). Unitized CAPEX is used throughout the main paper for the sake of interpretability and compatibility with prior literature. Table 6, Table 7, and Table 8 are analogous to Table 2, Table 3, and Table 4, respectively, except these tables use total CAPEX, rather than unitized CAPEX/W, as the dependent variable. As a result, the specifications in these tables only place q_i on the right-hand side of the equation. Even with this approach, however, the results do not change significantly. The magnitude and direction of the coefficients on the variables of interest, which are developer experience, turbine OEM experience, and turbine rating, are basically identical to their counterparts in Table 2, Table 3, and Table 4. The most notable change is that the coefficient on project capacity now indicates that increasing offshore wind farm size by 1 percent is associated with approximately a 1 percent increase in CAPEX. This close relationship is a result of the fact that the coefficient on project capacity now accounts for the direct relationship between size and total costs in addition to scale economies.

Table 6: Total CAPEX OLS Results

	(1)	(2)	(3)
VARIABLES		log(CAPEX)	
log(Developer Experience)	0.001	0.000	0.007
	(0.007)	(0.007)	(0.008)
log(OEM Experience)	-0.016**	-0.018**	-0.016**
	(0.007)	(0.007)	(0.007)
log(Average Turbine Rating)	-0.222***	-0.263***	-0.214***
	(0.075)	(0.076)	(0.081)
log(Project Capacity)	0.959***	0.951***	0.980***
	(0.035)	(0.034)	(0.030)
log(Steel Price)	0.090	-0.023	0.241**
	(0.242)	(0.126)	(0.109)
log(Copper Price)	0.119		
	(0.374)		
log(Wage)	0.171***	0.181***	0.122***
	(0.045)	(0.044)	(0.042)
log(Oil Price)	-0.164		
	(0.208)		
log(Cement Price)	-0.758		
	(0.581)		
Water Depth	0.013***	0.012***	0.013***
	(0.003)	(0.003)	(0.003)
Competitive Procurement	-0.054	-0.053	-0.097
	(0.124)	(0.123)	(0.135)
Average Onshore CAPEX/W			0.129
			(0.208)
Observations	124	124	124
Adjusted R-squared	0.971	0.971	0.966
Country Fixed Effects	Y	Y	Y
Time Fixed Effects	Y	Y	N

Notes: Table reports regression coefficients with robust standard errors. Developer and OEM Experience are measured in total MW of capacity developed and total MW of turbines sold prior to the observed wind farm, respectively. Steel price, copper price, oil price, cement price, and weighted average onshore CAPEX are aggregated globally and annually. Wage is aggregated by country and annually. The "Competitive Procurement" dummy equals 1 if the project received revenue under a competitive procurement, 0 otherwise. All CAPEX and price variables are in 2019 USD. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$.

Table 7: Total CAPEX Project Size Exogeneity Test Results

	(1)	(2)	(3)
VARIABLES		log(CAPEX)	
Developer Experience	-0.002	-0.002	0.001
	(0.005)	(0.005)	(0.006)
OEM Experience	-0.003***	-0.003***	-0.003**
	(0.001)	(0.001)	(0.001)
log(Average Turbine Rating)	-0.178**	-0.214***	-0.176**
	(0.073)	(0.073)	(0.078)
Observations	124	124	124
Adjusted R-squared	0.971	0.971	0.965
Country Fixed Effects	Y	Y	Y
Time Fixed Effects	Y	Y	N

Notes: Table reports regression coefficients with robust standard errors. The three specifications are identical to those in Table 6 except developer and OEM experience are measured in total projects developed and total projects to which turbines were sold prior to the observed wind farm, respectively. These experience variables are not logarithmically transformed. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$.

