Learning-by-Doing in Solar Photovoltaic Installations*

Bryan Bollinger[†]
Duke University

Kenneth Gillingham[‡]
Yale University

April 3, 2019

Abstract

The solar photovoltaic (PV) industry in the United States has been the recipient of billions of dollars of subsidies, motivated both by environmental externalities and spillovers across firms from learning-by-doing (LBD) in the installation of the technology. Using a dynamic model of demand and supply, this paper investigates installation cost reductions due to localized LBD using comprehensive data on all solar PV installations in California between 2002 to 2012, during a stage of initial growth of the PV market. We find that appropriable LBD can explain a decline in non-hardware costs of around 12 cents per watt, but we find evidence of only very small learning spillovers. This suggests that the California incentives are difficult to justify on short-run economic efficiency grounds.

Keywords: innovation; dynamic structural models; imperfect competition; diffusion; new technology; energy policy.
JEL classification codes: Q42, Q48, L13, L25, O33, O25.

*The authors would like to thank the participants at seminars at the Harvard Kennedy School, University of Chicago Booth, University of Michigan, London School of Economics, Yale University, University of Pennsylvania Wharton, Duke University, New York University, Stanford SITE Meetings, and the NBER EEE Summer Institute for their valuable comments and insights. We also acknowledge the excellent research assistance of Divita Bhandari and Hao Deng and support from the U.S. Department of Energy under contract DE-AC02-05CH11231. All errors are solely the responsibility of the authors.

[†]Fuqua School of Business, Duke University, 100 Fuqua Drive, Durham, NC 27708, phone: 919-660-7766, e-mail: bryan.bollinger@duke.edu.

[‡]Corresponding author: School of Forestry & Environmental Studies, Department of Economics, School of Management, Yale University, 195 Prospect Street, New Haven, CT 06511, phone: 203-436-5465, e-mail: kenneth.gillingham@yale.edu.

1 Introduction

Policies to promote solar photovoltaic (PV) system adoption are common throughout the world, as concerns over global climate change and energy independence continue to grow. In the United States, commercial and residential solar installations remain eligible to receive a 30% solar energy investment tax credit. This federal subsidy comes on top of individual state incentive programs, the most prominent of which was the California Solar Initiative (CSI), a $2 billion program begun in 2006 to provide substantial upfront rebates for rooftop solar installations (which have since been exhausted). Policies like the CSI are often justified based on both emissions reductions and the existence of *learning-by-doing* (LBD), whereby the cost of the technology declines as a function of cumulative experience with the technology.

This study investigates LBD in solar PV installations in California. For LBD to justify government intervention based on improving economic efficiency, (non-internalized) learning spillovers across firms must exist. Such learning spillovers are often called "non-appropriable" LBD, similar to the non-appropriable benefits from research and development, which are a standard justification for innovation policy. Using rich data on all solar installations in California in the early stage of the growth of the market from 2002 to 2012, we separately estimate the magnitude of appropriable LBD (internal learning) and non-appropriable LBD (external learning) in the cost of an installation. This time period is the crucial early period when one would expect the bulk of learning to occur. Since solar PV panels and inverters are traded on a global market, we focus on localized learning in non-hardware costs, which include labor, overhead, and marketing costs.

Two main challenges exist in identifying LBD in this industry. First, consumers' utility for a specific installer may increase with the installer's installed base as quality or perceived quality increases, which would increase the optimal markup for the installer, pushing price upward despite cost reductions that may result from LBD.[1] To address this

[1] It should be noted that quality improvements that result from installer experience also could be described as LBD; however, these improvements are firm-specific and thus would classify as internal learning. Throughout the paper, when we refer to LBD, we are referring to reductions in variable costs.

challenge, we estimate a dynamic nested Logit model of solar PV demand using conditional choice probabilities (CCP), following a similar approach as Hotz and Miller (1993). The second challenge in estimation is that with appropriable LBD, installers also have an incentive to lower prices in the short term to reduce costs (Benkard 2004), or to increase prices in order for costs to decline for other reasons, both of which can further obfuscate cost reductions from LBD. To tackle this second challenge, we estimate a dynamic model of installation pricing, using forward simulation as in Bajari, Benkard, and Levin (2007), aka BBL, to estimate future valuations. Thus, our approach is to directly estimate both the static and dynamic markups in the installers' first-order pricing condition.

We do indeed find that utility for an installer increases with the installer's installed base (i.e., cumulative installations) within the county, providing evidence of a quality or perceived quality effect of installers' local experience. We also find evidence of internal learning in contractor non-hardware costs, and this internal learning increases with competitors' installed bases within the county, providing evidence for external learning as well. However, the learning is small in magnitude; LBD can account for just a \$0.12 per watt decline in non-hardware costs over the data period. As a reference point, during this period, hardware costs declined from over \$7 per watt to less than \$3.50 per watt.[2] The difference between the transaction cost per watt and the hardware cost per watt, described in the industry as the "balance-of-system" (BOS) and which includes both the non-hardware costs (customer acquisition, labor, permitting, etc.) and the firm markup, declined less than a dollar, from \$3 per watt to a little over \$2 per watt, only 15% of which we can attribute to LBD. Thus it is hard to justify the substantial CSI incentives from an economic efficiency argument alone.

In addition to these substantive findings, we contribute to the literature on the estimation of dynamic models. We are one the the only papers to allow for dynamics on both the demand and supply side, and we provide methodological contributions to both. On the

[2]Hardware costs include the modules and inverters. The inverter converts electricity from direct current (DC) to alternating current (AC) and accounts for roughly 6 to 15 percent of the total cost. Both are traded on a global market, with manufacturing in Asia, Europe, and North America for use anywhere in the world (IEA 2009).

demand side, we estimate the dynamic model by representing consumers continuation values with the use of market-level conditional choice probabilities (CCPs), exploiting the fact that installing solar is a terminating action, as in DeGroote and Verboven (2019). Unlike DeGroote and Verboven (2019), we allow for correlated demand shocks through the use of an aggregated nested Logit model, in which we assume that installing solar with any of the active installers in the market (defined at the county level) are choices within one nest, and not installing is in the other. We estimate the nest parameter to be 0.57, indicating that much of the unobservable demand shock affects the decision of whether to install solar at the current time, not which installer to use. We account for across-market heterogeneity by including installer X market fixed effects, which allows consumers in different markets to have different price elasticities for each installer. This essentially assumes a representative consumer within each market. However, we also show that our effects are robust to within-market unobserved heterogeneity, in which higher utility consumers, the low-hanging fruit, adopt early, which changes the distribution of consumers over time. Our novel method decomposes the observed market-level CCPs into CCPs for the two latent types of consumers (adding more types is possible) as the distribution of non-adopters evolves. With very durable purchases such as solar PV, this is necessary if within-market heterogeneity is of concern and is to be accounted for correctly.

On the supply side, the dynamic pricing problem for each installer is fundamentally a forward-looking optimization problem, in which each installers optimal pricing decision depends on the number of ongoing installations and installed bases for every other installer operating in California, of which we observe over 3000. We develop a new, tractable estimation approach for installers pricing problem, recognizing that the only source of supply-side dynamics enters through the consumers choice probabilities. The benefit to the installer of lowering price in the short run is due to the increased probability of performing the installation, which can contribute to economies of scale in the short run, and increase market power and/or lower costs in the long run. Thus, for any given installation, and with non-parametric policy function estimates in-hand, we can forward simulate states of the world (a la BBL) for each possible installer who might have performed the installation (those installers operating in the market, plus the outside option of no installer performing the installation). As the installer who we observe in the data performing the installation lowers its price, it increases the probability of the set of forward simulations under that scenario (in which it performed the installation) being the expected evolution of the market. Correspondingly, it lowers the probabilities of the other sets of simulations representing the future, under the scenarios that one of them performed the installation. By forward simulating a set of market paths for each installer who might have performed the installation in the market, for each observation in the data, we also do not need to calculate the value functions for the entire state space (which is completely intractable given the fact there are four firm-specific variables in the state space for over 3000 firms). Instead, the value is calculated only for the permissible states that might be realized, starting from the set of observations we observe in the data. Once we perform the forward simulations, the expected profits the firm would make under the counterfactual outcomes of which installer performs the installation enters directly into the installers first-order pricing condition.

The rest of the paper is organized as follows. Section 2 provides a conceptual background on LBD in the economics literature. Section 3 describes the empirical setting: the California solar market. In Section 4, we develop our dynamic demand model and in Section 5 our supply model. Section 6 describes the data we use and discusses identification of our parameters. Section 7 presents our results. Finally, section 8 concludes.

2 Conceptual Background

The concept of LBD in economics dates to the early 1960s, beginning with theoretical work by Arrow (1962). Since then, economists have developed a wealth of theoretical findings based on LBD. For example, learning from cumulative experience has been shown to play a critical role in both the functioning of markets (e.g., Spence 1981; Fudenberg and Tirole 1983; Ghemawat and Spence 1985; Cabral and Riordan 1994; Besanko, Doraszelski, Kryukov, and Satterthwaite 2010) and in theories of endogenous growth (e.g., Stokey 1988; Young 1991, 1993; Jovanovic and Nyarko 1996). One notable theoretical finding is that when firms can appropriate the benefits from learning, they have an incentive to price dynamically by initially pricing below the short-run marginal cost in order to allow for future market dominance. We provide evidence suggestive of this force in our empirical context.

LBD also underpins an extensive empirical literature. Studies have estimated the speed of learning in a wide variety of contexts, including aircraft manufacturing (Alchian 1963; Benkard 2004), chemical processing (Lieberman 1984), semiconductor manufacturing (Irwin and Klenow 1994), agricultural technology (Foster and Rosenzweig 1995), shipbuilding (Thompson 2001; Thornton and Thompson 2001), oil drilling (Kellogg 2011; Covert 2014), and automobile manufacturing (Levitt, List, and Syverson 2013). LBD has also long been used to examine the cost of new energy technologies, beginning with Zimmerman (1982), and more recently as a common descriptive methodology for modeling technological change in renewable energy technologies.[3]

Given the importance for policy of differentiating between internal and external learning, it is not surprising that several empirical studies distinguish between the two. Learning spillovers across firms have been studied in several contexts (Zimmerman 1982; Irwin and Klenow 1994; Thornton and Thompson 2001; Kellogg 2011; Covert 2014). These spillovers have also been shown to influence market structure by undercutting barriers to entry (Ghemawat and Spence 1985) and at the same time may represent a classic positive externality (Stokey 1985; Melitz 2005; van Benthem, Gillingham, and Sweeney 2008; Gillingham and Sweeney 2010; Gillingham and Stock 2018). Both effects may be important in the solar PV market. van Benthem, Gillingham, and Sweeney (2008) perform an ex-ante welfare analysis of the CSI assuming non-appropriable LBD. They find that prior to the addition of Federal tax credits, the CSI can be justified on economic efficiency

[3]See Grubb, Khler, and Anderson (2002) and Gillingham, Newell, and Pizer (2008) for reviews of the modeling of endogenous technological change in climate policy models, and Nordhaus (2014) for an important critique of the naïve use of LBD in such models.

grounds based on the avoided environmental externalities and LBD spillovers–provided that learning follows the rates found in the literature *and* the learning is non-appropriable. This is important because the CSI was explicitly justified in the policy process based on both environmental grounds and learning. However, Borenstein (2008), van Benthem, Gillingham, and Sweeney (2008), and Burr (2014) clearly show that the CSI cannot easily be justified on economic efficiency grounds based on environmental externalities alone–non-appropriable learning is critical.

Learning can be expected to lower non-hardware costs for solar PV installations at a regional or localized level by improving labor productivity. Employees can increase the speed of installation with different types of roof layouts, discover ways to modify the hardware to facilitate installation, refine the site-visit software, and improve the processing of permits. Spillovers may occur through pathways such as hiring employees of other firms, watching competitor strategies, increased efficiency of permitting by building permit offices, and more widespread adoption of best practices as are publicized by industry organizations. Of course, labor markets may adjust in response to some of these pathways based on labor productivity, but if there are sticky wages and sufficiently high unemployment, as was the case in much of our empirical setting, LBD may still bring down labor costs.

To estimate how own experience (internal learning) and competitor experience (spillovers) reduce the non-hardware costs, we face a classic empirical challenge in industrial organization: we observe hardware costs (i.e., module and inverter costs) and the price of the system, but we do not separately observe the BOS and the markup. The optional static markup will change over time, and these changes may be corrected with the contractor cumulative number of installations if consumers perceive more experienced contractors as higher quality, thus allowing them to charge a higher markup. We would also expect firms that anticipate LBD to be pricing dynamically, so the markup would be lower in early periods and higher in later periods. For example, Benkard (2004) estimates a dynamic model of aircraft pricing, using marginal cost data, and shows that it may be optimal for firms to begin pricing considerably below static marginal costs.

3 Empirical Setting

3.1 Solar Policy

There has been a long history of government support for solar energy in both the United States in general and in California specifically. At the federal level, incentives for solar date back to the Energy Tax Act (ETA) of 1978. More recently, the Energy Policy Act of 2005 created a 30% tax credit for residential and commercial solar PV installations, with a $2,000 limit for residential installations. The Energy Improvement and Extension Act of 2008 removed the $2,000 limit and the American Recovery and Reinvestment Act of 2009 temporarily converted the 30% tax to a cash grant.

California's activity in promoting solar began as early as 1974 with the creation of the California Energy Commission (CEC). For several decades much of the emphasis was on larger systems. In 1997, California Senate Bill 90 created the Emerging Renewables Program, which directed investor-owned utilities to add a surcharge to electricity bills to promote renewable energy. The proceeds of this surcharge supported a $3 per watt rebate for distributed solar PV installations (Taylor 2008). Beginning in 1998 "net metering" allowed owners of solar PV systems to receive credit for electricity sold back to the grid. Moreover, from 2001 to 2005, a 15% state tax credit was granted for solar PV installations (CPUC 2009).

While the California rebate program put in place in 1997 was substantial, it was renewed on a year-by-year basis, leading to uncertainty in the solar market. The elements for a longer-term, more predictable policy originated in August 2004, with the announcement of the "Million Solar Roofs Initiative," a program with a goal of one million residential solar installations by 2015. In January 2006, the California Public Utilities Commission (CPUC) established the CSI, the $2.167 billion program aiming to install 1,940 MW of new solar by 2016 and "to transform the market for solar energy by reducing the cost of solar"

(CPUC 2009).

The CSI is a somewhat unusual subsidy policy in that it *counted on* LBD bringing down the cost of solar, for the subsidy declined in steps over time as the number of installed MW increases. As shown in Figure 1, the CSI used a separate step schedule for each of the three major investor-owned utilities in California: Pacific Gas & Electric (PG&E), Southern California Edison (SCE), and San Diego Gas & Electric (SDG&E).[4] Outside of these, there are also municipal utilities, such as the Los Angeles Department of Water and Power. The larger program that included the municipal utilities aimed to install 3,000 MW of solar PV by the end of 2016, for a total statewide budget of $3.3 billion. The number of installations in California exceeded expectations and the programs in all three utility regions are closed in 2015.

Adoption rates in CA increased quickly between 2002 and 2012, as shown in Figure 2. By the end of 2012, California accounted for nearly 50 percent of total US residential and commercial solar PV capacity installed in the U.S., making it the largest and most important market for distributed generation solar PV.[5] The vast majority of these systems were installed in 2002 or later, and thus our panel covers the major growth phase of the CA residential solar market. Over 80 percent of the systems installed in both California and the U.S. by the end of 2012 were under 10 kW, which is a common upper bound size for a small-scale residential or commercial system. This paper does not include the large-scale solar farms (Barbose, Darghouth, Weaver, and Wiser 2013).

Over the course of our time frame, the CA market has gradually become less concentrated although it also has seen the emergence of large players as well (e.g., SolarCity). This latter phenomenon was aided in creation of solar lease (third-party owned) products which were generally not available before 2008.

[4] SDG&E's CSI program was run by the California Center for Sustainable Energy (CCSE)

[5] This estimate is based on the detailed 2013 "Tracking the Sun" report by Lawrence Berkeley National Laboratory (LBNL), which includes roughly 72 percent of all grid-connected solar PV capacity from 1998 to 2012 (Barbose, Darghouth, Weaver, and Wiser 2013).

4 Solar PV Demand

4.1 Demand Model

It is essential to identify how markups change over time since they likely do not change by the same amount for all installers. Furthermore, the optimal markup itself may be a function of the installed bases if installed bases are a signal of (or proxy for) the quality of the installer. We thus estimate a demand model separately for each county-quarter in order to capture these evolving markups. We aggregate at the quarterly level to avoid zero shares for installers which are present in the market but happen to not perform any installations in a particular month (if they perform no installations that quarter we assume they are not active in the market). We assume a nested logit model of demand. The upper nest models whether to purchase solar or not, and the lower nest models the decision of which installer to use, i.e., all installers are in one group and the decision to not install ($j = 0$) is in the other. The mean current period utility of not installing is normalized to zero and installing solar is a terminating state. Following Berry (1994), we assume:

$$u_{ijt} = \mu_{mjt} + \varepsilon^{u}_{igt}(\sigma) + (1-\sigma)\epsilon^{u}_{ijt}, \tag{1}$$

in which the index j indicates the installer and g the nest group. We assume that ϵ^{u}_{ijt} is distributed iid as type one extreme value and $\varepsilon^{u}_{igt}(\sigma)$ has the unique distribution such that $[\varepsilon^{u}_{igt}(\sigma) + (1-\sigma)\epsilon^{u}_{ijt}]$ is distributed as type one extreme value (Cardell 1997).

