Demanding Innovation: The Impact of Consumer Subsidies on Solar Panel Production Costs*

Todd Gerarden†

March 31, 2018

Please click here for the most recent version.

Abstract

This paper analyzes the impacts of consumer subsidies in the global market for solar panels. Consumer subsidies can have at least two effects. First, subsidies shift out demand and increase equilibrium quantities, holding production costs fixed. Second, subsidies may encourage firms to innovate to reduce their costs over time. I quantify these impacts by estimating a dynamic structural model of competition among solar panel manufacturers. The model produces two key insights. First, ignoring long-run supply responses can generate biased estimates of the effects of government policy. Without accounting for induced innovation, subsidies increased global solar adoption 49 percent over the period 2010-2015, leading to over $15 billion in external social benefits. Accounting for induced innovation increases the external benefits by at least 22 percent. Second, decentralized government intervention in a global market is inefficient. A subsidy in one country increases long-run solar adoption elsewhere because it increases investment in innovation by international firms. This spillover underscores the need for international coordination to address climate change.

*I am grateful to Joe Aldy, Myrto Kalouptsidi, Ariel Pakes, Rob Stavins, and Jim Stock for their invaluable guidance and support. I thank Jonathan Baker, Evan Herrnstadt, Robin Lee, Jing Li, Alex MacKay, Marc Melitz, Mark Shepard, Sam Stolper, Rich Sweeney, Elie Tamer, and seminar participants at Harvard, the Northeast Workshop on Energy Policy and Environmental Economics, the London School of Economics, the Federal Reserve Bank of Dallas, the University of Toronto, Cornell, Georgia State University, UC Davis, the University of Maryland, and Resources for the Future for helpful comments and conversations. The data used in this paper were obtained through a data use agreement with IHS Markit. I gratefully acknowledge support from U.S. EPA STAR Fellowship Assistance Agreement no. FP-91769401-0 and Harvard's Joseph Crump Fellowship.

†Harvard University. Email: gerarden@fas.harvard.edu. Website: scholar.harvard.edu/gerarden.

1 Introduction

Can subsidies to consumers spur firms to innovate? In principle, government subsidies to consumers that address non-pecuniary externalities can also facilitate cost reductions for emerging technologies. While economic theory makes predictions about the potential impact of consumer subsidies on cost-reducing innovation, there is little empirical evidence on this relationship to inform policy. Crucially, analysis using static economic methods will understate the impacts of policies that induce technical change.

I study the short- and long-run impacts of consumer subsidies in the market for solar panels, where government intervention is widespread. Solar power is viewed as a key technology for mitigating climate change because it can displace conventional electricity sources that emit greenhouse gases. Over the period 2006-2015, solar panel prices fell by an order of magnitude due to a combination of innovation, input price reductions, and subsidies to producers. Manufacturing output increased, facilitating a 35-fold increase in new solar power capacity in 2015 relative to 2006 (IEA, 2016). Firms in China led this manufacturing expansion, collectively producing two-thirds of the world's solar panels over the period 2010-2015.[1]

To quantify the impacts of consumer subsidies, I formulate a dynamic structural model of firm competition based on Ericson and Pakes (1995). Consumer demand for undifferentiated solar panels is static but depends on subsidies and unobserved demand shocks, both of which vary over time. Incumbent firms compete in quantities (Cournot) to serve demand in each market. I employ a unique observable measure of technological innovation, the electrical conversion efficiency of solar panels, which I refer to as "technical efficiency." Increasing technical efficiency helps firms lower costs by reducing the materials needed to manufacture solar panels. This gives firms an incentive to make fixed investments in technical efficiency to lower their costs in order to increase their profits in the product market. I assume that firms condition only on the current industry state – firms' technical efficiencies, a common input price, and demand – and their own private shocks when making investment decisions, leading to a Markov-Perfect Nash Equilibrium.

To estimate the model, I use market-level data from four regional markets that span the globe: Germany, Japan, the United States, and a residual market, "Rest of the World." I estimate aggregate demand for solar panels (in Watts) in each market to

[1]Chinese government subsidies to producers are suspected to have contributed to the expansion of solar panel manufacturing in China and to have reduced equilibrium prices for solar panels. My empirical strategy accounts for such unobserved supply subsidies, but I do not directly analyze how subsidies to producers affected the solar panel market. This paper focuses on subsidies to consumers.

recover price elasticities and the impact of subsidies on demand. I use the demand estimates and the firm's first order condition for optimal production to recover marginal production costs. I then estimate the relationship between technical efficiency and costs. This relationship is identified using variation in technical efficiency and market shares: under Cournot competition, firms with lower costs have higher market shares, so a positive correlation between technical efficiency and market share implies that technical efficiency reduces costs.

I use a two-step estimator in the spirit of Bajari et al. (2007) to recover the fixed costs of investment.[2] First, I use estimates of the product market model described above, firm investment policies, and state transitions to forward-simulate potential industry paths, tracking firms' profits and investment decisions. This approach leverages the insight that a firm's value function is equivalent to the expected discounted sum of its future net profits. I use the simulated value function and the optimality condition of the firm's investment problem to estimate the fixed cost parameters via maximum likelihood.[3]

The fixed cost parameters are identified by variation in the frequency and the expected benefits of investment in different industry states. Observed firm-level improvements to technical efficiency are used to recover the frequency of investment in a given state, and the product market model predicts the benefits – in terms of future profits – of that investment. The fixed cost parameters rationalize firms' investment decisions: for example, if firms were to invest more frequently than observed at a given state they would increase their gross product market profits, so their decision not to invest more frequently implies a lower bound on the fixed cost of investment. Variation across industry states in investment patterns and profits implied by the product market model helps pin down the parameters. Finally, to allow for strategic interactions among a large number of firms, I employ an equilibrium refinement for counterfactual simulations grounded in the experience-based equilibrium of Fershtman and Pakes (2012) and the moment-based Markov equilibrium of Ifrach and Weintraub (2017).[4]

I find that subsidies to consumers have economically significant positive impacts

[2]Two-step estimators were first introduced for dynamic models by Hotz and Miller (1993) and later extended to dynamic games (e.g., Aguirregabiria and Mira, 2007; Bajari et al., 2007; Pakes et al., 2007).

[3]To be precise, the value function can only be simulated up to the unknown parameters of interest. Conveniently, the value function is linear in parameters, so I forward-simulate industry paths once and then search for the parameters that maximize the likelihood.

[4]I also use techniques used in recent empirical work to allow for continuous states (e.g., Ryan, 2012; Sweeting, 2013; Barwick and Pathak, 2015; Kalouptsidi, 2018).

on demand, consistent with prior work.[5] Firms with higher technical efficiencies have lower marginal production costs on average. This finding is robust to the inclusion of time fixed effects that proxy for unobserved supply shifters – such as production subsidies to Chinese firms – as well as proxies for learning-by-doing and economies of scale.[6] The results imply that improvements to technical efficiency constituted 3-32% of total cost reductions over the period 2010-2015, depending on the model specification. I interpret my preferred specification as a causal impact under the identifying assumption that there are no unobserved determinants of cost that are correlated with technical efficiency in the cross section, which is plausible given its robustness to potential confounding variables. I estimate fixed costs of improving technical efficiency that are in line with reported expenditures on research and development (R&D) and physical capital expenditures by a subsample of firms that are publicly owned.

I use the estimated model to quantify the impact of historical consumer subsidies. Results from counterfactuals that hold firms' costs fixed suggest that subsidies increased solar panel adoption (measured in Watts) by 49% over the period 2010-2015. These effects only account for demand responses and not for dynamic supply responses, implicitly assuming that innovation was exogenous rather than driven by profit motives. German subsidies explain over one-third of this increase in adoption, followed by national subsidies in Japan and in the United States.

I conduct a back-of-the-envelope analysis to quantify the external social benefits attributable to solar subsidies. Solar electricity generation creates a positive non-pecuniary externality if it displaces electricity generation from coal and natural gas, which emit harmful air pollution. I combine existing estimates of the external damages from coal and natural gas electricity – including both local air pollution and greenhouse gas emissions – with my estimates of the impact of subsidies on solar panel adoption to compute the total external social benefits of solar subsidies. My central estimate is a present discounted value of $15.4 billion (in 2017 dollars).

I then account for endogenous firm innovation by allowing firms to reoptimize in a counterfactual simulation without Germany's national subsidies. Removing these subsidies lowers firms' profits, leading firms to invest less frequently than under the

[5]See, for example: Hughes and Podolefsky (2015); Burr (2016); De Groote and Verboven (2016); Gillingham and Tsvetanov (2017).

[6]The subsidies to fixed and variable costs Chinese manufacturers allegedly benefit from are an example of unobserved supply shifters. The relationship between technical efficiency and cost is robust to including unobserved costs that vary by manufacturing location (i.e., inside China or outside China) intended to capture these unobserved subsidies to production.

baseline simulation with Germany's subsidies in place. Production costs are therefore higher than under the baseline simulation, and the difference between simulated costs with and without German subsidies grows over time. The change in investment activity is significant: 32% of the solar adoption due to increased technical efficiency would not have occurred in the absence of German subsidies. The vast majority of this marginal adoption occurs outside of Germany, highlighting the spillovers generated by decentralized government intervention in a global market.

Accounting for both demand and supply responses yields meaningfully different estimates of the external social benefits of subsidies than accounting only for demand responses. I repeat the back-of-the-envelope calculation described above, except that I now include external benefits from future solar adoption through 2040 attributable to subsidy-induced innovation during the sample period (2010-2015). This accounts for the fact that cost reductions due to past innovation affect future costs, and therefore future solar panel adoption.[7] The external benefits increase 22% to $18.8 billion under conservative assumptions. Furthermore, since this analysis focuses on just one margin of endogenous innovation, it is likely that this estimate considerably understates the full impacts of subsidies to solar adoption.[8]

This paper highlights two important points that apply to policy beyond this market. First, using short-run economic methods to analyze government policy can understate its true effects and lead to incorrect policy prescriptions. Policy and regulatory analyses often omit long-run considerations or use ad hoc methods to analyze them.[9] My results show that the resulting bias can be significant: dynamic supply responses such as induced innovation have the potential to become more important than short-run demand responses over long time horizons, as technological progress is cumulative.

[7] This is in contrast to the demand effect of subsidies, which will cease to affect equilibrium quantities once subsidies are removed. A level comparison of the impacts of subsidies through (1) only their demand effects, versus (2) both demand and supply effects, requires accounting for the different time horizon of the two mechanisms' impacts.

[8] This paper is only a first step toward a comprehensive welfare analysis of solar subsidies. The optimal coordinated policy depends on factors beyond the scope of this paper, such as the extent of spillovers in the R&D process. The growing literature on directed technical change makes clear that a Pigouvian tax (or subsidies to clean technologies that approximate a Pigouvian tax) may not be the optimal government response to climate change (Acemoglu et al., 2016; Aghion et al., 2016; Lemoine, 2017).

[9] Harrington et al. (2000) compare *ex-ante* and *ex-post* estimates of the direct costs of regulations and find that "unanticipated technological innovation appears to be an important factor" in cases where unit costs were overstated *ex-ante*. U.S. Office of Management and Budget (2011) directs agencies to identify "changing future compliance costs that might result from technological innovation" in their regulatory analyses. Yet practical analyses still fall short of this ideal. For example, in the regulatory impact analysis for the Mercury and Air Toxics Standards, U.S. Environmental Protection Agency (2011) acknowledges but does not quantify the effect of technological innovation on costs.

Furthermore, the benefits of induced innovation will remain even if policies are removed in the future. Second, decentralized government intervention in a global market is not economically efficient in general. Innovation spillovers across borders may lead governments to underinvest in new technologies like solar panels. This effect is distinct from, and compounds, the problem of free-ridership created by international spillovers in environmental benefits, underscoring the need for international coordination to address climate change.

Related Literature This paper builds on a large literature on induced innovation in energy markets (e.g., Newell et al., 1999; Popp, 2001, 2002; Jaffe et al., 2002). This paper also draws inspiration from research on how innovation responds to changes in public health policy (Finkelstein, 2004), expansion of prescription drug insurance (Blume-Kohout and Sood, 2013), and exogenous shifts in market size (e.g., Acemoglu and Linn, 2004; Dubois et al., 2015).[10]

The novelty of my approach relative to previous research on innovation is the use of a unique, observable, and verifiable measure of innovation: technical efficiency. Much of the innovation literature uses data on patents or product introductions, the value of which are highly variable and difficult to quantify *ex-ante*. In contrast, technical efficiency improvements are measured in common units, allowing me to directly model the impact of individual innovations on costs.

The fact that technical efficiency is observable is also a strength of this approach relative to the use of estimated measures such as production costs or productivity.[11] Over the past several years, manufacturers in China have benefited from manufacturing subsidies, and thus changes in estimated production costs over time may conflate cost-reducing innovation with changes in manufacturing subsidies. This could lead to overestimates of the extent of real resource cost reductions in the industry and therefore the impact of consumer subsidies on innovation. In contrast, my approach exploits observed consumer subsidies and an observed measure of investment outcomes. To prevent unobserved production subsidies that vary over time from biasing my estimates, I use cross-sectional variation to identify the causal impact of technical efficiency on production costs.

This paper also contributes to the growing literature on the economics of solar

[10] See also Schmookler (1966) and Pakes and Schankerman (1984).

[11] Another strand of literature estimates production functions using observed R&D spending. However, because firms use R&D to achieve multiple objectives, the benefits of R&D may not be fully captured by this approach (Popp, 2001). A related advantage of my approach is that it relies only on data from the product market and does not require data on inputs to production.

power (Baker et al., 2013). Several papers have found that past subsidies significantly expanded demand for solar systems (Hughes and Podolefsky, 2015; Burr, 2016; De Groote and Verboven, 2016; Gillingham and Tsvetanov, 2017). The consensus of this research is that many solar incentives are above the level justified by the static environmental benefits of adoption. Dynamic considerations such as innovation and learning-by-doing may justify these subsidies in theory (Arrow, 1962; Goulder and Mathai, 2000; van Benthem et al., 2008), but there is limited empirical evidence to assess and guide policy.[12,13]

While there is a large body of research on the market for solar systems, there is little economic research on the upstream industry of solar panel manufacturing. Some studies attempt to understand historical price reductions and forecast future prices at the industry level using learning curves, but this approach cannot separately identify endogenous and exogenous technological change (Nordhaus, 2014). Furthermore, this industry-level analysis conflates many underlying economic phenomena. Recognizing this, Nemet (2006) decomposes historical price reductions based on observable factors, finding that manufacturing scale, solar panel efficiency, and input prices are important explanatory variables. He concludes that learning-by-doing is not a major factor. Finally, Pillai and McLaughlin (2013) and Pillai (2015) use static economic models to study solar panel manufacturing. In contrast to prior studies of this industry, I estimate a dynamic game to answer a broad economic question.

Road Map The remainder of this paper is organized as follows: Section 2 describes the recent growth of the market for solar panels, the prevalence and types of consumer subsidies, how solar panels are manufactured, and the importance of technical efficiency to the industry. Section 3 details a structural model of the industry environment and manufacturer behavior. Section 4 introduces the data I use for estimation. Section 5 outlines my estimation approach, and Section 6 summarizes the estimates. Section 7 describes how I use these estimates to simulate counterfactual market outcomes and presents the results. Section 8 concludes.

[12] In a notable exception, Bollinger and Gillingham (2014) estimate the extent of learning-by-doing spillovers in solar panel installation, the industry downstream of the firms I study.

[13] This paper focuses on technological innovation through R&D rather than learning-by-doing. The model could be extended to incorporate learning-by-doing using methods developed by Benkard (2004).

2 Industry Background

2.1 Economic and Policy Environment

The solar industry has grown rapidly over the past several years due to a combination of demand and supply factors. Governments around the world have encouraged adoption of this technology through policies targeting solar power. For example, the United States Government expended one-third of all electricity subsidies ($5.3 billion) on solar power in fiscal year 2013 (U.S. Energy Information Administration, 2015). This excludes the value of state and local subsidies, which may have been as large or larger than Federal subsidies in some jurisdictions and time periods (Borenstein, 2017). Partly as a result of these subsidies, solar power was the largest source of electricity capacity additions in the United States in 2016. The United States is not unique in this regard. China, Germany, and Japan have provided subsidies for electricity generated by solar panels over the past decade that helped make these three countries the first, second, and third largest markets for solar panels in terms of cumulative quantities sold in 2015 (International Energy Agency, 2016).

The specific mix of policies employed by governments has varied across jurisdictions and over time. The most common policy mechanisms during the period 2010-2015 can be classified in two broad categories. The first is subsidies to solar adoption. An example of this is the Federal Investment Tax Credit in the United States, which provides a tax credit of 30% for solar investment costs. The second is subsidies to electricity generation from solar technology. Germany, Japan, China, and many other countries have offered payments for solar electricity in the form of "feed-in tariffs" set at the time of investment. Feed-in tariffs are prices paid for solar electricity fed onto the electric grid that are independent of the cost of electricity from alternative sources. Solar Renewable Energy Certificates from state-level Renewable Portfolio Standards in the United States are another example of incentives that target solar electricity generation rather than investment.

