Harnessing the Power of the Sun through Online Platforms: Evidence from the Rooftop Solar Market

Jackson Dorsey*

JOB MARKET PAPER

November 12, 2017

Please find the latest version HERE.

Abstract

Buyers in the residential solar photovoltaic (PV) market are often not fully informed about all competitors' prices when making a purchase decision. Consumers' lack of complete price information can incentivize sellers to charge higher markups, leading to less adoption. Large markups directly counteract public programs that subsidize solar PV. In this paper, I use proprietary data on consumers' purchase decisions and sellers' bids from an online platform to estimate a structural model of supply and demand for solar PV systems. The model allows me to estimate buyers' price elasticities and to infer the size of sellers' markups. I first show that markups make up a substantial portion of solar PV prices. I then use the estimated model to simulate changes in prices and solar panel adoption rates if consumers did not have access to the platform. I find that access to the platform increases solar panel adoption by 104%. There are two primary drivers of this result: 1) the market becomes more competitive which drives down quoted prices and, 2) reducing search costs connects buyers to higher quality sellers. I find that the improvement in competition alone would increase adoption by 30% and the improved access to high-quality sellers alone would increase adoption by 74%.

JEL Codes: D22, D44, L11, L15, Q40, Q58

Keywords: solar PV, renewable energy, multi-attribute auctions, search costs, online platforms, market power, quality, technology adoption

*Department of Economics, University of Arizona (e-mail: jacksondorsey@email.arizona.edu). I thank Ashley Langer, Derek Lemoine, Gautam Gowrisankaran, Stan Reynolds, Gary Solon, Derek Wolfson, Keith Meyers, Dev Bose, Alvin Murphy and seminar participants at Arizona State, the University of Arizona, NBER (EEE) Summer Institute, and the CU Environmental Economics Workshop for helpful comments, discussions, and suggestions. I also thank Luke Tarbi and his team at EnergySage Inc. for providing instrumental help with data access.

Solar photovoltaics (PV) enable homeowners to generate carbon-free electricity by bolting PV panels to their rooftop. Therefore, numerous local and national governments have enacted policies to reduce electricity sector emissions by encouraging investment in residential solar PV. Generous subsidy programs and plummeting PV hardware costs have bolstered a 50-fold increase in total U.S. solar PV capacity over the last decade. Nonetheless, total installation prices remain prohibitively high for most households to adopt the technology.[1] Overall installed prices are relatively high due to large balance-of-system (BOS) costs, which include the labor cost of installation, permitting fees, and installer markups. Figure 1 shows that BOS costs now account for over two-thirds of total PV prices (Barbose and Darghouth, 2016).[2]

Figure 1: Residential Solar PV Module Cost and Implied BOS Cost

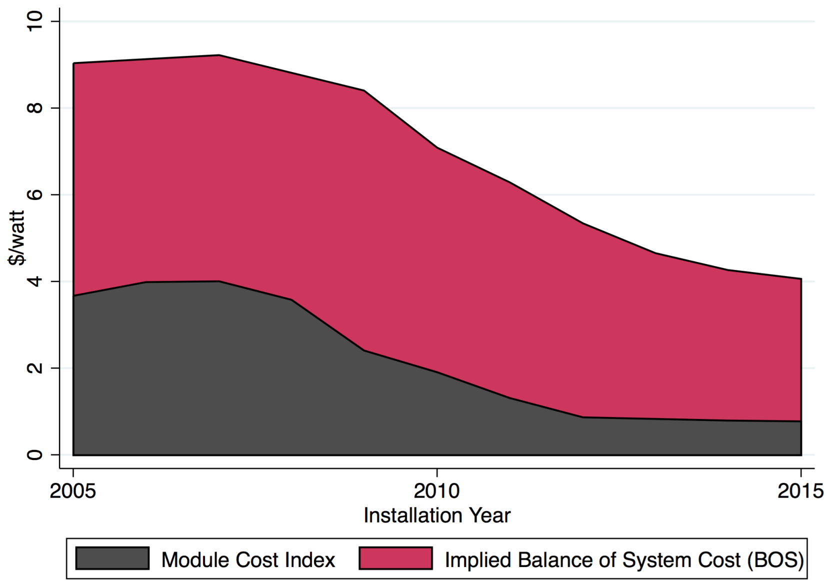

Data Source: Barbose and Darghouth (2016). Implied BOS costs are calculated as the total installed price minus the module price index, and therefore include installer profit margin.

One plausible explanation for high BOS costs is market power. Installers may exercise market power in this setting because most potential buyers are not fully informed about market prices. Buyers lack full information because sellers do not post their prices. Consequently, consumers need to engage in a time-consuming search process that involves calling sellers and scheduling on-site property visits to collect quotes. If collecting price quotes is costly for buyers, any installer asked to give a quote can expect to be bidding against few or no other sellers, thereby giving that installer incentive to charge a higher markup. High markups are particularly problematic because they counteract government subsidy programs that aim to expand solar PV adoption.

[1] As of 2017, solar PV accounted for less than 1% of total U.S. electricity generation.
[2] U.S. BOS costs also appear to be high relative to other countries, such as Germany.

A potential way to mitigate high markups is through an intermediary (e.g., an online platform) that connects buyers and sellers and provides ratings and reviews. Providing consumers with more information about prices and seller characteristics could increase output through two primary channels. For one, installers perceive a higher level of competition if buyers are informed about more sellers in the market. Therefore, each firm should quote lower prices and overall sales should increase. Also, if sellers are differentiated in their quality, giving a consumer access to more potential installers will improve the chances that she finds a higher quality installer and thus makes a purchase.

In this paper, I empirically evaluate how the introduction of an online platform affects markups, solar panel adoption, and welfare in the rooftop solar market. To do so, I collect proprietary data from a platform that links consumers to installers of solar PV systems. The data provides detailed information on purchase choices and seller bidding behavior, which I use to estimate a structural model of the solar PV market. Estimating the model involves combining discrete choice methods to estimate buyers' choice rule and adapting insights from the empirical auctions literature (Guerre et al., 2000; Krasnokutskaya and Seim, 2011) to infer sellers' marginal costs and bid preparation costs. After recovering the primitives of the model, I simulate counterfactual changes in adoption and welfare associated with providing buyers access to additional bids through an online platform.

In addition to measuring the effects of the online platform on prices, adoption, and welfare, I also provide new estimates of consumer price responses in the solar industry using a novel identification strategy. The existing literature treats solar PV as a homogeneous good, and uses variation in subsidy rates and panel costs across space and time to identify effects on overall adoption (Bollinger and Gillingham, 2014; Burr, 2014; De Groote and Verboven, 2016; Gillingham et al., 2016; Gillingham and Tsvetanov, 2014; Hughes and Podolefsky, 2015; Langer and Lemoine, 2017; Pless and van Benthem, 2017; Reddix II, 2015).[3] I exploit rich data on consumers' choice sets and model the adoption decision as a discrete choice between differentiated products. Identification of my model comes from variation in prices across alternatives within an individual's choice set, as well as variation in price quotes and choice sets (firms that submit bids) across ex-ante similar consumers. An advantage of my approach is that I can jointly estimate consumer tastes for other attributes such as panel quality and seller quality that could also be instrumental in consumers' decision to adopt solar. For instance, an installer's warranty package, years of experience, customer service, and star rating are all non-price factors that likely influence consumer purchase choices. To my knowledge, there are no existing estimates in the literature that quantify the importance of non-price attributes on consumers' solar PV adoption decision.

