Learning-by-doing and Market Structure: Empirical Evidence from the US Residential Solar Industry*

Jacob T. Bradt †

March 13, 2025

Abstract

Consumer subsidies are a common policy tool for supporting the adoption of clean energy technologies. Policymakers often justify these programs as a means of stimulating infant industries, arguing that greater demand increases industry learning-by-doing, which in turn reduces costs for potential entrants. This requires learning spillovers between firms that make experience-based cost reductions a public good. However, spillovers can reduce firms' incentives to expand output and lower costs. To evaluate this tradeoff, I estimate a dynamic structural model of the market for solar panel installations in California that endogenizes firms' entry and exit decisions and allows for learning-by-doing with knowledge spillovers. I find that a 1% increase in a firm's cumulative production leads to a 0.7% reduction in installation-specific costs and that learning spills over across firms. Counterfactual analysis reveals that a state-level consumer subsidy program increased solar adoption by 4% and installer entry by 9%, indicating that industry cost reductions outweigh any strategic incentives for firms to reduce learning. While consumer subsidies achieve industry growth, I find that standard industrial policies such as entry subsidies provide greater welfare gains.

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, Marcos Barrozo, Judd Boomhower, Andrés de Loera, Jackson Dorsey, Teevrat Garg, Todd Gerarden, Ken Gillingham, Ed Glaeser, Jeff Gortmaker, Catherine Hausman, Allan Hsiao, Mark Jacobsen, Akshaya Jha, Devesh Raval, Rob Stavins, Andrew Steck, Rich Sweeney, and Chris Timmins as well as seminar and conference participants at EARIE 2024, SITE 2024, ES-DSE 2024, BSE Summer Forum 2024, IIOC 2024, Cornell, The University of Texas at Austin, UC San Diego, The University of California Virtual Environmental Economics Seminar, 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.

†Harvard University, John F. Kennedy School of Government; jbradt@g.harvard.edu.

1 Introduction

Policymakers view subsidies to consumers as a win-win, providing constituents with resources to address externalities while also spurring growth and entry in target industries. This is apparent in efforts to decarbonize the United States economy. Facing political constraints on the use of first-best policy instruments such as a carbon tax, US policymakers appear to prefer decarbonization policies that target growth in specific clean energy technologies. For example, the Inflation Reduction Act (IRA) of 2022 provides $392 billion in tax credits and direct expenditures for clean technologies, including over $50 billion in tax credits for consumers (CBO, 2022). A number of provisions in the IRA, which represents the largest federal response to climate change to date, explicitly target growth in domestic employment and supply chains for clean energy technologies by including domestic sourcing and labor requirements (Bistline et al., 2023). This shift in political revealed preference towards second-best climate policies raises questions around how to best design and implement these politically-feasible policy instruments.

The policy debate surrounding consumer subsidies for clean technologies highlights the potential for industry growth, but can consumer subsidies encourage firms to enter infant industries? By increasing demand, consumer subsidies increase industry-wide experience-based cost reductions, or learning-by-doing, resulting in lower production costs for potential entrants. However, since a firm's rivals can benefit from a portion of the cost reductions the firm achieves, potential knowledge spillovers reduce the incentives incumbent firms face to lower costs by expanding output (Ghemawat and Spence, 1985). The net effect of consumer subsidies on growth and entry in target industries is therefore an empirical question, one which depends on the rate of learning-by-doing and the size of knowledge spillovers.

In this paper, I study consumer subsidies in the market for solar photovoltaics (PV) and their impact on growth in the PV installation industry. To do so, I build and estimate a dynamic structural model of the market for solar panel installations that endogenizes firms' entry and exit decisions and accounts for learning-by-doing with knowledge spillovers. Using data on the residential solar PV market in California from 2008 to 2013, I find evidence of substantial learning-by-doing and knowledge spillovers across firms. Counterfactual analysis suggests that a state-level consumer subsidy program expanded the solar PV industry.

Solar PV is a key technology for mitigating the catastrophic consequences of climate change due to its minimal life cycle greenhouse gas emissions and its ability to displace conventional, greenhouse gas-intensive electricity sources. Policymakers have provided substantial public subsidies for the adoption of solar PV, with many programs targeting residential consumers. The non-trivial design and construction of PV systems has given rise

to an industry of intermediary firms that offer PV installations as a service. These intermediary installers play a major role in the solar market: solar installation firms employed over 171,000 workers in the US in 2022, about 65% of total US employment in solar PV (Interstate Renewable Energy Council, 2023). Moreover, a growing share of the cost to end consumers is attributable to local installation costs. For example, Barbose et al. (2022) estimate that the share of US residential consumers' final costs attributable to installers rose from around 40% in 2006 to over 80% in 2016. Despite accounting for a growing share of final costs, the cost components attributable to installers have fallen in recent years, with existing anecdotal and empirical evidence suggesting potential installer learning-by-doing (Bollinger and Gillingham, 2019; Fu et al., 2016; Nemet, 2019).

In comparison to PV manufacturers, relatively little is known about the intermediary installer market despite its growing contribution to costs. California offers an ideal setting in which to study this component of the solar industry for two reasons. First, California is home to nearly half of all US residential PV systems, making it the largest market for residential solar in the US (Barbose et al., 2022). Second, California has a long history of generous consumer subsidies for solar PV, which in many cases sought to not only increase take-up of this clean technology, but also spur industry growth in the state. For example, the California Solar Initiative (CSI), which ran from 2007 to 2013 with a budget of $2.2 billion, provided direct cash rebates to households. In addition to targeting the decarbonization of California's electrical grid, an explicit goal of the CSI was to "establish a self-sufficient solar industry" (California State Senate, 2006).

To evaluate this goal and shed light on the efficacy of consumer subsidies as an unconventional form of industrial policy, I develop a dynamic structural model of the market for residential solar PV installations in California based on the theoretical framework for dynamic oligopoly of Ericson and Pakes (1995). A structural model is necessary to address this question for several reasons. First, the model allows me to estimate firms' profit margins, which together with the data on hardware costs isolates the component of production costs directly associated with installation. Second, I am ultimately interested in measuring—among other outcomes—changes in firms' exit and entry behavior, both of which are inherently dynamic decisions: they involve consideration of expected future earnings, which a model enables me to estimate. Finally, developing and estimating a structural model allows me to simulate industry outcomes under counterfactual policy environments.

The model endogenizes consumer demand for solar PV installations as well as installation firms' entry, exit, and quantity-setting decisions. Consumer demand for differentiated solar PV installations is static and follows the random coefficient nested logit model of Brenkers and Verboven (2006). Incumbent installers' installation-specific costs are a function of their

own cumulative production as well as the cumulative production of their rivals to allow for learning-by-doing and learning spillovers. In each geographic market, incumbent firms dynamically choose a quantity of installations to provide conditional on their marginal costs and their beliefs about future learning. I model firms' product market decisions as dynamic to capture the incentives to select a production level today based on its impact on own and rival costs in the future. Incumbent firms then choose whether to exit by comparing their expected discounted future profits with an idiosyncratic scrap value while a market-specific pool of potential entrants make one-shot entry decisions based on their expected discounted future profits and an idiosyncratic cost of entry. Firms' strategies lead to a Markov Perfect Equilibrium, which I assume is well-approximated by a Moment-based Markov Equilibrium concept (Ifrach and Weintraub, 2017).

I estimate the model using detailed, system-level data on prices, rebates, installed capacities, and hardware costs for 95% of all California residential PV systems installed between 2008 to 2013. I acquire the bulk of the system-level data from the Lawrence Berkeley National Laboratory's "Tracking the Sun" database, including information on the timing and location of each installation as well as the identity of the installing firm. I combine these data with information on system-level hardware costs in California, which I acquire from the California Public Utilities Commission. The hardware cost data are important as they allow me to isolate the component of installers' marginal costs in which learning-by-doing could occur, which I otherwise do not observe. I aggregate these system-level data to the county-quarter level for all installers that I observe operating between 2008 and 2013.

My approach to estimation builds on the family of two-step estimators of dynamic games and their various applications (Bajari et al., 2007; Collard-Wexler, 2013; Fowlie et al., 2016; Pakes et al., 2007; Ryan, 2012). The main components of the model that I estimate include: the demand system for residential solar PV installations, firms' marginal installation cost function, and the distributions of scrap values and entry costs. In the first stage, I estimate the static demand parameters and flexibly estimate the exit policy function and transition process of state variables from the data. I use the first stage estimates to obtain a flexible approximation of firms' value functions—which approximate expected discounted future profits—following recent work in the dynamic games literature (Barwick and Pathak, 2015; Barwick et al., 2021; Kalouptsidi, 2018; Sweeting, 2013). In the second stage, I form moments from the model's optimal quantity-setting and exit conditions to first recover the parameters governing production costs (i.e., learning) and exit. I then use these estimates to formulate the likelihood of observed entry decisions to recover the full set of dynamic parameters of interest.

The model estimates reveal two main findings. First, I find evidence of substantial

learning-by-doing in California's residential solar PV installation industry. Specifically, I estimate that a 1% increase in a firm's experience decreases marginal installation costs between 0.7 and 0.9%. This translates to an average reduction in marginal installation costs of $1.12 per watt (W) from the first quarter of 2008 to the final quarter of 2013. For context, the average total system cost in the sample declined from $8.63/W to $3.55/W, implying that installer learning accounts for roughly 25% of observed reductions in final costs.

Second, I find that learning spills over across installation firms. The learning benefits from a 1 unit increase in total industry experience generates 82% of the learning benefit of a 1 unit increase in a firm's own experience. This implies that learning spillovers from rivals are substantial and not far from individual learning in terms of their contribution to firms' experience-based cost reductions. To explore the underlying mechanisms, I estimate alternative specifications of the model that allow for differential spillovers across firms based on the source of rivals' experience. I find differences in the size of spillovers when I allow for differential contributions to learning from rivals based on their geographic market, suggesting that knowledge transfer between firms is facilitated by some mechanism that is stronger at the market-level. Examples might include: the movement of workers between firms, the visibility of rivals' installation practices, or passive learning by market-level regulatory regimes.

I use the estimated model to evaluate the implications of various counterfactual policy environments. These counterfactuals reveal two main findings. First, I find that the CSI, the state's main consumer subsidy policy, increased not only solar adoption by 4%, but also the number of operating firms by 9%. Counterfactual simulations that assume smaller learning spillovers than my model estimates reduce the CSI's impact. This suggests that knowledge transfer across firms is the main mechanism through which consumer subsidies lead to additional entry and industry growth. I also compare the existing CSI program to alternative rebate designs and other forms of climate policy, including a $30/ton carbon tax. Focusing solely on outcomes in the state's solar PV market, I find that the CSI's decreasing rebate structure results in higher welfare relative to alternative consumer subsidy designs and climate policies. Thus, if policymakers approach to decarbonizing the economy involves targeting growth in specific clean technologies—perhaps as a result of political constraints on first-best policy tools—consumer subsidies may be an effective means of doing so.

Second, I find that consumer subsidies may not be as effective as more conventional forms of industrial policy such as entry subsidies. I implement a set of counterfactual policies in which I remove the CSI subsidies and replace them with entry subsidies of varying sizes. The entry subsidies greatly increase not only the number of active incumbents, but also the number of solar PV installations. Moreover, each of the counterfactual entry subsidies greatly increases total welfare, mostly through greater consumer surplus and lower net-of-subsidy

total entry costs. The difference in performance of the CSI and entry subsidies is driven by the targeting of each type of subsidy: the CSI's consumer subsidies primarily work to decrease exit and raise the profits of inframarginal incumbents that would be active without the subsidies. This finding suggests that, though consumer subsidies may be an appealing second-best tool for policymakers to achieve decarbonization and industrial policy goals in specific technologies, other approaches may be more effective on these outcomes.

These findings contribute to a large literature on learning economies by estimating a model of learning-by-doing with endogenous market structure. Theoretical work shows that cumulative experience can impact market outcomes (Arrow, 1962; Besanko et al., 2010; Cabral and Riordan, 1994; Fudenberg and Tirole, 1983; Spence, 1981). Ghemawat and Spence (1985) use numerical simulations to show that non-appropriable learning can influence market structure by undercutting barriers to entry. A vast empirical literature seeks to explore these theoretical results by recovering learning curves in specific contexts, including the production of 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). Several studies have found learning spillovers in different contexts (Covert, 2015; Irwin and Klenow, 1994; Kellogg, 2011; Thornton and Thompson, 2001). In related work, Bollinger and Gillingham (2019) estimate learning-by-doing by solar PV installers in California and find non-zero spillovers across firms.[1] My findings build on this literature by endogenizing market structure, thereby connecting empirical findings of learning spillovers to early theoretical work.

This work also contributes to a large literature on industrial policy, which has grown in recent years as policymakers increase their use of these tools (Juhász et al., 2023). Empirical analysis of industrial policy focuses on various forms of trade policy, including research and development subsidies (Bloom et al., 2002; Hall and Van Reenen, 2000), production subsidies (Barwick et al., 2021; Kalouptsidi, 2018), place-based policies including trade adjustment assistance (Kline and Moretti, 2014), and environmental subsidies (Aldy et al., 2022). Despite the fact that policymakers often justify demand subsidies on industrial policy grounds, there is comparatively little research evaluating the effectiveness of such programs for achieving these objectives. The findings of my analysis expand the existing rigorous evidence to include less conventional forms of industrial policy.

These results emphasize the importance of evaluating a broad set of equilibrium out-

[1] While the ongoing work of Bollinger and Gillingham (2019) is quite similar to mine, there are several key differences worth noting. First and most notably, my model endogenizes firms' entry and exit decisions whereas Bollinger and Gillingham (2019) do not model these decisions and hold them fixed in the data. Second, Bollinger and Gillingham (2019) do not account for the potential for serially-correlated productivity shocks in their production cost estimation, whereas my estimation approach does.

comes when studying the impact of subsidy policies. A large and growing literature studies the economics of solar PV policies, focusing primarily on the impact of subsidies on adoption rates, finding that while consumer subsidies have increased the adoption of solar PV, these policies are not justified by the static environmental benefits of adoption alone (Borenstein, 2017; De Groote and Verboven, 2019; Gillingham and Tsvetanov, 2019; Hughes and Podolefsky, 2015). In contrast, Gerarden (2022) demonstrates that accounting for manufacturer innovation decisions can justify the level of solar PV subsidies. A broader set of papers examine dynamic issues in other clean technologies—for example, Armitage (2022) evaluates the timing of policy support for efficient lighting. My analysis is closely related to van Benthem et al. (2008) and Langer and Lemoine (2022), who derive theoretical results and demonstrate via numerical simulation that accounting for dynamic considerations such as persistent price reductions and passive learning-by-doing can rationalize existing PV subsidy design. My results contribute to this literature by providing empirical evidence that accounting for learning and evaluating market size and structure provides a more complete assessment of the impact of clean technology subsidies.

The rest of the paper is organized as follows. Section 2 provides an overview of the solar PV industry and policy environment. Section 3 discusses the data I use in my analysis and provides some descriptive results on solar PV installers in California. Section 4 presents the model. Sections 5 and 6 describe estimation and the model estimates. Finally, Section 7 presents results from counterfactual policy simulations while Section 8 concludes.

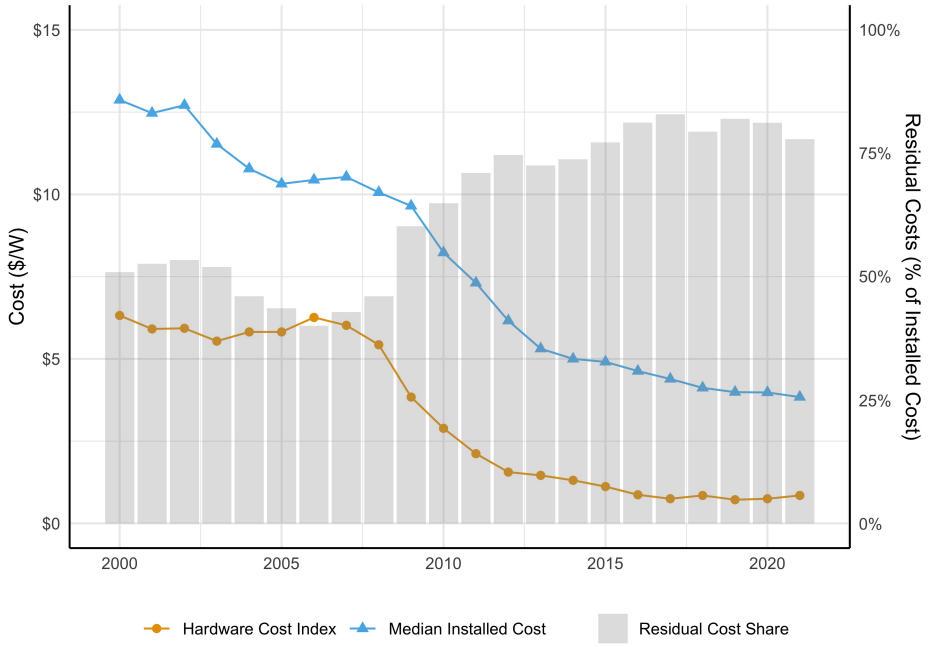
2 Economic and Policy Landscape

2.1 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—more colloquially referred to as panels—consist of interconnected solar cells, which convert sunlight into electricity via the photovoltaic effect. Since the creation of the first practical solar cell by Bell Labs in 1954, steady technological innovation and improved manufacturing efficiency have contributed to substantial reductions in the cost of solar modules (Nemet, 2019). From 1975 to 2021, the price of solar modules declined over 99%, from $115 to under $0.5 (2021 USD) per watt (W) of capacity (IRENA, 2022; Nemet, 2009). Unsurprisingly, adoption of solar technology increased dramatically during this period: global solar capacity grew from just under 1 gigawatt (GW) to over 1 terrawatt (TW) from 2000 to 2022, a 1000-fold increase (IRENA, 2023).

Solar PV is a modular technology that is manufactured at-scale. This allows for a wide range of applications, from utility-scale power generation to small, distributed residential

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



Notes: This figure shows estimates of the installed cost—i.e., the cost paid by the end consumer—and hardware component cost per watt of installed capacity for residential photovoltaic (PV) systems in the US from the Lawrence Berkeley National Lab's "Tracking the Sun" report (Barbose et al., 2022). Estimates of hardware costs include the cost of PV modules and inverters. Bars show the share of the median installed cost attributable to components other than hardware costs (i.e., installation labor, permitting costs, installer markups, etc.) over time.

systems, the latter of which are the focus of this paper. The installation of residential PV systems requires non-trivial design and construction. The fact that installation occurs in most cases on consumers' rooftops introduces site-specific features that require idiosyncratic design solutions.[2] The practical challenges associated with rooftop installation as well as the technical nature of a PV system's electrical components means that specialized labor is necessary for construction. These features combined with a convoluted regulatory environment that involves different permitting and inspection requirements across jurisdictions and, in many cases, generous but difficult-to-navigate incentive programs has given rise to an industry marketing PV installation as a service. These installation firms source hardware inputs; design and construct systems; and manage the necessary permitting, inspection, and other administrative processes for households.

While remarkable reductions in the cost of manufactured solar panels have occurred at a global scale, an increasing share of the cost to the end consumer is attributable to installers

[2]No more than 2% of solar PV adopting households installed ground-mounted systems in the US from 2000 to 2021 (Barbose et al., 2022).

operating at the local level. Figure 1 shows data collected by Lawrence Berkeley National Lab on solar PV hardware costs and median installed costs for residential consumers in the US over 2000-2021. As shown by the bars in Figure 1, the share of residential consumer's installation costs that is attributable to installers (e.g., installation labor costs, permitting costs, and installer markups) grew over this period, increasing from a low of around 40% in 2006 to a high of over 80% in 2016 (Barbose et al., 2022).

Despite accounting for a growing share of final costs, these cost components attributable to installers—so-called "soft" costs—have also fallen in recent years; however, not nearly as fast as hardware costs (Fu et al., 2016). Existing anecdotal and empirical evidence indicates that these soft cost reductions have been driven by a number of factors, including local installers improving processes and learning from one another as well as streamlined policies around PV installation (Bollinger and Gillingham, 2019; Nemet, 2019; Nemet et al., 2017). A key issue for these existing estimates of declines in soft costs and their drivers is the difficulty of separating installer costs and markups, with Bollinger and Gillingham (2019) a noteworthy exception in their explicit modelling of installer pricing incentives. Determining the magnitude and sources of any reductions in installation-specific costs remains an empirical question and is a key objective of this paper.

2.2 California's PV Policy Environment

California is home to nearly half of all US residential solar PV systems (Barbose et al., 2022). While California's sunny climate likely drives much of the observed high solar adoption rates, generous PV adoption policies also play a major role. Since the late 2000s, California households have been eligible for a wide range of adoption incentives at both the state and federal levels. I discuss several key incentive programs below. Figure A1 shows how adoption subsidies and the consumption benefits of residential solar compare with the average upfront installation cost over time for California households.

2.2.1 State Rebates and Federal Tax Credits

Under the California Solar Initiative (CSI), the largest direct rebate program available to consumers in the state, certain households were eligible for direct cash rebates based on the quantity of PV capacity they installed. The CSI program started in 2007 with a planned 10-year budget of $2.2 billion—although the program was exhausted by mid-2013—and provided rebates to customers of the state's three main investor-owned Utilities (IOUs): Pacific Gas and Electric (PG&E), Southern California Edison (SCE), and San Diego Gas and Electric (SDG&E), who together service the majority of California ratepayers. Funding for the CSI

came directly from the IOU ratepayers, so rather than a fiscal outlay, the program functioned as a transfer between households in the state.

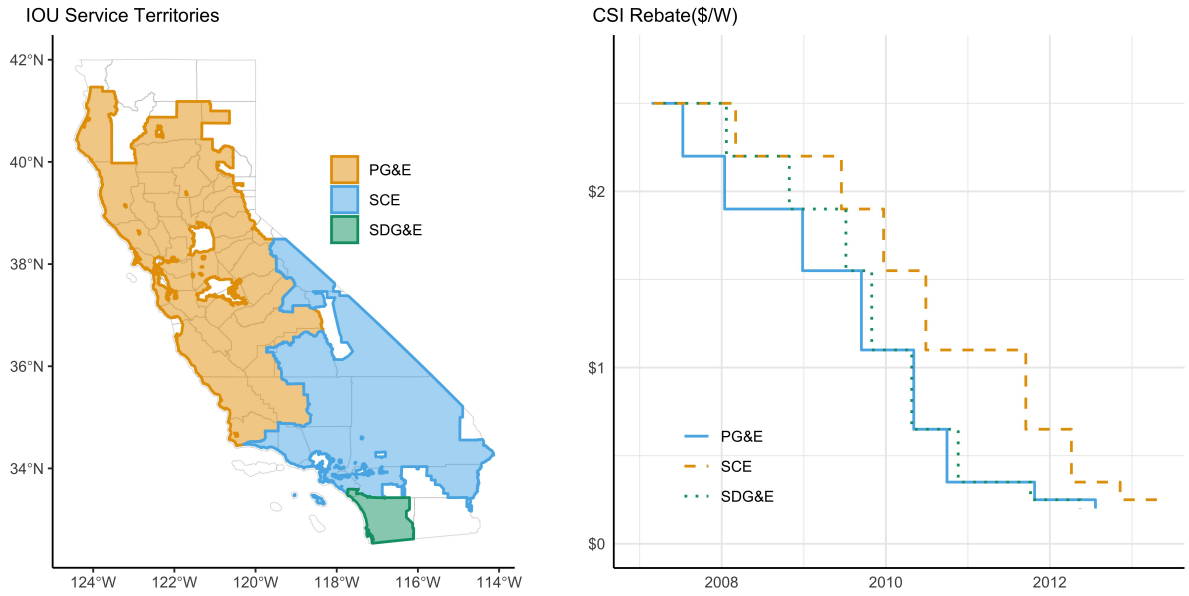
The schedule of rebates under the CSI was unique in that it was designed explicitly with learning-by-doing and industry cost savings in mind. Rebate levels started at $2.50 per watt (W) of installed system capacity and then gradually stepped down over 10 rate levels based on cumulative installed capacity in each IOU's service area (see Appendix Figure A2 for the full CSI rebate schedule). This subsidy design assumes that as industry experience accumulates, industry costs decline, thereby reducing the quantity of rebates necessary to incentivize adoption. Figure 2 shows the resulting spatial and temporal variation in rebates available under the CSI across the three main IOUs. As is clear from Figure 2, the rebate steps change at different times for each IOU, with consumers and firms facing sharp drops in rebate generosity at each step change based on the cumulative quantity of capacity installed throughout an IOU service territory. These sharp changes in rebate levels across IOUs and over time provides plausibly-exogenous variation in net-of-rebate prices that is useful in estimation as I discuss in Section 5.

An interesting feature of the CSI is the motivation for its creation. The CSI came out of the "Million Solar Roofs Initiative," a program announced by California's governor in August 2004 with the explicit goal of achieving one million residential solar installations by 2016. In addition to reducing emissions of harmful greenhouse gases and local air pollutants from California's electrical grid, an express-stated goal of the CSI was to "establish a self-sufficient solar industry" (California State Senate, 2006). To do so, the CSI aimed to "transform the market for solar energy by reducing the cost of solar" (CPUC, 2009). A primary motivation of this paper is to evaluate this objective by estimating the extent to which the CSI reduced the installed cost of solar PV systems and changed the structure of the installation industry.

Solar-installing households have also been eligible for a 30% non-refundable federal tax credit on the cost of an installed PV system since 2007. Prior to 2009, this investment tax credit (ITC) was capped at $2,000 for claiming households; however, since 2009, the residential ITC has been uncapped. Importantly, the federal ITC is applied net of any direct rebates a household receives. With net-of-subsidy system costs ranging anywhere from $20,000 to $60,000 (2013 USD) over 2000-2020, the federal ITC represents a substantial subsidy towards the upfront cost of PV adoption.[3]

[3]California had its own investment tax credit (ITC) program prior to 2006. The program gave claiming taxpayers a 15% non-refundable state tax credit prior to 2004, declining to a 7.5% credit for 2004-2005.

Figure 2. Spatial and Temporal Variation in CSI Rebates across IOUs



Notes: This figure shows the spatial and temporal variation in rebates available under the California Solar Initiative (CSI). The variation is driven by the spatial distribution of each of the three main investor owned utilities (IOUs), Pacific Gas & Electric (PG&E), Southern California Edison (SCE), and San Diego Gas & Electric (SDG&E), which is shown in the left panel, as well as the cumulative installed capacity of solar PV in each IOU's service territory. Since the level of the CSI rebate is a function of cumulative installed capacity (see Appendix Figure A2), this drives the variation in rebate levels over time across IOUs, which is shown in the right panel. County boundaries are shown in grey in the left panel. Sources: California Public Utilities Commission and US Census Bureau.

2.2.2 Net Energy Metering and Retail Electricity Pricing

Another way in which California subsidizes households' adoption of solar PV is through the structure of retail electricity rates and a policy known as net energy metering (NEM). California IOUs collect most of their residential revenue through increasing-block pricing, a volumetric rate structure that increases the marginal price of electricity as a household's total consumption increases within a billing period. Since these residential rates have little to no fixed monthly charge or non-volumetric charge, Borenstein (2017) points out that high-consumption households face greater returns from PV installation. NEM, which was first adopted in the 1990s, is a policy that requires the IOUs to purchase any excess electricity generated by grid-connected households with solar PV at their marginal retail electricity rate for the billing period. In effect, this policy allows households to run their electricity meters backwards anytime their PV system is generating more electricity than they are consuming. Given that marginal retail electricity rates far exceed marginal generation costs,

this represents a substantial subsidy to PV-adopting households (Borenstein, 2017).[4]

3 Data and Descriptive Evidence

3.1 Data Sources

In this section, I summarize my primary data sources and sample restrictions. A comprehensive discussion of the data used in my empirical analysis is available in Appendix B. I construct a dataset that tracks individual installation firms' average prices, market shares, hardware costs, and experience as well as entry and exit decisions across different markets—which I define at the county-level—and quarters.