Let the mean utility of installing solar using installer j in market m be given by:

$$\mu_{mjt} = \alpha(P_{mjt} - R_{mt}) + \theta^{m} b_{mj} + \theta^{-m} b_{-mj} + \theta^{E}\mathbf{Z}_{mt} + \omega^{u}_{j} + \eta^{u}_{m} + \zeta^{u}_{t} + \xi_{jmt}, \tag{2}$$

where P_{mjt} is the average installation price per watt for firm j in market m at time t, R_{mt} is the rebate, b_{mj} is the installer's installed base inside the market, b_{-mj} is the installer's installed base outside of the market. The $\mathbf{Z}_{mt}$ include controls for housing prices, the electricity rate, and the solar radiation for the county in that quarter, shown to positively affect solar adoption rates in Lamp (2014). We define ω_j^u, η_m^u, and ζ_t^u as fixed effects for contractor, market, and year, and ξ_{jmt} is a market-time unobserved shock for installer j with mean zero, unobserved by the econometrician. We assume that the installation decision is an exit decision and that all future value of the installation in captured in μ_{ijt}, making this an optimal stopping problem, as in Rust (1987).

Let the observed market state $\mathbf{x}_{mt}$ be defined as:

$$\mathbf{x}_{mt} \equiv \{P_{mjt}, R_{mt}, b_{mj}, b_{-mj}, \mathbf{Z}_{mt}, \omega_j^u, \eta_m^u, \zeta_t^u, \xi_{jmt}\}. \tag{3}$$

We assume that consumers do not anticipate the added hedonic utility they get from installing solar in a sunnier period (county fixed effects capture baseline sunlight levels). The value function for a household can be defined recursively by the following Bellman equation:

$$V(\mathbf{x}_{mt}, \varepsilon_{igt}^u . \epsilon_{ijt}^u) = \max_{j_{it}} \{u_{ijt} + \beta \mathbb{E}\left[V(\mathbf{x}_{mt}, \varepsilon_{igt+1}^u . \epsilon_{ijt+1}^u)\right]\}, \tag{4}$$

with discount rate β, and in which u_{ijt} depends on the current period shocks as shown in (1) and j_{it} is the consumer choice of which installer to select (including the option of not installing).

We can thus write each household's expected conditional value function as:

$$\begin{aligned} v(x_{mt}, j_t) &= \mu_{mjt} \\ v(x_{mt}, 0) &= \rho \int \int \max_{j'} \left(v(\mathbf{x}_{mt+1}, j_{t+1}) + \varepsilon_{igt+1}^u(\sigma) + (1-\sigma)\epsilon_{ijt+1}^u \right) \\ &\quad dF(\mathbf{x}_{mt+1} | x_{mt}, j_t) dG(\varepsilon_{igt+1}^u, \epsilon_{ijt+1}^u) \end{aligned} \tag{5}$$

in which we integrate over the transition distribution of the state variables $F(\mathbf{x}_{mt+1}|\mathbf{x}_{mt}, j_t)$ and the distribution of the unobservables, $G(\varepsilon_{igt+1}^u, \epsilon_{ijt+1}^u)$.

Because any decision to install is a termination decision, the future value function only appears for the decision to not install, and captures the continuation value. The share of households choosing installer j in market m conditional on installing is given by the familiar nested Logit expression:

$$s_{mjt}^{j/I} = \frac{\exp\left(v(\mathbf{x}_{mt}, j_t)/(1-\sigma)\right)}{D_{mt}^I}, \tag{6}$$

where we have the following inclusive value of installing solar:

$$D_{mt}^I = \sum_{j \neq 0} \exp\left(v(\mathbf{x}_{mt}, j_t)/(1-\sigma)\right) \tag{7}$$

The probability of (and share of people) installing is then given by:

$$s_{mt}^I = \frac{(D_{mt}^I)^{(1-\sigma)}}{\exp\left(v(x_{mt}, 0)\right)^{(1-\sigma)} + (D_{mt}^I)^{(1-\sigma)}}. \tag{8}$$

and not installing by:

$$s_{m0t} = \frac{1}{\exp\left(v(x_{mt}, 0)\right)^{(1-\sigma)} + (D_{mt}^I)^{(1-\sigma)}}. \tag{9}$$

4.2 Estimation

Calculating the conditional value function can be achieved by integrating over the transition distribution of the state variables and the distribution of the unobservables, although we can also calculate it in terms of the conditional choice probabilities (CCPs) for one of the terminating options, set without loss of generality as $j = 1$ (Hotz and Miller 1993; Arcidiacono and Ellickson 2011):

$$v(\mathbf{x}_{mt}, 0) = \beta \int \mu_{m1t+1}(\mathbf{x}_{mt+1}) - \psi(\boldsymbol{\delta}_{t+1}(\mathbf{x}_{mt+1}))dF(\mathbf{x}_{mt+1}|\mathbf{x}_{mt}), \tag{10}$$

where we know from Arcidiacono and Miller (2011) that:

$$\begin{aligned}\psi(\boldsymbol{\delta}_{t+1}(\mathbf{x}_{mt+1})) &= \gamma - (1-\sigma)\log\left(\delta_{m1t+1}\right) - \sigma \log\left(\delta^{j\neq 0}_{mt+1}\right) \\ &= \gamma - \log\left(\delta_{m1t+1}\right) - \sigma\left(\log\left(\delta^{j\neq 0}_{mt+1}\right) - \log\left(\delta_{m1t+1}\right)\right) \end{aligned} \tag{11}$$

and we define the probability of choosing the arbitrarily selected terminating option $j = 1$ and the probability of choosing from that nest, i.e., to install solar at all indicated by $j \neq 0$, with the following two expressions, respectively:

$$\delta^{j\neq 0}_{mt+1} \equiv \sum_{k=1}^{J} \delta(j_{mt+1} = k | x_{mt+1}) \tag{12}$$

$$\delta_{m1t+1} \equiv \delta(j_{mt+1} = 1 | x_{mt+1}) \tag{13}$$

The only integration that remains in (11) is that over the state transitions. We assume that consumers expect the state variables to evolve according to independent AR(1) processes, with the exception of the rebate. Given that the adoption of solar during this period exhibits large "pull-forward" effects as discussed in Rogers and Sexton, and is also evidenced by the spikes in demand shown in Figure 2, we assume that consumers have perfect foresight with regards to the rebate amount. Thus we can write:

$$\Delta v_{mjt} \equiv v(x_{mt}, j_t) - v(x_{mt}, 0) = \mu_{mjt} - \beta \mathbb{E}\left[\mu_{m1t+1} - \psi(\boldsymbol{\delta}_{mt+1}(\mathbf{x}_{mt+1}))\right] \tag{14}$$

This yields the following linear equation which can be used in estimation:

$$\begin{aligned}\log(s_{mjt}) - \log(s_{m0t}) + \beta\gamma - \beta\log\left(\delta_{m1t+1}\right) = \alpha(P^{\Delta}_{mjt} - R^{\Delta}_{mt}) + \theta^{m}\mathbf{b}^{\Delta}_{mj} + \theta^{-m}\mathbf{b}^{\Delta}_{-mj} \quad (15)\\ + \mathbf{Z}^{\Delta}_{mjt}\theta^{Z} + \omega^{\Delta u}_{j} + \eta^{\Delta u}_{m} + \zeta^{\Delta u}_{t} + \sigma\left(\log(s_{mj/It}) + \beta\mathbb{E}\left[\log\left(\delta^{j\neq 0}_{mt+1}\right) - \log\left(\delta_{m1t+1}\right)\right]\right) + \xi^{\Delta u}_{jmt}\end{aligned}$$

in which the Δ superscript designates that we subtract off the discounted value of the next period values for an arbitrary installer in the market whose probability of adoption we use to calcite the continuation value. We use a quarterly discount rate of 0.966

which correspond to an annual discount rate of 0.87, consistent with that estimated by De Groote and Verboven (2016). The demand estimates are robust to varying the discount rate.[6]. This expression only depends on the values of the current and next period state variables and the next period adoption probabilities. These probabilities are calculated at the county-quarter level which is essential since the model includes market-level unobservables.

We split the continuation value into its component that does not depend on σ, which we add to the right hand side of the equation, and the component which does, which can be include in the current within-group share term. For identification, we need instruments for the price and for the within-group share parameter. We use two cost shifters of the total installed cost: the mean rebate per watt, and the number of installations the installer has finished in other counties. The first is a straightforward cost shifters, as the rebate is given directly to the contractor. The second, the number of installations the installer has finished in other counties, might be expected to be a strong instrument for the within-group share parameter because if there are more finished installations in other counties, this frees up labor that can be moved across county borders, influencing the within-group share. At the same time, after inclusion of our time fixed effects, installations in other counties should not influence demand in the county of interest.

The number of installations the installer has finished in other counties should also impact the within-group share, since this cost shorter is installer-specific. Similarly, we would expect the number of installations the competitors have finished in other counties to also shift within-group share. Thus we include this third instrument as well, providing us with an over-identified model, which allows us to then test the over identifying restrictions.

Which firm is used to control for future utility does not matter in theory, but the challenge we face is that there is no one firm that is well represented in all markets in all years. Thus we use a novel strategy in which we average the values for all firms in the market that year. This will yield the same results asymptotically since we simply average

[6]We also use a 0.90 discount rate as is typical and corresponds to the value estimated by Bollinger (2015).

equation (15) for each possible choice of the focal firm. The final estimation equation is:

$$\log(s_{mjt}) - \log(s_{m0t}) + \beta\gamma - \beta\mathbb{E}\left[\overline{\log(\delta^j)}_{mt+1}\right] = \theta^m b_{mj} + \theta^{-m} b_{-mj} + \alpha(P^{\Delta}_{mjt} - R^{\Delta}_{mt})$$
$$+\mathbf{Z}^{\Delta}_{mjt}\theta^Z + \omega^{\Delta u}_j + \eta^{\Delta u}_m + \zeta^{\Delta u}_t + \sigma\left(\log(s_{j/It}) + \beta\left(\log\left(\delta^{j\neq 0}_{it+1}\right) - \overline{\log(\delta^j)}_{mt+1}\right)\right) + \xi^{\Delta u}_{jmt} \quad (16)$$

where we define:

$$\overline{\log(\delta^j)}_{mt+1} \equiv \frac{1}{|C_{mjt+1}|}\sum_{k\in C_{mjt+1}} \log\left(\delta_{ikt+1}\right) \quad (17)$$

in which C_{jt+1} is the set of installers active in the county, $|C_{mjt+1}|$ is the cardinality of C_{mjt+1}, and Δ superscript now designates that we subtract off the discounted value of the next period mean values for all installers in the market.

In order to calculate the expected next period probabilities, we assume that consumers expect AR(1) transitions for the shares and inside good shares and use the predicted values. We do this because some of the state variables affect all markets and thus we would not want to use only realized next period probabilities.[7]

We use aggregate data for our CCP estimation, just as was done by Derdenger and Kumar (2015) and De Groote and Verboven (2016), because this enables us to use the full dataset in estimation.[8] Furthermore, there is little to be gained from using disaggregated data since we the only household level state in our state space is whether the household has already installed solar (if they have, this excludes them from installing in the future). This approach does limit attempts to identify within-county unobserved heterogeneity, but since solar PV adoption is still early along the adoption curve in our empirical setting, the marginal consumer is likely not changing significantly. Dynamics, however, is a first order concern which is confirmed by our estimation results.

[7]As an alternative, we can use the predicted state transitions and estimate the model iteratively, using the previous iterations' estimates to estimate the next period probabilities as a function of state, and then integrate over the AR(1) transitions of the state variables.

[8]Including a separate observation for each household x month combination would make the estimation intractable.

We can calculate the derivative of market share with respect to price as follows:

$$\frac{\partial s_{mjt}}{\partial P_{mjt}} = \alpha \frac{1}{(1-\sigma)} s_{mjt}(1 - \sigma s_{mj/It} - (1-\sigma) s_{mjt}) \tag{18}$$

We can use the derivative of market share with respect to price to calcite the optimal static markup, $\frac{s_{mjt}}{\frac{\partial s_{mjt}}{\partial P_{mjt}}}$.

5 Solar Pricing

5.1 Model

To account for the markup that results from the dynamic pricing incentives in addition to the static markup, we develop a model of forward-looking solar PV contractor pricing. This is complicated by several factors. First, we need to control for firm heterogeneity in costs and in markups. Moreover, the drop in global module prices after 2008 did not correspond to as much of a drop in installation price, suggesting that there may be considerable time-varying market power at the contractor level.

5.1.1 Installer Profits

A contractor $j \in \mathcal{J}$ earns the following profits from installation i that it performs in market m at time t:

$$\pi_{ijt} = (p_{ijt} - c_{ijt} - w(S_{it}, \mathbf{q}_{mjt}, e_{ijt}(\mathbf{b}_{ijt}), W_{mt}), \varepsilon_{ijt}) S_{it}, \tag{19}$$

where p_{ijt} is the price per watt charged for the installation, c_{ijt} is the per-watt cost of the solar panels and inverters, and $w(\cdot)$ denotes the non-hardware costs, defined here as all costs minus the module and inverter costs. The non-hardware costs are a function of the system size S_{it} (in kilowatts), the number of ongoing installations by the contractor within and outside the county, $\mathbf{q}_{mjt}$, which capture economies of scale,[9] the contractor's knowledge or experience, $e_{ijt}(\mathbf{b_{ijt}})$ that is relevant for installation i, and the prevailing wage rate in installation's market, W_{mt}.

The installer's experience is a function of the vector of depreciated installed bases, $\mathbf{b_{ijt}}$, both its own and its competitors' in both the local market m, as well as in all other markets m. We track installations for all firms by market to allow local installations to have a different impact on non-hardware costs than installations performed farther away and to capture the evolution of the four installed bases that affect an installer in a specific market (own installations the market and own installations elsewhere to capture appropriable learning, and competitors' installations in the market and elsewhere to capture learning spillovers). The system size accounts for possibly installation-specific economies of scale.

Solar installation prices are typically set on an installation-by-installation basis since each potential installation has idiosyncracies that influence the cost. Furthermore, the size of every installation is generally set in large increments (i.e., with the addition or removal of a large panel) and is a function primarily of the available suitable roof space and the amount of electricity the consumer uses. Importantly, the system size is not a strategic choice variable for the installer.

We assume that an installer can quote an installation price to each potential customer. As with the demand model, we will use δ to designate demand probabilities. Let the probability that an installer is selected for installation i be $\delta_{ijt} \equiv \delta_{ijt}(p_{ijt}, \mathbf{p_{i-jt}}, \mathbf{b_{it}}, R_{mt})$, which is a function of the price offered for installation i by this contractor, p_{ijt}, the prices offered by other contractors $\mathbf{p_{i-jt}}$, and preferences for the installers based on the previous installed bases that affect installation i, $\mathbf{b_{it}}$ (where $\mathbf{b_{it}}$ is the stacked installed base variables for the set of installers), and the rebate, R_{mt}. We define $\boldsymbol{\delta}_{it}$ as the vector of all probabilities over installers. Let $\mathbf{p}_{jt}$ and $\mathbf{p}_{-jt}$ be the vector of prices across installations at time t for installer j and its competitors, respectively, and $\mathbf{p}_t$ be the full set of prices at time t. We drop the arguments from the marginal cost function for notational convenience, and with slight abuse of notation write $w(S_{it}, q_{mjt}, e_{ijt}(\mathbf{b_{ijt}}), W_{it})$ as w_{ijt}. Let Q_{mt}

[9] Defined as the number of non-completed installations.

be the number of available potential installations in market m at time t i.e., the aggregate demand and $g_{mt}(S_{it})$ be the probability distribution of sizes. We define $\mathbf{q}_{jt}$ be the vector of ongoing installations for firm j in each market m at time t, and $\mathbf{q}_t$ be the number of ongoing installations for all installers.

The observable (to the econometrician) state variables are $\mathbf{X}_t \equiv \{\mathbf{q}_t, \mathbf{p}_t, \mathbf{c}_t, \mathbf{b}_t, \mathbf{W}_t, \mathbf{Q}_t, \mathbf{g}_t(S)\}$, in which we stack the market and installer-market variables into vectors subscripted by t. $\boldsymbol{\varepsilon}_{it}$ is the vector of installer non-hardware cost shocks, ε_{ijt}. An installer's expected profits (conditional on the vector of state variables) at time t are:

$$\Pi_j(\mathbf{X}_t, \varepsilon_{ijt}, \mathbf{p}_t) = \sum_{m=1}^{M} \sum_{i=1}^{Q_{mt}} \pi_{ijt} \delta_{ijt}(\mathbf{X}_t, \boldsymbol{p}_{it}) - FC_j. \tag{20}$$

We define FC_j as the installer fixed costs, which are assumed not to vary with the size of the systems installed. We assume that the installer is at least breaking even in the medium-run and thus is not planning to exit the market.

Following Ericson and Pakes (1995) and BBL, prior to the realization of the non-hardware cost shocks, a forward-looking profit-maximizing firm has an expected value of:

$$EV_j(\mathbf{X}_t) = \mathbb{E}\left[\sum_{\tau=t}^{\infty} \rho^{(\tau-t)} \max_{p_{ijt}} \left[\Pi_j(\mathbf{X}_\tau, \varepsilon_{ijt}, \mathbf{p}_\tau)\right] |\mathbf{X}_t\right] \tag{21}$$

in which ρ is the installers' discount rate.