[3] Hughes and Podolefsky (2015) and Gillingham and Tsvetanov (2014) both use reduced-form approaches to estimate the elasticity of demand for residential solar systems and to quantify the adoption response to subsidy programs. Burr (2014), Reddix II (2015), De Groote and Verboven (2016), and Langer and Lemoine (2017) all develop dynamic discrete choice models to estimate demand for solar PV systems and to assess the welfare effects of different subsidy policies.

Supply-side behavior is also crucial in determining equilibrium prices and adoption rates. However, most previous work has not explicitly modeled or discussed the supply-side of the solar PV market. A few notable exceptions include Bollinger and Gillingham (2014), Pless and van Benthem (2017), and Gillingham et al. (2016) who all provide evidence of market power in the residential solar industry. In general, modeling the supply-side of the market has proved challenging because public solar PV datasets do not provide information on which installers are operating in specific geographic areas or which installers are in each consumer's consideration set. Therefore, researchers have either been unable to estimate the exact size of markups or needed to make strong assumptions about how to measure local market structure (which sellers are in a buyer's choice set).[4] In this paper, I improve on the existing literature by collecting new data from an online platform which allows me to observe the choice set of each buyer, including non-selected sellers' bids. I use the data to develop and estimate a model of seller bidding and participation in multi-attribute auctions. A multi-attribute auction is a procurement mechanism where bidders submit multi-dimensional bids (i.e., bid price, panel efficiency, and star-rating). The structure of the multi-attribute auction environment presents several empirical challenges which I discuss in the following sections. I build on methods recently developed by Krasnokutskaya et al. (2017) and Yoganarasimhan (2015) to recover the model parameters. To infer markups, I use each seller's first-order condition for an optimal price bid to decompose observed bid prices into a marginal cost and a markup component. I also estimate sellers' bid preparation costs using firms' observed participation decisions in the auctions. This is the first paper to specify a structural model of the supply and demand side of this market in a unified framework. The full model can provide a better understanding of how welfare and equilibrium adoption may change under different counterfactual policy environments.

My estimated demand model provides novel evidence of substantial differences in sellers' quality in the solar PV installation market. In particular, the average buyer would be willing to pay more than 19% more for a high-quality seller compared to a median quality installer.[6] I also find that the average buyer is willing to pay a significant premium for higher quality equipment (i.e., more efficient panels). On the supply side, I document substantial market power. Specifically, I find that average gross markups account for nearly 40% of system prices.[7] After accounting for overhead costs, my estimates imply that average net markups make up 23% of prices. However, I find that markups fall over time as the platform becomes more competitive.

---

[4] Bollinger and Gillingham (2014) develop a dynamic model of installer pricing under imperfect competition to motivate a reduced-form regression equation that allows them to estimate static and dynamic markups, the authors use county HHI to proxy for the level of competition in a local market. Pless and van Benthem (2017) measure the pass-through rate of solar subsidies to prices and find a pass-through rate of over 100%.[5] The authors demonstrate that this result is only possible in the presence of market power. Pless and van Benthem (2017) avoid making assumptions about which sellers are in buyers' choice sets, but they are not able to estimate the size of markups using their research design.

[6] I estimate firm fixed effects to determine the quality of sellers. Here, I define high-quality sellers as those with firm fixed effects in the 9th decile.

[7] Gross markup is the price minus the marginal cost of the installation and does not include installer overhead costs.

After estimating the model, I use the industry cost structure and demand primitives to simulate changes in prices and solar panel adoption rates if consumers did not have access to the platform. I find that access to the platform increases solar panel adoption by 104%. There are two primary drivers of this result: 1) the market becomes more competitive which drives down the lowest quoted price by 15% on average and, 2) buyers connect to higher quality sellers. I find that the improvement in competition alone would increase adoption by 30% and the improved access to high-quality sellers alone would increase adoption by 74%. I also find that the platform leads to an 89% increase in total welfare. Consumer surplus increases by 156% as consumers enjoy lower prices and access to more (and higher quality) sellers. Sellers benefit from a 69% increase in profits because the fall in prices is more than offset by the rise in the sales quantity. These results suggest that policymakers could increase solar PV adoption by developing their own platforms or by encouraging participation on existing platforms. In fact, the state of Connecticut recently introduced a state-sponsored platform to decrease residential solar prices and increase adoption. Platforms have also been used to promote competition in other industries such as healthcare.[8]

The majority of existing policies aimed at spurring the adoption of solar PV systems have used subsidies to reduce the cost of adoption. For example, the federal government has provided a 30% investment tax credit (ITC) for residential solar purchases since 2006.[9] Although these policies have been effective at increasing adoption,[10] they require considerable public expenditures and often are at risk of being removed. For instance, the federal production tax credit for wind energy has expired and been renewed several times over the past decade. The federal solar ITC has also faced political opposition is now scheduled to be eliminated by 2023.

To compare the impact of the platform to an existing subsidy policy, I run an additional counterfactual simulation where instead of giving consumers access to the platform, the government provides a subsidy equal to 30% of the system price, equivalent to the ITC. The ITC policy increases adoption by 52% compared to the 104% increase under the "platform counterfactual". The increase in adoption resulting from the ITC is more modest because it decreases prices but does not give buyers access to additional high-quality installers. Although the ITC policy provides significant increases in rents for producers and consumers in the market, these gains mostly come from wealth transfers via the subsidy. Namely, I find that total welfare only increases by 8% under the ITC policy after accounting for the cost of the subsidy. Because there are already existing platforms to connect buyers and sellers in the solar PV market, the cost of implementing a platform (or marketing the use of a platform) would likely come at a much lower price compared to a subsidy program.

In the following subsection, I review related literature. In Section 2, I discuss the details of

[8]For example, a primary aim of the Affordable Care Act was to reduce premiums and expand health insurance coverage for Americans. One principal mechanism for meeting that goal was the establishment of new individual health insurance marketplaces (platforms) where consumers can shop for, compare, and purchase plans.

[9]The solar ITC also provides a tax credit for commercial or utility solar PV investments.

[10]Hughes and Podolefsky (2015) and Burr (2014) find that one subsidy program, the California Solar Initiative, doubled the amount residential solar PV purchases in between 2007 and 2013.

the online platform, provide descriptive statistics, and show reduced-form evidence that use of the online platform correlates with lower solar PV prices paid. In Section 3, I develop a model of buyer and seller behavior in the solar PV market and then discuss the methods used to pair the model to the data. Section 5 discusses the results, and Section 6 concludes.

1.1 Related Literature

This paper relates to several literatures: first, it pertains to an extensive literature on competition in search markets and the role of intermediation. Second, it connects to recent work on the estimation of multi-attribute auctions. Finally, it relates to a growing literature in environmental economics on renewable energy markets and policy.

Search Costs, Price-transparency, and Intermediation: An influential literature starting with Stigler (1961) and has discussed equilibrium pricing behavior of firms and consumers' optimal search effort in markets where consumers lack perfect information about prices. Diamond (1971) and Stahl (1989) show that even small consumer search costs can lead to imperfectly competitive pricing by firms. In this paper, I empirically investigate how the introduction of an intermediary (online platform) affects markups and welfare in the residential solar market, a market with substantial search costs. Previous work has theoretically investigated the role of intermediaries in search markets (Gehrig, 1993; Hall and Rust, 2003; Spulber, 1996). More recently, several empirical studies examine the effect of introducing an intermediary or a technology that increases price transparency in other industries such as life insurance (Brown and Goolsbee, 2002), fisheries (Jensen, 2007), waste management (Salz, 2017), health care (Brown, 2017), and retail gasoline (Luco, 2016).