Table 8: Total CAPEX Selection Test Results

	(1)	(2)	(3)
VARIABLES		log(CAPEX)	
log(Developer Experience)	0.001	0.000	0.008
	(0.008)	(0.008)	(0.009)
log(OEM Experience)	-0.010	-0.021	-0.003
	(0.023)	(0.022)	(0.022)
log(Average Turbine Rating)	-0.261***	-0.302***	-0.267***
	(0.082)	(0.078)	(0.082)
Observations	124	124	124
Adjusted R-squared	0.971	0.971	0.965
Country Fixed Effects	Y	Y	Y
Time Fixed Effects	Y	Y	N
Turbine Manufacturer Fixed Effects	Y	Y	Y

Notes: Table reports regression coefficients with robust standard errors. The three specifications are identical to those in Table 6 except with turbine manufacturer fixed effects added in. Dummies were included for the turbine manufacturers Siemens Gamesa, Vestas, Adwen, Senvion, Goldwind, and Envision, which account for 95 percent of capacity installed in the sample. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$.

A.3 Construction Year Alternative Results

Table 9, Table 10, and Table 11 are analogous to Table 2, Table 3, and Table 4, respectively, except these tables use commodity price and wage data from the year in which construction commences (instead of the year of full commissioning) for the observed offshore wind farms. It could be argued that this approach to incorporating the control variable data is better because of the lead time in construction offshore wind farms, which leads to a lag between when capital components are ordered and a project is actually completed. Even with this approach, however, the results do not change significantly. The magnitude of the coefficients on the variables of interest is only slightly different. The most notable change is that the coefficient on developer experience is now negative in all three tables, which may suggest that developers have experienced some learning-by-doing. These estimates are not statistically significant, though.

Table 9: Construction Year OLS Results

VARIABLES	(1)	(2)	(3)
	log(CAPEX/W)		
log(Developer Experience)	-0.012	-0.013	-0.006
	(0.008)	(0.008)	(0.009)
log(OEM Experience)	-0.023***	-0.024***	-0.017**
	(0.008)	(0.007)	(0.007)
log(Average Turbine Rating)	-0.294***	-0.306***	-0.258***
	(0.087)	(0.086)	(0.078)
log(Project Capacity)	-0.036	-0.035	-0.013
	(0.033)	(0.033)	(0.031)
log(Steel Price)	-0.122	0.065	0.368***
	(0.290)	(0.115)	(0.082)
log(Copper Price)	0.102		
	(0.432)		
log(Wage)	0.238***	0.245***	0.198***
	(0.043)	(0.042)	(0.047)
log(Oil Price)	0.230		
	(0.190)		
log(Cement Price)	-0.373		
	(0.792)		
Water Depth	0.013***	0.012***	0.012***
	(0.003)	(0.003)	(0.003)
Competitive Procurement	-0.012	-0.023	-0.069
	(0.141)	(0.134)	(0.161)
Average Onshore CAPEX/W			0.011
			(0.201)
Observations	124	124	124
Adjusted R-squared	0.691	0.695	0.643
Country Fixed Effects	Y	Y	Y
Time Fixed Effects	Y	Y	N

Notes: Table reports regression coefficients with robust standard errors. Developer and OEM Experience are measured in total MW of capacity developed and total MW of turbines sold prior to the observed wind farm, respectively. Steel price, copper price, oil price, cement price, and weighted average onshore CAPEX are aggregated globally and annually. Wage is aggregated by country and annually. The “Competitive Procurement” dummy equals 1 if the project received revenue under a competitive procurement, 0 otherwise. All CAPEX and price variables are in 2019 USD. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$.

Table 10: Construction Year Project Size Exogeneity Test Results

	(1)	(2)	(3)
VARIABLES		log(CAPEX/W)	
Developer Experience	-0.006	-0.006	-0.003
	(0.005)	(0.005)	(0.007)
OEM Experience	-0.003**	-0.003***	-0.002*
	(0.001)	(0.001)	(0.001)
log(Average Turbine Rating)	-0.229***	-0.242***	-0.215***
	(0.084)	(0.084)	(0.075)
Observations	124	124	124
Adjusted R-squared	0.677	0.678	0.637
Country Fixed Effects	Y	Y	Y
Time Fixed Effects	Y	Y	N

Notes: Table reports regression coefficients with robust standard errors. The three specifications are identical to those in Table 9 except developer and OEM experience are measured in total projects developed and total projects to which turbines were sold prior to the observed wind farm, respectively. These experience variables are not logarithmically transformed. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$.