I collect information on individual solar PV installers from the Lawrence Berkeley National Laboratory's (LBNL) "Tracking the Sun" database (Barbose et al., 2022). LBNL collects system-level data on PV systems annually from state agencies and utilities that administer PV incentive programs, renewable energy credit registration systems, or grid interconnection processes. The database includes information on the date of installation, system size, total installed price, total pre-tax rebate value, customer type (i.e., residential or commercial), zip code, mounting type (i.e., ground- or roof-mounted), and installer name, as well as various technical details about the installed hardware, including the energy conversion efficiency (i.e., how much incoming solar radiation a panel converts into electrical power), make, and manufacturer of installed PV modules. The full sample includes data on over 2.5 million PV systems installed from 2000 to 2021, which is approximately 77% of the total estimated US market for PV systems over that period. The data for California cover over 98% of total installations in the state from 2000 to 2021(Barbose et al., 2022).[5]

I subset these system-level data in several ways before constructing the main datasets used in my empirical analysis. First, I restrict the data to only include systems installed in California and further remove all non-residential systems, all ground-mounted systems, any residential systems with capacity exceeding 20 kW, and any residential systems for which I do not observe the installed price or rebate information.[6] I also exclude self-installed systems

[4]Concerns around the generosity of NEM and the potential for utilities to recoup NEM-related costs from non-NEM households led to the introduction of an IOU-specific fixed fee and volumetric charge for NEM systems, which the CPUC collectively referred to as NEM 2.0. These fixed interconnection fees and volumetric charges were designed to help the IOUs cover costs associated with the NEM subsidy without affecting non-NEM rates. CPUC rolled out NEM 2.0 across the IOUs over 2016-2017. Further reforms to NEM are in progress as of August 2023.

[5]In the case of California, the California Public Utilities Commission (CPUC) publishes detailed system-level data on all grid interconnected PV systems and LBNL uses these data in constructing the Tracking the Sun database. As a result, coverage of the California market in the Tracking the Sun database represents the universe of grid-connected systems, which in turn accounts for the vast majority of installed PV systems.

[6]I choose to remove residential systems exceeding 20 kW in capacity as these are outliers likely associated

Table 1. Summary Statistics for Processed Installer-level Data

	Mean	SD	Min	Max
Number of Installations	2.59	4.02	1	138
Total Installed Capacity (kW)	14.36	20.30	0.92	583.95
Market Share (%)	0.00	0.01	0.00	0.35
Market Share: Inside (%)	4.66	8.04	0.13	100.00
Average Installed Price (2013 $/W)	6.69	1.79	1.54	12.03
Average Hardware Cost (2013 $/W)	4.25	1.54	0.00	10.58
Own Experience: In-market (kW)	128.37	275.11	0.00	3643.05
Own Experience: Out-of-market (kW)	1184.95	2706.12	0.00	18 355.82
Rival Experience: In-market (MW)	11.12	10.49	0.13	52.00
Rival Experience: Out-of-market (MW)	253.10	97.41	81.79	444.24
Rival Experience: Same Manufacturer (MW)	24.02	28.72	0.00	108.49
Rival Experience: Other Manufacturer (MW)	240.19	99.24	66.61	445.04
N		17,852		

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-quarter-level, so descriptive statistics pool observations across markets and quarters. Total installed capacity in a quarter and the measures of firms' own experience are in kilowatts (kW), whereas rivals' experience measures are in megawatts (MW), or 1000 kW.

as well as any systems owned by a third party.[7]

To estimate potential learning economies in installation-specific costs, it is necessary to remove hardware costs from the installed price of residential PV systems as this cost component happens upstream of installation firms. I acquire data on hardware costs associated with specific residential installations in California from the CPUC.[8] Though not a required component of rebate applications, the CPUC collected these data directly from installers and households during the CSI rebate application process. As a result, hardware cost data are unavailable for systems in my processed data that did not apply for the CSI rebate or otherwise did not report system hardware costs in their CSI rebate application. Coverage is most complete for the period 2008-2013, which corresponds to the main period of the CSI: I am able to successfully match non-zero hardware costs to over 79% of systems in the broader, processed LBNL data for California during this period. Given the importance of observing

with large condominium buildings. Dropping observations with missing pricing information mostly excludes observations prior to 2004.

[7]Third party ownership (TPO) in the residential PV market is a model that first emerged in California in 2007 and peaked several years later in 2012. The model is analogous to leasing a consumer product in other markets, where households sign a lease or power purchase agreement (PPA) with a third party solar company who owns the system and collects any available subsidies. While many installation firms in my final estimation dataset offer TPO options, I am unable to include this market in my analysis due to a lack of reliable data on the terms of TPO leases or PPAs. Any impacts on the TPO market or interactions between the household ownership market and the TPO market are therefore missing from my empirical analysis.

[8]I thank Galen Barbose from LBNL for making these data available to me and the CPUC for allowing me to use these data in my empirical analysis.

hardware costs in my empirical analysis, I subset the processed system-level data to include only those installations with valid hardware cost data that occurred between 2008 to 2013.

Before doing so, I use system-level data for the full period of 2000-2021 to determine the cumulative production of each installer within each county, where I measure cumulative production in terms of total installed system capacity, as well as the first and last quarter that an installer operates in a given county in the data.[9] This allows me to not only determine installation firms' cumulative experience operating in specific markets, but also identify firms that enter or exit specific markets during the main period of my analysis, 2008-2013. For each installer that I observe operating during 2008-2013, I then aggregate several fields from the system-level data to the county-quarter level, including average installed price per watt of system capacity, total installed capacity, average rebate per watt available to installers customers, and average hardware cost per watt.[10] I also calculate the average PV module efficiency, the number of distinct PV module types, and the name of the modal PV module manufacturer installed by a given installation firm at the county-quarter level.

The final installer-level dataset includes roughly 17,000 observations representing over 50,000 individual installations by 1,800 unique installers, which together operate across 44 California counties over 24 quarters. I present descriptive statistics for the main installer-level variables in Table 1. In addition to the main data sources that I discuss above, I bring in data from the US Census Bureau's American Community Survey (ACS) and Google's Project Sunroof to calculate the number of owner-occupied housing units that are suitable for solar adoption for each county-quarter in my data, which I use to estimate market shares each period. Additional data used in my analysis include demographic data from the ACS public use microdata samples (PUMS) for 2008-2013, which I use to estimate preference heterogeneity in my demand model; wage data from the Quarterly Census of Employment and Wages, which provide county-quarter-level wages for installation-adjacent fields for use as demand instruments; and retail electricity rates from the US Energy Information Administration's Form-861, which in combination with potential solar output data from the World Bank Group's Solar Atlas allow me to estimate households' NEM benefits. See Appendix B for further discussion of the data sources that I use to construct the main estimation data.

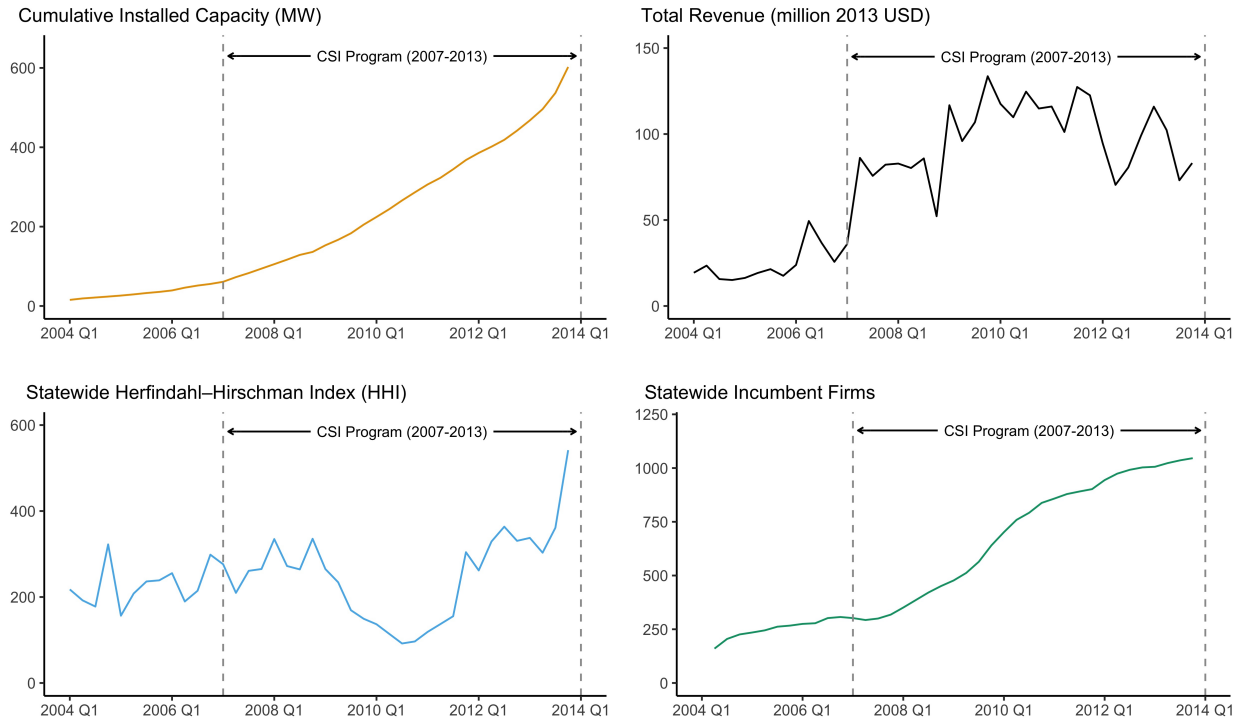
3.2 Descriptive Evidence

I begin my analysis of the supply-side implications of consumer subsidies by using my data to explore aggregate trends in the California residential PV industry around the period of the

[9] I use zip code to county crosswalk files from the US Department of Housing and Urban Development to map zip codes to counties, which serve as the spatial definition of a market in my setting.

[10] I normalize installed prices and rebates by system capacity as this allows for consistency across installations of various sizes.

Figure 3. California Residential PV Installation Activity, 2004-2013



Notes: This figure shows time series variation in several measures of activity for the California residential photovoltaic (PV) installation industry over the period 2004-2013, including cumulative installed capacity in megawatts (MW) (top left), total revenue in million 2013 US dollars (USD) (top right), the statewide Herfindahl-Hirschman Index (HHI) based on installed capacity (bottom left), and the number of statewide incumbent firms (bottom right). Data are for rooftop, household-owned installations and are taken from the Lawrence Berkeley National Lab's "Tracking the Sun" report public data file, which I discuss in detail in Section 3.1 (Barbose et al., 2022).

CSI. Figure 3 shows time series variation in several key measures of activity in the residential PV installation industry for the period 2004-2013, which corresponds to the main period of the CSI as well as the three years prior to the CSI for which I reliably observe industry data. As is the case throughout my empirical analysis, these aggregate trends focus on the household-owned, rooftop PV industry.

Similar to the rest of the global solar industry, the California PV industry experienced dramatic growth starting in the mid-2000s. As is clear from Figure 3, much of the observed growth in the California residential PV industry occurred during the main period of the CSI: from the start of 2007 to the end of 2013, total installed residential PV capacity increased nearly 10-fold, from 60 to 600 MW. Unsurprisingly, this period of rapid growth in residential PV capacity saw corresponding growth in the PV installation industry. Figure 3 shows time series variation in statewide installation revenue and the total number of operating

installation firms in the state, both of which increased rapidly during the main period of the CSI. Overall, installation industry concentration is low during this period, with a statewide Herfindahl-Hirschman Index (HHI) that ranges between 100 and 540 between 2007 and 2013.

While the trends in Figure 3 are suggestive, it is difficult to determine the role that California's PV policy environment played in driving the observed growth in the PV installation industry from aggregate time series data. Fortunately, it is possible to leverage plausibly-exogenous variation in PV subsidy policy across space and time to examine the possible impact of PV subsidies on growth in the installation industry. In particular, I use the variation across counties and quarters in CSI rebate levels shown in Figure 2 to estimate the relationship between the number of installation firms that enter a county and the generosity of consumer subsidies in that county in the quarter of entry. Since rebate levels depend on the history of past installations and unobserved factors that affect adoption—and in turn firm entry—may be correlated over time, I control for county-specific time-varying determinants of firm entry using a county-by-year fixed effect. Figure 4 shows the estimated relationship between the average number of entrants in a county and the CSI rebate level over the main period of the CSI, 2007-2013: there exists a strong, positive relationship between the CSI rebate level and the number of installation firms entering a market.[11]

The relationship in Figure 4 indicates that the CSI—and PV adoption incentives more broadly—likely contributed to the observed growth in the California PV installation industry; however, it is impossible to quantify the impact of these policies on market size and installer market structure from this result alone. Moreover, this result sheds little light on the underlying mechanisms through which the CSI and other PV adoption incentives served to expand the PV installation industry. In order to more precisely analyze the impact of consumer subsidies on the PV installation industry, including the mechanisms through which these impacts operate, I develop a model of PV installer entry and exit.

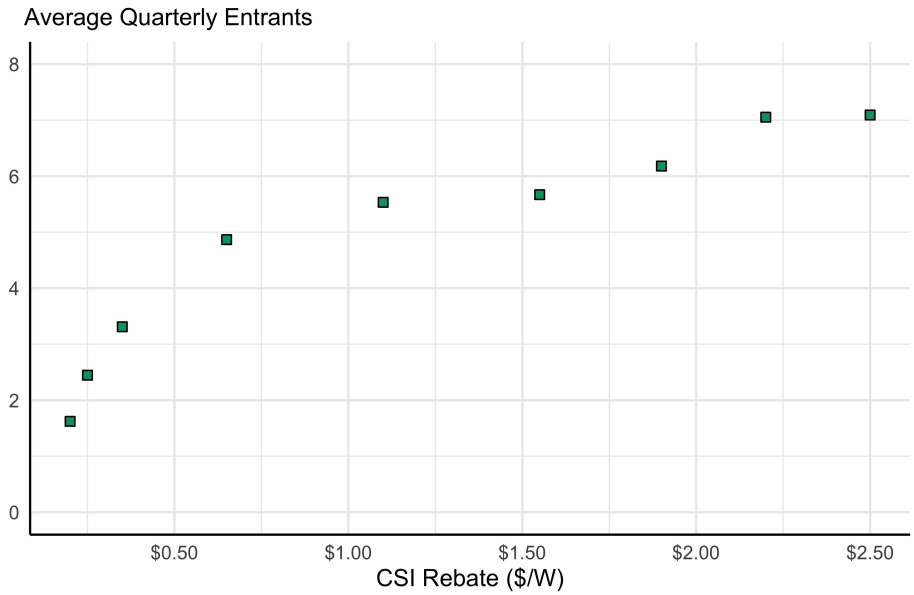
4 Model: Entry and Exit with Endogenous Learning-by-Doing

In this section, I introduce a model of firm entry, exit, and quantity-setting based on the theoretical framework for dynamic oligopoly of Ericson and Pakes (1995).

In each period t and market $m \in \{1, \ldots, M\}$, there are $j \in \{1, \ldots, J_{mt}\}$ incumbent firms

[11] The county-by-year fixed effect accounts for annual, market-specific variation in entry rates and aggregate demand that would influence the level of the CSI rebate. The relationship shown in Figure 4 therefore accounts for mean and county-specific, time-varying unobservables affecting entry rates and leverages plausibly-exogenous variation in county-level CSI rebates driven by IOU service territory boundaries, cumulative installed capacity in other counties within the same IOU service territory, and the cumulative installed capacity cutoffs in the CSI rebate schedule. This approach is similar to others in the literature (Bollinger and Gillingham, 2012; Hughes and Podolefsky, 2015; Ito, 2014).

Figure 4. Average County-level Entry and CSI Rebate Levels, 2007-2013



Notes: This figure shows the estimated relationship between the average number of entrants in a county and the California Solar Initiative (CSI) rebate level over the main period of the CSI, 2007-2013. The figure shows the mean number of quarterly entrants at each CSI rebate level observed in the data, conditional on a county-by-year fixed effect.

that face a set of static consumers $i \in \{1, \ldots, N_{mt}\}$ who demand differentiated solar PV installation services (Section 4.1). Incumbents dynamically choose a quantity of installations to provide conditional on their marginal costs (Section 4.2) and their beliefs of future learning (Section 4.3). Incumbent firms then choose whether to exit by comparing their expected discounted future profits with an idiosyncratic scrap value while a market-specific pool of potential entrants, $j \in \{1, \ldots, \bar{N}_m\}$, make one-shot entry decisions based on their expected discounted future profits and an idiosyncratic cost of entry (Section 4.4). At the end of each period, entry and exit decisions are implemented and the state evolves to the next period (Section 4.5). Firms' strategies lead to a Markov Perfect Equilibrium, which I assume is well-approximated by a Moment-based Markov Equilibrium concept (Section 4.6).

In practice, I define a discrete time period as a quarter and a market as a county. Firms have an infinite horizon and share a discount factor, β.

Incumbent installer j operating in market m at time t is differentiated by its state, which includes a common knowledge component s_{jmt} and a private component. The latter includes a shock to a firm's private selloff value, ϕ_{jmt} as well as a private, unobserved productivity

shock, κ_{jmt}. The former is defined by the vector

$$s'_{jmt} = \begin{bmatrix} E_{jmt} & \xi_{jmt} & h_{jmt} \end{bmatrix}$$

where to reflect learning-by-doing E_{jmt} is the installer's market-specific experience; ξ_{jmt} is a market- and time-specific measure of installation service quality, which I derive in practice from the demand system; and h_{jmt} is a market- and time-specific measure of hardware input costs. The state of market m at time t, s_{mt}, is therefore the union of all incumbent firms' common knowledge states and two additional aggregate state variables which vary across markets over time: an aggregate demand state, d_{mt}, and the market-level inclusive value, I_{mt}, which account for revenue potential and the intensity of competition among firms in a given market, respectively (Aguirregabiria et al., 2021). Potential entrants observe the market state and are differentiated by an idiosyncratic, private shock to entry costs, ω_{jmt}.

4.1 Demand for Solar Installations

I follow the random coefficient nested logit (RCNL) model of Brenkers and Verboven (2006) and Grigolon and Verboven (2014) to estimate consumer demand. The exposition of this model in my setting is inspired by Miller and Weinberg (2017). Incumbent firms in each period and market face a set of idiosyncratic, static consumers $i \in \{1, \ldots, N_{mt}\}$ who demand solar PV installation services. While there is evidence to suggest that consumers are forward-looking when purchasing PV systems, I assume that a static demand model provides a reasonable approximation to consumer behavior.[12] Each consumer purchases a solar PV installation from one of the observed incumbents ($j \in \{1, \ldots, J_{mt}\}$) or chooses to not install solar PV in this period ($j = 0$). The conditional indirect utility that consumer i receives from choosing installer j in market m in period t is

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

where X_{jmt} is a $K \times 1$ vector of observable product (firm) characteristics; 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; $\bar{\xi}_t$ allows the mean valuation of the indirect utility from installation to vary freely over time; and $\bar{\varepsilon}_{ijmt}$ is an idiosyncratic shock to preferences. I normalize prices and rebates by system capacity to ensure consistency when aggregating these variables across systems of different sizes.

[12] I develop and estimate a dynamic model of demand in Appendix C and find evidence to suggest that my static demand estimates are a reasonable reduced form for demand.

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. I normalize the indirect utility of non-installation such that $u_{i0,mt} = \varepsilon_{i0mt}$

I parameterize taste heterogeneity as follows: $\alpha_i^p = \alpha_p / y_i$ and $\alpha_i^k = \alpha_k + \sigma_k \log(y_i)$ for product attribute k, where α_p, α_k, σ_k are parameters to be estimated and y_i is observed consumer income.[13] It is therefore possible to re-write conditional indirect utility as

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

where

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

I express the market share of installer j in market m in period t as

$$ms_{jmt} = \frac{1}{N_{mt}} \sum_{i=1}^{N_{mt}} \frac{\exp\left((\upsilon_{jmt} + \mu_{jmt})/(1 - \eta)\right)}{\exp\left(I_{igmt}/(1 - \eta)\right)} \frac{\exp I_{igmt}}{\exp I_{imt}} \tag{2}$$

where I_{igmt} and I_{imt} are the McFadden (1977) inclusive values.[14]

4.2 Incumbent Cost Structure and Payoffs

After observing the market state, s_{mt}, incumbents privately observe a firm-, market-, and time-specific random productivity shock, κ_{jmt} and choose a quantity of installations to provide, q_{jmt}, at cost $mc_j(s_{mt})$ per watt of PV system capacity. Firms marginal production cost

[13] This parameterization of α_i^p approximates a Cobb–Douglas-style indirect utility function and is taken from Berry et al. (1999).

[14] By normalizing the mean indirect utility of the outside good to zero, $I_{i0mt} = 0$. The inclusive value of the inside goods is $I_{i1mt} = (1 - \eta) \log \sum_{j=1}^{J_{mt}} \exp\left((\delta_{jmt} + \mu_{ijmt})/(1 - \eta)\right)$ 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: $I_{mt} = \sum_{i=1}^{N_{mt}} I_{imt}$.

takes the following form:

$$mc_j(s_{mt};\theta^c) = h_{jmt} + w_j(s_{mt};\theta^c) + \kappa_{jmt} \tag{3}$$

where h_{jmt} is a firm's market-time specific marginal hardware cost, which is assumed to be exogenous; $w_j(s_{mt};\theta^c)$ is a firm's installation-specific marginal cost, which accounts for learning-by-doing by allowing installation costs to be a parametric function of accumulated experience—which enters s_{mt}—with parameters θ^c; and κ_{jmt} is a firm's unobserved productivity shock.[15] Installation-specific marginal cost, $w_j(s_{mt};\theta^c)$, captures costs associated with the sale, design, permitting, and construction for specific PV systems. Each term in (3) is measured in terms of per watt of system capacity.

I parameterize the relationship between accumulated production experience and installation-specific marginal costs following an unbounded learning model that is relatively standard in the learning-by-doing literature (Benkard, 2000; Covert and Sweeney, 2022; Levitt et al., 2013; Thornton and Thompson, 2001). Specifically, the relationship between installation costs and experience follows a power rule:

$$w_j(s_{mt};\theta^c) = c_0 \times \left(\tilde{E}_j\big(s_{mt};\theta^E\big)\right)^{\gamma} \tag{4}$$

where $\tilde{E}_j(s_{mt};\theta^E)$ is firm j's "effective" experience in market m in period t, which follows some parameterized function of the firm's own experience and the experience of its rivals; γ is a learning exponent that defines learning economies; c_0 is a parameter that scales effective experience into dollar amounts; and $\theta^c = (c_0, \theta^E, \gamma)$ collects the production cost parameters. Defining marginal installation cost as a function of effective experience allows me to test different models of experience accumulation. For example, one model of experience accumulation that I test is as follows:

$$\tilde{E}_j\big(s_{mt};\theta^E\big) = E_{jmt} + \theta_1^E \left(\sum_m \sum_{k\neq j} E_{kmt}\right) \tag{5}$$

where I allow industry cumulative production—the summation in the second term—to have a different and potentially incomplete marginal contribution to firm j's effective experience than firm j's own cumulative production. Note that the experience parameter θ_1^E is normalized with respect to the marginal contribution of a firm's own experience.[16] I report

[15] Where applicable throughout my exposition of the model, I index functions by j to illustrate that the corresponding function values are firm-specific, e.g., $mc_j(s_{mt})$. Since the state vector s_{mt} includes each firm's state, I use this notation to avoid duplicating a firm's state, s_{jmt}, in a function's arguments.

[16] In practice, I further normalize all experience terms by the total industry experience in the first period of

estimated experience parameters, θ^E, from several different models of experience accumulation in Section 6.

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

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

where $p_j(s_{mt}, q_{jmt})$ is firm j's market-time-specific price per watt, which is defined by the inverse demand curve corresponding to the demand model outlined in Section 4.1.

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

$$\begin{aligned} V_j(s_{mt}) &= \mathbb{E}_\phi\Big[\pi_j(s_{mt}) + \max\Big\{\phi_{jmt}, \beta\mathbb{E}\big[V_j(s_{mt+1})|s_{mt}, q^*_{mt}\big]\Big\}\Big] \\ &= \mathbb{E}_\phi\Big[\pi_j(s_{mt}) + \max\Big\{\phi_{jmt}, CV_j(s_{mt})\Big\}\Big] \end{aligned} \tag{6}$$

where $CV_j(s_{mt}) = \mathbb{E}\big[V_j(s_{mt+1})|s_{mt}, q^*_{mt}\big]$ denotes the continuation value, which equals the expected discounted future stream of profits with the expectation taken over the transition of the state variables; q^*_{mt} is the vector of optimal quantities chosen by incumbents in the market; and β is a common discount factor.

4.3 Product Market Game

Incumbent firms compete in the product market each period by choosing a quantity of differentiated installation services to provide. In particular, incumbents choose quantities in order to maximize the sum of current period profits and their expected continuation value, where expectations are taken over firms' beliefs of the evolution of the market state. Incumbent firm j in market m in period t has the following objective:

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

where $F(s_{mt+1}|s_{mt}, q_{mt})$ is the transition kernel for the state s_{mt} conditional on q_{mt}, the vector of quantity choices by incumbent firms in the market.

The optimal quantity choice for firm j in market m in period t will satisfy the following

the estimation sample, the first quarter (Q1) of 2008. This ensures the readability of the effective experience parameters, θ^E, and improves numerical stability in estimation.

first-order condition:

$$0 = \underbrace{\frac{\partial}{\partial q_{jmt}} \pi_j(s_{mt}, q_{jmt})}_{\text{marginal static profits}} + \underbrace{\frac{\partial}{\partial q_{jmt}} \beta \int V_j(s_{mt+1}) dF(s_{mt+1}|s_{mt}, q_{mt})}_{\text{dynamic ``markdown''}} \tag{7}$$

where the first term is the standard, static quantity-setting first-order condition and the second term captures both the incentive to raise production today to reduce future production costs as well as any strategic considerations regarding the impact of production today on a firm's rivals.[17] The first term in firm j's first-order condition can be written as:

$$\frac{\partial}{\partial q_{jmt}} \pi_j(s_{mt}, q_{jmt}) = p_j(s_{mt}, q_{jmt}) + \frac{\partial p_j(s_{mt}, q_{jmt})}{\partial q_{jmt}} q_{jmt} - mc_j(s_{mt}) \tag{8}$$

which is the standard, static quantity-setting first-order condition. This term accounts for the impacts of a firm's quantity choice on its benefits and costs in the current period: it trades off the marginal benefit and marginal cost of changes in quantities, accounting for the direct benefit and cost of producing a marginal unit in addition to the inframarginal impact of reducing the equilibrium price on all units it supplies.