Multi-attribute auctions: This papers also connects to a recent empirical literature on multi-attribute auctions. In multi-attribute auctions, buyers are allowed to deviate from allocation based solely on price (as in standard auctions) and to choose a seller based off of non-price attributes. A multi-attribute auction differs from a scoring auction because the auctioneer does not announce the decision rule used to pick the winning bidder and the auctioneer can also decide to cancel the auction (choose the outside option).[11] Multi-attribute auctions are used increasingly in both private and public sector procurement in setting where the buyer cares about non-price attributes. Governments use them to sell 3G licenses and in military contracting. In the private sector, multi-attribute auctions are now used by businesses to find freelance programmers and other workers (Krasnokutskaya et al., 2017; Yoganarasimhan, 2015) and for individuals to find contractors for roofing and other household jobs.[12] Estimating a multi-attribute auction is challenging because standard methods from the empirical auction literature (Guerre et al., 2000) cannot immediately be applied. The methods developed by (Guerre et al., 2000) do not directly apply to this setting because the auction allocation rule is unknown and needs to

[11]The marketing literature refers to multi-attribute auctions as "beauty contest" auctions.

[12]Google Adwords and Facebook ads auctions can also be viewed as multi-attribute auctions because the allocation rules depend on attributes other than price and are not announced to bidders.

be estimated. Furthermore, the standard differentiated products model of supply (Berry et al., 1995) also cannot be directly used because firms have incomplete information about their competitors in this setting (i.e., they don't know the number of bidders they will face or their identities). Furthermore, sellers receive project specific cost shocks in this environment. In recent work, Yoganarasimhan (2015) and Krasnokutskaya et al. (2017) develop methods for estimating multi-attribute auctions and apply them to online freelance labor markets. Krasnokutskaya et al. (2017) uses the model to estimate welfare gains in the programming market from increased globalization through new online platforms. I estimate a multi-attribute auction model of the residential solar market and use the model to simulate changes in adoption and welfare associated with access to an online platform.

Economics of Renewable Energy: This paper also contributes to a broader literature on the costs and benefits of renewable energy. Gowrisankaran et al. (2016) quantify the economic value of large-scale renewable energy with specific attention to the intermittent nature of renewable generation. In another set of related papers, Callaway et al. (2017), Cullen (2013), and Novan (2015) empirically measure the fossil-fueled generation that is displaced by an incremental unit of power from wind or solar within a system. More recently, Borenstein (2017) investigates the interacting role of electricity rate structure and government subsidies on consumer adoption of solar PV, as well as the distributional effects of subsidy programs. I focus on estimating the size of markups in the solar industry and measuring the welfare effects of providing consumers with additional price information via an online platform.

In the next section, I discuss the data used for the analysis.

2 Data and Setting

Shopping for a rooftop solar system can be a time-consuming endeavor. Installers typically do not post prices because installation costs can vary depending on input costs, location, rooftop characteristics, size of the system, the efficiency of the modules, and other factors. Therefore, inquiring buyers often need to make phone calls to individual installers and then schedule a site visit just to obtain a project proposal and a price quote. Because search is costly in this market, many buyers may only receive a limited number of price quotes, which increases the incentive for sellers to exercise market power. This market environment also poses a challenge for researchers aiming to document market power because it is hard to observe how many installers are actually in a buyer's choice set. Most publicly available data on solar installations only includes the chosen seller, so it is impossible to know which other installers the customer was considering.

In this study, I directly observe buyers' choice sets by collecting proprietary installer price quote and consumer purchase data from EnergySage Inc. EnergySage operates an online platform that connects potential solar customers to a network of solar PV installers. The platform is the first online medium that allows buyers to shop and compare solar system offers from different sellers in one place. The company has received funding and support from support the U.S. Department of Energy, New York State Energy Research & Development Authority, Connecticut Green Bank, and the Massachusetts Clean Energy Center. Furthermore, the platform became the official solar marketplace for the state of Connecticut through GoSolarCT.com in 2016.

The EnergySage platform allows households interested in installing solar PV to conduct their own multi-attribute auctions to select an installer for their project. Each EnergySage "auction" includes several stages. First, consumers create an account with the website and provide necessary information such as the physical address of the potential installation and a monthly electricity bill.[13] Second, registered installers[14] receive a notification of the project which includes details such as a Google Maps photo of the buyer's roof (depicted in Figure 8 of the Appendix), as well as the monthly electricity usage of the buyer. Installers are then able to submit a project quote to the buyer which includes the system price, panel brand, inverter brand, and details about the seller such as a customer rating and a description of their solar installation experience. Finally, after installers have submitted their bids, the potential consumers can select one of the quotes and move forward with the transaction, or they can opt not to purchase any of the offers.[15] Figure 9 in the appendix shows an example of the purchaser's comparison tool on the platform. A distinguishing feature of this environment is that buyers can base their selection off any criteria they choose and are not obligated to purchase the quote with the lowest price. Although many installers offer leases or power purchase agreements, 97% of buyers on EnergySage choose to purchase a system with cash or through a loan. The large skew towards purchases (instead of leases) is likely because EnergySage provides a calculation of the net present value of each offer and purchased systems nearly always offer a higher overall value to the buyer. Additionally, overall market shares for leased systems have been declining considerably in recent years. Finally, most bids on the EnergySage marketplace are from local installers or regional installers. SolarCity, the nation's largest installer, leases out the majority of their systems but is not active on the platform.

2.1 Descriptive Statistics

The analysis includes potential solar PV projects that occurred in 2015 and 2016 within the states of California, Connecticut, Massachusetts, and New York. I only consider residential projects and drop any projects that are smaller than 2KW or larger than 20KW in capacity.[16] I

[13] Alternatively, users can submit an estimate of their monthly electricity use.

[14] In order to submit bids, installers must first be pre-screened by EnergySage to ensure they are licensed, insured, and experienced.

[15] Buyers and sellers can communicate with each other via private messaging or phone calls before a selection is made. However, sellers cannot call buyer unless they are requested to do so by the buyer.

[16] Each installer provides a system size in their project quote; these quotes are usually very similar (a regression of system size on project ID fixed effects has an R^2 of 0.87). Therefore, I define each project's size as the mean size quote for that project. Nearly 99% of projects are between 2KW and 20KW in size.

also do not include projects where a lease or power purchase agreement was selected.[17]

Table 1: Project Summary Statistics

	Mean	SD	10-%tile	50-%tile	90-%tile
Number of Bids	3.95	1.90	1.00	4.00	6.00
Project Size (Watts)	7155.92	3138.63	3631.67	6560.00	11515.00
Distinct Panel Brands	2.80	1.28	1.00	3.00	4.00
# of Bids w/ Premium Panel	1.49	1.40	0.00	1.00	3.00
# Bids w/ Premium Plus Panel	0.16	0.41	0.00	0.00	1.00
# of Bids w/ Microinverter	2.93	1.85	1.00	3.00	5.00
# of Bids from Permanent Sellers	3.35	1.78	1.00	3.00	6.00
Observations	10545				

After dropping commercial projects and projects over 20KW in size, the EnergySage data set includes 10,545 potential projects or "auctions". Figure 2 maps the locations of all of the projects in the sample and Table 1 provides descriptive statistics of the projects.[18] The average project received just under four bids. However, there is noticeable variation in the number of bids across projects, 90% of projects received between 1 to 6 bids.[19] The size of potential projects also varies largely across projects. The median project had a 6,560-watt capacity, but the standard deviation was over 3000 watts, with 80% of projects between 3,631 watts and 11,515 watts. Buyers also often have quotes for several different types of panel (module) brands.[20] EnergySage also shows buyers a rating classification of each panel brand, premium panels have higher efficiency and better warranties.[21] More efficient panels are attractive because for a given physical system size[22] a more efficient panel will create more electricity. The average buyer receives 1.49 bids that are "premium" panels.[23] EnergySage also identifies the very highest quality panels as "premium plus",[24] however, these panels are much more expensive and less than 15% of households receive a bid offering a "premium plus" panel. Another vital component of a solar system is the inverter. The inverter converts the direct current (DC) output of the PV panel into alternating current (AC). String inverters are the cheapest and most commonly deployed inverter technology and can perform well if there is no shade at the project location at any time during the day. However, a system with a string inverter will only produce as much

[17] I drop these projects because comparing per-watt prices for leases vs. purchases is not straightforward. Furthermore, these projects compose less than 4% of choices and thus discluding them is unlikely to have major effects on the analysis. I also drop a handful of price quotes below $2/watt that appear to be miscoded.

