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

Todd D. Gerarden*

June 2022

Abstract

Private sector innovation is critical to mitigating and adapting to climate change. This paper studies innovation in solar energy technology, a key source of clean energy that has experienced rapid price declines over the past decade. To understand the causes and effects of innovation, I estimate a dynamic structural model of competition among solar panel manufacturers. The model captures important features of the industry, including the role of government subsidies for solar adoption, and I employ a unique measure of technological progress that is observable and verifiable. The results produce two main insights. First, ignoring innovation by firms can generate biased estimates of the effects of government policy. Second, decentralized government intervention in a global market generates spillovers: a subsidy in one country causes international firms to innovate more, leading to lower prices and increased adoption elsewhere. This spillover underscores the need for international coordination by governments and the private sector to address climate change.

---

*Charles H. Dyson School of Applied Economics and Management, Cornell University, Warren Hall, Ithaca, NY 14853. Email: gerarden@cornell.edu. Website: toddgerarden.com. I am grateful to Joe Aldy, Myrto Kalouptsidi, Ariel Pakes, Rob Stavins, and Jim Stock for their invaluable guidance and support. I thank Victor Aguirregabiria, Jonathan Baker, Chris Conlon, Jackson Dorsey, Ken Gillingham, Evan Herrnstadt, Robin Lee, Jing Li, Alex MacKay, Marc Melitz, Mark Shepard, Sam Stolper, Rich Sweeney, Elie Tamer, and numerous seminar participants 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. This paper has not been reviewed by the EPA and the views expressed in this paper are solely those of the author.

1 Introduction

Can subsidies to consumers spur firms to innovate? In principle, government subsidies to consumers that address 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 static and dynamic 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.

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 geographic markets and over time. Incumbent firms engage in Cournot competition in each geographic market. I employ a unique observable measure of technological innovation, the energy conversion efficiency of solar panels. Energy conversion efficiency is the fraction of incoming solar power that a solar panel converts into electrical power. It is a common, verifiable measure used by all firms in the solar industry. Increasing energy conversion efficiency helps firms lower costs by reducing the materials needed to produce solar panels. This gives firms an incentive to make costly fixed investments in energy conversion efficiency to lower their marginal costs in order to increase their future profits in the product market. I assume that firms condition only on the current industry state – firms' energy conversion 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 2010 through 2015 for four regional markets that span the globe: Germany, Japan, the United States, and the Rest of the World. I estimate demand for solar panels (in Watts) in each market to recover price elasticities and the impact of subsidies on demand. I use the demand estimates and firms' first order conditions for optimal production to recover marginal costs for the 15 firms that comprise approximately 70% of the global conventional solar panel market. I then estimate the relationship between these costs and energy conversion efficiencies. This relationship is identified using variation in energy conversion efficiencies and market shares: under Cournot competition, firms with lower costs have larger market shares, so a positive correlation between energy conversion efficiency and market share implies that energy conversion efficiency reduces costs (holding other factors fixed).

I use a two-step estimator similar to Bajari et al. (2007) to recover the fixed costs of investment. This approach leverages the insight that a firm's value function is equivalent to the expected discounted sum of its future net profits. First, I use estimates of the product market model, firm investment behavior, and state transitions to forward-simulate potential industry paths, tracking firms' profits and investment decisions. I then combine the simulated value function and the optimality condition for the firm's investment problem to estimate fixed costs via pseudo maximum likelihood.

Fixed costs are identified by variation in the frequency and the expected benefits of investment in different industry states. In the data, observed firm-level energy conversion efficiency improvements occur slightly less than once per year on average across states. The product market model predicts the benefits – in terms of future profits – of that investment. Fixed costs 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 model pins down the fixed costs.

The model estimates reveal two preliminary results. First, subsidies to consumers have significantly increased solar panel adoption, even before accounting for innovation. Without the demand response induced by subsidies, 51% less solar power generation capacity would have been adopted over the period 2010-2015.[1] Second, energy conversion efficiency improvements explain 3% to 32% of the reduction in marginal production costs over the period 2010-2015. Firms with higher efficiencies have lower costs on average. I interpret this relationship as a causal impact under the identifying assumption that there are no unobserved determinants of cost that are correlated with energy conversion efficiency in the cross section, which is plausible given its robustness to potential confounders. Estimated fixed costs of improving energy conversion efficiency 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 model to quantify the causes and effects of solar innovation by considering several counterfactual scenarios with different subsidy schemes. Two key findings emerge. First, I find that accounting for induced innovation by firms significantly increases the economic benefits attributable to subsidies. Without subsidies, solar adoption over the period 2010-2015 would have been 78% lower than observed after accounting for both demand and supply responses. This represents a 54% increase in solar adoption relative to the estimate above that only accounts for demand responses. As a result, solar subsidies produced significant increases in consumer surplus and external benefits due to avoided pollution from competing sources of electricity like natural gas.

Second, the economic benefits of solar subsidies extend beyond national boundaries. For example, in the absence of German subsidies, firms would have innovated less, and so their production costs would have been higher. Because solar manufacturers are multinational firms that compete in product markets throughout the world, these innovation decisions affect outcomes outside Germany. To quantify this spillover, I estimate how much additional solar adoption occurred in each regional market due to innovation induced by German subsidies. In total over the period 2010-2015, the model predicts that 86% of this additional solar adoption occured outside Germany.

[1] This is qualitatively and quantitatively consistent with prior studies from smaller geographic markets (e.g., Hughes and Podolefsky, 2015; De Groote and Verboven, 2019; Gillingham and Tsvetanov, 2019).

These results highlight two important points that apply beyond the solar market. First, using static economic methods to analyze markets experiencing rapid innovation can be misleading. It is important for firms and governments to recognize that innovation decisions can have larger impacts on profits and social objectives than short-term decisions about production or adoption of a technology. Second, international coordination may be needed to address innovation spillovers across borders when markets are interconnected.

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). The novelty of my approach relative to previous research on energy innovation is the use of a unique, observable, and verifiable measure of innovation: energy conversion 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, energy conversion efficiency improvements are measured in common units, allowing me to directly model the impact of individual innovations on costs.

The fact that energy conversion efficiency is observable is also a strength of this approach relative to the use of estimated measures such as production costs or productivity.[2] 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.

This paper also contributes to the growing literature on the economics of solar power (Baker et al., 2013). Several papers have found that past subsidies significantly increased adoption of solar systems (e.g., Hughes and Podolefsky, 2015; De Groote and Verboven, 2019; Gillingham and Tsvetanov, 2019). The consensus of this research is that many solar

[2]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.

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 (Goulder and Mathai, 2000; van Benthem et al., 2008), but there is limited empirical evidence to assess and guide policy.[3]

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) and Pillai (2015) decompose historical price reductions based on observable factors. In contrast, I develop and estimate a model that yields counterfactual predictions to quantify the impact of consumer subsidies on innovation.

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 energy conversion efficiency to the industry. Section 3 details the model of the industry environment and manufacturer behavior. Section 4 introduces the data used for estimation. Section 5 outlines my estimation approach, and Section 6 summarizes the estimates. Section 7 describes the counterfactual analysis. Section 8 concludes.

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

[3] In a notable exception, Bollinger and Gillingham (2019) estimate the extent of learning-by-doing spillovers in solar panel installation, the industry downstream of the firms I study.

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 category is subsidies that lower the upfront cost of adopting a solar system. An example of this is the Federal Investment Tax Credit (ITC) in the United States, which defrays solar investments costs by 30%. The second category is subsidies to electricity generation from solar systems after they are adopted. Germany, Japan, China, and many other countries have offered payments for solar electricity in the form of "feed-in tariffs." Feed-in tariffs are prices paid for solar electricity generation that are independent of the cost of electricity from alternative sources, and are determined at the time of investment rather than at the time of electricity generation. These feed-in tariffs are typically much higher than the prices that other electricity generation technologies receive (in the absence of other government support). Thus, feed-in tariff policies increase solar system revenues, providing an indirect subsidy to solar system adoption. Because demand for solar panels is derived from demand for solar systems, this indirect subsidy to solar systems leads to an outward shift in demand for solar panels in the upstream market.

Taken together, these government policies that affect the net present value of solar system adoption appear to significantly shift 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 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 was introduced prior to the sample period (2010-2015), but variation in its 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). Furthermore, previous empirical analyses of subnational consumer subsidies have found that a majority of observed solar adoption was attributable to subsidies (e.g., Hughes and Podolefsky, 2015; De Groote and Verboven, 2019; Gillingham and Tsvetanov, 2019).