Let us write the profile of Markov pricing strategies as $\boldsymbol{\sigma}(\mathbf{X}_t)$. Today's value function, post-realization of the cost shoes, can be written recursively using the following Bellman equation:

$$V(\mathbf{X}_t, \boldsymbol{\varepsilon}_{it}; \boldsymbol{\sigma}) = \Pi(\mathbf{X}_t, \varepsilon_{ijt}, \sigma(\mathbf{X}_t, \boldsymbol{\varepsilon}_{it})) + \rho\mathbb{E}\left[\int V(\mathbf{X}_{t+1}, \boldsymbol{\varepsilon}_{it}, \boldsymbol{\sigma}(\mathbf{X}_{t+1}, \boldsymbol{\varepsilon}_{it}))dF(\mathbf{X}_{t+1}|\boldsymbol{\sigma}(\mathbf{X}_{t+1}), \mathbf{X}_{t+1})\right]$$

The pricing profile $\mathbf{p} = \boldsymbol{\sigma}(\mathbf{X}_t, \boldsymbol{\varepsilon}_{it})$ is a Markov Perfect Equilibrium (MPE) if and only if each firm prefers σ_i to any alternative strategy, given the opponents' strategies $\boldsymbol{\sigma}_{-i}$. We assume that the data are generated by a single MPE.

We assume that the state variable transitions can be broken up into independent processes as follows:

$$f(\mathbf{X}_{t+1}|\mathbf{X}_t, \boldsymbol{p}_{it}) = f^W(\mathbf{W}_{t+1}|\mathbf{W}_t) f^c(\mathbf{c}_{t+1}|\mathbf{c}_t) \prod_m f^Q(Q_{mt+1}|Q_{mt}) f^S(S_{mt+1}|S_{mt}) \tag{22}$$

$$f^b(\mathbf{q}_{mt+1}, \mathbf{b}_{mt+1}|\mathbf{q}_{mt}, \mathbf{b}_{mt}, \boldsymbol{\delta}_m(\boldsymbol{X}_t, \boldsymbol{\sigma}(\mathbf{X}_t, \boldsymbol{\varepsilon}_{it}))) \tag{23}$$

We allow for market-specific wage rates and market x firm-specific hard costs, although we assume a common evolution since these are more global processes. We assume that hard costs, market-level wages, new demand, and average installations sizes (in market m at time t) evolve according to AR(1) processes. We define the evolution of the installed base for installer j in market m at time t as a function of the decay parameter κ and a random process describing the number of completed installations:

$$b'_{mj} = \kappa^1 b_{mj} + q^c_{mj}, \tag{24}$$

where the number of completed installations is distributed as a binomial distribution, $q^c_{jt} \sim bin(q_{jt}, \nu)$ where ν is the probability each ongoing installation in q_{jt} is completed that period.

The ongoing installations state variable evolves such that each new potential installation in Q_{mt+1} is performed by firm j with probability δ_{mjt} as determined by the demand model. Thus the vector of ongoing contracts for all firms in market m grows by the expected number of new installations and shrinks by the number of completed installations:

$$\mathbf{q}_{mt+1} = \mathbf{q}_{mt} - \mathbf{q}^c_{mt} + \sum_{q=1}^{Q_{mt}} MN(\boldsymbol{\delta}_m(X_t, \boldsymbol{p}_{it})) \tag{25}$$

in which MN indicates a multinomial distribution with probabilities given by the vector δ_m of length J_m, the number of installers in market m.

The state transitions obviously depend on the installers' pricing policy functions, which depend on the structural unobservable component to non-hardware costs. However, this dependency is completely through the choice probabilities for that installation, $\boldsymbol{\delta}_{it}$. Before solving for the installer's first order pricing condition, we first define the state vectors $\boldsymbol{X}_{t+1}^{q_k+}$ as the next period's state if installation i is assigned to installer k and $\boldsymbol{X}_{t+1}^{q_0+}$ if it is assigned to the outside option, rather than it being determined by $\boldsymbol{p}_{it}$ and $\boldsymbol{\varepsilon}_{it}$. Note that the transition probabilities for these objects do not depend on price because we are assigning the installation i to a specific installer, and for the reason they are also the same for all installers we might assign the installation to, i.e.:

$$f(\mathbf{X}_{t+1}^{q_k+}|\mathbf{X}_t, \boldsymbol{p}_{it}) = f(\mathbf{X}_{t+1}^{q_k+}|\mathbf{X}_t) = f(\mathbf{X}_{t+1}^{q_0+}|\mathbf{X}_t) = f(\mathbf{X}_{t+1}^{q_{k'}+}|\mathbf{X}_t)\,\forall k' \neq k \tag{26}$$

5.1.2 First Order Condition

The dynamics result from the fact that different prices will lead to different evolution of the ongoing installations vector, and these installations then enter the installed base variables after the installations are completed. Economies of scale, appropriable LBD, and an increase in market power through installed base could all lead to an incentive to lower price. In addition, due to the hard cost declines over time, there may also be value in increasing price so that more installations are performed when costs are lower. Thus the direction of the effect of accounting for dynamics on price is not clear a priori.

Differentiating with respect to price leads to the following first order condition for installation i[10]:

$$\begin{aligned}
\frac{\partial}{\partial p_{ijt}} V(\mathbf{X}_t; \boldsymbol{\sigma}) = 0 = \frac{\partial}{\partial p_{ijt}} \Pi(\mathbf{X}_t, \varepsilon_{ijt}, \boldsymbol{\sigma}(\mathbf{X}_t, \boldsymbol{\varepsilon}_t)) \\
+\rho \sum_{k \in \{0, C_{mt}\}} \frac{\partial \delta_k(\mathbf{X}_t, \varepsilon_{it})}{\partial p_j} \mathbb{E}_t \Big[\int \Big(V(\mathbf{X}_{t+1}^{q_k+}, \boldsymbol{\varepsilon}_{it+1}, \boldsymbol{\sigma}(\mathbf{X}_{t+1}^{q_0+}, \boldsymbol{\varepsilon}_{t+1})) \\
- V(\mathbf{X}_{t+1}^{q_0+}, \boldsymbol{\varepsilon}_{it+1}, \boldsymbol{\sigma}(\mathbf{X}_{t+1}^{q_0+}, \boldsymbol{\varepsilon})) \Big) dF^X(\mathbf{X}_{t+1}^{q_0+}|\mathbf{X}_t) \Big] \\
= \frac{\partial}{\partial p_{ijt}} \Pi(\mathbf{X}_t, \varepsilon_{ijt}, \boldsymbol{\sigma}(\mathbf{X}_t, \boldsymbol{\varepsilon}_t)) + \rho \sum_{k \in \{0, C_{mt}\}} \frac{\partial \delta_k(\mathbf{X}_t, \varepsilon_{it})}{\partial p_j} \int \left(EV(\mathbf{X}_{t+1}^{q_k+}) - EV(\mathbf{X}_{t+1}^{q_0+}) \right) dF^X(\mathbf{X}_{t+1}^{q_k+}|\mathbf{X}_t)
\end{aligned} \tag{27}$$

[10]Note that $\frac{dq_{mkt}}{d\delta_{ijt}}$ is equal to one for $k = j$ for installation i in market m, and zero otherwise.

The only link between current pricing and future value functions is dependent on the expected evolution of the market under the different alternative scenarios in which any installer in the market might receive the current installation, and how these probabilities change with the firm's pricing decision as determined by the derivative of demand for each installer in the market with respect to p_{ijt}. Thus, the dynamic pricing incentive is completely captured through the difference in the expected future values under these alternative scenarios. Because we assume independence in the non-hardware costs shocks across installations (after controlling for market and time fixed effects), the only dependency on the future valuation term is through the demand elasticity. Intuitively, the tradeoff from lowering price is between the decrease in profits today and the increase in the number of ongoing installations for the firm (decreasing those for its competitors), which in the short term may have positive or negative effects through the new ongoing installations (due to economies of scale or capacity constraints) and in the long term, i.e., once the contracts are completed, will impact firm profits through the installed base variables, which may affect both consumer utility and the non-hardware costs through LBD.

Note we have that:

$$\frac{d}{dp_{ijt}}\Pi(\mathbf{X}_t) = \delta_{ijt}(\mathbf{X}_t, \boldsymbol{\varepsilon}_{it})S_{it} + \left(p^*_{ijt} - c_{ijt} - w_{ijt}\right) S_{it}\frac{\partial \delta_{ijt}(\mathbf{X}_t, \boldsymbol{\varepsilon}_{it})}{\partial p_{ijt}}.$$

Thus we can write an installers optimal pricing equation as:

$$p^*_{ijt} = \underbrace{c_{ijt}}_{\substack{\text{hardware}\\ \text{costs}}} + \underbrace{w_{ijt}}_{\substack{\text{non-hardware}\\ \text{costs}}} \underbrace{-\frac{\delta_{ijt}(\mathbf{X}_t, \boldsymbol{\varepsilon}_{it})}{\frac{\partial \delta_{ijt}}{\partial p_{ijt}}}}_{\text{static markup}} \tag{28}$$

$$\underbrace{-\rho\frac{1}{S_{it}}\sum_{k\in\{0,\,C_{mt}\}} h_k(\boldsymbol{X}_t, \boldsymbol{\varepsilon}_{it})\int \left(EV(\mathbf{X}^{q_k+}_{t+1}) - EV(\mathbf{X}^{q_0+}_{t+1})\right) dF^X(\mathbf{X}^{q_0+}_{t+1}|\mathbf{X}_t)}_{\text{dynamic markup}}. \tag{29}$$

in which we define the function $h_k()$ as the ratio of the first derivatives of demand for the installation performed by installer k (where $k = 0$ indicates the outside option) with respect to the price set by installer j:

$$h_k(\boldsymbol{X}_t, \boldsymbol{\varepsilon}_{it}) \equiv \frac{\frac{\partial \delta_{ikt}(X_t, \boldsymbol{p}_{it})}{\partial p_{ijt}}}{\frac{\partial \delta_{ijt}(X_t, \boldsymbol{p}_{it})}{\partial p_{ijt}}}. \tag{30}$$

This expression is largely determined by the effect of own-price on own-demand probability, although this expression does account for the cross-derivatives as well. For the focal installer j, this ratio is equal to one.

To integrate over the future states, we can forward simulate once we estimate the state transition probabilities in a first stage. Because of the massive size of the state space (since every installer's installed base and ongoing contracts are relevant for every other installer), instead of calculating the value function for every possible combination of state variables, we instead perform a set of forward simulations starting from every observation in the data. In addition, we forward simulate the entire market (starting from any given observed installation) under each counterfactual scenario in which a competing installer in that market (or the outside option) receives the installation instead. As the focal installer changes the price, the probabilities of these counterfactuals occurring change as well, determined by $\frac{\partial \boldsymbol{\delta}_{it}}{\partial p_{ijt}}$.

The strategy of performing simulations starting from the realizations of the state vector captured in each observation makes estimation much more tractable, but it is still not feasible if this very large set of simulations has to be performed many times during estimation. Thus, we also linearize the non-hardware cost in the parameters (similar to BBL's linearization of the profit function) so the simulations only have to be performed once:

$$w(S_{it}, q_{mjt}, e_{ijt}(\mathbf{b_{ijt}}), W_{mt})) = \Psi(S_{it}, q_{mjt}, \mathbf{b_{ijt}}, W_{mt}) \cdot \beta + \varepsilon_{ijt}, \tag{31}$$

in which $\Psi(S_{it}, q_{mjt}, \mathbf{b_{ijt}}, W_{mt})$ is an vector of basis functions. By substitution we have:

$$
\begin{aligned}
EV(\mathbf{X}_t) &= \sum_{\tau=t}^{\infty} \rho^{(\tau-t)} \mathbb{E}\left[\sum_{m=1}^{M} \sum_{i=1}^{Q_{mt}} ((p_{ijt} - c_{ijt})S_{it})\delta_{ijt}|\right] \\
&- \sum_{\tau=t}^{\infty} \rho^{(\tau-t)} \mathbb{E}_t\left[\sum_{m=1}^{M} \sum_{i=1}^{Q_{mt}} \Psi(S_{it}, q_{mjt}, \mathbf{b_{ijt}}, W_{mt}) S_{it}\delta_{ijt}\right] \cdot \beta - FC_j + \varepsilon_{ijt}. \qquad (32)
\end{aligned}
$$

Thus we can rewrite an installers optimal pricing equation as:

$$
p^*_{ijt} = \underbrace{c_{ijt}}_{\substack{\text{hardware}\\\text{costs}}} + \underbrace{w_{ijt}}_{\substack{\text{non-hardware}\\\text{costs}}} \underbrace{-\frac{\delta_{ijt}}{\frac{\partial \delta_{ijt}}{\partial p_{ijt}}}}_{\text{static markup}} \qquad (33)
$$

$$
\underbrace{-\rho \frac{1}{S_{it}} \sum_{k \in \{0,\, C_{mt}\}} h_k(\boldsymbol{X}_t, \boldsymbol{\varepsilon}_{it}) \int \left(EV(\mathbf{X}^{q_k+}_{t+1}) - EV(\mathbf{X}^{q_0+}_{t+1})\right) dF^X(\mathbf{X}_{t+1}|\mathbf{X}_t, \sigma(\mathbf{X}_t))}_{\text{dynamic markup}} + \tfrac{1}{S}\varepsilon_{ijt}. \qquad (34)
$$

As previously discussed, the primary contributor to the future value term is through the effect or price on future demand for the focal installer. The derivative of own-demand with respect to price enters both the numerator and denominator, and so we have $h_j(\boldsymbol{X}_t, \boldsymbol{\varepsilon}_{it}) = 1$, which does not depend on the $\boldsymbol{\varepsilon}_{it}$; the unobservable in the dynamic markup term only affects $\boldsymbol{h}(.)$ through the cross-derivatives. With higher cost shocks, the probability of all installers receiving the installation goes down, and the probability of the focal installer receiving it goes up.

Taking the first order Taylor expansion, let us write

$$
h_k(\boldsymbol{X}_t, \boldsymbol{\varepsilon}_{it}) = h(\boldsymbol{X}_t, \boldsymbol{p}_{mt}, \nu_{ijt}) \approx \bar{h}(\boldsymbol{X}_t, \boldsymbol{p}_{mt}) + \frac{\partial \bar{h}(\boldsymbol{X}_t, \boldsymbol{p}_{mt})}{\partial p_{jm}} \nu \qquad (35)
$$

in which $\bar{h}(\boldsymbol{X}_t, \boldsymbol{p}_{mt}) \equiv h(\boldsymbol{X}_t, \boldsymbol{p}_{mt}, 0)$ in which the deviation between any installation's optimal price p_{ijt} and the average price for that installer in that market in quarter t is given by ν_{ijt}, which is due both to ε_{it} and to differences in the state variables for that installation relative to the average by that installer in that market at that time.

Rewriting (33), we have:

$$p^*_{ijt} = \underbrace{c_{ijt}}_{\substack{\text{hardware}\\\text{costs}}} + \underbrace{w_{ijt}}_{\substack{\text{non-hardware}\\\text{costs}}} \underbrace{-\frac{\delta_{ijt}}{\frac{\partial \delta_{ijt}}{\partial p_{ijt}}}}_{\text{static markup}} \tag{36}$$

$$\underbrace{-\rho\frac{1}{S_{it}}\sum_{k\in\{0,\,C_{mt}\}} \bar{h}(\boldsymbol{X}_t,\boldsymbol{p}_{mt})\int \left(EV(\mathbf{X}^{q_k+}_{t+1}) - EV(\mathbf{X}^{q_0+}_{t+1})\right) dF^X(\mathbf{X}_{t+1}|\mathbf{X}_t,\sigma(\mathbf{X}_t))\,.}_{\text{dynamic markup}}$$

$$\underbrace{-\rho\frac{1}{S_{it}}\sum_{k\in\{0,\,C_{mt}\}} \frac{\partial\bar{h}(\boldsymbol{X}_t,\boldsymbol{p}_{mt})}{\partial p_{jm}}(\boldsymbol{X}_t,\boldsymbol{p}_{mt})\int \left(EV(\mathbf{X}^{q_k+}_{t+1}) - EV(\mathbf{X}^{q_0+}_{t+1})\right) dF^X(\mathbf{X}_{t+1}|\mathbf{X}_t,\sigma(\mathbf{X}_t))\nu_{ijt} + \frac{1}{S}\varepsilon_{ijt}\,.}_{\text{error term}}$$

Although the structural cost shock $\boldsymbol{\varepsilon}_{it}$ leads to a difference in the average derivatives of demand for other installers with respect to installer j's price of installation i, its effect leads only to further the heteroskedaticity of the error.

5.2 Estimation

With our demand estimates in hand, the first step in the supply estimation is estimating the state transitions and policy functions. The estimation of the state transition probabilities are straightforward, and the structure we imposed on these transitions is theoretically motivated. For the policy function, we use the following flexible form:

$$\log(p_{ijt}) \quad = \quad (X^p \otimes X^p)\kappa + \xi_j + \eta_t + \epsilon_{ijt} \tag{37}$$

$$X^p \quad \equiv \quad \{\log(p_{ijt-1}),\, S_{ijt},\, q_{mjt},\, \sum_{k\neq j\in M} q_{mkt},\, R_{mt},\, \boldsymbol{b}_{ijt}\} \tag{38}$$

where $\otimes$ indicates the Kronecker product.

In simulating future prices, we model the evolution of the time fixed effects using an AR(1) process, and we assume that the household-specific unobserved shocks are common to installers (due to things like steeper roofs leading higher installation costs). In simulating future shocks, we use the standard deviation of the residuals from the estimation of the AR(1) process and assume normality.

With first-stage estimates of the transition probabilities and the parameters that govern the policy function in hand, we can use forward simulations for many possible realizations of all outcomes in future periods, as done in BBL, only in our context, we forward simulate for each household in the data to calculate the expected profits (not including the non-hardware costs) over time, where the expectation is taken over all of the installers who might get that household's installation. We can similarly calculate a term that represents the expected NPV of the non-hardware costs conditional on the learning parameters β^b, in which β^b enters multiplicatively, as shown in equation (36). This means that the forward simulation only has to occur once, although it must be done separately for every observation in the data since the pricing of any installation has downstream consequences for all installers in all markets.

