

Learning-by-doing Spillovers and Industrial Policy: Evidence from Residential Solar Installation^{*}

Jacob T. Bradt [†]

January 2026

Abstract

Learning-by-doing and knowledge spillovers fundamentally shape industry equilibrium when cumulative production experience reduces costs. I estimate these learning mechanisms in California's residential solar installation industry using a dynamic structural model with endogenous entry and exit. A 1% experience increase reduces costs by 0.21-0.36%, with in-market rivals' experience generating 74% of own learning benefits. Removing learning economies reduces installations by 10% and contracts market structure. Policy effectiveness depends critically on learning magnitude and spillovers: consumer subsidies leverage spillovers to expand adoption and industry size, while entry subsidies more effectively stimulate competition but at substantial fiscal cost.

Keywords: Clean Technology Policy, Learning-by-doing, Industry Dynamics, Industrial Policy

JEL Codes: L22, L52, Q42, Q54, Q55

^{*}I thank Joe Aldy, Myrto Kalouptsidi, Ariel Pakes, and Robin Lee for their guidance and support. I also thank Galen Barbose, Jackson Dorsey, Todd Gerarden, Ken Gillingham, Ed Glaeser, Jeff Gortmaker, Akshaya Jha, Devesh Raval, Nick Ryan, Rob Stavins, Andrew Steck, Rich Sweeney, Chris Timmins, and Tom Wollmann as well as seminar and conference participants at KU Leuven, EARIE 2024, SITE 2024, ES-DSE 2024, BSE Summer Forum 2024, IIOC 2024, Cornell, The University of Texas at Austin, UC San Diego, Harvard, and AERE Summer Conference 2023 for invaluable comments and discussion. I gratefully acknowledge funding from the Vicki Norberg Bohm Fellowship and the Harvard Pre-Doctoral Fellowship on Carbon Pricing and Alternative Instruments in Future U.S. Energy and Climate Policy. I am responsible for all remaining errors. A previous version of this work was circulated under the title "A Policy by Any Other Name: Unconventional Industrial Policy in the US Residential Solar Industry."

[†]The University of Texas at Austin, McCombs School of Business; jacob.bradt@austin.utexas.edu.

Learning-by-doing, the process through which cumulative production experience reduces firms' costs, is a fundamental force shaping industry evolution. When learning economies are substantial and knowledge spills over across firms, industry dynamics differ markedly from settings with static costs: firms face dynamic incentives to expand output today to reduce their own future costs, while recognizing that their production also reduces rivals' future costs through spillovers (Ghemawat and Spence, 1985). These learning dynamics shape equilibrium prices, quantities, entry, and exit in ways that depend critically on both the magnitude of learning and the extent of knowledge transfer across firms. Understanding these mechanisms is essential for explaining observed industry outcomes and evaluating policies that affect production incentives.

The presence of learning-by-doing and spillovers creates distinct equilibrium dynamics. First, learning links current production to future costs across all firms, generating dynamic production incentives beyond static profit maximization. Firms expand output not only for current profits but also to reduce future costs through experience accumulation. Second, spillovers create an externality: each firm's production reduces rivals' future costs, amplifying industry-wide learning but weakening individual incentives to invest in learning. Third, these forces interact with market structure: learning economies can facilitate entry by reducing costs for all firms, while spillovers may discourage entry if incumbents capture learning benefits from industry-wide production. Quantifying learning and spillovers is therefore necessary to understand how industries characterized by experience-driven cost reductions evolve and how policies that stimulate production affect equilibrium outcomes.

In this paper, I estimate learning-by-doing and knowledge spillovers in the solar photovoltaic (PV) installation industry and analyze how these learning mechanisms shape equilibrium outcomes and policy effectiveness. I develop and estimate a dynamic structural model that endogenizes firms' entry, exit, and production decisions while explicitly incorporating learning and spillovers into firms' cost functions and dynamic incentives. Using data on California's residential solar market from 2008 to 2013, I find substantial learning-by-doing—a 1% increase in experience reduces costs by 0.21-0.36%—with large spillovers that amplify industry-wide learning. Counterfactual analysis reveals these learning mechanisms are central to observed market outcomes: removing learning economies reduces installations by 10% and substantially contracts market structure. The magnitude and spillover structure of learning determine how policies affect equilibrium: consumer subsidies leverage learning to expand adoption and industry size, with effectiveness depending critically on spillover rates.

Solar PV is a key climate technology due to its minimal life cycle emissions and ability to displace fossil fuel generation. Policymakers have provided substantial subsidies for solar adoption, with many programs targeting residential consumers. The non-trivial design and

construction of PV systems has created an industry of intermediary installation firms that employed over 171,000 US workers in 2022—65% of total solar employment (Interstate Renewable Energy Council, 2023). Installers account for a growing share of final costs: Barbose et al. (2022) estimate installers’ share of residential costs rose from 40% in 2006 to over 80% in 2016. Despite this growing share, installation costs have fallen, with evidence suggesting installer learning-by-doing (Bollinger and Gillingham, 2019; Fu et al., 2016; Nemet, 2019).

Despite accounting for most residential solar costs, relatively little is known about installers compared to manufacturers. California offers an ideal setting to study learning dynamics in this industry. California hosts nearly half of all US residential PV systems (Barbose et al., 2022) and experienced substantial subsidy-driven growth during 2008-2013. The California Solar Initiative (CSI) provided \$2.2 billion in consumer rebates, creating variation in incentives that enables identification of learning parameters while also providing a natural application for evaluating how learning mechanisms affect policy outcomes.

I develop a dynamic structural model of California’s residential solar installation market based on Ericson and Pakes (1995) that explicitly incorporates learning-by-doing and spillovers into firms’ cost structures and decision-making. Combined with unique hardware cost data, it allows me to separate installation costs, where learning occurs, from hardware costs. Incumbent installers’ costs depend on both their own cumulative production and rivals’ cumulative production, creating two key incentives: firms benefit from expanding output today to reduce their own future costs, but they also recognize that their production reduces rivals’ future costs through spillovers. In each geographic market, incumbent firms choose output levels accounting for these learning effects and their impact on future competition. Consumer demand for differentiated installations follows the random coefficient nested logit model of Brenkers and Verboven (2006). Incumbent firms compare expected discounted future profits with idiosyncratic scrap values and make irreversible exit decisions, while potential entrants make one-shot entry decisions based on expected profits and entry costs. Firms’ strategies lead to a Markov Perfect Equilibrium, approximated by a Moment-based Markov Equilibrium (Ifrach and Weintraub, 2017).

I estimate the model using system-level data on prices, rebates, capacities, and hardware costs for 95% of California residential PV systems from 2008 to 2013. I obtain installation data from Lawrence Berkeley National Laboratory’s “Tracking the Sun” database, including timing, location, and installer identity. I combine these with system-level hardware costs from the California Public Utilities Commission. The hardware cost data are critical for isolating the installation cost component where learning occurs. I aggregate data to the county-half-year level for all installers operating from 2008 to 2013.

My estimation approach builds on two-step estimators of dynamic games (Bajari et al.,

2007; Collard-Wexler, 2013; Fowlie et al., 2016; Pakes et al., 2007; Ryan, 2012). In the first stage, I estimate demand parameters, exit policy functions, and state transitions. I use these estimates to flexibly approximate firms' value functions following Barwick and Pathak (2015); Barwick et al. (2025); Kalouptsidi (2018); Sweeting (2013). In the second stage, I form moments from optimal quantity-setting and exit conditions to recover production cost parameters (i.e, learning) and exit costs, then use entry likelihood to recover entry costs.

The model estimates reveal two main findings. First, I find substantial learning-by-doing: a 1% increase in effective experience decreases marginal installation costs by 0.21 to 0.36%. The implied Spence coefficient—proportional cost reduction from doubling experience—ranges from 0.13 to 0.22. While more modest than manufacturing learning curves, these estimates are consistent with a service-intensive process where learning occurs through improved coordination and permitting efficiency. Model-estimated average marginal costs closely match independent NREL estimates from 2008 to 2013.

Second, learning spills over substantially across firms. A 1 unit increase in total in-market rival experience generates 74% of the learning benefit of a firm's own experience increase. To explore mechanisms, I estimate specifications allowing differential spillovers by experience source. I find larger spillovers from rivals in the same geographic market, suggesting knowledge transfer operates through market-level mechanisms such as worker mobility, observable installation practices, or passive learning by regulators.

Counterfactual simulations reveal how learning mechanisms shape equilibrium outcomes and policy effects. First, learning-by-doing fundamentally determines market equilibrium: removing learning economies while holding policies fixed reduces installations by 10%, increases equilibrium prices persistently, and substantially contracts the number of active firms. This demonstrates that observed industry outcomes reflect learning dynamics, not just static production costs and demand. Second, spillover rates determine how production affects industry-wide costs: high spillovers amplify the cost-reducing effects of aggregate production, making demand expansion valuable for reducing future costs across all firms. Consumer subsidies exploit this by increasing production, with effectiveness scaling with spillover magnitude. Third, the interaction between learning and market structure creates feedback: subsidies stimulate production, reducing costs through learning, which facilitates entry and further production.

Finally, supply-side entry subsidies can leverage learning more effectively than demand subsidies, but at substantial fiscal cost. Replacing the CSI with entry subsidies of varying sizes dramatically increases active firms (by up to 62%) and installations. This expanded competition drives lower prices and expanded adoption. While entry subsidies generate significant benefits through increased consumer surplus and reduced exit, these gains are

offset by large government expenditures: entry subsidies reduce total welfare relative to the CSI by \$180 million to \$1.3 billion, primarily reflecting fiscal costs of \$2.1 billion to \$8.4 billion. These results suggest that while entry subsidies effectively expand market size and competition by directly stimulating learning, their substantial fiscal burden and different political economy must be weighed against demand-side approaches.

These findings contribute to our understanding of how learning-by-doing and spillovers shape industry equilibrium. Theoretical work establishes that cumulative experience affects market outcomes (Arrow, 1962; Besanko et al., 2010; Cabral and Riordan, 1994; Fudenberg and Tirole, 1983; Spence, 1981), with Ghemawat and Spence (1985) showing non-appropriable learning influences market structure when knowledge spills over across firms. A large empirical literature estimates learning curves in aircraft (Benkard, 2000, 2004), ships (Thompson, 2001, 2007; Thornton and Thompson, 2001), semiconductors (Irwin and Klenow, 1994), oil (Kellogg, 2011), automobiles (Levitt et al., 2013), and wind turbines (Covert and Sweeney, 2022), with several finding spillovers (Covert, 2015; Irwin and Klenow, 1994; Kellogg, 2011; Thornton and Thompson, 2001). Bollinger and Gillingham (2019) estimate learning by solar PV installers with spillovers.¹ I build on this literature by embedding learning and spillovers in a dynamic oligopoly model with endogenous entry and exit, allowing me to trace how learning mechanisms affect equilibrium prices, quantities, and market structure.

My results also contribute to evaluating policies that affect industries with learning economies. A large literature on solar PV subsidies focuses on adoption, finding consumer subsidies increase installations but may not be justified by static environmental benefits (Borenstein, 2017; De Groote and Verboven, 2019; Dorsey, 2024; Gillingham and Tsvetanov, 2019; Hughes and Podolefsky, 2015). Gerarden (2022) shows accounting for manufacturer innovation can justify subsidy levels. van Benthem et al. (2008) and Langer and Lemoine (2022) demonstrate via simulation that learning-by-doing can rationalize PV subsidy design. I contribute empirical evidence that learning mechanisms—both magnitude and spillover structure—fundamentally determine policy effects. Understanding these mechanisms is necessary for predicting how subsidies affect adoption, market structure, prices, and welfare. The results also speak to the growing industrial policy literature (Juhász et al., 2023), showing that learning dynamics determine how demand subsidies affect industry size and structure, complementing existing work on R&D subsidies (Bloom et al., 2002; Hall and Van Reenen, 2000) and production subsidies (Barwick et al., 2025; Kalouptsidi, 2018).

The rest of the paper is organized as follows. Section 1 provides an overview of the solar PV industry and policy environment. Section 2 discusses the data I use in my analysis and

¹My model endogenizes entry and exit whereas Bollinger and Gillingham (2019) hold these fixed. I also account for serially-correlated productivity shocks in production cost estimation.

provides some descriptive results on solar PV installers in California. Section 3 presents the model. Sections 4 and 5 describe estimation and the model estimates. Finally, Section 6 presents results from counterfactual policy simulations while Section 7 concludes.

1 Economic and Policy Landscape

1.1 Solar PV Industry


The global solar industry has grown rapidly since the first commercial application of PV technology on satellites in the 1950s. Solar modules consist of interconnected solar cells that convert sunlight into electricity via the photovoltaic effect. Since Bell Labs created the first practical solar cell in 1954, technological innovation and improved manufacturing efficiency have substantially reduced solar module costs (Nemet, 2019). From 1975 to 2021, solar module prices declined over 99%, from \$115 to under \$0.5 (2021 USD) per watt (IRENA, 2022; Nemet, 2009). Global solar capacity grew 1000-fold from just under 1 gigawatt to over 1 terawatt from 2000 to 2022 (IRENA, 2023).

Solar PV is a modular technology manufactured at-scale, enabling applications from utility-scale generation to small residential systems, the focus of this paper. Residential PV installation requires non-trivial design and construction. Rooftop installation, which accounts for 98% of residential use in the US, introduces site-specific features requiring idiosyncratic design. Practical installation challenges and technical electrical components require specialized labor. These features combined with convoluted regulatory environments involving varying permitting and inspection requirements across jurisdictions and often generous but difficult-to-navigate incentive programs have given rise to an industry marketing PV installation as a service. Installation firms source hardware inputs, design and construct systems, and manage permitting and inspection for households.

While solar panel costs have declined dramatically at a global scale, an increasing share of end consumer costs is attributable to local installers. Figure 1 shows Lawrence Berkeley National Lab data on solar PV hardware costs and median installed costs for US residential consumers over 2000-2021. The installer share (installation labor, permitting, and markups) more than doubled over 2006-2016, going from 40% to over 80% (Barbose et al., 2022).

Despite accounting for a growing share of final costs, these “soft” costs have fallen in absolute terms in recent years, though far less than hardware costs (Fu et al., 2016). Evidence suggests soft cost reductions stem from installer process improvements, inter-firm learning, and streamlined policies (Bollinger and Gillingham, 2019; Nemet, 2019; Nemet et al., 2017). Existing estimates struggle to separate installer costs and markups, though Bollinger and Gillingham (2019) addresses this. Determining the magnitude and sources of installation-

Figure 1. PV System Installed Cost Components, 2000-2021



Notes: This figure shows installed cost and hardware component cost per watt for residential PV systems in the US from Lawrence Berkeley National Lab’s “Tracking the Sun” report (Barbose et al., 2022). Hardware costs include PV modules and inverters. Bars show the share of installed cost attributable to non-hardware costs (installation labor, permitting, markups, etc.).

specific cost reductions is a key objective of this paper.

1.2 California’s PV Policy Environment

California hosts nearly half of all US residential solar PV systems (Barbose et al., 2022). While climate drives adoption, generous policies also play a major role. Since the late 2000s, California households have been eligible for adoption incentives at state and federal levels.

The California Solar Initiative (CSI), the state’s largest direct rebate program, ran from 2007 to 2013 with a \$2.2 billion budget, providing cash rebates to customers of three main investor-owned utilities (IOUs). The CSI rebate schedule was designed explicitly with learning-by-doing in mind: rebates started at \$2.50 per watt and stepped down over 10 rate levels based on cumulative installed capacity in each IOU service area. This design assumed industry experience would reduce costs, thereby reducing rebates needed to incentivize adoption. Appendix Figure A1 shows spatial and temporal variation in CSI rebates across IOUs, with rebate steps changing at different times based on cumulative capacity. These sharp changes provide plausibly-exogenous variation in net-of-rebate prices useful for estimation.

Solar-installing households have also been eligible for a 30% federal investment tax credit (ITC) since 2007 and net energy metering (NEM) crediting excess generation at retail rates.

See Appendix [A.2](#) for full policy details.

Beyond reducing emissions, an express goal of the CSI was to “establish a self-sufficient solar industry” ([California State Senate, 2006](#)). A primary motivation of this paper is evaluating this objective by estimating the extent to which the CSI reduced costs and changed industry structure.

2 Data and Descriptive Evidence

2.1 Data Sources

I construct a dataset tracking installation firms’ prices, market shares, hardware costs, experience, and entry-exit decisions across county-level markets and half-yearly periods from 2008 to 2013. This section summarizes key data sources and restrictions; see Appendix [A](#) for a more detailed discussion of these data.

I obtain installer data from Lawrence Berkeley National Laboratory’s “Tracking the Sun” database ([Barbose et al., 2022](#)), which compiles system-level data from state agencies and utilities. The database includes installation date, system size, prices, rebates, location, installer identity, and hardware specifications. California data cover over 98% of state installations. I apply several restrictions, focusing on residential rooftop systems below 20 kW with observed prices and rebates, excluding self-installed and third-party-owned systems.


I obtain hardware cost data from the California Public Utilities Commission, allowing me to separate hardware costs from installation costs. Coverage is most complete for 2008-2013: I match hardware costs to over 79% of systems in this period. Given hardware costs are essential for isolating the installation component where learning occurs, I restrict the sample to 2008-2013.

I aggregate system-level data to the county-half-year level for each installer, calculating average prices, installed capacity, rebates, and hardware costs. I use full 2000-2021 data to calculate each installer’s cumulative experience and identify entry and exit during 2008-2013. The final dataset includes roughly 10,000 observations representing over 94,000 installations by 500 installers across 33 counties over 12 half-yearly periods. Additional data sources include Census ACS data for market size estimation and demand heterogeneity, and EIA electricity rate data for estimating net metering benefits.

2.2 Descriptive Evidence

Figure [2](#) shows time series variation in key measures of California residential PV installation industry activity for 2004-2013, covering the main CSI period and three prior years.

Figure 2. California Residential PV Installation Activity, 2004-2013




Notes: This figure shows time series variation in California residential PV installation industry activity over 2004-2013: cumulative installed capacity in megawatts (top left), total revenue in million 2013 USD (top right), statewide Herfindahl-Hirschman Index based on installed capacity (bottom left), and number of statewide incumbent firms (bottom right). Data are for rooftop, household-owned installations from LBNL's "Tracking the Sun" database (Barbose et al., 2022).

Like the global solar industry, California's PV industry experienced dramatic growth starting in the mid-2000s. From 2007 to 2013, the main period of the CSI, total installed residential PV capacity increased nearly 10-fold, from 60 to 600 MW. This rapid capacity growth corresponded with expansion of the installation industry: both statewide revenue and the number of operating firms increased rapidly during this period. Installer industry concentration remained low, with a statewide Herfindahl-Hirschman Index between 100 and 540 from 2007 to 2013.

While these trends are suggestive, they provide limited insight into the mechanisms through which PV policy shaped industry growth. To explore potential learning spillovers, I estimate the relationship between rivals' cumulative experience in a county and an installer's own costs plus markups (measured as installation price per watt net of hardware costs). I control for county and year fixed effects to account for time-invariant market characteristics and aggregate time trends. Figure 3 shows a strong, negative relationship between rivals' experience and own costs plus markups over 2007-2013, consistent with learning spillovers

Figure 3. Rivals' Experience and Own Costs, Markups, 2007-2013



Notes: This figure shows the estimated relationship between rivals' normalized cumulative installation experience in a county and a firm's costs plus markups (installation price per watt net of hardware costs) over 2007-2013. The figure plots mean residuals conditional on county and year fixed effects.

reducing installation costs.²

While this descriptive evidence suggests learning spillovers may be important, it cannot quantify their magnitude or identify specific mechanisms. To precisely analyze how learning-by-doing and knowledge spillovers affect the PV installation industry and the impact of policy, I develop a structural model of installer entry, exit, and production.

3 Model: Entry and Exit with Endogenous Learning-by-Doing

I develop a model of firm entry, exit, and quantity-setting based on Ericson and Pakes (1995)'s dynamic oligopoly framework.

In each period t and market $m \in \{1, \dots, M\}$, there are $j \in \{1, \dots, J_{mt}\}$ incumbent firms facing static consumers $i \in \{1, \dots, N_{mt}\}$ who demand differentiated solar PV installation services. Incumbents dynamically choose installation quantities conditional on marginal costs and beliefs about future learning. Incumbents then choose whether to exit by comparing expected discounted future profits with an idiosyncratic scrap value, while a market-specific pool of potential entrants $j \in \{1, \dots, \bar{N}_m\}$ make one-shot entry decisions based on expected discounted future profits and an idiosyncratic entry cost. At each period's end, entry and exit decisions are implemented and the state evolves. Firms' strategies lead to a Markov Perfect Equilibrium, which I approximate using a Moment-based Markov Equilibrium concept.