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 energy conversion 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.[4]

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

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

Figure 1: Graphical Evidence that Subsidies Increase Demand

(a) Japan

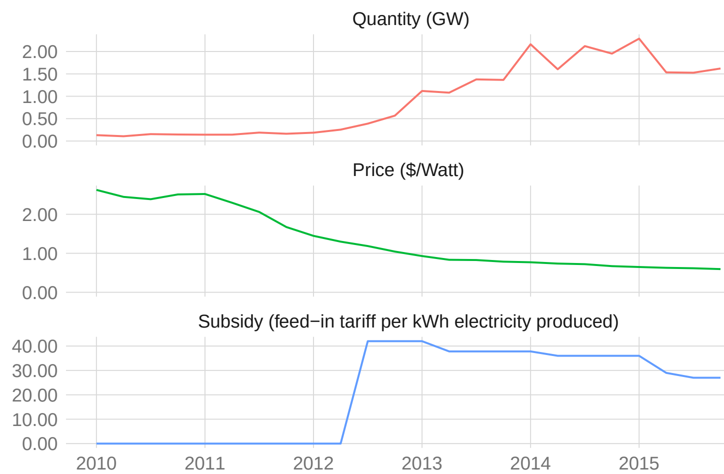

(b) Germany

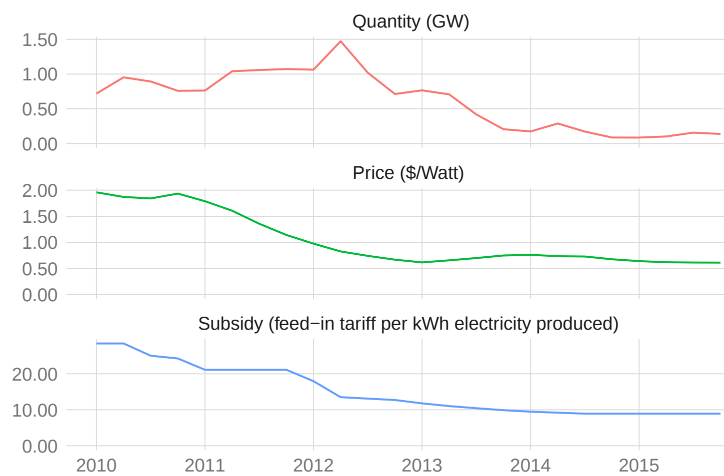

Notes: These two figures plot quarterly observations of each market's quantity of solar panel sales (GW), average price ($/W), and subsidy level per unit of electricity generation (yen/kWh for Japan, euro cents/kWh for Germany). Both figures are suggestive of the impact of subsidies on demand. Figure 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.

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 incentives until these innovations had occurred.

2.2 Solar Panel Manufacturing

Solar panels convert sunlight into electricity via the photovoltaic effect.[5] 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).[6] 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.[7] 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 there have been significant advances in firms' energy conversion efficiencies over time. Energy conversion efficiency is the fraction of incoming solar power

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

[6]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.

[7]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.

that a solar panel converts into electrical power, expressed in percentage points. The median energy conversion 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 energy conversion efficiency of their solar panels, although there is considerable variation in firms' energy conversion efficiencies in the cross section and in the relative position of firms over time (Figure 3).

Firms strive to increase their energy conversion efficiencies in order to reduce their production costs. Energy conversion 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).[8] Firms cite energy conversion 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 energy conversion 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."[9] 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.[10]

Solar panel manufacturers improve their energy conversion 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

[8] 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."

[9] Trina Solar also cited the role of energy conversion 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.

[10] 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.

Figure 2: Industry Energy Conversion Efficiency is Increasing over Time

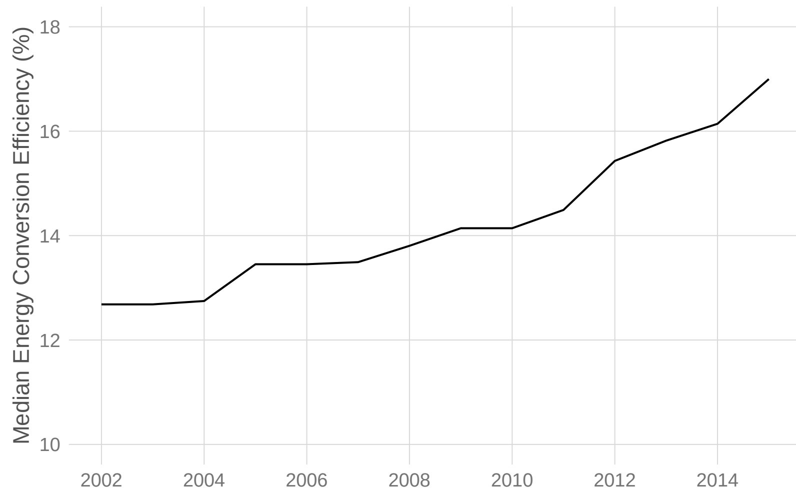

Notes: This plot shows the industry-wide progression of energy conversion efficiency based on data from the United States solar market. Energy conversion efficiency is the fraction of incoming solar power that a solar panel converts into electrical power. During 2010-2015, the median energy conversion 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 Energy Conversion Efficiencies over the Sample Period

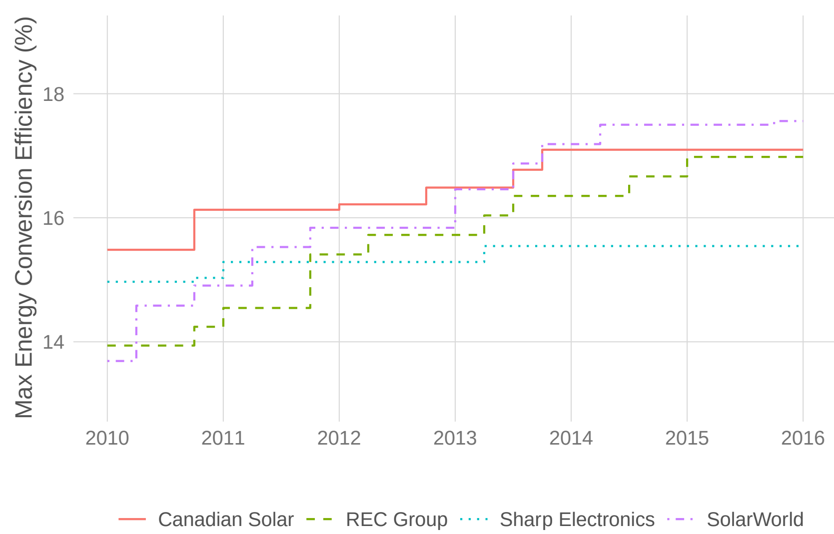

Notes: This plot shows the firm-level progression of energy conversion efficiency for a sample of firms based on data from the United States solar market. Energy conversion efficiency is the fraction of incoming solar power that a solar panel converts into electrical power. This figure shows that while firms throughout the industry have increased the energy conversion efficiency of their solar panels, there is variation in firm-level energy conversion 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).

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.[11]

Despite continual improvements to the production process and to energy conversion 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.[12]

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 market in each discrete time period (quarter). There are M regional markets. 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, energy conversion efficiency: $s_{it} = \eta_{it}$.[13] Firms share the industry state, s_t, which is comprised of:

- the distribution of firms' energy conversion efficiencies, $\eta_t = \begin{bmatrix} \eta_{1t} & \eta_{2t} & \dots & \eta_{It} \end{bmatrix}$;
- a common input price, w_t; and
- demand in each market, $d_t = \begin{bmatrix} d_{1t} & d_{2t} & \dots & d_{Mt} \end{bmatrix}$.

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

[11]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.

[12]Figure A.1 presents a side-by-side comparison of two solar panels for context.

[13]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 energy conversion efficiency (η_{it}) are equivalent and used interchangeably because firms are only differentiated on one dimension (energy conversion efficiency).

lower their future costs. When making this decision, each firm takes expectations 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 for Solar Panels

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).[14] I refer to this final good as "solar panels" throughout the paper.

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 level of demand in each market, d_{mt}, depends upon subsidies to consumers, S_{mt}, and follows an exogenous first-order Markov process. I use "demand state" to refer to d_{mt} throughout the paper because it is a state variable of the firms' dynamic game.[15]

The static demand specification implies that consumers are not forward-looking. In

[14] 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 A.1 compares two example solar panels to provide some context for these assumptions.

[15] An alternative approach would be to model demand as a function of prices net of subsidies, rather than gross prices, so that changes in government subsidies would correspond to movements along the demand curve rather than shifts in the level of the demand curve. The approach taken here has empirical advantages given that the primary subsidies under study take the form of payments for electricity generation rather than reductions in upfront price (see Section 5). However, it has the conceptual disadvantage that it treats upfront prices and payments for electricity generation differently, and therefore it is not micro-founded.