The difference in the simulated future valuations and costs if installer k were to get the installation versus no installer getting the installation are then multiplied by $h_k(\boldsymbol{X}_t, \boldsymbol{p}_m t)$ and summed together. Multiplying by $-\rho\frac{1}{S_{it}}$ gives us the estimate of the dynamic pricing term. The intuition for why inverse size enters into the term is that for bigger installations, the profit sacrifice of lowering price does not justify the added learning as much as it does for smaller installations.

The entire estimation procedure is as follows:

1. Estimate the demand model to calculate $\widehat{\delta}_{ikt}$ as a function of the states, $\mathbf{X}_t$.

2. Calculate $\widehat{\sigma}_{jt}$ and transition probabilities $\widehat{f}(.)$.

3. For each installation i:

   (a) Draw the random shocks for this simulation, rs, that will determine the evolution of the state variables for the next period.

(b) For each firm in market m at time t and for the outside option of no installation, assign the installation to that firm and update $\mathbf{q}_t$.

(c) Simulate the realization at $t+1$ for the state variables under each possible assignment of i to the installers in the market ($X^{j_m s}$), using the same random shocks for each.

(d) Calculate optimal prices $\widehat{\sigma}_{jt}(X^{j_m s})$ for all installers in all markets for each scenario.

(e) Calculate $\sum_{i=1}^{Q_{mt}}((p_{ijt}-c_{ijt})S_{it})\delta_{ij\tau}$ and $\sum_{i=1}^{Q_{mt}}\mathbf{b}_{ijt}S_{it}\delta_{ij\tau}\;\forall\,\tau>t$ for the simulation path, rs, $\forall k \in M_i$ who might have gotten installation.

(f) Repeat the last three steps for T^{sim} periods.

(g) Compute the NPV of the two values in step (e) and take the weighed sum over the potential installers (as well as the outside option) who might have gotten the installation using weights $\frac{\frac{\partial \delta_{ikt}}{\partial p_{ijt}}}{\frac{\partial \delta_{ijt}}{\partial p_{ijt}}}$ for $k \in \{0,\, C_{mt}\}$.

(h) Repeat for RS simulations.

For each installer who is assigned the installation (including the outside option of the installation not occurring), we simulate ten paths of the market transitions over 20 years, the standard life of a solar panel system (using the same unobservable shocks for the assignment of the installation to each installer). For example, if we have an installation with 19 active installers, there are 20 possible assignments of the installation including the no install option, and we perform 200 forward simulations from that observation to calculate the value function.

6 Identification and Data

6.1 Identification

Our identification strategy depends on our ability to both separately estimate the static and dynamic markup in order to isolate the non-hardware costs. For the former, we rely on calculations of the static markup that we get from estimating the dynamic demand model. For the dynamic markup, we are able to directly calculate the dynamic pricing incentive using the first order condition, using forward simulations as in BBL. By forward simulating states of the world from multiple starting points in which every competing installer is assigned that observation's installation, we can then incorporate the change in future value in the focal installer's first-order pricing equation by accounting for the effect of price on the likelihood of which installer gets the installation (if any), and thus the likelihood that each set of forward simulations are indicative of the market's evolution. Forward simulations starting from each observation in the data allow us to capture the effect of the pricing decision on the *entire* market in which there are over 3000 installers. Identification is aided by the fact that larger size installations have less of a dynamic pricing incentive, because the value of learning is smaller relative to the profits from the current installation.

One may be concerned about serial correlation leading to endogeneity due to a correlation between our installed base variables and the error. This is less of an issue in our setting for there is on average a six month lag between when an application for an installation is submitted (i.e., when the sale is made) and when the installation is completed. Thus, any serial correlation would have to be quite substantial. Examining the Durbin-Watson test statistic, we find serial correlation of only a few months, suggesting that this is not a concern.

Fundamentally, our coefficients of interest are identified from within-installer, within-county, and within-quarter variation in the installed base variables and the BOS across installations of different sizes. We can separately identify the effect of experience from economies of scale through the differing variation in the installed base variables and the on-going contracts variables.

6.2 Data

Our dataset, compiled by Lawrence Berkeley National Laboratory, is unique in that it includes both the price and the hardware costs for most of installations in California through 2012. Our data includes all installations in California that received an incentive payment. For the three investor-owned utilities, it covers both the earlier Emerging Renewables Program and the CSI. It also covers all municipal utility solar incentive programs. The data includes the type of installation (residential, commercial, government or nonprofit), price and size of the installation, whether the system is third-party owned or appraised value, the module and inverter costs, any financial incentives, PV installer and manufacturer information, average electricity rate in the zip code of the installation, and zip code of the installation.

The raw dataset has 135,654 observations. We include all of these installations in creating the installed base and ongoing installations variables. For 34,148 observations, the cost or price data appear unreliable, with reported prices of more than \$12/kW or less than \$1/kW or hardware costs of more than \$8/kW or less than \$0.30kW. Many of these appear to be extra zeros added or removed by the installer when the installation was reported, but it appeared difficult to correct them, and thus we opted to drop these observations. We focus on non-utility scale installations (removing the 100 installations greater in size than 100 kW) and we drop the 95 ground-mounted systems as well. Finally, for 26 installations have no wage data. This leaves us with 979,709 installations. 38 of these have no installer information and 11,623 only have an appraised rather than a transacted price, which we drop as well. This leaves 88,048 installations. Of these, we can only estimate the static markup for 76,838 (due to limited installations for some of the very small installers) which is our final estimation sample.

Table 1 provides summary statistics for the key variables in our final dataset. All dollar-valued variables are converted to real 2012 dollars. As we hypothesize that LBD occurs with installations at different levels, we create the installed base variables which are the cumulative number of installations by a specific contractor and/or in a specific county.[11]

The installed base variables are calculated first for a given installer at both the California-wide level and at the county level, assuming a continuous decay rate equivalent to an annual decay rate of 11%, as found by Benkard (2004) and Kellogg (2011). Then, since LBD spillovers are most likely to occur between competing contractors, we create a variable for the cumulative installations by the contractor's competitors within the county. To control for potential economies of scale or capacity constraints, we also create a variable for the contractor's on-going contracts, which is defined as the number of contacts that are in-progress between the contract signing and the actual installation.[12]

Since the dataset contains both residential and non-residential installations, we provide summary statistics for key variables in each of these categories respectively in Tables 2 and 3. Most of the observations are residential systems, with only 2,396 non-residential systems. Residential systems tend to be significantly smaller, with a mean size of 5.5 kW versus 20.38 kW, but slightly less expensive per watt on average than non-residential systems ($7.27 per W versus $7.46 per W).

We plot the number of installations over time in Figure 2. During the period of the CSI, we see rapid acceleration of solar PV adoption. This is not surprising; since 2002, the average installation price has declined from approximately ten dollars per watt to under six dollars per watt, with much of this decline occurring after 2009 (Figure 3).[13] This figure shows that while average solar PV prices have declined along with the hard costs, the BOS has not decreased as much as one might expect LBD were greatly lowering

[11] We also code up every merger and acquisition in the California solar PV market so that we can include the experience of both firms when they merge – our results are robust to the alternative assumption that learning is not transferred.

[12] The average time between signing of the contract and installation is roughly 120 days.

[13] For reference, we can compare the levelized cost (i.e., the present value cost of owning and operating the generation asset) of solar to other electricity generation sources. We assume a 30 year solar system lifespan, a 30 year mortgage with an interest rate of 3%, an inverter lifespan of 8 years, solar PV system output from Borenstein (2008), limited losses from soiling, and a PV panel decay for multi-crystalline silicon panels of 0.5% corresponding to the best available evidence (Osterwald, Adelstein, del Cueto, Kroposki, Trudell, and Moriarty 2006). Our calculations suggest that the 2009 residential system average cost of $8 per DC W corresponds to a levelized cost of roughly $0.30-$0.35 per kWh before any incentives, whereas centrally generated electricity sources, such as coal or natural gas had a 2009 levelized cost in the range of $0.05-$0.07. The cost of solar has dropped substantially since then.

non-hardware costs.

One possible justification for the lack of a large drop in BOS is simply that there is no learning. However, other factors must be accounted for before coming to this conclusion. For example, in 2008, there is is a large increase in third party systems. Figure 4 plots BOS for all systems as well as for just owned systems, and there is a drop of just over 1$/W in BOS for owned systems between 2008 and 2013. Third party systems are recorded as more expensive per watt, and so their increase in market share hides the BOS declines that are happening concurrently. Another explanation for less of a decline than expected with LBD is the competitive landscape.

Figure 5 provides a histogram after removing contractors who perform less than 10 installations and shows that most firms in this market still fall into the competitive fringe. This competitive fringe installs the majority of solar PV systems, but 31.2 percent of systems are installed by the top 10 installers (over the full time period of the data), so there is still significant concentration in the market. Table 4 provides summary statistics over the 3,017 installers that appear at any point in the dataset (there were 21 in 1998, 353 at the start of our panel in 2002, and 790 in 2012). On average, contractors operate in 2.4 counties and have performed 33.2 installations. As is clear in Figure 5, the distribution of installer size is very skewed with a dozen or so very large firms and a huge tail of tiny installers.

Another potential explanation for BOS not dropping as much as expected with LBD is the dynamic pricing incentive for firms. In the model, we showed that this incentive is larger for smaller installations than for large installations. We therefore plot BOS versus installation size in Figure 6. If firms are pricing dynamically, we would expect to see BOS decline more for large installations than for small installations, since in the early years before firms move down the learning curve, installers have an incentive to lower price in order to perform an installation and move father does the learning curve, but the profit reduction for the current installation is large if the installation is large. As is clear in the figure, the larger installations (with smaller relative dynamic pricing terms) exhibit larger

BOS declines over time.

7 Estimation and Results

7.1 Demand Estimates

For the demand model, we calculate market shares by collapsing the dataset so that the unit of observation is an installer-county-quarter. We calculate the share of new contracts for each installer at this unit of observation. For our model, we need an estimate of the potential market size. We begin with the number of owner-occupied homes and businesses, the latter taken from the UC Census County Business Patterns for 2012. To determine the fraction of potential adopters who would make up the relevant market, we use Google sunroof data to calculate the share of buildings for which adopting solar would lead to a positive net present value (using Google's assumed discount rate and current electricity prices). The share of buildings suited for solar ranges between 40% and 100% of the market.[14] Summary statistics for this dataset are shown in Table 5.

We estimate the demand model using both OLS and instrumental variable regression. We present our regression results in Table 6. We start with a static OLS regression to provide a benchmark. The price coefficient is negative and significant as expected. Both of the installer's installed base variables increase consumer utility of an installation, and the nest parameter is 0.726. We find positive and significant effects of monthly solar radiation as expected, but no statistically significant effect of electricity rates, likely due to the fact that much of the variation in rates is cross sectional and thus absorbed by the county fixed effects. Upon instrumenting for price using the exogenous rebate schedule, the estimated price coefficient is much larger, in magnitude. We also estimate positive effects of own-installations within the county on consumer utility.

These static estimates ignore the fact that consumers are forward-looking with respect to price and the other state variables. Columns 3 through 5 allow consumers to

[14] For the small number of counties without these data, we use 50%.

be forward-looking. Our preferred specification is column 5, in which we allow for forward looking behavior and instrument for both price and the (mechanically) endogenous within-group share. There is a positive effect of own-installations within the county and no effect of installations outside the county. In addition, there is a positive and significant coefficient on the average quarterly radiation.

We can estimate the demand elasticities as follows:

$$\begin{aligned} elast_{jt} &= \alpha\frac{1}{1-\sigma}(1-\sigma s_{j/It}-(1-\sigma)s_{jt})P_{mjt} \\ elast_{j/It} &= \alpha\frac{1}{1-\sigma}(1-s_{j/It})P_{mjt} \end{aligned} \tag{39}$$

in which s_{jt} is the share of consumers installing from installer j, $s_{j/It}$ is the share of consumers installing from installer j conditional on installing solar, and as before, P_{mjt} is the average price for installations performed by j in market m at time t.

We plot the estimated elasticity over time in Figure 7 where the unit of observation is the installer-county-quarter. We include fractional polynomial best fit lines as well. We break down the elasticity into the group elasticity and the within-group elasticity. The group-level elasticity is consumer's response to solar price on the decision to install at all, irrespective of which installer is chosen. The average group elasticity first decreases (becomes more elastic) and then increases (becomes less elastic), with a low of -1.2 and a high of -0.8. The group elasticities in the literature can be compared to others in the literature. Hughes and Podolefsky (2015) find an elasticity of -1.2 for CA, Gillingham and Tsvetanov (2017) find an elasticity of -0.65 for CT, and Rogers and Sexton (2014) find a rebate elasticity of 0.4 for CA. None of these papers allow for dynamic demand.

The second figure shows the elasticity conditional on installing reaches its largest magnitude in 2005. The installer elasticity is much higher than the group elasticity, as we would expect, with the annual average ranging between -4.0 and -2.5. Again we see the U-shaped pattern over time. After 2009, a handful of large installers see greatly increasing installed bases, decreasing the consumer elasticity and leading to more market power.

One thing that is notable in both graphs is the increased variation in these elasticities over time, which is indicative of increasing asymmetry in market power between large and small firms.

We can demonstrate this in a bin-scatter plot of the optimal static markup for the four quantile ranges of an installer's own installed base (Figure 8). Although in general the estimated static markups decline over time, the largest quantile of firms see their markups actually increase after 2009. We see the same pattern when we split firms by quantiles within each county. These results demonstrate the need to control for the changing market power of firms in assessing LBD. Time fixed effects can capture changing markups only when they change for all installers by the same amount, but this does not appear to be the case in our setting. The asymmetry in the trends, and specifically the increase in markups for the largest firms, helps to explain the discrepancy in the BOS trends in the raw data, which show BOS declining for the smaller firms over time but increasing for the largest firms.

To test to see whether unobserved, within-market heterogeneity needs to be accounted for, in Appendix A, we allow for an evolving distribution of heterogeneity as high type consumers adopt early and leave the market. Results are shown in Table A.1. We find that the results are largely unaffected by the inclusions of this added heterogeneity, largely because only a small fraction of potential adopters have adopted by the end of our panel.

We find that installers that move farther down the learning curve continue to price as if they had not. Investor reports for SolarCity confirm this story. For example, a 2013 SolarCity Investor presentation discusses the value creation due to the fact that they outpaced their cost reduction targets and experienced a reduction in labor hours for installations, in conjunction with the expanding size of the market. Profits increased between 2012 and 2013 from \$27.5 to \$39.4 million (21.6% of revenues to 24.0% of revenues).

7.2 Supply-Side Estimates

We estimate the supply-side pricing decision using equation (36) and show the results in Table 7. We cluster standard errors at the county level. We also include county and installer fixed effects, and thus inference results from within-installer and within-county variation over time. We begin in columns 1-3 with no controls for either the static markup or dynamic pricing incentives. In columns 4-6, we include the static markup control, and in columns 7-9, we include the controls for dynamic pricing. In each set of specifications, we start with a quadratic model, including the rebate amount and statewide installations as regressors. We then replace the rebate amount with utility x quarter fixed effects, and finally add the installed base interaction terms, which are necessary if learning from one installed base is a substitute or complement for learning through another. The quadratic specification with interactions allows for the learning based on the installed base to occur in a highly flexible manner.

Without controlling for the changing static markup, we see no effect of the installed base variables on non-hardware costs (columns 1-3), with the exception of a positive quadratic term for own installations within the county and an increasing effect of statewide installations on costs reductions. In the column (1) results, we also find that higher rebates lead to higher demand, as expected. The statewide installed base effect and rebate effect cannot be identified with the inclusion of utility-quarter fixed effects in columns (2) and (3).

When the effect of installed base on on markups is accounted for in columns 4-6, we find evidence supportive of county-level, appropriable LBD. Looking at the column (5) results, with no installations we find that for every 1000 installations the installer performs in the county, costs decline by \$0.95 per Watt. This marginal effect declines as installed base increases. Accounting for the quadratic term, we for every 1000 installations the installer performs in the county, costs decline on average by \$0.40-\$0.44 per Watt. We find similar average effects when we allow for installed base interactions in column (6). There is notably a significant, negative interaction effect between own-installations in the county and competitor installations within the county, indicating that firms with more own-experience are also affected more by learning spillovers. The significant positive effects between own installations inside and outside the county indicate that these installations serve as substitutes in their contribution to LBD. The combined effect of the four terms that include own installations within the county (installer installed base, installer installed base squared, the interaction between installer and competitor installations within the county, and the interaction between the installers installations within and outside the county) is significant at 5% (using an F-test of joint significance). The combined effect of the three terms that include own installations outside the county is also significant at 5%.

Focusing on the effect of competitor installations within the county, we find that the total effect of the three terms (competitor installations within the county, competitor installations within the county squared, and the interaction with the installer's own installations) is significant at 10%. Thus we find evidence for learning spillovers, albeit fairly weak evidence given the small magnitude of the effect. The joint significant of all the installed base coefficients is also significant at 5%. The results are largely unaffected when accounting for the dynamic pricing incentives in columns 7-9. The fact that the dynamic pricing incentive does not substantially alter the estimates of appropriable LBD is not surprising. We find that the dynamic pricing incentive leads to maximum price changes of less than three cents per Watt, (i.e. less than $150 for a typical installation); this is likely because we find that appropriable LBD is small.

