



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

Thomas Covert and Richard L. Sweeney

How much better have we gotten at manufacturing clean energy technologies, and how much of that progress is driven by learning by doing? We answer these questions for wind power using global data from 2000 to 2019, accounting for dramatic improvements in turbine productivity over this period. We find that quality-adjusted manufacturing costs fell by more than 75 percent, with learning by doing playing a major role: each doubling of a manufacturer's experience reduced manufacturing costs by over 25 percent. Knowledge spillovers across manufacturers are modest relative to own-firm learning, but substantial in aggregate.

How much better have we gotten at making clean tech?

Over the past quarter-century, clean energy technologies have transitioned from negligible niche markets to the primary drivers of global electric power and automobile markets. While it's clear that firms' ability to produce clean tech has improved, measuring exactly how much costs have come down is tricky for two reasons:

1. **We observe prices not costs.** Commonly reported metrics of clean tech improvement typically chart the evolution of the *prices* of these goods. However, in these imperfectly competitive industries, prices may differ from costs for strategic reasons.
2. **Device quality has improved a lot.** In order to truly capture manufacturing improvements, we need to account for the fact that today's clean tech is much better than early products. For

example, simply comparing the cost of an EV in 2010 to the cost today would neglect the fact that today's EVs are much better cars.

Wind energy provides a good setting to measure manufacturing progress

Wind turbines are a particularly useful technology to study because simple physics helps us understand both the value and the cost of making better products. The energy a wind turbine captures is proportional to the area swept by its blades. So bigger is better. Figure 1 shows that between 2000 and 2019, the average turbine radius doubled, while output per turbine more than quadrupled.

Given this, it's natural to wonder why turbines weren't bigger to begin with. The answer comes from another simple physical relationship: while the swept area of a turbine grows quadratically with its size, the volume of materials needed to build it grows cubically.




Figure 1. Global Turbine Size and Output Trends.

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

This "square-cube" relationship creates a tradeoff: bigger turbines generate more electricity, but they also cost more to manufacture. The optimal turbine size depends on how efficiently firms can manufacture turbine volume. As manufacturing becomes more efficient, building larger turbines becomes more attractive.

wind speeds into expected electricity production for each wind turbine using manufacturer provided wind turbine power curves. Finally, electricity output is converted into expected revenue using information on feed-in tariffs, power purchase agreements and renewable energy auctions.

Data and methods

We exploit these physical relationships, combined with empirical tools from industrial organization, to recover otherwise unobserved manufacturing costs and measure how they have changed over time. Unlike smaller consumer goods, where we can easily observe the price of each product, wind turbines are massive, expensive pieces of equipment, which are procured in a confidential environment. We model this procurement process as an auction, and leverage auction theory to translate developer choices into unobserved wind turbine manufacturer bids and costs.

We then estimate the model using the universe of commercial wind farm procurement decisions worldwide from 2000 to 2019. For each wind project, we obtain location-specific wind speed distributions from Vaisala. We then translate

Two decades of progress

Figure 2 presents the estimated annual average costs by year. The first panel reports the cost per MWh of capacity. Costs oscillated for the first decade, before ending the sample roughly 30% lower, suggesting relatively modest progress. However, this metric ignores the fact that, due to longer rotor blades, turbines became much more productive during this time. The second panel plots the average capacity factor (electricity output per MWh) over time, which shows that productivity per MWh increased by roughly 45%.

The bottom left panel reports the turbine-only levelized cost of energy, which is the capital cost amortized over lifetime output. This measure, which is a commonly reported way of comparing electricity costs across technologies with different capacity factors, declined by 55%. However, even this measure understates how much better manufacturers




Figure 2. Global average bid trends.




Figure 3. Average estimated bid per volume for top firms.

became at making large turbines. The final panel plots the average cost per unit of turbine material volume. This measure, which is our preferred measure of manufacturers' technical capabilities, fell by over 75%. This implies that producing a 2019 vintage turbine would have cost four times more using year 2000 manufacturing capabilities. If we look within firm, the declines in cost per volume look even larger.

How Much of the Decline is Learning by Doing?