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 recent evidence from the solar market. De Groote and Verboven (2019) estimate demand for solar systems (not solar panels) using both static and dynamic specifications. The estimated price coefficient from their static demand specification is roughly 30% smaller in absolute value than under dynamic demand. Similarly, Feger et al. (2021) estimate a dynamic model of solar system adoption and find that the installation cost sensitivity is about 30% smaller for a static specification than for their dynamic specification. Langer and Lemoine (2018) also estimate a dynamic model of solar system demand, and they evaluate the tradeoffs policymakers face between intertemporally price discriminating and taking advantage of technological progress.

In contrast to these recent papers that study how solar subsidies affect demand, my focus is on understanding how subsidies affect both demand and supply. Despite its potential shortcomings, the static model of demand I employ facilitates estimation of a dynamic model of supply, which is the focus and primary contribution of this paper.[16] 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 Figures 1b, A.2a, and A.2d for examples). Second, anecdotes from government and industry publications suggest that ongoing price reductions were not fully anticipated, even by industry insiders.[17] Finally, Gillingham and

[16] It is computationally demanding to jointly estimate and solve 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.

[17] 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

Tsvetanov (2019) present data on purchases and a consumer survey in the Connecticut residential solar market that shows limited evidence of forward-looking behavior.

3.2 Firm Cost Structure

Firms have constant marginal costs of production, $mc\,(\eta_{it}, w_t)$, that depend on firm-specific energy conversion efficiencies (η_{it}) and a common input price (w_t). Both energy conversion 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,

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

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 $\bar{\pi}_i(s_t)$.[18]

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

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

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

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.

[18] 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 $\bar{\pi}(\eta_{it}, s_t)$, since $s_{it} \equiv \eta_{it}$.

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.[19] The structural interpretation of these shocks is a random shock to the fixed cost of investment.

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 energy conversion efficiency is one-to-one in order to infer the unobserved investment decision. If a firm does not invest ($x_{it} = 0$), its energy conversion efficiency does not change. If a firm does invest ($x_{it} = 1$), the change in its energy conversion efficiency is ν_{it}, which is i.i.d. across firms and time with support $\nu_{it} \in (0, \infty)$. To summarize, energy conversion efficiency evolves according to

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

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 always yields a non-zero improvement in energy conversion efficiency. The stochastic nature of the investment outcome captures the uncertainty inherent in adapting R&D advances from laboratory activities to large-scale production.

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

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

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. Each firm's strategy, denoted $\zeta_i(s_t, \varepsilon_{it})$, 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,

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

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 and Descriptive Statistics

4.1 Data

Estimating the model outlined in Section 3 requires data on: the quantity and price of solar panels sold by each manufacturer; technological innovation, as measured by energy conversion efficiency; and government policies that encourage consumers to adopt solar panels. This section summarizes the sources for each of these data in turn.

First, I employ data on solar panel market outcomes from IHS Markit. IHS Markit's PV Module Intelligence Service provides data on the solar panel supply chain on a quarterly basis from January 2010 through December 2015. These data include firms' sales by region and production by country. I use the sales data to construct a dataset at the firm-market-quarter level that includes quantities sold 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. The analysis centers on four regional markets: Germany, Japan, the United States, and the Rest of the World. I complement these sales data with data on firms that vary over time but not across markets, including each firm's production capacity and actual production on a quarterly basis. For the purpose of demand estimation, I also construct data on total quantity and price in the Rest of the World's constituent submarkets: Australia, China, India, Italy, and Other (all other countries). Quarterly data on the global spot price of polysilicon, the primary input to solar panel production, also come from IHS Markit.

Second, I quantify technological innovation using firm-level energy conversion efficiencies from the public version of the Lawrence Berkeley National Laboratory's Tracking the Sun dataset.[20] The dataset contains the characteristics of installed solar systems throughout the United States. These characteristics include the manufacturer, model, and energy conversion efficiency of the solar panels utilized in each system. Energy conversion efficiency is the fraction of incoming solar power that a solar panel converts into electrical power. It is a common, verifiable measure reported by firms on their product specification sheets. 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 energy conversion efficiency, tracking the maximum energy conversion efficiency sold by each firm over time. This is the observed measure of energy conversion efficiency I refer to throughout the paper.

Third, data on market-specific subsidies to consumers that vary over time come from national governments via the International Energy Agency's Photovoltaic Power Systems Programme, and from other publicly available sources. As described in Section 2, the dominant form of government subsidies during the sample period were subsidies for solar electricity generation called feed-in tariffs. For Germany and Japan, I collect feed-in tariff

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

levels in local currency units per kilowatt-hour (kWh) of electricity generation. These are represented graphically in Figure 1. For the United States, which subsidizes investment in solar systems rather than solar electricity generation, I use the ITC and the tax advantage of accelerated depreciation. I use a subsidy of 40% of the solar panel purchase price to capture these two federal solar subsidies in the United States.[21] Finally, for the residual market, I collect data on feed-in tariffs in its constituent submarkets: Australia, China, India, and Italy. These feed-in tariffs are all in local currency units per kWh of electricity generation. For Other, which aggregates demand in all other markets, there is no clearly defined subsidy; thus, I do not collect data on subsidies in other countries.[22] Subsidies for all markets except Germany and Japan are visualized in Figure A.2.

Finally, I collect ancillary data on input prices for use as instrumental variables. The global price of silver is from the London Fix, retrieved from The Silver Institute. The global price of aluminum is from the International Monetary Fund's Primary Commodity Prices, retrieved from FRED. Figure A.3 plots these time series along with the price of polysilicon and the weighted average solar panel price. I describe the use of these ancillary data in more detail in Section 5.

4.2 Descriptive Statistics

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

[21]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. The use of aggregate national data on solar panel sales makes it difficult to separately identify the impacts of state or local policies such as rebates and the implicit subsidy of net metering.

[22]Solar panels in "Other" account for less than 25% of the global quantity sold during the sample period.

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.[23] 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 90th 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.

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). 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

[23]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 analysis. 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 A.1 provides additional summary statistics on the exposure of each firm to each market over the sample period.

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

Demand for Solar Panels I assume that demand for solar panels is log-linear, which is a common approach for estimating parameters that govern the total quantity of a good demanded.[24] Thus equation 1 can be written as

$$Q_{mt} = D_{mt}(S_{mt}) \cdot P_{mt}^{\alpha_{pm}},$$

where $d_{mt} = \ln D_{mt}(S_{mt})$. The empirical analog of this demand curve is

$$\ln Q_{mt} = \alpha_{0m} + \alpha_{sm} f(S_{mt}) + \alpha_{pm} \ln(P_{mt}) + \varepsilon_{mt}^{D}, \tag{2}$$

where Q_{mt} is the total wattage of solar panels purchased in market m and quarter t, S_{mt} is the level of the subsidy available to consumers of new solar systems, and P_{mt} is the price for solar panels.

As discussed in Section 2, feed-in tariffs are the most common type of subsidy in the solar market during the study period. Feed-in tariffs provide a fixed price for solar electricity that is determined at the time of investment. This provides an indirect subsidy to the adoption of solar systems. Given this, and the fact that demand for solar *panels* is derived from demand for solar *systems*, I model the effect of feed-in tariffs on the market for solar panels as a demand shifter rather than a change in the price level.[25]

I use data on feed-in tariff levels described in Section 4 for the subsidy variable, S_{mt},

[24] Prior influential papers that use this assumption include Porter (1983), Ryan (2012), Roberts and Schlenker (2013), and Kalouptsidi (2014). Another example is Hausman et al. (1994), which includes a multi-level demand system that accommodates differentiated products within a market, but in which the *overall* level of demand in the market is specified as log-linear.