²This descriptive relationship suggests rivals' experience may reduce firms' costs through mechanisms like worker mobility, visibility of rivals' practices, or regulator learning that benefits all firms in a market. The structural model in Section 3 formalizes and estimates these spillover effects.

A period is a half-year and a market is a county. Firms have an infinite horizon and share discount factor β . Incumbent installer j in market m at time t is differentiated by its state, which includes a common knowledge component s_{jmt} and a private component. The private component includes a shock to the firm's sell-off value ϕ_{jmt} and an unobserved productivity shock κ_{jmt} . The common knowledge component is

$$s'_{jmt} = [E_{jmt} \ \ \xi_{jmt} \ \ h_{jmt}]$$

where E_{jmt} is the installer's market-specific experience; ξ_{jmt} is installation service quality derived from the demand system; and h_{jmt} is hardware input costs. The market state s_{mt} is the union of all incumbent firms' common knowledge states plus two aggregate state variables: demand state d_{mt} and market-level inclusive value I_{mt} , which capture revenue potential and competition intensity (Aguirregabiria et al., 2021). Potential entrants observe the market state and are differentiated by an idiosyncratic entry cost shock ω_{jmt} .

3.1 Demand for Solar Installations

I estimate consumer demand using the random coefficient nested logit (RCNL) model of Brenkers and Verboven (2006) and Grigolon and Verboven (2014). Incumbent firms face static consumers $i \in \{1, \dots, N_{mt}\}$ who demand solar PV installation services. I assume a static demand model reasonably approximates consumer behavior.³ Each consumer purchases installation from an observed incumbent $j \in \{1, \dots, J_{mt}\}$ or does not install ($j = 0$). Conditional indirect utility from choosing installer j in market m in period t is

$$u_{ijmt} = \alpha_i^p(p_{jmt} - r_{jmt}) + \alpha'_i X_{jmt} + \xi_{jmt} + \bar{\xi}_j + \bar{\xi}_t + \bar{\varepsilon}_{ijmt} \quad (1)$$

where X_{jmt} is a $K \times 1$ vector of observable firm characteristics; p_{jmt} is retail price per watt; r_{jmt} is rebate per watt; ξ_{jmt} is firm's market-time-specific unobserved quality; $\bar{\xi}_j$ allows mean valuation of unobserved characteristics to vary by product; $\bar{\xi}_t$ allows mean utility from installation to vary over time; and $\bar{\varepsilon}_{ijmt}$ is an idiosyncratic preference shock. I normalize prices and rebates by capacity for consistency across system sizes.

Following Berry (1994), I decompose the idiosyncratic preference shock using nested logit distributional assumptions. Define two groups $g \in \{0, 1\}$, where $g = 1$ includes incumbent installers and $g = 0$ the singleton no-installation option. Then $\bar{\varepsilon}_{ijmt} = \zeta_{igmt} + (1 - \eta)\varepsilon_{ijmt}$ where ε_{ijmt} is i.i.d. Type 1 Extreme Value, ζ_{igmt} has the unique distribution such that $\bar{\varepsilon}_{ijmt}$ is

³I develop and estimate a dynamic demand model in Appendix B and find my static demand estimates provide a reasonable reduced form.

i.i.d. Type 1 Extreme Value, and $0 \leq \eta < 1$ is a nesting parameter proxying for within-group preference correlation. I normalize non-installation utility such that $u_{i0,mt} = \varepsilon_{i0mt}$.

Taste heterogeneity is parameterized as $\alpha_i^p = \alpha_p/y_i$ and $\alpha_i^k = \alpha_k + \sigma_k \log(y_i)$ for attribute k , where α_p , α_k , σ_k are parameters and y_i is consumer income.⁴ This allows me to re-write conditional indirect utility as

$$u_{ijmt} = v_{jmt} + \mu_{ijmt} + \zeta_{igmt} + (1 - \eta)\varepsilon_{ijmt}$$

where

$$v_{jmt} \equiv \alpha' X_{jmt} + \xi_{jmt} + \bar{\xi}_j + \bar{\xi}_t \quad \text{and} \quad \mu_{ijmt} \equiv \frac{\alpha^p}{y_i} (p_{jmt} - r_{jmt}) + \sigma' X_{jmt} \log(y_i)$$

Market share of installer j in market m in period t is

$$ms_{jmt} = \frac{1}{N_{mt}} \sum_{i=1}^{N_{mt}} \frac{\exp((v_{jmt} + \mu_{ijmt})/(1 - \eta))}{\exp(I_{igmt}/(1 - \eta))} \frac{\exp I_{igmt}}{\exp I_{imt}} \quad (2)$$

where I_{igmt} and I_{imt} are the [McFadden \(1977\)](#) inclusive values.⁵

3.2 Incumbent Cost Structure and Payoffs

After observing market state s_{mt} , incumbents privately observe productivity shock κ_{jmt} and choose installation quantity q_{jmt} at cost $mc_j(s_{mt})$ per watt. Marginal production cost is

$$mc_j(s_{mt}; \theta^c) = c_0 \times \left(h_{jmt} \right)^\tau \times \left(\tilde{E}_j(s_{mt}; \theta^E) \right)^\gamma \times e^{\kappa_{jmt}} \quad (3)$$

where h_{jmt} is exogenous hardware cost per watt; $\tilde{E}_j(s_{mt}; \theta^E)$ is firm j 's “effective” experience, a function of own and rivals’ experience with parameters θ^E ; γ is the learning exponent; τ is the hardware cost pass-through parameter; c_0 scales the cost function; and κ_{jmt} is an unobserved productivity shock.⁶ The parameter vector is $\theta^c = (c_0, \tau, \theta^E, \gamma)$.

This specification follows unbounded learning models standard in the learning-by-doing literature ([Benkard, 2000](#); [Covert and Sweeney, 2022](#); [Levitt et al., 2013](#); [Thornton and Thompson, 2001](#)). Defining marginal cost as a function of effective experience allows me to

⁴This α_i^p parameterization approximates Cobb–Douglas indirect utility ([Berry et al., 1999](#)).

⁵The inclusive value of the inside goods is $I_{i1mt} = (1 - \eta) \log \sum_{j=1}^{J_{mt}} \exp((\delta_{jmt} + \mu_{ijmt})/(1 - \eta))$ and the inclusive value of all goods is $I_{imt} = \log(1 + \exp I_{i1mt})$. Thus, to get the [McFadden \(1977\)](#) inclusive value at the market-time-level, I simply sum across individuals in the market: $I_{mt} = \sum_{i=1}^{N_{mt}} I_{imt}$.

⁶I index functions by j to indicate firm-specific values, e.g., $mc_j(s_{mt})$. Since state vector s_{mt} includes each firm’s state, this avoids duplicating s_{jmt} in function arguments.

test different models of experience accumulation. One model I test is

$$\tilde{E}_j(s_{mt}; \theta^E) = E_{jmt} + \theta_1^E \left(\sum_{k \neq j} E_{kmt} \right) \quad (4)$$

where market rivals' cumulative production has a potentially different marginal contribution to firm j 's effective experience than own cumulative production. The experience parameter θ_1^E is normalized with respect to own experience.⁷ I report estimated experience parameters from several models of experience accumulation in Section 5.

An incumbent active in period t in market m earns product market profits

$$\pi_j(s_{mt}, q_{jmt}; \theta^c) = (p_j(s_{mt}, q_{jmt}) - mc_j(s_{mt}; \theta^c)) q_{jmt}$$

where $p_j(s_{mt}, q_{jmt})$ is firm j 's price per watt, defined by the inverse demand curve from Section 3.1.

The ex-ante value function for incumbent j in market m at time t prior to realization of ϕ_{jmt} is

$$V_j(s_{mt}) = \mathbb{E}_\phi \left[\pi_j(s_{mt}) + \max \left\{ \phi_{jmt}, CV_j(s_{mt}) \right\} \right] \quad (5)$$

where $CV_j(s_{mt}) = \mathbb{E}[V_j(s_{mt+1})|s_{mt}, q_{mt}^*]$ is the continuation value (expected discounted future profits with expectation over state variable transitions); q_{mt}^* is the vector of optimal quantities; and β is the discount factor.

3.3 Product Market Game

Incumbents compete each period by choosing installation quantities to maximize current period profits plus expected continuation value. Firm j in market m in period t solves:

$$\max_{q_{jmt}} \left(\pi_j(s_{mt}, q_{jmt}) + \beta \int V_j(s_{mt+1}) dF(s_{mt+1}|s_{mt}, q_{mt}) \right)$$

where $F(s_{mt+1}|s_{mt}, q_{mt})$ is the transition kernel for state s_{mt} conditional on quantities q_{mt} .

The optimal quantity satisfies:

$$0 = \underbrace{\frac{\partial}{\partial q_{jmt}} \pi_j(s_{mt}, q_{jmt})}_{\text{marginal static profits}} + \underbrace{\frac{\partial}{\partial q_{jmt}} \beta \int V_j(s_{mt+1}) dF(s_{mt+1}|s_{mt}, q_{mt})}_{\text{dynamic "markdown"}}$$

⁷I normalize all experience terms by total industry experience in H1 2008, ensuring readability of θ^E parameters and improving numerical stability.

The first term is the standard static quantity-setting condition; the second captures the incentive to raise production today to reduce future costs and any strategic considerations regarding impacts on rivals.⁸ The first term can be written:

$$\frac{\partial}{\partial q_{jmt}} \pi_j(s_{mt}, q_{jmt}) = p_j(s_{mt}, q_{jmt}) + \frac{\partial p_j(s_{mt}, q_{jmt})}{\partial q_{jmt}} q_{jmt} - mc_j(s_{mt}; \theta^c) \quad (7)$$

This standard condition trades off marginal benefits and costs in the current period, accounting for both the direct effect of producing a marginal unit and the inframarginal impact of reducing equilibrium price on all units supplied.

The dynamic markdown describes the marginal effect of current production on discounted future profits. As noted by [Berry and Pakes \(2000\)](#) and [Covert and Sweeney \(2022\)](#), changes in q_{jmt} affect the transition distribution $F(s_{mt+1}|s_{mt}, q_{mt})$ only through experience evolution E_{jmt} , not the value function $V(\cdot)$ itself. Since $dF(s_{mt+1}|s_{mt}, q_{mt}) > 0$ for any q_{mt} , the dynamic markdown simplifies to:

$$\begin{aligned} \frac{\partial}{\partial q_{jmt}} \beta \int V_j(s_{mt+1}) dF(s_{mt+1}|s_{mt}, q_{mt}) &= \beta \int V_j(s_{mt+1}) \frac{\partial}{\partial q_{jmt}} dF(s_{mt+1}|s_{mt}, q_{mt}) \\ &= \beta \int V_j(s_{mt+1}) \left(\frac{\frac{\partial}{\partial q_{jmt}} dF(s_{mt+1}|s_{mt}, q_{mt})}{dF(s_{mt+1}|s_{mt}, q_{mt})} \right) dF(s_{mt+1}|s_{mt}, q_{mt}) \\ &= \beta \mathbb{E} \left[V_j(s_{mt+1}) \left(\frac{\frac{\partial}{\partial q_{jmt}} dF(s_{mt+1}|s_{mt}, q_{mt})}{dF(s_{mt+1}|s_{mt}, q_{mt})} \right) \middle| s_{mt}, q_{mt} \right] \end{aligned} \quad (8)$$

This simplified form shows firm j 's dynamic markdown is the expected benefits at a given future state realization s_{mt+1} multiplied by the marginal change in probability that this state is realized from changing q_{jmt} . Combining these forms for marginal static profits and dynamic markdown provides a feasible approach to writing firms' quantity-setting condition as a function of data and estimable parameters, outlined in Section 4.

3.4 Exit and Entry

After product market decisions, incumbents draw a private scrap value ϕ_{jmt} . The optimal exit policy follows a threshold form: firm j exits market m in period t if the scrap value exceeds its continuation value, $CV_j(s_{mt})$. With i.i.d. scrap values, the firm exits with probability $p_j^x(s_{mt})$:

$$p_j^x(s_{mt}) \equiv \Pr(\phi_{jmt} > CV_j(s_{mt})) = 1 - F_\phi(CV_j(s_{mt}))$$

⁸I follow [Covert and Sweeney \(2022\)](#) in referring to the second term as the “dynamic markdown.”

where F_ϕ is the CDF of ϕ_{jmt} . I assume $\phi_{jmt} \stackrel{\text{i.i.d.}}{\sim} \text{Exponential}(1/\sigma_\phi)$, implying:

$$p_j^x(s_{mt}; \sigma_\phi) = \exp\left(-\frac{CV_j(s_{mt})}{\sigma_\phi}\right) \quad (9)$$

Simultaneously, \bar{N}_m potential entrants observe state s_{mt} and a private i.i.d. entry cost ω_{jmt} before making a one-shot entry decision. If potential entrant $j \in \{1, \dots, \bar{N}_m\}$ enters, it pays ω_{jmt} and becomes an incumbent next period; otherwise it disappears with zero payoff. Entrants are endowed with quality and hardware cost drawn from the empirical distribution of observed states and start with zero experience, $E_{jmt} = 0$.

The optimal entry policy follows a threshold form: potential entrant j enters market m in period t if entry cost ω_{jmt} is below the value of entering:

$$\omega_{jmt} \leq VE_j(s_{mt}) \equiv \beta \mathbb{E}[V_j(s_{mt+1}) | s_{mt}, \chi_{jmt}^e = 1]$$

where χ_{jmt}^e equals 1 if potential entrant j enters and 0 otherwise, and the expectation is over the entrant's information set, which includes state s_{mt} . With i.i.d. ω_{jmt} , the entry probability is:

$$p_j^e(s_{mt}) \equiv \Pr\left(\omega_{jmt} \leq VE_j(s_{mt})\right) = F_\omega\left(VE_j(s_{mt})\right)$$

where F_ω is the CDF of ω_{jmt} . I assume $\omega_{jmt} \stackrel{\text{i.i.d.}}{\sim} \text{Exponential}(1/\sigma_\omega)$, implying:

$$p_j^e(s_{mt}; \sigma_\omega) = 1 - \exp\left(-\frac{VE_j(s_{mt})}{\sigma_\omega}\right) \quad (10)$$

3.5 State Transitions

I assume hardware cost (h_{jmt}) and installation quality (ξ_{jmt}) are exogenous and evolve stochastically according to a first-order Markov process. For hardware costs, this implies installers are price-takers in the upstream input market, reasonable given many module and inverter manufacturers. Data limitations preclude endogenizing quality.

The remaining state variables—aggregate demand (d_{mt}), market-level inclusive value (I_{mt}), and experience (E_{jmt})—are endogenous and evolve from the demand model and firm quantity-setting actions.

3.6 Equilibrium

A Markov-Perfect Equilibrium in market m at period t includes policies governing production, exit, and entry $(q_{mt}, p_j^x(s_{mt}), p_j^e(s_{mt}))$, value functions $V_j(s_{mt})$, and prices p_{jmt} such

that firms' decisions satisfy (6), (9), and (10). Equilibrium prices are generated by the inverse demand function and equate demand with supply. Incumbent value functions satisfy (5) and all firms employ equilibrium policy functions to form expectations. Equilibrium existence follows from [Ericson and Pakes \(1995\)](#) and [Doraszelski and Satterthwaite \(2010\)](#).

The state variable s_{mt} is high-dimensional when many firms are active. To reduce computational burden, I assume firms track only market-level moments of rivals' state variables rather than every rival's full state. [Ifrach and Weintraub \(2017\)](#) provide a detailed treatment of this Moment-based Markov Equilibrium approach, similar to oblivious equilibrium ([Benkard et al., 2015](#); [Weintraub et al., 2008](#)), which approximates Markov-Perfect Equilibrium in industries with many firms and has been widely employed ([Barwick et al., 2025](#); [Gerarden, 2022](#); [Jeon, 2022](#); [Vreugdenhil, 2023](#); [Wollmann, 2018](#)).

A remaining issue is nonstationarity in the regulatory environment. While numerous overlapping adoption subsidies evolve over the study window, explicitly modeling firms' beliefs about future subsidy distributions substantially complicates equilibrium computation. I follow standard practice and assume firms behave as if subsidy policy changes are unanticipated, one-time changes not repeated in the future ([Aguirregabiria et al., 2021](#); [Ryan, 2012](#)). As noted by [Barwick et al. \(2025\)](#), one approach to proxy for dynamic regulatory environments is using lower discount rates so future payoffs matter less for current decisions. I use this approach to test robustness of my permanent subsidy change assumption.

4 Estimation Strategy

I estimate the model in two stages. The first stage estimates the demand system, exit policy functions, and state transition processes. The second stage jointly estimates production cost parameters governing learning economies, exit costs, and entry costs.

4.1 First Stage: Demand Estimation, Exit Policies, and States

Demand Estimation I estimate demand following [Berry et al. \(1995\)](#)'s nested fixed point procedure and [Conlon and Gortmaker \(2020\)](#)'s best practices, adapted to the RCNL model of [Brenkers and Verboven \(2006\)](#) and [Grigolone and Verboven \(2014\)](#). I derive a GMM estimator from the moment condition $\mathbb{E}[Z'_D \xi(\theta_0^D)] = 0$, where $\theta_0^D = (\alpha_p, \alpha', \sigma, \eta)$ are demand parameters, $\xi(\theta^D)$ solves the system of market shares in (2), and Z_D are instruments. The GMM estimator is

$$\hat{\theta}^D = \arg \min_{\theta^D} (\hat{\xi}(\theta^D)' Z_D) W^{-1} (Z'_D \hat{\xi}(\theta^D))$$

where $\hat{\xi}(\theta^D)$ is the sample analog of $\xi(\cdot)$ and W is a positive definite weight matrix.

The installer characteristics that enter X_{jmt} include a measure of the efficiency of PV modules an installer offers, the number of distinct PV module types an installer offers, and the half-yearly average electricity price in a county. Including a firm fixed effect, $\bar{\xi}_j$, in (1) absorbs time-invariant installer characteristics and the period fixed effect, $\bar{\xi}_t$, accounts for aggregate trends in mean preferences for solar over time. In estimation, I take 200 draws of household income from the annual American Community Survey PUMS per county-period.

I adopt many of the best practices for differentiated demand estimation recommended by [Conlon and Gortmaker \(2020\)](#). I employ the standard two-step procedure for GMM estimation, adjusting the weight matrix in the second step to account for clustering at the county-level. I solve the system of market shares defined by (2) using SQUAREM with a damped version of the [Berry et al. \(1995\)](#) contraction mapping based on [Grigolon and Verboven \(2014\)](#). I calculate standard errors using the GMM formula, clustering observations at the county-level to allow for within-market correlation in unobserved quality.

Identification of demand parameters θ^D requires instruments addressing price endogeneity and identifying parameters governing consumer heterogeneity and demand curvature. Prices correlate with unobserved quality ξ_{jmt} as firms observe product valuations when setting prices ([Berry and Haile, 2014](#)). I use CSI rebates per watt and county-level electrician/roofing wages as cost shifters ([Gillingham and Tsvetanov, 2019](#); [Pless and Van Benthem, 2019](#)), plus rivals' non-price characteristics X_{-jmt} following [Berry et al. \(1995\)](#).

Identifying the nest parameter η and random coefficient parameters σ that govern demand curvature requires exogenous variation in the distribution of consumer valuations for inside goods. The number of active firms and lagged installations in other counties shift the distribution of inside-good market shares, identifying η . Following [Miller and Weinberg \(2017\)](#), mean income interacted with product characteristics X_{jmt} leverages demographic variation to identify σ . This approach identifies both demand elasticity (mean price sensitivity) and demand curvature (heterogeneity in price sensitivity), which jointly determine how firms' pricing responds to cost shocks and mergers.

Exit Policy Function I estimate exit probabilities using a logit regression:

$$\Pr(\chi_{jmt}^x = 1 | s_{mt}) = \frac{\exp(h_j(s_{mt}))}{1 + \exp(h_j(s_{mt}))}$$

where χ_{jmt}^x equals 1 if firm j exits market m in period t and $h_j(s_{mt})$ is a flexible function of states. Following [Gerarden \(2022\)](#), I select $h_j(s_{mt})$ via LASSO with k -fold cross validation,

then re-estimate the logit model with the selected regressors.⁹ I denote the resulting fitted exit probabilities as \hat{p}_{jmt}^x .