[18] Also see Appendix Figure 13 for a map of the installers' imputed locations.

[19] See Figure 10 in the appendix for a histogram of the number of bids per project.

[20] A PV module consists of many PV cells wired in parallel. A panel can consist of one or more modules and is the largest hardware component of a system in terms of size and cost.

[21] EnergySage designates LG Electronics panels as "premium".

[22] Physical size is distinct from capacity, if two panels with the same capacity but one is more efficient, the more efficient panel will be physically smaller.

[23] Each seller is only allowed to place a single bid. For example, a seller cannot place two different bids for different panel qualities.

[24] Panels from SunPower Corporation are "premium plus".

Figure 2: Potential Project Locations

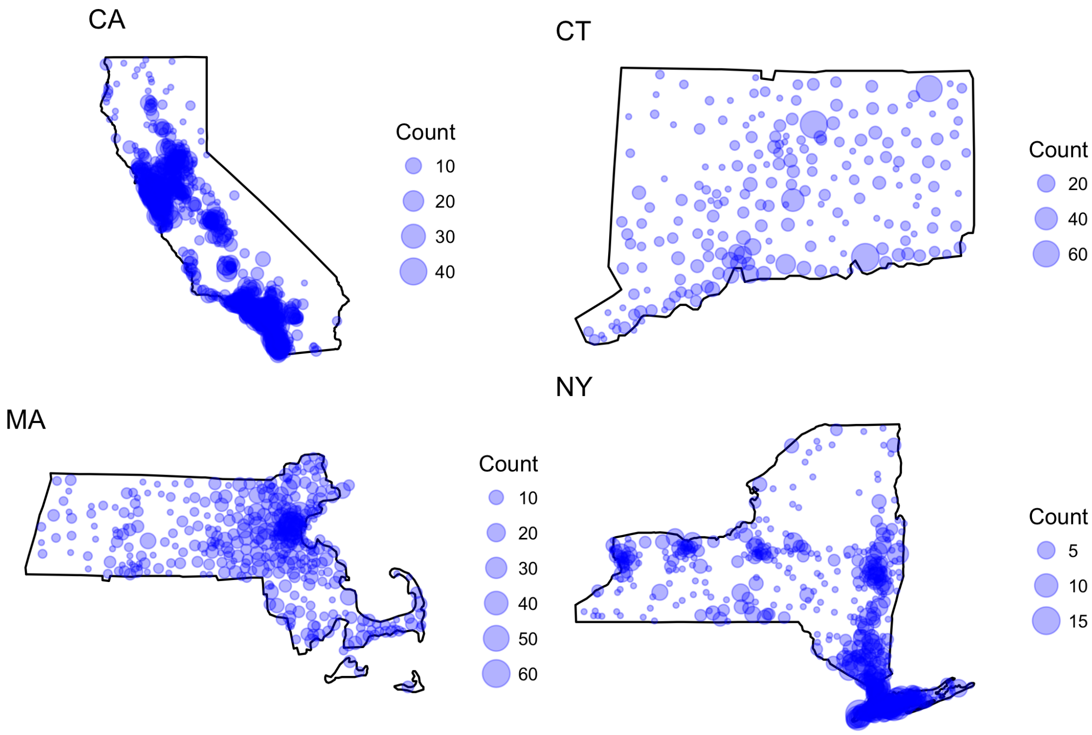

Notes: Count is the total number of potential projects within a ZIP code during the full sample.

as the least productive panel in an array, leading to reduced output if shade covers part of the roof. Microinverter and power optimizer technologies can help the system to perform better in partial shade conditions but typically are more expensive. Some quotes in the data include microinverter technologies (in the analysis, I include power optimizers in the microinverter category), the average household received just under three bids with microinverters.[25]

Projects also receive quotes from installers with varying levels of experience within the platform. The average buyer received 3.35 bids from "permanent" sellers, which I define as an installer that made at least 100 bids and at least one sale on the platform during the sample. Potential customers can see reviews and star ratings of each installer, how many years the installer has been installing solar, and how many total projects they have completed. Given this information, buyers can form perceptions about a seller's quality. Sellers' quality is thus an important component to account for in the empirical analysis. Unfortunately, the installer identities and star ratings over time are not available in the data. However, I do observe a unique installer identification number that can be used to track the behavior and performance of the same installer over time. I also observe a cross-section of each seller's star ratings and total installation experience. The seller ratings and experience information was collected in

[25]The data does not distinguish explicitly between microinverter and string-inverter bids but does list the inverter brand. I define a bid as having a microinverter if the inverter brand is Enphase Energy or SolarEdge Technologies. These two companies together controlled 95 percent of the module-level power electronics market in 2015.

September 2017, nine months after the sample period ended. I report summary statistics for the installers in Table Appendix 11. As of September 2017, 76% of sellers had ratings on EnergySage, and of them, 84% had a 5-star rating. The average seller had seven ratings, had completed 1,579 total residential installations, and had been installing solar for nine years. In the empirical section, I discuss my approach used to account for seller quality in the analysis.

The number of active installers participating in auctions in each state increases over time (see Appendix Figure 11). At the beginning of the sample, each state had between 17 and 28 distinct installers, and by the end of 2016 each state had between 25 to 58 different sellers submitting bids.[26] A substantial portion, 70%, of the 220 installers are transient sellers that made fewer than 100 total bids. However, since each permanent seller bids much more often, bids from transient sellers only compose about 13% of all price quotes.

Table 2: Summary Statistics - Bid Characteristics

Panel A: Full Sample	
Price ($/watt)	3.611
	(0.494)
Premium Panel (0,1)	0.386
	(0.487)
Premium Plus Panel (0,1)	0.0412
	(0.199)
Microinverter (0,1)	0.761
	(0.426)
Permanent Seller (0,1)	0.871
	(0.335)
Observations	40575

Panel B: CA South - 2016H1	Selected Bid (0,1)	
	0	1
Price ($/watt)	3.554	3.442
	(0.351)	(0.213)
Premium Panel (0,1)	0.542	0.714
	(0.498)	(0.454)
Premium Plus Panel (0,1)	0.0153	0.0357
	(0.123)	(0.187)
Microinverter (0,1)	0.792	0.881
	(0.406)	(0.326)
Permanent Seller (0,1)	0.928	0.976
	(0.258)	(0.153)
Observations	4273	

Across the 10,545 auctions, 40,575 bids were submitted during the two-year sample. Table 2 provides summary statistics of the submitted bid characteristics. Panel A shows that the mean price bid was $3.61 per watt,[27] additionally, 38% of bids included premium panels, and 76% of bids included microinverters. Panel B shows bid characteristics separately for selected and non-selected bids for a single market, Southern California in the first half of 2016. Not-surprisingly, quotes that are selected are lower in price on average. However, buyers also seem to care about other attributes. Selected quotes are more likely to be from permanent installers, more likely to have premium or premium plus panels, and are more likely to include microinverters. Indeed,

[26] Some installers participate in auctions in multiple states.