The dynamic markdown term describes the marginal effect of higher prices on a firm's value function, or discounted future profits. As noted by Berry and Pakes (2000) and Covert and Sweeney (2022), a change in a firm's choice of q_{jmt} will only affect the state transition distribution, $F(s_{mt+1}|s_{mt}, q_{mt})$, through the evolution of experience, E_{jmt}, not the value function $V(\cdot)$ itself. Noting further that for any possible value of q_{mt}, $dF(s_{mt+1}|s_{mt}, q_{mt}) > 0$, I can simplify the dynamic markdown term as follows:

$$\begin{aligned}
&\frac{\partial}{\partial q_{jmt}} \beta \int V_j(s_{mt+1}) dF(s_{mt+1}|s_{mt}, q_{mt}) \\
&\quad = \beta \int V_j(s_{mt+1}) \frac{\partial}{\partial q_{jmt}} dF(s_{mt+1}|s_{mt}, q_{mt}) \\
&\quad = \beta \int V_j(s_{mt+1}) \left(\frac{\frac{\partial}{\partial q_{jmt}} dF(s_{mt+1}|s_{mt}, q_{mt})}{dF(s_{mt+1}|s_{mt}, q_{mt})} \right) dF(s_{mt+1}|s_{mt}, q_{mt}) \\
&\quad = \beta \mathbb{E} \left[V_j(s_{mt+1}) \left(\frac{\frac{\partial}{\partial q_{jmt}} dF(s_{mt+1}|s_{mt}, q_{mt})}{dF(s_{mt+1}|s_{mt}, q_{mt})} \right) \middle| s_{mt}, q_{mt} \right]
\end{aligned} \tag{9}$$

The result is an intuitive, simplified form for the second term in firms' quantity-setting first-order condition. As shown above, firm j's dynamic markdown in market m in period t is

[17] I follow the terminology of Covert and Sweeney (2022) and refer to the second term as the "dynamic markdown."

its expectation of the benefits at a given realization of the future state, s_{mt+1}, multiplied by the marginal change in the probability that this state is realized resulting from a change in that firm's quantity choice, q_{jmt}. Combining these forms for the marginal static profits and dynamic markdown provides me with a feasible approach to writing firms' quantity-setting first-order condition as a function of data and estimable parameters. I outline my approach to doing so, including any necessary additional assumptions, in Section 5.

4.4 Exit and Entry

At the start of each period, incumbent firms first observe the market state, s_{mt}; draw a private productivity shock, κ_{jmt}; and then compete in the product market as outlined above. After product market decisions are implemented, incumbents draw a private scrap value, ϕ_{jmt}. The optimal exit policy follows a threshold form: a firm j will exit market m in period t if the scrap value that it draws is greater than its continuation value, $CV_j(s_{mt})$. The scrap value is distributed i.i.d., so the firm exits with probability $p_j^x(s_{mt})$:

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

where F_ϕ is the conditional distribution function (CDF) of the distribution of ϕ_{jmt}. In practice, I assume that ϕ_{jmt} is distributed i.i.d. exponential with scale parameter $1/\sigma_\phi$, i.e. $\phi_{jmt} \overset{\text{i.i.d.}}{\sim} \text{Exponential}(1/\sigma_\phi)$. This implies an exit probability of the following form:

$$p_j^x(s_{mt}; \sigma_\phi) = \exp\left(-\frac{CV_j(s_{mt})}{\sigma_\phi}\right) \tag{10}$$

At the same time that incumbents draw private scrap values, $\bar{N}_m$ potential entrants observe the common knowledge state s_{mt} and a private i.i.d. entry cost, ω_{jmt}, before making a one-shot entry decision. If potential entrant $j \in \{1, \ldots, \bar{N}_m\}$ chooses not to enter, it disappears with a payoff of zero; if it enters, it pays its entry cost and becomes an incumbent in the next period. I assume that entrants are endowed with random values of the non-deterministic state variables, quality and hardware cost, where values are drawn from the empirical distribution of observed states. Entrants start with zero experience, $E_{jmt} = 0$.

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

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

where χ^e_{jmt} equals 1 if potential entrant j chooses to enter market m in period t and 0

otherwise and where the expectation is over the potential entrants information set at the time of the one-shot entry decision, which includes the common knowledge state, s_{mt}. Given that ω_{jmt} is i.i.d., the probability that potential entrant j in market m in period t enters is:

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

where F_ω is the CDF of the distribution of ω_{jmt}. I assume that ω_{jmt} is distributed i.i.d. exponential with scale parameter $1/\sigma_\omega$, i.e. $\omega_{jmt} \overset{\text{i.i.d.}}{\sim} \text{Exponential}(1/\sigma_\omega)$. This implies an entry probability of the following form:

$$p_j^e(s_{mt}; \sigma_\omega) = 1 - \exp\left(-\frac{VE_j(s_{mt})}{\sigma_\omega}\right) \tag{11}$$

4.5 State Transitions

I assume that hardware cost (h_{jmt}) and installation quality (ξ_{jmt}) are exogenous and evolve stochastically over time according to a first-order Markov process. In the case of hardware costs, this assumption implies that PV installers are price-takers in the upstream market for PV system inputs, a reasonable assumption based on the number of PV module and inverter manufacturers relative to the number of installers. The assumption of a stochastic transition process for quality is perhaps tenuous; however, such a process could result from random outcomes to investment in quality. While it is possible—both theoretically and indeed computationally—to endogenize the quality process through investment decisions, this dynamic is not my primary focus and there are non-trivial data limitations to doing so, namely the lack of data on investment and installation quality.

The remaining state variables, aggregate demand (d_{mt}), market-level inclusive value (I_{mt}), and experience (E_{jmt}), are endogenous and evolve over one period as a result of the demand model outlined in Section 4.1 and firm quantity-setting actions outlined in Section 4.3.

4.6 Equilibrium

A Markov-Perfect Equilibrium of the model in a given market m in period t includes a set of policies governing production quantities, exit, and entry $\Big(q_{mt}, p_j^x(s_{mt}), p_j^e(s_{mt})\Big)$ as well as value functions $V_j(s_{mt})$, and prices p_{jmt} such that the firms' production, exit, and entry decisions satisfy (7), (10), and (11). Equilibrium prices are generated by the inverse demand function and set current demand equal to current supply. Moreover, incumbent value functions satisfy (6) in equilibrium and all firms employ the 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 a high-dimensional object due to the large number of firms in the industry. The above definition of Markov-Perfect Equilibrium assumes that firms keep track of the state variables of every rival, arguably a strong assumption in a setting in which there can be over 100 active incumbents in a given market. To reduce the computational burden of estimating model primitives and solving for the model equilibrium, I assume that firms do not keep track of the state variable of every rival, but instead track market-level moments of rivals' state variables each period. Ifrach and Weintraub (2017) provide a detailed treatment of this approach, which the literature refers to as a Moment-based Markov Equilibrium. This approach is similar in spirit to the oblivious equilibrium concept (Benkard et al., 2015; Weintraub et al., 2008), which approximates the Markov-Perfect Equilibrium in industries with many firms and has been widely employed in the empirical literature (Barwick et al., 2021; Gerarden, 2022; Jeon, 2022; Vreugdenhil, 2023).

A remaining issue in the equilibrium concept that I apply in this context is nonstationarity in the regulatory environment. While there are a number of overlapping adoption subsidies during the study window, several of which evolve from period to period, explicitly modeling firms' beliefs regarding the distribution of future subsidy environments substantially complicates equilibrium computation. Thus, I follow standard practice in the literature and assume that firms behave as if changes to existing subsidy policy are unanticipated, one-time changes that are not repeated in the future (Aguirregabiria et al., 2021; Ryan, 2012). As noted by Barwick et al. (2021), one approach to proxy for dynamic regulatory environments is to use lower discount rates so that future payoffs are less relevant for decisions in the current period. I use this approach to test the robustness of my assumption that subsidy changes are perceived as permanent in my setting.

5 Estimation Strategy

In this section, I outline the empirical approach that I take to recover key model parameters. The main primitives of the model that I aim to recover include: the demand system for rooftop solar PV installations, the installation-specific marginal cost function, the distribution of scrap values, and the distribution of entry costs.

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

5.1.1 Demand Estimation

Demand estimation follows the nested fixed point procedure of Berry et al. (1995) and best practices for differentiated demand estimation outlined by Conlon and Gortmaker (2020) as

adapted to the RCNL model of Brenkers and Verboven (2006) and Grigolon and Verboven (2014). The approach derives a generalized method of moments (GMM) estimator from the moment condition $\mathbb{E}[Z_D'\xi(\theta_0^D)] = 0$, where $\theta_0^D = (\alpha_p, \alpha', \sigma, \eta)$ is the vector of population demand parameters, $\xi(\theta^D)$ is the vector ξ_{jmt} that solves the system of market shares defined by (2) for a given set of parameters θ^D, and Z_D is a conformable matrix of valid instruments. This definition of $\xi(\theta^D)$ implies that this vector is the structural error term if evaluated at the population parameters (i.e., $\xi_{jmt} = \xi(\theta_0^D)$). The resulting GMM estimate is

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

where $\widehat{\xi}(\theta^D)$ 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 quarterly average electricity price in a county. Including a firm fixed effect, $\bar{\xi}_j$, in (1) absorbs time-invariant installer characteristics and the quarter 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 data per county-quarter.[18]

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 dampened 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 the demand parameters θ^D requires instruments for the parameter on price and each of the nonlinear parameters. The standard price endogeneity issue applies in this setting: prices are likely correlated with the structural error term as firms make product market decisions with knowledge of unobservable product- and market-specific consumer valuations. As outlined in Berry and Haile (2014), consumer heterogeneity in preferences for product characteristics leads to a simultaneity problem. This is due to the fact that the mean utilities that equate observed market shares to model-predicted shares depend

[18] I calculate per-capita household income as the total income divided by the number of household members. Since the ACS PUMS data are only available at an annual level, I draw from the same annual sample for each county-quarter within the same calendar year. County of residence is masked for select household observations from a subset of less populous California counties. For these counties, I draw from the full sample for the state in the relevant time period.

on the parameters governing consumer heterogeneity. In this setting, parameters that lead to consumer preference heterogeneity include the income terms (α_p and σ) and the nest parameter (η).

The first set of demand instruments that I use addresses the price endogeneity issue and includes the CSI rebate per watt and the county-level electrician/roofing wage rates, each of which represents a plausible marginal cost shifter. The CSI rebate, which varies across market and time and is used elsewhere in the literature estimating solar installation demand (Gillingham and Tsvetanov, 2019; Pless and Van Benthem, 2019), is analogous to a shift in firms' supply curve holding demand fixed assuming the standard statutory-incidence irrelevance result given that I model consumers as facing post-incentive prices. Electricians' and roofers' wages also serve as exogenous cost shifters, plausibly affecting demand only through their effect on the price of installing solar. I augment these price instruments with an additional set of instruments that includes the non-price characteristics of other goods in a given market, i.e., X_{-jmt}. This assumes that non-price characteristics of all products are mean-independent of quality shocks, ξ_{jmt}, which is a relatively standard assumption in the literature starting with Berry et al. (1995).

The second set of demand instruments identifies the nested logit parameter and includes the number of active firms within a given market and the lagged quantity of installations an installer completed in other counties. Identification of the nest parameter requires exogenous variation in the conditional shares of the inside goods.[19] The number of firms (or products in other settings with multi-product firms) is a relatively standard instrument in nested logit models and should be negatively correlated with the within-group share. The number of installations an installer finishes in other counties directly frees up labor that can be moved across counties, which also influences the within-group share. This instrument is used elsewhere in the literature estimating solar demand (Bollinger and Gillingham, 2019).

The final set of demand instruments identifies the parameters governing consumer heterogeneity in preferences for characteristics, σ. Assuming that the structural error term, ξ_{jmt}, is mean independent of income and product characteristics, identification of these parameters comes from correlation between local demographics and product shares. In particular, I follow Miller and Weinberg (2017) and use mean income interacted with the observed product characteristics in X_{jmt}, where mean income is calculated for each county-quarter using the

[19] To see this, consider a version of the RCNL model presented in Section 4.1 that eliminates individual level heterogeneity. This reduces the RCNL model to the nested logit model and yields an equation for market shares that is linear in its parameters:

$$\log(ms_{jmt}) - \log(ms_{0mt}) = \alpha_p(p_{jmt} - r_{jmt}) + \alpha' X_{jmt} + \bar{\xi}_j + \bar{\xi}_t + \eta \log(\bar{ms}_{jmt|g}) + \xi_{jmt}$$

where $\bar{ms}_{jmt|g} = ms_{jmt} / \sum_{j=1}^{J_{mt}} ms_{jmt}$ is the conditional within-group share.

200 draws from the American Community Survey PUMS.

5.1.2 Exit Policy Function

The exit policy function defines the equilibrium exit behavior of firms conditional on the industry state. I estimate firms' exit policy function using a logit regression:

$$\Pr(\chi^x_{jmt} = 1|s_{mt}) = \frac{\exp\left(h_j(s_{mt})\right)}{1 + \exp\left(h_j(s_{mt})\right)}$$

where χ^x_{jmt} 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.

Given that firms' exit probabilities approximate endogenous equilibrium objects, it is important to obtain consistent estimates of these functions in the first stage as this ensures the consistency of my second stage estimates of the dynamic parameters. I therefore follow the data-driven approach of Gerarden (2022) to determine the functional form of $h_j(s_{mt})$ that trades off flexibility with the issues associated with overfitting. In particular, I estimate a logit model of the discrete exit decision with a full set of candidate regressors, $h_j(s_{mt})$, via penalized maximum likelihood.[20] This first step identifies a set of non-zero regressors, $\tilde{h}_j(s_{jmt})$, which I then use in a second step to estimate a logit model of exit via maximum likelihood. Additional information is available in Appendix D. I denote the resulting fitted exit probabilities as $\hat{p}^x_{jmt}$.

5.1.3 State Space

As discussed in Section 4.6, the state variable s_{mt} is a high-dimensional object due to the large number of firms in the industry. In line with the Moment-based Markov Equilibrium concept of Ifrach and Weintraub (2017), I assume that firms track moments of the state variables of their rivals rather than the state variables of every rival. In particular, I assume that firms track within- and out-of-county averages of other firms' hardware cost (i.e., $\bar{h}_{kmt}$ and $\bar{h}_{klt}$ $\forall k \neq j, l \neq m$) and other firms' quality (i.e., $\bar{\xi}_{kmt}$ and $\bar{\xi}_{klt}$ $\forall k \neq j, l \neq m$). In the case of experience, firms track the total experience of other firms in a given market (i.e., $\bar{E}^m_{jmt} = \sum_{k \neq j} E_{kmt}$) as well as the total experience of other firms in other markets (i.e., $\bar{E}^o_{jmt} = \sum_{l \neq m} \sum_{k \neq j} E_{klt}$). Firms track their own state variables, $(E_{jmt}, \xi_{jmt}, h_{jmt})$ in addition to these moments of rivals' state variables.

[20] Candidate regressors include quadratic polynomials of the full set of state variables and their complete set of pairwise interactions as well as county and quarter fixed effects. I estimate the penalized logit via LASSO and select the sole tuning parameter via k-fold cross validation.

The aggregate states described in Section 4 remain unchanged. An aggregate demand state, d_{mt}, accounts for market-time variation in revenue potential. The market-level inclusive value, I_{mt}, measures the intensity of competition among firms in a given market and is calculated as the total inclusive value of all goods in a market-period based on the demand model. The result of these assumptions on the state space is an 11-dimensional state vector for each firm, which represents a substantial reduction in the dimensionality of the model.

5.1.4 State Transitions

When firms make entry, exit, and quantity decisions, they consider both their current payoffs as well as the value of operating in future periods, which is a function of future state variables; rivals' entry and exit decisions; and rivals' production decisions. As noted, I assume that the exogenous state variables—hardware cost, h_{jmt}, and installation quality, ξ_{jmt}—follow a first-order Markov process rather than explicitly modelling firms' beliefs. Following Aguirregabiria and Mira (2007) and much of the dynamic games literature, I model the evolution of the two exogenous states as first-order autoregressive (AR(1)) processes. I allow for market-specific intercepts in estimating the AR(1) transition processes to account for substantial heterogeneity in conditions across counties. The transition process for experience is a deterministic function of a firm's current experience level and quantity choice, with new entrants having an experience level of 0.

The aggregate states, aggregate demand and inclusive value, are complicated objects that are determined by equilibrium in the product market game. Following other work in the literature (see Aguirregabiria et al. (2021)), I assume that firms' beliefs over the transition of the aggregate states follow AR(1) processes similar to that for the exogenous states. The assumption that firms' beliefs over the transition of aggregate demand follow an AR(1) process is relatively standard in the literature. Gowrisankaran and Rysman (2012) and Barwick and Pathak (2015) make similar assumptions regarding agents' beliefs over the transition of inclusive value state variables.

5.2 Second Stage: Learning, Exit, and Entry Parameters

I now turn to my approach to recovering the dynamic parameters of the model, which include the production cost parameters that govern learning (θ^c), exit parameter (σ_ϕ), and entry parameter (σ_ω). I begin by outlining my approach to approximating firms' value functions before turning to the main estimating equations for the learning, exit, and entry parameters. As noted in my discussion of the model, β is a quarterly discount factor common to all firms.

I assume a quarterly discount factor that translates to an annual discount factor of 0.875.[21]

5.2.1 Value Function Approximation

In light of the fact that the conditions for optimal quantity-setting, exit, and entry defined by (7), (10), and (11) all depend on $V_j(s_{mt})$, estimation of the target structural parameters requires solving for the unknown value function. Since scrap values, ϕ_{jmt} are i.i.d. exponential, it is possible to write the value function prior to the realization of ϕ_{jmt} as

$$\begin{aligned} V_j(s_{mt}) &= \mathbb{E}_\phi[\pi_j(s_{mt}) + \max\{\phi_{jmt}, CV_j(s_{mt})\}] \\ &= \pi_j(s_{mt}) + p_j^x(s_{mt})\mathbb{E}_\phi[\phi_{jmt}|\phi_{jmt} > CV_j(s_{mt})] + (1 - p_j^x(s_{mt}))CV_j(s_{mt}) \\ &= \pi_j(s_{mt}) + p_j^x(s_{mt})\sigma_\phi + CV_j(s_{mt}) \end{aligned} \tag{12}$$

where the final line follows from the memoryless property of the exponential distribution, i.e., $\mathbb{E}_\phi[\phi_{jmt}|\phi_{jmt} > CV_j(s_{mt})] = \sigma_\phi + CV_j(s_{mt})$ 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, it is possible to obtain a flexible approximation of the value function implicitly defined by the Bellman equation (12) following recent work in the dynamic games literature (e.g., Barwick et al. (2021)). In particular, given the smoothness of the value function in this context, it is possible to approximate the value function arbitrarily well using L basis functions $b_j^l(s_{mt})$:

$$V_j(s_{mt}) \simeq \sum_{l=1}^{L} \lambda_l b_j^l(s_{mt}) \qquad\qquad CV_j(s_{mt}) \simeq \beta \sum_{l=1}^{L} \lambda_l \mathbb{E}[b_j^l(s_{mt+1})|s_{mt}] \tag{13}$$

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

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 produce non-trivial approximation errors in this setting. Second, given that the value function implicitly defined by the Bellman equation (12) is nonlinear in parameters, popular forward simulation approaches are computationally-infeasible in this setting.[22] I provide additional details on my approach to value function approximation, including my approach to constructing basis

[21] This is the same discount factor that Gerarden (2022) assumes for solar panel manufacturers and is similar in magnitude to that estimated by De Groote and Verboven (2019) for PV-adopting households. This discount factor is similar to those applied in other contexts: for example Ryan (2012) assumes an annual discount factor of 0.9.

[22] This non-linearity is due to the fact that static profits are a function of marginal production costs, which are nonlinear in the learning parameters.

functions of the state variables, approximating the expectation over the state transitions, and estimating $\{\hat{\lambda}_l\}_{l=1}^{L}$ in Appendix E. I report empirical relationships between the state variables and a set of final value function estimates in Figure E1.

5.2.2 Production and Exit Cost Estimation

Since both production cost and exit parameters are functions of firms' value functions—which itself is a function of these target parameters—it is necessary to jointly estimate both sets of parameters. I do so by deriving suitable moments from the conditions governing optimal quantity setting and exit as functions of data and parameters, which I then use to solve for the target parameters via non-linear GMM.

I re-express the condition for optimal quantity setting given in (7) as a function of parameters and data by combining the static markup (8) and dynamic markdown (9) with my functional form assumptions for marginal productions costs (3) as follows:

$$\begin{aligned} 0 = p_{jmt} + (\Delta_{mt}^{-1})_{(j,j)} \times ms_{jmt} - (h_{jmt} + w_j(s_{mt}; \theta^c) + \kappa_{jmt}) \\ + \beta \mathbb{E}[V_j(s_{mt+1}; \lambda) \times \Omega_j(s_{mt}, q_{mt})] \end{aligned} \tag{14}$$

where

- p_{jmt}, ms_{jmt}, and h_{jmt} are the price per watt, market share, and hardware cost per watt for firm j in market m in period t, each of which I take directly from the data;
- $(\Delta_{mt}^{-1})_{(j,j)}$ is the j-th diagonal element of the inverted matrix of own- and cross-price derivatives of demand, Δ_{mt}, estimated from the demand model in the first stage;
- $w_j(s_{mt}; \theta^c)$ is firm j's installation-specific cost per watt in market m in period t, which is a function of firm j's effective experience and parameters θ^c according to (4);
- κ_{jmt} is firm j's unobserved productivity shock;
- $V_j(s_{mt+1}; \lambda)$ is firm j's value in market m in period $t+1$, which I can approximate as a function of expected state variables—conditional on current states and quantity choices—and parameters λ as described above; and
- $\Omega_j(s_{mt}, q_{mt}) = \frac{\frac{\partial}{\partial q_{jmt}} dF(s_{mt+1}|s_{mt}, q_{mt})}{dF(s_{mt+1}|s_{mt}, q_{mt})}$ is the second term in the dynamic markdown that captures the sensitivity of state transitions to firm j's quantity choice.

Note that I omit function arguments for objects taken directly from the data or estimated in the first stage in (14) for ease of notation.

Two objects in (14) warrant additional discussion: $\Omega_j(\cdot)$ and κ_{jmt}. It is possible to write the object $\Omega_j(\cdot)$, which describes the sensitivity of state transitions to a firm's quantity choice, as a relatively simple function of data and parameters. Note that the only future state variable directly affected by firm j's current quantity decision is firm j's own effective experience state, which is a direct function of firms' current and past quantity choices. I can therefore write a closed-form for $\Omega_j(\cdot)$ by solving for the gradient of the transition probability distribution of firms' own effective experience with respect to their quantity choice. Define the transition kernel for firm j's effective experience state variable 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 incumbent firms across all markets in period t. The closed form expression for $dG(\cdot)$ and therefore $\Omega_j(\cdot)$ will vary based on the specific parametric model of experience accumulation that I assume. For example, taking the model of experience accumulation defined by (5), I know that

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

Assuming this model of experience accumulation gives the following form for the sensitivity of state transitions to firm j's quantity choice:

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

which I can plug into (14). This closed form expression allows me to calculate the expectation in (14) for a given guess of the target parameters using data on current states and quantity choices; estimated state transitions; and value function approximating coefficients, λ.

I now turn to the firm-specific, unobserved productivity shock, κ_{jmt}. From the model, the implied values of unobserved productivity shocks across firms, markets, and periods, can be computed as a function of the production cost and value function parameters to be estimated, (θ^c, λ). I assume that the unobserved productivity shock for firm j in market m evolves according to an AR(1) process, where the error

$$\nu_{jmt}(\theta^c, \lambda, \rho) = \kappa_{jmt}(\theta^c, \lambda) - \rho \kappa_{jmt-1}(\theta^c, \lambda) \tag{15}$$

is mean zero and independent such that

$$\mathbb{E}\big[\nu_{jmt}(\theta^c, \lambda, \rho)\big] = 0 \tag{16}$$

where ρ is an additional serial correlation coefficient to be estimated. The above defines a suitable moment from which I can recover estimates of the production cost parameters.

Relying on the moment given by (16) rather than moments from the unobservable productivity shock, κ_{jmt}, is appealing as it explicitly accounts for potential serial correlation in the productivity shocks. Serially correlated productivity is a common concern in the estimation of production functions (Olley and Pakes, 1996). In the present context, ignoring serial correlation would likely result in biased estimates of production cost parameters: firms with serially correlated, positive productivity shocks are likely to accumulate greater levels of experience while also having relatively low marginal production costs, resulting in an over-estimate of overall learning without accounting for this correlation. Indeed, the existing literature estimating passive learning finds substantial serial correlation (Argote and Epple, 1990; Benkard, 2000). Explicitly estimating the serial correlation coefficient on κ_{jmt}, ρ helps to not only produce unbiased estimates of learning parameters, but also account for serially correlated productivity shocks in counterfactual simulations in Section 7.

I derive a suitable moment condition for the estimation of the exit cost parameter, σ_ϕ, from the condition defining optimal exit. It is possible to form estimates of model-implied exit probabilities given in (10) using the basis function approximations to firms' continuation values. I therefore aim to find the exit cost parameter that minimizes the sum of squared differences between fitted exit probabilities estimated in the first stage and model-implied exit probabilities:

$$\min_{\sigma_\phi} \sum_{j,m,t} \Bigg(\underbrace{\hat{p}^x_{jmt} - \exp\left(-\frac{CV_j(s_{mt};\lambda)}{\sigma_\phi}\right)}_{\equiv \psi_{jmt}(\sigma_\phi,\lambda)} \Bigg)^2$$

The moment condition that corresponds to the above nonlinear least squares problem is

$$\mathbb{E}\left[\frac{\partial \psi_{jmt}(\sigma_\phi,\lambda)}{\partial \sigma_\phi}\psi_{jmt}(\sigma_\phi,\lambda)\right] = 0 \tag{17}$$

which has a closed form expression and defines a suitable moment from which I can estimate the exit cost parameter.

Stacking the moment conditions (16) and (17) provides a means of jointly estimating the target parameters, $\boldsymbol{\theta} = (\theta^c, \rho, \sigma_\phi)$. Note, however, that these moment conditions depend on the value function approximating coefficients, λ, which are functions of a subset of the target parameters (θ^c, σ_ϕ). I therefore adopt an approach to jointly estimate the production and exit cost parameters based on that of Sweeting (2013): for a given value of the target parameters, I solve for the approximating coefficients that minimize violations of the Bellman

equation and use these to estimate continuation values, which I use in turn to solve for the next iteration of target parameters. I then iterate this procedure until a convergence criterion is satisfied. Specifically, at a given iteration i and guess of the target parameters, $\hat{\boldsymbol{\theta}}^i$, I do the following:

1. Solve for the value function approximating coefficients, $\hat{\lambda}^i$, for the current values of the target parameters, $\hat{\boldsymbol{\theta}}^i$, and use these to calculate $CV_j^i(s_{mt}; \hat{\lambda}^i)$.

2. Use the resulting values for $CV_j^i(s_{mt}; \hat{\lambda}^i)$ to update estimates of the target parameters via two-step GMM, with the GMM estimator $\hat{\boldsymbol{\theta}}^{i+1} = \arg\min_{\boldsymbol{\theta}} \Psi(\boldsymbol{\theta})' W^{-1} \Psi(\boldsymbol{\theta})$, where $\Psi(\boldsymbol{\theta})$ is a vector of stacked moments from (16) and (17) and W is a positive definite approximation to the optimal weight matrix.

3. Calculate the L^1 norm of the difference between the new and starting sets of target parameters and check whether it is below a tolerance level, ϵ. If $\left|\left|\hat{\boldsymbol{\theta}}^{i+1} - \hat{\boldsymbol{\theta}}^i\right|\right| \not< \epsilon$, return to step 1 with new parameter values $\hat{\boldsymbol{\theta}}^{i+1}$; otherwise, the procedure stops.

On termination, the above procedure delivers a set of target parameter estimates, $\hat{\boldsymbol{\theta}}$, and value function approximating coefficients, $\hat{\lambda}$, the latter of which I can use to construct final estimates of firms' value functions. I calculate standard errors for the target parameter estimates via a non-parametric bootstrap of the above procedure in which I construct 200 samples by drawing entire market histories with replacement.