To get a better sense of the magnitudes of learning, we plot the estimated installed base effects for the observations in the data over time in Figure 9, for the six specifications that include the utility x quarter fixed effects. The solid lines indicate the average learning over time and the dotted lines indicate the 95th percent interval of learning across all observations in the data (using the point estimates). Without accounting for the dynamic markup, the total learning that occurs results in a price decline of approximately $0.12/W when using the more flexible quadratic with interactions. Thus, although we do find suggestive evidence of appropriable LBD, it is relatively small in magnitude, which may not be too surprising given the modest decline in BOS.

In Figure 10, we also plot the utility x quarter fixed effects for column (9), in which both markups are accounted for and we use the quadratic function of installed bases. As discussed, LBD that happens at the utility, state, or national level cannot be separately identified from other secular trends (other than with functional form assumptions). The PG&E fixed effects decline over time, whereas the SDG&E and SCE trends increase and then decrease, peaking in 2006 for SDG&E and 2009 for SCE. Much of the increase in 2007 can be explained by high silicon prices. It is possible that the overall decline in PG&E may be due to utility-level LBD–we cannot rule this out–but it is just as plausible that these effects would have occurred without the CSI policy.

In all the specifications when accounting for dynamic pricing incentives, we find significant local economies of scale, which are often confounded with LBD. Indeed the main effect of accounting for dynamics was the fact that these effects were insignificant without the dynamic pricing term. We estimate that a 20% increase in ongoing installations in the county leads to a non-hardware cost decline of one cent per watt. However, the economic benefits that result these economies scale should be internalized by the installer through its pricing decisions, and thus they also do not justify the large incentives on economy efficiency grounds. However, they do lead to larger environmental benefits that result from the CSI incentives than those that would have accrues without economies of scale.

7.3 Robustness Checks

To test the robustness of our findings, we re-estimate the model assuming no depreciation in installed bases over time. Demand estimates under this assumption are in Appendix C, in Table C.2 and the supply-side estimates in Table C.3. We also estimate the model in which installed base effects only occur through their interaction with the roofing wage rates, which assumes all learning is in labor costs. These results are in Table C.4. Finally, we estimate the model under the alternative assumptions that learning does not transfer with firm acquisition and that there is no learning depreciation (Table C.5).

Under all alternatives, we find evidence of small LBD, which reduces installation costs from between \$0.10 and \$0.20 per Watt when accounting for the static markup or both markups. As with our main results, it is critical to account for the changing static markup over time, otherwise an incorrect conclusion of no LBD would be reached. Accounting for the dynamic pricing incentive changes the magnitude of the estimates but not the qualitative conclusions.

7.4 Welfare effects of the CSI

We run a simple counterfactual scenario in which we remove the CSI incentives, assuming 100% pass-through, consistent with the findings in Pless and Benthem (2018). We find that removing the CSI leads to reductions in the number of installations of 22% in 2007, increasing to 28% in 2012, as shown in Figure 11. This is due to the effect not only on price, but on the local installed bases which affect consumer utility for solar.[15] The installed base at the end of 2012 is reduced from 110 thousand installations to 84.7 thousand installations, a drop of 23%. The increased consumer surplus as a result of the CSI is calculated to be \$491 million, in comparison to the costs of \$3.3 billion.

We also calculate the avoided environmental external costs, using county-level estimates of averted environmental damages in California, which we construct using the population-weighted zip code values calculated by Sexton, Kirkpatrick, Harris, and Muller (2018). The environmental benefits of new installations are plotted in Figure 12. The NPV of these averted damages (assuming similar adoptions rates under both regimes after 2012 once the subsidies disappear, which is approximately the case in years 2011 and 2012)) is \$875 million using an annual discount rate of 13% which we used for the installers, based on the findings of De Groote and Verboven (2016). With a 5% discount rate, the NPV of the averted damages is \$2.29 billion. Part of the reason the averted damages are not higher is that CA has a relatively clean energy mix, in comparison to other

[15] We assume invariance of the policy function that determines pre-incentive prices to this counterfactual environment, which we would expect to approximately hold with full pass through.

regions of the United States where the averted damages would be higher.

Even with the low level of discounting, the combined welfare benefits from the increase in consumer surplus and the avoided environmental damages are lower than the costs of the CSI program. If the LBD spillovers were larger, then the cost of the CSI would be better justified. Further, if the time fixed effects that we estimate are due to non-localized learning (i.e. learning at the state level) which would not have occurred in the absence of the CSI, then this would lead to greater estimates of both the consumer surplus and environmental benefits.

There may also be longer-run altruistic motivations. For example, Gerarden (2018) argues that from a global perspective, subsidies in individual regions can help foster innovation in panel manufacturing. Our analysis focuses on localized learning, as this was a major motivation for the CSI, but does not examine such broader innovation effects. However, our quantification of learning spillovers in the solar market is important for informing policymakers about the full effects of technology-oriented policies.

8 Conclusions

This paper develops a model of solar PV installer pricing to examine evidence for both appropriable LBD and non-appropriable LBD in the California solar PV market. We leverage a rich dataset of solar installations in California from 2002 to 2012 and develop a model of both dynamic supply and demand for solar installations in the California small-scale solar market. Our approach accounts for changing market power, economies of scale, capacity constraints, and firm dynamic pricing incentives. Disentangling these factors is particularly important in our setting for estimating localized learning in the non-hardware costs of a solar installation, which are combined with the markup in our data.

The results of our dynamic model of demand indicate that the markup is declining over time as the market has grown, but not for the largest installers. This is important for it immediately helps to explain why it appears that the BOS has not been declining much over time despite a decline in overall installed prices. It thus follows intuitively that our supply model provides evidence of LBD, albeit small LBD. The overall LBD we find is about $0.12 per watt (out of an average BOS of about $2.50 per watt). Following standard learning curves, our results suggest greater learning in the beginning of our time period and lesser learning later in the time period. Our results also provide evidence of learning spillovers from competitors to an individual firm. Perhaps the most interesting spillover coefficient suggests that firms with the largest cumulative number of installations showed even greater cost declines with competitor experience, suggesting that larger firms are better able to appropriate some of the learning from their competitors. This may occur from factors such as hiring of employees from competitors or watching how competitors install systems.

However, by running an illustrative counterfactual, we find that the CSI is very likely to be reducing short-run economic efficiency, even after accounting for the positive externality from the learning spillovers and environmental externalities. However, without learning, the finding would be even more stark and it would be extremely difficult to justify the CSI on economic efficiency grounds.

References

ALCHIAN, A. (1963): "Reliability of Progress Curves in Airframe Production," *Econometrica*, 31, 679–693.

ARCIDIACONO, P., AND P. B. ELLICKSON (2011): "Practical Methods for Estimation of Dynamic Discrete Choice Models," *Annual Review of Economics*, 3, 363–94.

ARCIDIACONO, P., AND R. A. MILLER (2011): "Conditional Choice Probabilty Estimation of Dynamic Discrete Choice Models with Unobserved Heterogeneity," *Econometrica, Vol.*, 79(6), 18231867.

ARROW, K. (1962): "The Economic Implications of Learning by Doing," *Review of Economic Studies*, 29, 155–173.

BAJARI, P., C. L. BENKARD, AND J. LEVIN (2007): "Estimating Dynamic Models of Imperfect Competition," *Econometrica*, 75(5), 1331–1370.

BARBOSE, G., N. DARGHOUTH, S. WEAVER, AND R. WISER (2013): "Tracking the Sun VI: A Historical Summary of the Installed Price of Photovoltaics in the United States from 1998 to 2012," *LBNL Working Paper 6350E. Lawrence Berkeley National Laboratory*.

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

BERRY, S. T. (1994): "Estimating Discrete-Choice Models of Product Differentiation," *The RAND Journal of Economics*, 25(2), 242–262.

BESANKO, D., U. DORASZELSKI, Y. KRYUKOV, AND M. SATTERTHWAITE (2010): "Leanrning-by-Doing, Organizational Forgetting, and Industry Dynamics," *Econometrica*, 78(2), 453–508.

BOLLINGER, B. (2015): "Green technology adoption: An empirical study of the Southern California garment cleaning industry," *Quantitative Marketing and Economics*, 13(4), 319358.

BORENSTEIN, S. (2008): "The Market Value and Cost of Solar Photovoltaic Electricity Production," *University of California Energy Institute, CSEM Working Paper 176*.

BURR, C. T. (2014): "Subsidies, Tariffs and Investment in the Solar Power Market," *University of Colorado-Boulder Working Paper*.

CABRAL, L., AND M. RIORDAN (1994): "The Learning Curve, Market Dominance, and Predatory Pricing," *Econometrica*, 62(5), 1115–1140.

CARDELL, N. S. (1997): "Variance components structures for the extreme-value and logistic distributions with application to models of heterogeneity," *Econometric Theory*, 13(2), 185–213.

COVERT, T. (2014): "Experiential and Social Learning in Firms: The Case of Hydraulic Fracturing in the Bakken Shale," *University of Chicago Working Paper*.

CPUC (2009): "California Solar Initiative Annual Program Assessment," *California Public Utilities Commission Go Solar California Report*.

DE GROOTE, O., AND F. VERBOVEN (2016): "Subsidies and myopia in technology adoption: evidence from solar photovoltaic systems," working paper.

DERDENGER, T., AND V. KUMAR (2015): "A CCP Estimator for Dynamic Discrete Choice Models with Aggregate Data," .

FOSTER, A., AND M. ROSENZWEIG (1995): "Learning by Doing and Learning from Others: Human Capital and Technical Change in Agriculture," *Journal of Political Economy*, 103(6), 1176–1209.

FUDENBERG, D., AND J. TIROLE (1983): "Learning-by-Doing and Market Performance," *Bell Journal of Economics*, 14, 522–530.

GERARDEN, T. (2018): "Demanding Innovation: The Impact of Consumer Subsidies on Solar Panel Production Costs," *Cornell University Working Paper*.

GHEMAWAT, P., AND A. M. SPENCE (1985): "Learning Curve Spillovers and Market Performance," *The Quarterly Journal of Economics*, 100, 839–852.

GILLINGHAM, K., R. NEWELL, AND W. PIZER (2008): "Modeling Endogenous Technological Change for Climate Policy Analysis," *Energy Economics*, 30(6), 2734–2753.

GILLINGHAM, K., AND J. STOCK (2018): "The Cost of Reducing Greenhouse Gas Emissions," *Journal of Economic Perspectives*, 32(4), 53–72.

GILLINGHAM, K., AND J. SWEENEY (2010): *Harnassing Renewable Energy*chap. Market Failure and the Structure of Externalities, p. Forthcoming. RFF Press.

GILLINGHAM, K., AND T. TSVETANOV (2017): "Estimating Demand for Solar Photovoltaics," .

GRUBB, M., J. KHLER, AND D. ANDERSON (2002): "Induced Technical Change in Energy and Environmental Modeling," *Annual Review of Energy and the Environment*, 27, 271308.

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

HUGHES, J. E., AND M. PODOLEFSKY (2015): "Getting Green with Solar Subsidies: Evidence from the California Solar Initiative,," *Journal of the Association of Environmental and Resource Economists*, 2, 235275.

IEA (2009): "Trends in Photovoltaic Applications: Survey Report of Selected IEA Countries Between 1992 and 2008," *Report IEA-PVPS T1-18*.

IRWIN, D., AND P. KLENOW (1994): "Learning-by-doing Spillovers in the Semiconductor Industry," *Journal of Political Economy*, 102(6), 1200–1227.

JOVANOVIC, B., AND Y. NYARKO (1996): "Learning by Doing and the Choice of Technology," *Econometrica*, 64(6), 1299–1310.

KELLOGG, R. (2011): "Learning by Drilling: Inter-Firm Learning and Relationship Persistence in the Texas Oilpatch," *Quarterly Journal of Economics*, 126(4), 1961–2004.

LAMP, S. (2014): "Projection Bias in Household Investment? The Case of Solar Photovoltaics in Germany," *Yale University Working Paper*.

LEVITT, S., J. LIST, AND C. SYVERSON (2013): "Toward an Understanding of Learning by Doing: Evidence from an Automotive Assembly Plant," *Journal of Political Economy*, 121(4), 643–681.

LIEBERMAN, M. (1984): "The Learning Curve and Pricing in the Chemical Processing Industries," *Rand Journal of Economics*, 15, 213–228.

MELITZ, M. (2005): “When and How Should Infant Industries be Protected?,” *Journal of International Economics*, 66, 177–196.

NORDHAUS, W. (2014): “The Perils of the Learning Model for Modeling Endogenous Technological Change,” 35(1), 1–13.

OSTERWALD, C., J. ADELSTEIN, J. DEL CUETO, B. KROPOSKI, D. TRUDELL, AND T. MORIARTY (2006): “Comparison of Degredation Rates of Individual Modules Held at Maximum Power,” in *Conference Record of the 2006 IEEE 4th World Conference on Photovoltaic Energy Conversion*.

PLESS, J., AND A. V. BENTHEM (2018): “Pass-Through as a Test for Market Power: An Application to Solar Subsidies,” *American Economic Journal: Applied Economics*, p. forthcoming.

ROGERS, E., AND S. SEXTON (2014): “Effectiveness of Subsidies for Residential Rooftop Solar Adoption: Lessons from California,” *North Carolina State University Working Paper*.

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

SEXTON, S. E., A. J. KIRKPATRICK, R. HARRIS, AND N. Z. MULLER (2018): “Heterogeneous Environmental and Grid Benefits from Rooftop Solar and the Costs of Inefficient Siting Decisions,” .

SPENCE, M. (1981): “The Learning Curve and Competition,” *Bell Journal of Economics*, 12, 49–70.

STOKEY, N. (1985): “The Dynamics of Industry-wide Learning,” in *Essays in Honor of Kenneth J. Arrow*, ed. by W. Heller, R. Starr, and D. Starrett. Cambridge University Press, Cambridge, U.K.

——— (1988): “Learning by Doing and the Introduction of New Goods,” *Journal of Political Economy*, 96(4), 701–717.

Taylor, M. (2008): "Beyond Technology-push and Demand-pull: Lessons from California's Solar Policy," *Energy Economics*, 30(6), 2829–2854.

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

Thornton, R. A., and P. 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, A., K. Gillingham, and J. Sweeney (2008): "Learning-by-doing and the California Solar Initiative," *Energy Journal*, 29(3), 131–151.

Young, A. (1991): "Learning by Doing and the Dynamic Effects of International Trade," *Quarterly Journal of Economics*, 106, 369–405.

——— (1993): "Invention and Bounded Learning by Doing," *Journal of Political Economy*, 101(3), 443–472.

Zimmerman, M. (1982): "Learning Effects and the Commercialization of New Energy Technologies: The Case of Nuclear Power," *Bell Journal of Economics*, 13, 297–310.

Table 1: Summary Statistics for California Solar Installations 2002-2012

Variable	Mean	Std. Dev.	Min.	Max.	N
pre-incentive price ($2012/W)	7.237	1.805	1.24	11.996	78875
hardware costs ($2012/W)	4.699	1.589	0.373	11.987	78875
system size (kW$)	6.268	6.644	0.6	99.75	78875
CA installed base (1000s)	31.225	15.733	1.482	64.126	78875
installer installed base (1000s)	0.335	0.61	-0.012	6.467	78875
installer installed base in county(1000s)	0.052	0.093	0	0.868	78875
competitors installed base in county (1000s)	1.664	1.699	0	10.112	78875
installer ongoing contracts in county (1000s)	0.07	0.121	0	1.057	78875
installer ongoing contracts (1000s)	0.349	0.661	0	5.951	78875
HHI in CA	0.036	0.022	0.011	0.323	78875
HHI in county	0.143	0.133	0.028	1	78875
market share in CA	0.022	0.037	0	0.565	78875
market share in county	0.127	0.162	0.001	1	78875
roofing wage rate (1000s 2012$)	39950.356	5936.93	11365.793	53518.617	78875
third party-owned system	0.241	0.428	0	1	78875
appraised value system	0	0	0	0	78875
average electricity rate (2012$)	0.153	0.011	0.095	0.177	78787
Radiation (W/m$\hat{2}$)	5395.802	2079.994	240.503	9534.576	78875
monthly average radiation (W/m-sq)	5460.296	1867.633	1559.458	8457.046	78875
deviation from monthly avg radiation (W/m-sq)	-96.898	329.144	-1346.515	883.237	78875
Census block group percent democrats	0.585	0.124	0.284	0.844	76101
residential rebate per W	1.306	1.128	0	7.778	78875
Zillow housing price index	496855.68	350280.551	57000	3807600	78513

Notes: An observation is an installation. All $ values are in 2012 values. The installed base is the cumulative number of installations by that point in time. The HHI refers to the Herfindahl-Hirschman Index.