[27] This is the average gross bid price and does not include the 30% ITC or any state rebates.

many buyers chose a quote that was not the lowest priced.[28] Because consumers appear to value non-price features of solar systems, I cannot model the consumer's decision as a first-price auction. A solar PV installation is a differentiated product, and therefore firm's optimal bidding strategy will depend on customers' preferences for different attributes of the installation.

One feature that is apparent in the EnergySage price quotes that is also salient in other solar PV sales data is the presence of significant price variation. System prices vary widely across both space and time.[29] Figure 12 in the appendix displays a kernel density plots of price quotes for each state. In Southern California, price quotes range from as low as \$3/watt to over \$4/watt. The average system size is 7,155 watts; this can mean differences in overall system prices of over \$7,000. These price differences likely result from differences in costs across installers, differences in unobserved installer quality, differences in unobserved project characteristics (i.e., different roof types), as well as differences in the amount of competition (number of bids) for specific projects. In the next subsection, I investigate the relationship between the amount of competition and prices, as well as differences between prices paid on and off the EnergySage platform.

2.2 Relationship Between Competition and Prices

Before moving to the structural model, it is useful to consider the reduced-form relationship between market structure and solar PV prices. First, I compare purchase prices of solar systems on EnergySage to prices paid for comparable installations off of the platform in the same geographic area. Next, I consider the link between the number of bidders in an auction and bid prices within the EnergySage platform.

2.2.1 Prices on the Platform vs. Prices Outside the Platform

The online platform gives consumers access to more potential installers than they would have in its absence. Likewise, installers are aware that customers on the platform likely see more competing price bids, which could cause themselves to place lower bid prices. To investigate the effect of the platform on equilibrium prices paid, I collect additional data from the Lawrence Berkeley National Lab and the National Renewable Energy Laboratory's Open PV dataset. The Open PV Project is a collaborative effort between government, industry, and the public to assemble data on Solar PV installations across the United States. The data comes from solar incentive programs, utilities, installers, and other groups.

To compare equilibrium prices paid on the platform to prices off the platform, I append the Open PV data for California, Connecticut, Massachusetts, and New York with the EnergySage

[28]This pattern is also apparent in almost any given market (not just Southern CA in 2016H1), however the pattern is not as clear when looking at selected versus non-selected bids for the full sample. This is because average prices and attributes are changing over time and so is the overall probability of purchase.

[29]See Tables 9 and 10 in the appendix for the project and bid summary statistics by state.

data for selected bids (EnergySage bids that were purchased by a consumer).[30] I estimate the following model:

$$P_i = \alpha \mathbb{1}[Platform]_i + \beta X_i + \nu_{st} + \lambda_z + \varepsilon_i \tag{1}$$

Where P_i is the price paid for system i in dollars per watt (or it's logarithm), $\mathbb{1}[Platform]_i$ is an indicator function that takes the value of one if the system was purchased via EnergySage and otherwise equals zero, X_i is a vector of system characteristics such as system size, panel quality dummies, and a microinverter dummy. Finally, ν_{st} and λ_z are state-time, and ZIP code fixed effects respectively, and ε_i is an idiosyncratic error.

Table 3: Effect of the EnergySage Platform on Installed Prices ($/watt)

	(1) Price	(2) Price	(3) ln(Price)	(4) ln(Price)
Platform (0,1)	-0.767***	-0.666***	-0.169***	-0.146***
	(0.0366)	(0.0350)	(0.00793)	(0.00754)
Controls	No	Yes	No	Yes
State-Time FE	Yes	Yes	Yes	Yes
Zipcode FE	Yes	Yes	Yes	Yes
N	135341	135341	135341	135341
R^2	0.163	0.238	0.159	0.243

Notes: Controls include system size, panel quality (rated by EnergySage), and a dummy for if the system includes a microinverter or power optimizer. All standard errors are listed in parenthesis. * indicates $p < 0.10$, ** indicates $p < 0.05$, and *** indicates $p < 0.01$

The regression results, displayed in Table 3, indicate a strong correlation between use of the platform and prices paid. In column 2, we see that average prices paid were over $ 0.66/watt lower on EnergySage than comparable systems off the platform. Column 4 gives estimates of the same model with the natural logarithm of price as the outcome variable; the results suggest that systems on EnergySage were sold at 14% lower prices.

There are several possible explanations for the difference in prices between EnergySage and off-line transactions. There could be differences in demand between EnergySage buyers and other buyers, perhaps EnergySage users are more price-elastic than the typical buyer. Another possibility is that EnergySage attracts installers that have lower costs. Finally, the degree of competition on EnergySage could be much higher. This is likely to be true if buyers outside the platform are only obtaining a single price quote. In this case, installers will be able to exercise market power and charge higher prices off the platform. In Section 3, I develop a structural model of buyer choice, seller bidding, and seller participation behavior which allows me to

[30] I only consider purchased residential systems between 2KW to 20 KW in size in both data sets. Also, it is likely that each of the EnergySage observations will also appear in the Open PV data. I attempt to deal with this issue by using a matching procedure to pair each observation in the EnergySage data with an observation that has similar observables in the Open PV data (same ZIP code, same time-period, similar price, similar size) and dropping the redundant observations. Since the EnergySage dataset is small relative to the Open PV dataset, the following regression results are similar even if I do not drop the redundant observations.

Figure 3: Average Number of Bids and Average Bid Prices

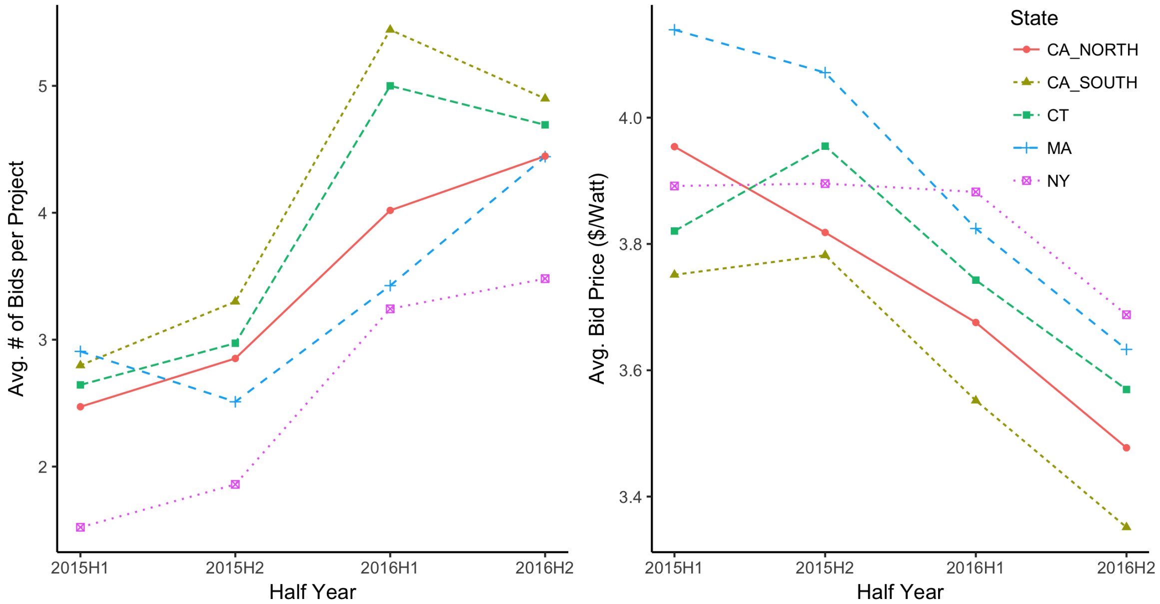

better identify the impact of increased competition on prices and adoption.