5.2.3 Entry Cost Estimation

Armed with estimates of the value function approximating coefficients, it is possible to form estimates of entry values for potential entrants for use in estimating the entry cost parameter. To do so, I compute the expected values of the state variables in the next period for all potential entrants using the observed aggregate state variables and assuming that entrants are endowed with random values of the non-deterministic state variables, quality and hardware cost, where values are drawn from the empirical distribution of observed states. Moreover, entrants enter with zero experience, $E_{jmt} = 0$. I calculate the expectation of future state variables for entrants conditional on entry in a similar fashion to incumbents: I average state values over 1000 draws from the state transition processes estimated in the first stage.

A challenge associated with estimating entry costs is that the number of potential entrants in each market is unobserved. I follow standard practice in the literature and take a data-driven approach to estimating the number of potential entrants in each market: I assume that the number of potential entrants is some multiple of either the median, mean, or maximum

number of observed entrants for each market over the sample period. I report estimates of the entry cost parameter for several different assumed values of potential entrants.

With estimates of entry values for all potential entrants, it is possible to recover estimates of the entry cost parameter, σ_ω. I do so via maximum likelihood estimation (MLE) given that an MLE estimator offers greater efficiency over a moment-based estimator matching fitted entry probabilities to model-implied entry probabilities. The log likelihood for entry is given by:

$$\begin{aligned}&\log(f(\chi^e_{jmt};\sigma_\omega)) \\ &\quad= \sum_{j,m,t}\left[\chi^e_{jmt}\log\left(1-\exp\left(\frac{-VE_j(s_{mt};\hat{\lambda})}{\sigma_\omega}\right)\right)-(1-\chi^e_{jmt})\left(\frac{VE_j(s_{mt};\hat{\lambda})}{\sigma_\omega}\right)\right] \end{aligned} \tag{18}$$

where χ^e_{jmt} equals 1 if potential entrant j chooses to enter market m in period t and 0 otherwise. The MLE estimator maximizes the log likelihood given by (18):

$$\hat{\sigma}_\omega = \arg\min_{\sigma_\omega} \sum_{j,m,t} \log(f(\chi^e_{jmt};\sigma_\omega))$$

I calculate standard errors for the target entry parameter estimate using a non-parametric bootstrap with 200 bootstrap samples clustered by county.

5.2.4 Identification of the Dynamic Parameters

Identification of the dynamic parameters builds on existing results in the literature as well as key model assumptions. As I discuss above, the primary concern with identifying the production cost parameters governing learning economies is the potential for serial correlation in the unobservable productivity shocks. Identification of the base cost, learning exponent, effective experience, and serial correlation parameters therefore relies on the validity of the model of firms' productivity shocks given by (15). If this model is accurate, then variation in observed prices, quantities, and exit decisions across incumbent firms with different hardware costs and experience vectors identifies these production cost parameters.[23]

It is important to note that identification of the effective experience parameters may not come directly from the data and existing assumptions. Of particular concern is the potential for unobservable factors that are correlated with both a firm's marginal installation cost and the experience of its rivals. One example of such a factor is an unobservable demand

[23]As a robustness check, I also estimate the model using moments from the unobservable productivity shock, κ_{jmt}, interacted with an alternative set of instruments, Z_{jmt}, that plausibly satisfy the relevance and exogeneity conditions. I discuss the results, which I present in Table A3, in Section 6.3.

shock, perhaps at the county, region, or industry-level, that would increase incumbent firms' own experience as well as the experience of their rivals in later periods. This would lead to decreased marginal installation costs for firms as well as higher experience levels for rival firms, which in the current model framework would potentially bias upward the effective experience parameters governing inter-firm spillovers. Given the challenges associated with identifying inter-firm spillovers, I choose to assume away this identification challenge in the current set of estimation results; however, to explore the robustness of my qualitative findings from the counterfactual simulations in Section 7, I simulate select counterfactuals with lower levels of knowledge spillovers.

Finally, identification of the exit and entry parameters as well as the value function approximating coefficients follows from Hotz and Miller (1993), who show that differences in choice-specific value functions are identified from observed choice probabilities.

6 Model Results

I present the main model estimates in this section. I begin by presenting the first stage estimates (Section 6.1) before turning to my estimates of the dynamic model primitives (Section 6.2). I then discuss several robustness checks (Section 6.3).

6.1 First Stage Estimates

6.1.1 Consumer Demand

I present parameter estimates and standard errors clustered by county for the consumer demand system in Table 2. Columns (1) and (2) estimate nested logit (NL) versions of the demand system that remove individual-level heterogeneity in taste parameters. Columns (3) through (5) correspond to different versions of the full RCNL model outlined in Section 4.1. To ensure comparability in price coefficient estimates across the NL and RCNL models, I divide price by county-quarter mean income in the NL models. All specifications include both year and firm fixed effects.

The price coefficients (α_p) are precisely estimated and have the expected sign across all five specifications. Median own-price elasticities range from -1.09 to -1.43 and are smaller in magnitude when including observable firm attributes. Market price elasticities are substantially lower than own-price elasticities, which suggests that substitution primarily occurs across installers rather than on the extensive margin (i.e., between installing and not installing solar PV). This is evident from the nesting parameter (η), which is relatively large and precisely estimated across all specifications. An easier way to interpret the large nest-

Table 2. Estimated Demand System Parameters

		NL-1	NL-2	RCNL-1	RCNL-2	RCNL-3
	Parameter	(1)	(2)	(3)	(4)	(5)
Price/Income	α_p	−0.959	−0.944	−0.599	−0.551	−0.862
		(0.303)	(0.306)	(0.22)	(0.215)	(0.352)
Nesting Parameter	η	0.901	0.901	0.901	0.901	0.902
		(0.030)	(0.028)	(0.031)	(0.028)	(0.045)
Firm Attributes						
High Efficiency	α_1		0.023		0.027	−0.27
			(0.015)		(0.015)	(0.936)
# Modules	α_2		0.076		0.071	0.712
			(0.015)		(0.014)	(1.025)
Avg. Electricity Price	α_3		−7.962		−7.433	21.114
			(1.633)		(1.535)	(141.37)
Income Interactions						
log(Income) × Constant	σ_1					1.966
						(5.867)
log(Income) × High Efficiency	σ_2					0.065
						(0.234)
log(Income) × # Modules	σ_3					−0.185
						(0.318)
log(Income) × Avg. Electricity Price	σ_4					−7.599
						(36.735)
Income Distribution				Yes	Yes	Yes
Firm, Year FE		Yes	Yes	Yes	Yes	Yes
Median Own Price Elast.		−1.4	−1.39	−1.43	−1.33	−1.09
Median Outside Diversion		10.12%	10.03%	10.07%	10.07%	10.06%
J-Statistic		51.21	62.3	51.57	62.92	53.69

Notes: Estimation follows the procedure outlined in Section 5.1. There are 22,713 observations at the firm-county-quarter level. The nested logit models NL-1 and NL-2 divide price by county-quarter mean income whereas the random coefficients nested logit models RCNL-1, RCNL-2, and RCNL-3 use the full sample of incomes drawn from the ACS PUMS to estimate price sensitivities and other income interaction terms. High Efficiency is a binary variable that equals 1 if the installer offers high efficiency modules, # Modules is the number of module types offered by the installer, and Avg. Electricity Price is the county-quarter average electricity price. I report medians among all firm-county-quarter observations of own price elasticities and outside diversion rates. Standard errors clustered by county are reported in parentheses.

ing parameter estimates is through diversion ratios, which imply that not installing is the second-best choice for around 10% of consumers.

The coefficients on observable, time- and market-varying firm attributes are enlightening; however, these coefficients are imprecisely estimated in the RCNL model that includes income interactions. Consumers like installers that offer high efficiency PV modules, though this appears to be primarily driven by higher income households. Consumers also appear to like installers that offer a greater variety of solar PV module types; however, this appears to be driven by lower income households. Low income households appear quite sensitive to average electricity prices, deriving greater utility from PV installation in periods of high electricity

prices, whereas the PV installation decisions of high income households are insensitive to these prices. Overall, these results should be interpreted with caution given the lack of precision of many of these estimates.

Given that I do not endogenize firms' decisions over specific time-varying product attributes in my model of installers and given the desirable substitution patterns of the RCNL model, I use the demand system estimates in column (3) in the second stage of estimation as well as 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 A3.[24] 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 C. 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 C1.

6.1.2 Exit Policy Function

As outlined in Section 5.1.2, I select a set of candidate regressors to include in the logit regression of the discrete exit decision via LASSO. Of a total of 242 candidate regressors, this process selects 206 non-zero regressors. The resulting logit regression has a deviance of 12.34%. As shown in Appendix Figure D2, fitted exit probabilities for incumbents that exit are larger than those for incumbents that continue: the average fitted exit probability is 14.88% for exiting incumbents and 7.72% for continuing incumbents.

6.1.3 State Transitions

I report estimates of the AR(1) transition processes for the aggregate state variables, including the demand state and the inclusive value state as well as the county-quarter average price per watt, in Appendix Table A1. I report estimates of the AR(1) transition processes for the firm-level state variables, including quality and hardware cost as well as a firm's price per watt, in Appendix Table A2. For each state variable, I report results from two separate specifications, one each with and without county-specific intercepts. The estimated transition process for each state variable is stationary.

[24]Overall, my elasticity estimates are similar in magnitude to those found elsewhere in the literature; however, most existing elasticity estimates report the elasticity of the adoption decision with respect to price rather than specific firms' own-price elasticities. Hughes and Podolefsky (2015) and Bollinger and Gillingham (2019) estimate static elasticities of -1.2 and -1.0 for California households. Gillingham and Tsvetanov (2019) estimate a static elasticity of -0.65 for Connecticut households. Estimates from De Groote and Verboven (2019) imply an upper bound on the static elasticity of -6.6. Gerarden (2022) estimates a static elasticity of demand for PV *modules* of -1.48.

6.2 Second Stage Estimates

6.2.1 Production and Exit Cost Estimates

I report results from the joint estimation of the main production and exit cost parameters in Table 3. I allow for three distinct models of spillovers across rival firms and report the resulting production cost, effective experience, and exit parameter estimates in each column.

The parameter that determines learning economies in marginal installation costs, γ, is negative, large-in-magnitude, and precisely-estimated across all sets of estimates in Table 3. I report estimates from a model that assumes no serial correlation in firm's productivity shocks (i.e., $\rho = 0.0$) in column (1): as expected, the resulting learning exponent is larger than the analogous specification that accounts for serial correlation in column (2), highlighting the importance of modeling serial correlation in this unobservable. The learning exponent estimates in columns (2) through (4) imply that a 1% increase in effective experience decreases marginal installation costs by between 0.73 and 0.88%. To aid in interpretation of the learning exponent, I report the "Spence coefficient," which describes the proportional reduction in marginal installation costs for a doubling of effective experience (Spence, 1981). The Spence coefficient ranges from 0.398 to 0.459 across columns (2) through (4), which implies the presence of considerable learning-by-doing in installation-specific costs.

Importantly, I estimate substantial serial correlation in firms' productivity shocks, κ_{jmt}. Across columns (2) through (4), estimates of the productivity shock serial correlation parameter, ρ, range from 0.55 to 0.86. Though large in magnitude, these estimates are qualitatively similar to other estimates of serial correlation in learning curve estimation: Benkard (2000) reports first-order serial correlation coefficients between 0.73 and 0.97 in his learning curve estimates for commercial aircraft production. My estimates of the serial correlation parameter, ρ, suggest that ignoring the potential for serial correlation would produce biased estimates. Indeed, while my estimates of the learning exponent are large in magnitude, failure to account for serial correlation biases these estimates further away from zero as evidenced by the results in column (1). This is due to the fact that firms with serially correlated, positive productivity shocks are likely to have greater experience and relatively low marginal production costs. Failure to account for this high degree of serial correlation in counterfactuals could also lead to incorrect conclusions about counterfactual policies. I therefore focus on the results in columns (2) through (5) moving forward.

The parameters on rival firms' experience levels describe the marginal contribution of other firms' cumulative production to a firm's effective experience. I estimate the production cost and exit parameters using three different parameterizations of effective experience, each of which implies a distinct model of learning spillovers across firms. All three parameter-

Table 3. Estimated Production and Exit Cost Parameters

		Base Specifications		Alt. Spillovers		Forgetting
	Parameter	(1)	(2)	(3)	(4)	(5)
Production Cost Parameters						
Base Cost	c_0	2.067	2.145	2.041	2.006	2.179
		(0.064)	(0.046)	(0.065)	(0.158)	(0.006)
Learning Exponent	γ	−1.098	−0.733	−0.887	−0.731	−0.649
		(0.085)	(0.083)	(0.098)	(0.099)	(0.017)
Productivity Serial Correlation	ρ		0.838	0.860	0.554	0.979
			(0.099)	(0.271)	(0.285)	(0.023)
Effective Experience						
Industry Experience:						
Total	θ_1^E	0.854	0.817			0.760
		(0.014)	(0.059)			(0.010)
In-market	θ_2^E			0.747		
				(0.002)		
Out-of-market	θ_3^E			0.687		
				(0.043)		
Same Manufacturer	θ_4^E				0.749	
					(0.008)	
Other Manufacturer	θ_5^E				0.740	
					(0.088)	
Forgetting Parameter	δ					0.954
						(0.009)
Exit Parameter						
Mean Scrap Value	σ_ϕ	4.585	2.488	3.451	4.550	1.351
		(0.656)	(1.148)	(1.262)	(1.348)	(0.267)
N		11,581	11,581	11,581	11,581	11,581
Spence Coefficient $(1-2^\gamma)$		0.532	0.399	0.459	0.398	0.362

Notes: Estimation follows the procedure outlined in Section 5.2. I normalize experience variables by the industry total experience level in the first quarter of the sample (Q1 2008). The base cost parameter is therefore the marginal installation cost when a firm's effective experience equals the industry total experience in Q1 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 (see discussion in Section 6.2.1). 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. Bootstrapped standard errors clustered by county using 200 replications are reported in parentheses.

izations normalize the marginal contribution of a firm's own experience to one, which implies that all parameters on rival firms' experience can be interpreted as marginal contributions relative to a firm's own experience.

The first parameterization groups the experience of all rival firms together, which is consistent with a model of learning spillovers where there are no differences across firms. Column (2) reports the baseline estimates from this parameterization: the learning benefits from a 1 unit increase in total industry experience generates 82% of the learning benefits

from a 1 unit increase in own-firm experience. This implies that learning spillovers from rivals are substantial and not far from individual learning in terms of their contribution to firms' experience-based cost reductions.

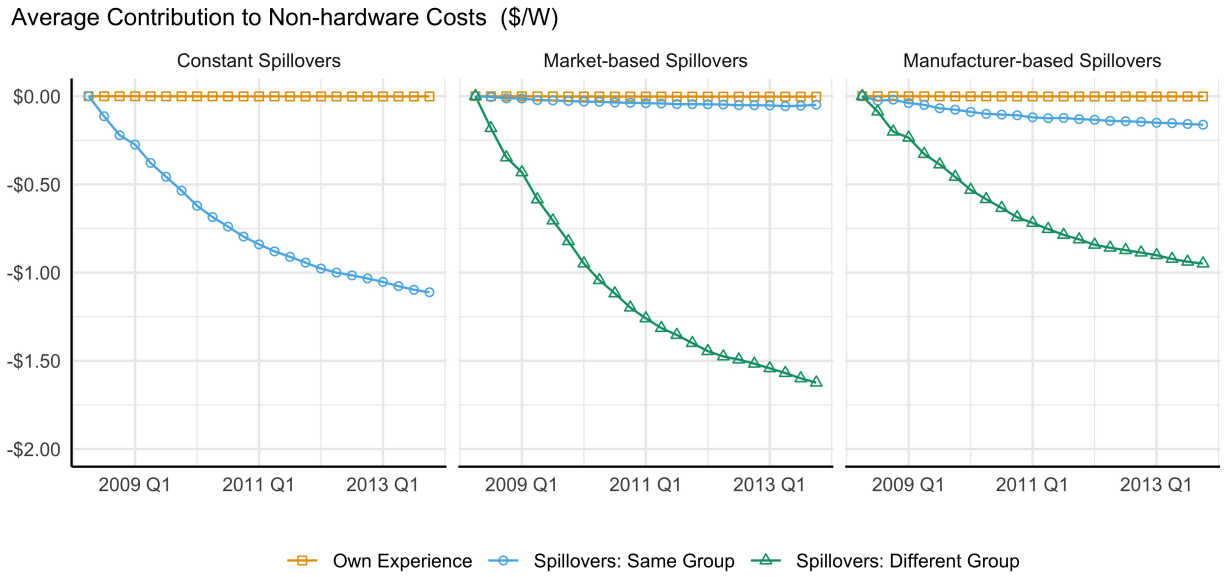
To tease apart the mechanisms underlying this finding of non-trivial spillovers, I estimate two alternative parameterizations of effective experience. The first allows for differential spillovers based on whether or not firms are within the same market. This alternative parameterization is consistent with models of learning spillovers based on geographic proximity, where knowledge transfer between firms is driven by the movement of workers between firms, the visibility of rivals' installation practices, or other forms of diffusion facilitated by proximity.[25] This parameterization is also consistent with a model of regulator learning: since firms operating in the same county are likely to face similar regulatory regimes in the form of local permitting requirements and utility rebate processing practices, any passive learning by regulators that influences permitting or rebate processing costs associated with PV installations would show up as greater learning benefits from rivals within the same county.[26] I report estimates from this parameterization of effective experience in column (3) of Table 3. I find that the marginal contribution of in-market rivals' experience is larger relative to that of out-of-market rivals; however, the difference between the two is minor. Thus, the effective experience parameters in column (3) offer suggestive yet inconclusive evidence in favor of one of these possible models of learning spillovers.

The second alternative parameterization of effective experience allows for differential spillovers based on whether or not firms install modules from the same manufacturer. This parameterization is consistent with a model of learning where knowledge transfer and learning-by-doing is facilitated by PV module manufacturers, perhaps through manufacturer-

[25] I explore the plausibility of this model of learning diffusion by estimating the empirical relationship between total PV installation-related employment, average PV installation-related wages, and the number of residential PV installations within a county using data from the US Census Bureau's Quarterly Census of Employment and Wages. As shown in Figure A4, I find strong, positive relationships between PV installation-related employment/wages and the number of PV installations. While this does not directly speak to inter-firm worker mobility, it does suggest that there is some degree of worker movement in the PV installation-related labor pool. I am in the process of gaining access to a large sample of public employment profiles from the social networking site LinkedIn that will allow me to more directly measure the extent of worker movement between solar installers.

[26] I explore the plausibility of this model of regulator learning in my context using several additional data sources. First, I examine differences in CSI rebate processing times across IOUs, each of which is responsible for administering CSI rebates within their service area. As shown in Figure A5, there is suggestive evidence that IOU rebate processing times decreased over the sample period, which would influence installation-specific costs by decreasing overall project completion times and indirectly increasing installers' installation capacity in a given period. This suggestive evidence would result in greater spillovers across firms within the same county to the extent that IOU service territories overlap county boundaries. To further explore this model, I examine PV-related permitting times for San Diego County, the single largest county by number of installations in my data. Figure A6 shows suggestive evidence that—at least in the case of San Diego County—regulators' permitting times decreased over the sample period.

Figure 5. Cumulative Contribution of Experience to Installation Costs



Notes: This figure shows the average cumulative contribution of different forms of experience to installation-specific costs over the full estimation sample period, 2008 to 2013. Each panel uses a set of estimates from Table 3: the left panel uses estimates from column (2), the center panel uses estimates from column (3) and the right panel uses estimates from column (4). The left panel therefore does not distinguish spillovers across firms, whereas the center and right panel allow for different spillovers based on whether or not firms are within the same market or install modules from the same manufacturer, respectively. Average contributions are calculated by evaluating the marginal contribution of each source of experience using the estimates from Table 3 at the average observed experience levels across the estimation sample in a given period.

led training or improvements in module design that directly impact the installation process. I report estimates from this parameterization in column (4) of Table 3. I cannot reject the hypothesis that the marginal contribution of rivals installing modules from the same manufacturer is equivalent to that of rivals installing modules from other manufacturers.

While the effective experience parameters in Table 3 describe contributions of different types of experience on the margin, the *total* contribution of each form of cumulative production depends on both the parameter estimates and the empirical distributions of experience. Figure 5 plots the average cumulative contribution of different forms of experience to marginal installation costs over the full estimation sample period. Despite all spillover estimates in Table 3 implying smaller marginal contributions to learning relative to own-firm experience, learning spillovers in practice drive the bulk of estimated installation-specific cost reductions. I calculate that the average reduction in installation costs over the full sample period resulting from a firm's own experience is between \$0.001 and \$0.002. By contrast, using the parameter estimates from column (2) of Table 3, I find that the average reduction

in installation costs over the sample period resulting from other firms' experience is \$1.11. Moreover, despite experience from in-market rivals having a larger marginal contribution to learning than that from out-of-market rivals, the difference in magnitude of these different experience levels means that knowledge spillovers from out-of-market rivals have a larger impact on cost reductions overall.

I estimate a version of the model that allows firms' experience-based cost reductions to depreciate over time. Such "forgetting" models are relevant in several empirical settings characterized by boom and bust cycles—including aircraft manufacturing, ship building, auto manufacturing, and oil extraction—likely due to prolonged down time during busts (Benkard, 2000; Kellogg, 2011; Levitt et al., 2013; Thompson, 2007). It is unlikely that knowledge depreciation plays a major role in the current setting, which is characterized by substantial, persistent growth; however, factors such as worker departures from one quarter to the next could result in incomplete knowledge retention. To test this, I estimate a version of the model with constant industry spillovers that assumes that each component of effective experience accumulates according to a perpetual-inventory process: for example, firm j's own experience in market m in period t is defined by $E_{jmt} = \delta(E_{jmt-1} + q_{jmt-1})$.[27] The parameter δ defines the rate of forgetting or, more precisely, knowledge retention.[28] The results from this model, which I report in Column (5) of Table 3 indicate minimal forgetting from one period to the next, with an estimated quarterly retention parameter of 0.954. Allowing for forgetting does reduce the estimated magnitude of learning economies and spillovers; however, the qualitative findings of the full retention specifications appear robust to this alternative model with forgetting.

Though the magnitude of learning-by-doing implied by the production cost estimates in Table 3 is substantial, it aligns with existing estimates of installation-specific costs during this period. Figure 6 compares the average installation-specific, non-hardware costs implied by the parameter estimates reported in column (2) of Table 3 with comparable, publicly-available estimates from the National Renewable Energy Laboratory (NREL). Fu et al. (2016) construct these estimates from publicly-available data on installations, corporate filings from

[27] An alternative parameterization of forgetting could be $E_{jmt} = \delta E_{jmt-1} + q_{jmt-1}$. I choose to assume incomplete retention of knowledge based on the prior periods' production as this is consistent with a model of forgetting driven by worker turnover.

[28] The full parameterization of effective experience in this specification with constant spillovers and forgetting becomes:

$$\tilde{E}_j\left(s_{mt}; \theta^E, \delta\right) = \delta(E_{jmt-1} + q_{jmt-1}) + \theta_1^E \left(\sum_m \sum_{k \neq j} \delta(E_{kmt-1} + q_{kmt-1}) \right)$$

This assumes that the rate of forgetting is constant across own and rival cumulative production.

publicly-traded installers, and existing engineering studies of PV installation costs. I report the sum of Fu et al. (2016)'s cost estimates for the following categories: installation labor costs per watt as well as permitting, inspection, and installation costs per watt.[29] As is clear from Figure 6, the installation-specific costs implied by my model estimates reasonably match not only the magnitude of Fu et al. (2016)'s estimates, but also the rate of learning over time implied by their estimates. Moreover, this is true both within my estimation sample as well as in the three years after my estimation sample.

Figure 6 also shows the average marginal non-hardware costs implied by the model assuming that firms are completely static in their quantity-setting decisions. I calculate these costs from the standard, static quantity-setting first order condition using observed prices, observed hardware costs, and estimated price elasticities. Since I only estimate the demand model in-sample, I therefore cannot estimate these static costs out-of-sample as I do in the case of the dynamic model, which uses observed experience levels to make out of sample predictions. While the installation-specific costs implied by my dynamic model reasonably match the magnitude and shape of Fu et al. (2016)'s estimates over time, the static estimates are nearly twice as large as both the dynamic model estimates and the NREL benchmark. This emphasizes the importance of accounting for learning-by-doing and dynamic incentives in firms' quantity-setting decisions in my model: without accounting for these factors, I would substantially overestimate non-hardware costs and underestimate the rate of learning.

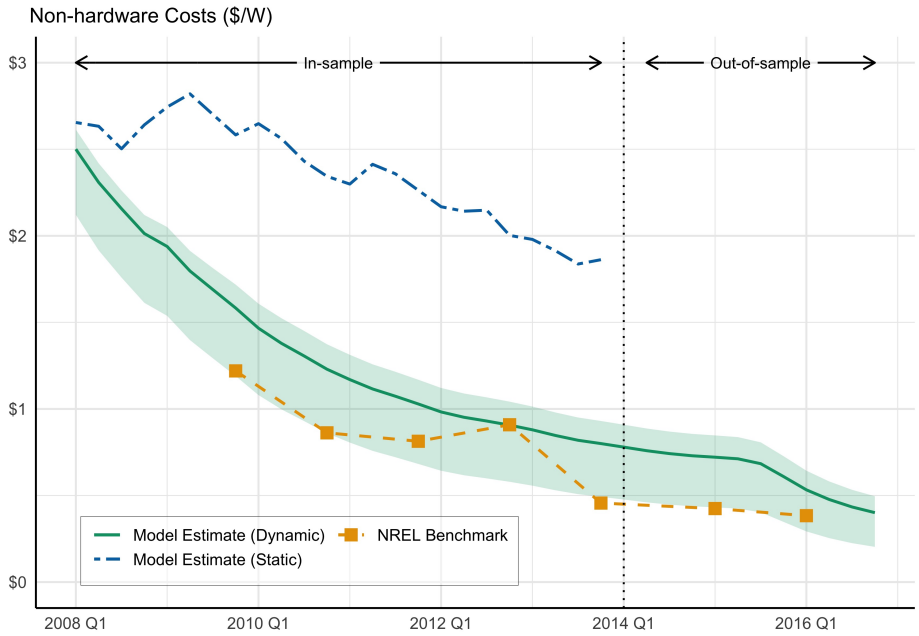
The remaining production and exit cost parameters are precisely estimated and consistent across specifications. The mean scrap value, σ_ϕ, is between \$248,800 and \$455,000. To ensure the readability of the effective experience parameters and improve numerical stability in estimation, I normalize all experience terms by the total industry experience in the first period of the estimation sample, Q1 of 2008. It is therefore possible to interpret the base cost parameter as the marginal installation cost when a firm's effective experience equals the industry total experience in Q1 2008. Estimates of c_0 range between \$2.0 and \$2.2 per watt.

6.2.2 Entry Cost Estimate

The final target parameter in the second stage of estimation is the mean entry cost. Table 4 reports estimates of the entry cost parameter using different data-driven approaches to defining the number of potential entrants in each market. In particular, I assume that the number of potential entrants is either one or two times the median, mean, or maximum number of observed entrants for each market over the sample period. This data-driven approach

[29] In their report, Fu et al. (2016) also include estimates of installer overhead and net profit in their estimate of overall soft costs—or non-hardware costs—per watt. I exclude these categories in Figure 6 as markups are strictly not a component of installers' installation-specific costs and the estimates of overhead include fixed business operating costs that are not a relevant to my installation-specific costs estimates.