Taken together, government policies significantly shifted out demand for solar panels. Figure 1 provides graphical evidence of the impacts of individual subsidies on demand. In Japan, demand for solar panels was fairly low throughout 2010 and 2011, but increased after a feed-in tariff was introduced in the wake of the Fukushima Daiichi nuclear disaster. Furthermore, the equilibrium quantity of solar panels sold fell as the feed-in tariff was lowered in 2015, despite the fact that prices were slowly but steadily declining over that year (Figure 1a). The German feed-in tariff predates the sample period, but variation in the feed-in tariff level over time appears to affect demand: equilibrium quantities fell after the second quarter of 2012 as feed-in tariffs were lowered, despite the fact that prices were falling through early 2013 and remained fairly stable thereafter (Figure 1b). Previous empirical analyses of subnational consumer subsidies found that a majority of solar adoption was attributable to subsidies (e.g., Hughes and Podolefsky, 2015; Burr, 2016; De Groote and Verboven, 2016; Gillingham and Tsvetanov, 2017).

Supply side developments also contributed to the growth of this market. The global average price of solar panels fell 75% over the period 2010-2015. Previous analyses attribute historic solar panel price reductions to improvements in technical efficiency by solar panel manufacturers and reductions in the price of the primary input (Nemet, 2006; Pillai, 2015). The expansion of manufacturing activity in China, aided by government subsidies to manufacturers, has also played an important role in the industry's evolution. Chinese manufacturers produced roughly two-thirds of all solar panels between 2010 and 2015.[14]

While solar panels are the main input to solar systems that generate electricity, the cost of complementary inputs also plays a role in determining the size of the global solar panel market. Recognizing this, the U.S. Department of Energy established its SunShot Initiative in 2011 to encourage cost reductions for both hardware and non-hardware inputs to solar systems. The costs of complementary hardware (e.g., mounting hardware and inverters) and installation have fallen over the past several years, and downstream business models have also contributed to the growth of the solar market. However, solar panels remain central to efforts to reduce the cost of solar electricity. The SunShot Initiative recently announced that the solar industry had met its 2020 cost targets for utility-scale solar systems "largely due to rapid cost declines in solar photovoltaic hardware" (U.S. Department of Energy, 2017). The SunShot Initiative's statement cited a detailed bottom-up cost analysis by Fu et al. (2017) that identified solar panel prices as the primary driver of recent solar system cost reductions. This research does not specifically attribute these solar panel price reductions to public or private efforts to reduce the cost of solar panels, either through government subsidies on the supply side (e.g., for production and R&D) or the demand side, which is the focus of this study.

These demand and supply developments are inextricably linked. By increasing demand, government policies may have increased the returns to innovation, encouraging firms to invest in response. If these policy-induced demand shocks did not induce innovation, governments may have been better off had they delayed the use of

[14] Author's calculation based on IHS Markit's PV Suppliers Tracker (2016Q2).

Figure 1: Graphical Evidence that Subsidies Increase Demand

(a) Japan

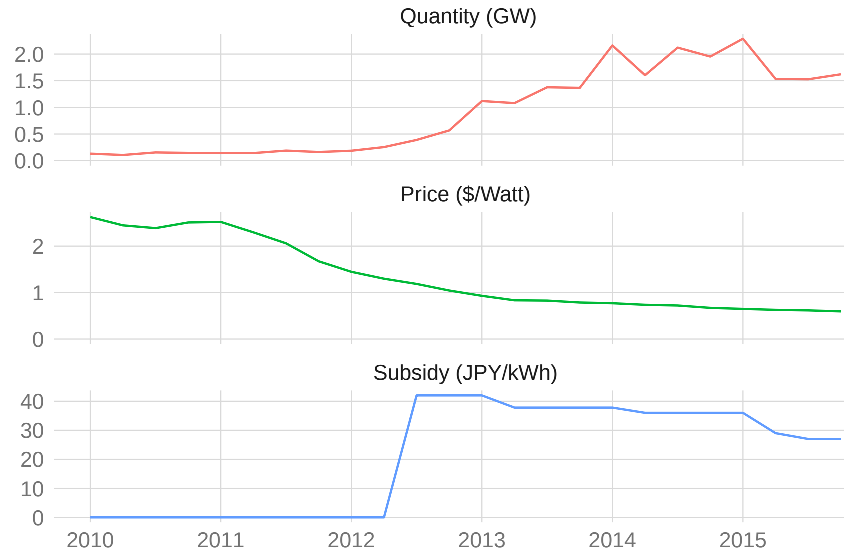

(b) Germany

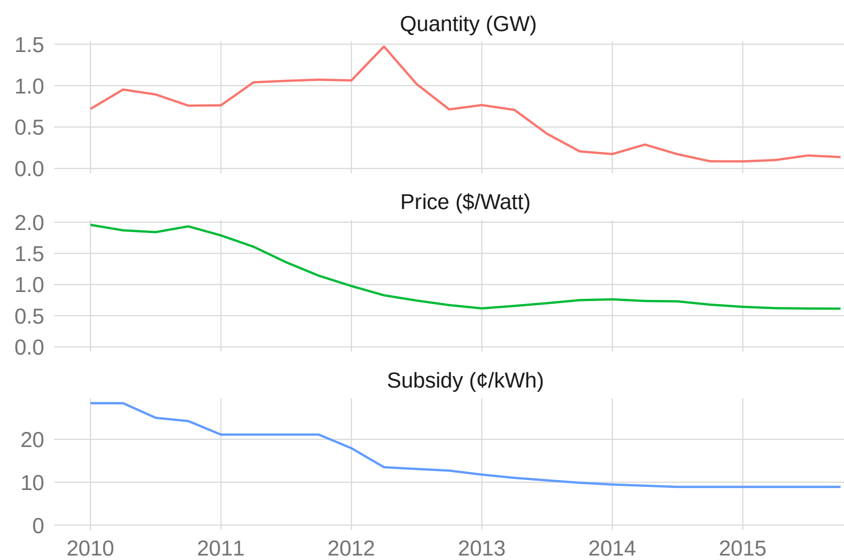

Notes: These two figures plot quarterly observations of each market's quantity of solar panel sales (GW), average price ($/W), and subsidy level (feed-in tariff in local currency units). Both figures are suggestive of the impact of subsidies on demand. Figure 1a shows that equilibrium quantities were low throughout 2010 and 2011, but increased after a feed-in tariff was introduced in the wake of the Fukushima Daiichi nuclear disaster. Furthermore, the equilibrium quantity of solar panels sold fell as the feed-in tariff was lowered in 2015, despite prices slowly but steadily declining over that year. Figure 1b shows that equilibrium quantities fell in Germany after 2012Q2 as the feed-in tariff was lowered, even though prices were falling through early 2013 and remained fairly stable thereafter.

Data Source: IHS Markit and the International Energy Agency.

incentives until these innovations had occurred.

2.2 Solar Panel Manufacturing

Solar panels convert sunlight into electricity via the photovoltaic effect.[15] I study manufacturers of p-type silicon-based photovoltaic panels who collectively produced approximately 88% of the world's solar panels over 2010-2015 in terms of electricity generating capacity (Watts).[16] I refer to their products as "conventional solar panels" to distinguish them from solar panels made from alternative semiconducting materials.

Conventional solar panels are produced from highly purified silicon. Manufacturers process silicon to create solar cells that generate electricity when exposed to light.[17] These cells are arrayed and assembled into a panel that consists of a frame, backing, circuitry, lamination to protect cells from the elements, and a glass front.

The fundamental technology produced by manufacturers has not changed significantly over the past decade, but firms have worked to increase their technical efficiencies and decrease their production costs. The median technical efficiency of solar panels installed in the United States rose from 14.1% in 2010 to 17.0% in 2015 (Figure 2; Barbose and Darghouth, 2017). Firms throughout the industry have increased the technical efficiency of their solar panels, although there is considerable variation in the firm-level frontier of (i.e., maximum) technical efficiency in the cross section and in the relative position of firms over time (Figure 3).

Firms strive to increase their technical efficiencies in order to reduce their production costs. Technical efficiency determines electricity output from a solar panel holding its physical size fixed, so efficiency improvements can lower materials costs for a given amount of electricity output. This is important because materials comprised two-thirds of solar panel costs during the period under study (Powell et al., 2012).[18] Firms cite

[15] The terms *solar panel* and *solar module* are used interchangeably. I use *solar panel* throughout this paper because it is more familiar.

[16] Author's calculation based on IHS Markit's PV Integrated Market Tracker (2016Q2). The 88% figure includes all manufacturers of p-type silicon-based photovoltaic panels. I use a subset of these manufacturers to estimate the model due to data limitations. See Section 4 for more details.

[17] Silicon is first doped with boron and formed into blocks of monocrystalline and multicrystalline material called ingots. These ingots are sliced into very thin wafers roughly $6'' \times 6''$. The wafers are then doped with phosphorous to create a layer of n-type silicon that forms a junction. The addition of contacts for conducting electricity and chemical processing renders photovoltaic cells.

[18] Powell et al. (2012) state that "improved solar cell conversion efficiency is a major driver for c-Si module [panel] cost reduction, as cost scales inversely with efficiency for all area-dependent cost components." Green (2016) claims that "efficiency... is probably the key both to future photovoltaic electricity cost reduction and to commercialization of new technologies" and that observed efficiency improvements by existing manufacturers "contribute increasingly significantly to ongoing cost reduction."

Figure 2: Industry Technical Efficiency is Increasing over Time

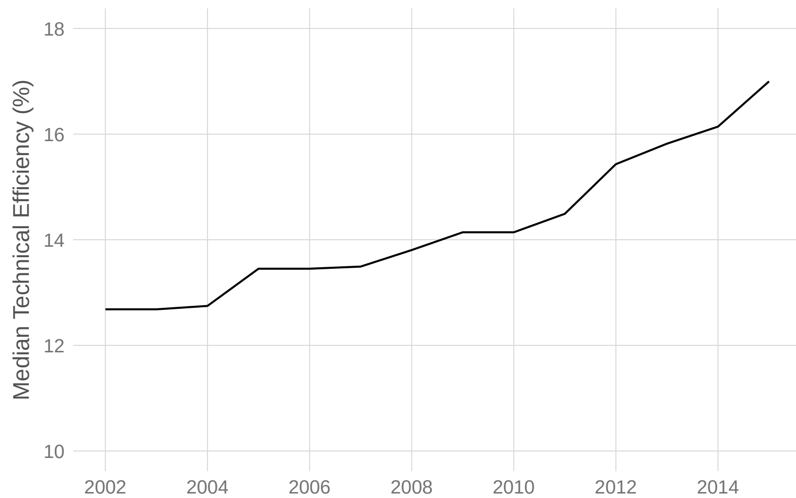

Notes: This plot shows the industry-wide progression of technical efficiency based on data from the United States solar market. During the sample period of 2010-2015, the median technical efficiency of solar panels installed in the U.S. rose from 14.1% in to 17.0%. These data include thin-film and high-efficiency n-type silicon solar panels, whereas I focus on conventional p-type silicon solar panels. *Source:* Barbose and Darghouth (2017).

Figure 3: Firm-Specific Technical Efficiencies over Sample Period

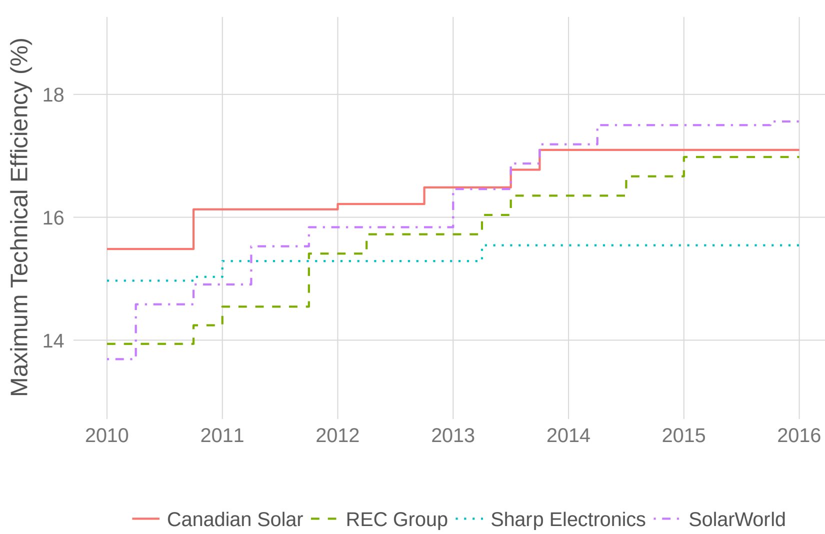

Notes: This plot shows the firm-level progression of technical efficiency for a sample of firms based on data from the United States solar market. This figure shows that while firms throughout the industry have increased the technical efficiency of their solar panels, there is variation in firm-level technical efficiencies in the cross section and in the relative position of firms over time.
Data Source: Lawrence Berkeley National Laboratory's Tracking the Sun dataset (openpv.nrel.gov).

technical efficiency as a source of cost reductions in press releases and SEC filings. These statements also corroborate the role of R&D in enabling advances in technical efficiency. For example, Trina Solar's 2014 Form 20-F states: "To reduce raw material costs, we continue to focus our research and development on improving solar cell conversion efficiency and enhancing manufacturing yields."[19] Industry analysts report that large manufacturers have "top-notch" in-house R&D labs, and that this is true of Chinese firms as well as Western firms. These labs are specific to individual firms, and they do not directly share intellectual property.[20]

Solar panel manufacturers improve their technical efficiency through investments in R&D and physical capital. These firms primarily use established, off-patent technologies and innovate to improve their current implementations or commercialize alternative implementations of these technologies. Research and development advances are operationalized through investment in physical capital, either through production line upgrades or the installation of new production lines.[21]

Despite continual improvements to the production process and to technical efficiencies, the firms I study produce a highly commoditized product. Solar panels come in standardized form factors, with most solar panels composed of either 60 or 72 cells. The smaller size is most commonly used for residential rooftop applications, while the larger size is typically used in commercial and utility applications. The electrical properties of solar panels are also standardized. Figure 4 presents a side-by-side comparison of two solar panels.

3 Model

I model the solar panel manufacturing industry as an imperfectly competitive oligopoly. There are I incumbent firms who compete in quantities (Cournot) in each regional

[19] Trina Solar also cited the role of technical efficiency improvements in its 2013 Form 20-F. Yingli Solar has cited efficiency as a key means by which to achieve cost reductions in promotional materials.

[20] Phone interview with Jade Jones, Senior Analyst at GTM Research (May 26, 2017). While firms conduct their own in-house research, they may benefit indirectly from knowledge generated by their competitors. For example, the observation that a competitor successfully commercialized an existing technology may spur others to do the same. This is consistent with widespread adoption of specific technologies, such as the ongoing shift toward PERC technology. In addition, past and current government-sponsored research in many countries may benefit both domestic and international manufacturers.

[21] This process could be described by a theoretical model in which capital investment is necessary to capitalize on innovations resulting from R&D (e.g., Lach and Rob, 1996). Lach and Schankerman (1989) provide empirical evidence consistent with this type of model using data on U.S. manufacturing firms. I abstract from these details and model the joint process of investing in R&D and physical capital due to data constraints.

Figure 4: Comparison of Two Solar Panels

Output (Watts)	275	330
Size (cells)	60	72
Technical Efficiency (%)	16.8	17.0
Relative Price ($/W)	**1.00**	**1.01**

Notes: Individual solar panels are rated at different output levels, come in a few different standardized physical sizes (measured here by the number of solar cells), and can be different colors. Despite the differences between these two example solar panels in physical size, power output, and appearance, their prices are very similar when measured using the industry convention of $/Watt.
Data Source: Prices and product details retrieved for two models made by Canadian Solar from https://ressupply.com on September 24, 2017.

market in each discrete time period (quarter). There are M regional markets. Aggregate demand for undifferentiated solar panels in each market is static but depends on subsidies, which vary over time. Firms have an infinite horizon and share a common discount factor β. Each firm is differentiated by its state, technical efficiency: $s_{it} = \omega_{it}$.[22] Firms share the industry state, s_t, which is comprised of:

- the distribution of individual firms' technical efficiencies, $\omega_t = [\omega_{1t} \quad \omega_{2t} \quad ... \quad \omega_{It}]$;
- a common input price, w_t; and
- aggregate demand in each market, $d_t = [d_{1t} \quad d_{2t} \quad ... \quad d_{Mt}]$.

In each period, firms first observe the industry state and realize private shocks to investment. They then compete in the product market and choose whether to invest to lower their future costs. When making this decision, each firm takes expectations

[22]The notation s_{it} is common in the literature to denote an individual agent's state. In this context, the firm's state (s_{it}) and its technical efficiency (ω_{it}) are equivalent and used interchangeably because firms are only differentiated on one dimension (technical efficiency).

over the outcome of its investment decision, future demand, the evolution of the input price, and investment by its competitors. Finally, investments are implemented and their outcomes are realized at the beginning of the next period.