[25] This is consistent with the fact that feed-in tariffs do not drive a wedge between the prices consumers pay and the prices producers receive for solar panels, in contrast to subsidies that lower the upfront cost of adopting solar technology. An alternative approach would be to use information on feed-in tariff levels along with assumptions about electricity generation over time, consumer time preferences, and the cost of solar panels as a share of solar systems in order to estimate the net present value of feed-in tariff revenues per Watt of solar panel capacity. This could be used to construct a consumer price that is net of feed-in tariff revenues, which could replace the gross price in equation 2. This would reduce the number of parameters in the model and would utilize variation in feed-in tariffs, as well as prices, to estimate the price elasticity of demand. In contrast, the approach taken here is more flexible in that it does not require arbitrary technical and behavioral assumptions.

for the markets of Germany, Japan, and Rest of the World. Constructing a single measure of the feed-in tariff level for the residual market is challenging because it depends on feed-in tariffs in many individual countries, and the functional form of that dependence is unknown. To circumvent this challenge, I incorporate additional data on total quantities, prices, and feed-in tariffs for the most important submarkets in the residual market: Australia, China, India, Italy, and Other (all other countries).[26] When estimating equation 2, I allow for submarket-specific intercepts (α_{0m}) and subsidy parameters (α_{sm}), but I restrict the price elasticity of demand (α_{pm}) to be common across the residual market's submarkets. This restriction ensures an internally-consistent approach to estimating production costs and aggregating residual market demand for use in estimating the dynamic parameters. All feed-in tariffs are in local currency units per kWh of electricity generation.[27]

The second type of subsidy described in Section 2 are policies that lower the upfront cost of adopting a solar system. This is the primary approach used in the United States. The ITC provided a tax credit worth 30% of investment costs during the sample period. Because the level of the ITC was constant over time, it does not provide any variation to identify the subsidy's effect on demand. For this reason, I model the United States subsidy as a fraction of the solar panel price rather than as a separate additive term. This captures the dependence of the subsidy's value on the cost of solar systems, rather than on their electricity output as in the case of feed-in tariffs. This approach uses variation in solar panel prices to estimate both the price elasticity of demand and the impact of subsidies on demand. This is valid as long as consumers respond in the same way to changes in the net price of solar panels regardless of whether they are caused by changes in the gross price or in the subsidy level. As described in Section 4, I use a time-invariant subsidy of 40% of the solar panel purchase price to capture the combined effects of the

[26] Australia, China, India, and Italy are the four largest submarkets within the residual market in terms of solar panel adoption over the sample period that are individually tracked by IHS Markit. Data on feed-in tariffs are available for each of these four countries. IHS Markit aggregates other countries into regional groups for which there is no clearly-defined feed-in tariff. For this reason, I combine all other countries into a single market, Other, and do not model the effects of subsidies for that market.

[27] The choice of currency and energy units for the feed-in tariffs are not important, as their effects are scaled by estimable parameters that are not of direct interest. The product of each feed-in tariff and its coefficient gives the impact of that subsidy on that market's demand.

ITC and accelerated depreciation. Figures 1 and A.2 plot the total quantity, price, and subsidy level over time for all eight markets used to estimate equation 2.

Equation 2 allows for alternative parameterizations of the relationship between quantity demanded and feed-in tariff levels. A natural choice would be $f(S_{mt}) = \ln(S_{mt})$. A practical challenge to using this parameterization is that subsidies are not available in every time period and market. In Japan, for example, the feed-in tariff is zero until the middle of the sample period as shown in Figure 1a. Estimating this model would require dropping a non-trivial share of observations. This is undesirable for two reasons. First, it would sacrifice useful identifying variation. Second, it would prevent estimation of the demand state, d_{mt}, for the affected markets and time periods. This would severely hamper estimation of the state transition processes and dynamic parameters.

Instead, I approximate natural logarithms by using the inverse hyperbolic sine function: $f(S_{mt}) = \text{arcsinh}(S_{mt})$. This transform allows retention of zero-valued observations. To provide confidence that the inverse hyperbolic sine function is a reasonable approximation, I also estimate models using $\ln(S_{mt})$ on the subsample for which subsidies are non-zero.

I instrument for price to account for potential endogeneity using two sets of instruments. In the first, I use the prices of three inputs: polysilicon, silver, and aluminum. Polysilicon and silver are the first and second largest materials cost shares over the sample period.[28] These instruments vary over time but not across markets because polysilicon, silver, and aluminum are traded globally. Each of these input prices is correlated with the weighted average solar panel price in the time series, as can be seen in Figure A.3. The untestable identifying assumption is that these instruments satisfy the exclusion restriction: that increases in input prices affect equilibrium quantities only through their effects on solar panel prices. This is the econometric rationale for using a cost shifter to instrument for price when estimating demand, and this is a plausible assumption in this industry. The primary threat to instrument validity is reverse causality: if a regional market demand shock for solar panels were large enough to increase global prices for these inputs, the input prices would no longer be a valid instrument for price *in that regional*

[28]Polysilicon 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. Aluminum is used for pastes and mounting frames. See Section 2 for a more detailed description of the manufacturing process.

market. This scenario is unlikely because each market is a relatively small fraction of global demand for solar panels, and the solar market is only a fraction of global demand for these inputs.[29] However, if some or all of the regional markets were to experience correlated demand shocks, this could invalidate the vector of input price instruments. Since this assumption is fundamentally untestable, I also consider a variant on these instruments in which I only use the prices of silver and aluminum. This approach helps to mitigate possible concerns about reverse causality for the price of polysilicon, which is most likely to be affected by solar demand shocks.

Furthermore, 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. This is similar in spirit to the input price instruments, in that prices in other markets should reflect common cost shifters. In contrast to the use of input prices, constructing instruments using product market prices has the potential to capture common cost shifters that are not directly observable. 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 could be violated by demand shocks that affect multiple markets if they are not captured by the market-specific subsidies to consumers included as regressors.

Finally, the demand states of the model, d_{mt}, are the demand curve intercepts,

$$d_{mt} = \alpha_{0m} + \alpha_{sm} \operatorname{arcsinh}(S_{mt}) + \varepsilon^{D}_{mt}. \tag{3}$$

They 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 for Germany, Japan, and the U.S. For the Rest of the World, I compute d_{mt} by aggregating the demand states of its five constituent submarkets.

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

[29]The fraction of overall demand for each input that comes from the solar industry varies across the inputs. For polysilicon, demand from solar manufacturers grew over time to dominate global polysilicon consumption during the sample period. For silver, the solar industry consumed approximately 10% of annual silver production (Powell et al., 2012). For aluminum, the fraction is even smaller.

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

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

where α_{pm} is the price elasticity of demand.[30]

I parameterize these inferred costs to quantify the impact of energy conversion efficiency improvements on production costs. I use a parametric form that is motivated by the economics of the industry. The measure of energy conversion 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:

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

where $\tilde{\eta}_{it}$ is observed firm-specific energy conversion efficiency, w_t is the common input price, β_0 is an unobserved time-invariant common scale parameter, and ε_{imt}^S is an unobserved shock at the firm-market-time level. The economic model formulated in Section 3 is in terms of cost-indexed energy conversion efficiency, η_{it}. There is a one-to-one mapping between observed energy conversion efficiency and cost-indexed energy conversion efficiency: $\eta_{it} = \tilde{\eta}_{it}^{\beta_1}$. The econometric model is used to estimate the parameter

[30] As described above, the subsidy is modeled as a constant fraction of the solar panel price for the United States. Nevertheless, this condition still holds: the effect of the United States subsidy is equivalent to an intercept shift for the demand curve in equation 2 because the solar panel price enters logarithmically.

β_1 governing this relationship.[31,32] For w_t, I use the observed price of polysilicon, which is the primary input to solar panel production. In some specifications, I use time fixed effects to estimate the common effects of input prices more flexibly. These fixed effects absorb the effects of polysilicon prices as well as other cost shifters that vary over time but not across firms. I estimate the cost function in logs,

$$\ln(mc_{imt}) = \beta_0 + \beta_1 \ln(\tilde{\eta}_{it}) + \beta_2 \ln(w_t) + \varepsilon^S_{imt}. \tag{4}$$

I include firm fixed effects in some specifications to capture unobserved factors that shift costs independently of variation in energy conversion efficiencies and input prices. I use ordinary least squares to estimate equation 4 under the assumption that ε^S_{imt} is i.i.d. over firms, markets, and time.[33] This model quantifies the relationship between energy conversion efficiency and marginal cost. The sign and magnitude of β_1 dictate whether and how much energy conversion efficiency improvements lower costs, and therefore the incentive firms face to innovate in response to changes in demand.

[31]Observed energy conversion efficiency, $\tilde{\eta}_{it}$, is the fraction of incoming solar power that a solar panel converts into electrical power, while the energy conversion efficiency measure used in the economic model, η_{it}, is in terms of production costs. The latter measure is used for convenience: I leverage cost-indexed energy conversion 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 energy conversion efficiency, which is necessary for firms to make investment decisions in the real world.

[32]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 energy conversion 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 – energy conversion efficiency in the U.S.

[33]Energy conversion 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.