State Space Following the Moment-based Markov Equilibrium concept of [Ifrach and Weintraub \(2017\)](#), firms track moments of rivals' states rather than every rival's full state vector. Firms observe their own states $(E_{jmt}, \xi_{jmt}, h_{jmt})$ plus within- and out-of-county averages of rivals' hardware cost and quality, total rival experience in-county ($\bar{E}_{jmt}^m = \sum_{k \neq j} E_{kmt}$) and out-of-county ($\bar{E}_{jmt}^o = \sum_{l \neq m} \sum_{k \neq j} E_{klt}$), plus aggregate demand d_{mt} and inclusive value I_{mt} . This yields an 11-dimensional state vector per firm.

State Transitions I model exogenous states (hardware cost h_{jmt} and quality ξ_{jmt}) as AR(1) processes with county-specific intercepts following [Aguirregabiria and Mira \(2007\)](#). Experience evolves deterministically: $E_{jmt+1} = E_{jmt} + q_{jmt}$, with entrants having $E_{jmt} = 0$. Following [Aguirregabiria et al. \(2021\)](#), [Gowrisankaran and Rysman \(2012\)](#), and [Barwick and Pathak \(2015\)](#), I assume firms believe aggregate states (demand d_{mt} and inclusive value I_{mt}) follow AR(1) processes.

4.2 Second Stage: Learning, Exit, and Entry Parameters

I now estimate production cost parameters governing learning (θ^c), exit costs (σ_ϕ), and entry costs (σ_ω). I assume a half-yearly discount factor corresponding to an annual discount factor of 0.875 following [Gerarden \(2022\)](#) and [De Groote and Verboven \(2019\)](#).¹⁰

Value Function Approximation Estimation requires solving for firms' value functions $V_j(s_{mt})$, which enter the optimality conditions (6), (9), and (10). Since scrap values ϕ_{jmt} are i.i.d. exponential, the value function prior to realizing ϕ_{jmt} is

$$\begin{aligned} V_j(s_{mt}) &= \mathbb{E}_\phi[\pi_j(s_{mt}) + \max\{\phi_{jmt}, CV_j(s_{mt})\}] \\ &= \pi_j(s_{mt}) + p_j^x(s_{mt})\sigma_\phi + CV_j(s_{mt}) \end{aligned} \quad (11)$$

where I use the memoryless property of the exponential distribution ([Pakes et al., 2007](#)).

Following [Barwick et al. \(2025\)](#), I approximate the $V_j(s_{mt})$ using L basis functions $b_j^l(s_{mt})$:

$$V_j(s_{mt}) \simeq \sum_{l=1}^L \lambda_l b_j^l(s_{mt}) \quad CV_j(s_{mt}) \simeq \beta \sum_{l=1}^L \lambda_l \mathbb{E}[b_j^l(s_{mt+1})|s_{mt}] \quad (12)$$

⁹Candidate regressors include quadratic polynomials of state variables, their pairwise interactions, and county and period fixed effects. See Appendix C for details.

¹⁰[Ryan \(2012\)](#) assumes an annual discount factor of 0.9.

where λ_l are coefficients to be estimated. This approach avoids computationally-intensive state discretization and is feasible despite the value function's nonlinearity in parameters.¹¹

Production and Exit Cost Estimation Since production cost and exit parameters are functions of firms' value functions, which themselves depend on these target parameters, I jointly estimate both sets via non-linear GMM. I derive moments from the conditions governing optimal quantity setting and exit.

I re-express the optimal quantity condition (6) by combining the static markup (7) and dynamic markdown (8) with marginal production costs (3):

$$0 = p_{jmt} + (\Delta_{mt}^{-1})_{(j,j)} \times ms_{jmt} - mc_j(s_{mt}; \theta^c) + \beta \mathbb{E}[V_j(s_{mt+1}; \lambda) \times \Omega_j(s_{mt}, q_{mt})] \quad (13)$$

where p_{jmt} and ms_{jmt} are price and market share (from data); $(\Delta_{mt}^{-1})_{(j,j)}$ is the diagonal element of the inverse own- and cross-price derivative matrix (from first-stage demand estimation); $mc_j(s_{mt}; \theta^c)$ is marginal production cost per watt as a function of effective experience, hardware costs, and parameters θ^c via (3); $V_j(s_{mt+1}; \lambda)$ is the approximated value function; and $\Omega_j(s_{mt}, q_{mt}) = \frac{\partial}{\partial q_{jmt}} \frac{dF(s_{mt+1}|s_{mt}, q_{mt})}{dF(s_{mt+1}|s_{mt}, q_{mt})}$ captures the sensitivity of state transitions to firm j 's quantity choice.

The dynamic markdown term $\Omega_j(\cdot)$ has a closed form since the only future state variable directly affected by firm j 's current quantity is its own effective experience. Define the transition kernel for effective experience as $dG(\tilde{E}_{jmt+1}|E_t, q_t)$, where E_t and q_t are vectors of cumulative production and quantities for all incumbents. Under the experience accumulation model (4), $dG(\tilde{E}_{jmt+1}|E_t, q_t) = (E_{jmt} + q_{jmt}) + \theta_1^E \left(\sum_m \sum_{k \neq j} (E_{kmt} + q_{kmt}) \right)$, which yields:

$$\Omega_{jmt} \equiv \frac{\frac{\partial}{\partial q_{jmt}} dG(\tilde{E}_{jmt+1}|E_t, q_t)}{dG(\tilde{E}_{jmt+1}|E_t, q_t)} = \frac{1}{(E_{jmt} + q_{jmt}) + \theta_1^E \left(\sum_m \sum_{k \neq j} (E_{kmt} + q_{kmt}) \right)}$$

This closed form allows me to calculate the expectation in (13) using current states and quantities, estimated state transitions, and value function coefficients λ .

I recover productivity shocks κ_{jmt} from (13) as a function of production cost and value function parameters (θ^c, λ) . Following standard practice for production function estimation (Benkard, 2000; Olley and Pakes, 1996), I assume κ_{jmt} follows an AR(1) process:

$$\nu_{jmt}(\theta^c, \lambda, \rho) = \kappa_{jmt}(\theta^c, \lambda) - \rho \kappa_{jmt-1}(\theta^c, \lambda) \quad (14)$$

¹¹See Appendix E for basis function construction, expectation approximation, and coefficient estimation details.

To form moment conditions for estimation, I interact the innovation with instruments Z_{jmt} :

$$\mathbb{E}[Z'_{jmt}\nu_{jmt}(\theta^c, \lambda, \rho)] = 0 \quad (15)$$

where ρ is the serial correlation coefficient to be estimated. I discuss instrument selection in Section 4.3. Accounting for serial correlation prevents biased learning estimates since firms with persistent positive productivity shocks accumulate more experience while having lower marginal costs, leading to overestimated learning without this correction.

For the exit cost parameter σ_ϕ , I derive a moment from the optimal exit condition by minimizing squared differences between fitted exit probabilities (from first-stage estimation) and model-implied exit probabilities $\exp\left(-\frac{CV_j(s_{mt}; \lambda)}{\sigma_\phi}\right)$. Defining $\psi_{jmt}(\sigma_\phi, \lambda)$ as this difference, the moment condition is:

$$\mathbb{E}\left[\frac{\partial\psi_{jmt}(\sigma_\phi, \lambda)}{\partial\sigma_\phi}\psi_{jmt}(\sigma_\phi, \lambda)\right] = 0 \quad (16)$$

Stacking moments (15) and (16) allows joint estimation of target parameters $\boldsymbol{\theta} = (\theta^c, \rho, \sigma_\phi)$. Since these moments depend on value function coefficients λ , which themselves depend on (θ^c, σ_ϕ) , I follow Sweeting (2013) and iterate between solving for λ given $\boldsymbol{\theta}$ and updating $\boldsymbol{\theta}$ via two-step GMM until convergence. Specifically, at iteration i with parameter guess $\hat{\boldsymbol{\theta}}^i$: (1) solve for $\hat{\lambda}^i$ that minimizes Bellman violations; (2) update $\hat{\boldsymbol{\theta}}^{i+1}$ via GMM using stacked moments with weight matrix W ; (3) check L^1 norm convergence. I calculate standard errors via non-parametric bootstrap with 200 samples drawn by resampling entire market histories.

Entry Cost Estimation Using estimated value function coefficients, I compute entry values for potential entrants by calculating expected next-period state variables using observed aggregate states and assuming entrants draw quality and hardware cost from the empirical distribution of observed states with zero initial experience. I average over 1000 draws from estimated state transition processes.

Since the number of potential entrants is unobserved, I follow standard practice and assume it is some multiple of the median, mean, or maximum observed entrants per market over the sample period, reporting estimates for several specifications. I estimate the entry cost parameter σ_ω via maximum likelihood (which is more efficient than moment-based approaches) using the log likelihood:

$$\begin{aligned} & \log(f(\chi_{jmt}^e; \sigma_\omega)) \\ &= \sum_{j,m,t} \left[\chi_{jmt}^e \log \left(1 - \exp \left(\frac{-VE_j(s_{mt}; \hat{\lambda})}{\sigma_\omega} \right) \right) - (1 - \chi_{jmt}^e) \left(\frac{VE_j(s_{mt}; \hat{\lambda})}{\sigma_\omega} \right) \right] \end{aligned} \quad (17)$$

where χ_{jmt}^e equals 1 if potential entrant j enters market m in period t . Standard errors are computed via non-parametric bootstrap with 200 samples clustered by county.

4.3 Identification of the Dynamic Parameters

Identification of production cost parameters relies on the AR(1) productivity shock model (14). Given this model, variation in prices, quantities, and exit decisions across firms with different hardware costs and experience vectors identifies the base cost, learning exponent, effective experience, and serial correlation parameters.

The instruments Z_{jmt} in (15) include predicted own experience, predicted rival experience, lagged hardware costs, a time trend, and lagged installation-adjacent wages. I construct predicted experience instruments by estimating a reduced-form demand model regressing log market shares on rebates, solar radiation (GHI), electricity prices, installer fixed effects, year fixed effects, and commuting zone fixed effects, instrumenting inside shares with the number of firms and rival experience. Fitted values yield predicted quantities and thus predicted cumulative experience. Predicted rival experience is the sum of other firms' predicted experience in the same county-period. These instruments are valid since ν_{jmt} is serially uncorrelated by construction and orthogonal to the exogenous demand and cost shifters. The instruments provide identifying power: predicted experience correlates with marginal costs but not current innovations (being based on exogenous shifters); lagged hardware costs and wages provide exogenous cost variation; the time trend captures aggregate technological change. See Appendix D for details on construction, validity, and relevance.

Identifying spillover parameters is challenging due to potential unobservable factors (e.g., county-level demand shocks) correlated with both a firm's costs and rivals' experience, which could bias spillover estimates upward. Given these challenges, I explore counterfactual robustness by simulating scenarios with lower spillovers in Section 6. Identification of exit and entry parameters follows Hotz and Miller (1993), who show choice-specific value function differences are identified from observed choice probabilities.

5 Model Results

I present the main model estimates in this section. I begin by presenting the first stage estimates (Section 5.1) before turning to my estimates of the dynamic model primitives (Section 5.2). I then discuss several robustness checks (Section 5.3).

Table 1. Estimated Demand System Parameters

	Param	NL-1 (1)		RCNL-1 (2)		RCNL-2 (3)	
		Est	SE	Est	SE	Est	SE
Price/Income	α_p	-8.513	1.063	-7.718	1.187	-14.69	3.661
Nesting Parameter	η	0.636	0.026	0.632	0.033	0.630	0.071
Total Experience	α_1					-10.835	13.061
Solar Irradiance	α_2					1.270	0.615
$\log(\text{Income}) \times \text{Constant}$	σ_0					2.917	0.816
$\log(\text{Income}) \times \text{Total Experience}$	σ_1					2.192	2.437
$\log(\text{Income}) \times \text{Solar Irradiance}$	σ_2					-0.282	0.132
Firm, Year, CZ FE		Yes		Yes		Yes	
Income Distribution				Yes		Yes	
Median Own Price Elast.		-3.76		-3.46		-2.80	
Median Outside Diversion		36.70%		37.06%		37.23%	

Notes: Estimation follows the procedure outlined in Section 4.1. There are 10,247 observations at the firm-county-half-year level. The nested logit model NL-1 divides price by county-year-half mean income whereas the random coefficients nested logit models RCNL-1 and RCNL-2 use the full sample of incomes drawn from the ACS PUMS to estimate price sensitivities and other income interaction terms. Total Experience is a firm's cumulative quantity of installations in a county and Solar Irradiance is a measure of observed solar radiation in a given county-year. Standard errors are clustered at the county-level.

5.1 First Stage Estimates

Consumer Demand Table 1 presents parameter estimates and standard errors clustered by county for the consumer demand system. Column (1) estimates a nested logit (NL) model that removes individual-level heterogeneity in taste parameters. Columns (2) and (3) correspond to different versions of the full RCNL model outlined in Section 3.1. To ensure comparability across models, I divide price by county-quarter mean income in the NL specification. All specifications include firm, year, and commuting zone fixed effects.

The price coefficients (α_p) are precisely estimated with the expected sign across all specifications. Median own-price elasticities range from -2.80 to -3.76 and are smaller when including observable firm attributes and income interactions. Market price elasticities are substantially lower than own-price elasticities, indicating substitution primarily occurs across installers rather than on the extensive margin. The nesting parameter (η) is large and precisely estimated across all specifications. Diversion ratios imply that not installing is the second-best choice for around 37% of consumers.

The coefficients on observable firm attributes are imprecisely estimated in the RCNL model with income interactions. Higher income households appear to derive greater utility from installers with more cumulative experience and from high solar irradiance areas, while lower income households appear to derive lower utility from high solar irradiance. These

results should be interpreted with caution. Given that I do not endogenize firms' product attribute decisions and given the desirable substitution patterns of the RCNL model, I use column (2) estimates in the second stage and all counterfactual simulations. I plot the distribution of own price elasticities of demand estimated using my preferred demand model and compare these to estimated elasticities from the literature in Appendix Figure G3.¹² To test the robustness of my assumption that a static demand model offers a reasonable approximation to consumer behavior, I develop and estimate a dynamic demand model in Appendix B. Estimates from this dynamic demand model are similar in magnitude—the median elasticity is -2.07—and demonstrate a time pattern consistent with analogous static estimates as shown in Figure B1.

Exit Policy Function and State Transitions As outlined in Section 4.1, I select candidate regressors for the logit regression of the discrete exit decision via LASSO. Of 242 candidate regressors, this process selects around 80 non-zero regressors with a deviance of 12.34%. As shown in Appendix Figure C2, fitted exit probabilities are larger for exiting incumbents than continuing ones: 14.88% versus 7.72%.

I report AR(1) transition process estimates for aggregate state variables (demand state, inclusive value, and county-quarter average price) in Appendix Table G1 and for firm-level state variables (quality, hardware cost, and firm price) in Appendix Table G2. For each state variable, I report specifications with and without county-specific intercepts. All estimated transition processes are stationary.

5.2 Second Stage Estimates

Production and Exit Cost Estimates Table 2 reports results from joint estimation of production and exit cost parameters. I allow for three distinct spillover models across rival firms, with corresponding estimates in each column.

The learning parameter, γ , is negative and precisely estimated across all specifications. The learning exponent estimates in columns (1) through (3) imply a 1% increase in effective experience decreases marginal installation costs by 0.21 to 0.36%. The “Spence coefficient,” which describes the proportional cost reduction from doubling effective experience (Spence, 1981), ranges from 0.134 to 0.223, implying meaningful learning-by-doing in installation-specific costs. While more modest than manufacturing learning curves, these estimates are

¹²Overall, my elasticity estimates are similar in magnitude to those found elsewhere in the literature, which range from -0.65 (Gillingham and Tsvetanov, 2019) to -6.6 De Groote and Verboven (2019) for analogous (i.e., static) estimates. Most similar to my preferred estimates is Dorsey (2024), who estimates a mean own price elasticity for installations of between -2.43 and -2.26.

Table 2. Estimated Production and Exit Cost Parameters

Param	(1)		(2)		(3)		(4)		
	Est	SE	Est	SE	Est	SE	Est	SE	
Production Cost Parameters									
Learning Exponent	γ	-0.363	0.150	-0.207	0.176	-0.243	0.130	-0.267	0.135
Hardware Cost	τ	-0.654	0.050	-0.234	0.131	-0.471	0.087	-0.579	0.105
Base Cost	c_0	0.643	0.256	0.610	0.207	0.496	0.233	0.576	0.235
Serial Correlation	ρ	0.420	0.007	0.470	0.019	0.436	0.009	0.427	0.015
Effective Experience Parameters									
Rival Experience:									
In-market	θ_1^E	0.735	0.253			0.710	0.240	0.653	0.302
Same Manuf.	θ_2^E			0.382	0.206	0.282	0.185	0.070	0.183
Forgetting Parameter	δ						1.048	0.377	
Exit Parameter									
Mean Scrap Value	σ_ϕ	49.698	3.062	48.431	2.792	49.745	2.231	52.296	2.354
Market, Firm, Time FE		Yes		Yes		Yes		Yes	
Spence Coef. $(1 - 2^\gamma)$		0.223		0.134		0.155		0.169	

Notes: Estimation follows the procedure outlined in Section 4.2. There are 7,351 observations at the installer-county-half-year level. I normalize experience variables by the industry total experience level in the first half year of the sample (H1 2008). All effective experience parameters can be interpreted as marginal experience contributions relative to a firm’s own experience. The “forgetting parameter,” δ , describes the rate of learning depreciation from one period to another. The mean scrap value parameter is measured in 100,000 2013 USD. The “Spence Coefficient” describes the proportional reduction in cost from a doubling of effective experience. Standard errors are calculated using the Bayesian Bootstrap with bootstrap weights clustered by county (Rubin, 1981). Bootstrap weights for each county are drawn according to a Dirichlet distribution with $\alpha = 1$ across 200 bootstrap samples.

consistent with a service-intensive process where learning occurs through improved coordination and permitting efficiency.


I estimate nontrivial serial correlation in firms’ productivity shocks, κ_{jmt} . The serial correlation parameter, ρ , ranges from 0.42 to 0.47 across columns (1) through (4). Though moderate, these estimates suggest ignoring serial correlation would bias learning estimates. Firms with serially correlated positive productivity shocks likely have greater experience and relatively low marginal costs, which could be misattributed to learning-by-doing if not properly modeled. All specifications account for serial correlation.

The rival experience parameters describe the marginal contribution of other firms’ cumulative production to a firm’s effective experience. I estimate three parameterizations of effective experience, each implying a distinct spillover model. Rival experience parameters represent marginal contributions relative to own experience.


The first parameterization groups all rival experience together. Column (1) reports baseline estimates: a 1 unit increase in in-market rival experience generates 74% of the learning benefits from a 1 unit increase in own experience. Learning spillovers from rivals within the

Figure 4. Estimated Marginal Costs and Learning Contributions

(a) Mean Marginal Costs



(b) Cumulative Learning Contributions



Notes: Figure 4a compares the average non-hardware costs implied by the dynamic model estimates reported in column (2) of Table 2 with comparable, publicly-available estimates from the National Renewable Energy Laboratory, or NREL (Fu et al., 2016). The figure also shows static non-hardware costs implied by the model. Static non-hardware costs are calculated from the standard, static quantity-setting first order condition using observed prices and estimated price elasticities. The shaded area shows the 95% bootstrap confidence interval for the dynamic model estimates. Figure 4b shows the cumulative contribution of own and rivals' experience to marginal costs over time using estimates from column (2) of Table 2. Cumulative experience-based cost reductions are calculated at the 10th, 50th, and 90th percentiles of observed experience components in each period. Costs are in 2013 dollars per watt.

same geographic market are substantial, though somewhat smaller than individual learning.

To identify spillover mechanisms, I estimate two alternative parameterizations. The first allows for differential spillovers based on whether firms install modules from the same manufacturer, consistent with manufacturer-facilitated learning through training or module design improvements. Column (2) reports estimates: rivals installing the same manufacturer's modules contribute 0.38 versus lower contributions from other manufacturers' modules, though the difference is not statistically significant. This provides suggestive but inconclusive evidence for manufacturer-facilitated spillovers.

The second parameterization allows for differential spillovers based on both geographic market and manufacturer. Column (3) shows in-market rivals contribute more to learning (0.71) than same-manufacturer rivals (0.28). This is consistent with geographic proximity-based spillovers driven by worker movement between firms, visibility of rivals' practices, or other local diffusion mechanisms.¹³ This parameterization is also consistent with regulator

¹³I explore this by estimating the relationship between PV installation-related employment, wages, and

learning: firms in the same county face similar permitting requirements and utility rebate processes, so regulatory learning would show up as greater benefits from in-market rivals.¹⁴

While experience parameters describe marginal contributions, the *total* contribution depends on both parameter estimates and empirical experience distributions. Figure 4 plots average cumulative contributions to marginal installation costs over the full sample period. Despite smaller marginal spillover effects relative to own experience, learning spillovers drive the bulk of estimated cost reductions. The average reduction from own experience is modest, while spillovers from rivals' experience (using column 2 estimates) generate more substantial reductions. The distribution of experience levels means even smaller marginal spillover effects have meaningful aggregate impacts on industry-wide cost reductions.