3.1 Demand

Consumers demand electricity, a prototypical homogeneous good. Solar panels are one potential source of electricity. I assume that consumers do not have preferences over solar panels *per se*, but instead have preferences over the electricity generating capacity of undifferentiated solar panels (in Watts).[23] I refer to this final good as "solar panels" throughout the paper.

Aggregate demand for solar panels in each market is static and depends on market-specific factors,

$$Q_{mt} = Q_m(P_{mt}; d_{mt}(S_{mt})), \tag{1}$$

where Q_{mt} is the quantity of solar panels (in Watts) and P_{mt} is the price (in \$/Watt). The demand curve $Q_m(\cdot)$ is indexed by m to allow the shape of the demand curve to vary across markets. The demand state in each market, $d_{mt}(S_{mt})$, depends upon subsidies to consumers, S_{mt}, and follows an exogenous first-order Markov process.

The static demand specification implies that consumers are not forward-looking. In reality, potential purchasers of a durable good who expect prices to fall over time, as has been the case for solar panels, may delay their purchase. Static demand estimation may therefore understate the magnitude of the true price elasticity (Aguirregabiria and Nevo, 2013). The static demand specification also implies that consumers do not exit the market after purchasing solar panels. This rules out changes in the distribution of consumers over time, such as if early adopters are less price sensitive than late adopters. In this hypothetical case, static estimation may again understate the price elasticity of demand (Conlon, 2012; Gowrisankaran and Rysman, 2012).

These theoretical insights are consistent with the limited evidence on demand estimation from the solar market. De Groote and Verboven (2016) estimate demand for solar systems (not solar panels) using both static and dynamic specifications. The

[23] The assumption that this final good is undifferentiated implies that consumers do not have preferences over specific brands, in keeping with the commoditized nature of the product outlined in Section 2. The assumption that consumers do not have preferences over the number of solar panels they purchase is in keeping with industry convention; firm sales are denominated in Watts rather than the number of solar panels (with prices denominated in \$/Watt). Figure 4 compares two example solar panels to provide some context for these assumptions.

estimated price coefficient from their baseline model, which abstracts from consumer heterogeneity, is roughly 40% higher under dynamic demand than static demand. Allowing for consumer heterogeneity further magnifies the price coefficient, to roughly 20% higher than the dynamic model without heterogeneity. More generally, however, it is possible that other regional solar markets could display forms of demand or unobserved heterogeneity that would bias static price coefficients away from zero.

Despite its potential shortcomings, the static model of demand facilitates estimation of a dynamic model of supply, which is the focus and contribution of this paper relative to previous studies of the solar market.[24] While there were significant reductions in solar panel prices during the sample period, there are a few features of this market that ameliorate concern over the potential impact of using a static demand specification. First, changes in prices and subsidies were countervailing in some markets (see Figure 1b for an example).[25] Second, anecdotes from government and industry publications suggest that ongoing price reductions were not fully anticipated, even by industry insiders.[26] Finally, Gillingham and Tsvetanov (2017) provide support for their assumption of static demand using data on purchases and a consumer survey in the residential solar market in Connecticut, in the United States.

3.2 Firm Cost Structure

Firms have constant marginal costs of production, $mc\left(\omega_{it}, w_t\right)$, that depend on firm-specific technical efficiencies (ω_{it}) and a common input price (w_t). Both technical efficiency and the input price are fixed from the perspective of the firm at the time of product market competition and are not choice variables in the firm's static optimization problem. I assume that product market competition is Cournot, with firms choosing quantities to maximize profits,

$$\max_{q_{imt}} \; \left[P_{mt}(Q_{mt}; d_{mt}(S_{mt})) - mc\left(\omega_{it}, w_t\right)\right] q_{imt},$$

[24]It is computationally challenging to estimate models of dynamic demand and dynamic supply, and as a result it is common in the industrial organization literature to use a parsimonious demand model to facilitate estimation of a dynamic supply model.

[25]De Groote and Verboven (2016) cite this fact as a reason that price coefficients were not even more different between their static and dynamic specifications.

[26]As one example, the U.S. Energy Information Administration (2016) acknowledged that "EIA, like many other industry trend watchers, did not anticipate the sharp decline in solar PV costs seen over the past several years." Creutzig et al. (2017) chronicle "a history of widespread underestimation of the growth in PV deployment" and attribute these underestimates to faster than expected solar panel price declines, among other things.

where $P_{mt}(Q_{mt}; d_{mt}(S_{mt}))$ is the inverse demand curve corresponding to equation 1. Equilibrium product market profits for each firm depend only on the firm's state and the industry state (which includes the input price and demand states). I denote these equilibrium product market profits $\bar{\pi}_i(s_t)$.[27]

Firms choose technical efficiency dynamically. If costs are decreasing in technical efficiency, firms will have an incentive to invest in R&D and physical capital to increase future technical efficiency and improve their competitive position. Firms make a discrete decision whether to invest (x_{it}). The firm's per-period payoff,

$$\pi_i(x_{it}, s_t; \varepsilon_{it}) = \bar{\pi}_i(s_t) - \gamma x_{it} + \sigma \varepsilon_{it}(x_{it}),$$

is comprised of three terms. Each firm earns profits from the product market, $\bar{\pi}_i(s_t)$, which do not depend on the firm's investment choice. The second term consists of a nonrandom fixed cost, γ, which is paid only if the firm invests (in which case $x_{it} = 1$). Finally, firms receive private choice-specific shocks, $\varepsilon_{it}(x_{it})$, which are independent and identically distributed (i.i.d.) according to the Type I extreme value distribution.[28] The structural interpretation of these shocks is a random shock to the fixed cost of investment.[29] The shocks are scaled by the parameter σ.

3.3 State Transitions

I assume the input price (w_t) and demand states (d_t) are exogenous and evolve according to independent first-order Markov processes. Each firm's state evolves stochastically over one period. I assume the relationship between investment and the evolution of technical efficiency is one-to-one in order to infer the unobserved investment decision. If a firm does not invest ($x_{it} = 0$), its technical efficiency does not change. If a firm does invest ($x_{it} = 1$), the change in its technical efficiency is ν_{it}, which is i.i.d. across firms and time with support $\nu_{it} \in (0, \infty)$. To summarize, technical efficiency evolves according to

$$\omega_{it+1} = \omega_{it} + x_{it}\nu_{it}.$$

[27] Firms share a common profit function which is indexed by i to illustrate its dependence on firm i's state. Profits could equivalently be expressed as $\bar{\pi}(s_{it}, s_t)$ or $\bar{\pi}(\omega_{it}, s_t)$, since $s_{it} \equiv \omega_{it}$.

[28] This distributional assumption is common in the discrete choice literature due to the analytic form it implies for conditional choice probabilities.

[29] Modeling the innovation process as an outcome of fixed cost investments is guided by the industry background discussed in Section 2.2.

Although the outcomes of R&D activities are inherently stochastic, firms must upgrade existing capital or install new capital to implement the advances realized through R&D. This is the economic basis for the assumption that the investment I model always yields a non-zero improvement in technical efficiency. The stochastic nature of the investment outcome captures the uncertainty inherent in adapting R&D advances from laboratory pilot lines to large-scale production.

The distribution of technical efficiency across firms (ω_t) evolves over one period as a result of firm actions and the resulting realizations of technical efficiency improvements.

3.4 Equilibrium

I assume firms use symmetric pure strategies that depend only on the current state and their private information, leading to a Markov-Perfect Nash Equilibrium (Maskin and Tirole, 1988). Each firm's strategy, denoted $\zeta_i(s_t, \epsilon_{it})$, is a mapping from states and private shocks to actions (i.e., quantities sold in each market and a binary investment decision). The firm's value function at the time of its investment decision is,

$$V_i\left(s_t; \zeta_i, \zeta_{-i}, \varepsilon_{it}\right) = \max_{x_{it} \in \{0,1\}} \bar{\pi}_i(s_t) - \gamma x_{it} + \sigma \varepsilon_{it}(x_{it}) \quad + \beta E\left[V_i\left(s_{t+1}; \zeta_i, \zeta_{-i}, \varepsilon_{it+1}\right) \middle| s_t, x_{it}\right],$$

where the expectation is taken with respect to i's investment outcome (ν_{it}), investment by i's competitors, future realizations of the exogenous demand and input price states, and i's own future cost shocks. ζ_{-i} denotes the strategies of firms other than firm i. Markov-Perfect Nash Equilibrium requires that each firm's strategy is optimal given the common strategy used by its competitors,

$$V_i\left(s_t; \zeta_i, \zeta_{-i}, \varepsilon_{it}\right) \geq V_i\left(s_t; \zeta_i', \zeta_{-i}, \varepsilon_{it}\right),$$

for all firms (i), states (s), shocks (ε), and alternative strategies (ζ').

4 Data

I employ data on the global solar panel market from IHS Markit and Lawrence Berkeley National Laboratory. IHS Markit's PV Module Intelligence Service provides data on the solar panel supply chain on a quarterly basis from January 2010 through March 2016. These data include sales by regional market and production activities by country of production. I use the sales data to construct a dataset at the firm-market-quarter level that includes sales for all firms and prices for a subset of firms. Total firm sales are denominated in Watts (W) of electricity generating capacity and prices are in dollars per Watt ($/W). The data are at the firm level and do not include sales of individual solar panel models. I focus on four regional markets: Germany, Japan, the United States, and the Rest of the World. I complement these with data on firms that vary over time but not across markets, including each firm's production capacity and quarterly production throughout the supply chain (i.e., wafers, cells, and modules). Quarterly data on the global spot price of polysilicon, the primary input to solar panel production, also comes from IHS Markit.

The measure of each firm's technology comes from Lawrence Berkeley National Laboratory's Tracking the Sun dataset.[30] The dataset contains the characteristics of installed solar systems throughout the United States. These characteristics include the manufacturer, model, and electrical conversion efficiency of the solar panels utilized in each system. After removing missing data, the dataset contains over 425,000 systems installed between January 2010 and December 2015. I use these data to construct summary statistics of the state of each firm's technology on a quarterly basis. In my empirical analysis, I focus on the frontier of electrical conversion efficiency, tracking the maximum electrical conversion efficiency sold by each firm over time. This is the observed measure of technical efficiency I refer to throughout the paper.

Data on market-specific subsidies to consumers that vary over time come from the International Energy Agency's Photovoltaic Power Systems Programme and national governments. For Germany and Japan, I construct a panel dataset of the feed-in tariff level in each market over time (represented graphically in Figure 1). For the United States, I use the Investment Tax Credit (ITC) and the tax advantage of accelerated depreciation, both of which depend on the cost of the solar system rather than the amount of electricity it generates. I use a subsidy of 40% of the solar panel purchase price to summarize federal solar subsidies in the United States.[31] Finally, for the residual market, I use the feed-in tariff from Italy as a proxy for subsidies in all other markets. Italy was chosen because the Italian government offered generous feed-in

[30]The August 2016 version of this dataset was downloaded from https://openpv.nrel.gov.

[31]The ITC is a tax credit for 30% of investment costs. Borenstein (2017) estimates that accelerated depreciation is equivalent to a 12.6% to 15.2% reduction is the cost of a solar system after state incentives and the ITC. I approximate this by assuming accelerated depreciation is worth 10% of total costs. While accelerated depreciation benefits are not available to households that purchase their own solar panels, they are available to businesses that purchase solar panels either for their own use or for leasing to households. I also experimented with using state-level subsidies, including incentives from the California Solar Initiative, but my use of aggregate national data on solar panel sales makes it difficult to separately identify the impacts of these state-level subsidies. My analysis also omits the value of net metering.

tariffs during the first half of the sample period, and because Italy was a large solar market: in each year from 2010 to 2012, Italy was either the first- or second-largest national solar market in the world (International Energy Agency, 2013).[32] The feed-in tariffs are all units of local currency per unit of electricity output.[33]

I also collect ancillary data to use as instrumental variables from multiple sources. The global price of silver comes from the London Fix via The Silver Institute. Data on international trade restrictions come from the United States Government and the European Commission. I describe the use of these ancillary data in more detail in Section 5.

Table 1 contains summary statistics for the four markets. Germany, Japan, and the United States comprise 48% of cumulative sales over the period 2010-2015 (in terms of electricity generating capacity). There are 15 firms in the sample that comprise approximately 70% of the global conventional solar panel market. The collective market share of these firms was stable over the sample period despite entry and exit in the competitive fringe, which is omitted from this analysis due to data constraints. While some of the 15 firms are not active in every market at the beginning of the sample period, all 15 firms are active in all markets during the sample period. Instances of firms not being active in a market are rare, as summarized by the average number of firms in each market over all time periods.[34] The final three columns summarize the distribution of market shares across firms within a given time period (pooling across time periods). The median market share in each region ranges from 2.7-5.5%, but market shares for some firms in some periods are much larger: the 90$^{\text{th}}$ percentile ranges from 14.4-19.5%. However, this is not driven by a small number of large firms that dominate the industry: the cumulative global market share of each firm ranges from 2-13%. These market shares motivate the use of a model of imperfect competition to characterize the product market.

[32] I also experimented with creating an index of subsidies from major solar markets, but I found that the Italian subsidies alone outperformed the indices I tried in terms of model fit.

[33] Feed-in tariffs are paid per unit of electricity output, not per unit solar capacity installed, as described in Section 2. The units of the feed-in tariffs have no impact on the model, as they are all scaled by estimable parameters that have no economic interpretation in isolation. The product of each feed-in tariff and its coefficient gives the impact of that subsidy on demand, and this product is invariant to changes in the feed-in tariff's units.

[34] The model outlined in Section 3 does not endogenize market entry but instead implicitly assumes firms compete in every market in every time period. While this assumption is not required for demand or production cost estimation, it affects equilibrium profits under the model and therefore may affect the dynamic parameter estimates and counterfactual simulations. These effects should be negligible, both because instances of firms not being active in all four markets are infrequent and because they occur at the beginning of the sample when costs were high, and therefore equilibrium quantities and profits were low relative to later periods.

Table 1: Summary Statistics by Market

	Sales	Active Firms		Market Shares		
Market	(GW)	Mean	Max	10%	50%	90%
Germany	14.8	14.9	15	0.5	4.6	15.2
Japan	22.3	14.7	15	0.2	2.7	19.5
ROW	59.8	15.0	15	1.2	5.5	14.4
USA	18.0	13.7	15	0.3	4.9	18.7

Notes: This table presents summary statistics for the four regional markets studied. Cumulative sales are measured in gigawatts (10^9 Watts). Germany, Japan, and the United States comprised roughly half of global sales during the period 2010-2015. All 15 firms are active in every market during the sample period, although there are some periods in which some firms are not active in every market. Market shares are constructed in each market and time period, and then pooled across time periods within each market to summarize the distribution of market shares. There are 24 time periods for each market.
Data Source: Author's calculations using data described in Section 4.

5 Model Estimation

5.1 First Step: Estimate Product Market Model, State Transitions, and Investment Policy Function

Demand I assume the demand curve in equation 1 is a constant elasticity demand curve and estimate the model in logs,

$$\ln Q_{mt} = \alpha_{0m} + \alpha_{1m} \ln P_{mt} + \alpha_{2m} S_{mt} + \varepsilon^{D}_{mt}. \tag{2}$$

I use data on feed-in tariff levels described in Section 4 for the subsidy variable, S_{mt}, for the markets of Germany, Japan, and Rest of the World. For the United States, I model the subsidy as a function of the solar panel price, because the primary federal solar incentives depend on the cost of a solar system rather than its electricity output. See Section 4 for further details.

I instrument for price to account for potential endogeneity using two sets of instruments. In the first, I use the price of two inputs – polysilicon and silver – that comprise the first and second largest materials cost shares over the sample period.[35] These two instruments vary over time but not across markets because polysilicon and

[35]Silicon forms the core of conventional solar cells, and silver is used to construct contacts on the front and back of conventional solar cells to create an electrical circuit. See Section 2 for a more detailed description of the manufacturing process.

silver are traded globally. I augment them with indicator variables for the presence of trade restrictions that create differential variation in prices across markets and are assumed to be uncorrelated with unobserved demand shocks.

I construct a second set of instruments for the price of solar panels in market m using the average price of solar panels in other markets. These instruments are valid under the assumption that supply shocks are correlated across markets but demand shocks are not (Hausman, 1996; Nevo, 2001). This assumption would be violated by demand shocks that affect multiple markets and are not captured in the market-specific subsidies to consumers included as regressors.

The demand states are the demand curve intercepts,

$$d_{mt} = \alpha_{0m} + \alpha_{2m} S_{mt} + \varepsilon^{D}_{mt}, \tag{3}$$

which must be estimated because the demand parameters (α) and shocks (ε^{D}) are not observed. I recover d_{mt} using the coefficients from estimating equation 2.