I use forward simulation to construct the endogenous distribution of energy conversion 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' energy conversion 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, denoted x_{it}, by estimating a logit model via penalized maximum likelihood:

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

where $f(\eta_{it}, s_t)$ is a flexible function of the firm's technical efficiency, η_{it}, and the industry state, s_t.[34] In particular, $f(\eta_{it}, s_t)$ includes cubic polynomials in the following variables: the firm's energy conversion efficiency, the mean and standard deviation of the industry

[34] See Section 3 for definitions of each term.

distribution of energy conversion efficiency, the input price, and the four demand states; as well as first-order interactions between the firm's energy conversion efficiency and all other variables listed previously. The tuning parameter λ is selected by leave-one-out cross-validation.[35] I then model the investment choice using a logit model and estimate the parameters via maximum likelihood,

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

where $\tilde{f}(\eta_{it}, s_t)$ 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 E. 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

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

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

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

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

where $\theta = \left[\begin{array}{ccc}1 & \gamma & \sigma\end{array}\right]'$.[36] $V_i(s_t;\zeta)$ 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 $W_i(s_t;\zeta)$.[37]

[36]The parameter σ is a scale parameter on the private choice-specific shocks. It is normalized rather than estimated due to empirical challenges to separately identifying γ and σ. See Appendix F.

[37]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,

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

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 energy conversion 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 to generate the discounted sum of product market profits, investment costs, and random shocks, up to the parameters θ.[38] I simulate 500 of these industry paths from each initial state and take the mean to approximate $W_i(s_t;\zeta)$. I base β, the quarterly discount factor, on an annual discount factor of 0.875, similar to discount factors used in prior work.[39]

To estimate θ, I use $\widehat{W}_i(s_t;\zeta)$ to implement a pseudo maximum likelihood estimator based on the firm's optimal investment decision. Appendix F provides more details.

6 Estimation Results

6.1 First Step Estimates

Demand for Solar Panels Table 2 presents estimates of the price elasticities of demand ($\hat{\alpha}_{pm}$ from equation 2). The subsidy coefficients ($\hat{\alpha}_{sm}$) are presented in Table B.1. Each column of Tables 2 and B.1 presents estimates from a different estimation strategy, and each row corresponds to a different market. There are two main takeaways from Table 2. First, the estimated price elasticities across all estimation strategies and markets range between -1 and -4.4, with most coefficients falling between -1.3 and -2.5.[40] Second,

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.

[38] 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 previous periods.

[39] It is standard to assume a discount factor when estimating dynamic models because it is difficult to separately identify the discount factor from other model parameters in practice (see, e.g., Rust, 1987). To give two examples from the empirical literature on dynamic games, Ryan (2012) assumes an annual discount factor of 0.9 while Igami (2017) uses 0.8. To provide a measure of how sensitive the estimates are to this necessary assumption, I present investment cost estimates using alternative discount factors in Table F.1.

[40] These estimates lie within the wide range of previous estimates in the literature. Gillingham and Tsvetanov (2019) estimate a static demand elasticity of -0.65. Coefficient estimates from De Groote and

the ordinary least squares and instrumental variables estimates are not statistically distinguishable, and the point estimates are quite similar in most cases.[41]

Table 2: Estimated Demand Elasticities

	OLS	IV: Input Prices	IV: Metal Prices	IV: Other Prices
	(1)	(2)	(3)	(4)
Germany	-2.53^{***}	-2.65^{***}	-4.39^{***}	-2.47^{***}
	(0.50)	(0.56)	(1.13)	(0.51)
Japan	-1.46^{***}	-1.09^{***}	-1.07^{***}	-1.26^{***}
	(0.31)	(0.29)	(0.32)	(0.27)
ROW	-1.34^{***}	-1.30^{***}	-1.29^{***}	-1.36^{***}
	(0.24)	(0.24)	(0.22)	(0.24)
USA	-1.49^{***}	-1.39^{***}	-1.48^{***}	-1.48^{***}
	(0.26)	(0.25)	(0.32)	(0.26)
Region FE	X	X	X	X
asinh(Subsidy)	X	X	X	X
Min. F-stat		18.96	5.98	53.6
Within R^2	0.7	–	–	–
Observations	192	192	192	192
Note:				*p<0.1; **p<0.05; ***p<0.01

Notes: This table presents estimated price elasticities of demand (i.e., $\hat{\alpha}_{pm}$ from equation 2). Each row corresponds to a different market (m) and each column presents estimates from a different estimation strategy. The first column presents estimates of equation 2 using ordinary least squares. The second column presents IV estimates that use the prices of polysilicon, silver, and aluminum as instruments for the price of solar panels. The third column uses only the prices of silver and aluminum as instruments for the price of solar panels. 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. Min. F-stat is the minimum F-statistic across market-specific tests of the hypothesis that the excluded instruments are jointly irrelevant. Heteroskedasticity-consistent standard errors are in parentheses.

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

To assess the performance of the inverse hyperbolic sine transform, I also estimate equation 2 with subsidies in logs on the subsample with no zeros. The resulting price elasticity estimates are presented in Table B.2. In addition, Table B.3 presents estimates using $\ln(1 + S_{mt})$ to preserve the full sample. Both sets of estimates are similar to those

Verboven (2019) imply an upper bound on the static elasticity of roughly -6.6. These papers estimate demand for residential solar systems, whereas I estimate demand for solar panels.

[41]The results in the third column should be viewed with caution, as the minimum F-statistic is less than 10, indicating the instruments are weak for one or more regional markets.

found in Table 2.

The specification that uses prices from other markets as instruments for price in a given market serves as the baseline specification used to estimate the supply model. The subsidy-inclusive demand states ($\hat{d}_{mt}$) are derived using the estimated parameters from equation 2 and the definition of the demand states in equation 3. Counterfactual demands had the subsidies not been in place are recovered by subtracting the estimated impact of subsidies on demand ($\hat{d}_{mt} - \hat{\alpha}_{sm}\text{arcsinh}\left(S_{mt}\right)$). The resulting demand states with and without subsidies are plotted in Figure A.4.[42] Both are in terms of the natural logarithm of quantity, as they represent the demand curve intercept from the log-linear estimating equation.

Production Costs Graphical evidence suggests that energy conversion efficiency reduces costs: Figure 4a shows that there is a positive cross-sectional relationship between energy conversion efficiency and market share in the raw data. This positive correlation is consistent with higher energy conversion efficiency lowering costs for firms, which would lead to larger market share under a model of Cournot competition.

Figure 4b presents an analogous plot of the cross-sectional relationship between energy conversion efficiency and production cost. Costs are recovered from the demand estimates and model of competition as described in Section 5. Higher energy conversion efficiency is associated with lower cost in the cross-section, after eliminating time series variation that could induce spurious correlation and undermine identification.[43]

Table 3 presents the coefficient on energy conversion efficiency across several alternative specifications of equation 4. The negative relationship between cost and energy conversion 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 energy conversion efficiency and a constant as regressors

[42]Figure B.1 plots a disaggregated version with each market used for demand estimation before aggregation of the ROW submarkets Australia, China, India, Italy, and Other into a single ROW market for supply estimation.

[43]Appendix D presents graphical summaries of firms' market shares and estimated marginal costs over time. It also contains discussion and descriptive analysis regarding variation in these outcomes across markets for a given firm and time period.

Figure 4: Cross-Sectional Relationship between Energy Conversion Efficiency and:

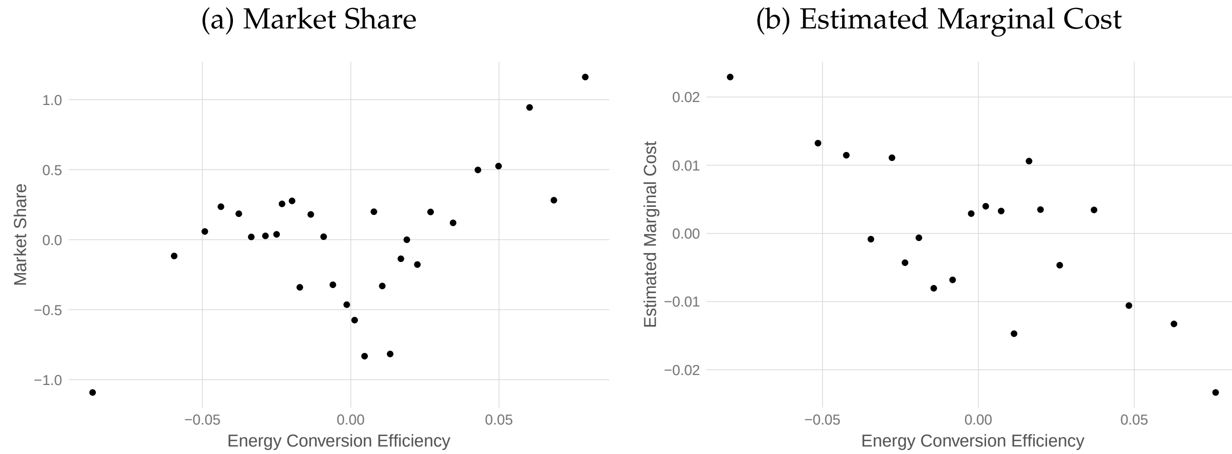

Notes: Figures are binned scatterplots of residual variation in market share (a) and cost (b) versus energy conversion efficiency within each market and time period. Each point is a local mean of underlying data. Figure a is constructed using data without any economic assumptions. Marginal costs in Figure b are inferred from firms' first order conditions for optimal production using data and estimated demand parameters. Plots are consistent with higher energy conversion efficiency lowering production costs.
Data Source: Author's calculations using data described in Section 4 and model described in Section 5.