Table 11: Construction Year Selection Test Results

	(1)	(2)	(3)
VARIABLES		log(CAPEX/W)	
log(Developer Experience)	-0.014	-0.014	-0.007
	(0.009)	(0.009)	(0.010)
log(OEM Experience)	-0.011	-0.017	-0.001
	(0.020)	(0.019)	(0.020)
log(Average Turbine Rating)	-0.330***	-0.343***	-0.301***
	(0.094)	(0.094)	(0.083)
Observations	124	124	124
Adjusted R-squared	0.684	0.686	0.639
Country Fixed Effects	Y	Y	Y
Time Fixed Effects	Y	Y	N
Turbine Manufacturer Fixed Effects	Y	Y	Y

Notes: Table reports regression coefficients with robust standard errors. The three specifications are identical to those in Table 9 except with turbine manufacturer fixed effects added in. Dummies were included for the turbine manufacturers Siemens Gamesa, Vestas, Adwen, Senvion, Goldwind, and Envision, which account for 95 percent of capacity installed in the sample. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$.

A.4 LCOE Version of Sweeney and Covert Model

In this section, I modify the Sweeney and Covert model to find the optimal turbine size assuming the developer purchasing the turbines is seeking to minimize the LCOE rather than maximize profits at a given project site. As a reminder, LCOE is a measure of the average present-value cost of each unit of electricity produced by a power plant over its lifetime. In the competitive procurement regimes that secured revenue offtake structures for 5 of the 124 wind farms in the sample and for most of the offshore wind farms currently under development or construction, developers win by submitting the lowest LCOE-based bid. In these cases, cost (and specifically LCOE) minimization is more important than profit maximization.

To roughly approximate the LCOE, I simply divide the CAPEX associated with the turbine by its projected power output and ignore the discount rate and OPEX inputs in the conventional LCOE formula. As mentioned in Section 2.1, the largest component of LCOE is CAPEX, so this approximation is still valid. Finding the turbine size that minimizes this estimate of the LCOE involves writing an equation like

$$argmin_r \frac{\omega r^3 + I}{r^2 v^3}. \tag{12}$$

Taking the FOC and solving for r^*, I find that

$$r^* = \sqrt[3]{\frac{2I}{\omega}} \; s.t. \; \frac{\omega r^{*3} + I}{r^{*2} v^3} \leq p. \tag{13}$$

The constraint ensures that the capital investment in the wind turbine at least breaks even over the lifetime of the wind farm. The constraint checks that the LCOE is no greater than the electricity price the developer receives assuming the developer wins the auction. This result is similar to Equation (11) in that r^* is decreasing in ω. One difference is that r^* no longer depends explicitly on v, which is interesting because v is a major determinant of a wind farm's output over its lifetime. The exclusion of v from r^* is perhaps an artifact of the simplicity of the model. The key improvement from the profit maximization model is that r^* is now increasing in I, meaning that, as per-turbine installation costs increase, it is better to have larger turbines. This follows from the logic described in Section 7.2, as, given a fixed capacity for the entire wind farm, larger turbines require fewer installations in aggregate.

References

Abernathy, William J. and Kenneth Wayne. 1974. "Limits of the Learning Curve" *Harvard Business Review* 52 (9).

Anderson, John W., Gordon W. Leslie, and Frank A. Wolak. 2019. "Measuring the Impact of Own and Others' Experience on Project Costs in the U.S. Wind Generation Industry." National Bureau of Economic Research Working Paper 26114.

Benkard, C. Lanier. 2000. "Learning and Forgetting: The Dynamics of Aircraft Production." *The American Economic Review* 90 (4): 1034-54.

Bollinger, Bryan and Kenneth Gillingham. 2019. "Learning-by-Doing in Solar Photovoltaic Installations." *SSRN Electronic Journal.*

BVG Associates. 2019. "Guide to an Offshore Wind Farm." *BVG Associates.* https://bvgassociates.com/cases/guide-offshore-wind-farm/

Crabtree, Christopher J., Donatella Zappalá, and Simon I. Hogg. 2015. "Wind Energy: UK Experiences and Offshore Operational Challenges." *Proceedings of the Institution of Mechanical Engineers. Part A, Journal of power and energy* 227 (7): 727-46.

