Winds of Change: Technical Progress and Learning in Wind Power

Thomas R. Covert, University of Chicago Booth School of Business and NBER
Richard L. Sweeney, Boston College

Private industry can reduce the cost of clean energy by developing inherently lower cost technologies (innovating) and by accumulating manufacturing experience with existing technologies (learning-by-doing). As in other sectors, manufacturers will have socially inefficient innovation incentives if new technologies spill over to competitors, or if the gains to business stealing are too high (Van Reenen et al., 2013). And while firms may be able to fully capture the benefits of their own manufacturing experience, the prospect of innovative competitors means that the industry as a whole may innovate *too often*, and in doing so, insufficiently "amortize" the fixed cost of innovations through learning (Stein, 1997). In the energy sector, the welfare costs of these distortions are compounded by the presence of unpriced externalities, which clean energy technologies are specifically designed to avoid.

In this project, we will quantify these inefficiencies and compute optimal supply side policy to correct them in the global market for wind turbines. To do this, we will estimate a structural model of the industry that features both innovation and learning-by-doing. The key structural parameters in this model are the learning scale economies within a wind turbine design, the cost of developing an improved wind turbine design, and the extent to which this cost differs for innovators, who are the first to develop a new turbine design, and for followers, who may benefit from innovation spillovers. We will identify these parameters using decades of rich spatial and temporal variation in demand-side incentives for wind turbines which have encouraged their widespread adoption. Although the market for wind turbines is global, demand exhibits strong home market bias (Coşar et al., 2015), which we use to isolate demand shocks specific to firms headquartered in a given country.

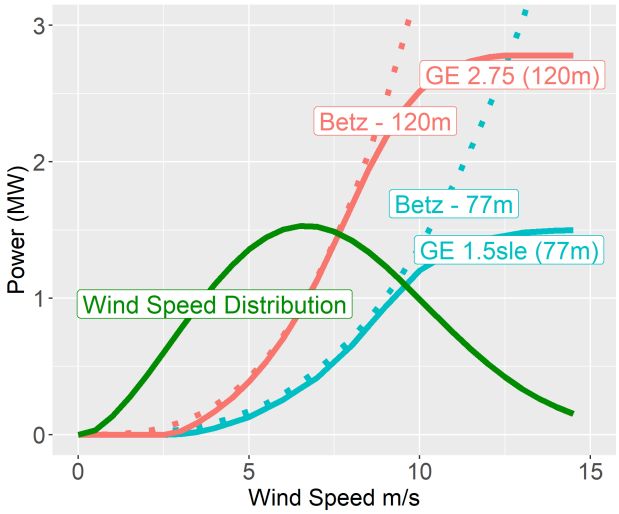
We will use counterfactual predictions from this model to measure how policy has affected the speed and composition of technological progress in wind turbines, and to estimate the gains from alternative, better targeted, policies. Our first counterfactual will compute how much more expensive wind turbines would be, both due to learning effects and innovation effects, in a world without the observed pattern of demand side policies. This number is important because recent research has found that the high costs of these policies cannot be justified by their contemporaneous emissions reductions alone: in order for these demand side policies to be welfare positive, they must have a substantial *causal* impact on learning and innovation. We will also generate counterfactual predictions with policy that intentionally targets innovation and learning, rather than just adoption. This exercise will reveal the costs and characteristics of wind turbines in a world where policy makers augment existing demand side policies with supply side tools explicitly designed to foster technological progress.

We focus on the wind turbine industry because it is now large, covering more than ten percent of global power generation capacity, and because costs have declined significantly over time. Crucially, the technological progress underpinning the industry's rapid growth is relatively simple: *better* wind turbines, those that generate more output per unit of cost, are generally *bigger* wind turbines. This is a consequence of Betz law, which states that the energy produced by a wind turbine of radius r, spinning at a velocity v, is proportional to r^2v^3. Consistent with this physical relationship, since 2000, wind turbine rotors have more than doubled in size, while expected output per turbine has increased nearly fourfold.

The constraints on this growth in rotor size are also physical in nature. Galileo's "square-cube" law, adapted to this setting, states that the volume of materials necessary to produce a rotor of radius r should be proportional to r^3. We use this relationship, combined with Betz law, to characterize the optimal size of a turbine as a function of power prices p, wind speed v, and manufacturing cost per unit of volume c. In particular, the rotor size that maximizes variable profits $pr^2v^3 - cr^3$ is $r^* = \frac{2}{3}\frac{pv^3}{c}$. This tells us that if demand side policy increases wind output prices, firms will respond by producing larger turbines, even without improvements in c coming from learning or innovation. However, it also tells us that if we see the size of turbines increase over time, *conditional* on wind speeds and power prices, it must be that c, the wind turbine manufacturers' costs per unit of turbine volume, have decreased.

We use this insight about optimal turbine sizing to infer firms' unobserved technology states c, from data on the observed distribution of p, v and r. We now have a database of the near-universe of commercial wind farms, worldwide, including each farm's exact location, entry date, and, for most farms, the exact wind turbine model. We combine this data with detailed information about the distribution of wind speeds at each site from Vaisala, a leading commercial provider of wind resource information, and with an extensive list of wind turbine "power curves" which plot out the power generation of a given turbine model across different wind speeds (i.e., the empirical analogue to Betz's law). For a subset of these projects, we also observe the *price* the manufacturer charged the developer. With this data, we are able to compute the expected output at any site under any turbine model (Figure 1). To infer a manufacturer's marginal cost of size, we assume that the turbine it ultimately installed at a project maximizes the surplus generated by that project. Thus, by computing expected revenue for each possible turbine within a manufacturer i's portfolio at time t, we can recover revealed preference estimates of c_{it}.

Figure 1: Example Power Curves



These costs $\{c_{it}\}$ form the core state variables for each wind turbine manufacturer in our dynamic model of pricing, learning, and innovation. In each time period, firms play a Nash pricing game for each possible wind farm project. Because costs decline with experience, firms have an incentive to *underprice*, as in Benkard (2004), and we thus use observed pricing choices to infer learning scale economies. Learning benefits are concave, so firms will eventually elect to pay a fixed "innovation" cost to move to an improved state of technology, and reset their learning progress. For technology leaders, extending the frontier temporarily alleviates price competition in the stage game. However, it also reduces other firms' costs of achieving the frontier state, as rivals can observe the design and material decisions made by the leader. We use this tradeoff to infer the unobserved costs of innovating from the observed timing of new wind turbine introductions.

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

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

Reenen, John Van, Nicholas Bloom, and Mark Schankerman, "Identifying Technology Spillovers and Product Market Rivalry," *Econometrica*, 2013, *81* (4), 1347–1393.

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