Production Costs I infer marginal costs from the firm's first order condition for optimal production and the demand estimates under the maintained assumption that firms compete in quantities (Cournot) with constant marginal costs and non-binding capacity constraints. Firm i's first order condition for market m and time t is

$$\frac{d\pi_{imt}}{dq_{imt}} = P_{mt} + \frac{dP_{mt}}{dq_{imt}} q_{imt} - mc_{imt} = 0.$$

The firm equalizes the marginal benefit and marginal cost of increasing the quantity it sells, accounting for the direct benefit and cost of producing one more unit (P_{mt} and mc_{imt}) as well as the inframarginal impact of depressing the equilibrium price on all its units ($dP_{mt}/dq_{imt} \cdot q_{imt}$). Under constant elasticity demand, marginal costs are

$$mc_{imt} = P_{mt} \left(1 + \frac{1}{\alpha_{1m}} \frac{q_{imt}}{Q_{mt}} \right),$$

where α_{1m} is the price elasticity of demand.

I parameterize these inferred costs to quantify the impact of technical efficiency improvements on production costs. I use a parametric form that is motivated by the economics of the industry. The measure of technical efficiency used in this study is similar to productivity in that it acts as a multiplier on the cost of materials in a manner similar to that of total factor productivity in a production function:

$$mc_{imt} = \tilde{\omega}_{it}^{\beta_1} w_t^{\beta_2} \exp(\beta_0 + \varepsilon_{it}^S)$$

where $\tilde{\omega}_{it}$ is observed firm-specific technical efficiency, w_t is the observed common input price, β_0 is an unobserved time-invariant common scale parameter, and ε_{it}^S is an unobserved firm-specific shock. The economic model formulated in Section 3 is in terms of cost-indexed technical efficiency, ω_{it}. There is a one-to-one mapping between observed technical efficiency and cost-indexed technical efficiency: $\omega_{it} = \tilde{\omega}_{it}^{\beta_1}$. The econometric model is used to estimate the parameter β_1 governing this relationship.[36,37] I estimate this cost function in logs,

$$\ln(mc_{imt}) = \beta_0 + \beta_1 \ln(\tilde{\omega}_{it}) + \beta_2 \ln(w_t) + \varepsilon_{it}^S. \tag{4}$$

I include time period and firm fixed effects in some specifications to capture unobserved factors that shift costs independently of variation in the input price and innovations in technical efficiency. I use ordinary least squares to estimate equation 4 under the assumption that ε_{it}^S is i.i.d. over firms and time.[38] This model quantifies the relationship between technical efficiency and marginal cost. The sign and magnitude of β_1 dictate whether and how much technical efficiency improvements lower costs, and therefore the incentive firms face to innovate in response to changes in demand.

[36]Observed technical efficiency, $\tilde{\omega}_{it}$, is the fraction of solar radiation that a solar panel converts into electricity, while the technical efficiency measure used in the economic model, ω_{it}, is in terms of production costs. The latter measure is used for convenience: I leverage cost-indexed technical efficiency along with properties of the underlying economic game to simplify the state space when solving the model for counterfactuals in Section 7. This is without loss of generality; the two measures are interchangeable under the assumption that firms know the mapping from observed to cost-indexed technical efficiency, which is necessary for firms to make investment decisions in the real world.

[37]I assume that each firm's cost is determined by its maximum efficiency and focus on this measure throughout the paper. This is consistent with my focus on technological innovation that advances the capabilities of the firm, and it is also a practical solution to the unavailability of data on the full distribution of technical efficiencies within a firm due to a lack of product-level data from all markets. In Section 6 I assess the robustness of this assumption by estimating an alternative model in which production costs depend on the mean – rather than the maximum – technical efficiency in the U.S.

[38]Technical efficiency and the input price are both fixed from the perspective of the firm at the time of product market competition. The use of time period fixed effects helps to account for the possibility of correlation in errors across firms within a time period. The i.i.d. errors assumption rules out serial correlation in errors within a firm. Allowing for an unobserved, serially correlated component of costs would substantially complicate the analysis.

State Transitions I estimate a vector autoregression model for the evolution of the exogenous states:

$$\begin{bmatrix} w_t \\ d_t \end{bmatrix} = R_0 + R_1 \begin{bmatrix} w_{t-1} \\ d_{t-1} \end{bmatrix} + \xi_t \tag{5}$$

where both R_0 and ξ_t are diagonal by assumption. I use ordinary least squares to separately estimate this model for each exogenous state (Hamilton, 1994).

I use forward simulation to construct the endogenous distribution of technical efficiencies by aggregating individual firm states. The evolution of individual firms' states and the distribution of firms' states are characterized by the investment policy function.

Investment Policy Function The investment policy function characterizes the investment behavior of firms conditional on their own state and the industry state. Consistent estimates of the policy function are necessary for estimation of the dynamic parameters. The ideal approach would be to use a nonparametric estimation strategy to capture this unknown and potentially complex function. This is infeasible in my setting, however, both because states are continuous and because firms' technical efficiencies are increasing over time. Instead, I adopt a data-driven approach to approximate the investment policy function using a parametric specification that balances the benefits of a very flexible specification with the potential pitfalls of overfitting.

I do this in two stages. First, I begin with a large set of candidate regressors and use lasso for variable selection. I model the discrete decision to invest by estimating a logit model via penalized maximum likelihood:

$$\min_{\mu} \frac{1}{N} \sum_{i,t} -x_{it} f(\omega_{it}, s_t)'\mu + \ln\left[1 + \exp\left(f(\omega_{it}, s_t)'\mu\right)\right] + \lambda \|\mu\|_1.$$

The candidate regressors in $f(\omega_{it}, s_t)$ are cubic polynomials in the following variables: the firm's technical efficiency, the mean and standard deviation of the industry distribution of technical efficiency, the input price, and the four demand states; as well as first-order interactions between the firm's technical efficiency and all other variables listed previously. The tuning parameter λ is selected by leave-one-out cross-validation.[39] I then model the investment choice using a logit model and estimate the

[39]The selected regressors are unchanged using k-fold cross-validation with $k = 5$ and $k = 10$.

parameters via maximum likelihood,

$$\min_{\tilde{\mu}} \frac{1}{N} \sum_{i,t} -x_{it}\tilde{f}(\omega_{it}, s_t)'\tilde{\mu} + \ln\left[1 + \exp\left(\tilde{f}(\omega_{it}, s_t)'\tilde{\mu}\right)\right],$$

where $\tilde{f}(\omega_{it}, s_t)$ contains only the non-zero regressors selected in the first stage. This two-step approach to policy function estimation is inspired by the attractive properties of ordinary least squares after model selection via lasso (Belloni and Chernozhukov, 2013).

I assume the outcome of the investment process is drawn from a stationary distribution conditional on making the decision to invest, as described in Section 3. I use nonparametric tests to assess this assumption in Appendix D. The test results are broadly consistent with my modeling assumption, although they have low power due to small sample size. I fit the distribution of investment levels using a gamma distribution. The gamma distribution captures the skewed nature of investment levels in the data and outperforms other candidate distributions – weibull, lognormal, and beta – in terms of the Akaike information criterion.

5.2 Second Step: Estimate Dynamic Parameters

I estimate the parameters of the investment cost function using a forward simulation estimator based on Bajari et al. (2007). This approach simulates industry paths based on the theoretical model and estimates from the first step in order to find parameters that make observed investment behavior optimal.

Firm *i*'s per-period payoff first introduced in Section 3 is

$$\pi_i(x_{it}, s_t; \varepsilon_{it}) = \bar{\pi}_i(s_t) - \gamma x_{it} + \sigma \varepsilon_{it}(x_{it}). \tag{6}$$

The firm's *ex-ante* value function, before realizing its private shocks, can be written as an expected discounted sum of per-period payoffs,

$$V_i\left(s_t; \zeta\right) = E\left[\sum_{\tau=0}^{\infty} \beta^{\tau} \pi_i(x_{it+\tau}, s_{t+\tau}; \varepsilon_{it+\tau})\right], \tag{7}$$

where the expectation is over current and future values of the private shocks (ε_{it}) and future values of the states (s_t). The dependence of per-period payoffs on strategies (ζ) is subsumed into x_{it}. I follow Bajari et al. (2007) by rewriting equation 6 as the inner product of two vectors and substituting it into equation 7 to give

$$V_i\left(s_t;\zeta\right) = E\left[\sum_{\tau=0}^{\infty}\beta^{\tau}\left[\begin{array}{ccc}\bar{\pi}_i(s_{t+\tau}) & -x_{it+\tau} & \varepsilon_{it+\tau}(x_{it+\tau})\end{array}\right]\right]\cdot\theta = W_i(s_t;\zeta)\cdot\theta, \qquad (8)$$

where $\theta = [\begin{array}{ccc}1 & \gamma & \sigma\end{array}]'$. $V_i(s_t;\zeta)$ is linear in parameters because the per-period payoff is linear in parameters. As a result, $W_i(s_t;\zeta)$ does not depend on θ and only needs to be simulated once for a given strategy profile.

I use forward simulation to approximate $W_i(s_t;\zeta)$ under the optimal strategy profile. For each initial state s_t, I construct a vector containing the elements of per-period payoffs using parameter estimates from the first step at $\tau = 0$. Product market profits, $\bar{\pi}_i(s_t)$, are treated as known and computed using the closed-form solution to the product market game.

The firm's investment decision is stochastic due to the presence of a private shock. I compute the probability of investment based on the first-stage policy function estimates. I then draw from the estimated policy function's error distribution. Each firm's draw determines that firm's investment. I use this information to increment the second and third elements of $W_i(s_t;\zeta)$.[40]

I construct the next period's state using the firms' simulated investment decisions. If a firm does not invest, its state does not change; if it invests, the change in its state (ν_{it}) is drawn from the distribution of technical efficiency improvements. Collectively, all firms' actions determine the industry's endogenous state in the next period. Finally, I use the estimated transition process for the exogenous states to predict the input price and demand states in the next period. I repeat this for 200 periods (50 years) to generate one discounted sum of product market profits, investment costs, and random shocks, up to the parameters θ.[41] I simulate 250 of these industry paths from each

[40] I first integrate out the private shocks from firm i's per-period payoff in equation 6 so that the second and third terms of $W_i(s_t;\zeta)$ are given by predictions from the policy function estimated in the first stage,

$$E_{\varepsilon_{it}}\left[\pi_i(x_{it}, s_t; \varepsilon_{it})\right] = \bar{\pi}_i(s_t) - \gamma p_i(x_{it} = 1|s_t) + \sigma\left[\varkappa - \sum_{x_{it}\in\{0,1\}} p_i(x_{it}|s_t)\ln\left(p_i(x_{it}|s_t)\right)\right],$$

where $p_i(x_{it}|s_t)$ is the probability that firm i chooses action x_{it} at state s_t, and $\varkappa$ is Euler's constant. The final term follows from the assumption that ε_{it} is drawn i.i.d. from the Type I extreme value distribution.

[41] The length of the forward simulation is arbitrary and is selected to ensure that the discounted profits from the terminal period are small relative to the discounted sum of profits over all periods. With an annual discount rate of 0.9, the discount factor for the 200$^{\text{th}}$ quarter – equivalent to the 50$^{\text{th}}$ year – is approximately 0.005, so that one dollar of profit in that period is worth only 0.5 cents from the perspective of the firm in period 0.

initial state and take the mean to approximate $W_i(s_t;\zeta)$. The discount factor, $\beta \approx 0.974$, corresponds to an annual discount factor of 0.9 used in prior literature (e.g., Ryan, 2012).

To estimate θ, I use $\widehat{W}_i(s_t;\zeta)$ to implement a maximum likelihood estimator based on the firm's optimal investment decision. Appendix E describes my approach in detail.

6 Estimation Results

6.1 First Step Estimates

Demand Table 2 presents the estimated price elasticities of demand ($\hat{\alpha}_{1m}$ from equation 2). Each column presents estimates from a different model specification, and each row corresponds to a different market (m). There are two main takeaways. First, the estimated price elasticities across all models and markets range between -1.3 and -2.2, with most coefficients falling between -1.3 and -1.5.[42] Second, the ordinary least squares and instrumental variables estimates are not statistically distinguishable, and the point estimates are quite similar in most cases. The similar point estimates could be because secular reductions in costs drive much of the observed price variation during the sample period and help trace out the demand curve even without instrumental variables.

To assess the robustness of these estimates, I estimate equation 2 with subsidy levels in logs rather than in levels for the markets with feed-in tariffs.[43] The results, presented in Table B1, are not statistically distinguishable from the baseline estimates presented in Table 2.

The specification that uses prices from other markets as instruments for price serves as the baseline specification used to estimate the supply model. I plot the resulting demand states with and without subsidies in Figure A1 to illustrate the impact of subsidies on demand. The higher, solid lines are the subsidy-inclusive demand

[42] My estimates lie within the wide range of previous estimates in the literature. Gillingham and Tsvetanov (2017) estimate a static demand elasticity of -0.65 using microdata from Connecticut. Coefficient estimates from De Groote and Verboven (2016) imply a static elasticity of close to -6.3 based on aggregate data from Flanders. Burr (2016) uses microdata from California to estimate long-run elasticities ranging from -1.6 to -4.7 across different time periods and model specifications. These papers estimate demand for residential solar systems, whereas I estimate demand for solar panels.

[43] This applies to Germany, Japan, and the residual market. The subsidy in the U.S. market is modeled as a fraction of the solar panel price since the primary subsidy available for the entire U.S. market is the Investment Tax Credit, which is worth 30% of solar system costs.

Table 2: Estimated Demand Elasticities

	Model		
	OLS	IV: Input Price	IV: Other Prices
Germany	-1.82^{**}	-2.20^{**}	-1.88^{**}
	(0.78)	(0.97)	(0.83)
Japan	-1.63^{***}	-1.37^{***}	-1.51^{***}
	(0.24)	(0.22)	(0.22)
Rest of World	-1.39^{***}	-1.32^{***}	-1.41^{***}
	(0.22)	(0.25)	(0.22)
USA	-1.49^{***}	-1.38^{***}	-1.48^{***}
	(0.26)	(0.26)	(0.26)
Min. F-stat		34.49	88.23
Observations	96	96	96
Adjusted R^2	0.74	0.73	0.74
Note:			*p<0.1; **p<0.05; ***p<0.01 Models estimated within market.

Notes: This table presents estimated price elasticities of demand (i.e., $\hat{\alpha}_{1m}$ from equation 2). Each row corresponds to a different market (m) and each column presents estimates from a different model specification. The first column presents estimates of equation 2 using ordinary least squares. The second column includes the following instruments for price: the price of polysilicon, the price of silver, and indicator variables for the presence of trade restrictions. The final column uses the prices of solar panels in markets other than the market of interest as instruments for the price in the market of interest.
Data Source: Author's calculations based on the model estimation described in Section 5 using data described in Section 4.

states ($\hat{d}_{mt}$). These are derived using the estimated parameters from equation 2 and the definition of the demand states in equation 3. The lower, dashed lines are the counterfactual demand had the subsidies not been in place. I recover counterfactual demand by subtracting the estimated impact of subsidies on demand ($\hat{d}_{mt} - \hat{\alpha}_{2m} S_{mt}$). Both are in terms of the natural logarithm of quantity, as they represent the demand curve intercept from the constant elasticity specification. The shaded area represents the portion of demand attributable to the subsidies included in this analysis.

Production Costs I first provide graphical evidence that suggests technical efficiency reduces costs: Figure 5 shows that there is a positive correlation between technical efficiency and market share in the raw data. This plot displays demeaned data within each market and time period, so that it only reflects cross-sectional variation. Each point on the graph is a local mean of the underlying data. This positive correlation is consistent with higher technical efficiency lowering costs for firms, which would lead to higher market share under a model of Cournot competition.

Figure 5: Cross-Sectional Relationship between Technical Efficiency and Market Share

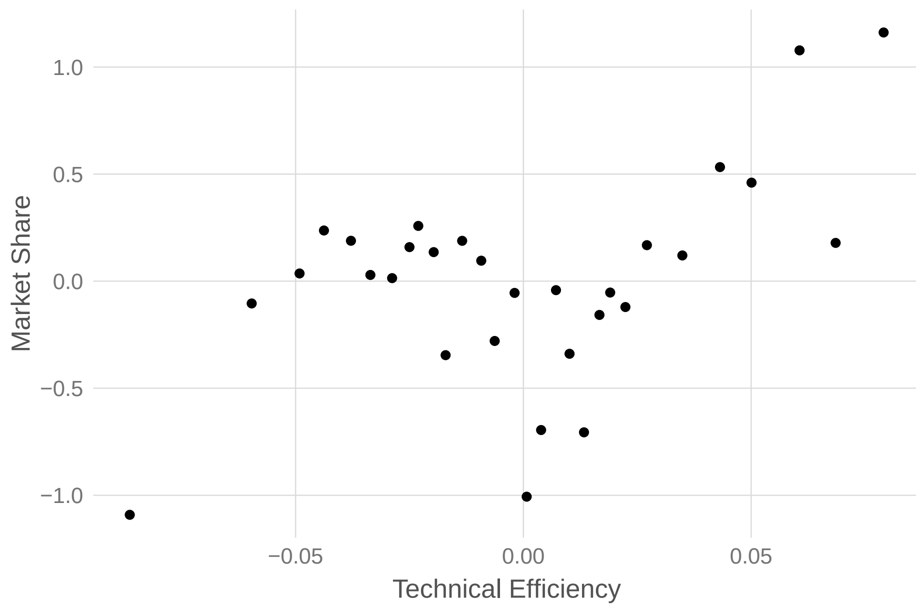

Notes: This figure is a binned scatterplot of residual variation in market shares and technical efficiencies within each market and time period. Each point on the graph is a local mean of the underlying data. The plot is constructed using raw data without any economic assumptions. The upward slope of these points is consistent with higher technical efficiency lowering costs for firms: under Cournot competition, this would yield higher market shares for firms with higher technical efficiencies.
Data Source: Author's calculations using data described in Section 4.