I estimate a version allowing experience-based cost reductions to depreciate over time. “Forgetting” models are relevant in boom-bust settings like aircraft, shipbuilding, and oil extraction (Benkard, 2000; Kellogg, 2011; Levitt et al., 2013; Thompson, 2007). While knowledge depreciation is unlikely to be major in this persistent-growth setting, worker departures could cause incomplete knowledge retention. I test this with a perpetual-inventory process where firm j ’s experience in market m in period t is $\tilde{E}_{jmt} = \delta(\tilde{E}_{jmt-1} + f(q_{mt-1}))$, where $f(\cdot)$ is the function that maps lagged quantities into effective experience.¹⁵ The parameter δ defines knowledge retention. Column (4) shows minimal forgetting, with a quarterly retention parameter of 1.05 (not statistically different from 1.0), suggesting knowledge depreciation is not significant in California’s persistent-growth solar installation industry.

The learning-by-doing magnitude implied by Table 2 estimates aligns with existing estimates of installation-specific costs. Figure 4 compares average non-hardware costs implied by column (2) estimates with comparable estimates from the National Renewable Energy Laboratory (NREL). Fu et al. (2016) construct these from installation data, corporate filings, and engineering studies. I report the sum of their installation labor, permitting, inspection, and installation costs per watt.¹⁶ My model estimates match both the magnitude and rate of learning implied by Fu et al. (2016)’s estimates within my estimation sample period.

Figure 4 also shows average marginal non-hardware costs assuming firms are completely installations within a county using data from the Quarterly Census of Employment and Wages. Figure G4 shows strong, positive relationships between employment/wages and installations, suggesting worker movement in the PV installation labor pool.

¹⁴I explore this using CSI rebate processing times across IOUs. Figure G5 shows suggestive evidence that IOU rebate processing times decreased over the sample period, which would reduce installation-specific costs. Figure G6 shows permitting times also decreased in San Diego County.

¹⁵This functional form assumes incomplete retention of prior production knowledge, consistent with forgetting driven by worker turnover.

¹⁶Fu et al. (2016) also include installer overhead and net profit in soft costs. I exclude these as markups are not components of installation-specific costs and overhead includes fixed operating costs not relevant to my estimates.

static in quantity-setting. I calculate these from the static first-order condition using observed prices, hardware costs, and estimated price elasticities. Since I only estimate demand in-sample, I cannot estimate static costs out-of-sample as I can with the dynamic model using observed experience levels. While my dynamic model matches Fu et al. (2016)'s estimates, static estimates are substantially larger than both dynamic estimates and the NREL benchmark. This emphasizes the importance of accounting for learning-by-doing and dynamic incentives: without these, I would overestimate non-hardware costs.

The remaining parameters are precisely estimated and consistent across specifications. The mean scrap value, σ_ϕ , is \$4.8–5.2 million. The hardware cost pass-through parameter, τ , ranges from -0.23 to -0.65, suggesting incomplete pass-through. I normalize experience terms by total industry experience in H1 2008 for readability and numerical stability. The base cost parameter, c_0 , ranges from \$0.50 to \$0.64 per watt, interpretable as marginal installation cost when firm effective experience equals industry total experience in H1 2008.

Entry Cost Estimate Table 3 reports mean entry cost estimates using different data-driven approaches to defining potential entrants in each market. I assume the number of potential entrants is one or two times the median, mean, or maximum observed entrants per market over the sample period. This data-driven approach is common, with most studies using multiples of the maximum observed entrants (Barwick and Pathak, 2015; Barwick et al., 2025; Seim, 2006). Unsurprisingly, estimated entry costs increase with the pool of potential entrants, ranging from \$18.0 to \$44.7 million (in units of \$100,000). While large relative to mean scrap values, these are unconditional mean entry costs. Conditional on entering, average entry costs range from \$5.7 to \$6.2 million, somewhat larger than but broadly consistent with existing estimates and available information on publicly-traded installers.¹⁷

5.3 Robustness

I explore robustness of my main estimates in several ways. First, I test sensitivity to alternative quarterly discount factors, β . My main estimates assume a quarterly discount factor corresponding to an annual discount factor of 0.875, equal to Gerarden (2022)'s for PV manufacturers and similar to De Groote and Verboven (2019)'s for PV-adopting households. I re-estimate using annual discount factors of 0.9 and 0.8, both found in the literature (Igami,

¹⁷Feldman et al. (2013) estimate upfront costs for developing, constructing, and arranging third-party-financed residential PV systems, finding fixed business expenses of \$600,000/year for a representative firm in 2012, or roughly \$5 million in perpetuity assuming an annual discount factor of 87.5%. SolarCity Corp. had a market capitalization of \$5.6 billion at end of 2013 Q4, having installed roughly 75 MW of residential capacity that year. Scaling by average annual installed capacity in my data (0.35 MW) yields roughly \$2 million valuation for county-level operations of an average firm.

Table 3. Estimated Entry Cost Parameter

Parameter	(1)	(2)	(3)	(4)	
Mean Entry Cost	σ_ω	179.770 (5.921)	186.381 (3.395)	185.970 (17.926)	446.580 (35.445)
Potential Entrant Def.		$2 \times \text{median}(\bar{N}_{mt})$	$2 \times \text{mean}(\bar{N}_{mt})$	$1 \times \max(\bar{N}_{mt})$	$2 \times \max(\bar{N}_{mt})$
N		3,179	3,250	3,216	6,432
N^e		260	268	268	536
$\bar{\omega}_{jmt} \text{entry}$		57.457	57.731	57.715	62.127

Notes: Estimation follows the procedure outlined in Section 4.2. The entry cost parameter is measured in 100,000 2013 USD. Each column corresponds to a different approach to defining the market-specific, time-invariant number of potential entrants based on observed quantities of entrants, \bar{N}_{mt} : column (1) uses twice the median of \bar{N}_{mt} , column (2) uses twice the mean of \bar{N}_{mt} , column (3) uses the maximum observed value of \bar{N}_{mt} , and column (4) uses twice the maximum of \bar{N}_{mt} . N is the total number of observations used in estimation and N^e is the number of potential entrants per year across all counties based on the assumed potential entrant definition. $\bar{\omega}_{jmt}|\text{entry}$ is the mean entry cost conditional on a firm choosing to enter. Bootstrapped standard errors clustered by county using 200 replications are reported in parentheses.

2017; Ryan, 2012). The latter also serves as a proxy for dynamic regulatory environments, rendering future payoffs less relevant for current decisions, testing my assumption that firms perceive policy changes as permanent.

Table G3 reports estimates using alternative discount factors alongside baseline estimates. Results are qualitatively consistent, though implied costs increase with the discount factor.¹⁸ This is consistent with the model: higher discount factors imply higher expected values from future operations, requiring higher cost estimates to rationalize observed patterns. Qualitative consistency suggests assuming stationary policy beliefs is reasonable.

Given that I define experience as cumulative production, some estimated learning may reflect static scale economies. Separately identifying learning-by-doing and scale economies is challenging in learning curve estimation. I test sensitivity by estimating a version with a static, contemporaneous firm size measure as a state variable in the value function approximation. Table G4 compares estimates with and without firm size, finding qualitatively consistent results. Similar to Benkard (2000), estimated learning rates increase slightly when including scale.

6 Counterfactual Analysis

Having recovered estimates of the main model parameters, I can simulate market outcomes under counterfactual policy environments, which requires a method for solving for the

¹⁸Mean scrap value and base cost parameters increase with β as these are positively related to firm costs. Effective experience and learning exponent parameters decrease: for given experience levels, lower parameter values imply higher costs.

model's equilibrium. I begin this section by describing my approach to solving the model (Section 6.1). I then compare the fit of model-predicted outcomes under the baseline policy environment with observed data (Section 6.2). I next quantify the importance of learning-by-doing in determining equilibrium market structure, prices, and quantities (Section 6.3). Finally, I discuss results from three sets of counterfactual policy scenarios that interact with learning economies in different ways: consumer subsidies that stimulate demand and thus learning indirectly (Section 6.4), supply-side entry subsidies that reduce barriers to firm entry (Section 6.5), and alternative climate policies including a carbon tax (Section 6.6).

Results from these counterfactual simulations provide three main findings. First, learning-by-doing plays a central role in shaping market outcomes: removing installer learning economies while maintaining baseline subsidies reduces cumulative installations by 10%, increases equilibrium prices throughout the sample period, and substantially reduces the number of active firms. Second, I find that the CSI contributed to growth in the installation industry through its interaction with learning dynamics: the CSI increases the number of solar PV installations by 4% and increases the number of operating installers by roughly 9% relative to a world with no CSI. The CSI's decreasing rebate structure—which exploits learning-driven cost reductions—produces higher welfare than alternative rebate designs that do not account for learning economies. Third, I find that supply-side entry subsidies can more effectively leverage learning economies than consumer subsidies: replacing the CSI with entry subsidies of varying sizes results in substantially more installations and firm entry by directly stimulating competition and production, though at significant fiscal cost.


6.1 Counterfactual Solution Method

My approach to solve for the model's equilibrium builds on Sweeting (2013), adapting parametric policy iteration (Benitez-Silva et al., 2000) to allow for value function approximation. Appendix F provides detailed implementation. Solving the model involves two steps: first, solving for the new Bellman equation, policy functions, and product market equilibrium in a period; second, simulating the industry forward one period. I initiate this procedure at the observed data in the first half-year of 2008 and iterate through the last half-year of 2013.

I implement the first step via a fixed point algorithm that produces conditional exit probabilities and value function approximating coefficients, from which I calculate conditional entry probabilities. The second step draws from these probabilities to implement discrete exit and entry decisions, then draws from estimated state transition processes for the next period's states.

This process yields a single industry path. Since I take single draws from conditional probabilities and state transitions each period, I repeat this process 60 times for each counterfac-

Figure 5. Baseline Model Fit



Notes: This figure shows equilibrium prices, cumulative quantities, exit probabilities, and entry probabilities, comparing model-implied values to observed data at the industry-level from 2008 to 2013. Model-implied values are averaged across 60 distinct, forward-simulated industry paths with baseline policies in place. Prices are in 2013 dollars per watt and cumulative quantities are in megawatts.

tual scenario and average outcomes across all runs. The model generates policy-relevant outcomes including quantities, prices, marginal costs, entry, and exit, allowing me to calculate firms' profits and consumer surplus. I calculate changes in environmental damages by combining model-predicted quantities with geographically-differentiated estimates of marginal environmental benefits from [Sexton et al. \(2021\)](#). See Appendix F for details.

6.2 Model Fit under Baseline Policies

Before turning to counterfactual simulations, I verify the model fits observed data under baseline policies. Figure 5 shows the model accurately matches prices (declining from \$8 to \$5 per watt), cumulative installed capacity, and exit probabilities. Entry is matched well overall, though slightly under-predicted early and over-predicted late in the sample.

6.3 The Role of Learning-by-doing

To quantify the importance of learning-by-doing, I simulate industry equilibrium removing installer learning economies while maintaining baseline subsidies. This isolates the role of experience-driven cost reductions in determining prices, quantities, and market structure.

Figure 6 shows learning-by-doing plays a substantial role in market outcomes. Without installer learning, cumulative installed capacity is significantly lower throughout the sample period. By 2013, the gap between learning and no-learning scenarios is around 48 MW—10% of total installed base with learning—indicating experience-driven cost reductions are critical for market expansion.

The middle panel reveals the mechanism: equilibrium prices are consistently higher without learning-by-doing. While prices decline in both scenarios due to falling hardware costs, the relative price premium in the no-learning scenario remains high and grows slightly over time. Though the absolute difference narrows as both prices fall, the percentage difference increases, indicating learning economies become more important as the industry matures. This persistent and growing relative price premium dampens demand and limits adoption. Market structure also differs substantially. The right panel shows the number of active firms is consistently lower without learning-by-doing, with the gap growing over time. This reflects two mechanisms. First, without learning economies, returns to production are purely static, reducing firms' incentives to enter and expand output. Second, lower overall market size supports fewer firms in equilibrium.


The welfare implications are large. Total welfare in the no-learning scenario is \$49 million lower than baseline, driven primarily by \$287 million in lost consumer surplus, \$263 million in lost product market profits, and \$20 million in reduced environmental benefits. These losses are partly offset by lower subsidy outlays, lower entry costs, and greater scrap values.

These results demonstrate that learning-by-doing generates substantial benefits for consumers and the industry. Cost reductions from accumulated experience enable lower prices that expand the market while simultaneously supporting entry and growth in active firms. This highlights the central role learning economies play in the solar installation industry and motivates analyzing how different policies interact with these learning dynamics.

6.4 Demand Subsidy Counterfactuals

Having established the importance of learning-by-doing, I now examine how consumer subsidies interact with learning dynamics. Consumer subsidies like the CSI stimulate demand directly, inducing learning-by-doing indirectly through increased production. The CSI's declining rebate structure may exploit learning-driven cost reductions, making the temporal

Figure 6. Equilibrium with and without Learning-by-doing



Notes: This figure shows equilibrium cumulative quantities, prices, and total incumbent firms with and without installer learning-by-doing with baseline subsidy policies in place. The baseline subsidy policies correspond to those outlined in Section 1.2: full California Solar Initiative (CSI) subsidies, 30% federal investment tax credit (ITC), and net-metering. Each counterfactual removes the CSI subsidies and replaces them with the indicated demand subsidy. Model-implied values are averaged across 60 distinct, forward-simulated industry paths with baseline policies in place. Prices are in 2013 dollars per watt and cumulative quantities are in megawatts.


structure of subsidies potentially important for welfare.

I simulate industry equilibrium under three demand subsidy counterfactuals: complete removal of the CSI, a flat rebate equal to the quantity-weighted average CSI rebate, and an increasing rebate schedule that inverts the baseline CSI's declining steps. All scenarios maintain the federal ITC and net metering while varying only the state-level CSI design.

Figure 7 shows removing the CSI substantially reduces market outcomes across all dimensions. By 2013, cumulative installed capacity falls by 42 MW (approximately 8%) while active firms decline by as much as 4% at peak relative to baseline. The middle panel reveals the mechanism: equilibrium prices rise throughout the period without CSI subsidies, dampening demand and limiting industry growth. This price increase reflects both the direct effect of removing consumer rebates and the indirect effect of reduced learning-by-doing from lower cumulative production.

The alternative CSI designs yield more nuanced results. The flat rebate produces outcomes similar to the baseline declining rebate, with only modest differences in quantities, prices, and market structure, suggesting the temporal structure of rebates matters less than their overall generosity. The increasing rebate schedule, motivated by Langer and Lemoine (2022)'s finding that rising subsidies enable optimal price discrimination, performs worse than both baseline and flat rebate designs. This indicates that with learning-by-doing, capturing cost reductions through declining subsidies may be more important than price discrimination benefits.

Figure 7. Changes from Baseline under Demand Subsidy Counterfactuals



Notes: This figure shows changes in equilibrium cumulative quantities, prices, and total incumbent firms relative to the baseline scenario of existing subsidy policies under three alternative demand subsidy counterfactuals. The baseline subsidy policies correspond to those outlined in Section 1.2: full California Solar Initiative (CSI) subsidies, 30% federal investment tax credit (ITC), and net-metering. Each counterfactual removes the CSI subsidies and replaces them with the indicated demand subsidy. Model-implied values are averaged across 60 distinct, forward-simulated industry paths with baseline policies in place. Prices are in 2013 dollars per watt and cumulative quantities are in megawatts.

As shown in Table 4, removing the CSI reduces total welfare by \$328.7 million relative to baseline, driven primarily by \$220.0 million in lost consumer surplus and \$227.4 million in reduced static profits. The flat rebate design increases welfare by \$230.9 million relative to baseline, suggesting smoothing rebates over time could improve outcomes. However, the increasing rebate substantially reduces welfare by \$1,108.7 million compared to baseline. These results indicate that while the temporal structure of rebates matters, declining or flat schedules that exploit learning-by-doing cost reductions outperform increasing schedules.

6.5 Supply Subsidy Counterfactuals

While consumer subsidies leverage learning indirectly through demand stimulation, supply-side industrial policies can potentially accelerate learning more directly. Entry subsidies reduce barriers to market participation, stimulating firm entry and intensifying competition. This increased competition expands production and triggers learning-by-doing that benefits the entire industry through spillovers. Since learning economies generate dynamic benefits from cumulative production, policies that directly increase the number of active producers may be more effective than demand subsidies at exploiting these learning dynamics. I simulate counterfactual scenarios removing the CSI and replacing it with entry subsidies of varying sizes: one-quarter, one-half, and three-quarters of the estimated mean conditional entry cost.

Figure 8 shows entry subsidies substantially alter market outcomes relative to baseline

Table 4. Estimated Changes in Welfare under Counterfactual Policy Scenarios


Scenario	Welfare Components (\$M)						
	ΔCS	ΔEB	$\Delta Profit$	$\Delta \phi$	$-\Delta \omega$	$-\Delta G$	$\Delta Total$
<i>Demand Subsidy Counterfactuals:</i>							
Remove CSI	-220.0	-17.6	-227.4	-199.0	76.3	258.9	-328.7
Flat CSI	119.4	12.1	112.7	138.8	22.1	-174.2	230.9
Increasing CSI	661.9	64.7	609.3	-11.6	-1,554.3	-878.7	-1,108.7
<i>Entry Subsidies:</i>							
1/4*Mean Entry Cost	-50.7	-8.2	-106.0	1,965.7	197.4	-2,178.1	-179.9
1/2*Mean Entry Cost	180.0	4.7	51.2	4,519.8	-372.3	-5,699.3	-1,316.0
3/4*Mean Entry Cost	521.7	23.7	284.7	8,141.4	781.4	-10,866.6	-1,113.7
<i>Alternative Climate Policies:</i>							
Carbon Tax (\$30/ton)	15.6	2.9	10.7	269.4	337.5	-43.7	592.4
Remove ITC	-958.6	-73.0	-906.9	-84.2	2,004.7	975.2	957.1
10% ITC	-664.1	-52.0	-634.4	-84.6	1,454.8	712.2	732.0
26% ITC	-177.4	-13.6	-162.3	234.2	127.9	174.9	183.6

Notes: This table reports model-predicted changes in welfare components relative to the baseline scenario of existing consumer subsidy policies under three sets of counterfactuals: alternative demand subsidy designs, varying levels of entry subsidies, and alternative climate policies. The baseline subsidy policies correspond to those outlined in Section 1.2: full CSI subsidies, 30% federal ITC, and net-metering. The alternative CSI designs, entry subsidies, and carbon tax all remove the full CSI subsidy program. The three different ITC counterfactuals keep the full CSI subsidy program in-place. Each counterfactual welfare component reported in the table represents an average across 60 forward-simulated industry paths under the given counterfactual. ΔCS is the change in consumer surplus, ΔEB is the change in environmental benefits, $\Delta Profit$ is the change in firm flow profits, $\Delta \phi$ is the change in scrap values, and $\Delta \omega$ is the change in entry costs. ΔG is the change in government expenditures and $\Delta Total$ is the sum across welfare components assuming a marginal cost of public funds of 1. All values are reported in million 2013 USD.

CSI policies. The right panel demonstrates entry subsidies dramatically increase active firms, with effects scaling with subsidy magnitude. By 2013, the largest entry subsidy (three-quarters of mean entry cost) raises active firms by approximately 62% relative to baseline. This expanded competition drives the price effects shown in the middle panel: equilibrium prices fall substantially under entry subsidies as increased firm entry intensifies competition. The left panel reveals this combination of lower prices and expanded choice sets significantly boosts cumulative installed capacity, with effects again scaling with subsidy size.

Table 4 reveals a more nuanced welfare picture. While entry subsidies reduce total welfare relative to baseline CSI by \$179.9 million (smallest subsidy) to \$1,316.0 million (middle subsidy), this primarily reflects substantial fiscal costs. Entry subsidies generate significant benefits through increased consumer surplus (\$180.0 million for the middle subsidy) and reduced scrap values from lower exit (\$4,519.8 million for the middle subsidy). However, these gains are offset by large negative changes in entry costs and government expenditures, highlighting the fiscal burden of subsidizing firm entry.

Figure 8. Changes from Baseline under Supply Counterfactuals



Notes: This figure shows changes in equilibrium cumulative quantities, prices, and total incumbent firms relative to the baseline scenario of existing subsidy policies under three alternative demand subsidy counterfactuals. The baseline subsidy policies correspond to those outlined in Section 1.2: full California Solar Initiative (CSI) subsidies, 30% federal investment tax credit (ITC), and net-metering. Each counterfactual removes the CSI subsidies and replaces them with the indicated demand subsidy. Model-implied values are averaged across 60 distinct, forward-simulated industry paths with baseline policies in place. Prices are in 2013 dollars per watt and cumulative quantities are in megawatts.