and produces a large negative coefficient. The second specification, which includes the common, time-varying price of polysilicon (w_t), produces a significantly attenuated coefficient on energy conversion efficiency. This highlights 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 polysilicon price is not included in this regression because it only varies over time and so its effect is no longer identified. This further attenuates the coefficient on energy conversion efficiency. The final columns 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 energy conversion efficiency. These specifications generate slightly more negative coefficients on energy conversion 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 energy conversion efficiency is higher. The model estimates in columns 2 through 5 suggest that improvements to energy conversion efficiency constituted 3% to 32% of total cost reductions over the period 2010-2015.

Table 3: Relationship between Marginal Cost and Energy Conversion Efficiency

	(1)	(2)	(3)	(4)	(5)
$\ln(\tilde{\eta}_{it})$	-5.43^{***}	-0.91^{***}	-0.17^{*}	-1.81^{***}	-0.32^{**}
	(0.16)	(0.10)	(0.09)	(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 energy conversion efficiency ($\tilde{\eta}_{it}$). 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 energy conversion 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 other unobservable cost shifters (captured by time period fixed effects).

Data Source: Author's calculations using data described in Section 4 and model described in Section 5.

The qualitative results in Table 3 are robust to a range of alternative specifications summarized in Appendix C. 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 C.1a). 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 C.1b and C.1c). 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 energy conversion 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 – energy conversion efficiency in the U.S. market in Table C.1d. 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 firms' manufacturing capacity may be constrained in the short run. If a firm's production were capacity-constrained, its quantity choice would also be constrained. This could lead to a violation of the first order condition used for estimation of the firm's production cost, resulting in biased estimates. Furthermore, because firms' quantity choices are interdependent, this could bias production cost estimates for other firms as well. To address this possibility, I first quantify the prevalence of capacity constraints. Figure C.1 plots module capacity utilization over time for the full sample. In aggregate, the industry does not appear to suffer from severe capacity constraints. At the firm level, utilization of all capacity is rare but does occur: firms produce at their capacity in 2.2% of observations. Given this, I restrict attention to the time periods in which no firms were operating at full capacity and re-estimate equation 4. The results are presented in Table C.1e. This restricted sample is significantly smaller than the full sample, so the coefficients are less precise. However, the relationship between production cost and energy conversion efficiency is robust, suggesting that instances of capacity constraints do not induce a significant bias in the results.[44]

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,

[44]In principle, the model developed in Section 3 could be extended to account for capacity constraints in the short run and firms' investment in manufacturing capacity in the long run. However, since the production cost estimates are robust to omitting periods in which capacity constraints bind, and full capacity utilization is infrequent, the model abstracts from this aspect of firms' decision-making.

that would significantly complicate the analysis. This estimate identifies the causal impact of energy conversion efficiency on the cost of production under the assumption that there are no omitted factors that are correlated with both energy conversion 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 A.2 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 energy conversion efficiency (η_{it}), and the interaction of the firm's energy conversion efficiency and the industry average energy conversion efficiency ($\eta_{it}\bar{\eta}_t$). Table A.3 presents estimates of the second-stage logit model using these covariates. The distribution of investment levels is fit using a gamma distribution as discussed in Section 5.

6.2 Second Step Estimates

Investment cost estimates are presented in Table 4. The fixed cost point estimate is \$108.6 million.[45] In the data, investments occur slightly less often than once per year, so these estimates imply annual investment costs directed toward improving energy conversion 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.

[45] Table F.1 presents investment cost estimates using alternative discount factors. As the discount factor increases, the fixed cost of investing in energy conversion efficiency improvements also increases. This is because a higher discount factor puts more weight on future profits that result from technological advancement, and so a higher fixed cost is necessary to rationalize observed innovation patterns.

Table 4: Investment Cost Parameter Estimates

Parameter	Point Estimate	Confidence Interval
γ	108.6	(55.5, 201.5)
σ	30.0	–

Notes: This table presents estimates of the fixed cost of energy conversion efficiency improvements (γ). The scale parameter on the private shocks (σ) is held fixed rather than estimated due to the practical challenge of separately identifying γ and σ. The value $\sigma = 30$ is based on the central estimate derived from the original sample. Confidence intervals for γ 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 using data described in Section 4 and model described in Section 5.

7 Counterfactual Analysis

7.1 Solving the Model

This subsection first describes how the product market equilibrium is computed, and then summarizes the process of solving for firms' innovation behavior and simulating the model forward over time. Appendix G provides a more detailed description.

Solving for the Product Market Equilibrium Since the game in each product market and time period only depends on the current state, it can be solved independently from, or in concert with, the broader dynamic model. This requires estimates of: demand states $\hat{d}_{mt}$, the price elasticity of demand $\hat{\alpha}_{pm}$, the sum of firms' costs $\sum_i \widehat{mc}_{it}$, and the number of firms I. Given these, equilibrium price and quantity predictions are computed as follows:

$$\widehat{P}_{mt} = \frac{\sum_i \widehat{mc}_{it}}{I + \frac{1}{\hat{\alpha}_{pm}}} \qquad\qquad \widehat{Q}_{mt} = \exp(\hat{d}_{mt})\widehat{P}_{mt}^{\hat{\alpha}_{pm}}.$$

Solving for Firms' Innovation Decisions To assess the long-run impacts of consumer subsidies, accounting for both demand and supply responses, I use the estimates to solve the dynamic model for different subsidy scenarios. I adapt the methods of Sweeting (2013) to this context, and solve the model using parametric policy iteration (Benitez-Silva et al., 2000). This procedure is similar to standard policy iteration methods in the literature, except that it relies on the use of value function approximation due to the continuous nature of the state space.

The model is initialized at the set of observed states, with an initial guess for the vector of conditional choice probabilities in each of those states, P^1. Then, each iteration j consists of the following steps:

1. Compute *ex-ante* expected flow profits, given choice probabilities P^j.
2. Simulate the states that can be reached in one period from each observed state by taking 1000 draws from the state transition processes, given choice probabilities P^j.
3. Compute $\widehat{\lambda}^{P^j}$ (the parameters of the value function approximation).
4. Use $\widehat{\lambda}^{P^j}$ to compute choice-specific value functions for each choice (invest or not) for each firm, and then use these to compute new choice probabilities P^{j+1}.
5. If $\max|P^{j+1} - P^j| < 10^{-4}$, the iterations stop. If not, iteration $j+1$ begins at step 1.

This procedure produces conditional choice probabilities in each state, P^*, which describe investment behavior.

Simulating the Model Forward over Time To perform a comprehensive counterfactual analysis, I start by solving the model for the observed states as described above. Then, to simulate the model forward over time, I start from the first quarter of 2010 and compute the expectation of the next state using the conditional choice probabilities P^* and parameters of the AR(1) process for d_t and w_t. This simulated state is used to solve the model again. I repeat this procedure 80 times to simulate the model forward 20 years.

Summarizing the Counterfactual Results Simulating the model forward over time results in a trajectory of states over time for each counterfactual scenario.[46] These states can can be used to compute equilibrium quantities and prices as described above, as well as firms' profits and consumer surplus. Finally, I use the quantities to perform a back-of-the-envelope calculation to quantify the change in external benefits attributable to the subsidies. I base this on existing estimates of the external damages from natural gas

[46]Figure G.1 summarizes the evolution of energy conversion efficiencies over time for the three scenarios.

electricity generation from Muller et al. (2011). Details of the external benefits calculations are provided in Appendix G.

Implementing the Counterfactual Scenarios I solve the full model three times. First, for the *Baseline* scenario, I solve the model using estimates of the subsidy-inclusive demand states, $\hat{d}_{mt}$, and the parameters that govern their transition processes. Second, for the *No Subsidies - Dynamic* scenario, I solve the model using the counterfactual demand states, $\hat{d}_{mt} - \hat{\alpha}_{sm}\text{arcsinh}\left(S_{mt}\right)$, and the parameters that govern their counterfactual transition processes. Finally, I construct a *No German Subsidies - Dynamic* scenario that only removes subsidies in Germany to highlight how domestic policy can produce international spillovers through the innovation responses of multinational firms.