Figure 6 presents an analogous plot of the cross-sectional relationship between technical efficiency and production cost. Costs are recovered from the demand estimates and competition model as described in Section 5. The downward slope of these points shows that higher technical efficiency is associated with lower costs in the cross-section, after eliminating time series variation that could induce spurious correlation and undermine identification.

Table 3 presents the coefficient on technical efficiency across several alternative specifications of equation 4. The negative relationship between cost and technical efficiency is evident in all specifications despite the fact that the identifying assumptions under each model are different due to their use of different variation. The first specification includes only technical efficiency and a constant as regressors and produces a large negative coefficient. The second specification, which includes the common,

Figure 6: Cross-Sectional Relationship between Technical Efficiency and Cost

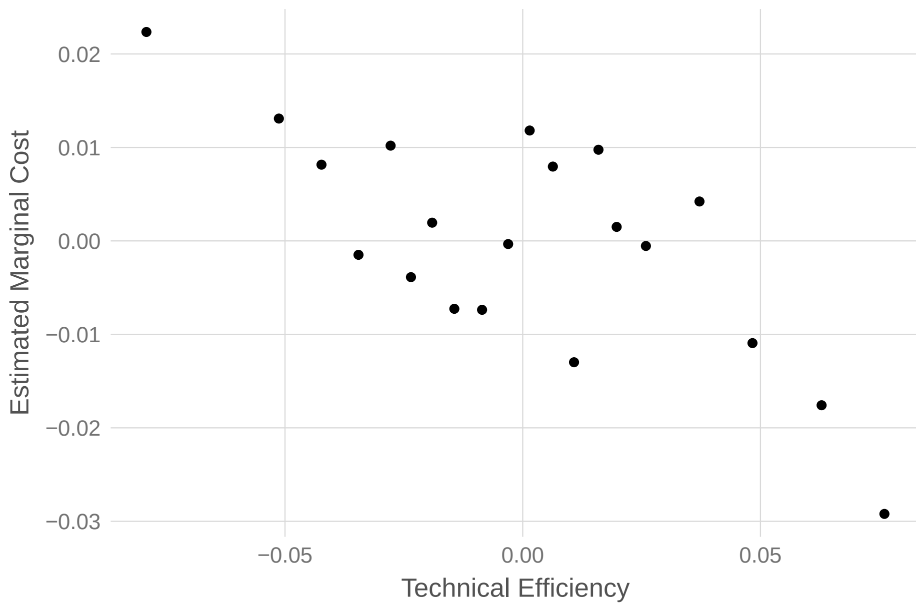

Notes: This figure is a binned scatterplot of residual variation in estimated marginal costs and technical efficiencies within each market and time period. Each point on the graph is a local mean of the underlying data. Marginal costs are inferred from the firm's first order condition for optimal production using data and estimated demand parameters, and technical efficiencies are observed in the data. The downward slope of these points indicates that higher technical efficiency is associated with lower marginal cost in the cross section.
Data Source: Author's calculations based on the model estimation described in Section 5 using data described in Section 4.

time-varying input price (w_t), produces a significantly attenuated coefficient on technical efficiency, highlighting the importance of accounting for time series variation in input prices. The third specification includes time period fixed effects to flexibly capture time series variation in unobserved costs. The input price is not included in this regression because it only varies over time and is therefore no longer identified. This further attenuates the coefficient on technical efficiency. In the final columns, I replicate the second and third specifications with the addition of firm fixed effects to capture unobserved firm-specific factors that affect cost and may be correlated with technical efficiency. These specifications generate slightly more negative coefficients on technical efficiency than the specifications without firm fixed effects. While there is variation in magnitudes across all the specifications, each model suggests that marginal costs are lower when technical efficiency is higher. The model estimates in columns 2-5 imply that improvements to technical efficiency constituted 3-32% of total cost reductions over the period 2010-2015.

Table 3: Relationship between Marginal Cost and Technical Efficiency

	(1)	(2)	(3)	(4)	(5)
$\ln(\tilde{\omega}_{it})$	-5.44^{***}	-0.92^{***}	-0.18^{*}	-1.80^{***}	-0.31^{*}
	(0.16)	(0.10)	(0.10)	(0.13)	(0.16)
$\ln(w_t)$		0.66^{***}		0.60^{***}	
		(0.01)		(0.01)	
Time Period FE			X		X
Firm FE				X	X
Observations	1,352	1,352	1,352	1,352	1,352
Adjusted R^2	0.47	0.89	0.92	0.90	0.92

Note: *p<0.1; **p<0.05; ***p<0.01
Data include 24 periods (T) for 4 markets (M).

Notes: This table presents coefficients from alternative specifications of the model in equation 4. The dependent variable is the natural logarithm of estimated marginal cost. The first row contains the coefficients on the regressor of interest, observed technical efficiency. The second row contains the coefficients on the common, time-varying price of polysilicon (w_t). There is a robust negative relationship between cost and technical efficiency (both in natural logarithms). The attenuation of coefficients from the first specification to all other specifications highlights the importance of conditioning on other factors that vary over time, such as observable input prices (w_t) and unobservable cost shifters (time period fixed effects).
Data Source: Author's calculations based on the model estimation described in Section 5 using data described in Section 4.

The qualitative results in Table 3 are robust to a range of alternative specifications summarized in Appendix C.1. I first allow the unobserved costs captured by time period fixed effects to vary for firms that manufacture in China and those that do not. Firms in China allegedly benefit from government subsidies that are unobserved and vary over time; interacting time period fixed effects with manufacturer location may account for this to the extent that the timing and the magnitude of subsidies are common across Chinese manufacturers. This more flexible specification has a negligible effect on the coefficient of interest (Table C1a). I also assess the robustness of the main results to the inclusion of measures of production capacity and past production, which proxy for economies of scale and experience (Tables C1b and C1c). Including these additional covariates attenuates the coefficient of interest. However, these additional variables have little impact on the final model, which includes time period and firm fixed effects. This suggests that differences in the coefficient of interest between the baseline specification and alternative specifications in the less restrictive models may be driven by firm-specific factors that are correlated with, but not due to, economies of scale and experience.

This paper focuses on the firm's maximum technical efficiency as an observable measure of innovation that leads to lower production costs, as discussed in Section 5. To assess the importance of this restriction, I present results from an alternative model in which production costs depend on the mean – rather than the maximum – technical efficiency in the U.S. market in Table C1d. The coefficient of interest is qualitatively similar in all cases and statistically indistinguishable in the specification used for estimation of the dynamic parameters.

A final possibility I consider is that consumers demand technical efficiency.[44] By abstracting from this mechanism, my approach may understate firms' returns to increasing technical efficiency. This could bias the investment cost estimates downward. While the full impact of relaxing this assumption on the counterfactual simulations is difficult to assess without estimating a model that allows for product differentiation, I discuss this possibility in more detail and present additional robustness analysis in Appendix C.2.

I use the model in column 3 of Table 3 for dynamic estimation because it relies only on cross-sectional variation and is therefore a conservative estimate, and because it does not include other dynamic choice variables, such as production capacity and past production, that would significantly complicate the analysis. This estimate identifies the causal impact of technical efficiency on the cost of production under the assumption that there are no omitted factors that are correlated with both technical efficiency and cost. I use the time period fixed effects recovered from estimation of the model in column 3 as the common, time-varying input price for estimating the dynamic model (i.e., w_t).

State Transitions Table A1 presents estimates of the state transition process for the input price and demand states (equation 5). The estimated transition process is stationary, as all eigenvalues of R lie within the unit circle.

Investment Policy Function The first-stage lasso procedure selects a small subset of the candidate regressors: a constant, the firm's technical efficiency (ω_{it}), and the

[44]Quantities are measured in Watts, so increasing technical efficiency – holding fixed the number of solar panels a firm produces – increases the firm's output in the model. This mechanical impact of technical efficiency on quantities captures much of the benefit of increased technical efficiency from the perspective of the firm and is the focus of this analysis.

interaction of the firm's technical efficiency and the industry average technical efficiency ($\omega_{it}\bar{\omega}_t$). As discussed in Section 5, I fit the distribution of investment levels using a gamma distribution. Figure A2 provides an assessment of the fit of the investment policy function and the parameterization of the investment outcome by comparing the industry average technical efficiency over time from the data to the analogous summary statistic from the forward simulation procedure used in estimation. The investment policy function's predictions track industry investment reasonably well despite its parsimony.

6.2 Second Step Estimates

Investment cost estimates are presented in Table 4. All numbers are in millions of dollars. The fixed cost point estimate is \$107 million.[45] In the data, investments occur slightly less often than once per year, so these estimates imply annual investment costs directed toward improving technical efficiency of about \$95 million. This is in line with accounting data: according to annual reports for a subset of the firms in this sample that are publicly traded, median R&D expenditures are about \$20 million and median annual capital expenditures are about \$200 million over the sample period.

Table 4: Investment Cost Parameter Estimates

Parameter	Point Estimate	Confidence Interval
γ	107.1	(61.1, 436.8)
σ	21.3	(10.2, 256.7)

Notes: This table presents estimates of the fixed cost of investing in technical efficiency improvements (γ) and the scale parameter on the private choice-specific shocks (σ). Confidence intervals are constructed via bootstrap, resampling residuals from each stage of estimation prior to forward simulation 500 times. All numbers are in millions of dollars.
Data Source: Author's calculations based on the model estimation described in Section 5 using data described in Section 4.

[45]Confidence intervals are constructed via bootstrap, resampling residuals from each stage of estimation prior to forward simulation 500 times.

7 Counterfactual Simulations

7.1 Short-Run Impacts of Subsidies

I first treat past advances in technical efficiency in the industry as exogenous and compare market equilibria with and without consumer subsidies. This provides a quantitative assessment of the "short-run" impacts of subsidies on equilibrium quantities and associated external social benefits accounting for only demand responses. This preliminary exercise requires only the demand and production cost estimates from Section 6.1 and does not require solving the dynamic model. I recover counterfactual demand states without subsidies from the results of demand estimation and solve for the product market equilibrium with and without subsidies.

Solar Adoption Subsidies increased the quantity of solar panels sold (in Watts) by 49% globally over the period 2010-2015 relative to sales in the absence of subsidies. The subsidies increased total solar panel market revenues by 60%, partly because subsidies in Germany and Italy encouraged substantial solar adoption in earlier years when costs (and therefore prices) were higher. Figure 7 plots model predictions for solar panel sales over time with and without subsidies. Figure 8 presents a regional decomposition of the solar adoption attributable to subsidies. Based on the model's estimates, German subsidies contributed 36% of the increase in demand attributable to subsidies, with Japanese subsidies contributing 30% and subsidies in the United States contributing 27%.

External Benefits The external social benefits attributable to consumer subsidies depend on the quantity of solar panels adopted due to subsidies, the amount of electricity the solar panels produce, the external damages associated with alternative electricity generation sources that solar electricity displaces, and the social discount rate. To construct a back-of-the-envelope estimate of the external benefits attributable to consumer subsidies, I consider a range of values for each of these.[46] Electricity generation estimates for each regional market come from PVWatts, a publicly available engineering tool that predicts solar electricity generation for different locations and solar system configurations.[47] I construct estimates of lifetime solar generation based

[46] This calculation of external benefits builds on Gillingham and Tsvetanov (2017).

[47] PVWatts is developed by the National Renewable Energy Laboratory and accessible via http://pvwatts.nrel.gov.

Figure 7: Short-Run Counterfactual: Impact of Subsidies on Solar Adoption

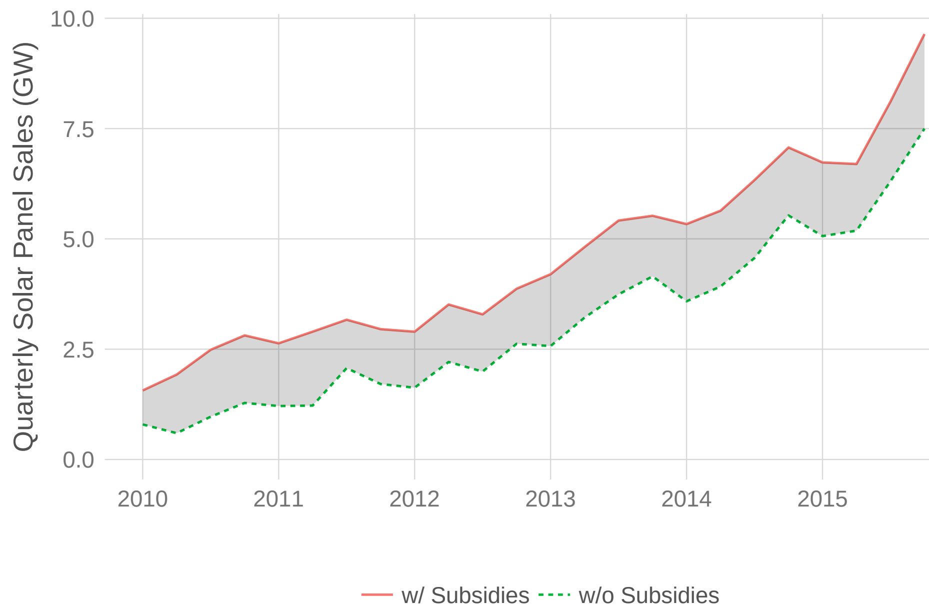

Notes: This figure plots model-predicted solar panel sales over time with and without subsidies holding firms' production costs fixed as estimated. The higher, solid line represents the model predictions based on historical subsidies. The lower, dashed line represents the counterfactual equilibrium quantities based on solving for the product market equilibrium after removing subsidies. The shaded area represents the portion of sales attributable to the subsidies included in this analysis.
Data Source: Author's calculations based on the model estimation described in Section 5 and counterfactual exercise described in Section 7 using data described in Section 4.

on potential electricity output in one location in each of the four markets.[48] Estimates of the external damages attributable to electricity generation from coal and natural gas, including both local air pollution and greenhouse gases, come from Muller et al. (2011).

My central estimate of the external benefits due to the solar subsidies studied in this paper is \$15.4 billion. Alternative assumptions imply external benefits that range from \$7.0 to \$56.1 billion (all in 2017 dollars).[49] This range reflects uncertainty about the parameters underlying this calculation. In addition, the subsidies studied in this

[48]I use the PVWatts default weather data for Germany, Japan, and the United States. I use the default weather data from a search for "Italy" to calculate electricity generated from subsidies in the residual market because the Italian feed-in tariff is used in demand estimation for the residual market.

[49]The central estimate is based on avoided damages from natural gas electricity generation (\$0.0203/kWh in 2017), a solar system lifetime of 25 years, and a discount rate of 3%. The low end estimate is based on natural gas, a lifetime of 20 years, and a discount rate of 7%. The high end estimate is based on avoided damages from coal electricity generation (\$0.0516/kWh in 2017), a lifetime of 30 years, and a discount rate of 3%.

Figure 8: Short-Run Counterfactual: Decomposition of Subsidy-Induced Adoption

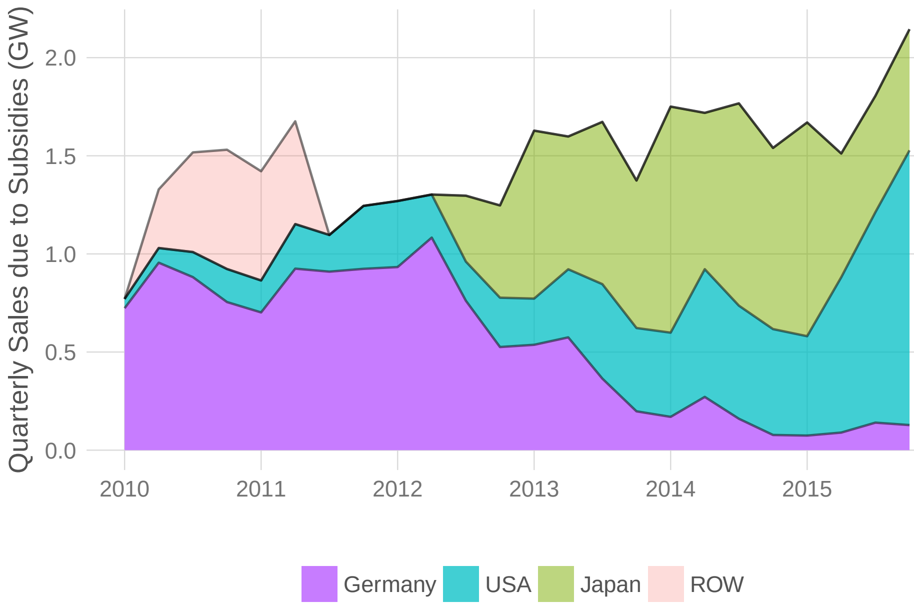

Notes: This plot decomposes subsidy-induced solar panel sales over time by the regional market of sale, holding firms' production costs fixed as estimated. In other words, this plot breaks out the shaded area represented in Figure 7 by market.
Data Source: Author's calculations based on the model estimation described in Section 5 and counterfactual exercise described in Section 7 using data described in Section 4.

paper do not capture all national and subnational policies that support solar. Thus, these estimates are meant to be illustrative of the potential order of magnitude of external benefits generated by solar subsidies. This is not a comprehensive assessment of the subsidies' benefits: a complete assessment of solar subsidies requires accounting for their long-run impacts allowing for both demand and supply responses, which I address in the next section.