Danish Wind Industry Association. 2003. "The Power of the Wind: Cube of Wind Speed." *Danish Wind Industry Association.* http://xn–drmstrre-64ad.dk/wp-content/wind/miller/windpower%20web/en/tour/wres/enrspeed.htm

Durakovic, Adnan. 2020a. "German Cabinet Approves 40 GW by 2040 Offshore Wind Target." *OffshoreWIND.Biz.* https://www.offshorewind.biz/2020/06/03/german-cabinet-approves-40-gw-by-2040-offshore-wind-target/

Durakovic, Adnan. 2020b. "UK Prime Minister: Offshore Wind to Power Every Home by 2030." *OffshoreWIND.Biz.* https://www.offshorewind.biz/2020/10/06/uk-prime-minister-offshore-wind-to-power-every-home-by-2030/

Energy Information Administration. 2020. "Electricity Explained." *Energy Information Administration.* https://www.eia.gov/energyexplained/electricity/electricity-in-the-us-generation-capacity-and-sales.php

Evans, Simon. 2019. "Analysis: Record-low Price for U.K. Offshore Wind Cheaper Than Existing Gas Plants by 2026." *Carbon Brief.* https://www.carbonbrief.org/analysis-record-low-uk-offshore-wind-cheaper-than-existing-gas-plants-by-2023

Garcia Da Fonseca, Leila and Daniel Liu. 2019. "A Closer Look at Siemens Gamesa's Deal for Senvion." *Greentech Media.*

https://www.greentechmedia.com/articles/read/closer-look-at-siemens-gamesas-deal-for-senvion

Gottlieb, Magnus H., Anders H. Nymark, Emil D. Grann, and Johannes Bøggild. 2019. "Making Green Energy Affordable." *Ørsted.* https://orsted.com/en/about-us/whitepapers/making-green-energy-affordable

Ibenholt, Karin. 2002. "Explaining Learning Curves for Wind Power." *Energy Policy* 30 (13): 1181-89.

Jamasb, Tooraj. 2007. "Technical Change Theory and Learning Curves: Patterns of Progress in Electricity Generation Technologies." *The Energy Journal* 28 (3): 51-71.

Jansen, Malte, Iain Staffell, Lena Kitzing, Sylvain Quoilin, Edwin Wiggelinkhuizen, Bernard Bulder, Iegor Riepin, and Felix Müsgens. 2020. "Offshore Wind Competitiveness in Mature Markets Without Subsidy." *Nature Energy* 5 (8): 614-22.

Jennings, Tom, Helen A. Tipper, Jonathan Daglish, Michael Grubb, and Paul Drummond. 2020. "Policy, Innovation, and Cost Reduction in U.K. Offshore Wind." *The Carbon Trust.* https://www.carbontrust.com/resources/policy-innovation-and-cost-reduction-in-uk-offshore-wind

Karnitschnig, Matthew. 2014. "Germany's Expensive Gamble on Renewable Energy." *The Wall Street Journal* https://www.wsj.com/articles/germanys-expensive-gamble-on-renewable-energy-1409106602

Lehn-Christiansen, Tom. 2017. "The World's First Offshore Wind Farm is Retiring." *Ørsted* https://orsted.com/en/media/newsroom/news/2017/03/the-worlds-first-offshore-wind-farm-is-retiring

McKenna, Phil. 2017. "America's First Offshore Wind Energy Makes Landfall in Rhode Island." *Inside Climate News* https://insideclimatenews.org/news/01052017/block-island-wind-farm-deepwater-wind-renewable-energy-climate-change/

Musial, Walter, Philipp Beiter, Paul Spitsen, Jake Nunemaker, and Vahan Gevorgian. 2019. "2018 Offshore Wind Technologies Market Report." *U.S. Department of Energy.* https://www.energy.gov/eere/wind/downloads/2018-offshore-wind-market-report

Nemet, Gregory F. 2019. *How Solar Energy Became Cheap: A Model for Low-Carbon Innovation.* New York, NY: Routledge.