2.2.2 Competition and Prices within the EnergySage Platform

For projects within EnergySage, we can see a few notable trends in Figure 3. First, projects receive more bids on average over time. At the same time, bid prices have consistently declined. Additionally, bid prices are lower on average in states where consumers receive more bids. However, these trends could be explained by numerous factors and are not proof on their own that more competition will lead to substantial reductions in price. For example, declining prices over time could be caused by falling hardware input costs. Also, differences in prices across states could be driven by differences in labor costs or by differences in demand. I run the following regression to better isolate the effect of increased auction participation on bid prices:

$$B_{ij} = \alpha\mathbb{E}[Bids_j] + \beta X_{ij} + \gamma_j + \nu_{st} + \lambda_z + \varepsilon_{ij} \tag{2}$$

Where B_{ij} is the bid price in dollars per watt submitted by installer j in auction i. $\mathbb{E}[Bids_i]$ is the expected number of bids for auction i. I obtain the expected number of bids by taking the mean number of bids for projects in the same state and same quarter as project i. X_{ij} is a vector of bid and project characteristics such as system size, panel quality dummies, and a microinverter dummy. Finally, γ_j, ν_s, and λ_t are installer, state, and time (quarter of sample) fixed effects respectively, and ε_{ij} is an idiosyncratic error. I run several specifications with different combinations of the fixed effects.

I report the regression results in Appendix Table 12. In the first specification, I only include quarter-of-sample fixed effects and find a large negative relationship between the expected number of bids and prices. This means that after controlling for time-period shocks across states such as hardware input costs, states with more bids have lower prices. However, in Columns 2-4, we see that this strong negative relationship is not robust to the inclusion of state fixed effects or installer fixed effects. In Column 2, we can see that the specification with state and time-period fixed effects indicates a small but positive relationship between the number of bids and prices. In column 4, I include installer, state, and time-period fixed effects and the coefficient on the expected number of bids is almost zero and not statistically significant. In Appendix Table 13, I find similar results when using the realized number of bids (instead of expected bids) as the dependent variable. In these regressions, the estimated coefficients are negative but very small in magnitude.

These results highlight a clear simultaneity problem. In particular, installers are making multiple decisions; whether to register for the platform, whether to enter an auction for a specific project, and the price to bid conditional on "entering" an auction. If more consumers sign up for the platform or consumers' willingness-to-pay increases in a state, this will likely cause more installers to register for the platform in that state and could also encourage registered installers to bid on more projects. At the same time, if consumers have a higher willingness-to-pay, this will also encourage sellers to bid higher prices. This simultaneity issue will lead to biased estimates of the effects of competition (number of bids) on prices. The simultaneity problem, in this empirical context, is difficult to correct using control variables and fixed effects. To better isolate factors contributing to differences in prices and purchase decisions, I turn next to the structural model. Estimating a structural model will allow me first to identify seller's cost structure and buyer's preferences, then with seller's cost structure and buyer's preferences in hand, I can estimate how prices and adoption would change if access to the platform were removed.

3 Structural Model

Each buyer i seeks to procure installation services for a single indivisible project using a multi-attribute auction. Throughout the paper, i references both an individual buyer and their respective project. Buyer i's project is distinguished by its project type τ_i, which depends on the project's location, size, and time-period. For each project of type τ, there is a set $\mathcal{N}(\tau)$ of potential sellers that choose whether or not to submit a bid for the project.[31]

Each seller j is differentiated by their quality group index which belongs to a discrete set of $k+1$ values, $\mathcal{Q}_j = \{t, p_1, .., p_k\}$. A seller that only submits a limited number of total bids is considered to be a "transient" seller and belongs to quality group t. All other sellers that submitted a sufficiently large number of total bids and made at least one sale are considered to be "permanent" sellers and belong to one of k quality groups $p_1, .., p_k$, with p_k denoting the

[31]In the empirical implementation, I consider any seller that entered at least one bid for a project of type τ a potential seller for project τ_i.

highest quality. Each seller's quality group index is observable to both the buyer and the other potential sellers. If a seller chooses to participate in the auction for project i he then also selects a price bid B_{ij}. Each seller's bid is also characterized by a vector of non-price characteristics $\mathbf{x}_{ij}$ such as panel quality and inverter type. $\mathbf{x}_{ij}$ is allowed to vary across projects for a given seller.

3.1 Demand: Buyer's Choice Problem

The allocation rule in a multi-attribute auction comes from the buyer's choice problem. In this subsection, I outline the buyer's choice problem. Then, in a later subsection, I describe how I estimate buyers' preferences using discrete choice methods.

Let $\mathcal{K}_i \subset \mathcal{N}(\tau_i)$ be the set of sellers that decide to participate in the auction for project i. Buyer i then chooses between the project bids and an unspecified outside option (k^0) to maximize their utility. Household i's utility from selecting project j is given by:

$$u_{ij} = \alpha B_{ij} + \mathbf{x}'_{ij}\beta + \delta_\tau + \zeta_{ig} + (1-\lambda)\varepsilon_{ij} \tag{3}$$

Here B_{ij} is the bid price for option j, $\mathbf{x}_{ij}$ contains observable characteristics of the system, such as the panel brand's quality, inverter type, the size of the system, and the installer quality group index. δ_τ is a demand shifter for projects of type τ that allows utility for all of the "inside options" to vary depending on location, time-period, and project size. ε_{ij} is an independent and identically distributed random term that is assumed to follow a type-one extreme value distribution. ζ_{ig} is also an idiosyncratic term but is assumed to be constant for each buyer across all the "inside options". ζ_{ig} follows the unique distribution distributed such that $\zeta_{ig} + (1-\lambda)\varepsilon_{ij}$ is also an extreme value random variable. This utility specification gives rise to the well-known nested logit model. The nested logit model allows for more flexible substitution patterns in comparison to the standard logit model because it accommodates correlation in preferences for products within pre-specified groups. Here, I specify one group to be the "outside option," and the other group to contain all of the project bids. Some households may register for the platform just out of curiosity about solar PV prices and may not be serious about making a purchase. Likewise, there may be customers that are very adamant about buying a solar PV system. Therefore, these consumers would be unlikely to select the outside option even if some of the options in their choice set were removed. The nested logit model allows for these types of individuals. As λ approaches zero, there is no correlation in preferences for the "inside option", and the model reduces to the standard logit model. As λ goes to one, there is a perfect correlation in preferences for each "inside option". Finally, the overall level of utility is not identified, so I normalize the utility of the outside option to equal zero plus an error term, this normalization is standard in the literature.

Given the utility specification in equation 3, the probability that household i chooses project j is:

$$Prob_{ij} = \frac{exp\Big(\frac{\alpha B_{ij}+\mathbf{x}'_{ij}\beta+\delta_\tau}{1-\lambda}\Big)\sum\limits_{k\neq k^0} exp\Big(\frac{\alpha B_{ik}+\mathbf{x}'_{ik}\beta+\delta_\tau}{1-\lambda}\Big)^{-\lambda}}{1+\sum\limits_{k\neq k^0} exp\Big(\frac{\alpha B_{ik}+\mathbf{x}'_{ik}\beta+\delta_\tau}{1-\lambda}\Big)^{1-\lambda}} \tag{4}$$

Where the sum is over all bids $k \in \mathcal{K}_i$ in the individual's choice set except the outside good k^0.