7.2 Long-Run Impacts of Subsidies

To assess the long-run impacts of consumer subsidies, accounting for both demand and supply responses, I solve the dynamic model utilizing the estimates from Section 6. This requires some simplification of the state space due to computational constraints and for economic plausibility. In industries with a large number of firms, such as solar panel manufacturing, it would be computationally demanding for each firm to track and predict the current and future state of every one of its individual competitors. As an alternative, I assume that firms track moments of the distribution of firm states.

This approach is inspired by the experience-based equilibrium concept of Fershtman and Pakes (2012) and the moment-based Markov equilibrium of Ifrach and Weintraub (2017). I assume that firm strategies depend only on their own state and a set of functions of the state space, $\tilde{s}_t$, that is of lower dimension than the full state space, s_t. I leverage the fact that under Cournot competition, market equilibrium depends on the sum of firms' costs rather than the full distribution of firms' costs, and I assume that firm strategies depend only on their own state and the modified state space, $\tilde{s}_t \equiv (\sum_i \omega_{it}, w_t, d_t)$. This modified state space contains all the information needed to fully determine the product market equilibrium, conditional on the parameters that govern demand and production costs.[50]

I use value function approximation methods to solve for the equilibrium of the model with and without historical subsidies to consumers. For each case, I use a third-order polynomial expansion of the state space to approximate the value function, and I use parametric policy iteration to solve the dynamic model.[51] The resulting conditional choice probabilities, which describe investment behavior, can be used to simulate how the firms' technical efficiencies would have evolved in the absence of consumer subsidies and what share of the increase in industry average technical efficiency is attributable to consumer subsidies. I use this information to solve for the counterfactual product market equilibrium in each regional market over time, producing an estimate of the impact of consumer subsidies on solar panel adoption. Finally, I use these results to perform a back-of-the-envelope calculation using existing estimates of the environmental benefits of solar panel adoption to quantify the change in external benefits attributable to the subsidies, and I compare these long-run external benefits to the short-run external benefits presented above.

Solar Adoption I focus on a counterfactual that removes only German subsidies to highlight the potential for the impacts of policies to spill over across countries in a global market. Removing subsidies lowers firms' profits, leading firms to invest less frequently, production costs to be higher than under the baseline simulation, and the difference between simulated costs with and without German subsidies to grow over time. The change in investment activity is significant: 32% of the solar adoption due to increased technical efficiency would not have occurred in the absence of German

[50] To be precise, this property holds for Cournot competition when firms have constant marginal costs and the equilibrium is an interior solution in which all firms produce non-zero quantities and no production constraints bind. See Bergstrom and Varian (1985) for a clear derivation. This is a special case of a broad class of aggregative games.

[51] See Sweeting (2013) for a detailed summary of the procedure.

subsidies. This translates to a change in solar adoption over 2010-2015 of 0.7%, as under the conservative baseline production cost specification technical efficiency only constitutes about 3% of the sample period cost reductions. The vast majority of this marginal adoption occurs outside of Germany.

A central insight of this paper is that innovation responses by firms can swamp the contemporaneous impacts of subsidies over long time horizons. Furthermore, the innovation induced by past subsidies continues to generate benefits in the future, even if subsidies are phased out. This is especially relevant given that national governments have reduced subsidies in recent years, including in the major markets of Germany, Japan, and China. The U.S. Investment Tax Credit is currently scheduled to be phased down over the period 2019-2022. As a result, comparing solar adoption during 2010-2015 with and without endogenous innovation by firms will understate the long-run impacts of subsidies. To provide a level comparison, I estimate the discounted external social benefits attributable to innovation over a longer time horizon.

External Benefits Accounting for endogenous innovation by firms yields meaningfully different estimates of the external benefits of subsidies than accounting only for demand responses. The results are summarized in Table 5. The first row summarizes estimates of the external social benefits attributable to subsidies, accounting only for their short-run effects through demand responses. The second row's estimates include these benefits and the additional benefits that accrue due to supply responses induced by German subsidies that lowered solar panel production costs. The columns reflect a range of assumptions used to construct the present discounted value of external benefits from solar panel adoption.[52]

The external benefits increase 22% to $18.8 billion under the baseline assumptions. The magnitude of external benefits is sensitive to the assumptions used in this back-of-the-envelope calculation, ranging from $8 to $70 billion between the low and high

[52]The estimates in the first column are based on the avoided external damages of electricity generation from natural gas, a 25-year lifetime for solar panels, and a discount rate of 3%. Electricity generation estimates come from NREL's PVWatts tool using one location for each of the four regional markets in this study. The second column uses electricity output based on the location with the least sunlight (Germany), external damages from natural gas, a 20-year lifetime for solar panels, and a discount of 7%. The final column uses electricity output based on the location with the most sunlight (the United States), external damages from coal, a 30-year lifetime for solar panels, and a discount of 3%. For each specification, the dynamic impacts include both the present value of external benefits due to solar adoption induced by subsidies over 2010-2015, and also the present value of additional solar panels installed for x years after the sample years. For x, I use the assumed lifetime for solar panels under each scenario. The benefits from these future adoptions are discounted at the same rate used to discount benefits from solar adoption during the sample period. For these calculations, I hold market structure, costs, and subsidies fixed as of the final period in the data (2015Q4).

Table 5: Counterfactual External Benefits (billion $)

	Baseline	Low	High
Short-Run (without innovation)	15.4	7.0	56.1
Long-Run (with innovation)	18.8	8.0	69.5
Difference (%)	22.0	14.4	23.9

Notes: This table presents external benefits based on two sets of counterfactuals. The first row summarizes estimates of the external social benefits attributable to subsidies, accounting only for their short-run effects through demand responses. The second row's estimates include these effects and the additional benefits that accrue due to solar panel cost reductions induced by German subsidies. The columns reflect three different sets of assumptions used to construct the present discounted value of external benefits from solar panel adoption. See Section 7 for details on the underlying assumptions.
Data Source: Author's calculations based on model simulation described in Section 7 using model estimates described in Section 6.

cases. However, the increase in benefits due to subsidy-induced innovation is less sensitive, ranging from 14.4% to 24%.

These counterfactual simulations are based on the most conservative specification of the production cost model in Table 3. The long-run impacts of subsidies may exceed the short-run impacts under alternative specifications. For example, estimates from the cost specification that includes both time and firm fixed effects suggest that the causal effect of technical efficiency on costs may be larger than that used in this counterfactual simulation. This could imply that the counterfactual results understate the impact of subsidies on innovation, and therefore understate the external social benefits generated by subsidies. Furthermore, since this analysis focuses on just one margin of endogenous innovation, it is likely that these estimates understate the full impacts of subsidies to solar adoption.

8 Conclusion

I study the impact of consumer subsidies for solar panels on solar adoption and innovation by firms. I estimate a dynamic model to recover structural parameters governing the product market and firms' investments in innovation. I find that subsidies to consumers have a significant impact on demand and firms' revenues, and that improvements to technical efficiency significantly reduce production costs.

I use these estimates to evaluate the short- and long-run impact of solar subsidies. First, I conduct a short-run counterfactual simulation that allows demand to respond to subsidies but holds production costs fixed. I find that solar panel subsidies increased equilibrium quantities by roughly 50% over the period 2010-2015. This subsidy-induced solar adoption generates external social benefits of over $15 billion by replacing electricity from conventional, polluting sources with solar electricity. Second, I conduct a long-run counterfactual simulation that allows for both demand and supply responses, endogenizing firm investment in technical efficiency with and without observed subsidies in Germany. I find that accounting for induced innovation increases the external social benefits attributable to subsidies by at least 22%. This estimate includes future benefits from solar adoption over the 2010-2040 time period to account for the persistent effects of innovation. Much of these benefits are generated by solar adoption outside Germany. These findings demonstrate that consumer subsidies can indeed induce innovation but that national governments may not be able to appropriate the resulting innovation when markets are global in scope.

While the insights in this paper extend to similar policies, including a Pigouvian tax on emissions, I do not explicitly consider the implications of other forms of government intervention in the market for solar panels. The prevalence of subsidies to solar panel manufacturers raises the question of what the relative benefits of subsidies to consumption and production are in practice, both in terms of economic efficiency and distribution. Subsidies to production have led to trade actions against Chinese manufacturers by the European Union and the United States, which I am studying in a separate project. However, the results in this paper highlight that trade disputes could affect innovation incentives and thereby have long-run impacts in addition to the short-run impacts of trade restrictions on consumers, producers, and the environment.

The welfare implications of my results depend on additional assumptions. If policymakers are constrained to use consumer subsidies to solar power rather than a technology-neutral Pigouvian tax, a globally coordinated subsidy equal to the marginal external benefits of solar generation may be optimal if innovation is entirely appropriated by firms. However, this is unlikely to be the case, both because firms may learn from their competitors and because poor future performance could lead firms to exit the market. Both of these forces would lead firms to underinvest in innovation relative to the social optimum. To address this, I plan to incorporate a range of assumed innovation spillovers into the model to quantify how welfare changes as a function of both subsidies and innovation spillovers.

The optimal coordinated policy to address climate change may require a mix of measures that extend beyond carbon pricing. For example, public support for R&D

could be justified by the existence of knowledge spillovers or the impact of targeted R&D on the direction of innovation (Acemoglu et al., 2012). There is a growing literature on the optimal mix of R&D subsidies and carbon pricing to address climate change (e.g., Acemoglu et al., 2016; Aghion et al., 2016; Lemoine, 2017), and this remains an important area for future research.

References

Acemoglu, D., P. Aghion, L. Bursztyn, and D. Hemous (2012). The Environment and Directed Technical Change. *American Economic Review 102*(1), 131–166.

Acemoglu, D., U. Akcigit, D. Hanley, and W. Kerr (2016). Transition to Clean Technology. *Journal of Political Economy 124*(1), 52–104.

Acemoglu, D. and J. Linn (2004). Market Size in Innovation: Theory and Evidence from the Pharmaceutical Industry. *The Quarterly Journal of Economics 119*(3), 1049–1090.

Aghion, P., A. Dechezleprêtre, D. Hémous, R. Martin, and J. Van Reenen (2016). Carbon Taxes, Path Dependency, and Directed Technical Change: Evidence from the Auto Industry. *Journal of Political Economy 124*(1), 1–51.

Aguirregabiria, V. and P. Mira (2007). Sequential Estimation of Dynamic Discrete Games. *Econometrica 75*(1), 1–53.

Aguirregabiria, V. and A. Nevo (2013). Recent Developments in Empirical IO: Dynamic Demand and Dynamic Games. In D. Acemoglu, A. Arellano, and D. Dekel (Eds.), *Advances in Economics and Econometrics: Volume 3, Econometrics: Tenth World Congress*. New York: Cambridge University Press.

Arrow, K. J. (1962). The Economic Implications of Learning by Doing. *The Review of Economic Studies 29*(3), 155–173. 09621.

Bajari, P., C. L. Benkard, and J. Levin (2007). Estimating Dynamic Models of Imperfect Competition. *Econometrica 75*(5), 1331–1370.

Baker, E., M. Fowlie, D. Lemoine, and S. S. Reynolds (2013). The Economics of Solar Electricity. *Annual Review of Resource Economics 5*(1), 387–426.

Barbose, G. and N. Darghouth (2017). Tracking the Sun 10: The Installed Price of Residential and Non-Residential Photovoltaic Systems in the United States. Technical Report LBNL-2001062, Lawrence Berkeley National Laboratory.

Barwick, P. J. and P. A. Pathak (2015). The costs of free entry: an empirical study of real estate agents in Greater Boston. *The RAND Journal of Economics 46*(1), 103–145.

Belloni, A. and V. Chernozhukov (2013). Least squares after model selection in high-dimensional sparse models. *Bernoulli 19*(2), 521–547.

Benkard, C. L. (2004). A Dynamic Analysis of the Market for Wide-Bodied Commercial Aircraft. *The Review of Economic Studies 71*(3), 581–611. 00235.

Bergstrom, T. C. and H. R. Varian (1985). When Are Nash Equilibria Independent of the Distribution of Agents' Characteristics? *The Review of Economic Studies 52*(4), 715–718.

Blume-Kohout, M. E. and N. Sood (2013). Market size and innovation: Effects of Medicare Part D on pharmaceutical research and development. *Journal of Public Economics 97*, 327–336.

Bollinger, B. and K. Gillingham (2014). Learning-by-Doing in Solar Photovoltaic Installations. Working Paper.

Borenstein, S. (2017). Private Net Benefits of Residential Solar PV: The Role of Electricity Tariffs, Tax Incentives, and Rebates. *Journal of the Association of Environmental and Resource Economists 4*(S1), S85–S122.

Burr, C. (2016). Subsidies, Tariffs and Investments in the Solar Power Market. Working Paper, University of Colorado at Boulder.

Conlon, C. T. (2012). A Dynamic Model of Prices and Margins in the LCD TV Industry. Working Paper.

Creutzig, F., P. Agoston, J. C. Goldschmidt, G. Luderer, G. Nemet, and R. C. Pietzcker (2017). The underestimated potential of solar energy to mitigate climate change. *Nature Energy 2*, nenergy2017140.

De Groote, O. and F. Verboven (2016). Subsidies and Myopia in Technology Adoption: Evidence from Solar Photovoltaic Systems. Working Paper, University of Leuven.

Dubois, P., O. de Mouzon, F. Scott-Morton, and P. Seabright (2015). Market size and pharmaceutical innovation. *The RAND Journal of Economics 46*(4), 844–871.

Ericson, R. and A. Pakes (1995). Markov-Perfect Industry Dynamics: A Framework for Empirical Work. *The Review of Economic Studies 62*(1), 53–82.

Fershtman, C. and A. Pakes (2012). Dynamic Games with Asymmetric Information: A Framework for Empirical Work. *The Quarterly Journal of Economics 127*(4), 1611–1661.

Finkelstein, A. (2004). Static and Dynamic Effects of Health Policy: Evidence from the Vaccine Industry. *The Quarterly Journal of Economics 119*(2), 527–564.

Fu, R., D. Feldman, R. Margolis, M. Woodhouse, and K. Ardani (2017). US Solar Photovoltaic System Cost Benchmark: Q1 2017. Technical Report NREL/TP-6A20-68925, National Renewable Energy Laboratory (NREL), Golden, CO, United States.

Gillingham, K. and T. Tsvetanov (2017). Hurdles and Steps: Estimating Demand for Solar Photovoltaics. Working Paper.

Goulder, L. H. and K. Mathai (2000). Optimal CO2 Abatement in the Presence of

Induced Technological Change. *Journal of Environmental Economics and Management 39*(1), 1–38.

Gowrisankaran, G. and M. Rysman (2012). Dynamics of Consumer Demand for New Durable Goods. *Journal of Political Economy 120*(6), 1173–1219. 00000.

Green, M. A. (2016). Commercial progress and challenges for photovoltaics. *Nature Energy 1*, 15015.

Hamilton, J. D. (1994). *Time Series Analysis*. Princeton, New Jersey: Princeton University Press.

Harrington, W., R. D. Morgenstern, and P. Nelson (2000). On the Accuracy of Regulatory Cost Estimates. *Journal of Policy Analysis and Management 19*(2), 297–322.

Hausman, J. A. (1996). Valuation of new goods under perfect and imperfect competition. In *The economics of new goods*, pp. 207–248. University of Chicago Press.

Hotz, V. J. and R. A. Miller (1993). Conditional Choice Probabilities and the Estimation of Dynamic Models. *The Review of Economic Studies 60*(3), 497–529.

Hughes, J. E. and M. Podolefsky (2015). Getting Green with Solar Subsidies: Evidence from the California Solar Initiative. *Journal of the Association of Environmental and Resource Economists 2*(2), 235–275.

Ifrach, B. and G. Y. Weintraub (2017). A Framework for Dynamic Oligopoly in Concentrated Industries. *The Review of Economic Studies 84*(3), 1106–1150.