Figure 6. Comparison of Model-implied and Public Estimates of Installation Costs



Notes: This figure compares the average non-hardware costs implied by the dynamic model estimates reported in column (2) of Table 3 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, observed hardware costs, and estimated price elasticities. The shaded area shows the 95% bootstrap confidence interval for the dynamic model estimates.

is common in the literature, with most studies using some multiple of the maximum observed entrants (Barwick and Pathak, 2015; Barwick et al., 2021; Seim, 2006). Unsurprisingly, the estimated entry cost parameter increases as the pool of potential entrants increases, with estimates of the mean entry cost ranging from $2.7 to $8.1 million. While these estimates may seem large—particularly in comparison to the estimates of mean exit cost—it is important to note that this is the unconditional mean entry cost. Conditional on entering, the average entry cost ranges between $943,900 and $1.05 million. Moreover, these estimates of the average entry cost conditional on entry are similar in magnitude to existing estimates as well as available information on publicly-traded installers during the estimation period.[30]

[30] Feldman et al. (2013) use engineering and accounting data to construct estimates of the upfront costs associated with developing, constructing, and arranging third-party-financed residential PV systems and find that fixed business expenses amount to $600,000/year for a representative firm in 2012. This translates to roughly $5 million in perpetuity, assuming an annual discount factor of 87.5%. SolarCity Corp., which was acquired by Tesla in 2016 at a cost of $2.6 billion had a market capitalization of roughly $5.6 billion at the end of 2013 Q4, having installed roughly 75 megawatts (MW) of residential capacity during that calendar year. Scaling this down using the average annual installed capacity across all installer-county observations in my data—approximately 35 kilowatts (or 0.35 MW)—this translates to a roughly $2 million valuation for

Table 4. Estimated Entry Cost Parameter

	Parameter	(1)	(2)	(3)	(4)
Mean Entry Cost	σ_ω	27.286	28.665	34.931	81.240
		(1.628)	(1.696)	(3.153)	(6.687)
Potential Entrant Def.		2×median($\bar{N}_{mt}$)	2×mean($\bar{N}_{mt}$)	1×max($\bar{N}_{mt}$)	2×max($\bar{N}_{mt}$)
N		8,763	9,014	10,193	20,386
N^e		311	327	383	766
$\bar{\omega}_{jmt}$|entry		9.439	9.514	9.785	10.505

Notes: Estimation follows the procedure outlined in Section 5.2. 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. $\bar{\omega}_{jmt}$|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.

6.3 Robustness

I explore the robustness of my main estimates of the model parameters in several ways. First, I re-estimate the model using a different set of moments. In particular, rather than relying on the moment from the innovation in firms' unobserved productivity shocks, I use these productivity shocks themselves interacted with a set of instruments, Z_{jmt}, that satisfy the relevance and exogeneity conditions to form a set of moments from which I can recover valid estimates of the production cost and experience parameters. As I discuss in Section 5.2, estimation from the unobservable productivity shock alone would likely produce biased estimates. I therefore interact the following instruments with the unobservable: realized consumer rebates; observed firm-level deviations from average utility interconnection processing times; current and lagged cost shifters (aluminum and polysilicon commodity prices; county-quarter specific wage rates); and current and lagged demand shifters (retail electricity prices, observed solar insolation). To further account for the endogeneity concern due to potentially serially-correlated productivity shocks, I include a common time trend in firm's marginal costs in this alternative estimator. I report estimates from this approach alongside my baseline production and exit cost parameter estimates in Table A3. Results are remarkably robust across these estimators: estimates of learning economies and exit costs are similar across each approach. While I estimate slightly smaller spillovers across rival firms from the alternative estimator, the magnitude of these spillovers is still quite large.

I also test the sensitivity of my main model estimates to alternative assumptions on the quarterly discount factor, β, which is common to all firms. In my main estimates and the

the county-level operations of an average firm in my data.

counterfactual analysis that follows, I assume a quarterly discount factor that corresponds to an annual discount factor of 0.875. This discount factor is equal to that of Gerarden (2022)'s model of solar PV manufacturers and is similar in magnitude to that estimated by De Groote and Verboven (2019) for PV-adopting households. I re-estimate the model using two alternatives corresponding to annual discount factors of 0.9 and 0.8, each of which can be found in the literature (Igami, 2017; Ryan, 2012). The latter discount factor also serves as a proxy for dynamic regulatory environments as it renders future payoffs less relevant for decisions in the current period. Thus, estimating the model with an annual discount factor of 0.8 serves as a test of the robustness of my assumption that firms perceive changes to the subsidy policy regime as permanent in my setting.

I report estimates using these alternative discount factors alongside my baseline estimates in Table A4. Overall, results are qualitatively consistent across alternative discount factors; however, the costs implied by the resulting parameter estimates increase with the discount factor.[31] This is consistent with the model: a higher discount factor implies higher expected values from future operations, which means that higher cost estimates are necessary to rationalize the observed patterns of entry, exit, and quantity choices. The qualitative consistency across discount factors suggests that assuming stationary beliefs over the policy environment is reasonable in this setting.

Given that I define a firm's experience as a function of their cumulative production, it is possible that some of the estimated learning economies actually represent static scale economies. Separately identifying learning-by-doing and static scale economies is a challenge throughout the empirical literature estimating learning curves. I test the sensitivity of my main model estimates to the presence of static economies of scale by estimating a version of the preferred specification that includes a static, contemporaneous measure of firm size as a state variable when approximating firms' value function. Table A5 compares model estimates with and without the measure of firm size as a state variable and finds qualitatively consistent results. Similar to Benkard (2000), the estimated learning rate increases slightly when including scale as a state variable.

7 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

[31] Specifically, the mean scrap value and base cost parameters all increase with the discount factor, β, as these parameters are positively related to firm costs. The effective experience and learning exponent parameters each decrease with the discount factor: for a given realization of firm and industry experience levels, lower parameter values imply higher costs.

the model's equilibrium. I begin this section by briefly describing my approach to solving the model and summarizing counterfactual results (Section 7.1). I then compare the fit of model-predicted outcomes under the baseline policy environment with observed data (Section 7.2) before discussing results from three sets of counterfactual scenarios: I first quantify the impacts of the California Solar Initiative and evaluate alternative consumer subsidy designs (Section 7.3). I then evaluate a set of alternative industrial policies, including an entry subsidy (Section 7.4), and alternative climate policies, including a carbon tax (Section 7.5).

Results from these counterfactual simulations provide three findings. First, I find that the CSI contributed to growth in the installation industry: 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. Second, I find that the CSI's decreasing rebate structure results in the highest welfare—which includes consumer and producer surplus as well as net environmental damages—in comparison to alternative rebate designs and other forms of climate policy, including a $30/ton carbon tax, when focusing solely on outcomes in the state's solar PV market. This suggests that some sort of technology-specific policy with a design similar to that of the CSI may be desirable as part of a broader portfolio of climate policies. Finally, I find that consumer subsidies may not be as effective as more conventional forms of industrial policy such as entry subsidies: replacing the CSI with entry subsidies of varying sizes results in more installation and entry as well as higher welfare.

7.1 Counterfactual Solution Method

My approach to solve for the model's equilibrium builds on the method of Sweeting (2013), which adapts parametric policy iteration (Benitez-Silva et al., 2000) to allow for value function approximation. I provide a detailed description of this approach in Appendix F. Solving the model involves two steps: first, solving for the new Bellman equation, policy functions, and product market equilibrium in a given period and second simulating the industry forward one period. In each of the counterfactual scenarios that I describe below, 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 sample—the last quarter of 2013.

I implement the first step of this counterfactual solution method via a fixed point algorithm. This fixed point algorithm, which I describe in greater detail in Appendix F, produces conditional exit probabilities and value function approximating coefficients in a given period, which I can use to calculate conditional entry probabilities. 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 the model in that period.

The above process results in a single industry path over the full counterfactual period. Given that I take single draws from the conditional exit and entry probabilities as well as the state transition processes each time I simulate the industry forward one period, I repeat the above process multiple times, averaging market outcomes across a number of distinct, forward-simulated industry paths. In practice, I repeat the process of simulating the model forward over the 6 years in the estimation sample 60 distinct times for each counterfactual scenario and then average key outcomes across all 60 model runs.

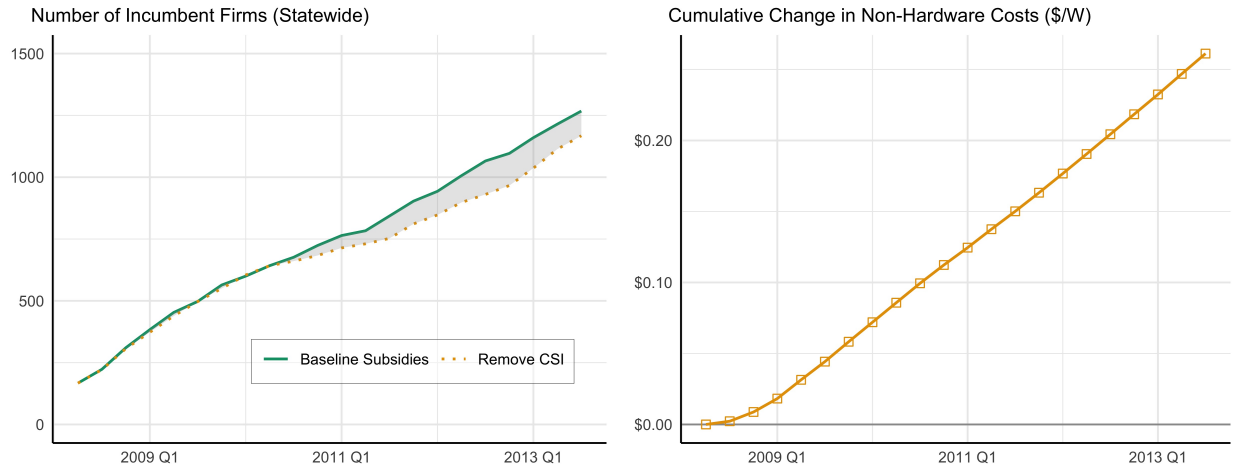
A number of policy- and welfare-relevant outcomes are endogenous to the model, including quantities, prices, and marginal costs as well as firm entry and exit. These objects allow me to calculate firms' profits and consumer surplus. Given that a key policy justification for incentivizing solar PV adoption is to address environmental externalities associated with legacy electricity generation sources, I use the quantities of PV adoption predicted for each counterfactual scenario to conduct a back-of-the-envelope calculation of any changes in environmental damages. I combine my model-predicted quantities with geographically-differentiated estimates of the marginal environmental benefits of additional solar capacity in the US from Sexton et al. (2021). Additional details are available in Appendix F.

7.2 Model Fit under Baseline Policies

Before turning to results from counterfactual policy simulations, I compare the fit of model-predicted outcomes under the baseline policy environment with observed data. This exercise complements the discussion of model fit in Section 6.2 and is particularly important as the predicted equilibrium outcomes from simulating the model under the baseline, realized policy environment serve as points of comparison for each of the counterfactual policy scenarios.

As shown in Figure A7, the model-predicted, statewide entry and exit probabilities reasonably match those observed in the data. In particular, the model-predicted, statewide exit probability matches both the level and shape of the empirical exit probability from 2008 to 2013. In the case of the statewide entry probability, the model matches the overall number of entries well; however, the model appears to under-predict entry early on and over-predict entry late in the second half of the period from 2008 to 2013. Overall, the model matches observed entry and exit patterns reasonably well.

Figure 7. Impact CSI Subsidies on Market Size and Installation Costs



Notes: This figure shows the model-predicted number of active firms with and without the California Solar Initiative (CSI) subsidies (left) and the cumulative increase in marginal installation costs resulting from removing the CSI subsidies. Counterfactual and baseline outcomes reported above are averaged across 60 distinct, forward-simulated industry paths. The cumulative change in non-hardware costs is the cumulative difference in average marginal installation costs per watt between the counterfactual policy scenario without the CSI subsidies and the baseline scenario with the full CSI subsidies.

7.3 Counterfactual Results: California Solar Initiative

Figure 7 summarizes the effect of the CSI's consumer subsidies on the number of active firms over time. The figure shows the model-predicted number of active firms statewide with and without the CSI as well as the cumulative difference in marginal non-hardware costs from removing the CSI relative to the baseline of existing subsidies. The CSI increases the number of installed PV systems by 4% and the number of active firms at the end of the estimation period by around 9% relative to a scenario that removes the CSI entirely. As shown in Figure 7, the effect of the CSI on the number of active firms occurs gradually over time. This matches the gradual reduction in installed PV systems resulting from CSI removal and is likely due to the incremental impacts on reduced knowledge transfer. As shown in the right panel of Figure 7, the cumulative impact of CSI removal on marginal non-hardware costs, which captures the cumulative increase in marginal production costs for a firm active throughout the estimation period, steadily increases over time.

I further explore the role of learning spillovers in driving this gradual reduction in the number of active firms by re-simulating the industry under both the baseline subsidy policy regime and the CSI removal counterfactual, but with knowledge spillovers that are half as large as the main model estimates. Figure A8 reports the change in the model-predicted number of active firms from CSI removal relative to the baseline for each level of learning

spillover. Overall, cutting the size of learning spillovers in half reduces the impact of CSI removal on the number of active firms and delays the timing of CSI removal's impacts taking effect. However, the qualitative result remains the same: CSI removal with smaller learning spillovers results in a gradual reduction in the number of incumbent firms.

These results offer two main findings: first, the qualitative finding of CSI removal reducing the number of active firms is robust to smaller spillovers, suggesting that the challenges associated with identifying learning spillovers that I discuss in Section 5.2.4 do not drive this qualitative finding. Second, the fact that the magnitude and timing of the CSI removal scale with the size of learning spillovers indicates that non-appropriable learning serves as the main mechanism through which consumer subsidies impact the number of incumbents.

I report estimates of the change in welfare in Table 5. Overall, removing the CSI reduces welfare, primarily through reduced consumer surplus and reduced revenues, as well as higher production and exit costs. To explore the implications of the CSI's unique rebate design that decreased over time, I estimate two additional CSI counterfactuals: a flat rebate equal to the quantity-weighted average rebate amount under the baseline CSI design and a rebate that increases over time, inverting the steps of the baseline CSI design. The latter counterfactual is motivated by the finding of Langer and Lemoine (2022) that an increasing subsidy can allow policymakers to optimally price discriminate. I find that an increasing subsidy decreases welfare relative to the decreasing, baseline CSI subsidy, suggesting that in this case taking advantage of cost reductions from learning-by-doing may be optimal. Interestingly, the welfare losses from a flat subsidy are minor. Figure 8 plots the change in the number of active firms as well as the change in the number of solar PV installations relative to the baseline of full subsidies under each of these alternative CSI counterfactuals.

7.4 Counterfactual Results: Alternative Industrial Policies

The CSI's consumer subsidies increase the quantity of solar installations and the number of active firms, suggesting that they are effective at achieving the objective of growing the state's solar installation industry. How do the CSI's consumer subsidies compare to more common industrial policies, such as entry subsidies? I examine this by simulating a set of counterfactual policies in which I remove the CSI subsidies and replace them with entry subsidies of varying sizes. In particular, I separately replace the CSI subsidies with entry subsidies equal to one quarter, one half, and three quarters of the preferred estimate of the mean entry cost.[32]

As shown in Table 5, replacing the CSI subsidies with these entry subsidies greatly

[32]The preferred mean entry cost is \$2.7 million, so the entry subsidies that I implement are \$675,000, \$1.35 million, and \$2.025 million per entrant, approximately.

Table 5. Estimated Changes in Welfare under Counterfactual Policy Scenarios

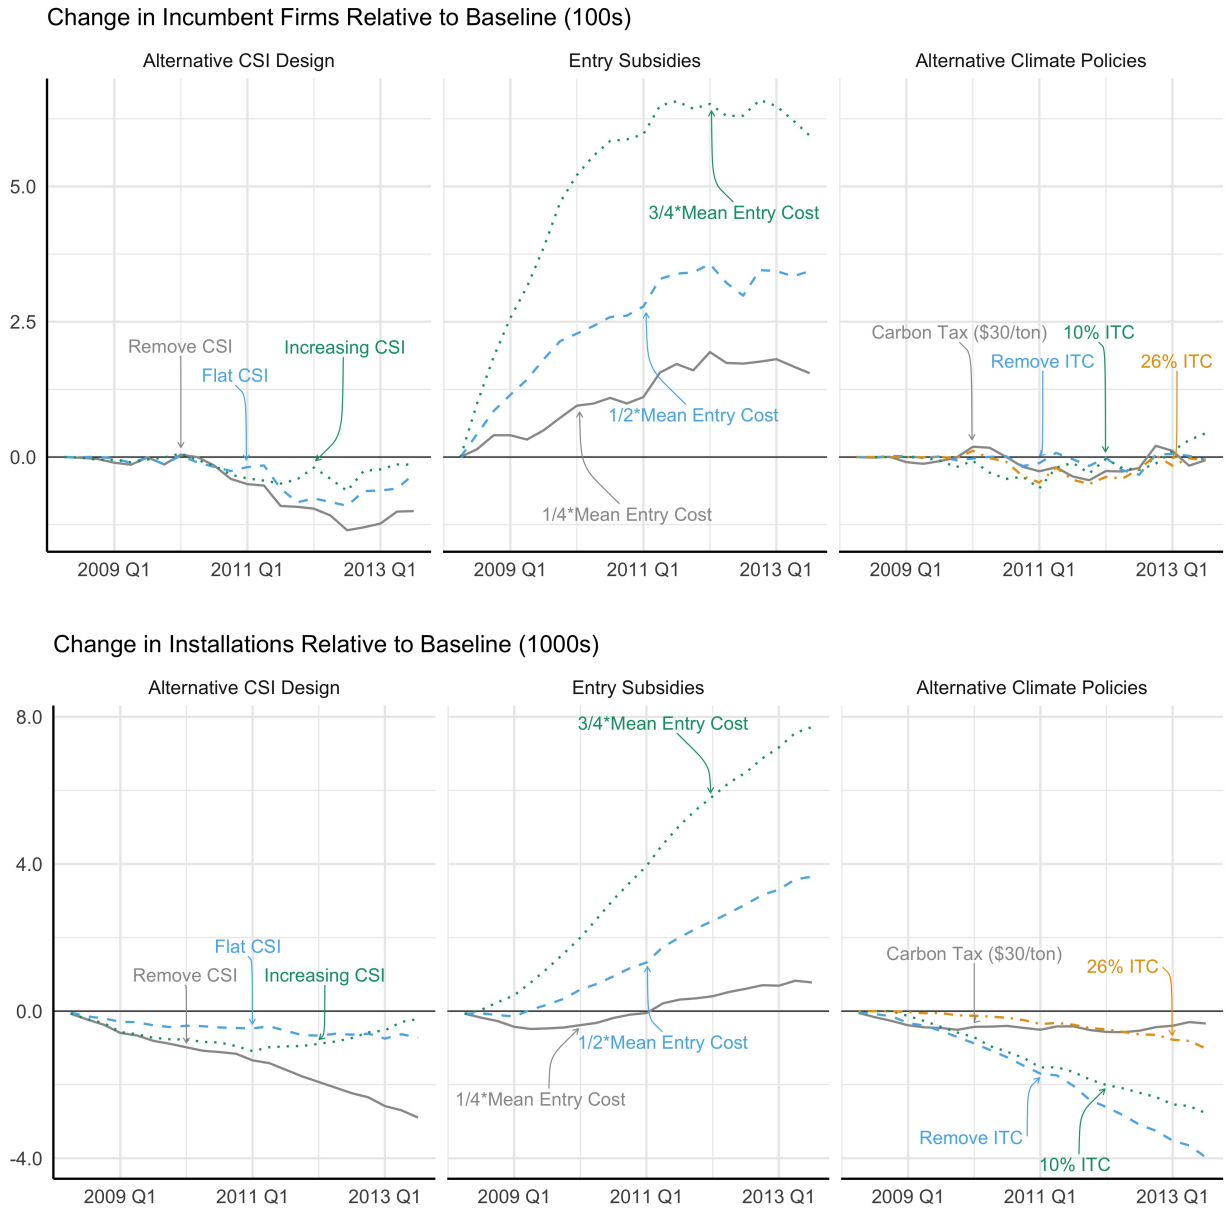
	Welfare Components ($M)						
Scenario	ΔCS	ΔEB	ΔRev.	$-\Delta Cost_{prod}$	$-\Delta\phi$	$\Delta\omega$	ΔTotal
CSI Counterfactuals:							
Remove CSI	-22.8	-6.6	-20.9	-18.8	-63.3	115.9	-16.6
Flat CSI	-0.7	-1.8	-8.5	-7.7	-39.5	57.6	-0.6
Increasing CSI	-50.0	-1.4	-11.5	-10.8	-28.6	58.1	-44.2
Entry Subsidies:							
1/4*Mean Entry Cost	134.7	1.2	-6.2	-5.7	120.7	-43.9	200.7
1/2*Mean Entry Cost	472.6	7.7	13.4	13.1	247.6	240.3	994.7
3/4*Mean Entry Cost	785.1	17.3	36.4	33.7	508.7	1089.5	2470.7
Alternative Climate Policies:							
Carbon Tax ($30/ton)	-45.1	-1.2	-9.0	-8.3	-8.9	10.4	-62.1
Remove Federal ITC	-2.3	-8.6	-48.3	-44.3	-22.3	21.5	-104.3
10% Federal ITC	-7.0	-6.2	-34.2	-31.1	-18.6	15.3	-77.7
26% Federal ITC	16.9	-2.0	-12.6	-11.2	-12.8	9.0	-12.7

Notes: This table reports model-predicted changes in welfare components as well as total welfare relative to the baseline scenario of existing consumer subsidy policies under three sets of counterfactuals: alternative California Solar Initiative (CSI) subsidy designs, varying levels of entry subsidies, and alternative climate policies. The baseline subsidy policies correspond to those outlined in Section 2.2: full CSI subsidies, 30% federal investment tax credit (ITC), and net-metering policy. 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 and alter the design of the federal ITC. Each counterfactual welfare component reported in the table represents an average across 60 distinct, forward-simulated industry paths under the given counterfactual. ΔCS is the change in consumer surplus, ΔEB is the change in environmental benefits, ΔRev. is the change in firm revenue, $\Delta\text{Cost}_{\text{prod}}$ is the change in production costs, $\Delta\phi$ is the change in scrap values, and $\Delta\omega$ is the change in entry costs. All values are reported in million 2013 USD.

increases total welfare, mostly through increased consumer surplus and lower net-of-subsidy total entry costs. Figure 8 shows that entry subsidies greatly increase not only the number of active incumbents, but also the number of solar PV installations. The effects of the entry subsidies on welfare, market size, and installations scale with the size of the subsidy.

While both the CSI and the counterfactual entry subsidies lead to industry growth, the two function in different ways. The additional demand induced by the CSI's consumer rebates leads to greater learning-by-doing and lower production costs relative to a world with no state-level consumer subsidies. This in turn makes it more profitable to be an incumbent, raises firm values, and leads to an increase in net entry. In the case of the counterfactual entry subsidies, the resulting reductions in entry costs increases pressure from potential entrants, lowers firm values, and greatly increases churn in the industry (i.e., increases both entry and exit) relative to the equilibrium with CSI subsidies as shown in Figure A9. The net expansion in the number of operating firms resulting from entry subsidies increases competition in the product market, which reduces firms' net markups (i.e., markups net of any markdowns) by

Figure 8. Change in Active Firms and Installation Quantities under Counterfactual Policies



Notes: This figure shows the model-predicted change in the number of active firms (top) and installation quantities (bottom) relative to the baseline scenario of existing consumer subsidy policies under three sets of counterfactuals: alternative California Solar Initiative (CSI) subsidy designs (left), varying levels of entry subsidies (center), and alternative climate policies (right). The baseline subsidy policies correspond to those outlined in Section 2.2: full CSI subsidies, 30% federal investment tax credit (ITC), and net-metering policy. 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 and alter the design of the federal ITC. Each counterfactual outcome plotted in the figure represents an average across 60 distinct, forward-simulated industry paths under the given counterfactual.

8 to 25% relative to the baseline CSI subsidies depending on the size of the entry subsidy. Figure A10 shows changes in average net markups under the entry subsidy and CSI removal counterfactuals. This combined with the fact that consumers face a substantially expanded choice set increases demand for solar installations, which leads to lower production costs through learning as shown in Figure A11.[33] This works to further amplify the effects of the entry subsidies in a similar fashion to the CSI's consumer subsidies, which helps to explain the finding that entry subsidies outperform the CSI subsidies in terms of number of installations, number of operating firms, and aggregate welfare.

The fact that the counterfactual entry subsidies raise welfare, result in greater entry, and generate more take-up of solar PV relative to the CSI suggests that some form of entry policy may be preferable to consumer subsidies in this setting. It is important to note, however, that these counterfactual entry subsidies have substantial fiscal costs, ranging from $2.1 billion to $8.4 billion. While these are similar in magnitude to the $2.2 billion in consumer subsidies under the CSI, the CSI was paid for by California's electricity ratepayers and therefore represents a transfer between all ratepayers and CSI-adopting households. Without specifying a revenue-generating policy for these counterfactual entry subsidies, it is difficult to fully compare them to the CSI. Moreover, there are likely political costs associated with implementing entry subsidies that may not apply consumer subsidies (e.g., the former are seen as a windfall to industry), therefore making the latter more appealing in practice.

7.5 Counterfactual Results: Alternative Climate Policies

While the CSI's consumer subsidies appear to increase both solar take-up and solar industry size, this policy is still second-best relative to an optimal carbon price to address the global externality problem that is climate change. However, as the CSI counterfactual results indicate, consumer subsidies for clean technologies may still be desirable given their ability to not only mitigate emissions externalities, but also address technology market failures resulting from knowledge spillovers in learning-by-doing. Particularly at carbon pricing levels observed in practice, which are often far below the socially-optimal level, policies that address these market failures in specific technologies may outperform carbon pricing alone, suggesting that some combination of carbon pricing and industrial policy may be desirable.

To explore this dynamic empirically, I implement a counterfactual that compares a hypothetical carbon tax of $30/ton of carbon dioxide emissions to the CSI's consumer subsidies.[34]

[33] An expanded choice set leads to greater demand through the parametric assumption of a logit demand model, which implicitly assumes consumers have a taste for variety.

[34] A carbon tax of $30/ton is roughly equal to the average price of emissions allowances in the secondary market under California's existing cap-and-trade program from 2020 to 2023 (for additional information, see: https://ww2.arb.ca.gov/our-work/programs/cap-and-trade-program). California's cap-

I do so by estimating the likely impact of a $30/ton carbon tax on retail electricity prices, which I can then translate into a change in the net present value of PV adoption due to households energy consumption and net-metering receipts.[35]

Overall, a $30/ton carbon tax results in marginally less PV installation and active installation firms relative to the CSI as shown in the right panel of Figure 8. This is due to the fact that the impact of a $30/ton carbon tax on the present discounted value of solar adoption, which operates through the effect of the tax on retail electricity rates, is minor in comparison to the CSI subsidies. As a result of the lower adoption and entry rates in this counterfactual, a $30/ton carbon tax results in lower welfare relative to the baseline CSI subsidies. It is important to note that the carbon tax in this counterfactual will have wide-ranging impacts beyond the solar industry, for which these counterfactual simulations do not account. Thus, this counterfactual exercise does not invalidate a carbon tax as a first-best climate policy. However, the results from this counterfactual indicate that additional policies to address specific technology market failures, such as those resulting from learning spillovers, may be necessary when implementing carbon pricing at levels observed in practice.