3.2 Supply: Seller Bidding and Participation Decision

The supply-side model has several fundamental differences from a standard differentiated products model. First, firms must make an explicit decision about whether to submit a price quote to each potential buyer. Second, the sellers do not have information about exactly how many competing suppliers will make bids to the customer. Moreover, the suppliers do not have perfect information about the identity and characteristics of the competitors they will face, nor about the price quotes those competitors will submit. Firms cannot see the exact identity of installers that offer bids for a particular project. However, they can observe how many have been submitted ex-post. They also see which other firms participate on the platform in their area. Therefore, it is reasonable to assume that the suppliers know the distribution of possible competition they are likely to face for a given project. The model of firm behavior outlined below accounts for these features of the online platform.

I model suppliers bidding behavior as a two-stage process. In the first stage, each potential bidder $j \in \mathcal{N}(\tau_i)$ must decide whether or not to enter the auction for the project i. At the time of entry, firms do not know their exact marginal cost of completing the project, but they know the distribution of possible costs they could incur. They also know the probabilities of each of their opponents entering the auction, the characteristics of those opponents and the distribution of possible prices those opponents would submit conditional on entry. Additionally, they know the mean utility of the buyer (but not the random component of utility).[32] Therefore, each firm can form an expectation about their profits conditional on the decision to enter the auction. If firm j decides to enter the auction for the project, they incur a bid preparation cost $S(\tau_i, \mathcal{Q}_j) + \eta_{ij}$, where $\eta \sim Lognormal(0, \sigma^2(\mathcal{Q}_j))$. The bid preparation cost contains a deterministic component that depends on the project type and the seller's quality group, and a random component that's variance depends on the seller's quality group. I assume that the random component is i.i.d. across projects and firms and is private information of each potential bidder. If a firm decides to enter auction i they learn the non-price characteristics of their bid $\mathbf{x}_{ij}$ and the marginal cost of completing the project c_{ij}.

To make the model empirically tractable, I assume that the the non-price characteristics, x_{ij} including the size of the project bid are not strategic choices of bidders. This assumption means that firms are not choosing non-price characteristics such as panel quality and inverter type strategically when placing a bid. While this assumption is made primarily for tractability,

[32] In particular, I assume that sellers know all of the parameters of the buyer's utility function, α, β, δ, and λ.

the assumption also finds support in the data. For example, companies will typically use the same equipment for many consecutive projects. They may change module brands occasionally, but the hardware available to them to complete a given project is likely predetermined by their existing inventory. The practical interpretation of this assumption is that sellers need to check their existing product stock (which is predetermined) before knowing the exact non-price characteristics of their bid. They learn the non-price components of the bid by incurring the bid preparation costs.

The second part of the assumption means that the size of the project is pre-determined by the buyer and is not a choice variable for the seller. On the EnergySage platform, buyers submit a monthly electricity bill and sellers then choose the exact size of the system. In practice, the system size quotes are very similar across installers for a given project because almost all sellers size the system so that it will cover 100% of the buyer's annual electricity use. A regression of installers' system size quote on project dummies has an R^2 of 0.87, which means almost all of the system size choice can be explained by the project characteristics and is not likely a critical strategic variable.[33] Analogous to the other non-price characteristics, each firm learns the size of their project bid after incurring the bid preparation cost. Sellers know consumers' taste for size and can optimize their bid price based on the size draw that they receive.

When firms make their entry decision they do not know their marginal cost or non-price characteristics, but they do know the joint distribution from which their marginal cost and non-price characteristics will be drawn, $F_{CX|\mathcal{Q}_j,\tau_i}(c, \mathbf{x}|\mathcal{Q}_j, \tau_i)$. The distribution depends on both the sellers quality group and the project type. After the firms make their entry decisions in stage one, each firm's marginal cost and non-price characteristics are drawn from $F_{CX|\mathcal{Q}_j,\tau_i}$ and the installer then decides on a price bid during the second stage.

3.2.1 Sellers' Bid Pricing Problem

It will be helpful to first consider the firm's problem in the second stage after marginal costs and non-price characteristics are realized. Conditional on entering an auction, the firm j solves the following problem when setting a bid price for project i:

$$\max_{B_{ij}} \quad [B_{ij} - c_{ij}]\mathcal{P}_{ij}(B_{ij}, \mathbf{x}_{ij}, \mathcal{Q}_j|\tau_i) \tag{5}$$

Where B_{ij} is firm j's price bid, c_{ij} is firm j's marginal cost, and $\mathbf{x}_{ij}$ are the firm's non-price characteristics for project i. $\mathcal{P}_{ij}(B_{ij}, \mathbf{x}_{ij}, \mathcal{Q}_j|\tau_i)$ is the equilibrium probability of winning the auction conditional on placing placing a bid price of B_{ij}, having non-price characteristics $\mathbf{x}_{ij}$, and belonging to quality group $\mathcal{Q}_j$. The equilibrium expected probability of being selected is a function of the type of project τ_i. We work with expected probabilities because the seller does not know exactly which competitors he will face nor the bids of those competitors.

[33]In theory, it would be possible to include size as an additional choice variable of the seller. Then the seller would have two first-order conditions for an optimal bid (size and price). However, this would likely lead to multiple equilibria for sellers in the bidding stage as numerous price-size pairs may satisfy the first-order-conditions.

When formulating firms expectations, I assume that all sellers submit bids simultaneously. This assumption means that each firm independently submits a bid at the same instant. Therefore, the installers do not know the exact number of bidders they will be competing against nor the identities of their competitors. Thus, firms' expectations (about the probability of winning) will only be a function of the project type, conditional on the price and non-price characteristics of their bid. In practice, firms on the platform submit bids at slightly different times. Although the identities of competing bidders is not visible to auction participants, firms can see how many bids have already been submitted for a given auction. Therefore, it is possible that firms could update their expectations based on the number of bids that have already been submitted. The assumption of simultaneous bidding is made primarily to simplify computation in the empirical exercise. However, I provide evidence that the assumption is a reasonable approximation of firms' behavior. In Table 14, I regress bid price on the order that a bid was submitted, controlling for the total number of bids, installer fixed effects, state fixed effects, and time-period fixed effects. The coefficients on "order of bid" is small and not significant. This suggests firms are not making significant changes in bidding strategy based on the order they submitted a bid.

Under the assumption of simultaneous bidding, a firm's expected probability of winning $\mathcal{P}_{ij}$ can be expanded as follows:

$$\begin{aligned}&\mathcal{P}_{ij}(B_{ij}, \mathbf{x}_{ij}, \mathcal{Q}_j|\tau_i) = \mathbb{E}[Prob_{ij}(B_{ij}, \mathbf{x}_{ij}, \mathcal{Q}_j; \mathbf{B}_{i,-j}, \mathbf{X}_{i,-j}, \boldsymbol{\mathcal{Q}}_{-j}|\tau_i)] = \\ &\int Prob_{ij}(B_{ij}, \mathbf{x}_{ij}, \mathcal{Q}_j; \mathbf{B}_{i,-j}, \mathbf{X}_{i,-j}, \boldsymbol{\mathcal{Q}}_{-j}|\tau_i) \cdot G(\mathbf{B}_{i,-j}, \mathbf{X}_{i,-j}, \boldsymbol{\mathcal{Q}}_{-j}|\tau_i) \cdot d(\mathbf{B}_{i,-j}, \mathbf{X}_{i,-j}, \boldsymbol{\mathcal{Q}}_{-j}, |\tau_i)\end{aligned} \tag{6}$$