We then ask how much of these large improvements in manufacturing efficiency can be attributed to learning by doing. The belief that clean tech costs would come down as firms gained experience was an important motivation for early renewable subsidy schemes. We find that learning benefits are substantial: Doubling manufacturing experience reduces costs by 26 to 29 percent. These estimates are robust across models that allow for persistent cost differences between manufacturers.

Spillovers: Small per Turbine, Large in Aggregate

Importantly, we find that learning extends beyond the specific turbine model being produced. At least 3 percent of the experience gained producing one turbine model spills over

onto other models produced by the same firm, while roughly 0.5 percent spills over onto turbine models produced by other firms. Put another way, when GE makes one additional unit of turbine X, it learns 33 times more about how to make the next unit of turbine X as it would have from producing one of its other models, and 200 times more than it would have from observing another firm produce a turbine.

The timing of experience also matters for spillovers. Estimates reveal that the majority of the within-firm spillovers, and all of the across-firm spillovers, occurs prior to a turbine model entering production. This is consistent with firms using manufacturing experience to improve the design of next generation turbines to economize on turbine material costs.

While spillovers across firms are small in relative terms, compared to the value of turbine-specific experience, they have economically large implications in aggregate. That is because there are dozens of turbine models on the market in a given year, and the stock of other firms' experience is typically hundreds of times larger than an individual turbine's cumulative sales. Figure 4 presents the log cost of turbine volume with and without cross-firm spillovers for three large manufacturers. When Goldwind, the largest Chinese manufacturer in our data, entered the market in 2005, its limited experience alone would have implied




Figure 4. New Turbine Cost of Size with and without Spillovers.

costs nearly ten times those of the market leader, Vestas. Thanks to spillovers from the rest of the industry, its costs were only about 1.6 times as high. Conversely Nordex, a large German manufacturer, was an industry leader in 2000, but saw its market share gradually decline over time. Had it not benefited from spillovers, its cost disadvantage relative to Vestas would have been three times larger by the end of the sample.

Implications

1. **It is important to account for quality changes when measuring clean tech progress.** While costs per unit of capacity came down over time, today's wind turbines are dramatically better devices than two decades ago. The same is true for batteries, solar panels, electrolyzers, and electric vehicles, though the devices do not offer physics as convenient as wind's.

2. **Learning is important.** The prospect of learning by doing was a motivation for renewable energy subsidies. We find that a doubling of manufacturing experience reduces wind turbine manufacturer cost of volume by more than 25%. Although we do not conduct a counterfactual exercise that allows us to say how much these high costs would have been without these subsidies, the learning magnitudes estimated can rationalize a substantial fraction of the observed cost declines.

3. **Spillovers provide another justification for subsidies.** Subsidies to the adoption of green technology can be justified solely by the existence of unpriced fossil fuel externalities. However, the existence of learning by doing spillovers provides another rationale. We find that one firm's production today reduces other firms' production costs tomorrow, creating future value that firms will not be fully appropriated.

References


Link to the full working paper discussed in this brief:

Covert, T., and Sweeney, R. L. (2026), "Winds of Change: Technical Progress and Learning by Doing in the Wind Industry," [MIT CEEPR Working Paper 2026-19](#), October 2026

About the Authors



Thomas Covert is a Data Scientist at Airbnb, working on pricing guidance models using tools from industrial organization and machine learning. Prior to Airbnb, he was an Assistant Professor at the University of Chicago Booth School of Business, with research focusing on energy economics. He received an S.B. in mathematics from MIT and a PhD in Business Economics from Harvard.



Richard L. Sweeney is an Associate Professor of Economics at Boston College. Sweeney's research lies at the intersection of industrial organization and environmental economics. His current work is centered around understand and measuring innovation in energy markets. Professor Sweeney graduated from Boston College in 2004, and received his Ph.D. in Public Policy from Harvard Kennedy School in 2015.



About the Center for Energy and Environmental Policy Research (CEEPR)

Since 1977, CEEPR has been a focal point for research on energy and environmental policy at MIT. CEEPR promotes rigorous, objective research for improved decision making in government and the private sector, and secures the relevance of its work through close cooperation with industry partners from around the globe. CEEPR is jointly sponsored at MIT by the MIT Energy Initiative (MITEI), the Department of Economics, and the Sloan School of Management.