International Energy Agency (2013). Trends in Photovoltaic Applications 2013. Technical Report Report IEA-PVPS T1-23:2013, International Energy Agency, Paris.

International Energy Agency (2016). Trends in Photovoltaic Applications 2016. Technical Report Report IEA-PVPS T1-30:2016, International Energy Agency, Paris.

Jaffe, A. B., R. G. Newell, and R. N. Stavins (2002). Environmental Policy and Technological Change. *Environmental and Resource Economics 22*(1-2), 41–70.

Kalouptsidi, M. (2018). Detection and Impact of Industrial Subsidies: The Case of Chinese Shipbuilding. *The Review of Economic Studies 85*(2), 1111–1158.

Lach, S. and R. Rob (1996). R&D, Investment, and Industry Dynamics. *Journal of Economics & Management Strategy 5*(2), 217–249.

Lach, S. and M. Schankerman (1989). Dynamics of R & D and Investment in the Scientific Sector. *Journal of Political Economy 97*(4), 880–904.

Lemoine, D. (2017). Innovation-Led Transitions in Energy Supply. Working Paper 23420, National Bureau of Economic Research.

Maskin, E. and J. Tirole (1988). A Theory of Dynamic Oligopoly, I: Overview and Quantity Competition with Large Fixed Costs. *Econometrica 56*(3), 549–569.

Muller, N. Z., R. Mendelsohn, and W. Nordhaus (2011). Environmental Accounting for Pollution in the United States Economy. *The American Economic Review 101*(5), 1649–1675.

Nemet, G. F. (2006). Beyond the learning curve: factors influencing cost reductions in photovoltaics. *Energy Policy 34*(17), 3218–3232.

Nevo, A. (2001). Measuring Market Power in the Ready-to-Eat Cereal Industry. *Econometrica 69*(2), 307–342.

Newell, R. G., A. B. Jaffe, and R. N. Stavins (1999). The Induced Innovation Hypothesis and Energy-Saving Technological Change. *The Quarterly Journal of Economics 114*(3), 941–975.

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

Pakes, A., M. Ostrovsky, and S. Berry (2007). Simple estimators for the parameters of discrete dynamic games (with entry/exit examples). *The RAND Journal of Economics 38*(2), 373–399.

Pakes, A. and M. Schankerman (1984). An Exploration into the Determinants of Research Intensity. In *R&D, Patents, and Productivity*, pp. 209–232. University of Chicago Press.

Pillai, U. (2015). Drivers of cost reduction in solar photovoltaics. *Energy Economics 50*, 286–293.

Pillai, U. and J. McLaughlin (2013). A model of competition in the solar panel industry. *Energy Economics 40*, 32–39.

Popp, D. (2002). Induced Innovation and Energy Prices. *The American Economic Review 92*(1), 160–180.

Popp, D. C. (2001). The effect of new technology on energy consumption. *Resource and Energy Economics 23*(3), 215–239.

Powell, D. M., M. T. Winkler, H. J. Choi, C. B. Simmons, D. B. Needleman, and T. Buonassisi (2012). Crystalline silicon photovoltaics: a cost analysis framework for determining technology pathways to reach baseload electricity costs. *Energy & Environmental Science 5*(3), 5874–5883.

Ryan, S. P. (2012). The Costs of Environmental Regulation in a Concentrated Industry.

Econometrica 80(3), 1019–1061.

Schmookler, J. (1966). *Invention and Economic Growth*. Harvard University Press.

Sweeting, A. (2013). Dynamic Product Positioning in Differentiated Product Markets: The Effect of Fees for Musical Performance Rights on the Commercial Radio Industry. *Econometrica 81*(5), 1763–1803.

U.S. Department of Energy (2017). Energy Department Announces Achievement of SunShot Goal, New Focus for Solar Energy Office.

U.S. Energy Information Administration (2015). Direct Federal Financial Interventions and Subsidies in Energy in Fiscal Year 2013. Technical report, Washington, D.C.

U.S. Energy Information Administration (2016). Wind and Solar Data and Projections from the U.S. Energy Information Administration: Past Performance and Ongoing Enhancements. Technical report.

U.S. Environmental Protection Agency (2011). Regulatory Impact Analysis for the Final Mercury and Air Toxics Standards. Technical Report EPA-452/R-11-011, Office of Air Quality Planning and Standards, Research Triangle Park, NC.

U.S. Office of Management and Budget (2011). Memorandum for the Heads of Executive Departments and Agencies and of Independent Regulatory Agencies on Executive Order 13563, "Improved Regulation and Regulatory Review". Memorandum M-11-10, Office of Information and Regulatory Affairs, Washington, D.C.

van Benthem, A., K. Gillingham, and J. Sweeney (2008). Learning-by-Doing and the Optimal Solar Policy in California. *The Energy Journal 29*(3), 131–151.

Appendix A Additional Figures and Tables

Figure A1: Estimated Demand States with and without Subsidies

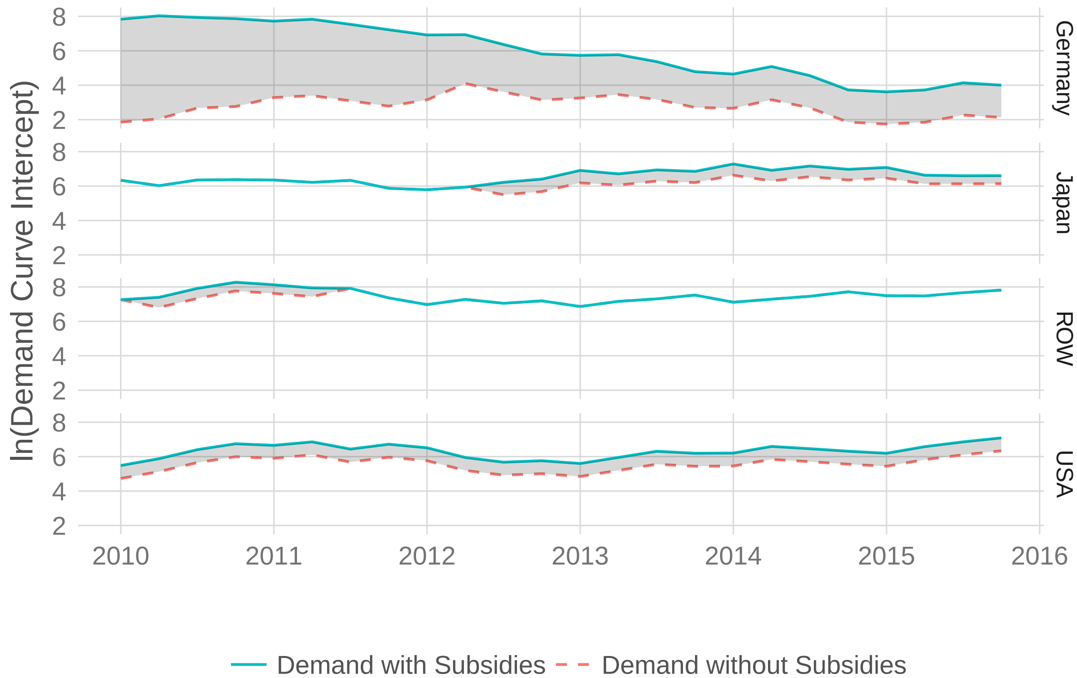

Notes: This figure plots estimated demand states with and without subsidies. The higher, solid lines are the subsidy-inclusive demand states (i.e., $\hat{d}_{mt}$) recovered from estimating equation 2 and inserting the estimates into equation 3. The lower, dashed lines are the counterfactual demand states had the subsidies not been in place. Both are in terms of the natural logarithm of quantity, as they represent the demand curve intercept from equation 2. The shaded area represents the portion of demand attributable to the subsidies included in this analysis. ROW denotes the residual market ("Rest of the World").

Data Source: Author's calculations based on the model estimation described in Section 5 using data described in Section 4.

Figure A2: Fit of Forward Simulation: Average Technical Efficiency

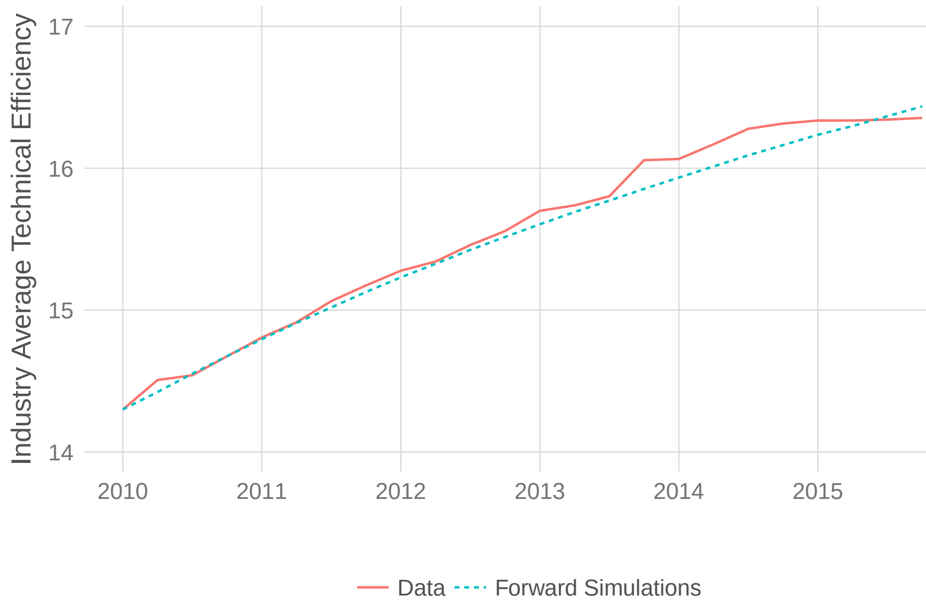

Notes: This figure provides an assessment of the fit of the investment policy function by comparing the industry average technical efficiency (i.e., the unweighted average of firms' maximum conversion efficiencies) over time from the data to the analogous summary statistic from the forward simulation procedure used in estimation of the dynamic parameters.

Data Source: Author's calculations based on the model estimation described in Section 5 using data described in Section 4.

Table A1: Estimates of Exogenous State Transitions

	w	Germany	Japan	ROW	USA
Constant	-0.035	0.736	2.885	4.227	1.605
	(0.017)	(0.396)	(1.129)	(1.545)	(0.817)
lag(w)	0.960				
	(0.028)				
lag(Germany)		0.742			
		(0.138)			
lag(Japan)			0.531		
			(0.183)		
lag(ROW)				0.428	
				(0.210)	
lag(USA)					0.722
					(0.148)
$\hat{\Sigma}_{\xi}$	0.0032	0.16	0.054	0.075	0.076
Observations	23	23	23	23	23

Note: Estimated separately via OLS with input cost (w) in logs.

Notes: This table presents coefficients from estimation of the vector autoregression model for exogenous states described by equation 5. The estimates are consistent with a stationary process for exogenous states, as the point estimates of the lag coefficients are all less than one in absolute value.

Data Source: Author's calculations based on the model estimation described in Section 5 using data described in Section 4.

Appendix B Robustness: Demand Estimates

Table B1: Estimated Demand Elasticities:

Subsidy Measures (Feed-in Tarrifs) in Logs rather than Levels

	Model		
	OLS	IV: Input Price	IV: Other Prices
Germany	-2.53^{***}	-2.82^{***}	-2.59^{***}
	(0.51)	(0.61)	(0.53)
Japan	-1.47^{***}	-1.08^{***}	-1.30^{***}
	(0.31)	(0.27)	(0.28)
Rest of World	-1.43^{***}	-1.36^{***}	-1.44^{***}
	(0.22)	(0.25)	(0.22)
USA	-1.49^{***}	-1.38^{***}	-1.48^{***}
	(0.26)	(0.26)	(0.26)
Min. F-stat		22.52	58.23
Observations	96	96	96
Adjusted R^2	0.81	0.80	0.80
Note:			*p<0.1; **p<0.05; ***p<0.01 Models estimated within market.

Notes: This table presents estimated price elasticities of demand (i.e., $\hat{\alpha}_{1m}$ from equation 2). Each row corresponds to a different market (m) and each column presents estimates from a different model specification. The first column presents estimates of equation 2 using ordinary least squares. The second column includes the following instruments for price: the price of polysilicon, the price of silver, and indicator variables for the presence of trade restrictions. The final column uses the prices of solar panels in markets other than the market of interest as instruments for the observed price in the market of interest. The coefficients are not statistically distinguishable from the baseline estimates presented in Table 2. The coefficients for the U.S. market are unchanged because the subsidies for that market are a function of the solar panel price.

Data Source: Author's calculations based on the model estimation described in Section 5 using data described in Section 4.

Appendix C Robustness: Marginal Cost Estimates

C.1 Alternative Specifications of the Production Cost Model

Table C1: Results under Alternative Production Cost Specifications

(a) Includes Time Period x Chinese Manufacturer FEs

	(1)	(2)	(3)	(4)	(5)
$\ln(\tilde{\omega}_{it})$	-5.44^{***}	-0.92^{***}	-0.18^{*}	-1.80^{***}	-0.32^{**}
	(0.16)	(0.10)	(0.10)	(0.13)	(0.16)
$\ln(w_t)$		0.66^{***}		0.60^{***}	
		(0.01)		(0.01)	
Period x China FE			X		X
Firm FE				X	X
Observations	1,352	1,352	1,352	1,352	1,352
Adjusted R^2	0.47	0.89	0.91	0.90	0.91

Note: $^{*}p<0.1$; $^{**}p<0.05$; $^{***}p<0.01$
Data include 24 periods (T) for 4 markets (M).

(b) Includes Scale Measure

	(1)	(2)	(3)	(4)	(5)
$\ln(\tilde{\omega}_{it})$	-5.13^{***}	-0.88^{***}	-0.15	-1.77^{***}	-0.26
	(0.16)	(0.10)	(0.10)	(0.14)	(0.16)
$\ln(w_t)$		0.65^{***}		0.60^{***}	
		(0.01)		(0.01)	
Time Period FE			X		X
Firm FE				X	X
Observations	1,352	1,352	1,352	1,352	1,352
Adjusted R^2	0.49	0.89	0.91	0.90	0.92

Note: $^{*}p<0.1$; $^{**}p<0.05$; $^{***}p<0.01$
Data include 24 periods (T) for 4 markets (M).

(c) Includes Experience Measure

	(1)	(2)	(3)	(4)	(5)
$\ln(\tilde{\omega}_{it})$	-3.19^{***}	-0.62^{***}	-0.10	-1.16^{***}	-0.29^{*}
	(0.17)	(0.10)	(0.10)	(0.15)	(0.16)
$\ln(w_t)$		0.61^{***}		0.54^{***}	
		(0.01)		(0.01)	
Time Period FE			X		X
Firm FE				X	X
Observations	1,352	1,352	1,352	1,352	1,352
Adjusted R^2	0.60	0.89	0.91	0.91	0.92

Note: $^{*}p<0.1$; $^{**}p<0.05$; $^{***}p<0.01$
Data include 24 periods (T) for 4 markets (M).

(d) Mean instead of Maximum Technical Efficiency

	(1)	(2)	(3)	(4)	(5)
$\ln(\tilde{\omega}_{it})$	-6.34^{***}	-1.18^{***}	-0.21^{*}	-1.65^{***}	-0.35^{**}
	(0.16)	(0.12)	(0.13)	(0.14)	(0.16)
$\ln(w_t)$		0.63^{***}		0.60^{***}	
		(0.01)		(0.01)	
Time Period FE			X		X
Firm FE				X	X
Observations	1,239	1,239	1,239	1,239	1,239
Adjusted R^2	0.57	0.88	0.91	0.90	0.91

Note: $^{*}p<0.1$; $^{**}p<0.05$; $^{***}p<0.01$
Data include 24 periods (T) for 4 markets (M).

Notes: These tables present coefficients from alternative specifications of the model in equation 4. The dependent variable is the natural logarithm of estimated marginal cost. The first row contains the coefficients on the regressor of interest, observed technical efficiency. The second row contains the coefficients on the common, time-varying price of polysilicon (w_t). The finding that technical efficiency lowers production cost is robust: in model 3, the specification used to estimate the dynamic model, the coefficient on technical efficiency in each of these alternative models is statistically indistinguishable from the baseline specification in Table 3.
Data Source: Author's calculations based on the model estimation described in Section 5 using data described in Section 4.

C.2 Potential Demand Impacts of Technical Efficiency

This paper studies the impact of consumer subsidies on cost-reducing innovation by focusing on endogenous technical efficiency improvements by firms to lower their future costs. The baseline model treats all solar panels as homogeneous conditional on their electricity output. This model captures the mechanical effect of increased technical efficiency on output: if a firm increases the technical efficiency of its solar panels and therefore can produce more electricity from an individual solar panel, consumers will be willing to pay more for that individual solar panel. However, this approach abstracts from potential demand effects of increased technical efficiency that are not captured by electricity generation potential. This implies that, for example, consumers are indifferent between 10 solar panels with 16.5% efficiency and 11 solar panels with 15% efficiency, holding the physical size of the solar panels constant.