As a final set of counterfactual exercises, I keep the CSI subsidies in place but vary the level of the federal investment tax credit (ITC). This is relevant to the ITC extension included in the recently-passed Inflation Reduction Act of 2022, which increased the net-of-subsidy ITC for PV-adopting households from 26% to 30% through 2032.[36] I implement separate counterfactuals that keep the CSI subsidies in place but replace the 30% baseline ITC with net-of-rebate ITC rates of 0%, 10%, and 26%. Overall, these alternative ITC levels have limited impacts on installer entry despite having relatively substantial impacts on the number of solar installations. This is perhaps due to the fact that firms are able to charge higher net markups under these scaled-back ITC counterfactuals, since the remaining adopting households are quite price inelastic.

---

and-trade program, a form of carbon pricing, went into effect in 2013, the final year of my sample period. This counterfactual therefore asks what would be the outcome from removing the CSI and implementing a far more stringent carbon pricing policy for the full sample period, 2008 to 2013.

[35] In particular, I take data on the observed emissions rate for gas-fired boilers—the main fossil generator in the state—in California over 2008-2013 from the US Energy Information Administration, which allows me to translate a $30/ton carbon tax into an increased cost of generation per kilowatt-hour. I then assume full pass through of the tax to retail rates which, assuming a system capacity of 5 kW, an annual production amount equal to the statewide average, and a system lifespan of 25 years, I can translate into a per-watt subsidy to PV-adopting households. A $30/ton carbon tax roughly translates to an adoption subsidy equal to $0.61/watt of system capacity.

[36] Prior to the passage of the Inflation Reduction Act, the 30% federal ITC had been scaled back to 26% starting in 2020.

8 Conclusion

Policymakers face a growing imperative to invest in climate mitigation. Troubling trends in natural disasters and extreme weather events continue to emphasize the need to avoid the most catastrophic consequences of a changing climate. Despite economists' decades of advocacy in favor of carbon pricing as the first-best policy tool to address the global externality problem of climate change, policymakers increasingly turn to non-pricing forms of climate policy, including subsidies for the adoption of clean technologies, to achieve their decarbonization goals. As policymakers respond to growing calls to decarbonize the economy, rigorous evidence on the full suite of impacts of these policy tools is necessary to ensure an informed policy debate.

Despite playing a prominent role in the policy debates surrounding consumer subsidy programs, particularly in the context of clean technology subsidies, little empirical evidence exists on whether consumer subsidies can effectively increase the size of target industries. At issue is a tradeoff fundamental to infant industries characterized by learning economies and knowledge spillovers, where incumbent firms face a tension between reducing their own future costs and raising their rivals' future costs when making production decisions today. I shed light on this tradeoff in a policy-relevant setting by developing a model of the residential solar PV installation industry in California, which I use to simulate equilibrium market outcomes under a suite of counterfactual policies.

I find that learning economies and knowledge spillovers are substantial in this setting. As a result, the model implies that consumer subsidies are an effective tool for increasing not only the number of solar installations, but also the number of active firms. However, in comparison to more conventional forms of industrial policy such as entry subsidies, consumer subsidies may not prove as effective: I find that welfare, the quantity of solar installations, and the amount of new entry are all greater for an entry subsidy of a similar magnitude to existing consumer subsidies in California. Thus, while consumer subsidies may be an appealing second-best tool for policymakers to achieve their decarbonization and industrial policy goals, other approaches to addressing the twin externality problem of un-priced pollution and non-appropriable learning may be preferable.

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.

Aldy, Joseph, Todd Gerarden, and Richard Sweeney. 2022. "Investment versus Output Subsidies: Implications of Alternative Incentives for Wind Energy." *Journal of the Association of Environmental and Resource Economists*.

Argote, Linda, and Dennis Epple. 1990. "Learning Curves in Manufacturing." *Science*, 247(4945): 920–924.

Armitage, Sarah. 2022. "Technology Adoption and the Timing of Environmental Policy: Evidence from Efficient Lighting." *Working Paper*.

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. 2021. "Industrial Policy Implementation: Empirical Evidence from China's Shipbuilding Industry." *Working Paper*.

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.

Bistline, John, Neil Mehrotra, and Catherine Wolfram. 2023. "Economic Implications of the Climate Provisions of the Inflation Reduction Act." National Bureau of Economic Research Working Paper 31267.

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. 2012. "Peer Effects in the Diffusion of Solar Photovoltaic Panels." *Marketing Science*, 31(6): 900–912.

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.

CBO. 2022. "Summary Estimated Budgetary Effects of Public Law 117-169, to Provide for Reconciliation Pursuant to Title II of S. Con. Res. 14 (September 7, 2022)." Congressional Budget Office and Joint Committee on Taxation.

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*.

CPUC. 2009. "California Solar Initiative Annual Program Assessment." California Public Utilities Commission.

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.

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.

Ifrach, 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.

Ito, Koichiro. 2014. "Do Consumers Respond to Marginal or Average Price? Evidence from Nonlinear Electricity Pricing." *American Economic Review*, 104(2): 537–563.

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." *NBER WP*.

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.

Kline, Patrick, and Enrico Moretti. 2014. "Local Economic Development, Agglomeration Economies, and the Big Push: 100 Years of Evidence from the Tennessee Valley Authority." *The Quarterly Journal of Economics*, 129(1): 275–331.

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.

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.

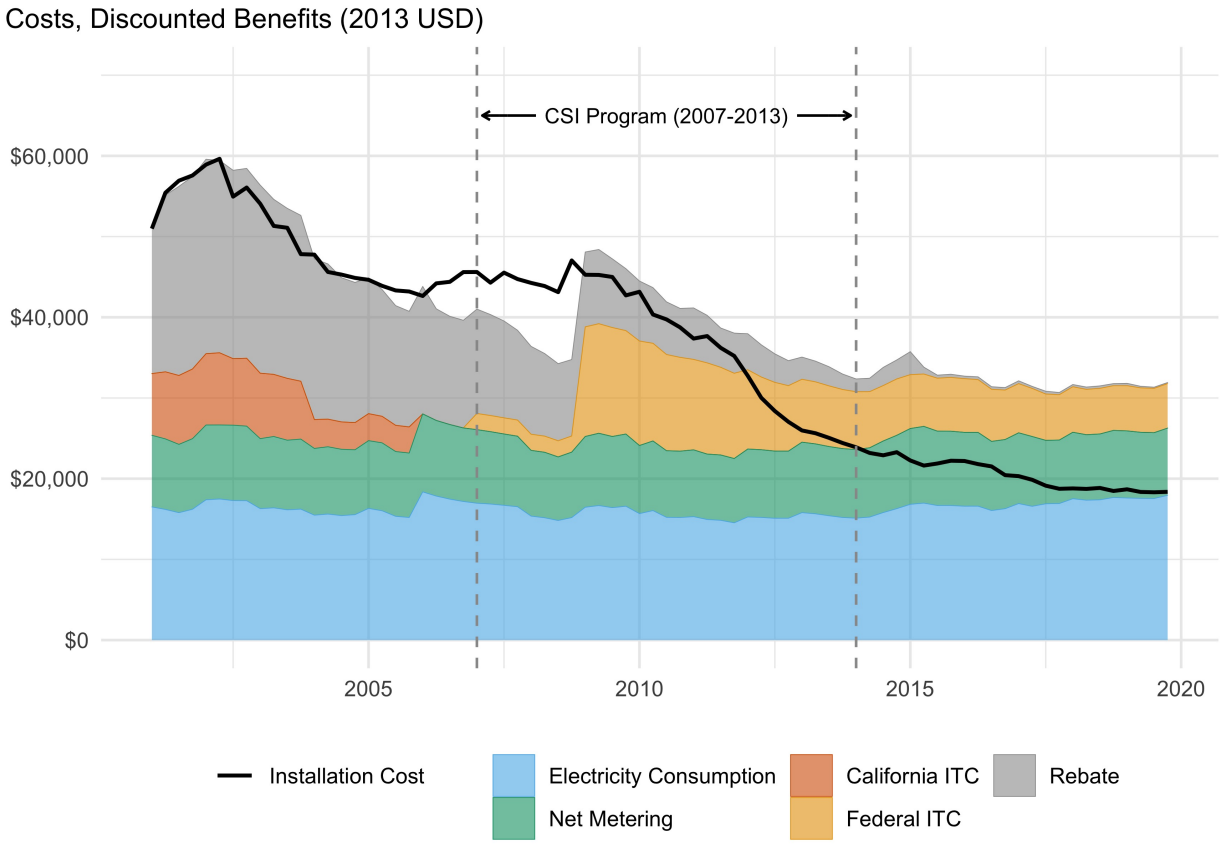
van Benthem, Arthur, Kenneth Gillingham, and James Sweeney. 2008. "Learning-by-Doing and the Optimal Solar Policy in California." *The Energy Journal*, 29(3): 131–151.

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.

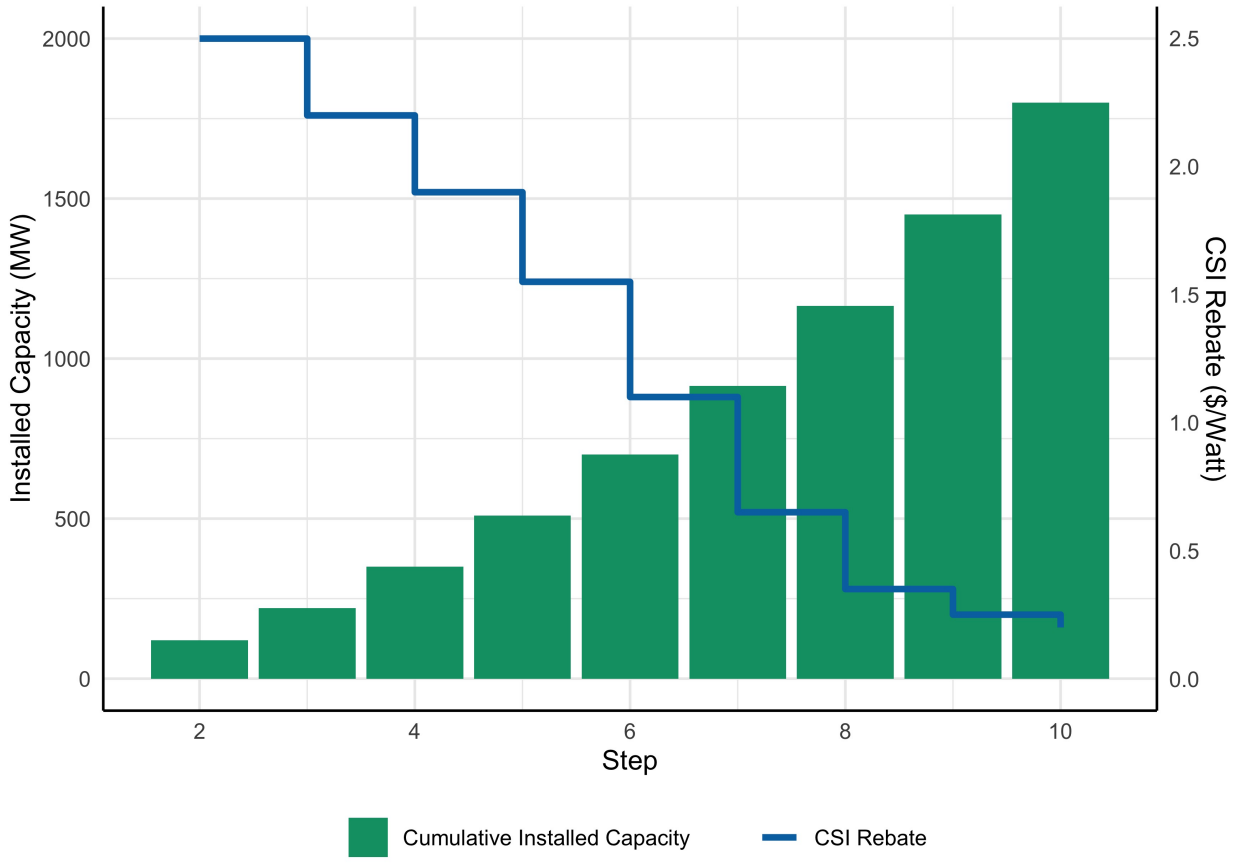
A Supplemental Figures and Tables

Figure A1. Costs and Discounted Benefits of a 5 kW PV System in California, 2000-2020



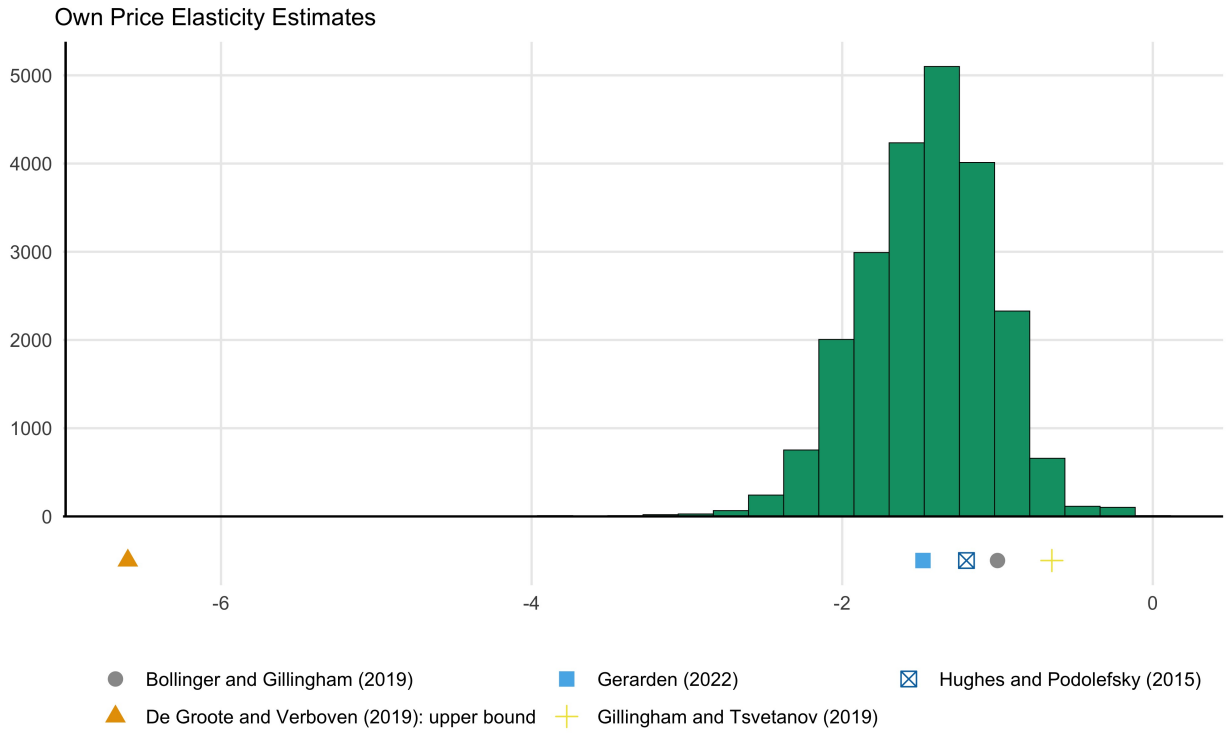
Notes: This figure shows the evolution of the upfront installation cost and present discounted benefits of a representative 5 kilowatt (kW) residential PV system in California from 2000 to 2020. Data on average installation costs and rebates in California come from the Lawrence Berkeley National Lab's "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 A2. CSI Rebate Rate Structure



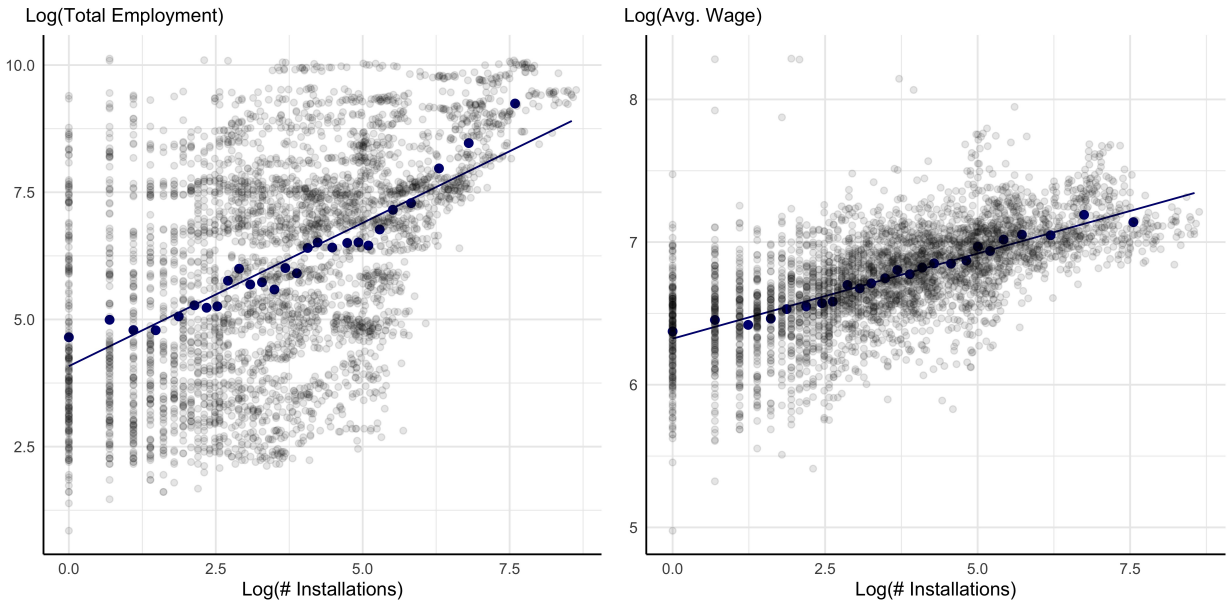
Notes: This figure shows the rebate llevels under the California Solar Initiative (CSI) as a function of cumulative-installed capacity. This figure is inspired by a similar figure that appears in Pless and Van Benthem (2019).

Figure A3. Estimated Own Price Elasticities



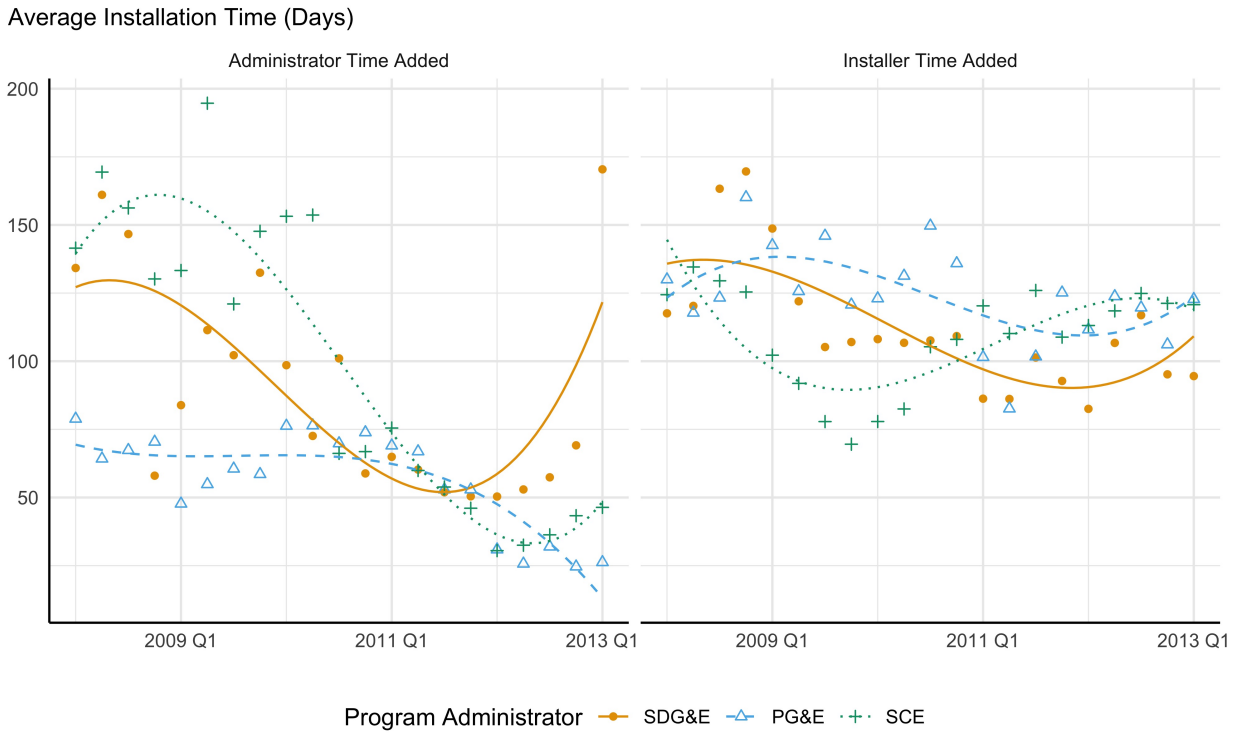
Notes: This figure shows the distribution of estimates of the own price elasticities of demand for all firm-county-quarter observations in the main estimation sample estimated using the random coefficients nested logit model reported in Column (3) of Table 2. Estimates of the adoption decision elasticity from the literature are reported below the horizontal axis (Bollinger and Gillingham, 2019; De Groote and Verboven, 2019; Gerarden, 2022; Gillingham and Tsvetanov, 2019; Hughes and Podolefsky, 2015).

Figure A4. Relationship between Total PV Installation-related Employment, Wages and Installations



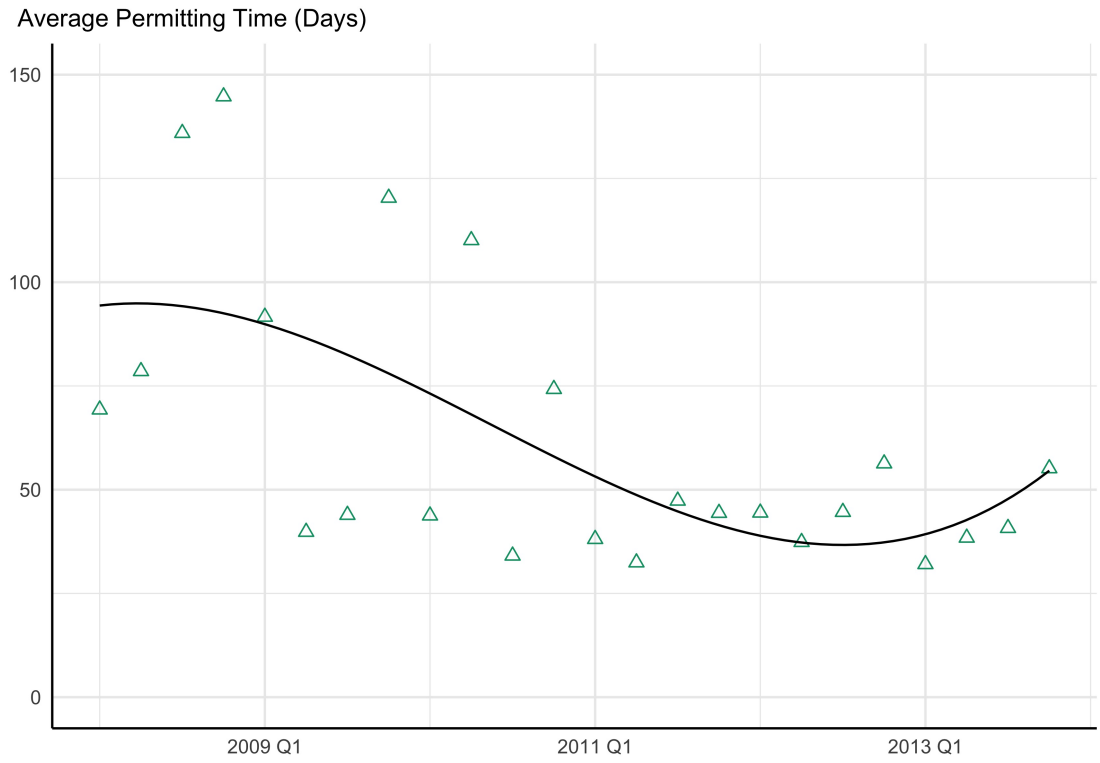
Notes: This figure shows the relationship between total quarterly PV installation-related employment and the average PV installation-related wage in a county and the quarterly total number of residential PV installations within that county. Data on employment levels and wages are from the US Census Bureau's 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 A5. Average CSI Project Completion Time Added



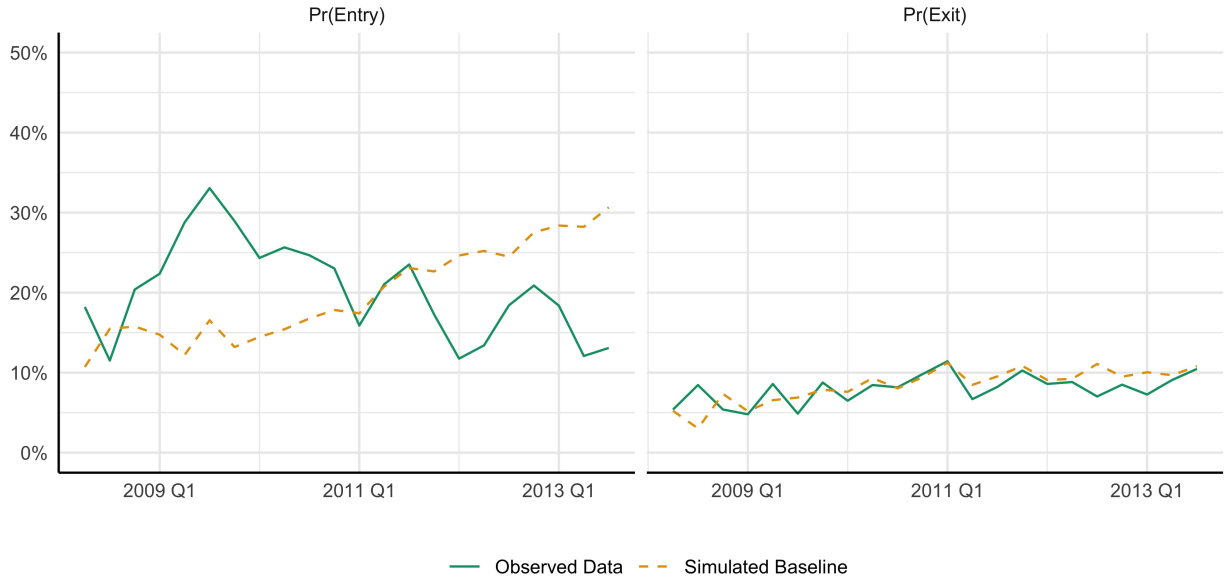
Notes: This figure shows the quarterly average observed time added to residential solar PV installations that apply for California Solar Initiative (CSI) rebates separately for rebate program administrators and installers. Quarterly averages are reported separately by each of the three main investor-owned 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 A6. Average Permitting Time for PV Projects, San Diego County



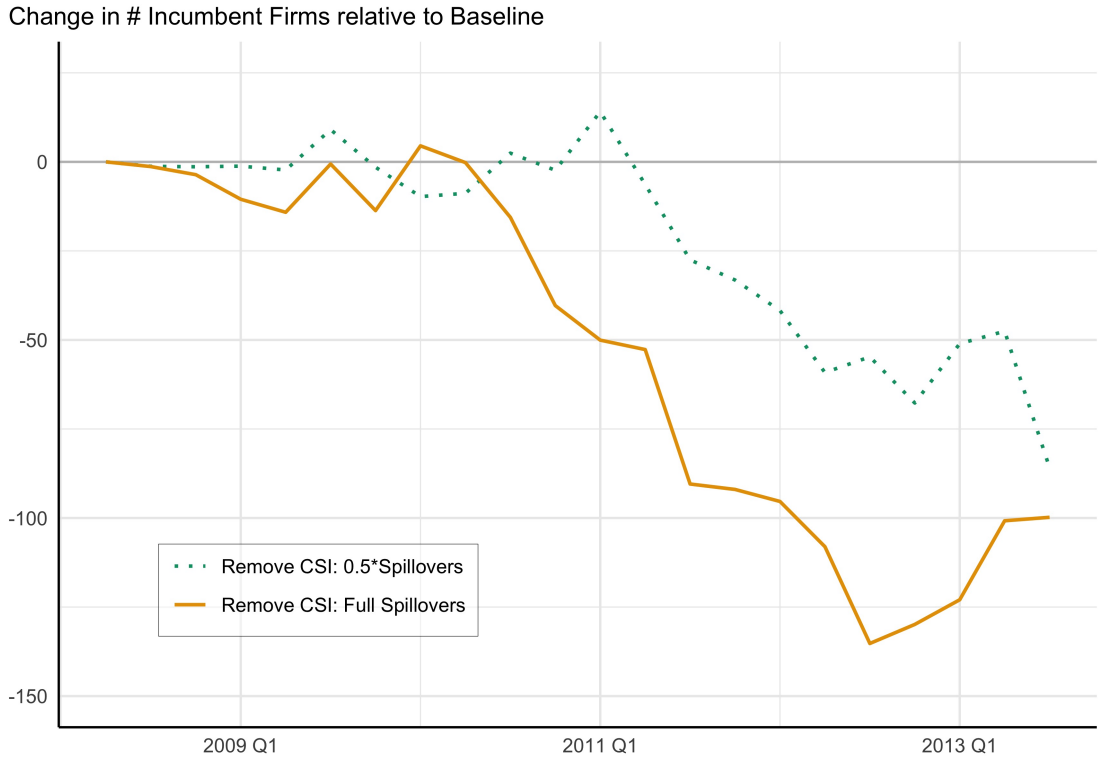
Notes: This figure shows the quarterly average time-to-completion for solar photovoltaic (PV)-related permits in San Diego County, California. Permit-level data on historical developments permits are available from https://data.sandiego.gov/datasets/development-permits-set1/ (last accessed August 10, 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.