Table 2: Installation price and size, residential

Variable	Mean	Std. Dev.	Min.	Max.
size	5.559	3.808	0.6	98.88
price	7.237	1.798	1.24	11.996
hard costs (2012$ per W)	4.699	1.585	0.379	11.987
N		75,653		

Table 3: Installation price and size, non-residential

Variable	Mean	Std. Dev.	Min.	Max.
size	20.23	18.16	1.17	99.75
price	7.37	1.904	1.82	11.985
hard costs (2012$ per W)	4.855	1.674	0.373	10.617
N		2,396		

Table 4: Installations by contractor

Variable	Mean	Std. Dev.	Min.	Max.
Contractor number of installations	26.144	162.963	1	4664
Contractor MW of installations	0.164	0.969	0.001	28.219
Contractor number of counties	2.37	3.508	1	53
N		3,017		

Table 5: Demand Data Summary Statistics

Variable	Mean	Std. Dev.	N
log odds ratio	-7.493	1.539	32162
log within-group share	-3.848	1.515	32167
price ($/W)	5.805	1.45	31716
contractor installed base in county	0.018	0.062	32167
contractor installed base outside county	0.331	0.975	32167
house value ($1000K)	462.508	219.709	31964
electricity rate ($/kWh)	0.151	0.009	32167
average monthly radiation	5.292	1.938	32167

Table 6: Demand Results with Depreciation

	(1) OLS static	(2) IV static	(3) OLS dynamic	(4) IV dynamic I	(5) IV dynamic II
price/W ($/W)	-0.013***	-5.288*	-0.010***	-0.280***	-0.233***
	(0.003)	(2.864)	(0.003)	(0.076)	(0.017)
log within-group share	0.733***	1.102***	0.894***	0.910***	0.719***
	(0.021)	(0.217)	(0.012)	(0.015)	(0.045)
contractor installed base in county	1.339***	4.112	0.309***	0.400***	0.743***
	(0.292)	(2.687)	(0.085)	(0.152)	(0.109)
contractor installed base outside county	0.053***	-1.021+	0.003	-0.035***	0.002
	(0.012)	(0.608)	(0.006)	(0.013)	(0.008)
electricity rate ($/kWh)	-10.666	84.456	2.680	5.871	7.352**
	(7.885)	(82.690)	(7.932)	(8.338)	(2.919)
average monthly radiation	0.038***	-0.122	0.029***	0.032***	0.030***
	(0.010)	(0.086)	(0.007)	(0.006)	(0.004)
contractor fixed effects	Y	Y	Y	Y	Y
county fixed effects	Y	Y	Y	Y	Y
R-squared	0.743	0.008	0.934	0.787	0.777
N	24081	24081	24081	24081	24062

Notes: Robust standard errors clustered on county and installer in parentheses.
*** indicates significant at the 1% level, ** at the 5% level, * at the 10% level

Table 7: Learning Estimates

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
		No Markup			Static Markups			Both Markups	
installer installed base within county (1000s)	-0.619	-0.583	-0.078	-0.987**	-0.954*	-0.736*	-0.988**	-0.955*	-0.738*
	(0.511)	(0.513)	(0.346)	(0.491)	(0.493)	(0.396)	(0.491)	(0.493)	(0.390)
installer installed base within county squared	0.691	0.697	1.064**	0.999	1.019*	1.256**	1.002	1.020*	1.258**
	(0.583)	(0.579)	(0.432)	(0.611)	(0.596)	(0.505)	(0.609)	(0.595)	(0.505)
installer installed base outside county (1000s)	0.021	-0.000	-0.065	0.037	0.020	-0.038	0.037	0.020	-0.038
	(0.058)	(0.067)	(0.072)	(0.063)	(0.069)	(0.072)	(0.063)	(0.069)	(0.072)
installer installed base outside county squared	-0.008	0.003	-0.011	-0.010	0.000	-0.015	-0.010	0.000	-0.015
	(0.014)	(0.016)	(0.015)	(0.015)	(0.016)	(0.015)	(0.015)	(0.016)	(0.015)
competitor installed base within county (1000s)	0.069	0.023	0.046	0.042	-0.013	0.001	0.042	-0.013	0.001
	(0.044)	(0.045)	(0.047)	(0.045)	(0.043)	(0.045)	(0.045)	(0.043)	(0.045)
competitor installed base within county squared	-0.010*	-0.002	-0.001	-0.008	-0.000	0.001	-0.008	-0.000	0.001
	(0.005)	(0.003)	(0.004)	(0.005)	(0.003)	(0.004)	(0.005)	(0.003)	(0.004)
installer installed base inside X outside county			0.637***			0.646***			0.647***
			(0.189)			(0.199)			(0.199)
installer X competitor installed base inside county			-0.446***			-0.328**			-0.328**
			(0.120)			(0.143)			(0.143)
installed base (1000s)	0.027***			0.035***			0.035***		
	(0.008)			(0.009)			(0.009)		
installed base squared	-0.000***			-0.001***			-0.001***		
	(0.000)			(0.000)			(0.000)		
roofing wage rate ($10,000)	0.064	0.029	0.021	0.068	0.029	0.023	0.068	0.029	0.023
	(0.044)	(0.035)	(0.035)	(0.045)	(0.037)	(0.038)	(0.045)	(0.037)	(0.038)
non-residential	0.031	0.101*	0.096*	0.031	0.099*	0.095	0.031	0.099*	0.095
	(0.062)	(0.057)	(0.057)	(0.062)	(0.058)	(0.057)	(0.063)	(0.058)	(0.057)
third-party owned	-0.157***	-0.158***	-0.145***	-0.164***	-0.161***	-0.149***	-0.164***	-0.161***	-0.149***
	(0.024)	(0.024)	(0.024)	(0.025)	(0.025)	(0.025)	(0.025)	(0.025)	(0.025)
size (kW)	-0.264***	-0.265***	-0.264***	-0.256***	-0.259***	-0.259***	-0.256***	-0.259***	-0.259***
	(0.015)	(0.015)	(0.014)	(0.015)	(0.015)	(0.014)	(0.015)	(0.015)	(0.014)
installer ongoing installations (1000s)	-0.032**	-0.040***	-0.038***	-0.033**	-0.042***	-0.040***	-0.017	-0.016	-0.018
	(0.014)	(0.013)	(0.014)	(0.015)	(0.015)	(0.015)	(0.015)	(0.015)	(0.014)
installer ongoing installations in county (1000s)	0.006	0.008	0.005	-0.017	-0.016	-0.018	-0.033**	-0.043***	-0.041***
	(0.013)	(0.013)	(0.012)	(0.015)	(0.015)	(0.014)	(0.015)	(0.015)	(0.015)
rebate ($/W)	0.263***			0.248***			0.248***		
	(0.045)			(0.049)			(0.049)		
County FEs	Y	Y	Y	Y	Y	Y	Y	Y	Y
Installer FEs	Y	Y	Y	Y	Y	Y	Y	Y	Y
Utility X Quarter FE	N	Y	Y	N	Y	Y	N	Y	Y
R-squared	0.576	0.581	0.583	0.573	0.579	0.580	0.573	0.579	0.580
N	76846	76846	76846	76846	76846	76846	76837	76837	76837

Notes: Robust standard errors clustered on county in parentheses.
*** indicates significant at the 1% level, ** at the 5% level, * at the 10% level

Figure 1: The California Solar Initiative incentive steps.