In reality, consumers may be willing to pay more to get the same electricity generation capacity from a smaller number of solar panels due to costs for space, mounting hardware, electrical equipment, and labor that scale with the area of the solar system (Fu et al., 2017). By focusing only on costs and omitting this potential demand effect, the main results may understate the impact of technical efficiency on profits. This would, in turn, bias my investment cost estimates downward and potentially affect the conclusions of counterfactual analysis.

To provide a sense of the importance of this assumption, I use a normalization to account for this demand effect based on the optimization problem of a downstream firm that incorporates insights from the cost model in Fu et al. (2017). This captures the benefits of higher efficiency solar panels to consumers in a reduced-form manner.[53] The resulting cost estimates are contained in Table C2. The estimated coefficients on technical efficiency are larger than the baseline coefficients in Table 3, as expected. However, the additional impact of allowing for this demand effect is smaller than the magnitude of the pure cost effect estimated in the baseline model.

[53]This is meant to provide some insight into the potential ramifications of relaxing the assumption that solar panels are undifferentiated conditional on output. An alternative would be to estimate a differentiated products model, directly estimating consumers' preferences for technical efficiency (and potentially other characteristics such as warranty and brand). Incorporating this into the full model would require modification of firms' optimization problem.

Table C2: Relationship between Marginal Cost and Technical Efficiency:

Prices and Quantities Normalized by Efficiency

	(1)	(2)	(3)	(4)	(5)
$\ln(\tilde{\omega}_{it})$	-5.80^{***}	-1.17^{***}	-0.29^{***}	-2.20^{***}	-0.43^{***}
	(0.16)	(0.10)	(0.10)	(0.14)	(0.17)
$\ln(w_t)$		0.67^{***}		0.60^{***}	
		(0.01)		(0.01)	
Time Period FE			X		X
Firm FE				X	X
Observations	1,352	1,352	1,352	1,352	1,352
Adjusted R^2	0.48	0.88	0.92	0.90	0.92

Note: *p<0.1; **p<0.05; ***p<0.01
Data include 24 periods (T) for 4 markets (M).

Notes: This table presents coefficients from alternative specifications of the model in equation 4. These are based on a data normalization that account for the possible demand effects of increasing technical technical efficient. The coefficients on technical efficiency are larger in magnitude than the baseline coefficients in Table 3, consistent with a modest impact of technical efficiency on demand.

Data Source: Author's calculations based on the model estimation described in Section 5 using data described in Section 4.

Appendix D Investment Outcomes

The model in Section 3 assumes that firms make a discrete choice of whether to invest, and that the improvement in technical efficiency conditional on investing, ν_{it}, is i.i.d. across firms and time. I use data on technical efficiency to assess whether this assumption is reasonable. To summarize, the tests that follow fail to reject this assumption, although they have low power due to small sample size.

Comparison of Efficiency Improvements by Firm The assumption on ν_{it} implies that all firms draw from the same distribution. To assess this visually, I plot observed changes in technical efficiency by firm in Figure D1a. I formalize this comparison by applying the Kolmogorov-Smirnov test to each pair of firms under the null hypothesis that each pair of sets of observed changes in technical efficiency are drawn from the same distribution. I compute p-values, pool them across all the pairwise combinations, and evaluate the distribution of p-values at three quantiles:

10%	50%	90%
0.08	0.36	0.87

While this pooled comparison of individual tests is informal, the distribution of p-values is not heavily skewed toward zero. I also use a Kruskal-Wallis rank sum test for a joint (rather than pairwise) comparison under the null hypothesis that all firm-level samples originate from the same distribution. The result is a p-value of 0.09.

Comparison of Efficiency Improvements by Time Period I plot observed changes in technical efficiency by date in Figure D1b. I apply the Kolmogorov-Smirnov test to each pair of time periods under the null hypothesis that each pair of sets of observed changes in technical efficiency are drawn from the same distribution. I compute p-values, pool them across all the pairwise combinations, and evaluate the distribution of p-values at three quantiles:

10%	50%	90%
0.29	0.68	0.95

I use a Kruskal-Wallis rank sum test for a joint (rather than pairwise) comparison under the null hypothesis that all firm-level samples originate from the same distribution. The result is a p-value of 0.24.

Serial Correlation Perhaps the most important economic implication of the assumption on ν_{it} is that it rules out serial correlation in a given firm's outcomes over time. To assess this, I plot the autocorrelation function by firm in Figure D2. This visual test relies on a small number of non-missing observations in the time series for each firm because firms rarely invest, so it is only suggestive. Still, I informally interpret the plots as failing to reject the null that ν_{it} is not serially correlated within a firm.

Figure D1: Investment Realizations

(a) Investment Realizations (ν_{it}) by Firm (i)

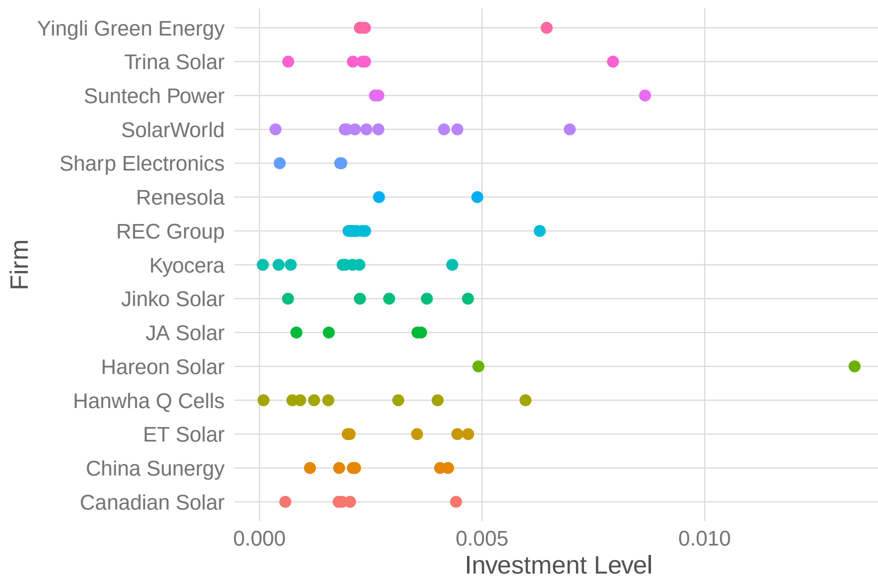

(b) Investment Realizations (ν_{it}) by Date (t)

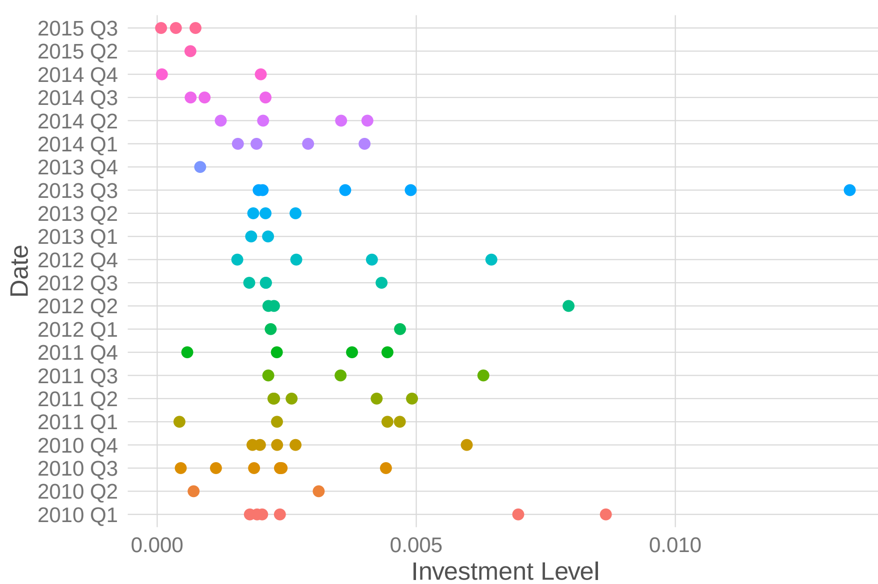

Figure D2: Autocorrelation of Investment Level (ν_{it}) within a Firm

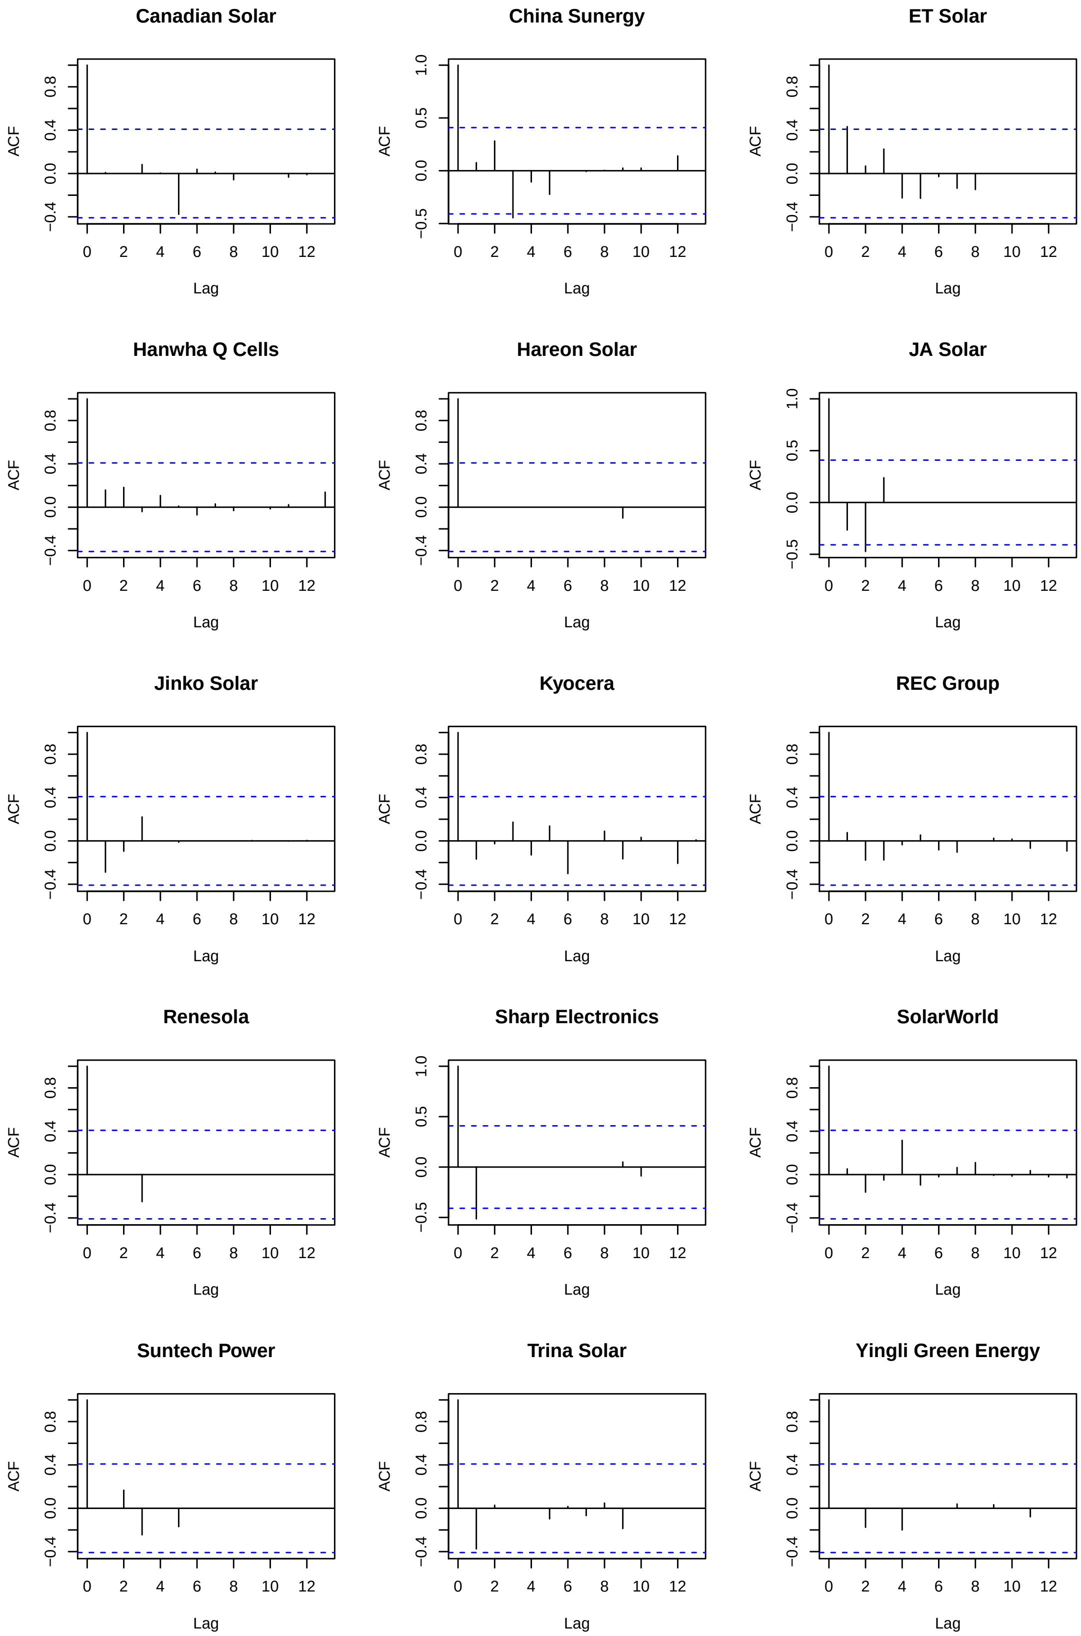

Appendix E Dynamic Parameter Estimation

The firm's investment optimization problem is:

$$V_i\left(s_t\right) = \max_{x_{it}\in\{0,1\}} \ \bar{\pi}_i(s_t) - \gamma x_{it} + \sigma\varepsilon_{it}(x_{it}) + \beta E\left[\left. V_i\left(s_{t+1}\right)\right| s_t, x_{it}\right],$$

where the dependence of the value function on strategies (ζ), private shocks (ε_{it}), and parameters (θ) are omitted for clarity. Under the assumption that the choice-specific error terms – $\varepsilon_{it}(x_{it})$ – are i.i.d. Type I extreme value, the *ex-ante* probability of investment is:

$$\Pr(x_{it} = 1) = \frac{\exp(v_1)}{\exp(v_0) + \exp(v_1)}$$

where

$$v_0 = \frac{1}{\sigma}\left(\bar{\pi}_i(s_t) + \beta E\left[\left. V_i\left(s_{t+1}\right)\right| s_t, x_{it} = 0\right]\right) \quad \text{and}$$
$$v_1 = \frac{1}{\sigma}\left(\bar{\pi}_i(s_t) - \gamma + \beta E\left[\left. V_i\left(s_{t+1}\right)\right| s_t, x_{it} = 1\right]\right).$$

The log-likelihood is:

$$\ln\mathcal{L}(\theta|x) = \sum_i\sum_t (1 - x_{it})v_0 + x_{it}v_1 - \ln\left(\exp(v_0) + \exp(v_1)\right)$$

Evaluating the log-likelihood requires knowledge of $E\left[\left. V_i\left(s_{t+1}\right)\right| s_t, x_{it}\right]$. I recover this expectation using a combination of forward simulation and value function approximation as described in Section 5.[54] I rewrite $V_i(s_t)$ as $W_i(s_t;\zeta)\cdot\theta$ and forward-simulate expected profits, investments, and private shocks for all the states observed in the data in a preliminary step to estimate $W_i(s_t;\zeta)$. I then use $\widehat{W}_i(s_t;\zeta)$ to evaluate $\widehat{V}_i(s_t)$ at a given choice of θ and model the value function as a flexible parametric function of the underlying state variables, $\widehat{V}_i(s_t;\theta) = f(s_t)'\lambda$. For $f(s_t)$, I use a third-degree polynomial expansion of the vector $[\ \omega_{it} \quad \bar{\omega}_t \quad \mathrm{sd}(\omega_t) \quad d_{mt} \quad w_t\ ]$, where I use the mean and standard deviation to summarize the distribution of technical efficiencies (ω_t).

[54]In principle this could be done entirely via forward simulation, without the use of value function approximation, by drawing from the set of possible next period states for each state observed in the data and forward-simulating the value function (up to parameters) for each of these draws. The computational burden of this approach grows linearly with the number of draws used to compute the multidimensional integral defined by the expectation operator, so I use value function approximation to allow for a large number of draws.

This is similar to the specification used for the first-stage reduced-form investment policy function. I estimate λ via ordinary least squares.

I use this parametric approximation to recover $E\left[V_i\left(s_{t+1}\right) | s_t, x_{it}\right]$ by drawing from the distribution of states that can be reached in one period from each state observed in the data – conditional on the firm's investment choice, x_{it} – and use that to evaluate the log-likelihood. Finally, I search over θ, repeating these steps for each candidate parameter value, to maximize the log-likelihood.