The mechanisms through which entry subsidies affect outcomes differ fundamentally from consumer demand subsidies. Entry subsidies directly reduce entry barriers, stimulating firm entry and intensifying product market competition. This increased competition reduces markups and prices, expanding demand and triggering learning-by-doing that further lowers costs. In contrast, the CSI's consumer rebates stimulate demand directly, which then induces entry indirectly through higher firm values. While both approaches leverage learning economies, entry subsidies operate through the supply side whereas consumer subsidies work through the demand side.

These results suggest that while entry subsidies can effectively expand market size and competition, their substantial fiscal costs (ranging from approximately \$2.1 billion to \$8.4 billion in government expenditures) must be weighed against the CSI's more modest fiscal burden of \$2.2 billion. Moreover, the CSI was funded by California electricity ratepayers rather than general tax revenue, representing a transfer between households rather than a true fiscal outlay. The political economy of directly subsidizing firm entry versus consumer adoption also differs substantially, with entry subsidies potentially facing greater resistance as perceived industry windfalls.

6.6 Alternative Technology Policies

To provide broader policy context, I examine two additional counterfactuals: a carbon tax that removes the CSI and instead increases retail electricity prices, and variations in the

federal investment tax credit (ITC) level while maintaining the CSI.

I simulate a counterfactual carbon tax of \$30/ton of CO₂ emissions, roughly equal to California’s cap-and-trade allowance prices from 2020-2023. I estimate the tax’s impact on retail electricity prices using California’s gas-fired generation emissions rates, translating this into equivalent PV adoption incentives through net-metering benefits (approximately \$0.61/watt). This removes the CSI’s direct rebates but indirectly subsidizes solar by raising electricity prices.¹⁹

Table 4 shows the carbon tax increases total welfare by \$592.4 million relative to the CSI baseline, primarily reflecting reduced government expenditures (\$43.7 million savings) and increased scrap values (\$269.4 million) from lower exit rates. However, this comparison is incomplete because the carbon tax generates substantial benefits beyond the solar industry—including reduced emissions across all sectors—which my model does not capture. Thus, while the results suggest carbon pricing can promote solar adoption at levels comparable to technology-specific subsidies, the broader climate benefits make direct welfare comparisons inappropriate.

I also examine varying the federal ITC while maintaining the CSI. I simulate ITC rates of 0%, 10%, and 26% (compared to the 30% baseline), relevant given the ITC’s extension to 30% through 2032 under the Inflation Reduction Act of 2022.²⁰ Table 4 shows eliminating the ITC increases welfare by \$957.1 million, primarily through savings in government expenditures (\$975.2 million) that outweigh consumer surplus losses (\$958.6 million). The 10% and 26% ITC rates show similar patterns with smaller magnitudes. These results reflect the ITC’s high fiscal cost relative to the CSI: removing the federal subsidy saves nearly \$1 billion despite reducing adoption, whereas the CSI costs only \$258.9 million. However, this welfare comparison ignores the broader environmental benefits of increased solar adoption across all states, making it an incomplete evaluation of the ITC’s social value.

7 Conclusion

Learning-by-doing and knowledge spillovers fundamentally shape industry equilibrium when cumulative production experience reduces costs and knowledge transfers across firms. This paper estimates these learning mechanisms in California’s residential solar PV installation industry and analyzes how they determine equilibrium outcomes. The results demonstrate that learning dynamics are central to observed market structure: removing learning economies reduces installations by 10% and substantially contracts the number of active firms. The

¹⁹See Appendix F for calculation details. The cap-and-trade program began in 2013; this counterfactual removes the CSI and implements carbon pricing throughout 2008-2013.

²⁰Prior to the Inflation Reduction Act, the 30% federal ITC had been scaled back to 26% starting in 2020.

magnitude and spillover structure of learning determine how policies affect equilibrium, with consumer subsidies leveraging learning to expand adoption and industry size.

The estimates reveal substantial learning-by-doing: a 1% increase in experience reduces installation costs by 0.21-0.36%. Knowledge spills over substantially across firms, with in-market rivals' experience generating 74% of the learning benefit from own experience. These spillovers amplify industry-wide learning but create a tension: while spillovers magnify cost reductions from aggregate production, they weaken individual firms' incentives to expand output. The equilibrium that emerges reflects this tradeoff, with learning and spillovers jointly determining prices, quantities, entry, and exit.

Policy effectiveness depends critically on these learning mechanisms. Consumer subsidies stimulate production, triggering learning that reduces costs for all firms through spillovers. This creates feedback between learning and market structure: subsidies expand production, reducing costs through learning, which facilitates entry and further production. The magnitude of this effect depends on spillover rates—with large spillovers, subsidies generate substantial industry expansion; with small spillovers, effects are muted. Supply-side entry subsidies can more effectively leverage learning by directly stimulating competition and production, though at substantial fiscal cost: while entry subsidies dramatically increase active firms (by up to 62%) and installations, they reduce total welfare by \$180 million to \$1.3 billion relative to consumer subsidies, primarily reflecting fiscal costs of \$2.1 billion to \$8.4 billion compared to \$2.2 billion for the CSI.

These findings contribute to understanding how learning mechanisms affect industries with experience-driven cost reductions and inform policy design in settings where production generates learning externalities. The results emphasize that accounting for learning dynamics—both magnitude and spillover structure—is necessary for predicting how policies affect equilibrium outcomes in industries characterized by substantial learning economies.

References

Aguirregabiria, Victor, Allan Collard-Wexler, and Stephen P. Ryan. 2021. “Dynamic games in empirical industrial organization.” In *Handbook of Industrial Organization*. Vol. 4 of *Handbook of Industrial Organization, Volume 4*, , ed. Kate Ho, Ali Hortaçsu and Alessandro Lizzeri, 225–343. Elsevier.

Aguirregabiria, Victor, and Pedro Mira. 2007. “Sequential Estimation of Dynamic Discrete Games.” *Econometrica*, 75(1): 1–53.

Arrow, Kenneth J. 1962. “The Economic Implications of Learning by Doing.” *The Review of Economic Studies*, 29(3): 155–173.

Bajari, Patrick, C. Lanier Benkard, and Jonathan Levin. 2007. “Estimating Dynamic Models of Imperfect Competition.” *Econometrica*, 75(5): 1331–1370.

Barbose, Galen, Naïm Darghouth, Eric O’Shaughnessy, and Sydney Forrester. 2022. “Tracking the Sun: Pricing and Design Trends for Distributed Photovoltaic Systems in the United States, 2022 Edition.” Lawrence Berkeley National Laboratory.

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

Barwick, Panle Jia, Myrto Kalouptsidi, and Nahim Bin Zahur. 2025. “Industrial Policy Implementation: Empirical Evidence from China’s Shipbuilding Industry.” *The Review of Economic Studies*, 92(6): 3611–3648.

Benitez-Silva, Hugo, John Rust, Gunter Hitsch, Giorgio Pauletto, and George Hall. 2000. “A Comparison Of Discrete And Parametric Methods For Continuous-State Dynamic Programming Problems.” *Computing in Economics and Finance 2000*.

Benkard, C. Lanier. 2000. “Learning and Forgetting: The Dynamics of Aircraft Production.” *American Economic Review*, 90(4): 1034–1054.

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

Benkard, C. Lanier, Przemyslaw Jeziorski, and Gabriel Y. Weintraub. 2015. “Oblivious equilibrium for concentrated industries.” *The RAND Journal of Economics*, 46(4): 671–708.

Berry, Steve, and Ariel Pakes. 2000. “Estimation from the Optimality Conditions for Dynamic Controls.” *Working Paper*.

Berry, Steven, James Levinsohn, and Ariel Pakes. 1995. “Automobile Prices in Market Equilibrium.” *Econometrica*, 63(4): 841–890.

Berry, Steven, James Levinsohn, and Ariel Pakes. 1999. “Voluntary Export Restraints on Automobiles: Evaluating a Trade Policy.” *American Economic Review*, 89(3): 400–430.

Berry, Steven T. 1994. “Estimating Discrete-Choice Models of Product Differentiation.” *The RAND Journal of Economics*, 25(2): 242–262.

Berry, Steven T., and Philip A. Haile. 2014. “Identification in Differentiated Products Markets Using Market Level Data.” *Econometrica*, 82(5): 1749–1797.

Besanko, David, Ulrich Doraszelski, Yaroslav Kryukov, and Mark Satterthwaite. 2010. “Learning-by-Doing, Organizational Forgetting, and Industry Dynamics.” *Econometrica*, 78(2): 453–508.

Bloom, Nick, Rachel Griffith, and John Van Reenen. 2002. “Do R&D tax credits work? Evidence from a panel of countries 1979–1997.” *Journal of Public Economics*, 85(1): 1–31.

Bollinger, Bryan, and Kenneth Gillingham. 2019. “Learning-by-Doing in Solar Photovoltaic Installations.” *Working Paper*.

Borenstein, Severin. 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.

Brenkers, Randy, and Frank Verboven. 2006. “Liberalizing a Distribution System: The European Car Market.” *Journal of the European Economic Association*, 4(1): 216–251.

Cabral, Luis M. B., and Michael H. Riordan. 1994. “The Learning Curve, Market Dominance, and Predatory Pricing.” *Econometrica*, 62(5): 1115–1140.

California State Senate. 2006. “S.B. 1 - Electricity: solar energy: net metering.” 2005–2006 Legislative Session.

Collard-Wexler, Allan. 2013. “Demand Fluctuations in the Ready-Mix Concrete Industry.” *Econometrica*, 81(3): 1003–1037.

Conlon, Christopher, and Jeff Gortmaker. 2020. “Best practices for differentiated products demand estimation with PyBLP.” *The RAND Journal of Economics*, 51(4): 1108–1161.

Covert, Thomas R. 2015. “Experiential and Social Learning in Firms: The Case of Hydraulic Fracturing in the Bakken Shale.” *Working Paper*.

Covert, Thomas R., and Richard L. Sweeney. 2022. “Winds of Change: Estimating Learning by Doing without Cost or Input Data.” *Working Paper*.

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

Doraszelski, Ulrich, and Mark Satterthwaite. 2010. “Computable Markov-perfect industry dynamics.” *The RAND Journal of Economics*, 41(2): 215–243.

Dorsey, Jackson. 2024. “Solar Market Frictions: The Role of Platforms and Policies.” *The Review of Economics and Statistics*, 1–45.

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

Feldman, David, David Friedman, and Robert Margolis. 2013. “Financing, Overhead, and Profit: An In-Depth Discussion of Costs Associated with Third-Party Financing of Residential and Commercial Photovoltaic Systems.” National Renewable Energy Laboratory, U.S. Department of Energy NREL/TP-6A20-60401, Golden, CO.

Fowlie, Meredith, Mar Reguant, and Stephen P. Ryan. 2016. “Market-Based Emissions Regulation and Industry Dynamics.” *Journal of Political Economy*, 124(1): 249–302.

Fudenberg, Drew, and Jean Tirole. 1983. “Learning-by-Doing and Market Performance.” *The Bell Journal of Economics*, 14(2): 522–530.

Fu, Ran, Donald Chung, Travis Lowder, David Feldman, Kristen Ardani, and Robert Margolis. 2016. “U.S. Solar Photovoltaic System Cost Benchmark: Q1 2016.” National Renewable Energy Laboratory, U.S. Department of Energy NREL/TP-6A20-66532, Golden, CO.

Gerarden, Todd D. 2022. “Demanding Innovation: The Impact of Consumer Subsidies on Solar Panel Production Costs.” *Management Science*, Forthcoming.

Ghemawat, Pankaj, and A. Michael Spence. 1985. “Learning Curve Spillovers and Market Performance.” *The Quarterly Journal of Economics*, 100: 839–852.

Gillingham, Kenneth, and Tsvetan Tsvetanov. 2019. “Hurdles and steps: Estimating demand for solar photovoltaics.” *Quantitative Economics*, 10(1): 275–310.

Gowrisankaran, Gautam, and Marc Rysman. 2012. “Dynamics of Consumer Demand for New Durable Goods.” *Journal of Political Economy*, 120(6): 1173–1219.

Grigolon, Laura, and Frank Verboven. 2014. “Nested Logit or Random Coefficients Logit? A Comparison of Alternative Discrete Choice Models of Product Differentiation.” *The Review of Economics and Statistics*, 96(5): 916–935.

Hall, Bronwyn, and John Van Reenen. 2000. “How effective are fiscal incentives for R&D? A review of the evidence.” *Research Policy*, 29(4): 449–469.

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

Hughes, Jonathan E., and Molly 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.

Ifraim, Bar, and Gabriel Y. Weintraub. 2017. “A Framework for Dynamic Oligopoly in Concentrated Industries.” *The Review of Economic Studies*, 84(3): 1106–1150.

Igami, Mitsuru. 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.

Interstate Renewable Energy Council. 2023. “National Solar Jobs Census 2022.” Interstate Renewable Energy Council.

IRENA. 2022. “Renewable Power Generation Costs in 2021.” International Renewable Energy Agency, Abu Dhabi.

IRENA. 2023. “Renewable Capacity Statistics 2023.” International Renewable Energy Agency, Abu Dhabi.

Irwin, Douglas A., and Peter J. Klenow. 1994. “Learning-by-Doing Spillovers in the Semiconductor Industry.” *Journal of Political Economy*, 102(6): 1200–1227.

Jeon, Jihye. 2022. “Learning and investment under demand uncertainty in container shipping.” *The RAND Journal of Economics*, 53(1): 226–259.

Juhász, Réka, Nathan J. Lane, and Dani Rodrik. 2023. “The New Economics of Industrial Policy.”

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

Kellogg, Ryan. 2011. “Learning by Drilling: Interfirm Learning and Relationship Persistence in the Texas Oilpatch.” *The Quarterly Journal of Economics*, 126(4): 1961–2004.

Langer, Ashley, and Derek Lemoine. 2022. “Designing Dynamic Subsidies to Spur Adoption of New Technologies.” *Journal of the Association of Environmental and Resource Economists*, 9(6): 1197–1234.

Levitt, Steven D., John A. List, and Chad Syverson. 2013. “Toward an Understanding of Learning by Doing: Evidence from an Automobile Assembly Plant.” *Journal of Political Economy*, 121(4): 643–681.

McFadden, Daniel. 1977. “Modelling the Choice of Residential Location.” Cowles Foundation for Research in Economics, Yale University Cowles Foundation Discussion Paper 477.

Miller, Nathan H., and Matthew C. Weinberg. 2017. “Understanding the Price Effects of the MillerCoors Joint Venture.” *Econometrica*, 85(6): 1763–1791.

Nemet, Gregory F. 2009. “Demand-pull, technology-push, and government-led incentives for non-incremental technical change.” *Research Policy*, 38(5): 700–709.

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

Nemet, Gregory F., Eric O’Shaughnessy, Ryan Wiser, Naïm Darghouth, Galen Barbose, Ken Gillingham, and Varun Rai. 2017. “Characteristics of low-priced solar PV systems in the U.S.” *Applied Energy*, 187: 501–513.

Olley, G. Steven, and Ariel Pakes. 1996. “The Dynamics of Productivity in the Telecommunications Equipment Industry.” *Econometrica*, 64(6): 1263–1297.

Pakes, Ariel, Michael Ostrovsky, and Steven Berry. 2007. “Simple estimators for the parameters of discrete dynamic games (with entry/exit examples).” *The RAND Journal of Economics*, 38(2): 373–399.

Pless, Jacquelyn, and Arthur A. Van Benthem. 2019. “Pass-Through as a Test for Market Power: An Application to Solar Subsidies.” *American Economic Journal: Applied Economics*, 11(4): 367–401.

Rubin, Donald B. 1981. “The Bayesian Bootstrap.” *The Annals of Statistics*, 9(1): 130–134. Publisher: Institute of Mathematical Statistics.

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

Seim, Katja. 2006. “An Empirical Model of Firm Entry with Endogenous Product-Type Choices.” *The RAND Journal of Economics*, 37(3): 619–640.

Sexton, Steven, A. Justin Kirkpatrick, Robert I. Harris, and Nicholas Z. Muller. 2021. “Heterogeneous Solar Capacity Benefits, Appropriability, and the Costs of Suboptimal Siting.” *Journal of the Association of Environmental and Resource Economists*, 8(6): 1209–1244.

Spence, A. Michael. 1981. “The Learning Curve and Competition.” *The Bell Journal of Economics*, 12(1): 49–70.

Sweeting, Andrew. 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.

Thompson, Peter. 2001. “How Much Did the Liberty Shipbuilders Learn? New Evidence for an Old Case Study.” *Journal of Political Economy*, 109(1): 103–137.

Thompson, Peter. 2007. “How Much Did the Liberty Shipbuilders Forget?” *Management Science*, 53(6): 908–918.

Thornton, Rebecca Achee, and Peter Thompson. 2001. “Learning from Experience and Learning from Others: An Exploration of Learning and Spillovers in Wartime Shipbuilding.” *American Economic Review*, 91(5): 1350–1368.

van Benthem, Arthur, Kenneth Gillingham, and James Sweeney. 2008. “Learning-by-Doing and the Optimal Solar Policy in California.” *The Energy Journal*, 29(3): 131–151.

Vreugdenhil, Nicholas. 2023. “Booms, Busts, and Mismatch in Capital Markets: Evidence from the Offshore Oil and Gas Industry.” *Working Paper*.

Weintraub, Gabriel Y., C. Lanier Benkard, and Benjamin Van Roy. 2008. “Markov Perfect Industry Dynamics With Many Firms.” *Econometrica*, 76(6): 1375–1411.

Wollmann, Thomas G. 2018. “Trucks without Bailouts: Equilibrium Product Characteristics for Commercial Vehicles.” *American Economic Review*, 108(6): 1364–1406.

Online Appendix for “Learning-by-Doing, Spillovers, and Market Structure in the U.S. Residential Solar Industry”

Jacob T. Bradt

The following appendices are **for online publication only**:

- Appendix Section [A](#): Data Appendix
- Appendix Section [B](#): A Dynamic Model of Demand for Solar Installations
- Appendix Section [C](#): Estimating the Exit Policy Function
- Appendix Section [D](#): Identification: Predicted Experience Instruments
- Appendix Section [E](#): Value Function Approximation
- Appendix Section [F](#): Counterfactual Solution Method
- Appendix Section [G](#): Supplemental Figures and Tables

A Data Appendix

A.1 Data Sources

My primary data source is Lawrence Berkeley National Laboratory’s “Tracking the Sun” database, which compiles system-level data on PV installations from state agencies and utilities administering incentive programs, renewable energy credit registration, or grid interconnection ([Barbose et al., 2022](#)). The database includes installation date, system size, installed price, rebates, customer type, zip code, mounting type, installer name, and technical hardware details including module efficiency, make, and manufacturer. LBNL processes source data by removing missing entries, standardizing names, and integrating equipment specifications. The database covers 2.5+ million systems from 2000 to 2021, representing approximately 77% of the US market. These data are available at <https://emp.lbl.gov/tracking-the-sun>.

I obtained system-level hardware cost data for California residential PV systems from the California Public Utilities Commission (CPUC). Though not uniformly required, installers voluntarily provided hardware costs for many systems applying for CSI rebates, particularly during 2008-2013. I acquired data for over 60,000 systems from CPUC via LBNL contacts and successfully matched these to over 79% of processed LBNL installations during this period using unique identifiers. These hardware cost data are critical for separating installation costs from hardware costs. Data are available upon request at jbradt@g.harvard.edu.

I also use the CPUC CSI Working Dataset, which contains complete CSI rebate applications from California’s three main investor-owned utilities, including non-accepted applications. This dataset contains detailed rebate eligibility information that I use to determine precise dates of CSI rebate changes for each IOU. These data are available at <https://www.californiadgstats.ca.gov/downloads/>.

To estimate potential market size for residential solar adoption, I combine data from Google Project Sunroof with Census housing counts. Project Sunroof estimates the share of buildings per county where solar adoption yields positive net present value by combining

satellite imagery, 3D modeling, property-level shade calculations, NREL weather data, Clean Power Research electricity rates, and Aurora Solar cost data. Property-level variation in returns from shading and roof orientation makes this approach preferable to assuming all households are potential adopters. Data are available at <https://sunroof.withgoogle.com/>.

I use solar irradiance data from the World Bank Global Solar Atlas, which provides long-term annual average PV power potential of 1 kW systems at 250-meter resolution. I combine these data with Census county boundaries to estimate average annual production potential per county. I use these estimates to calculate consumption and net energy metering benefits (Section A.4) and to quantify counterfactual power output changes for avoided climate and air pollution damages. Data are available at <https://globalsolaratlas.info/download>.