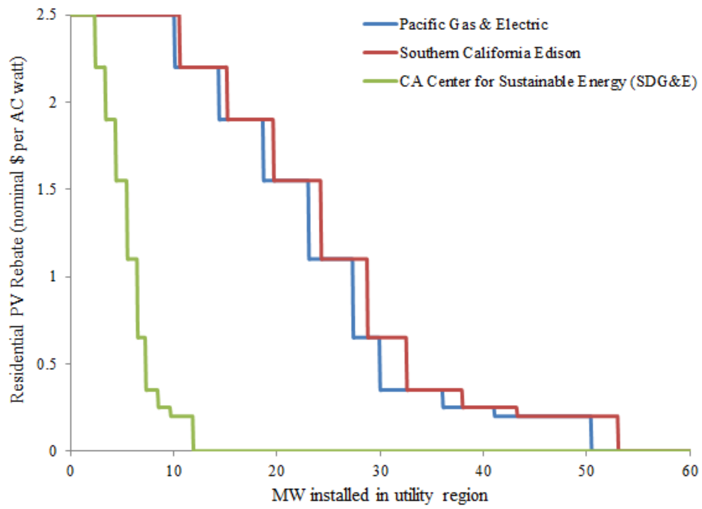

Figure 2: Average requested installations per month

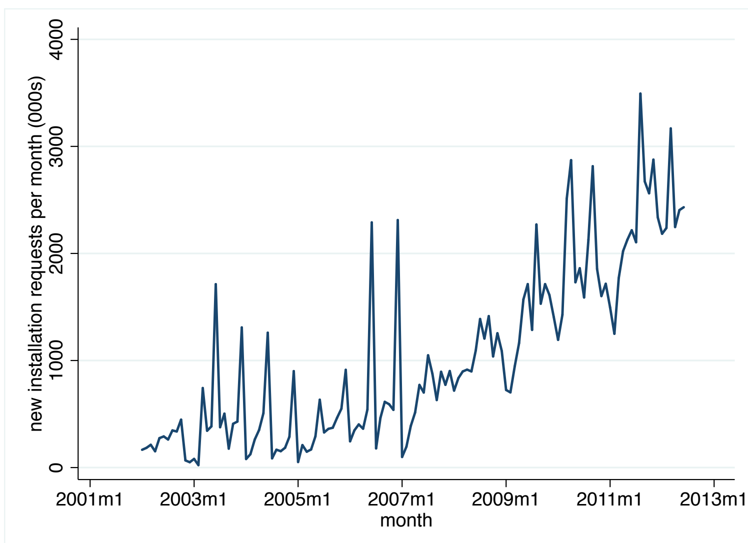

Figure 3: CA solar prices over time

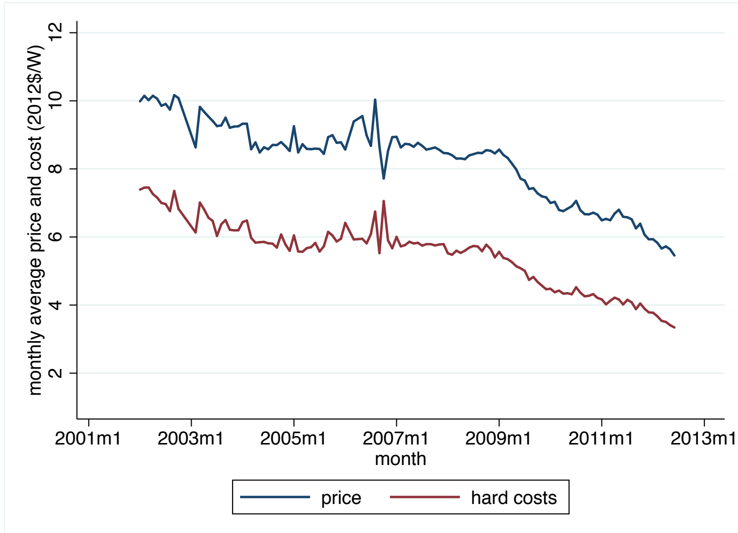

Figure 4: BOS all vs. owned systems

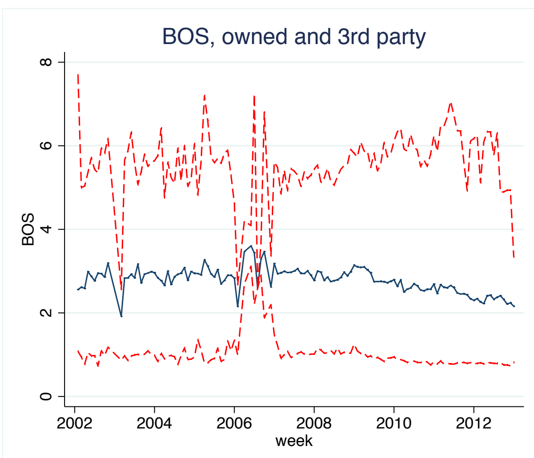

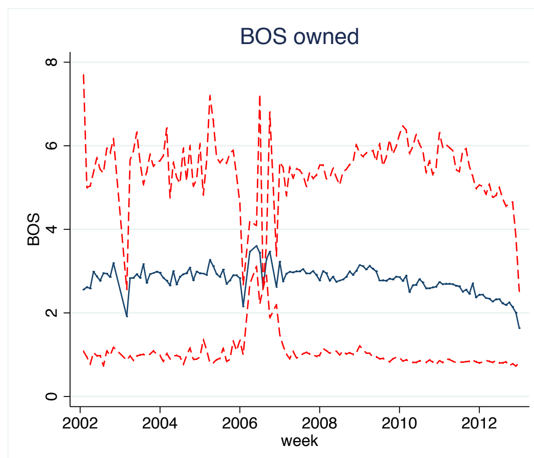

Figure 5: Distribution of firms by firm size

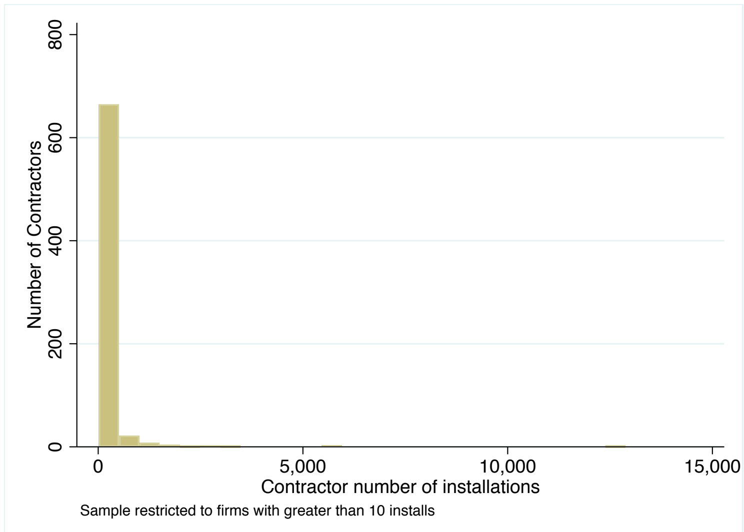

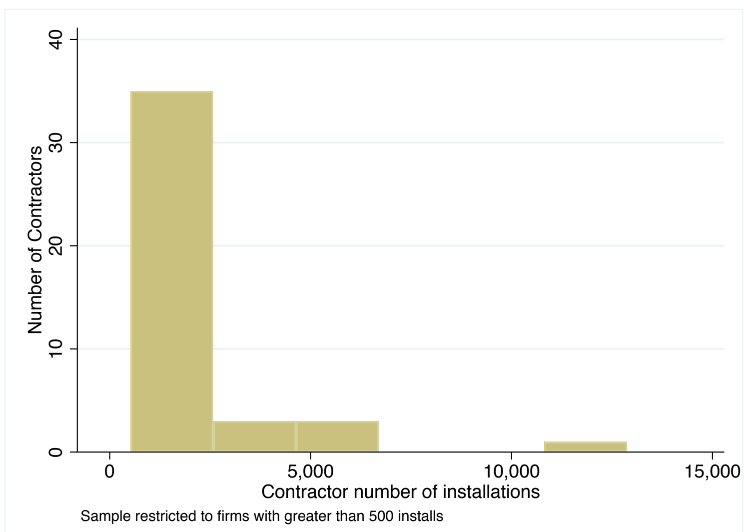

Figure 6: Balance-of-System (BOS) over time by installation size

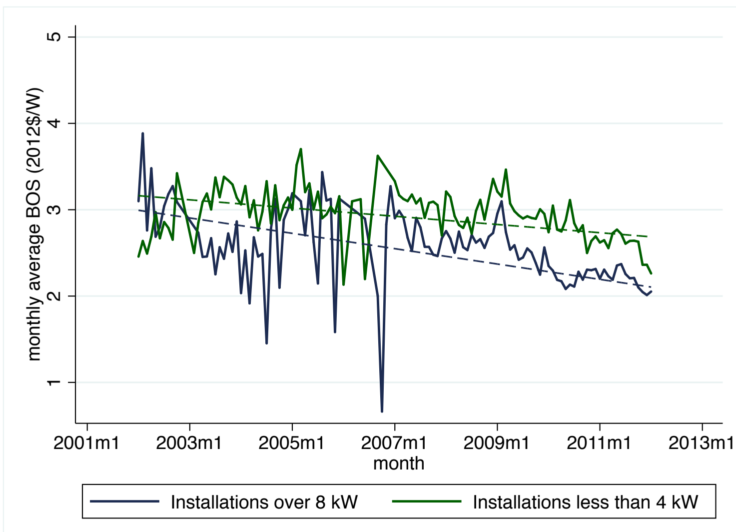

Figure 7: Demand Elasticity over Time

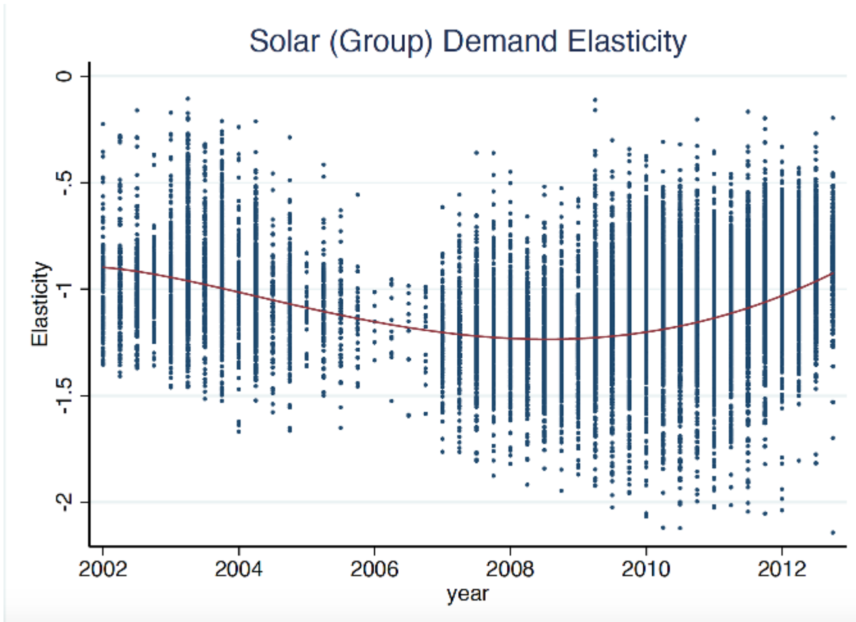

(a) Solar (group) demand elasticity

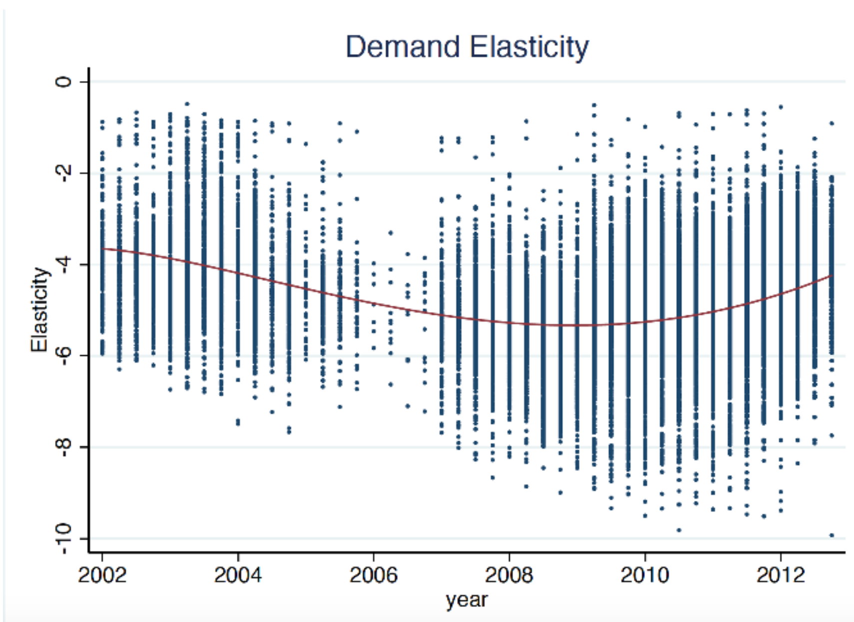

(b) Installer demand elasticity

Figure 8: Optimal Static Markup over Time, By Installer Size

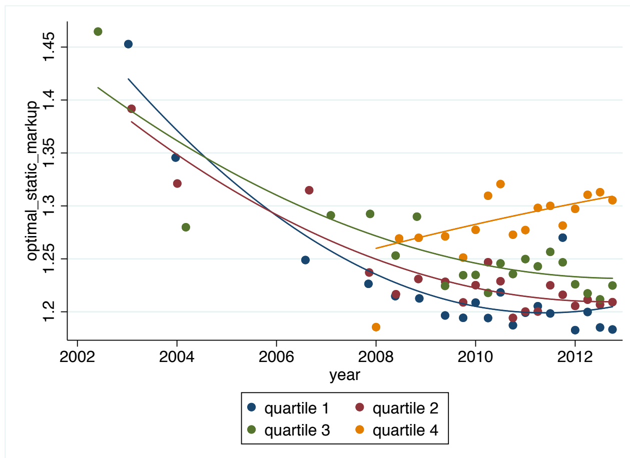

(a) by quartile of contractor installed base

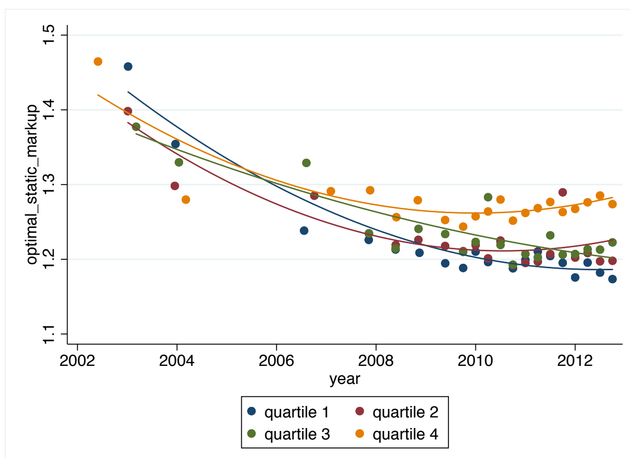

(b) by quartile of contractor county installed base

Figure 9: Estimated Non-Hardware Costs over Time

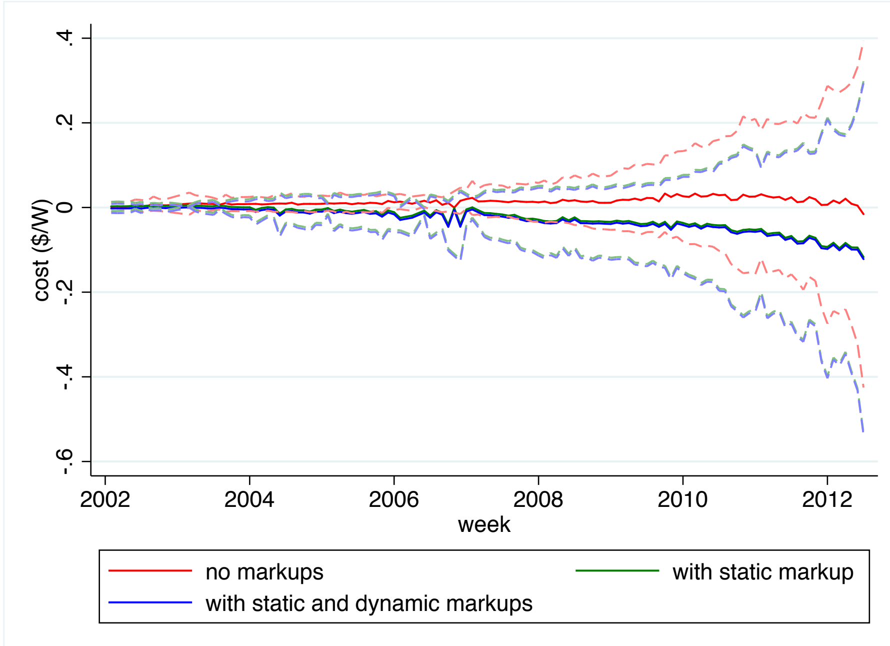

Figure 10: Quarter Fixed Effects

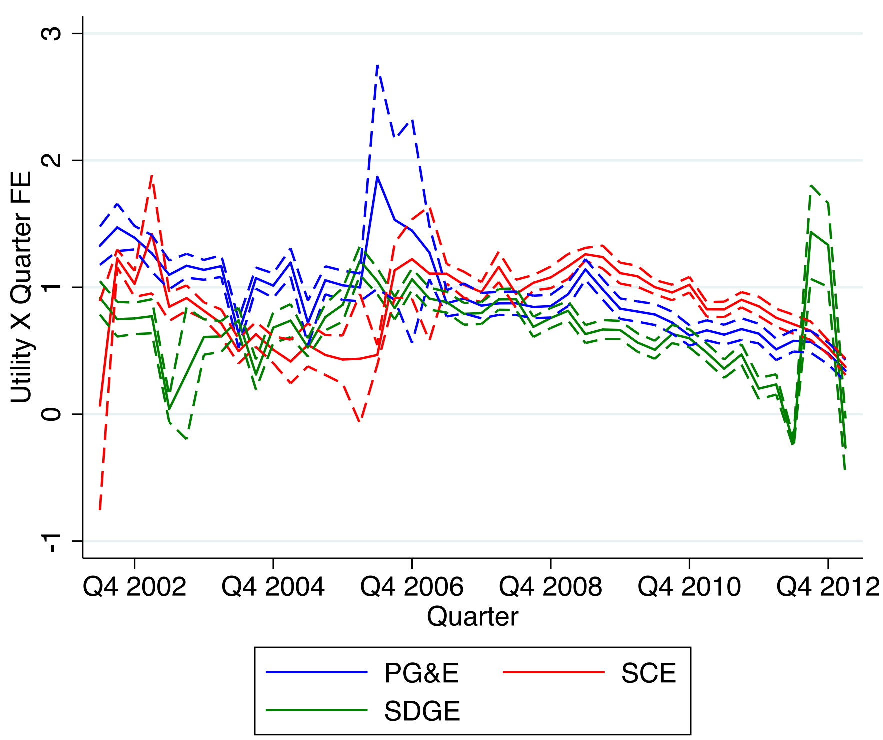

Figure 11: Cumulative Installations with vs. without CSI

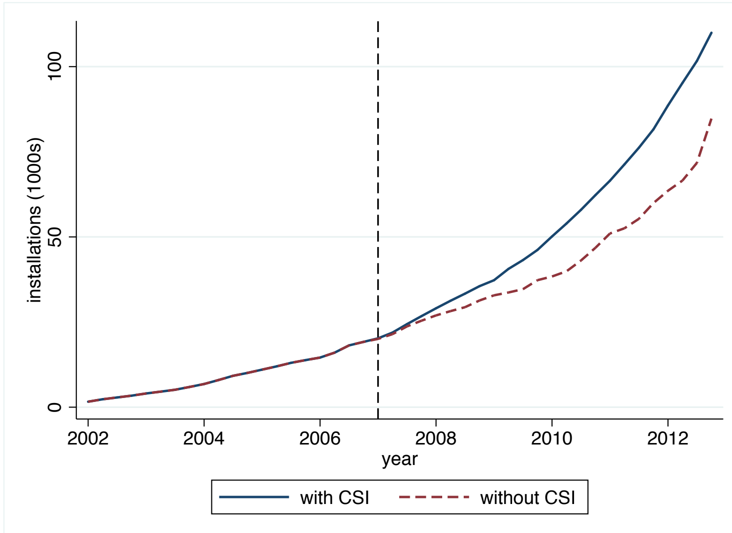

Figure 12: Quarterly Environmental Benefits of New Installations due to CSI

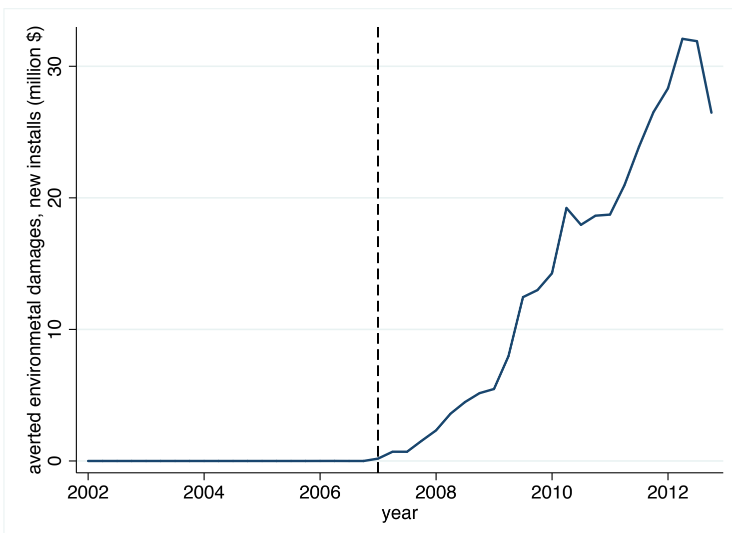

Appendix A: Demand Model with Evolving Heterogeneity Distribution

Let demand now be given by:

$$u_{ijt} = \mu_{mjt} + \eta_i + \varepsilon^u_{igt}(\sigma) + (1-\sigma)\epsilon^u_{ijt},$$

in which η_i leads to within-market individual or segment level heterogeneity in utility for solar (recall there are also market-installer fixed effects in μ_{mjt} that shift the entire distribution for the market). We define s_{ijt} and s_{ijt} as the probability of choice j and 0, respectively, for consumer (or consumer type) i. We assume that the heterogeneity in utility for solar is not firm specific, and thus the η_i drives heterogeneous substitution to the outside good.

We can write the log odds equation for any consumer in the dynamic model as:

$$\log s_{ijt} - \log s_{i0t} = \Delta v_{ijt} + \eta_i, \tag{40}$$

where similar to before, we define:

$$\Delta v_{ijt} \equiv v(x_{mt}, j_t) - v(x_{mt}, 0) = \mu_{mjt} - \beta \mathbb{E}\left[\mu_{m1t+1}\right] + \beta \mathbb{E}\left[\psi(p_{it+1}(x_{mt+1}))\right] \tag{41}$$

The only difference than in the homogenous case is that the next period purchase probabilities are individual-specific.

Aggregating, we have that:

$$\sum_{i \in M_m} w_{it} \left(\log s_{ijt} - \log s_{i0t}\right) = \Delta v_{mjt} + \sum_{i \in M_m} w_i \eta_i.$$

where the probability of consumer type i, i.e. the share of type i, is w_{it} at time t. If the initial distribution of η_i is symmetric and centered at zero, we have that $\sum_{i \in M_m} w_{i0}\eta_i = 0$.

Now we do not observe s_{ijt} and s_{ijt} for every consumer – instead we observe the aggregated market share for the market. However, we know that:

$$\begin{aligned}
\sum_{i \in M_m} s_{ijt} w_{it} &= s_{mjt} \\
\sum_{i \in M_m} s_{i0t} w_{it} &= s_{m0t} \\
s_{ij/It} &= s_{mj/It} \; \forall i \in M_m
\end{aligned} \tag{42}$$

Note that the observed within-group share for each installer j conditional on observing is the same for all consumers, since the heterogeneity scales utility for all installers by $\exp(\eta_i)$.

The distribution of consumer types, represented by the w_{it}, does evolve over time. As installations occur in the market, those consumers with larger η_i are more likely to install earlier, leaving more consumers with lower η_i. Because of this, we need to account for the change in substitution to the outside good over time.

Let us consider a simple two-segment market with $\eta_H = \eta_l$ and $\eta_L = -\eta_l$ such that the probability of the high type is w_{mt}. We can aggregate across segments to get the combined market shares:

$$s_{mHjt} w_{mt} + s_{mLjt}(1 - w_t) \;\; = \;\; s_{mjt}, \tag{43}$$

for each installer j. The log-odds equation for each segment is:

$$\log s_{mljt} - \log s_{ml0t} = \Delta v_{mljt} + \sigma \log(s_{mj/It}) + \eta_l,$$

for $l = L, H$. We can subtract of the low segment from the high segment to get:

$$(\log s_{mHjt} - \log s_{mH0t}) - (\log s_{mLjt} - \log s_{mL0t}) = 2\eta, \tag{44}$$

We also know of course that:

$$\sum_j s_{mHjt} = 1 - s_{mH0t}$$
$$\sum_j s_{mLjt} = 1 - s_{mL0t} \tag{45}$$
$$\tag{46}$$

Across (43), (44), and (45), we have $2j+2$ equations with $2j+2$ unknowns for each market m, if we know η. We first solve for the market shares of each type of consumer (H and L) in the first period, given the starting probability distribution determined by w_{m0}. Given the share of installations that we observe, we can calculate the share that comes from each type in expectation, and then we update $w_{mt} = w_0$ for the next period, repeating for all periods.

This approach allows us to control for the evolution of market heterogeneity. After each period, the share of remaining non-adopters of the high type will decline, while the share of low types will increase (relative to the new market size of non-adopters). We test whether this evolution of consumer types over time matters for our estimation by using different combinations of w_0 and η. Results are shown in Table A.1.

We expect that the largest differences from the model assuming a representative consumer (for each market) will be situations in which the share of high types is low to begin with, because the number of high types becomes exhausted over the course off the panel, resulting in only fewer high types remaining. We find that even for the case in which heterogeneity is large ($eta = 1$) and the share of low types is initially low (1%), the results are not significantly different than those when assuming a representative consumer. This is likely due to the fact that solar adoption is fairly low still in 2012 relative to the potential market size, and thus the distribution of types within markets does not change enough to alter the results, given the extensive market, time, and installer fixed effects we also include.

We would expect our results to possibly differ the most in the last specification, in

Table A.1: Demand Results with Evolving Heterogeneity

	(1)	(2)	(3)	(4)	(5)	(6)
w_{m0}	0.5	0.1	0.01	0.5	0.1	0.01
η	0.5	0.5	0.5	1.0	1.0	1.0
price/W ($)	-0.234***	-0.237***	-0.237***	-0.225***	-0.234***	-0.241***
	(0.045)	(0.045)	(0.045)	(0.044)	(0.045)	(0.046)
log within-group share	0.717***	0.714***	0.715***	0.727***	0.715***	0.709***
	(0.075)	(0.075)	(0.075)	(0.074)	(0.076)	(0.075)
contractor installed base in county	0.747**	0.758***	0.755***	0.716**	0.752**	0.774***
	(0.227)	(0.230)	(0.228)	(0.222)	(0.229)	(0.230)
contractor installed base outside county	0.003	0.003	0.002	0.002	0.003	0.003
	(0.015)	(0.015)	(0.015)	(0.015)	(0.015)	(0.015)
electricity rate ($/kWh)	7.373	7.429	7.429	7.164	7.331	7.463
	(7.080)	(7.107)	(7.064)	(7.069)	(7.163)	(7.039)
average monthly radiation	0.030***	0.030***	0.030***	0.030***	0.030***	0.030***
	(0.006)	(0.006)	(0.006)	(0.006)	(0.006)	(0.006)
contractor fixed effects	Y	Y	Y	Y	Y	Y
county fixed effects	Y	Y	Y	Y	Y	Y
R-squared	0.776	0.771	0.772	0.790	0.774	0.764
N	24062	24062	24062	24062	24062	24062

Notes: Robust standard errors clustered on county and installer in parentheses.
*** indicates significant at the 1% level, ** at the 5% level, * at the 10% level

which there is only a small segment of consumers who have much higher utility for solar, and it is these consumers that are adopting early, such that a much smaller proportion of high types remain by the end of the panel. However, across all specification, we find that the results are very similar to the results when assuming homogenous consumers. The time fixed effects presumably are sufficient to capture much of the effect of declining average utility for solar as the ”low-hanging fruit” are picked. Another reason we see a negligible effect from adding the within-market heterogeneity is that the market sizes are still much bigger than the installed bases within the county by the end of the panel, and thus there are still enough high types remaining.