New York Power Authority. 2019. "Offshore Wind- A European Perspective." *New York Power Authority.* https://www.nypa.gov/news/press-releases/2019/20190807-key-learnings-offshore-wind-transmission-models

Nordhaus, William D. 2014. "The Perils of the Learning Model for Modeling Endogenous Technological Change." *The Energy Journal* 35 (1): 1-13.

Rubin, Edward S., Inês M.L. Azevedo, Paulina Jaramillo, and Sonia Yeh. 2015. "A Review of Learning Rates for Electricity Supply Technologies." *Energy Policy* 86: 198-218.

Scheu, Matti and Alexander Stegelmann. 2019. "OPEX Benchmark – An Insight into Operational Expenditures of European Offshore Wind Farms." *PEAK Wind.* https://peak-wind.com/insights/opex-benchmark-an-insight-into-operational-expenditures-of-european-offshore-wind-farms/

Siemens Gamesa Renewable Energy. 2017. "Company History." *Siemens Gamesa Renewable Energy.* https://www.siemensgamesa.com/en-int/about-us/company-history

Snyder, Brian and Mark J. Kaiser. 2009. "A Comparison of Offshore Wind Power Development in Europe and the U.S.: Patterns and Drivers of Development." *Applied Energy* 86 (10): 1845-56.

Stehly, Tyler, Philipp Beiter, and Patrick Duffy. 2020. "2019 Cost of Wind Energy Review." *National Renewable Energy Laboratory.* https://www.osti.gov/biblio/1756710-cost-wind-energy-review

Stein, Jeremy. 1997. "Waves of Creative Destruction: Firm-Specific Learning-by-Doing and the Dynamics of Innovation." *The Review of Economic Studies* 64 (2): 265-88.

Steinberg, Julie and Joe Wallace. 2021. "Money Managers Look to Blue Seas for Green Investments." *The Wall Street Journal* https://www.wsj.com/articles/money-managers-look-to-blue-seas-for-green-investments-11611060569

Sweeney, Richard L. 2020. "Winds of Change: Technical Progress and Learning in Wind Power." *Harvard Seminar in Environmental Economics and Policy.* 4 November 2020.

Taylor, Michael, Pablo Ralon, Harold Anuta, and Sonia Al-Zoghoul. 2020. "Renewable Power Generation Costs in 2019." *IRENA.* https://www.irena.org/publications/2020/Jun/Renewable-Power-Costs-in-2019

Van Benthem, Arthur, Kenneth Gillingham, and James Sweeney. 2008. "Learning-by-Doing and the Optimal Solar Policy in California." *The Energy Journal* 29 (3): 131-51.

Van Der Zwaan, Bob, Rodrigo Rivera-Tinoco, Sander Lensink, and Paul Van Den Oosterkamp. 2012. "Cost Reductions for Offshore Wind Power: Exploring the Balance between Scaling, Learning and RD." *Renewable Energy* 41: 389-93.

Van Hoof, Jeroen and Jan W. Velthuijsen. 2018. "Unlocking Europe's Offshore Wind Potential: Moving Towards a Subsidy Free Industry." *PricewaterhouseCoopers.* https://www.pwc.nl/en/publicaties/unlocking-europes-offshore-wind-potential.html

Vieira, Mário A., Brian Snyder, Elsa Henriques, and Luis Reis. 2019. "European Offshore Wind Capital Cost Trends up to 2020." *Energy Policy* 129: 1364-71.

Windpower Monthly. 2004. “Danish wind giants head for merger – Vestas and NEG Micon going for global supremacy.” *Windpower Monthly.* https://www.windpowermonthly.com/article/956160/danish-wind-giants-head-merger-vestas-neg-micon-going-global-supremacy

Wright, T.P. 1936. “Factors Affecting the Cost of Airplanes,” *Journal of Aeronautical Sciences* 3 (4): 122–8.