I obtain household demographic data from the Census American Community Survey (ACS) Public Use Microdata Sample (PUMS) for California 2008-2013. The PUMS provides a 1-in-100 national sample identifying location to public use microdata area (100,000+ persons) or county when population exceeds 100,000. I draw 200 households per county-period to model income-based preference heterogeneity in demand. For small counties not identified in PUMS, I draw from the statewide sample. Since PUMS data are only available at an annual level, I draw from the same annual sample for each county-period within the same calendar year. Data are available at <https://www.census.gov/programs-surveys/acs/microdata.html>.

I use Census ACS 5-year Estimates to obtain annual owner-occupied housing units per county for 2009-2013, which I combine with Project Sunroof data to calculate potential PV adopter market size and firm market shares. For 2008, which is unavailable in the ACS 5-year estimates, I impute housing counts using the 2000 Decennial Census and 2009-2013 ACS estimates. Data are available at <https://www.census.gov/data/developers/data-sets/acs-5year.html>.


I obtain electricity rate data from the Energy Information Administration's Form 861, which provides annual data on US electric utilities. Following [Borenstein and Bushnell \(2022\)](#), I use total retail revenues and kilowatt-hour sales for residential customers to calculate IOU-specific average retail electricity rates for California's three main IOUs for 2008-2013. I use these rates to calculate consumption and NEM benefits (Section A.4). Data are available at <https://www.eia.gov/electricity/data/eia861/>.

Finally, I use wage data from the Bureau of Labor Statistics Quarterly Census of Employment and Wages (QCEW) to construct demand instruments. The QCEW provides quarterly average wages for private sector electricians and roofers by California county for 2008-2013. I construct county-level weighted averages using county employment levels and use these as demand instruments. Data are available at <https://www.bls.gov/cew/>.

A.2 Detailed Policy Descriptions

This section provides comprehensive details on the three main policy instruments subsidizing residential solar PV adoption in California during the 2008-2013 period: the California Solar Initiative (CSI), the federal Investment Tax Credit (ITC), and Net Energy Metering (NEM).

Figure A1. Spatial and Temporal Variation in CSI Rebates across IOUs



Notes: This figure shows spatial and temporal variation in CSI rebates. Variation is driven by the spatial distribution of the three main IOUs (left panel) and cumulative installed capacity in each IOU service territory. Since CSI rebate levels are a function of cumulative installed capacity (see Appendix Figure G2), this drives variation in rebate levels over time across IOUs (right panel). County boundaries shown in grey. Sources: California Public Utilities Commission and US Census Bureau.

California Solar Initiative (CSI) The CSI was California’s largest direct rebate program for solar PV, authorized by Senate Bill 1 in 2006 with a \$2.2 billion budget. The program ran from 2007 to 2013, providing upfront cash rebates to customers of California’s three main investor-owned utilities (IOUs): Pacific Gas & Electric (PG&E), Southern California Edison (SCE), and San Diego Gas & Electric (SDG&E).

The CSI rebate schedule was designed explicitly with learning-by-doing in mind. Rebates started at \$2.50 per watt and stepped down over 10 rate levels based on cumulative installed capacity in each IOU service area. This declining structure assumed that industry experience would reduce installation costs over time, thereby reducing the rebates needed to incentivize adoption. Each IOU progressed through the rebate schedule independently based on cumulative capacity in its service territory, creating spatial and temporal variation in rebate levels that provides plausibly-exogenous variation in net-of-rebate prices useful for estimation.

The CSI was funded through electricity ratepayer transfers rather than general tax revenue. Each IOU collected funds from its ratepayers to finance rebates for systems installed within its service territory. This funding mechanism means the CSI represents a transfer between California households rather than a true fiscal outlay from government coffers, though it still imposed costs on ratepayers.

The CSI had explicit industrial policy goals beyond emissions reductions. The authorizing legislation stated the program aimed to “establish a self-sufficient solar industry” (California State Senate, 2006), making evaluation of the program’s effect on market structure central to assessing its success.

Federal Investment Tax Credit (ITC) The federal ITC provides a tax credit equal to a percentage of installed system costs for solar PV systems. The Energy Policy Act of 2005 established a 30% tax credit for residential solar systems, which became effective in 2006. Initially, the credit was capped at \$2,000 per system. The Emergency Economic Stabilization Act of 2008 removed this cap starting in 2009, substantially increasing the subsidy for typical residential systems.

The ITC is applied to the net-of-rebate installed cost. For a system with gross installed price p and upfront rebate r , the ITC provides a credit of $0.3 \times (p - r)$. This means the ITC and CSI rebates interact: higher CSI rebates reduce the base to which the ITC applies, creating a marginal subsidy rate of $0.3 + 0.7 \times CSI$ rather than $0.3 + CSI$.

The ITC has been extended and modified several times. It was scheduled to step down to 26% in 2020 and 22% in 2021 before expiring for residential systems in 2022. However, the Inflation Reduction Act of 2022 extended the 30% credit through 2032. During my sample period (2008-2013), the credit remained at 30% throughout, with the only major change being the removal of the \$2,000 cap in 2009.

California also offered a state-level solar tax credit from 2001-2005, but this had expired before my sample period begins.

Net Energy Metering (NEM) Net energy metering requires utilities to credit solar PV owners for excess electricity generation at the retail electricity rate. Under California's NEM program, households with solar systems can offset their electricity consumption with their own generation. When generation exceeds consumption, the excess is fed back to the grid and credited at the retail rate, which is substantially higher than the wholesale rate utilities would otherwise pay for electricity.

California's retail electricity rates feature increasing-block pricing, where the marginal price per kilowatt-hour increases with total consumption. Households consuming above certain thresholds face higher rates for marginal consumption. This pricing structure interacts with solar PV adoption in important ways: households in higher consumption tiers receive greater per-kilowatt-hour savings from solar generation, making adoption more attractive for high-consumption households. Additionally, excess generation is credited at the household's marginal retail rate, which depends on their consumption tier.

The value of NEM benefits depends on several factors: the household's retail electricity rate (which varies by IOU and consumption level), the system's annual electricity generation (which depends on solar irradiance and system size), and the household's consumption patterns. I calculate NEM benefits using estimates of excess PV output from [Darghouth et al. \(2011\)](#), county-level solar irradiance data from the World Bank Global Solar Atlas, and IOU-specific average retail rates from EIA Form 861 data. I assume a 25-year system lifespan and discount future electricity bill savings at an annual rate of 12.5% following [De Groote and Verboven \(2019\)](#).

The NEM program has evolved over time. The original NEM 1.0 rules that applied during my sample period were modified in 2016 (NEM 2.0) and again in 2022 (NEM 3.0), with later versions reducing the credit rates for excess generation. During 2008-2013, however, the basic NEM 1.0 structure remained constant.

Together, these three policy instruments provided substantial subsidies for residential

solar adoption in California. The combination of upfront CSI rebates, federal tax credits, and ongoing NEM benefits reduced the effective price of solar systems by 50-70% in many cases, driving rapid adoption growth during this period.

A.3 Sample Restrictions

I restrict LBNL Tracking the Sun data to California residential rooftop systems under 20 kW with observed price and rebate data. Excluded: non-residential systems, ground-mounted systems, systems over 20 kW (likely multi-family condominiums), self-installed systems, third-party-owned systems, and price-per-watt outliers (below 1st or above 99th percentile).

A.4 Constructing Data for Demand Estimation

I aggregate system-level data to firm-county-period level, subsetting to installations with observed hardware costs. Normalize prices and rebates by installed capacity to aggregate comparable goods, abstracting from scale economies (reasonable given commoditized PV module technology).

Market Shares: Combine Project Sunroof data (share of buildings with positive NPV from solar) with Census owner-occupied housing counts to estimate potential adopter pool, avoiding principal-agent problems in renter-occupied housing. Account for PV durability by removing prior adopters (using full 2000-2021 LBNL history) from potential market size each period.

Prices and Rebates: Calculate weighted county-period-firm averages, weighting each installation by its share of the firm's total installed capacity within that county-period.

Investment Tax Credit (ITC): Assuming full capitalization, ITC benefit per watt is:

$$ITC_{ijmt} = \begin{cases} \min \left\{ 1000/q_{ijmt}, 0.3(p_{ijmt} - r_{ijmt}) \right\} & \text{if } t < 2009 \text{ H1} \\ 0.3(p_{ijmt} - r_{ijmt}) & \text{if } t \geq 2009 \text{ H1} \end{cases}$$


where p_{ijmt} is installed price, r_{ijmt} is upfront rebate per watt, and q_{ijmt} is system capacity.

Net Energy Metering (NEM): Calculate NPV of NEM benefits using [Darghouth et al. \(2011\)](#)'s excess PV output estimates, World Bank Global Solar Atlas power potential data (Figure A2), and EIA Form 861 retail electricity rates (Figure A2). Assumptions: 25-year system lifespan, 12.5% annual household discount rate ([De Groote and Verboven, 2019](#)). Aggregate ITC and NEM with CSI rebates as total public incentives per watt.

All dollar values converted to 2013 real dollars using CPI.²¹ Firm attributes include count of distinct module types and indicator for above-75th-percentile module efficiency.

²¹US Bureau of Labor Statistics, Consumer Price Index for All Urban Consumers: All Items in U.S. City Average [CPIAUCSL], retrieved from FRED, Federal Reserve Bank of St. Louis.

Figure A2. Spatial Variation in PV Power Potential and Retail Electricity Rates



Notes: These maps show spatial variation in the long-term annual average photovoltaic (PV) power potential of a 1 kW capacity PV system from the World Bank Group's Global Solar Atlas (left) and average retail electricity rates by county based on data from the Energy Information Administration's Form 861 (right).

A.5 Constructing Data for Supply Model

Supply data start from the same installer-county-period dataset used for demand estimation, with added fields for average hardware cost per watt and experience/quantity variables. Hardware cost is calculated as a firm-county-period weighted average and deflated using the CPI.

I compute experience variables using the full LBNL history (2000-2021), defining experience as cumulative PV system capacity installed (watts). For each installer-county-period, I calculate cumulative installed capacity plus rival experience measures: cumulative capacity for rivals in the same county, other counties, same module manufacturer, and other manufacturers. All experience fields are normalized by statewide cumulative installed watts as of H1 2008 to ensure numerical stability, as own and rival experience can differ by orders of magnitude. I also calculate corresponding period-specific quantities in both absolute and normalized watts for use in estimation.

Table A1. Summary Statistics for Processed Installer-level Data

	Mean	SD	Min	Max
Number of Installations	9.18	30.92	1.00	1560.00
Total Installed Capacity (kW)	48.53	164.46	1.23	8129.72
Market Share	0.01	0.01	0.00	0.26
Market Share: Inside	3.86	6.62	0.01	83.00
Average Installed Price (2013 \$/W)	6.49	1.59	1.25	10.90
Average Hardware Cost (2013 \$/W)	4.12	1.47	0.92	10.34
Own Experience: In-market (kW)	213.40	522.51	0.00	15 045.89
Own Experience: Out-of-market (kW)	2899.14	7682.26	0.00	78 300.92
Rival Experience: In-market (MW)	16.04	17.20	0.29	85.68
Rival Experience: Out-of-market (MW)	323.60	164.15	81.30	665.95
Rival Experience: Same Manufacturer (MW)	30.96	36.34	0.00	144.37
Rival Experience: Other Manufacturer (MW)	308.69	160.58	71.82	668.80
N	10,247			

Notes: This table presents summary statistics for the processed installer-level dataset that I use in my empirical analysis. The unit of observation is at the installer-county-period-level (half-yearly), so descriptive statistics pool observations across markets and periods. Total installed capacity in a period and the measures of firms' own experience are in kilowatts (kW), whereas rivals' experience measures are in megawatts (MW), or 1000 kW.

B A Dynamic Model of Demand for Solar Installations

A large literature uses static discrete choice models to estimate demand for durable goods, including vehicles (Berry et al., 1995, 1999) and commercial airplanes (Benkard, 2004). However, several papers implement dynamic discrete choice models of demand for solar PV, finding that static models can deliver biased demand estimates (Bollinger and Gillingham, 2019; De Groote and Verboven, 2019; Feger et al., 2022; Langer and Lemoine, 2022).

While there is evidence to suggest that consumers are forward-looking in this setting, I assume that a static demand model provides a reasonable approximation to consumer behavior. This simplifying assumption buys substantial computational gains, eliminating the need to jointly solve dynamic supply and demand in counterfactual analysis. Given my emphasis on the supply-side impacts of solar incentives, the use of a static model as a reduced form for demand is likely reasonable; however, it is worth assessing the extent to which this introduces bias in my estimates.

To do so, I develop and estimate a dynamic discrete choice model of solar adoption based on De Groote and Verboven (2019) and Bollinger and Gillingham (2019). For the purposes of this exercise, I omit individual-level heterogeneity in preferences (i.e., remove random coefficients) and omit observable, time- and firm-specific attributes.

B.1 Model Setup

Similar to the static model in Section 3.1, incumbent firms in each period and market face a set of idiosyncratic consumers, $i \in \{1, \dots, N_{mt}\}$, who demand solar PV installation services. Each consumer observes a market- and time-specific state, s_{mt} , and either purchases a solar PV installation from one of the observed incumbents ($j \in \{1, \dots, J_{mt}\}$) or chooses to not

install solar PV in this period ($j = 0$). The market and time specific state is the union of active incumbents' prices; available rebates, including the present discounted value of the future stream of net metering rents; and qualities:

$$\cup_{j \in \mathcal{J}_{mt}} [p_{jmt} \ r_{jmt} \ \xi_{jmt}]$$

Consumers are also differentiated by an idiosyncratic random utility shock that is alternative specific, $\bar{\varepsilon}_{ijmt}$. The conditional indirect utility that consumer i receives from choosing installer j in market m in period t is:

$$u_{ijmt}(s_{mt}) = \underbrace{\alpha(p_{jmt} - r_{jmt}) + \xi_{jmt} + \bar{\xi}_j + \bar{\xi}_t + \bar{\varepsilon}_{ijmt}}_{\equiv \delta_{jmt}} \quad (B1)$$

where p_{jmt} is the retail price per watt of system capacity; r_{jmt} is a market-time-varying rebate or subsidy per watt of system capacity; ξ_{jmt} is a firm's market-time-specific unobserved quality; $\bar{\xi}_j$ allows the mean valuation of unobserved product characteristics to vary freely by product; and $\bar{\xi}_t$ allows the mean valuation of the indirect utility from installation to vary freely over time. As in the main text, I normalize prices and rebates by system capacity to ensure consistency when aggregating these variables across systems of different sizes.

Consumers that do not adopt solar in market m in period t receive a flow utility u_{0mt} and experience the option value of adopting in the future:

$$u_{i0mt}(s_{mt}) = \underbrace{u_{0mt} + \beta \mathbb{E}[V(s_{mt+1}|s_{mt})]}_{\equiv \delta_{0mt}} + \bar{\varepsilon}_{i0mt} \quad (B2)$$

where β is a common, quarterly discount factor.

As with the static demand model in the main text, I decompose the idiosyncratic preference shock using the distributional assumptions of the nested logit model following [Berry \(1994\)](#). For each market and in each period, define two groups, $g \in \{0, 1\}$, where $g = 1$ includes the full set of incumbent installers and $g = 0$ the no-installation option. Then

$$\bar{\varepsilon}_{ijmt} = \zeta_{igmt} + (1 - \eta)\varepsilon_{ijmt}$$

where ε_{ijmt} is independent and identically distributed (i.i.d.) Type 1 Extreme Value, ζ_{igmt} has the unique distribution such that $\bar{\varepsilon}_{ijmt}$ is i.i.d. Type 1 Extreme Value, and $0 \leq \eta < 1$ is a nesting parameter that proxies for the degree of preference correlation within a group.

This assumption on the structure of the idiosyncratic preference shocks results in predicted market shares that follow the usual nested logit closed form, which I include in Section 3.1. Following [Berry \(1994\)](#), it is possible to invert predicted market shares as follows:

$$\log(ms_{jmt}(s_{mt})) - \log(ms_{0mt}(s_{mt})) = \delta_{jmt} - \delta_{0mt} + \eta \log(\bar{m}s_{jmt|g}(s_{mt})) \quad (B3)$$

where ms_{jmt} is firm j 's predicted share of market m in period t ; ms_{0mt} is the predicted outside share in market m in period t ; and $\bar{m}s_{jmt|g}$ is firm j 's conditional within-group share in market m in period t (i.e., the firm's inside share).

B.2 Ex Ante Value Function

Taking (B3) to the data requires a closed form for households' ex ante value function. This requires an approach to handling the expectation operator in (B2), which integrates over uncertainty about the next period state variables. One standard approach (applied in the dynamic supply model in the text) is to specify and estimate an explicit stochastic process for the state variables. However, for the purposes of flexibility in this dynamic demand model, I follow [De Groote and Verboven \(2019\)](#) and decompose the expected ex ante value function into a realized value function and a short run prediction error:

$$e_{mt} \equiv V(s_{mt+1}) - \mathbb{E}\left[V(s_{mt+1}|s_{mt})\right] \quad (B4)$$

where I assume that households' expectations are on average correct (i.e., households have rational expectations) such that e_{mt} is mean zero.

[Hotz and Miller \(1993\)](#) and [Arcidiacono and Miller \(2011\)](#) show that it is possible to express continuation values as functions of the conditional choice probabilities for one of the terminating options, say $j = j'$. Normalizing the flow utility of non-adoption to the product of Euler's constant and the common discount factor, i.e., $u_{0mt} + 0.577\beta = 0$, [Arcidiacono and Miller \(2011\)](#) show that the assumption of a nested logit error structure provides a helpful closed form expression for the value function. This combined with the rational expectations assumption results in the following closed form for the mean utility of non-adoption:

$$\begin{aligned} \delta_{0mt} &= \mathbb{E}\left[V(s_{mt+1}|s_{mt})\right] \\ &= \beta \left(\delta_{j'mt+1}(s_{mt+1}) - \log(ms_{j'mt+1}(s_{mt+1})) \right. \\ &\quad \left. - \eta \left(\log(ms_{gmt+1}(s_{mt+1})) - \log(ms_{j'mt+1}(s_{mt+1})) \right) - e_{mt} \right) \end{aligned} \quad (B5)$$

where $ms_{gmt+1}(s_{mt+1})$ is the inside group share (i.e., the probability of adoption) in market m and period t .

B.3 Estimating Equation

Combining (B1), (B3), and (B5) and rearranging terms gives the following equation:

$$\begin{aligned} \log\left(\frac{ms_{jmt}}{ms_{0mt}}\right) - \beta \log(ms_{j'mt+1}) \\ &= \alpha \left((p_{jmt} - \beta p_{j'mt+1}) - (r_{jmt} - \beta r_{j'mt+1}) \right) \\ &\quad + \eta \left(\log(\bar{ms}_{jmt|g}) + \beta \left(\log(ms_{gmt+1}) - \log(ms_{j'mt+1}) \right) \right) \\ &\quad + \underbrace{\left(\xi_{jmt} - \beta \xi_{j'mt+1} \right)}_{\equiv \tilde{\xi}_{jmt}} + \underbrace{\left(\bar{\xi}_j - \beta \bar{\xi}_{j'} \right)}_{\equiv \tilde{\xi}_j} + \underbrace{\left(\bar{\xi}_t - \beta \bar{\xi}_{t+1} \right)}_{\equiv \tilde{\xi}_t} + \underbrace{\beta e_{mt}}_{\text{mean zero}} \end{aligned} \quad (B6)$$

Table B1. Estimated Demand Parameters from a Dynamic Model

	(1)
‘Dynamic’ Price/Income	−16.900 (1.890)
Nesting Parameter	0.423 (0.059)
Firm FE	Yes
Year FE	Yes
CZ FE	Yes
Observations	9,264
R ²	0.732
F-test (1st stage), ‘Dynamic’ Price/Income	424.2
F-test (1st stage), Nesting Parameter	133.7

Notes: Estimation follows the procedure outlined in Appendix B. I divide prices and rebates by county-quarter mean income whereas for comparison to the static demand estimates in Table 1. Standard errors clustered by county are reported in parentheses.

The above estimating equation (B6) is a function of data—including current period and lead values of market shares, prices, and rebates—and the target parameters, (α, η) , which I can estimate via ordinary least squares.


Two issues remain before implementing (B6): the choice of reference option, j' , and the treatment of the scalar unobservable terms, $(\tilde{\xi}_{jmt}, \tilde{\xi}_j, \tilde{\xi}_t)$. In terms of the reference option, an easy solution would be to use a universal installer that appears across markets. Unfortunately, there is no single installer that operates in each county-quarter in my data. As a result, I follow Bollinger and Gillingham (2019) and use a market-specific next-period average for the values of market shares and state variables of the reference option that enter (B6). This revised estimating equation converges asymptotically to (B6).