Figure A7. Comparison of Simulated and Actual Entry and Exit Probabilities



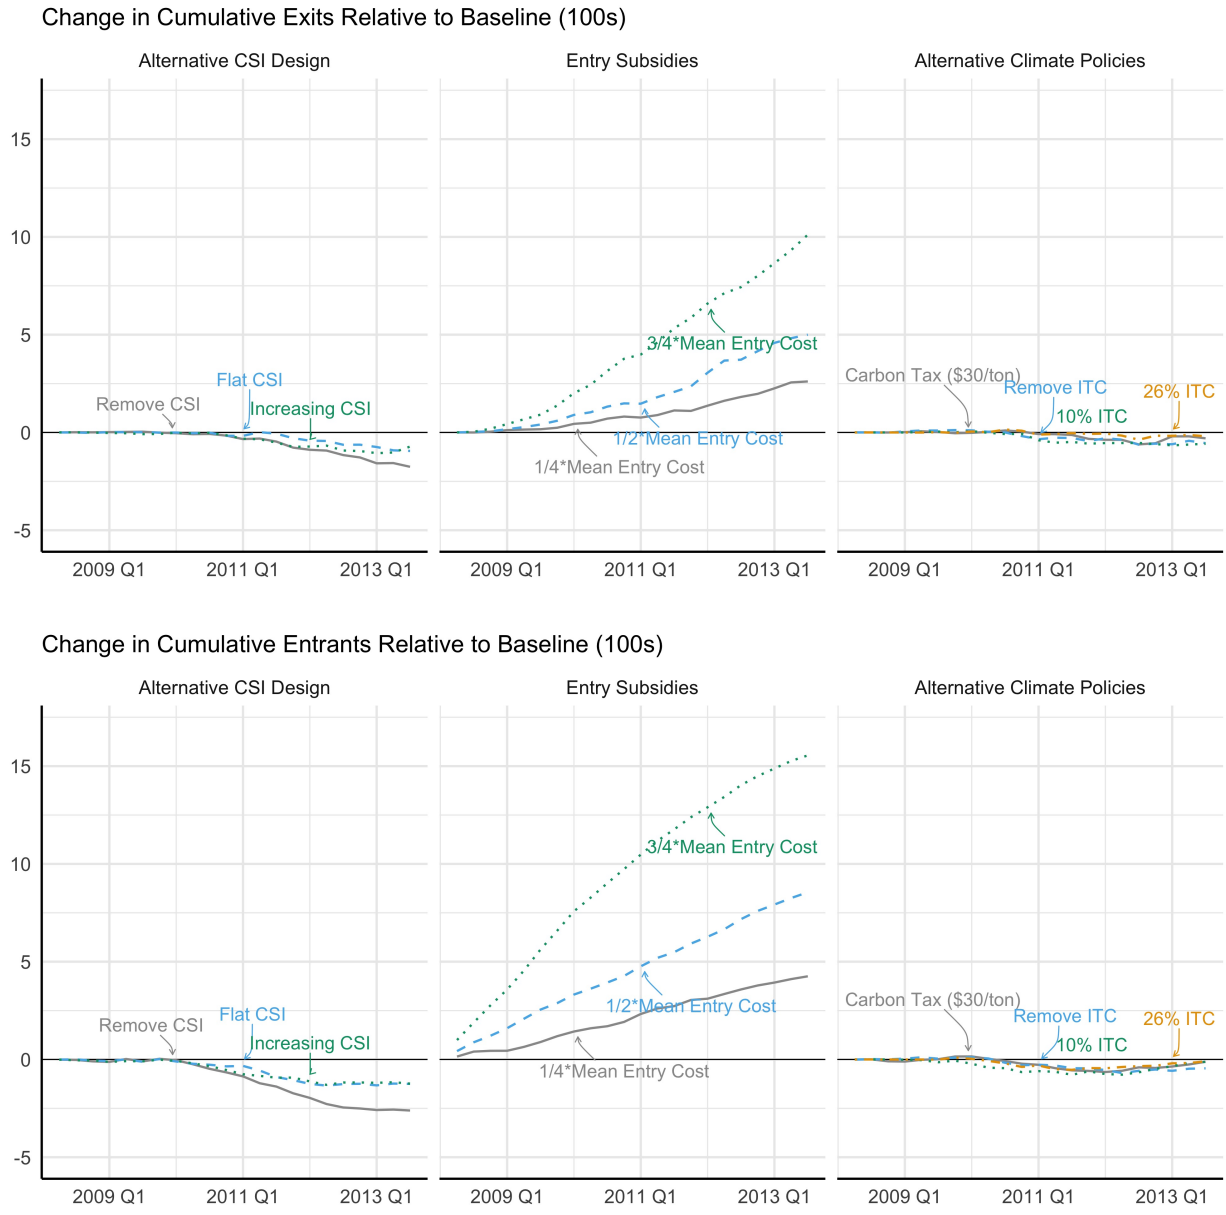
Notes: This figure compares the model-predicted, statewide entry (left panel) and exit (right panel) probabilities with the statewide entry and exit probabilities observed in the data. The baseline subsidy policies correspond to those outlined in Section 2.2 and observed in practice: full CSI subsidies, 30% federal investment tax credit (ITC), and net-metering policy. The simulated probabilities plotted in the figure represent averages across 60 distinct, forward simulated industry paths under the baseline subsidy policies, each of which is generated from the model as described in Section 7.1 and Appendix F. To generate the probabilities of entry plotted in the left panel, I assume the number of potential entrants for each county is equal to twice the maximum number of observed quarterly entrants for that county.

Figure A8. Change in Active Firms from Removing CSI Subsidies with Different Spillover Sizes



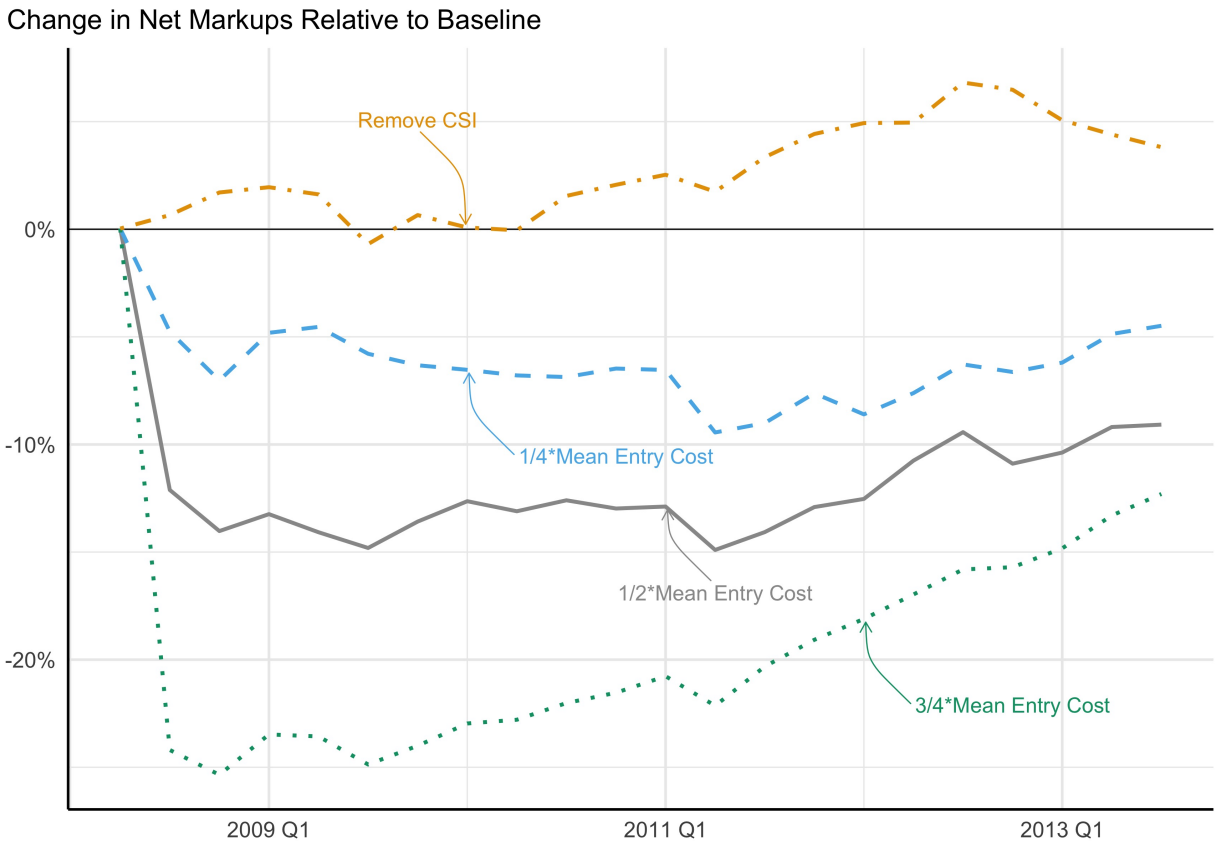
Notes: This figure shows the model-predicted change in the number of active firms relative to the baseline scenario of existing consumer subsidy policies under a counterfactual of no CSI subsidies for two levels of learning spillovers: the full spillovers estimated in the preferred specification in Table 3 and 1/2 these full spillovers. The baseline subsidy policies correspond to those outlined in Section 2.2: full CSI subsidies, 30% federal investment tax credit (ITC), and net-metering policy. I simulate the baseline of full subsidies and the counterfactual of no CSI subsidies separately for each spillover size. Each counterfactual outcome plotted in the figure represents an average across 60 distinct, forward-simulated industry paths under the given counterfactual.

Figure A9. Cumulative Change in Exits and Entries under Counterfactual Policies



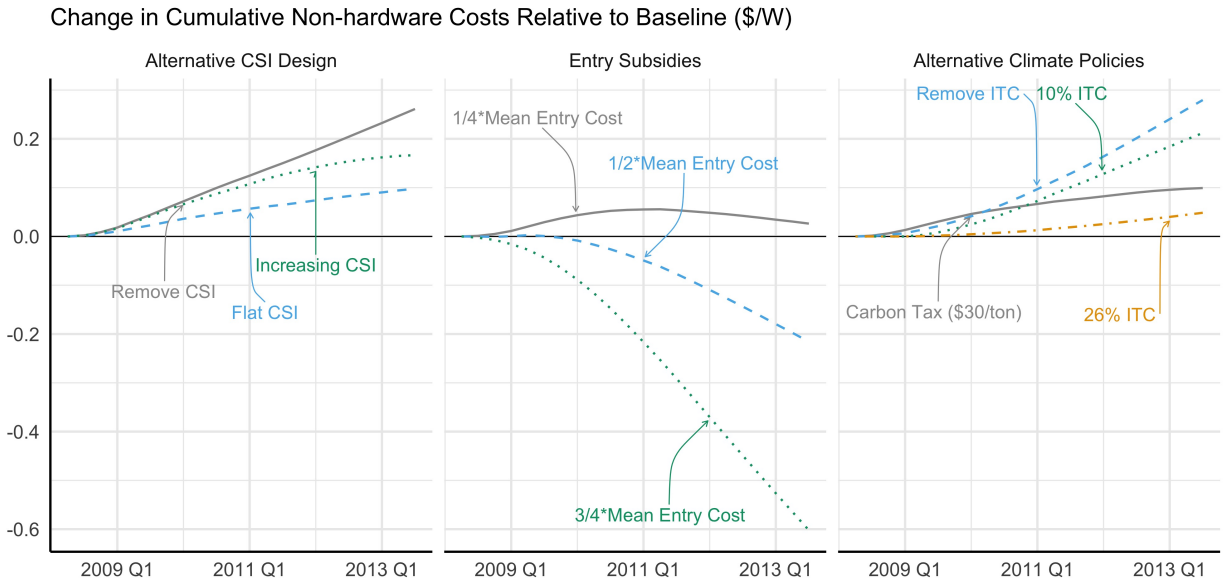
Notes: This figure shows the model-predicted change in the cumulative number of exiting firms (top) and entering firms (bottom) relative to the baseline scenario of existing consumer subsidy policies under three sets of counterfactuals: alternative California Solar Initiative (CSI) subsidy designs (left), varying levels of entry subsidies (center), and alternative climate policies (right). The baseline subsidy policies correspond to those outlined in Section 2.2: full CSI subsidies, 30% federal investment tax credit (ITC), and net-metering policy. 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 and alter the design of the federal ITC. Each counterfactual outcome plotted in the figure represents an average across 60 distinct, forward-simulated industry paths under the given counterfactual.

Figure A10. Change in Net Markups under Select Counterfactual Policies



Notes: This figure shows the model-predicted average net markups (i.e., average sum of markups and markdowns) relative to the baseline scenario of existing consumer subsidy policies under four counterfactuals. The four counterfactuals included in this figure each remove the CSI and either do not replace the CSI's consumer subsidies or replace them with entry subsidies of various sizes. The baseline subsidy policies correspond to those outlined in Section 2.2: full CSI subsidies, 30% federal investment tax credit (ITC), and net-metering policy. Each counterfactual outcome plotted in the figure represents an average across 60 distinct, forward-simulated industry paths under the given counterfactual.

Figure A11. Cumulative Change Non-hardware Costs under Counterfactual Policies



Notes: This figure shows the model-predicted cumulative change in non-hardware costs relative to the baseline scenario of existing consumer subsidy policies under three sets of counterfactuals: alternative California Solar Initiative (CSI) subsidy designs (left), varying levels of entry subsidies (center), and alternative climate policies (right). The baseline subsidy policies correspond to those outlined in Section 2.2: full CSI subsidies, 30% federal investment tax credit (ITC), and net-metering policy. 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 and alter the design of the federal ITC. Each counterfactual outcome plotted in the figure represents an average across 60 distinct, forward-simulated industry paths under the given counterfactual.

Table A1. Estimated Transition Processes for Aggregate State Variables

	Demand (Installations)		Avg. Price ($/W)		Inclusive Value	
	(1)	(2)	(3)	(4)	(5)	(6)
Intercept	2.853		0.4100		0.0026	
	(1.932)		(0.1057)		(0.0008)	
$(\text{Demand})_{t-1}$	0.9724	0.7191				
	(0.0380)	(0.1205)				
$(\text{Avg. Price})_{t-1}$			0.9126	0.9077		
			(0.0156)	(0.0165)		
$(\text{Inclusive Value})_{t-1}$					0.8393	0.6380
					(0.0655)	(0.1156)
County FE		Yes		Yes		Yes
Observations	885	885	885	885	885	885
R^2	0.84	0.86	0.81	0.81	0.62	0.66
Within R^2		0.43		0.80		0.34

This table reports estimates of the first-order autoregressive (AR(1)) transition processes for three county-quarter aggregate state variables, demand (number of installations), 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 A2. Estimated Transition Processes for Firm-specific State Variables

	Own Quality		Hardware Cost ($/W)		Price ($/W)	
	(1)	(2)	(3)	(4)	(5)	(6)
Intercept	-1.164		0.3997		0.9885	
	(0.1895)		(0.0295)		(0.0806)	
$(\text{Own Quality})_{t-1}$	0.8438	0.5275				
	(0.0253)	(0.0780)				
$(\text{Hardware Cost})_{t-1}$			0.8715	0.8666		
			(0.0069)	(0.0069)		
$(\text{Price})_{t-1}$					0.8256	0.8127
					(0.0125)	(0.0118)
County FE		Yes		Yes		Yes
Observations	5,862	5,862	5,862	5,862	5,862	5,862
R^2	0.70	0.75	0.76	0.77	0.69	0.69
Within R^2		0.24		0.76		0.67

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 A3. Comparison of Estimators for Production and Exit Cost Parameter

	Parameter	Baseline Moment (1)	Alternative Moments (2)
Production Cost Parameters			
Base Cost	c_0	2.145	2.110
		(0.046)	(0.334)
Learning Exponent	γ	-0.733	-0.734
		(0.083)	(0.138)
Productivity Serial Correlation	ρ	0.838	
		(0.099)	
Common Time Trend	t		0.147
			(0.039)
Effective Experience			
Industry Experience: Total	θ_1^E	0.817	0.505
		(0.059)	(0.119)
Exit Parameter			
Mean Scrap Value	σ_ϕ	2.488	2.208
		(1.148)	(0.373)
N		11,581	18,272
Spence Coefficient $(1-2^\gamma)$		0.399	0.399
Production Cost Moment(s)		$\mathbb{E}[\nu_{jmt}] = 0$	$\mathbb{E}[Z'_{jmt} \cdot \kappa_{jmt}] = 0$

Notes: Estimation in column (1) follows the procedure outlined in Section 5.2. Estimation of the production cost parameters in column (2) relies on moments from the unobservable productivity shock, κ_{jmt}, interacted with a set of instruments, Z_{jmt}, that satisfy the relevance and exogeneity conditions: i.e., $\mathbb{E}[Z'_{jmt} \cdot \kappa_{jmt}] = 0$. As a result of not relying on the innovation in firms' serially-correlated unobservable productivity shock, ν_{jmt}, data do not need to be first-differenced in implementing the alternative estimation approach in Column (2). Instruments in Z_{jmt} include: realized consumer rebates; observed firm-level deviations from average utility interconnection processing times; current and lagged cost shifters (aluminum and polysilicon commodity prices; county-quarter specific wage rates); and current and lagged demand shifters (retail electricity prices, observed solar insolation). I normalize experience variables by the industry total experience level in the first quarter of the sample (Q1 2008). The base cost parameter is therefore the marginal installation cost when a firm's effective experience equals the industry total experience in Q1 2008. All effective experience parameters can be interpreted as marginal experience contributions relative to a firm's own experience. 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. Bootstrapped standard errors clustered by county using 200 replications are reported in parentheses.

Table A4. Estimated Production and Exit Cost Parameter Estimates for Different Discount Factors

Annual Discount Factor:	Parameter	$\beta = 0.8$ (1)	$\beta = 0.875$ (2)	$\beta = 0.9$ (3)
Production Cost Parameters				
Base Cost	c_0	2.004	2.145	2.303
Learning Exponent	γ	-0.979	-0.733	-0.426
Productivity Serial Correlation	ρ	0.680	0.838	1.051
Effective Experience				
Industry Experience: Total	θ_1^E	0.993	0.817	0.560
Exit Parameter				
Mean Scrap Value	σ_ϕ	1.383	2.488	4.429
N		11581	11581	11581
Spence Coefficient $(1 - 2^\gamma)$		0.256	0.399	0.493

Notes: Estimation follows the procedure outlined in Section 5.2. Different annual discount factors are used in each column. I normalize experience variables by the industry total experience level in the first quarter of the sample (Q1 2008). The base cost parameter is therefore the marginal installation cost when a firm's effective experience equals the industry total experience in Q1 2008. All effective experience parameters can be interpreted as marginal experience contributions relative to a firm's own experience. 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.

Table A5. Estimated Production and Exit Cost Parameter Estimates with Scale as a State Variable

	State Variables:	Baseline	+Firm Size
	Parameter	(1)	(2)
Production Cost Parameters			
Base Cost	c_0	2.145	2.075
		(0.046)	(0.046)
Learning Exponent	γ	-0.733	-0.846
		(0.083)	(0.176)
Productivity Serial Correlation	ρ	0.838	0.842
		(0.099)	(0.621)
Effective Experience			
Industry Experience: Total	θ_1^E	0.817	0.878
		(0.059)	(0.084)
Exit Parameter			
Mean Scrap Value	σ_ϕ	2.488	3.174
		(1.148)	(0.848)
N		11,581	18,272
Spence Coefficient $(1 - 2^\gamma)$		0.399	0.444

Notes: Estimation follows the procedure outlined in Section 5.2. Column (1) corresponds to the baseline preferred specification of the model in Column (2) of Table 3. 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 quarter of the sample (Q1 2008). The base cost parameter is therefore the marginal installation cost when a firm's effective experience equals the industry total experience in Q1 2008. All effective experience parameters can be interpreted as marginal experience contributions relative to a firm's own experience. 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.

B Data Appendix

B.1 Additional Information on Data Sources

I provide information on each of the data sources used in my empirical analysis below.

- *Lawrence Berkeley National Lab's (LBNL) Tracking the Sun Database:* provides system-level data on PV systems annually from state agencies and utilities that administer PV incentive programs, renewable energy credit registration systems, or grid interconnection processes. The public use database includes information on the date of installation, system size, total installed price, total pre-tax rebate value, customer type, zip code, mounting type, and installer name, as well as various technical details about installed hardware, including the energy conversion efficiency (i.e., how much incoming solar radiation a panel converts into electrical power), make, and manufacturer of installed PV modules. These data also include a number of unique identifier fields which are maintained from the original databases sourced from state agencies and utilities, which allows me to link individual systems to other data sources that I outline below. LBNL processes the source data prior to publishing the public use data, including removing systems with missing size or installation date fields; standardizing installer, module, and inverter manufacturer names; and integrating publicly available equipment specification data with system-level data. The full sample includes data on over 2.5 million PV systems installed from 2000 to 2021, covering both residential and non-residential systems. Barbose et al. (2022) estimate that the database covers approximately 77% of the total estimated US market for PV systems over 2000-2021. I use these data to generate the main datasets used to estimate the demand and supply models. Additional information on these data is available at `https://emp.lbl.gov/tracking-the-sun` (last accessed 8/22/2023).

- *California Public Utilities Commission (CPUC) California Solar Initiative (CSI) Hardware Cost Data:* provides system-level data on PV system hardware cost for a subset of California residential PV systems that applied for the CSI rebate. Installation firms submitting CSI rebate paperwork to the relevant administering investor owned utility (IOU) on behalf of installing households were not uniformly required to include information on the total cost of all system hardware installed; however, these data were voluntarily provided by installers for a large number of systems, particularly during the period 2008-2013, which corresponds to the main period of the CSI. I acquire these system-level hardware cost data from the CPUC via contacts at LBNL for over 60,000 California systems and I am able to successfully match non-zero hardware costs to over

79% of systems in the broader, processed LBNL Tracking the Sun data for this period using unique CPUC-generated identifiers contained in both datasets. Please contact me for additional information on these data: jbradt@g.harvard.edu.

- *CPUC CSI Working Dataset:* provides system-level data on the universe of CSI rebate applications from the three main IOUs, including all non-accepted applications. These data are one of the underlying sources for LBNL's Tracking the Sun public use data; however, these data contain more detailed information on the CSI rebate for which applications are eligible, which I use to determine the precise dates of CSI rebate changes for each of the three main IOUs. Information on these data is available at https://www.californiadgstats.ca.gov/downloads/ (last accessed 8/22/2023).

- *Google Project Sunroof:* provides estimates of the share of buildings in a given county for which adopting solar would lead to a positive net present value. These data combine satellite imagery, three-dimensional modeling, and shade calculations at a property-level with weather data from NREL, retail electricity rates from Clean Power Research, and solar cost data from Aurora Solar Software to estimate for a range of different assumed system lifespans and discount rates the share of households for which solar has a positive net present value. I use these data in combination with data on the number of owner-occupied housing units in a county to estimate the size of the market for residential solar PV installations in California counties. Given that the returns from solar installation can vary significantly from property to property based on shading and roof orientation, approximating the universe of households that are suitable for solar is far more appealing than assuming that all households are potential adopters. These data therefore should provide more reasonable estimates of installation firms' market shares for use in estimating the demand model. Additional information on these data is available at https://sunroof.withgoogle.com/ (last accessed 8/22/2023).

- *World Bank Group's Global Solar Atlas:* provides estimates of the long-term annual average photovoltaic power potential of a 1 kilowatt (kW) capacity PV system in raster data format at a 250-meter resolution. I use these data in combination with administrative boundary data for California Counties from the US Census Bureau to estimate the average annual power production potential for a residential PV system installed in each county. This helps me calculate the approximate consumption and net energy metering (NEM) benefits for each system in my system-level data (see additional details in Section B.3 below). I also use these data to calculate the changes in PV electrical power output for each counterfactual scenario, which I then use to

determine avoided climate and local air pollution damages. Information on these data is available at https://globalsolaratlas.info/download (last accessed 8/22/2023).

- *US Census Bureau American Community Survey (ACS) Public Use Microdata Sample (PUMS):* provides household-level demographic data for a sample of households in California each year in my main estimation period, 2008 to 2013. ACS PUMS provide a 1-in-100 national random sample of the population and identify households geographic location down to the public use microdata area, which contains at least 100,000 persons. When the number of residents in a county exceeds 100,000, PUMS also identify a household's county of residence. I obtain PUMS data for the state of California for each year from 2008 to 2013 and randomly draw 200 households for each county-quarter in that year in order to allow for preference heterogeneity by income in my demand model. When counties are sufficiently large to be identified in the PUMS, I draw 200 households from the sample identified as residing in that county. In cases where counties are not sufficiently large to be identified in the PUMS, I draw 200 households from the entire state. Additional information on these data is available at https://www.census.gov/programs-surveys/acs/microdata.html (last accessed 8/22/2023).

- *US Census Bureau ACS 5-year Estimates and 2000 Decennial Census:* provides annual estimates of the number of owner-occupied housing units for each California county for 2009-2013. I use these estimates in combination with data from Google's Project Sunroof to calculate the size of the potential market of residential PV adopters in each California county for the period 2008 to 2013, which I use in turn to estimate each installation firm's market share in estimating the demand model. Unfortunately, ACS estimates are unavailable for the year 2008, so I use the count of owner-occupied housing units in each California county from the 2000 Decennial Census in combination with the 5-year ACS estimates for 2009-2013 to impute values for the year 2008. Additional information on these data is available at https://www.census.gov/data/developers/data-sets/acs-5year.html (last accessed 8/22/2023).