To understand the static and dynamic impacts of consumer subsidies in the market for solar panels, I construct a fourth scenario by combining the *Baseline* simulated costs with the simulated demands from the *No Subsidies* scenario. This does not require solving the full model again, as product market outcomes in each market and time period can be computed directly using these costs and demands. This produces estimates of counterfactual outcomes without consumers subsidies while holding cost reductions fixed. In other words, it treats the endogenous innovation that occurs with subsidies in place as exogenous. I refer to this scenario as *No Subsidies - Static*.

7.2 Counterfactual Results

Solar Panel Adoption and Prices over Time Figure 5 summarizes the effects of subsidies on solar adoption and prices over time. In the *Baseline* scenario, the total quantity of solar panels increases significantly over time as the price falls. This is similar to the patterns in the raw data.[47]

In the *No Subsidies - Static* scenario, the quantity of solar panels adopted increases over time, but at a much lower rate. The prices are essentially unchanged, because this scenario holds firms' production costs fixed at their *Baseline* levels. The dark shaded area in

[47] For realism, I use estimated demand states for the three scenarios to construct Figure 5 rather than using the simulated demand states from solving the full model. The results using the simulated demand states are very similar. Appendix G provides a detailed description of how Figure 5 was constructed.

Figure 5a is the estimated solar panel adoption attributable to consumer subsidies through their effects on demand. Absent subsidies, the quantity of solar panels sold (in Watts) over the period 2010-2015 would have been 51% lower. German subsidies contributed 24% of that difference, with Japanese subsidies contributing 24%, U.S. subsidies contributing 18%, and subsidies in the residual market contributing the remaining 34%.

Figure 5: Counterfactual: Impact of Subsidies on Solar Adoption and Prices

(a) Quantities

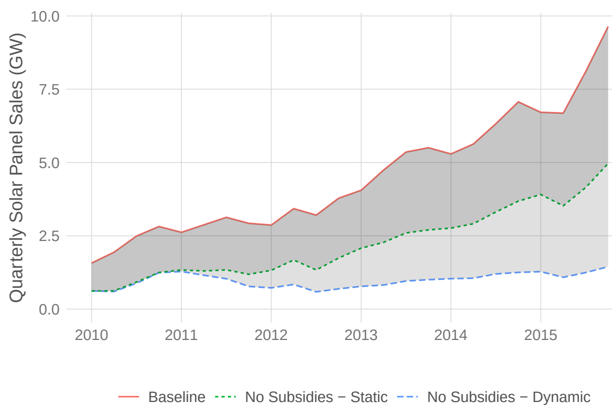

(b) Prices

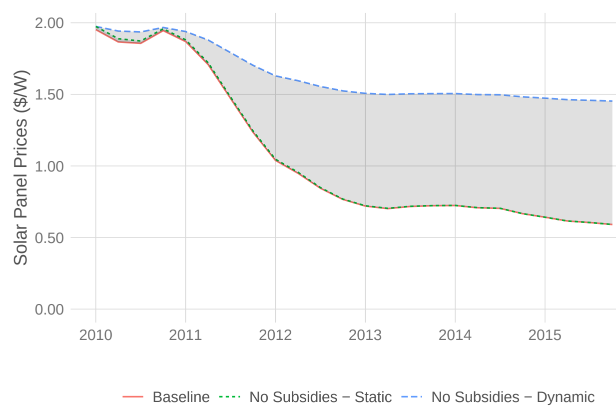

Notes: This figure plots model-predicted quantities (a) and prices (b) over time with and without subsidies. The quantities are the global total. The prices are a global average, weighted by the quantities in each market. There are three lines in each plot: *Baseline* represents model predictions based on historical subsidies and production costs. *No Subsidies - Static* represents counterfactual outcomes after removing subsidies but holding production costs fixed. *No Subsidies - Dynamic* represents counterfactual outcomes after removing subsidies and allowing for induced innovation by firms. This estimate extrapolates the energy conversion efficiency improvements predicted by the model to apply to all production costs, under the assumption that subsidies induce the same proportional reductions in cost through efficiency as through other means.
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.

In the *No Subsidies - Dynamic* scenario, I plot the results of a calculation that extrapolates the model's predictions for energy conversion efficiency improvements to apply to other forms of cost reduction as well. The key assumption I make is that subsidies induce the same proportional reductions in cost through efficiency as through other means. The result is a more comprehensive, albeit somewhat speculative, estimate of the effects of subsidies on innovation by firms than an estimate that only allows for endogenous cost reductions through one margin (energy conversion efficiency).[48]

[48] If this assumption is wrong, the *No Subsidies - Dynamic* lines would be biased. For example, if other

The results of the *No Subsidies - Dynamic* scenario are striking. The time series of solar panel adoption is roughly flat after removing induced innovation. In total, the quantity of solar panels sold over the period 2010-2015 would have been 78% lower than under the *Baseline* scenario. This represents a 54% increase in the solar adoption attributable to subsidies relative to the *No Subsidies - Static* estimates that fail to account for induced innovation. This is because firms invest less in innovation when the demand they face is lower, and so their production costs decline less over time. Figure 5b shows that the resulting price reductions are only about one-third as large as in the *Baseline* simulation.

These results highlight the importance of accounting for firm responses in studying the impacts of subsidies on solar adoption. Even over a relatively short time period, induced innovation yields significant reductions in costs and increases in adoption. Furthermore, the innovation induced by subsidies over the sample period will continue to generate benefits over time, even if subsidies are altered or removed in the future.

Summary of the Economic Benefits of Solar Subsidies Table 5 provides a more comprehensive summary of the economic benefits of subsidies, over a longer period of time. All numbers are computed by solving the model over 80 quarters, from 2010 through 2029, computing outcomes in each time period, and then constructing the present discounted value in 2010 terms.

The first panel of Table 5 reports absolute changes in economic benefits relative to the *Baseline* scenario. Static results correspond to the difference between *No Subsidies - Static* and *Baseline*. Dynamic - Conservative results correspond to the difference between *No Subsidies - Dynamic* and *Baseline*. Dynamic - Preferred results follow the same approach that is used in Figure 5 to extrapolate energy conversion efficiency improvements to apply to all production costs. The second panel of Table 5 reports the difference between the results of each dynamic scenario and the static scenario in relative terms. The columns of Table 5 report changes in consumer surplus (CS), firms' product market profits ($\bar{\pi}$), and

cost reductions are less responsive to changes in demand than efficiency improvements, the counterfactual outcomes would lie in between the lines *No Subsidies - Static* and *No Subsidies - Dynamic*. If, on the other hand, other cost reductions are *more* responsive to changes in demand than efficiency improvements, the counterfactual outcomes would be even more extreme than *No Subsidies - Dynamic*. See Appendix G for the details of this calculation.

external environmental benefits (EB_{CO_2} and EB_{ALL}). The final two columns present the sum of consumer surplus, profits, and the two different measures of external benefits.

Table 5: Change in Benefits from Removing All Subsidies, Relative to Baseline Scenario

	CS	$\bar{\pi}$	EB_{CO_2}	EB_{ALL}	$Total_{CO_2}$	$Total_{ALL}$
Change Relative to Baseline Scenario						
Static ($B)	-25	-2.5	-13	-47	-40	-75
Dynamic - Conservative ($B)	-26	-2.5	-13	-49	-41	-77
Dynamic - Preferred ($B)	-69	-3.3	-24	-90	-96	-162
Percent Difference from Static						
Static (%)	–	–	–	–	–	–
Dynamic - Conservative (%)	4	1.5	3	3	4	4
Dynamic - Preferred (%)	174	32.9	90	90	139	116

Notes: This table presents counterfactual outcomes based on three sets of counterfactuals. The *Static* rows summarize the effects of removing subsidies, accounting only for their effects through demand responses. The *Dynamic - Conservative* rows account for both demand responses and endogenous changes in energy conversion efficiency improvements. The *Dynamic - Preferred* rows extrapolate the energy conversion efficiency improvements predicted by the model to apply to all production costs, under the assumption that subsidies induce the same proportional reductions in cost through efficiency as through other means. All values are in present discounted value terms, in 2010 dollars. Details of the calculations are in Appendix G.
Data Source: Author's calculations based on model simulation described in Section 7 using model estimates described in Section 6.