To account for the scalar unobservable terms, I include firm and time fixed effects when implementing (B6). This leaves the scalar unobservable $\tilde{\xi}_{jmt}$, which generates a set of analogous endogeneity concerns to that associated with the static demand model in the text. I therefore estimate my revised estimating equation (B6) via two-stage least squares, with the same set of demand instruments as those that I outline in Section 4.1 of the text.

B.4 Results

I report estimates of the main target parameters in Table B1. Overall, the results are qualitatively consistent with the static demand estimates reported in Column (1) of Table 1. Note that I divide prices and rebates by the county-quarter mean income for the sake of direct comparison with the static models in the main text. The estimate of the nesting parameter from the dynamic model is slightly larger, but indistinguishable from that from the analogous static model. The first-stage F -statistics for both the nesting parameter and price parameter are large in magnitude.

Figure B1. Comparison of Static and Dynamic Demand Elasticities



Notes: This figure compares the model-predicted, short-run price elasticities of demand using the static demand estimates from Column (1) of Table 1 (left panel) and the dynamic demand estimates from Table B1. The overlaid line shows a cubic b-spline fitted to the short-run price elasticity estimates.

To assess the performance of the static demand model that I use in the text as a reduced form for dynamic demand, I plot the implied short-run elasticities from the dynamic estimates in Table B1 alongside those from the analogous static model. As shown in Figure B1, the static model reasonably matches both the level and shape of demand elasticities over time. While the short-run elasticities from the dynamic demand estimates are in general larger in magnitude—the median short-run elasticity from the dynamic estimates is -4.75 compared with -3.76 for the analogous static estimates—both are within the range of previously published static estimates (see Figure G3) and the difference between the two is relatively minor. Figure B1 suggests that the reduced form demand model used in the main analysis in the text offers a reasonable approximation.

C Estimating the Exit Policy Function

I estimate firms' exit policy function using a logit regression:

$$\Pr(\chi_{jmt}^x = 1 | s_{mt}) = \frac{\exp(h_j(s_{mt}))}{1 + \exp(h_j(s_{mt}))}$$

where χ_{jmt}^x equals 1 if firm j exits market m in period t and 0 otherwise and $h_j(s_{mt})$ is a flexible function of the states. Obtaining consistent estimates of the exit policy function is important for consistent estimation of the dynamic parameters in the second stage. I therefore follow the data-driven approach of [Gerarden \(2022\)](#) to determine the functional form

of $h_j(s_{mt})$ when estimating firms' exit policies. This approach has the benefit of optimizing the tradeoff between a flexible specification and the challenges associated with overfitting.

In particular, I begin by identifying a large set of candidate regressors to use in $h_j(s_{mt})$. These include quadratic polynomials of the full set of state variables and the complete set of pairwise interactions between these terms.²² I also include county and quarter fixed effects. I then use LASSO for variable selection. Specifically, I model the discrete decision to exit using the following penalized maximum likelihood:

$$\min_{\mu} - \left[\frac{1}{N} \sum_{j,m,t} \chi_{jmt}^x h_j(s_{mt}; \mu) - \log \left(1 - \exp \left(h_j(s_{mt}; \mu) \right) \right) \right] + \lambda \|\mu\|_1$$

I select the tuning parameter, λ , by leave-one-out k -fold cross validation with $k = 10$. Figure C1 shows the resulting estimated binomial deviance for different values of λ .

This identifies a set of non-zero regressors, $\tilde{h}_j(s_{mt}; \tilde{\mu})$. I then model the discrete exit decision using a logit model with the non-zero regressors selected in the first stage and estimate the parameters on the remaining regressors via maximum likelihood:

$$\min_{\tilde{\mu}} - \left[\frac{1}{N} \sum_{j,m,t} \chi_{jmt}^x \tilde{h}_j(s_{mt}; \tilde{\mu}) - \log \left(1 - \exp \left(\tilde{h}_j(s_{mt}; \tilde{\mu}) \right) \right) \right]$$

The final step logit model has an estimated binomial deviance of 12.34%. The resulting parameter estimates allow me to fit exit probabilities for each incumbent observed in my data. Figure C2 shows the density of fitted exit probabilities for incumbents that I observe continue and incumbents that I observe exit.

D Identification: Predicted Experience Instruments

As discussed in Section 4.3, identification of the production cost parameters governing learning economies relies on the validity of the instrumented moment conditions given in equation (15). A key concern in forming these moment conditions is ensuring that the instruments are both relevant (correlated with the endogenous variables) and valid (uncorrelated with the innovation term ν_{jmt}). This appendix provides additional detail on the construction and justification of the predicted experience instruments used in GMM estimation.


D.1 Instrument Construction

I construct predicted experience instruments by exploiting exogenous variation in demand shifters and market structure. The construction proceeds in three steps:

Step 1: Reduced-Form Demand Model. I estimate a reduced-form demand model


²²I include quadratic polynomials of the following variables and their pairwise interactions: prices, own experience, other firms' experience within a county, other firms' experience outside a county, hardware cost, the average hardware costs of other firms within a county, quality, the average quality of other firms within a county, the county-quarter inclusive value, and the aggregate demand state.

Figure C1. First-step Tuning Parameter Selection via k -fold Cross-validation



Notes: This figure shows binomial deviance for different values of λ , where binomial deviance is calculated via leave-one-out k -fold cross-validation with $k = 10$. The vertical dashed line shows the value of λ that corresponds to the minimum estimated binomial deviance.

Figure C2. Density of Fitted Exit Probabilities



Notes: This figure shows the density of the resulting fitted exit probabilities separately for incumbents that I observe continue and incumbents that I observe exit.

that regresses log market shares on exogenous demand and cost shifters:

$$\log(s_{jmt}/s_{0mt}) = \alpha_r r_{jmt} + \alpha_{ghi} \text{GHI}_{mt} + \alpha_{elec} \text{ElecPrice}_{mt} + \bar{\alpha}_j + \bar{\alpha}_t + \bar{\alpha}_c + \hat{\alpha}_s \log(s_{jmt}^{inside}) + \varepsilon_{jmt}$$

where r_{jmt} is the rebate per watt, GHI_{mt} is solar radiation in the county, ElecPrice_{mt} is the average electricity price, $\bar{\alpha}_j$, $\bar{\alpha}_t$, and $\bar{\alpha}_c$ are installer, year, and commuting zone fixed effects respectively, and s_{jmt}^{inside} is the inside market share. To address the endogeneity of inside shares, I instrument using the number of firms in the market (N_{mt}) and the total experience of rival firms in the county ($\sum_{k \neq j} E_{kmt}$).

Step 2: Predicted Quantities. Using the fitted values from this demand model, I construct predicted market shares \hat{s}_{jmt} for each firm-county-period. I then convert these to predicted quantities:

$$\hat{q}_{jmt} = \hat{s}_{jmt} \times \text{PotentialMarket}_{mt} \times \text{MeanSystemSize}_{mt}$$

Step 3: Predicted Experience. I construct the predicted own experience instrument as the predicted quantity normalized by total industry experience in the first half-year of the sample (H1 2008):

$$z_{jmt}^e = \frac{\hat{q}_{jmt}}{\text{IndustryExperience}_{H1 \text{ 2008}}}$$

The predicted rival experience instrument is the sum of all other firms' predicted experience in the same county-period:

$$z_{jmt}^{spill} = \sum_{k \neq j} z_{kmt}^e$$

D.2 Instrument Validity

The validity of these instruments rests on two key assumptions:

Assumption 1: Exogeneity of Demand Shifters. The demand shifters used in the reduced-form model (rebates, solar radiation, electricity prices, and market structure) must be orthogonal to the innovation term ν_{jmt} in the productivity shock process. This is plausible because:

- Rebate levels are determined by cumulative installed capacity in each utility service territory based on a predetermined schedule, not by individual firm productivity shocks
- Solar radiation is a geographic characteristic determined by weather patterns
- Electricity prices are set by utility rate cases and regulatory proceedings
- The number of firms and rival experience in period t are predetermined relative to the innovation in period t

Assumption 2: Serial Uncorrelation of Innovations. The innovation term $\nu_{jmt} = \kappa_{jmt} - \rho \kappa_{jmt-1}$ must be serially uncorrelated by construction. This is ensured by the AR(1) specification of the productivity shock process, where any serial correlation in κ_{jmt} is captured by the parameter ρ .

Given these assumptions, the predicted experience instruments satisfy the moment condition $\mathbb{E}[Z'_{jmt}\nu_{jmt}] = 0$ because they are constructed using only exogenous demand shifters and predetermined state variables.

D.3 Instrument Relevance

Figure D1 provides evidence on instrument relevance by showing the relationship between the predicted experience instruments and actual experience measures. The figure displays binscatter plots (30 bins) of residualized relationships after controlling for year fixed effects.

The top-left panel shows a strong, positive relationship between the predicted own experience instrument (z_{jmt}^e) and actual own experience (E_{jmt}), demonstrating strong first-stage relevance. The fitted line has a steep positive slope, indicating that the predicted experience instrument based on exogenous demand shifters is highly correlated with actual experience accumulation.

The second panel from the left in the top row shows the relationship between the predicted rival experience instrument (z_{jmt}^{spill}) and actual own experience. While this relationship is weaker than the own-experience relationship, it remains positive, reflecting the fact that firms operating in markets with many rivals (which drives predicted rival experience) also tend to have accumulated more experience themselves due to market selection effects.

The remaining panels show that the predicted experience instruments have weak or no relationship with own hardware costs and mean rival quality after controlling for year effects. This is reassuring because it suggests the instruments provide identifying variation in experience that is not confounded by variation in other cost components or quality measures.

The figure also displays both current and lagged versions of each instrument (in orange circles and blue triangles, respectively). The lagged instruments, which are included in the actual instrument matrix Z_{jmt} used in estimation, show similar patterns of relevance, ensuring that the instruments provide valid identifying variation even when lagged one period.


D.4 Overidentification

The GMM estimation uses five instruments (z_{jmt}^e , z_{jmt}^{spill} , lagged hardware costs, a time trend, and lagged installation-adjacent wages) to estimate four production cost parameters (c_0 , γ , θ^E , and ρ) plus the exit cost parameter (σ_ϕ). The fact that the model fits the data well (as shown in Section 5) and that the parameter estimates are stable across different instrument specifications provides informal support for the validity of the instruments.

E Value Function Approximation

In light of the fact that the conditions for optimal quantity-setting, exit, and entry all depend on $V_j(s_{mt})$, estimation of the target structural parameters requires solving for the unknown value function. As discussed in the text, I approximate the value function via B-spline basis functions. Value function approximation is appealing in my setting for several reasons. First, given the high dimensionality of the model's continuous state space, conventional approaches that rely on discretization of the states remain computationally-intensive and can

Figure D1. Instrument Relevance for Predicted Experience Instruments



Notes: This figure shows the relationship between the predicted experience instruments used in GMM estimation of production and exit cost parameters and actual experience measures. The predicted experience instruments are constructed by first estimating a reduced-form demand model that regresses log market shares on rebates, solar radiation (GHI), average electricity prices, installer fixed effects, year fixed effects, and commuting zone fixed effects, instrumenting for inside shares using the number of firms and rival experience in the county. Fitted values from this demand model are used to construct predicted quantities and thus predicted cumulative experience for each firm-county-period. The figure shows binscatter plots (30 bins) of residualized relationships after controlling for year fixed effects. The top row shows the relationship between predicted own experience (left) and predicted rival experience (second from left) instruments and actual own experience, demonstrating strong first-stage relevance. The remaining panels show relationships with own hardware cost and mean rival quality. Both current and lagged versions of each instrument are shown, with current instruments in orange circles and lagged instruments in blue triangles.

produce non-trivial approximation errors in this setting. Second, given that the value function implicitly defined by the Bellman equation in this setting is nonlinear in parameters, popular forward simulation approaches are computationally-infeasible in this setting. This non-linearity is due to the fact that static profits are a function marginal production costs, which are nonlinear in the learning parameters. Moreover, [Barwick and Pathak \(2015\)](#) and [Kalouptsidi \(2018\)](#) provide Monte Carlo evidence to suggest that value function approximations perform well in estimating dynamic games.

As shown in the main text, given my assumption that scrap values, ϕ_{jmt} are i.i.d. expo-

nential, it is possible to write the value function prior to the realization of ϕ_{jmt} as

$$\begin{aligned} V_j(s_{mt}) &= \mathbb{E}_\phi[V_j(s_{mt}, \phi_{jmt})] = \mathbb{E}_\phi[\pi_j(s_{mt}) + \max\{\phi_{jmt}, CV_j(s_{mt})\}] \\ &= \pi_j(s_{mt}) + p_j^x(s_{mt})\mathbb{E}_\phi[\phi_{jmt}|\phi_{jmt} > CV_j(s_{mt})] + (1 - p_j^x(s_{mt}))CV_j(s_{mt}) \\ &= \pi_j(s_{mt}) + p_j^x(s_{mt})\sigma_\phi + CV_j(s_{mt}) \end{aligned} \quad (\text{E1})$$

where the final line follows from the fact that $\mathbb{E}_\phi[\phi_{jmt}|\phi_{jmt} > CV_j(s_{mt})] = \sigma_\phi + CV_j(s_{mt})$ as shown by [Pakes et al. \(2007\)](#).

Having obtained estimates of the static demand parameters, exit policy functions, and state transition processes in the first step of estimation, it is possible to obtain a flexible approximation of the value function implicitly defined by the Bellman equation (E1) following recent work in the dynamic games literature (e.g., [Barwick et al. \(2025\)](#)). In particular, given the smoothness of the value function in this context, it is possible to approximate the value function arbitrarily well using L B-spline basis functions $b_j^l(s_{mt})$:

$$V_j(s_{mt}) \simeq \sum_{l=1}^L \lambda_l b_j^l(s_{mt}) \quad CV_j(s_{mt}) \simeq \beta \sum_{l=1}^L \lambda_l \mathbb{E}[b_j^l(s_{mt+1})|s_{mt}] \quad (\text{E2})$$

where $b_j^l(s_{mt})$ are basis functions of the state variables and λ_l are coefficients to be estimated. Plugging (E2) into (E1) gives

$$\sum_{l=1}^L \lambda_l b_j^l(s_{mt}) \simeq \pi_j(s_{mt}; \theta^c) + p_j^x(s_{mt})\sigma_\phi + \beta \sum_{l=1}^L \lambda_l \mathbb{E}[b_j^l(s_{mt+1})|s_{mt}] \quad (\text{E3})$$

where $\theta^c = (c_0, \theta^E, \gamma)$ are the production cost parameters governing learning. From (E3), it is possible to recover estimates $\hat{\lambda}_l$ using data, estimated exit policy functions, and estimated state transitions for a given set of parameter values (θ^c, σ_ϕ) :

$$\{\hat{\lambda}_l\}_{l=1}^L = \arg \min_{\lambda_l} \left\| V_j(s_{mt}; \lambda) - \pi_j(s_{mt}; \theta^c) - \hat{p}_j^x(s_{mt})\sigma_\phi - CV_j(s_{mt}; \lambda) \right\|_2 \quad (\text{E4})$$

where I am choosing approximating coefficients, $\{\hat{\lambda}_l\}_{l=1}^L$, that minimize the L^2 norm of violations of the Bellman equation (E1).

Firm value functions are a function of a high-dimensional state vector. To ease the computational burden associated with approximating firm value functions, I follow the model's simplifying assumption about the approximating equilibrium concept and use moments of the state variables of a firm's rivals when forming approximating basis functions. In particular, I form basis functions of the following variables to approximate firms' value functions:

- E_{jmt} : firm j 's own experience in market m in period t
- $\bar{E}_{jmt}^m = \sum_{k \neq j} E_{kmt}$: total experience of firm j 's rivals in market m in period t
- $\bar{E}_{jmt}^o = \sum_{l \neq m} \sum_{k \neq j} E_{klt}$: total experience of firm j 's rivals in markets outside of market m in period t

- h_{jmt} : firm j 's average hardware costs in market m in period t
- \bar{h}_{kmt} and \bar{h}_{klt} , $\forall k \neq j, l \neq m$: within- and out-of-county averages of firm j 's rival firms' hardware costs in period t
- ξ_{jmt} : firm j 's quality in market m in period t
- $\bar{\xi}_{kmt}$ and $\bar{\xi}_{klt}$, $\forall k \neq j, l \neq m$: within- and out-of-county averages of firm j 's rival firms' quality in period t
- d_{mt} : aggregate demand in market m in period t
- I_{mt} : inclusive value in market m in period t

I augment the basis functions formed with these 11 variables with the full set of county fixed effects when approximating firms' value functions in order to account for differences in expected discounted returns across counties not captured by the above variables.


To select the basis function of the above 11 variables, I test how well B-splines of different orders with different percentile knots fit observed revenue data, since value functions measure expected discounted profits which are functions of revenues. I ultimately select third-order B-splines (i.e., quadratic piecewise polynomials with 3 interior knots). I approximate the expectation in (E3) by averaging state values over 1000 draws of the estimated state transition processes.

A key remaining issue is the set of state realizations on which to evaluate the approximate Bellman equation E3. Similar to [Sweeting \(2013\)](#) and [Barwick et al. \(2025\)](#), I construct a sample of state realizations that includes both all states observed in the data as well as a set of states randomly drawn to span the support of the state variables. In particular, I draw 50,000 additional realizations of the state variables where I independently draw at random each state variable from its empirical support. I then used the fitted exit policy function to predict exit probabilities at each of these additional realizations of the state and estimate simple linear fits of prices and quantities on observed realizations of the state to allow me to predict profits at these simulated states. I ensure that these additional realizations of the state are uniformly distributed across counties and quarters in my estimation period. I estimate $\{\hat{\lambda}_l\}_{l=1}^L$ via E4 using the full set of observed and simulated realizations of the state.

These additional states ensure that I obtain a good approximation of the value function in estimation for two reasons. First, some states (for example, hardware costs and experience) are correlated in the observed data, which makes it difficult to separately identify the basis function coefficients on these variables. Second, parts of the state space are relatively sparse: for instance, certain counties have relatively few observations spanning small regions of the state space in the realized data. These simulated states are therefore quite important in providing a good approximation of the value function.

Figure E1 shows binned scatterplots and third order polynomial fits of the relationship between nine state variabels and the final value function estimates from the main production and exit cost parameter estimtes reported in column (1) of Table 2.

Figure E1. Relationship between State Variables and Value Function Estimates



Notes: This figure shows binned scatterplots and third order polynomial fits of the relationship between nine state variables and the final value function estimates from the main production and exit cost parameter estimates reported in column (2) of Table 2. Value function approximation follows the procedure discussed in detail in Appendix E. Value function estimates are reported in millions of 2013 USD.

F Counterfactual Solution Method

Simulating counterfactual policy environments requires a method for solving the model. My approach to solve the model builds on the method of [Sweeting \(2013\)](#), which adapts parametric policy iteration ([Benitez-Silva et al., 2000](#)) to allow for value function approximation.

My approach is similar to other counterfactual solution methods in applied settings in the literature (Barwick et al., 2025; Gerarden, 2022).

As was the case in estimation, the high-dimensionality and continuous nature of the state space presents a challenge in solving the model. As a result, I maintain the approach to value function approximation outlined in Appendix E.²³ Solving the model involves two steps: first, solving for the new Bellman equation, policy functions, and product market equilibrium, and second simulating the industry forward one period. In each counterfactual scenario, I initiate this two-step procedure at the observed data in the first period of my main estimation sample—the first quarter of 2008—and then repeat the two-step procedure until I reach the end of the main estimation—the last quarter of 2013.

F.1 Solving a Single Period

I implement the first step of this counterfactual solution method via a fixed point algorithm. For a given iteration of this fixed point algorithm, i , I take the following steps:

1. Compute static profits at each state, $\pi_j^i(s_{mt}; \hat{\theta}^c)$, where $\hat{\theta}^c = (\hat{c}_0, \hat{\theta}^E, \hat{\gamma})$ are the preferred production cost parameter estimates, using equilibrium prices, $p_j^i(s_{mt})$; market shares, $ms_j^i(s_{mt})$; and continuation values, $CV_j^i(s_{mt})$ from the previous fixed point iteration.
2. Solve for the value function approximating coefficients, $\hat{\lambda}^{i+1}$ using

$$V_j^{i+1}(s_{mt}; \lambda^{i+1}) = \pi_j^i(s_{mt}; \hat{\theta}^c) + \hat{\sigma}_\phi p_j^{x,i}(s_{mt}) + CV_j^{i+1}(s_{mt}; \lambda^{i+1})$$

where $\pi_j^i(s_{mt}; \hat{\theta}^c)$ is from step 1, $\hat{\sigma}_\phi$ corresponds to the preferred exit parameter estimate, $p_j^{x,i}(s_{mt})$ is the equilibrium exit policy function from the previous iteration, and I form the expectation in $CV_j^{i+1}(s_{mt}; \lambda^{i+1})$ by averaging state values over 1000 draws from the state transition processes estimated in the first stage.