Appendix B: Dynamic Model Estimation Details

State transitions

Table B.1: State Transitions

	(1) log labor rate	(2) county new installs	(3) log cost per W	(4) log average size
lagged DV	0.7217***	0.7733***	0.9019***	0.0805***
	(0.0098)	(0.0124)	0.0040	(0.0079)
log average size			-0.0234***	
			(0.0030)	
county fixed effects	Y	Y	N	N
contractor X county FE	N	N	Y	Y
R-squared	0.9090	0.8010	0.8933	0.4556
N	1,579	2,155	17,685	17,685
S.E. of residuals	0.0021	0.0636	0.0577	.01631

Notes: Robust standard errors clustered on county and installer in parentheses.
*** indicates significant at the 1% level, ** at the 5% level, * at the 10% level

Policy function

The policy function regression is a flexible quadratic including all interactions between lag price, the three installed base variables, installer ongoing contracts (both in the county and outside), rebate levels, and labor costs. The R squared of the policy function regression is 0.7513 and the standard deviation of the residuals of policy function regression is 0.1313. We include time fixed effects in this regression to ensure that we do not over parametrize it. In the forward simulations, for periods not yet observed we use predicted values of the fixed effects, calculated from regressing the time fixed effects on the previous period fixed effect.

To get a sense of firms' strategy functions, we plot the price policy as a function of time and own-county installed base for the observations in the data in Figure B.1. We also show how the price changes if the own-installed base in the county were to be increased by one, which has direct implications for the current-period incentive to lower price in order to increase the installed base.

Figure B.1: Pricing policy function

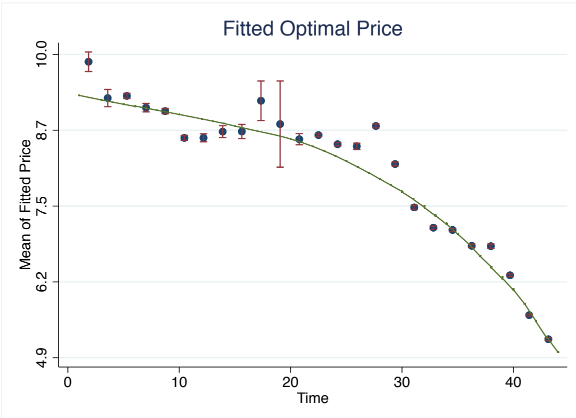

(a) Price as a function of time

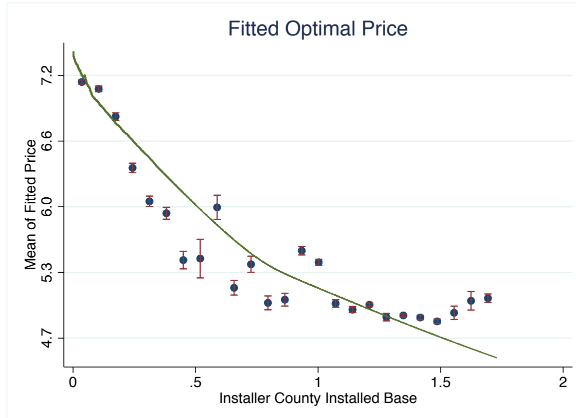

(b) Price as a function of own county installations

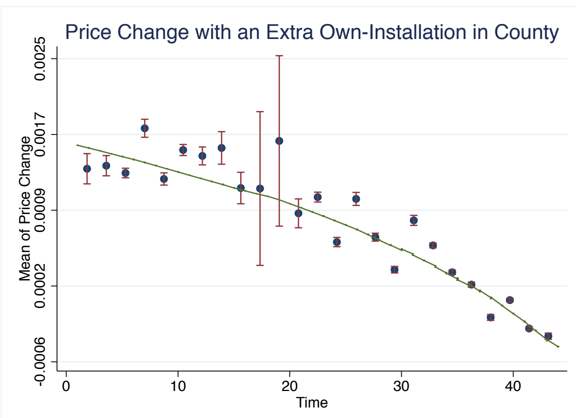

(c) Change in price with extra own county installation as a function of time

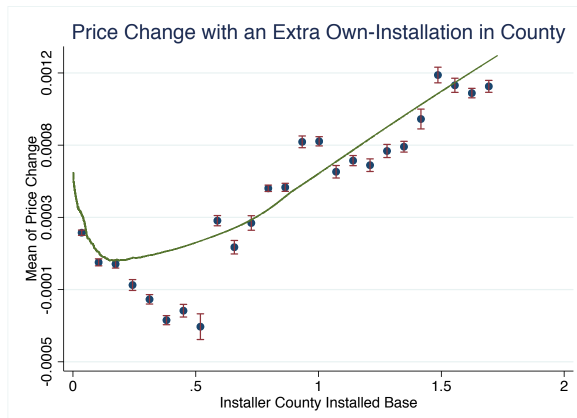

(d) Change in price with extra own county installation as a function of own county installations

Optimal prices of course go down over time, and they also go down with own county installed base in the data, but of course this can simply be reflecting the fact that installed bases increase with time. With an increase in own county installed base, the optimal price increases, especially in the earlier period of the data, which can reflect both the fact that consumers are more willing to pay for installations by that installer and also that the dynamic pricing incentive has declined now that the firm is farther down the learning curve (assuming convexity of the learning curve). When we graph this price difference as a function of own county installed base, we see that at larger installed bases the price premium actually increases more in the data.

Appendix C: Additional Results

Table C.1: First Stage Demand Model

	price ($/W)	log within-group share
contractor installed base in county	0.854***	2.020***
	(0.200)	(0.141)
contractor installed base outside county	-0.057***	0.070***
	(0.016)	(0.011)
electricity rate ($/kWh)	22.967***	2.914
	(8.013)	(5.642)
average monthly radiation	0.037***	-0.010
	(0.010)	(0.006)
rebate ($/W)	0.345***	0.076
	(0.084)	(.059)
installer installations finished outside county	-0.883***	0.650***
	(0.102)	(0.072)
competitor installations finished outside county	-0.162***	-0.053***
	(0.011)	(0.008)
contractor X county fixed effects	Y	Y
F-statistic	54.51	37.82
R-squared	0.0300	0.600
N	24062	24062

Notes: Robust standard errors clustered on county and installer in parentheses.
*** indicates significant at the 1% level, ** at the 5% level, * at the 10% level

Table C.2: Demand Results with No Depreciation

	(1) OLS static	(2) IV static	(3) OLS dynamic	(4) IV dynamic I	(5) IV dynamic II
price ($/W)	-0.011***	-6.292*	-0.015***	-0.254***	-0.222***
	(0.003)	(3.610)	(0.003)	(0.092)	(0.045)
log within-group share	0.737***	1.155***	0.894***	0.911***	0.568***
	(0.017)	(0.259)	(0.013)	(0.018)	(0.059)
contractor installed base in county	1.275***	4.247	0.354***	0.414***	0.839***
	(0.259)	(2.827)	(0.094)	(0.119)	(0.165)
contractor installed base outside county	0.026*	-1.139	0.009	-0.022	0.018
	(0.012)	(0.698)	(0.006)	(0.014)	(0.013)
electricity rate ($/kWh)	-14.095	91.569	1.620	3.750	10.120*
	(10.212)	(109.348)	(5.637)	(7.383)	(5.556)
average monthly radiation	-0.010	-0.217*	0.023**	0.026***	0.016**
	(0.011)	(0.116)	(0.009)	(0.008)	(0.007)
contractor fixed effects	Y	Y	Y	Y	Y
county fixed effects	Y	Y	Y	Y	Y
R-squared	0.673	0.004	0.914	0.793	0.701
N	23999	23999	23999	23999	23980

Notes: Robust standard errors clustered on county and installer in parentheses.
*** indicates significant at the 1% level, ** at the 5% level, * at the 10% level

Table C.3: Learning Estimate Robustness Checks: Depreciation of Installed Bases

	(1)	(2)	(3)	(4)	(5)	(6)
	With Static Markup			With Both Markups		
installer installed base within county (1000s)	-0.714**	-0.676**	-0.630**	-0.711**	-0.673**	-0.626**
	(0.291)	(0.294)	(0.282)	(0.292)	(0.295)	(0.284)
installer installed base within county squared	0.773***	0.777***	0.987***	0.770***	0.773***	0.984***
	(0.266)	(0.261)	(0.233)	(0.268)	(0.263)	(0.233)
installer installed base outside county (1000s)	0.029	0.026	-0.001	0.029	0.026	-0.001
	(0.039)	(0.040)	(0.043)	(0.039)	(0.040)	(0.043)
installer installed base outside county squared	-0.006	-0.003	-0.006	-0.006	-0.003	-0.006
	(0.006)	(0.006)	(0.006)	(0.006)	(0.006)	(0.006)
competitor installed base within county (1000s)	0.020	-0.026	-0.020	0.020	-0.026	-0.020
	(0.027)	(0.022)	(0.024)	(0.027)	(0.022)	(0.024)
competitor installed base within county squared	-0.003	0.001	0.001	-0.003	0.001	0.001
	(0.002)	(0.001)	(0.001)	(0.002)	(0.001)	(0.001)
installer installed base inside X outside county			0.169**			0.169**
			(0.074)			(0.075)
installer X competitor installed base inside county			-0.070***			-0.071***
			(0.026)			(0.026)
installed base (1000s)	0.017***			0.017***		
	(0.004)			(0.004)		
roofing wage rate ($10,000)	0.063	0.018	0.014	0.063	0.018	0.014
	(0.047)	(0.041)	(0.040)	(0.047)	(0.041)	(0.040)
non-residential	-0.062	0.001	-0.001	-0.063	0.001	-0.001
	(0.067)	(0.063)	(0.063)	(0.068)	(0.063)	(0.063)
third-party owned	-0.184***	-0.182***	-0.173***	-0.184***	-0.182***	-0.173***
	(0.027)	(0.025)	(0.025)	(0.027)	(0.025)	(0.025)
size (kW)	-0.012***	-0.011***	-0.011***	-0.012***	-0.011***	-0.011***
	(0.001)	(0.001)	(0.001)	(0.001)	(0.001)	(0.001)
installer ongoing installations (1000s)	-0.042**	-0.048***	-0.048***	-0.042**	-0.048***	-0.048***
	(0.018)	(0.017)	(0.017)	(0.018)	(0.017)	(0.017)
installer ongoing installations in county (1000s)	-0.046***	-0.045***	-0.045***	-0.046***	-0.045***	-0.045***
	(0.011)	(0.011)	(0.011)	(0.011)	(0.011)	(0.011)
rebate ($/W)	0.220***			0.220***		
	(0.041)			(0.041)		
County FEs	Y	Y	Y	Y	Y	Y
Installer FEs	Y	Y	Y	Y	Y	Y
Quarter X Utility FE	N	Y	Y	N	Y	Y
R-squared	0.562	0.568	0.569	0.562	0.568	0.569
N	71728	71728	71728	71727	71727	71727

Notes: Robust standard errors clustered on county in parentheses.
*** indicates significant at the 1% level, ** at the 5% level, * at the 10% level

Table C.4: Learning Estimate Robustness Checks: Interactions with Wage Rate

	(1)	(2)	(3)	(4)	(5)	(6)
	With Static Markup			With Both Markups		
installer installed base within county (1000s)	-0.247**	-0.237**	-0.171**	-0.247**	-0.237**	-0.171**
	(0.101)	(0.099)	(0.084)	(0.100)	(0.099)	(0.084)
installer installed base within county squared	0.256*	0.256**	0.294***	0.256*	0.256**	0.294***
	(0.128)	(0.122)	(0.107)	(0.128)	(0.122)	(0.107)
installer installed base outside county (1000s)	0.008	0.004	-0.011	0.008	0.004	-0.011
	(0.014)	(0.015)	(0.015)	(0.014)	(0.015)	(0.015)
installer installed base outside county squared	-0.003	-0.000	-0.004	-0.003	-0.000	-0.004
	(0.004)	(0.004)	(0.004)	(0.004)	(0.004)	(0.004)
competitor installed base within county (1000s)	0.007	-0.007	-0.003	0.007	-0.007	-0.003
	(0.012)	(0.010)	(0.010)	(0.012)	(0.010)	(0.010)
competitor installed base within county squared	-0.002	0.000	0.001	-0.002	0.000	0.001
	(0.001)	(0.001)	(0.001)	(0.001)	(0.001)	(0.001)
installer installed base inside X outside county			0.163***			0.163***
			(0.052)			(0.052)
installer X competitor installed base inside county			-0.080**			-0.080**
			(0.038)			(0.038)
installed base (1000s)	0.036***			0.036***		
	(0.009)			(0.009)		
roofing wage rate ($10,000)	0.071*	0.039	0.034	0.071*	0.039	0.034
	(0.042)	(0.033)	(0.034)	(0.042)	(0.033)	(0.034)
non-residential	0.030	0.099*	0.095	0.030	0.099*	0.095
	(0.063)	(0.058)	(0.057)	(0.063)	(0.058)	(0.057)
third-party owned	-0.164***	-0.160***	-0.149***	-0.164***	-0.160***	-0.149***
	(0.026)	(0.025)	(0.025)	(0.026)	(0.025)	(0.025)
size (kW)	-0.256***	-0.259***	-0.259***	-0.256***	-0.260***	-0.260***
	(0.015)	(0.015)	(0.014)	(0.015)	(0.015)	(0.014)
installer ongoing installations (1000s)	-0.032**	-0.042***	-0.040***	-0.032**	-0.042***	-0.040***
	(0.015)	(0.015)	(0.015)	(0.015)	(0.015)	(0.015)
installer ongoing installations in county (1000s)	-0.016	-0.015	-0.018	-0.016	-0.015	-0.018
	(0.014)	(0.014)	(0.014)	(0.014)	(0.014)	(0.014)
rebate ($/W)	0.248***			0.248***		
	(0.050)			(0.050)		
County FEs	Y	Y	Y	Y	Y	Y
Installer FEs	Y	Y	Y	Y	Y	Y
Quarter X Utility FE	N	Y	Y	N	Y	Y
R-squared	0.573	0.579	0.580	0.573	0.579	0.580
N	76846	76846	76846	76837	76837	76837

Notes: Robust standard errors clustered on county in parentheses.
*** indicates significant at the 1% level, ** at the 5% level, * at the 10% level

Table C.5: Learning Estimate Robustness Checks: No Learning Transfer with Acquisition

	(1)	(2)	(3)	(4)	(5)	(6)
	With Static Markup			With Both Markups		
installer installed base within county (1000s)	-0.987**	-0.887*	-0.710*	-0.986**	-0.886*	-0.707*
	(0.447)	(0.449)	(0.392)	(0.447)	(0.448)	(0.391)
installer installed base within county squared	1.100*	1.037*	1.293***	1.103*	1.039*	1.288***
	(0.580)	(0.562)	(0.466)	(0.579)	(0.562)	(0.465)
installer installed base outside county (1000s)	0.012	0.012	-0.042	0.012	0.011	-0.042
	(0.065)	(0.069)	(0.073)	(0.065)	(0.069)	(0.073)
installer installed base outside county squared	-0.007	0.000	-0.015	-0.007	0.000	-0.015
	(0.016)	(0.017)	(0.016)	(0.016)	(0.017)	(0.016)
competitor installed base within county (1000s)	0.030	-0.027	-0.014	0.030	-0.027	-0.014
	(0.041)	(0.043)	(0.045)	(0.041)	(0.043)	(0.045)
competitor installed base within county squared	-0.007	0.002	0.003	-0.007	0.002	0.003
	(0.005)	(0.003)	(0.004)	(0.005)	(0.003)	(0.004)
installer installed base inside X outside county			0.631***			0.628***
			(0.201)			(0.200)
installer X competitor installed base inside county			-0.301**			-0.299**
			(0.117)			(0.117)
installed base (1000s)	0.037***			0.037***		
	(0.009)			(0.009)		
roofing wage rate ($10,000)	0.079*	0.022	0.017	0.078*	0.022	0.017
	(0.046)	(0.038)	(0.038)	(0.046)	(0.038)	(0.039)
non-residential	0.037	0.099*	0.095	0.037	0.099*	0.095
	(0.063)	(0.059)	(0.059)	(0.063)	(0.059)	(0.059)
third-party owned	-0.168***	-0.161***	-0.150***	-0.168***	-0.161***	-0.150***
	(0.026)	(0.026)	(0.026)	(0.026)	(0.026)	(0.026)
size (kW)	-0.252***	-0.255***	-0.255***	-0.252***	-0.255***	-0.255***
	(0.015)	(0.015)	(0.015)	(0.015)	(0.015)	(0.015)
installer ongoing installations (1000s)	-0.042**	-0.050***	-0.048**	-0.042**	-0.050***	-0.048**
	(0.018)	(0.017)	(0.018)	(0.018)	(0.017)	(0.018)
installer ongoing installations in county (1000s)	-0.025**	-0.024**	-0.026**	-0.025**	-0.024**	-0.026**
	(0.011)	(0.011)	(0.010)	(0.011)	(0.011)	(0.010)
rebate ($/W)	0.220***			0.220***		
	(0.044)			(0.044)		
County FEs	Y	Y	Y	Y	Y	Y
Installer FEs	Y	Y	Y	Y	Y	Y
Quarter X Utility FE	N	Y	Y	N	Y	Y
R-squared	0.570	0.576	0.577	0.570	0.576	0.577
N	72654	72654	72654	72654	72654	72654

Notes: Robust standard errors clustered on county in parentheses.
*** indicates significant at the 1% level, ** at the 5% level, * at the 10% level