Three findings jump out from Table 5. First, within each scenario, consumer surplus and external benefits dominate the total benefits. Second, the external benefits are large, although the precise magnitudes are sensitive to assumptions. The estimated climate benefits from solar subsidies, EB_{CO_2}, range from $13 to $24 billion across scenarios. In the columns that also include local air pollution benefits, EB_{ALL}, the total external benefits range from $47 to $90 billion.[49]

Finally, and most importantly, accounting for endogenous innovation by firms yields significantly different estimates of the economic benefits of subsidies than accounting only

[49]The local air pollution benefit estimates are based on historical data from the United States, so they may not be representative of the benefits from avoided local air pollution in other regions and time periods. In addition to assumptions about the avoided damages from local and global air pollution, the magnitude of the external benefits are sensitive to the assumed lifetime of solar panels and the marginal sources of electricity that they displace. If solar panels last longer, the external benefits would be greater, and vice versa. Similarly, if solar panels displace electricity generation from coal rather than natural gas, the external benefits would be significantly larger. However, if they also displace some cleaner sources of electricity such as nuclear power, the benefits could be smaller. Appendix G contains more details on these calculations.

for demand responses. In total, the economic benefits increase 4% under the conservative scenario, and 116% to 139% under the preferred scenario. While the magnitude of the difference varies across outcomes and scenarios, the qualitative result is clear and consistent across cases.

International Spillovers through Innovation Together, the findings on induced innovation and the global nature of the solar market suggest that domestic policy could produce international spillovers through the innovation responses of multinational firms. To analyze this, I compare outcomes under the *No German Subsidies - Dynamic* scenario to the *Baseline* scenario. The *No German Subsidies - Dynamic* scenario removes German subsidies but keeps all other subsidies in place. The qualitative results from solving the model for this scenario are similar to the *No Subsidies - Dynamic* scenario: firms invest less frequently, leading production costs to be higher over time, than in the *Baseline* scenario. The quantitative results are smaller than the *No Subsidies - Dynamic* scenario, as expected.[50]

I use the results of this simulation to compute how this change in innovation affects solar panel adoption over time in each market. To isolate this international spillover through innovation from the direct effects of German subsidies on demand, I compute the difference between equilibrium quantities in each market using the production costs from the *Baseline* and *No German Subsidies - Dynamic* scenarios, but holding demand fixed at its estimated level with subsidies in place, $\hat{d}_{mt}$.[51] Then, in each period, I compute the share of this change in solar panel adoption that accrues to each market. Finally, I plot these market shares over time in Figure 6. Initially, the majority of this marginal solar adoption is captured within the German market. Over the course of the sample period, however, the vast majority of this marginal adoption shifts to other markets. In total over the period 2010-2015, the model predicts that 86% of the marginal solar adoption attributable to innovation induced by German subsidies occurs outside Germany.

[50]Figure G.1 shows that the energy conversion efficiency improvements implied by the model for the *No German Subsidies - Dynamic* scenario lie in between the *Baseline* and *No Subsidies - Dynamic* scenarios.

[51]As in Figure 5 , I use estimated demand states rather than using the simulated demand states from solving the full model for realism. The results using the simulated demand states are very similar.

Figure 6: Decomposition of Solar Adoption due to German-Induced Innovation

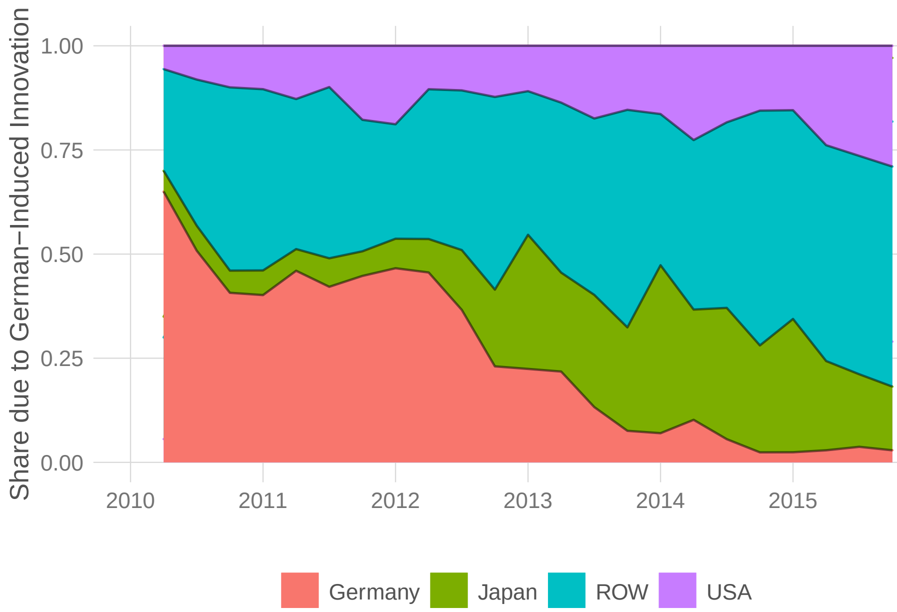

Notes: This figure plots the shares of model-predicted changes in quantities that accrue to each market over time due to innovation induced by German subsidies. The quantities are computed using production cost estimates from the *Baseline* and *No German Subsidies - Dynamic* scenarios. For realism, demand in each market is held fixed at its estimated value including subsidies.

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.

8 Conclusion

This paper studies innovation in solar energy technology, a key source of clean energy that has experienced rapid price declines over the past decade. I estimate a dynamic model of competition among solar panel manufacturers to understand the causes and effects of innovation. I find that subsidies to consumers have a significant impact on firms' incentives for innovation, and that induced innovation can generate cross-country spillovers due to the globally interconnected nature of the solar market. These results highlight the importance of accounting for firm responses in studying the economic benefits of subsidies on solar adoption.

While the insights of this paper extend to any change in the market environment that affects consumer demand, including a Pigouvian tax on emissions, studying the role of supply-side policy and its implications for market structure and innovation remains an important avenue for future research. In addition, this paper focuses on only one observable form of innovation to aid identification. As a result, future work is needed to understand other choices managers make to improve their technologies and remain competitive in the rapidly evolving market for solar panels.

This paper highlights two important points that apply beyond the solar market. First, using static economic methods to analyze markets experiencing rapid innovation can be misleading. It is important for firms and governments to recognize that innovation decisions can have larger impacts on profits and social objectives than short-term decisions about production or adoption of a technology. Second, international coordination may be needed to address innovation spillovers across borders when markets are interconnected.

References

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.

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.

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

Benitez-Silva, H., G. Hall, G. J. Hitsch, G. Pauletto, and J. Rust (2000). A Comparison of Discrete and Parametric Approximation Methods for Continuous-State Dynamic Programming Problems. *Manuscript, Yale University*.

Bollinger, B. and K. Gillingham (2019). 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.

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 (2019). Subsidies and Time Discounting in New Technology Adoption: Evidence from Solar Photovoltaic Systems. *American Economic Review 109*(6), 2137–2172.

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

Feger, F., N. Pavanini, and D. Radulescu (2021). Welfare and Redistribution in Residential Electricity Markets with Solar Power. Discussion Paper DP12517, Centre for Economic Policy Research.

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 (2019). Hurdles and steps: Estimating demand for solar photovoltaics. *Quantitative Economics 10*(1), 275–310.

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.

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

Hausman, J., G. Leonard, and J. D. Zona (1994). Competitive Analysis with Differentiated Products. *Annales d'Économie et de Statistique* (34), 159–180.

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.

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.

Igami, M. (2017). Estimating the Innovator's Dilemma: Structural Analysis of Creative Destruction in the Hard Disk Drive Industry, 1981–1998. *Journal of Political Economy 125*(3), 798–847.

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. (2014). Time to Build and Fluctuations in Bulk Shipping. *American Economic Review 104*(2), 564–608.

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.

Langer, A. and D. Lemoine (2018). Designing Dynamic Subsidies to Spur Adoption of New Technologies. Working Paper 24310, National Bureau of Economic Research.

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.

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

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

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

Porter, R. H. (1983). A Study of Cartel Stability: The Joint Executive Committee, 1880-1886. *The Bell Journal of Economics 14*(2), 301–314.

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.

Roberts, M. J. and W. Schlenker (2013). Identifying Supply and Demand Elasticities of Agricultural Commodities: Implications for the US Ethanol Mandate. *American Economic Review 103*(6), 2265–2295.

Rust, J. (1987). Optimal Replacement of GMC Bus Engines: An Empirical Model of Harold Zurcher. *Econometrica 55*(5), 999–1033.

Ryan, S. P. (2012). The Costs of Environmental Regulation in a Concentrated Industry. *Econometrica 80*(3), 1019–1061.

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.

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.