- *US Energy Information Administration (EIA) Form 861:* provides annual data on every electric utility in the US that allows me to calculate the average residential electricity rate for each of the three main IOUs in California. Following the approach of Borenstein and Bushnell (2022), I take annual data on total retail revenues and kilowatt-hour sales for residential customers for each of the three California IOUs for 2008-2013 to calculate an IOU-specific measure of average retail electricity rates. I then use these estimates to calculate the approximate consumption and net energy metering (NEM) benefits for each system in my system-level data (see additional details

in Section B.3 below). Additional information on these data is available at https://www.eia.gov/electricity/data/eia861/ (last accessed 8/23/2023).

- *US Bureau of Labor Statistics (BLS) Quarterly Census of Employment and Wages:* provides quarterly data on average wages for private sector electricians and roofers for all California counties for the period 2008-2013. I use quarterly county-level employment levels to construct a weighted average of observed electrician and roofer wages for each California county during this period. These data are used as instruments in the demand model. Additional information on these data is available at https://www.bls.gov/cew/ (last accessed 8/23/2023).

B.2 Sample Restrictions

I now discuss the full set of sample restrictions that I apply to the main system-level data obtained from LBNL's Tracking the Sun public data file. These data are the primary source for constructing the estimation data for the demand and supply model. As discussed in Section 3.1, I restrict the LBNL data to only include systems installed in California and further remove all non-residential systems, all ground-mounted systems, any residential systems with capacity exceeding 20 kW, and any residential systems for which I do not observe the installed price or rebate. Removing residential systems exceeding 20 kW in capacity only eliminates a small subset of outlier residential systems that are likely installed at multi-family condominium buildings. I also exclude self-installed systems given my focus on the residential installation market and any systems owned by a third party. I further exclude those installations with a price per watt that is less than the first or greater than the ninety-ninth percentile of the prices per watt of all installations in the sample to remove outliers.

B.3 Constructing Data for Demand Estimation

Estimating the demand model requires me to aggregate the processed system-level data to the firm-county-quarter-level. Given the importance of the system-level hardware cost data in estimating the supply model, I subset the system-level data to those installations for which I observe hardware cost data when constructing aggregate data for demand estimation. While this process of aggregating the system-level data is relatively straightforward, it is worth describing steps taken for certain key data fields. In particular, it is worth outlining the way in which I estimate county-quarter-level market shares, average installed price per watt, the per watt value of public PV adoption incentives, and a set of firm-level attributes. Normalizing prices and rebate amounts by installed system capacity is important in aggregating data for demand estimation as it ensures I am aggregating comparable goods:

averaging total installed prices across systems of vastly different sizes could lead to misleading conclusions about firms' prices. While this eliminates possible scale economies in firms' installed prices, the relatively commoditized nature of the underlying PV module technology makes this a reasonable approach to handling installations of different sizes.

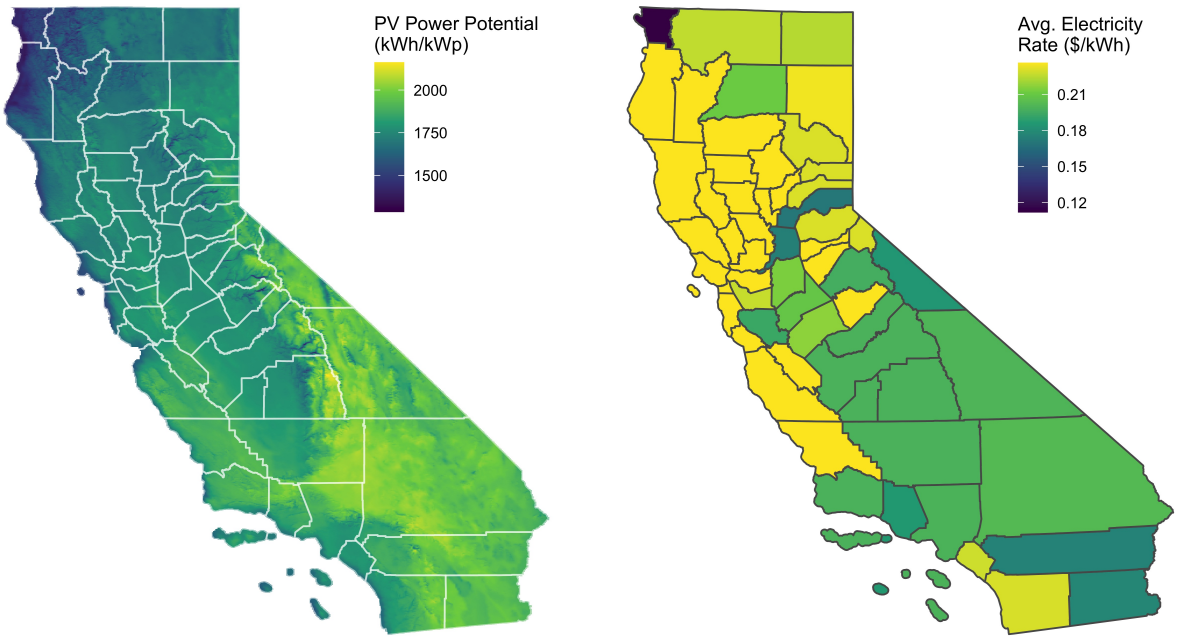
Calculating market shares requires me to first estimate the pool of potential PV adopters. As discussed in Section B.1 above, given that the returns from solar installation can vary significantly on a property-level based on shading and roof orientation, it is necessary to approximate the universe of households that are suitable for solar as assuming that all households are potential adopters would vastly underestimate firms' market shares. I therefore combine data from Google Project Sunroof on the share of buildings in each California county for which adopting solar has a positive net present value with annual estimates of the number of owner-occupied housing units in each county from the US Census Bureau. I focus on owner-occupied housing to abstract away from standard principal-agent problems inherent to energy-efficient technology adoption decisions in renter-occupied housing. Importantly, in estimating the demand model I account for the durable nature of PV system installations by removing adopting households from my estimate of the potential market size. To do so, I take advantage of the full history of processed system-level data for California going back to 2000 that is available in the LBNL Tracking the Sun public use data file. Having constructed reasonable estimates of the potential market of residential PV adopters, I can then use observed installation counts to calculate firms' market shares.

Turning to the price and rebate data, calculating county-quarter averages for all installers from the existing fields in the processed system-level data is relatively straightforward. One important step that I take in calculating these averages is to construct weights for each installation based on that installation's share of the installing firm's total installed capacity within a county in a given quarter. I then calculate weighted averages at the county-quarter-firm-level that should be representative of the prices offered and rebates available to a firms' customers across a county in a given period.

While I directly observe upfront CSI rebates in the system-level data, I do not observe the two other primary forms of public PV adoption incentives, the federal investment tax credit (ITC) and net energy metering (NEM). Fortunately, calculating the per watt ITC benefit for all installations in my data is relatively straightforward, assuming that all households fully capitalize the tax credit. For household i choosing installer j in market m in period t, the per watt ITC benefit, ITC_{ijmt} can be written as:

$$ITC_{ijmt} = \begin{cases} \min\left\{1000/q_{ijmt}, 0.3(p_{ijmt} - r_{ijmt})\right\} & \text{if } t < 2009 \text{ Q1} \\ 0.3(p_{ijmt} - r_{ijmt}) & \text{if } t \geq 2009 \text{ Q1} \end{cases}$$

Figure B1. 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).

where p_{ijmt} is the total installed price paid by household i, r_{ijmt} is the per watt upfront rebate received by household i, and q_{ijmt} is household i's installed system capacity.

Constructing system-level estimates of the upfront benefit from NEM policy is slightly more complicated. Doing so requires me to make some strong assumptions about household energy consumption, system lifespan, and household discount rates. I rely on estimates of excess PV output as a share of system output for California from Darghouth et al. (2011) to calculate the quantity of electricity sold back to the grid by each system in my data. I then calculate expected annual system output using estimates of the long-term annual average photovoltaic power potential of a 1 kW capacity PV system from the World Bank Group's Global Solar Atlas and system-level data on overall system capacity (see Figure B1). Combining these system-level estimates of expected annual output with Darghouth et al. (2011)'s findings on excess generation and data on average retail electricity rates from EIA Form 861 (see Figure B1), I can calculate the net present value of expected NEM benefits assuming a system lifespan of 25 years and an annual household discount rate of 12.5%, which is approximately equal to that estimated by De Groote and Verboven (2019). Having

obtained system-level estimates of the ITC and NEM benefits per watt, I then calculate weighted averages for each installer-county-quarter observation as described above. I then aggregate all public incentives per watt in estimating the demand model.

Finally, I convert all dollar-denominated fields to 2013 real dollars using the consumer price index (CPI) for all urban consumers.[37] I construct a set of attributes that vary over time for inclusion in the demand model. These include counts of the number of distinct module types offered by an installer as well as an indicator of whether the installer's average installed module efficiency exceeds the 75th percentile in the sample.

B.4 Constructing Data for Supply Model

I now describe the process of constructing the data used in estimating the supply model. These data start from the same installer-county-quarter-level dataset used to estimate the demand model, with several additional fields. In particular, I add to this dataset the following fields: average hardware cost per watt as well as a set of experience and quantity variables. Average hardware cost is calculated in the same manner as the average installed price and rebate fields above, namely as a firm-county-quarter weighted average. Hardware costs are similarly deflated using the CPI.

I compute a set of experience variables using the full history of processed system-level data for California going back to 2000 that is available in the LBNL Tracking the Sun public use data file. Experience in my model can be defined in several ways; however the primary way in which I define experience is in terms of cumulative PV system capacity installed. I therefore use data going back to 2000 to calculate cumulative installed capacity for each installer-county-quarter combination. I also calculate the following measures of rivals experience: cumulative installed capacity for rivals in the same county, rivals in other counties, rivals installing PV modules from the same manufacturer, and rivals installed PV modules from other manufacturers. Importantly, I normalize all experience fields, which are measured in watts, by the total, statewide cumulative installed watts at the beginning of my main estimation period, the first quarter of 2008. This helps ensure numerical stability in estimating the supply model given that firm-level cumulative installed capacity in a given county and, for example, the cumulative installed capacity of rival firms in other counties can differ by several orders of magnitude. For each of these experience fields, I also calculate the corresponding quantity of installed capacity in each period. I calculate these quantities in absolute watts as well as normalized by the statewide cumulative installed watts in 2008 Q1 as both the raw and normalized quantities are also used in estimation.

[37] US 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.

C 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.

C.1 Model Setup

Similar to the static model in Section 4.1, incumbent firms in each period and market face a set of idiosyncratic consumers, $i \in \{1, \ldots, N_{mt}\}$, 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 ($j \in \{1, \ldots, J_{mt}\}$) 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:

$$\cup_{j \in \mathcal{J}_{mt}} \begin{bmatrix} p_{jmt} & r_{jmt} & \xi_{jmt} \end{bmatrix}$$

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:

$$u_{ijmt}(s_{mt}) = \underbrace{\alpha(p_{jmt} - r_{jmt}) + \xi_{jmt} + \bar{\xi}_j + \bar{\xi}_t}_{\equiv \delta_{jmt}} + \bar{\varepsilon}_{ijmt} \tag{C1}$$

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:

$$u_{i0mt}(s_{mt}) = \underbrace{u_{0mt} + \beta \mathbb{E}\Big[V(s_{mt+1}|s_{mt})\Big]}_{\equiv \delta_{0mt}} + \bar{\varepsilon}_{i0mt} \tag{C2}$$

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 4.1. Following Berry (1994), it is possible to invert predicted market shares as follows:

$$\log\left(ms_{jmt}(s_{mt})\right) - \log\left(ms_{0mt}(s_{mt})\right) = \delta_{jmt} - \delta_{0mt} + \eta \log\left(\bar{ms}_{jmt|g}(s_{mt})\right) \tag{C3}$$

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 $\bar{ms}_{jmt|g}$ is firm j's conditional within-group share in market m in period t (i.e., the firm's inside share).

C.2 Ex Ante Value Function

Taking (C3) to the data requires a closed form for households' ex ante value function. This requires an approach to handling the expectation operator in (C2), 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:

$$e_{mt} \equiv V(s_{mt+1}) - \mathbb{E}\Big[V(s_{mt+1}|s_{mt})\Big] \tag{C4}$$

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:

$$\begin{aligned}\delta_{0mt} &= \mathbb{E}\Big[V(s_{mt+1}|s_{mt})\Big] \\ &= \beta\Bigg(\delta_{j'mt+1}(s_{mt+1}) - \log\big(ms_{j'mt+1}(s_{mt+1})\big) \\ &\quad - \eta\Big(\log\big(ms_{gmt+1}(s_{mt+1})\big) - \log\big(ms_{j'mt+1}(s_{mt+1})\big)\Big) - e_{mt}\Bigg)\end{aligned} \tag{C5}$$

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

C.3 Estimating Equation

Combining (C1), (C3), and (C5) and rearranging terms gives the following equation:

$$\begin{aligned}&\log\left(\frac{ms_{jmt}}{ms_{0mt}}\right) - \beta\log(ms_{j'mt+1}) \\ &\quad = \alpha\Big((p_{jmt} - \beta p_{j'mt+1}) - (r_{jmt} - \beta r_{j'mt+1})\Big) \\ &\qquad + \eta\Big(\log(\bar{ms}_{jmt|g}) + \beta\big(\log(ms_{gmt+1}) - \log(ms_{j'mt+1})\big)\Big) \\ &\qquad + \underbrace{\Big(\xi_{jmt} - \beta\xi_{j'mt+1}\Big)}_{\equiv\tilde{\xi}_{jmt}} + \underbrace{\Big(\bar{\xi}_j - \beta\bar{\xi}_{j'}\Big)}_{\equiv\tilde{\xi}_j} + \underbrace{\Big(\bar{\xi}_t - \beta\bar{\xi}_{t+1}\Big)}_{\equiv\tilde{\xi}_t} + \underbrace{\beta e_{mt}}_{\text{mean zero}}\end{aligned} \tag{C6}$$

Table C1. Estimated Demand Parameters from a Dynamic Model

	(1)
'Dynamic' Price/Income	−0.448
	(0.160)
Nesting Parameter	0.968
	(0.036)
Firm FE	Yes
Year FE	Yes
Observations	16,825
R^2	0.799
F-test (1st stage), Nesting Parameter	46.2
F-test (1st stage), 'Dynamic' Price/Income	2,256.1

Notes: Estimation follows the procedure outlined in Appendix C. I divide prices and rebates by county-quarter mean income whereas for comparison to the static demand estimates in Table 2. Standard errors clustered by county are reported in parentheses.

The above estimating equation (C6) 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.

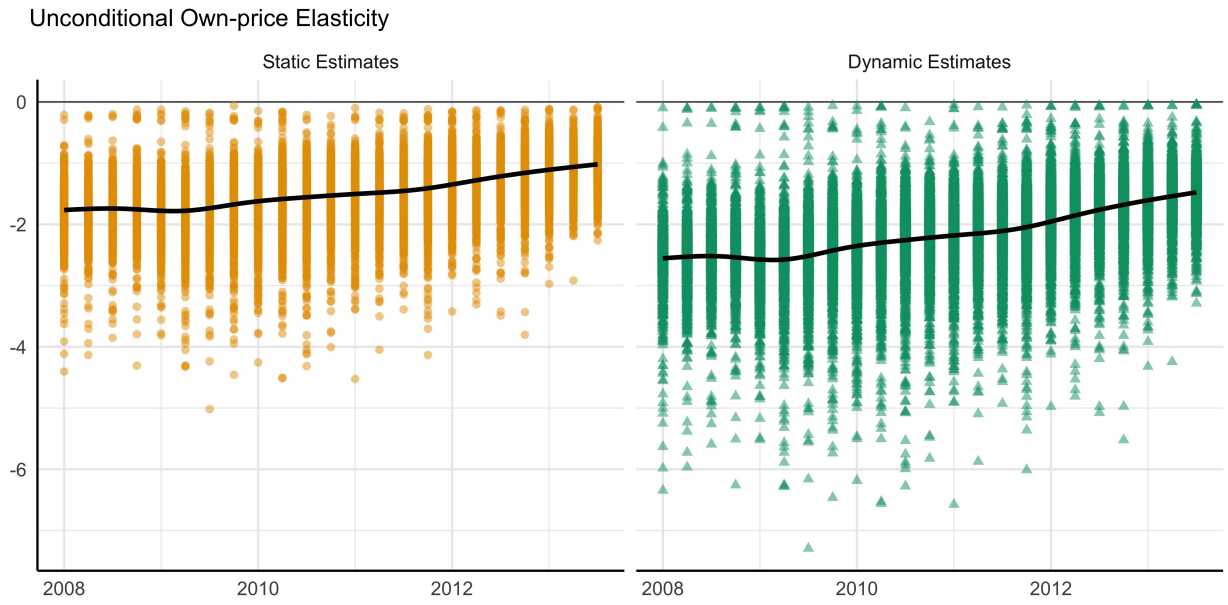
Two issues remain before implementing (C6): the choice of reference option, j', and the treatment of the scalar unobservable terms, $(\tilde{\xi}_{jmt}, \tilde{\xi}_j, \tilde{\xi}_t)$. In terms of the reference option, an easy solution would be to use a universal installer that appears across markets. 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 (C6). This revised estimating equation converges asymptotically to (C6).

To account for the scalar unobservable terms, I include firm and time fixed effects when implementing (C6). 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 (C6) via two-stage least squares, with the same set of demand instruments as those that I outline in Section 5.1.1 of the text.

C.4 Results

I report estimates of the main target parameters in Table C1. Overall, the results are qualitatively consistent with the static demand estimates reported in Column (1) of Table 2. Note that I divide prices and rebates by the county-quarter mean income for the sake

Figure C1. 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 2 (left panel) and the dynamic demand estimates from Table C1. The overlaid line shows a cubic b-spline fitted to the short-run price elasticity estimates.

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.

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 C1 alongside those from the analogous static model. As shown in Figure C1, 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 -2.07 compared with -1.42 for the analogous static estimates—both are within the range of previously published static estimates (see Figure A3) and the difference between the two is relatively minor. Figure C1 suggests that the reduced form demand model used in the main analysis in the text offers a reasonable approximation.

D Estimating the Exit Policy Function

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

$$\Pr(\chi^x_{jmt} = 1|s_{mt}) = \frac{\exp\left(h_j(s_{mt})\right)}{1 + \exp\left(h_j(s_{mt})\right)}$$

where χ^x_{jmt} 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.[38] 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:

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

I select the tuning parameter, λ, by leave-one-out k-fold cross validation with $k = 10$. Figure D1 shows the resulting estimated binomial deviance for different values of λ.

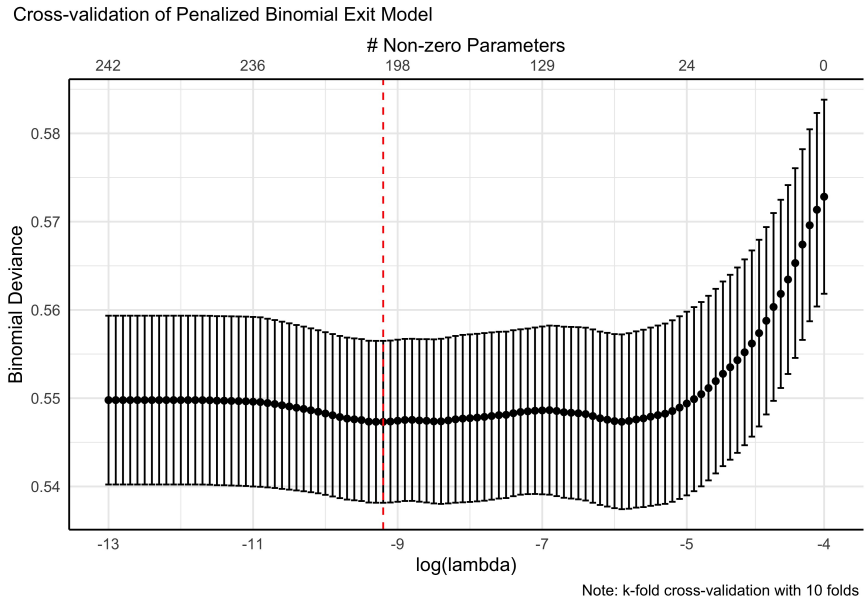
This identifies a set of non-zero regressors, $\tilde{h}_j(s_{mt};\tilde{\mu})$. I then model the discrete exit decision using a logit model with the non-zero regressors selected in the first stage and estimate the parameters on the remaining regressors via maximum likelihood:

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

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 D2 shows the density of fitted exit probabilities for incumbents that I observe continue and incumbents that I observe exit.

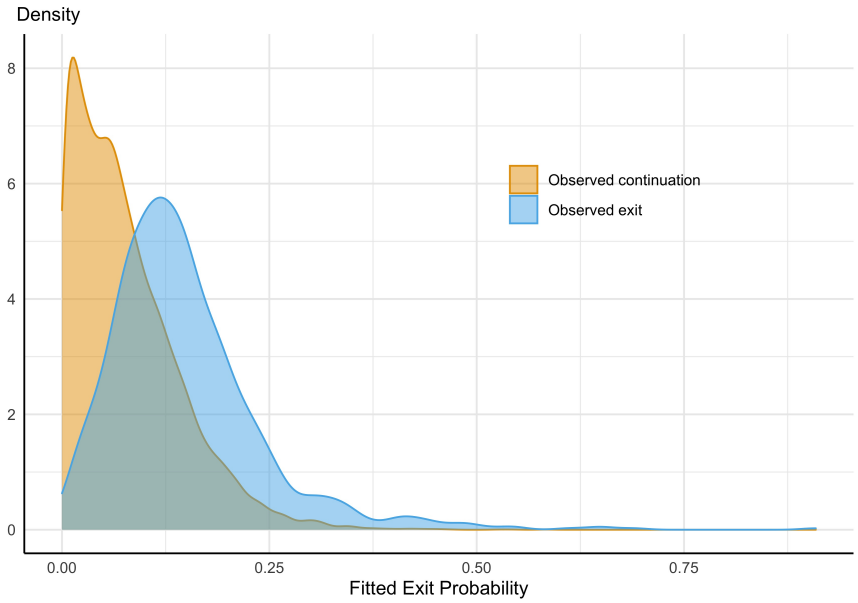
[38] I include quadratic polynomials of the following variables and their pairwise interactions: prices, own experience, other firms' experience within a county, other firms' experience outside a county, hardware cost, the average hardware costs of other firms within a county, quality, the average quality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

Figure D1. 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 D2. 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.

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

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

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. (2021)). 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 $b_j^l(s_{mt})$:

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

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

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

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):

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

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.

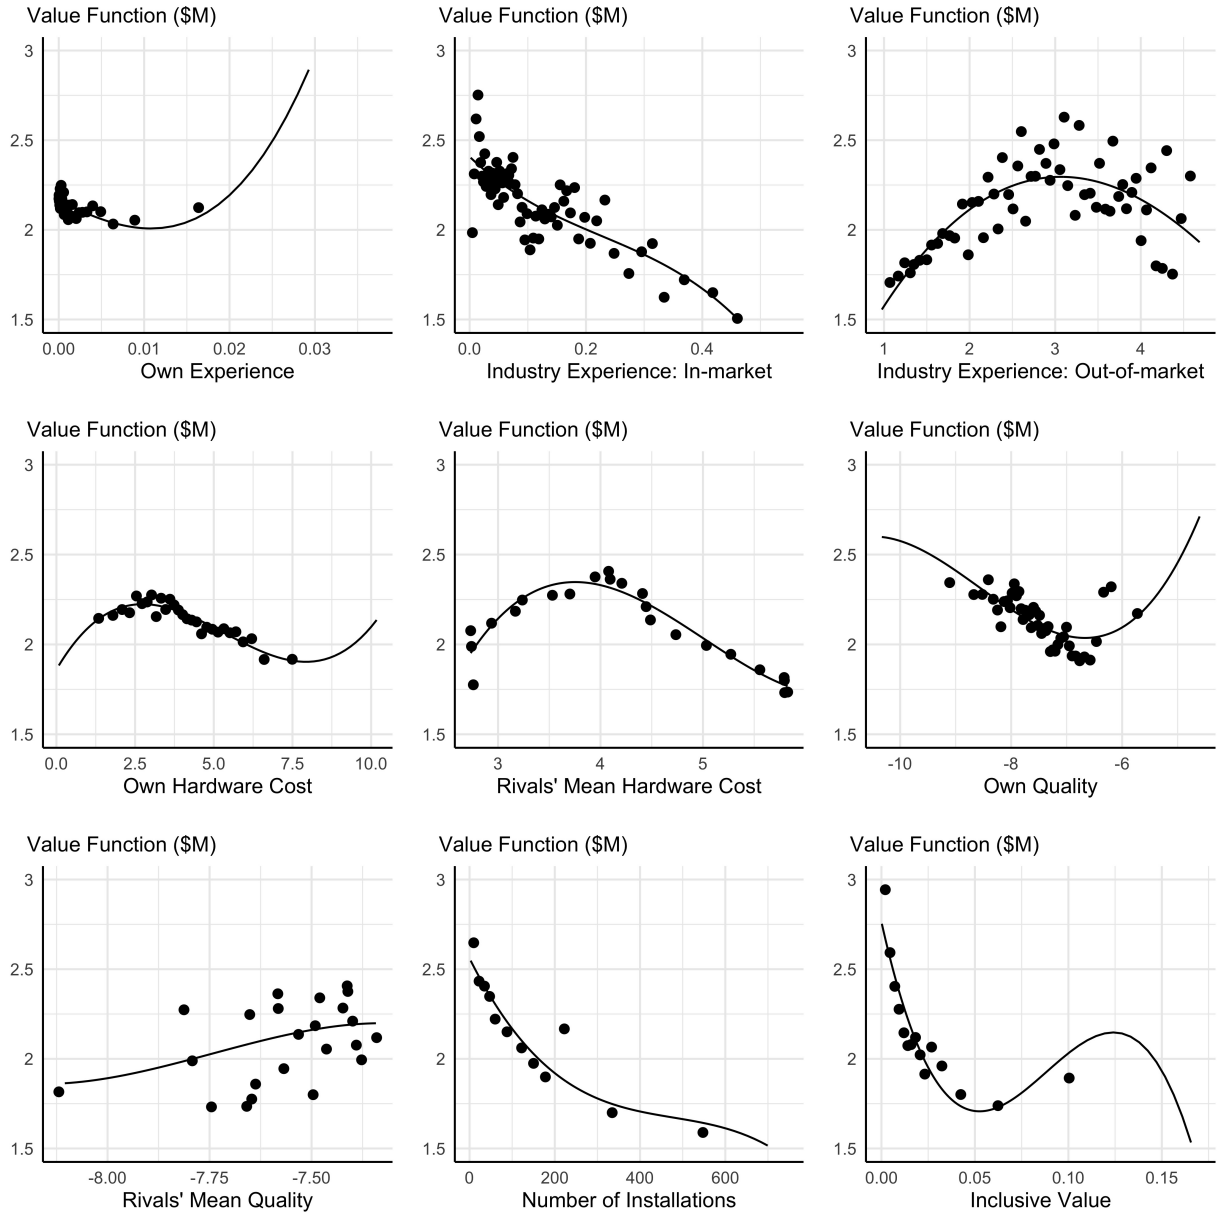
To select the basis function of the above 11 variables, I test how well B-splines of different orders with different percentile knots fit observed revenue data, since value functions measure expected discounted profits which are functions of revenues. I ultimately select third-order B-splines (i.e., 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. (2021), 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 (2) of Table 3.

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 3. 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., 2021; 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.[39] 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.

[39] I 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.

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.[40]

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.[41]

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

[40] In 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}.

[41] Naturally, 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.

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 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 A2), 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 B, 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 2) 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 meter-

ing (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.

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. 2021. "Industrial Policy Implementation: Empirical Evidence from China's Shipbuilding Industry." *Working Paper*.

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.

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.

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.

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.