3. Update the exit policy function, $p_j^{x,i+1}(s_{mt})$, using $\hat{CV}_j^{i+1}(s_{mt}; \hat{\lambda}^{i+1})$ and the closed form solution for firms' exit probabilities.
4. Update equilibrium market shares, $ms_j^{i+1}(s_{mt})$, and prices, $p_j^{i+1}(s_{mt})$, by fixed point iteration using the closed form for market shares from the demand model and firms' quantity-setting first order condition, the latter of which uses $\hat{CV}_j^{i+1}(s_{mt}; \hat{\lambda}^{i+1})$ in calculating the optimal dynamic markdown term.²⁴

²³I make several minor changes to the value function approximation approach used in estimation. In particular, I do not use simulated states to estimate the approximating coefficients. This is due to the fact that the approach to solving the model requires finding new equilibrium exit policies and equilibrium in the product market via fixed point iteration. Whereas it was simple to use first stage estimates to fit simulated exit policies and product market variables at simulated states, doing so for a large number of simulated states presents a computational challenge in solving the model, so I prioritize approximating the value function at the observed states.

²⁴In practice, convergence of the fixed point iteration on firms' quantity-setting first order condition is reliable and rapid. I iterate this procedure until the norm of the difference of the price vectors from successive iterations falls below 10^{-10} .

5. Check whether $\|p^{x,i+1}(s_{mt}) - p^{x,i}(s_{mt})\| < tol$, where $p^{x,i+1}(s_{mt})$ is the stacked vector of firms' exit policy functions and $tol = 10^{-4}$. If this condition is met, the iterations stop; if not, iteration $i + 1$ starts with step 1 above.

The above procedure produces conditional exit probabilities at each state in a given period as well as value function approximating coefficients. I use these value function approximating coefficients and a set of assumptions about the states of potential entrants to calculate conditional entry probabilities for that period. In particular, I use the resulting value function approximating coefficients and expected values of the state variables in the next period for all potential entrants to calculate conditional entry probabilities, where expected values of the state variables in the next period are calculated using the observed aggregate state variables and assuming that entrants are endowed with random values of the non-deterministic state variables drawn from the empirical distribution of observed states.²⁵

F.2 Simulating Forward

Armed with conditional exit and entry probabilities for incumbents and potential entrants in a given period from the first step, I can then implement the second step of solving the model: simulating the industry forward one period. In particular, I take a single draw from the conditional exit and entry probabilities and then implement firms' resulting, discrete exit and entry decisions. For the next period's new incumbents and potential entrants, I then take a single draw from the state transition processes estimated in the first stage of estimation. The simulated industry then proceeds to the next period and the fixed point algorithm outlined above is used to solve for policy functions in the next period. As noted above, I begin the counterfactual solution process at the observed data in period 1 of my estimation sample, 2008 Q1. I then repeat this process of solving a single period and simulating forward until the final period of my estimation sample, 2013 Q4. I therefore repeat this procedure 24 times, simulating the model forward 6 years or 24 quarters. Given that each time I simulate the industry forward I take single draws from the conditional exit and entry probabilities as well as the state transition processes, I repeat this process of simulating the model forward 6 years multiple times and average the results across the full set of forward-simulated industries. In practice, I repeat this process of simulating the model forward 6 years 60 distinct times for each counterfactual scenario and then average key outcomes across all 60 model runs.

One important idiosyncrasy in this step that is worth noting is how I model installed capacity. Given that my price and rebate fields are denominated on a per watt basis, I need to know firms' total installed capacity in watts in order to calculate firm profits; however, my demand model only predicts adoption, not the size of individual installed systems. In estimation, I observe total installed capacity in the data; however, nothing in my model allows me to predict this field in solving counterfactuals. Moreover, given that I define experience in terms of cumulative installed capacity (in watts), knowing watts of capacity installed each period is important for updating experience levels each period. I therefore take the simple approach of assuming that each installation predicted by my demand model

²⁵Naturally, entrants enter with zero experience. As in estimation, I calculate the expectation of future state variables for entrants conditional on entry by averaging state values over 1000 draws from the state transition processes estimated in the first stage.

has a capacity equal to the sample average system capacity in my processed estimation data, which is around 4-5 kW.

F.3 Subsidy Levels

Given that the main policy counterfactuals of interest involve adjusting the subsidy environment, it is worth discussing how I treat subsidy levels in solving counterfactuals. In the case of CSI rebates, despite the fact that these subsidy levels are conditional on cumulative installed capacity (see Figure G2), which itself is defined by lagged demand for solar PV installations, I choose to not endogenize the timing of CSI rebate changes when solving counterfactual scenarios with the CSI in place. While I could easily endogenize the CSI step changes in my model given that my model predicts demand for solar PV installations each period, I choose not to given that I do not explicitly model system capacity as discussed above. Moreover, the fact that I omit self-installations and subset the estimation data as described in Appendix A, even with my assumption about counterfactual installed system capacity described above, I would be guaranteed to under-predict cumulative installed capacity and therefore implement counterfactual subsidy levels that are higher than they likely should be. I therefore hold fixed the date of each CSI rebate step change from the observed data (shown in Figure A1) in any counterfactual that implements the CSI.

Implementing the federal investment tax credit (ITC) is relatively straightforward as this is just a fixed proportion of the post-rebate price. Implementing net energy metering (NEM) counterfactuals is more challenging given that utilities directly recover NEM payments through retail electricity rates. As a result, I hold NEM fixed in place in all counterfactuals.

F.4 Counterfactual Environmental Damages


Simulating the model forward as described above generates a set of key outcomes for each counterfactual scenario, including market structure outcomes (number of entries, number of exits, market concentration, etc.) as well counts of installed systems and measures of consumer surplus predicted from the demand model and total profits and cost components predicted from the supply model.

Given that a key policy justification for incentivizing the adoption of solar PV is to reduce electricity generation from legacy, alternative electricity generation sources such as coal and natural gas-fired power plants, I use the quantities of solar PV adoption predicted for each counterfactual scenario to conduct a back-of-the-envelope calculation of any avoided environmental damages from the solar PV subsidies. The external social benefits of solar PV subsidies are a function of the quantity of solar PV adopted due to subsidies, the amount of electricity produced by these systems, and the external damages associated with alternative electricity generation sources displaced by this additional solar capacity. I use estimates of the marginal environmental benefits of additional solar capacity in the US from [Sexton et al. \(2021\)](#). These estimates account for both the marginal external damages from harmful local air pollutants as well as carbon dioxide. Using rich data on electricity generation, solar insolation, and air pollution transport, [Sexton et al. \(2021\)](#) produce spatially-differentiated estimates of the marginal environmental benefits of additional solar capacity that account

for substantial heterogeneity in solar generation, displaced pollution emissions, and marginal costs of electricity over space and time. These off-the-shelf estimates therefore allow me to account for variation across the state of California in not only the lifetime generation potential of additional solar capacity, but also characteristics of the electricity grid.


G Supplemental Figures and Tables

Figure G1. Costs and Discounted Benefits of a 5 kW PV System in California, 2000-2020




Notes: This figure shows the evolution of the upfront installation cost and present discounted benefits of a representative 5 kilowatt (kW) residential PV system in California from 2000 to 2020. Data on average installation costs and rebates in California come from the Lawrence Berkeley National Lab's "Tracking the Sun" database (Barbose et al., 2022). Data on net retail electricity rates, expected PV output, and net metering policy come from the Energy Information Administration's Form 816, the World Bank Group's Global Solar Atlas, and California Public Utilities Commission materials, respectively. I assume a real interest rate of 3% to calculate present values and assume a system lifespan of 25 years, annual household PV energy consumption totaling 6000 kW-hours, and PV power potential equal to the average for the state in order to calculate electricity consumption and net metering benefits.

Figure G2. CSI Rebate Rate Structure




Notes: This figure shows the rebate levels under the California Solar Initiative (CSI) as a function of cumulative-installed capacity. This figure is inspired by a similar figure that appears in [Pless and Van Benthem \(2019\)](#).

Figure G3. Estimated Own Price Elasticities




Notes: This figure shows the distribution of estimates of the own price elasticities of demand for all firm-county-quarter observations in the main estimation sample estimated using the random coefficients nested logit model reported in Column (3) of Table 1. Estimates of the adoption decision elasticity from the literature are reported below the horizontal axis (Bollinger and Gillingham, 2019; De Groot and Verboven, 2019; Gerarden, 2022; Gillingham and Tsvetanov, 2019; Hughes and Podolefsky, 2015).

Figure G4. Relationship between Total PV Installation-related Employment, Wages and Installations




Notes: This figure shows the relationship between total quarterly PV installation-related employment and the average PV installation-related wage in a county and the quarterly total number of residential PV installations within that county. Data on employment levels and wages are from the US Census Bureau's Quarterly Census of Employment and Wages and include the quarterly number of workers in the roofing and electrician industries. The blue circles depict binned means; the blue line shows the linear relationship between the log of total PV installation-related employment/wages and the log of total installations; and the black points represent the raw data.

Figure G5. Average CSI Project Completion Time Added



Notes: This figure shows the quarterly average observed time added to residential solar PV installations that apply for California Solar Initiative (CSI) rebates separately for rebate program administrators and installers. Quarterly averages are reported separately by each of the three main investor-owned utilities (IOUs): San Diego Gas and Electric (SDG&E), Pacific Gas and Electric (PG&E), and Southern California Edison (SCE). Rebate-level data for the CSI obtained from the California Public Utilities Commission provide dates for detailed rebate processing milestones, which allow me to attribute cumulative time added to rebate processing due to the IOUs (left panel), which serve as the rebate program administrators, and individual installers (right panel). The lines show cubic b-splines, which I estimate separately for each IOU.

Figure G6. Average Permitting Time for PV Projects, San Diego County



Notes: This figure shows the quarterly average time-to-completion for solar photovoltaic (PV)-related permits in San Diego County, California. Permit-level data on historical developments permits are available from <https://data.sandiego.gov/datasets/development-permits-set1/> (last accessed August 10, 2023). These permit-level data provide dates for key project milestones, including the date a permit application is received and the date a permit is approved. To identify PV-related permits, I use an existing category of permit types in the data that distinguishes PV-related electrical permits; however use of this category appears to become widespread in the data starting in 2012. I therefore identify electrical permits in earlier years that are likely for PV-related projects by matching keywords (e.g., “PV,” “solar,” and various iterations of these terms) in a detailed project description field. The overlaid line shows a cubic b-spline fitted to the average time-to-completion data.

Table G1. Estimated Transition Processes for Aggregate State Variables

	Potential Market		Avg. Price (\$/W)		Inclusive Value	
	(1)	(2)	(3)	(4)	(5)	(6)
Intercept	-36.81 (14.71)		-0.200 (0.061)		0.260 (0.048)	
(Potential Market) _{t-1}	0.997 (0.0001)	0.994 (0.035)				
(Avg. Price) _{t-1}			0.982 (0.009)	0.983 (0.009)		
(Inclusive Value) _{t-1}					0.777 (0.055)	0.7112 (0.087)
County FE		Yes		Yes		Yes
Observations	363	363	363	363	363	363
R ²	0.99	0.99	0.91	0.91	0.56	0.58
Within R ²		0.74		0.91		0.44

This table reports estimates of the first-order autoregressive (AR(1)) transition processes for three county-quarter aggregate state variables, demand (number of potential adopters), average price per watt, and the inclusive value. The table reports two separate specifications for each state variable, one each with and without county-specific intercepts. Standard errors clustered at the county-level are reported in parentheses.

Table G2. Estimated Transition Processes for Firm-specific State Variables

	Own Quality		Hardware Cost (\$/W)		Price (\$/W)	
	(1)	(2)	(3)	(4)	(5)	(6)
Intercept	-1.342 (0.173)		0.263 (0.024)		0.729 (0.043)	
(Own Quality) _{t-1}	0.794 (0.026)	0.736 (0.042)				
(Hardware Cost) _{t-1}			0.878 (0.006)	0.876 (0.005)		
(Price) _{t-1}					0.840 (0.006)	0.835 (0.006)
County FE		Yes		Yes		Yes
Observations	6,527	6,527	6,527	6,527	6,527	6,527
R ²	0.70	0.71	0.77	0.77	0.72	0.72
Within R ²		0.54		0.76		0.71

This table reports estimates of the first-order autoregressive (AR(1)) transition processes for three firm-county-quarter state variables, own quality (which is derived from the demand system estimates), hardware cost per watt, and price. The table reports two separate specifications for each state variable, one each with and without county-specific intercepts. Standard errors clustered at the county-level are reported in parentheses.

Table G3. Estimated Production and Exit Cost Parameter Estimates for Different Discount Factors

Annual Discount Factor:	Parameter	$\beta = 0.9$ (1)	$\beta = 0.875$ (2)	$\beta = 0.85$ (3)
<i>Production Cost Parameters</i>				
Learning Exponent	γ	-0.274 (0.157)	-0.363 (0.150)	-0.606 (0.327)
Hardware Cost	τ	-0.673 (0.027)	-0.654 (0.050)	-0.591 (0.094)
Base Cost	c_0	0.153 (0.342)	0.643 (0.256)	0.637 (3.524)
Productivity Serial Correlation	ρ	0.420 (0.003)	0.420 (0.007)	0.424 (0.015)
<i>Experience Parameters</i>				
Rival Experience: In-market	θ_1^E	0.641 (0.281)	0.735 (0.253)	1.540 (3.570)
<i>Exit Parameter</i>				
Mean Scrap Value	σ_ϕ	82.655 (4.668)	49.698 (3.062)	36.597 (2.960)
Installer, County, Year FE		Yes	Yes	Yes
N		7,351	7,351	7,351
Spence Coef. $(1 - 2^\gamma)$		0.173	0.223	0.343

Notes: Estimation follows the procedure outlined in Section 4.2. There are 7,351 observations at the installer-county-half-year level. I normalize experience variables by the industry total experience level in the first half year of the sample (H1 2008). All effective experience parameters can be interpreted as marginal experience contributions relative to a firm's own experience. The "forgetting parameter," δ , describes the rate of learning depreciation from one period to another. The mean scrap value parameter is measured in 100,000 2013 USD. The "Spence Coefficient" describes the proportional reduction in cost from a doubling of effective experience. Standard errors are calculated using the Bayesian Bootstrap with bootstrap weights clustered by county (Rubin, 1981). Bootstrap weights for each county are drawn according to a Dirichlet distribution with $\alpha = 1$ across 200 bootstrap samples.

Table G4. Estimated Production and Exit Cost Parameter Estimates with Scale as a State Variable

	Parameter	State Variables:	
		Baseline	+Firm Size
Production Cost Parameters			
Learning Exponent	γ	-0.363 (0.150)	-0.371 (0.162)
Hardware Cost	τ	-0.654 (0.050)	-0.602 (0.035)
Base Cost	c_0	0.643 (0.256)	0.748 (0.233)
Productivity Serial Correlation	ρ	0.420 (0.007)	0.426 (0.004)
Experience Parameters			
Rival Experience: In-market	θ_1^E	0.735 (0.253)	0.755 (0.438)
Exit Parameter			
Mean Scrap Value	σ_ϕ	49.698 (3.062)	43.640 (1.859)
Installer, County, Year FE		Yes	Yes
N		7,351	7,351
Spence Coef. $(1 - 2^\gamma)$		0.223	0.227

Notes: Estimation follows the procedure outlined in Section 4.2. Column (1) corresponds to the baseline preferred specification of the model in Column (1) of Table 2. Column (2) re-estimates this specification but adds firms' current period size (i.e., production quantity) as a state variable to account for static economies of scale. I normalize experience variables by the industry total experience level in the first half year of the sample (H1 2008). All effective experience parameters can be interpreted as marginal experience contributions relative to a firm's own experience. The "forgetting parameter," δ , describes the rate of learning depreciation from one period to another. The mean scrap value parameter is measured in 100,000 2013 USD. The "Spence Coefficient" describes the proportional reduction in cost from a doubling of effective experience. Standard errors are calculated using the Bayesian Bootstrap with bootstrap weights clustered by county (Rubin, 1981). Bootstrap weights for each county are drawn according to a Dirichlet distribution with $\alpha = 1$ across 200 bootstrap samples.

References

Arcidiacono, Peter, and Robert A. Miller. 2011. “Conditional Choice Probability Estimation of Dynamic Discrete Choice Models With Unobserved Heterogeneity.” *Econometrica*, 79(6): 1823–1867.

Barbose, Galen, Naïm Darghouth, Eric O’Shaughnessy, and Sydney Forrester. 2022. “Tracking the Sun: Pricing and Design Trends for Distributed Photovoltaic Systems in the United States, 2022 Edition.” Lawrence Berkeley National Laboratory.

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

Barwick, Panle Jia, Myrto Kalouptsidi, and Nahim Bin Zahur. 2025. “Industrial Policy Implementation: Empirical Evidence from China’s Shipbuilding Industry.” *The Review of Economic Studies*, 92(6): 3611–3648.

Benitez-Silva, Hugo, John Rust, Gunter Hitsch, Giorgio Pauletto, and George Hall. 2000. “A Comparison Of Discrete And Parametric Methods For Continuous-State Dynamic Programming Problems.” *Computing in Economics and Finance 2000*.

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

Berry, Steven, James Levinsohn, and Ariel Pakes. 1995. “Automobile Prices in Market Equilibrium.” *Econometrica*, 63(4): 841–890.

Berry, Steven, James Levinsohn, and Ariel Pakes. 1999. “Voluntary Export Restraints on Automobiles: Evaluating a Trade Policy.” *American Economic Review*, 89(3): 400–430.

Berry, Steven T. 1994. “Estimating Discrete-Choice Models of Product Differentiation.” *The RAND Journal of Economics*, 25(2): 242–262.

Bollinger, Bryan, and Kenneth Gillingham. 2019. “Learning-by-Doing in Solar Photovoltaic Installations.” *Working Paper*.

Borenstein, Severin, and James B. Bushnell. 2022. “Do Two Electricity Pricing Wrongs Make a Right? Cost Recovery, Externalities, and Efficiency.” *American Economic Journal: Economic Policy*, 14(4): 80–110.

California State Senate. 2006. “S.B. 1 - Electricity: solar energy: net metering.” 2005–2006 Legislative Session.

Darghouth, Naïm R., Galen Barbose, and Ryan Wiser. 2011. “The impact of rate design and net metering on the bill savings from distributed PV for residential customers in California.” *Energy Policy*, 39(9): 5243–5253.

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

Feger, Fabian, Nicola Pavanini, and Doina Radulescu. 2022. “Welfare and Redistribution in Residential Electricity Markets with Solar Power.” *The Review of Economic Studies*, 89(6): 3267–3302.

Gerarden, Todd D. 2022. “Demanding Innovation: The Impact of Consumer Subsidies on Solar Panel Production Costs.” *Management Science*, Forthcoming.

Gillingham, Kenneth, and Tsvetan Tsvetanov. 2019. “Hurdles and steps: Estimating demand for solar photovoltaics.” *Quantitative Economics*, 10(1): 275–310.

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

Hughes, Jonathan E., and Molly 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.

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

Langer, Ashley, and Derek Lemoine. 2022. “Designing Dynamic Subsidies to Spur Adoption of New Technologies.” *Journal of the Association of Environmental and Resource Economists*, 9(6): 1197–1234.

Pakes, Ariel, Michael Ostrovsky, and Steven Berry. 2007. “Simple estimators for the parameters of discrete dynamic games (with entry/exit examples).” *The RAND Journal of Economics*, 38(2): 373–399.

Pless, Jacquelyn, and Arthur A. Van Benthem. 2019. “Pass-Through as a Test for Market Power: An Application to Solar Subsidies.” *American Economic Journal: Applied Economics*, 11(4): 367–401.

Rubin, Donald B. 1981. “The Bayesian Bootstrap.” *The Annals of Statistics*, 9(1): 130–134. Publisher: Institute of Mathematical Statistics.

Sexton, Steven, A. Justin Kirkpatrick, Robert I. Harris, and Nicholas Z. Muller. 2021. “Heterogeneous Solar Capacity Benefits, Appropriability, and the Costs of Sub-optimal Siting.” *Journal of the Association of Environmental and Resource Economists*, 8(6): 1209–1244.

Sweeting, Andrew. 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.