Recall that $Prob_{ij}$ is the probability that buyer i selects firm j's bid conditional on realized set of competitors submitting a vector of price bids $\mathbf{B}_{i,-j}$, having a stacked vector of non-price characteristics $\mathbf{X}_{i,-j}$, and having quality indices $\boldsymbol{\mathcal{Q}}_{-j}$. G represents the joint probability distribution of $\mathbf{B}_{i,-j}, \mathbf{X}_{i,-j}$, and $\boldsymbol{\mathcal{Q}}_{-\boldsymbol{j}}$ occurring in equilibrium conditional on the project being of type τ_i. Since each firm's entry draw and marginal cost draw is assumed to be i.i.d., we can express G as the product of the probabilities that each individual competing firm l decides to enter the auction and then bids B_{il} and has non-price characteristics x_{il}. I define the optimal bid function as $B^*_{il}(c_{il}|\mathbf{x}_{il}, \mathcal{Q}_l, \tau_i)$ and $H(\mathcal{Q}_l, \tau_i)$ is the probability that a potential seller l that is of quality $\mathcal{Q}_l$ enters an auction of type τ_i. Then we have:

$$\begin{aligned}&G(\mathbf{B}_{i,-j}, \mathbf{X}_{i,-j}, \boldsymbol{\mathcal{Q}}_{-\boldsymbol{j}}|\tau_i) \cdot d(\mathbf{B}_{i,-j}, \mathbf{X}_{i,-j}, \boldsymbol{\mathcal{Q}}_{-j}|\tau_i) = \\ &\prod_{l \in \mathcal{N}(\tau_i) \setminus \{j\}} H(\mathcal{Q}_l, \tau_i) \cdot \int F_{CX|\mathcal{Q}_l,\tau_i}\big(B^{*-1}(B_{il}|\mathbf{x}_{il}, \mathcal{Q}_l, \tau_i), \mathbf{x}_{il}|\mathcal{Q}_l, \tau_i\big) d(B_{il}, \mathbf{x}_{il})\end{aligned} \tag{7}$$

Where B^{*-1} represents the inverse bid function. The expression inside the product is the probability that firm l enters the auction multiplied by the probability that firm l bids B_{il} and has non-price characteristics x_{il}.

Firm i's first-order condition for an optimal bid is given by:

$$(B_{ij} - c_{ij})\frac{\partial \mathcal{P}_{ij}(B_{ij}, \mathbf{x}_{ij}, \mathcal{Q}_j|\tau_i)}{\partial B_{ij}} + \mathcal{P}_{ij}(B_{ij}, \mathbf{x}_{ij}, \mathcal{Q}_j|\tau_i) = 0 \tag{8}$$

Given a vector of non-price characteristics, the optimal bid function $B^*_{il}(c_{il}|\mathbf{x}_{il}, \mathcal{Q}_l, \tau_i)$ is defined implicitly by equation 8.

3.2.2 Sellers' Participation Decision

Now consider the firm's decision of whether or not to enter an auction. Each firm will enter if the expected marginal profits conditional on entering are larger than the fixed cost of bid preparation $S(\tau_i, \mathcal{Q}_j) + \eta_{ij}$. When firm j decides to enter auction i, they only know the project type and their own quality index, and their private entry cost draw. Firm j's expected profits conditional on entering the auction for project i can be expressed as follows:

$$\mathbb{E}[\pi_{ij}|\mathcal{Q}_j, \tau_i] = \int \Big[(B^*_{ij}(c_{ij}|\mathbf{x}_{ij}, \mathcal{Q}_j, \tau_i) - c_{ij})\mathcal{P}_{ij}(B^*_{ij}, \mathbf{x}_{ij}, \mathcal{Q}_j|\tau_i)\Big] dF_{CX|\mathcal{Q}_j,\tau_i}(c_{ij}, \mathbf{x}_{ij}|\mathcal{Q}_j, \tau_i) \tag{9}$$

Recall that $F_{CX|\mathcal{Q}_j,\tau_i}(c, \mathbf{x}|\mathcal{Q}_j, \tau_i)$ is the joint distribution of unknown marginal shocks and non-price characteristics whose realization is not known to the firm at the time of entry. Therefore, the firm will enter the auction as long as:

$$\mathbb{E}[\pi_{ij}|\mathcal{Q}_j, \tau_i] \geq S(\tau_i, \mathcal{Q}_j) + \eta_{ij} \tag{10}$$

Under the assumption that η_{ij} follows a lognormal distribution, the probability that firm j enters the auction for project i is:

$$H(\mathcal{Q}_j, \tau_i) = \Phi\left(\frac{ln\Big(\mathbb{E}[\pi_{ij}|\mathcal{Q}_j, \tau_i] - S(\tau_i, \mathcal{Q}_j)\Big)}{\sigma(\mathcal{Q}_j)}\right) \tag{11}$$

Where Φ represents the cumulative distribution function for a standard normal random variable.

Summary: Timing of the game

1. A potential buyer i initiates a multi-attribute auction by announcing the project type τ_i to all potential entrants $\mathcal{N}(\tau_i)$.

2. Each potential seller $j \in \mathcal{N}(\tau_i)$ receives a private entry cost shock η_{ij}. Each potential entrant then compares their total entry cost $S(\tau_i, \mathcal{Q}_j)+\eta_{ij}$ to the expected marginal profit

conditional on entering the auction $\mathbb{E}[\pi_{ij}|\tau_i, \mathcal{Q}_j]$. Each potential bidder chooses to enter if and only if expected marginal profits are larger than their entry cost.

3. Each seller that enters auction i receives a private marginal cost draw c_{ij} and also learns their non-price characteristics $\mathbf{x}_{ij}$. Sellers do not observe which other competitors have entered the auction. Each entrant then chooses a bid price B_{ij}.

4. Buyer i chooses from each of the project bids or the outside option.

3.3 Equilibrium

For each seller j a strategy consists of two functions: a participation strategy $\mathcal{Q}_j \times \tau_i \times \mathbb{R}_+ \to \{0, 1\}$, and a bidding strategy $\mathcal{Q}_j \times \tau_i \times \mathbb{R}_+ \times \mathbf{x}_{ij} \to \mathbb{R}_+$. Specifically, sellers use information about the project type, their quality group, and their entry cost shock to determine the binary choice of whether or not to enter. In the bidding stage, firms consider the project type, their quality group, their marginal cost draw, and their non-price characteristics to form a price bid. I follow the convention in the literature by focusing on type-symmetric pure strategy Bayesian equilibrium (Krasnokutskaya et al., 2017). That is, all sellers of the same quality index use the same participation strategy in equilibrium, and all sellers of the same quality index and same non-price-characteristics use the same bidding strategy in equilibrium. An equilibrium in the participation stage is a strategy profile such that all sellers satisfy the inequality in 10, given the strategies of other firms. An equilibrium in the bidding stage requires that all firms satisfy equation 8 given the other installer's strategies. Krasnokutskaya et al. (2017) prove the existence of a type-symmetric pure strategy Bayesian equilibrium of this game. However, there is no guarantee of a unique equilibrium in the participation stage.

3.4 Estimation and Implementation Details

I estimate the structural parameters in three steps. First, I solve for the demand parameters using a two-stage grouped fixed effects approach. Second, I use the estimated demand parameters to simulate each firm's first-order conditions for each bid in the data and recover bid-specific markups. Finally, I use the estimates from the first two steps to calculate each potential bidders expected marginal profits from entering each auction in the data and estimate the entry cost parameters using observed entry decisions. I discuss the details of each step in the following subsections.

3.4.1 Demand Estimation and Seller Quality Groups

I estimate the demand parameters via maximum likelihood. Each buyer's utility is allowed to depend on bid price (\$/watt), panel quality (dummies for premium and premium plus modules), inverter type (microinverter dummy), the capacity of the system, state and time fixed effects (demand shifters for all the inside options). The price that enters the buyer's utility is equal 70%

of the installer's quoted price to account for the 30% ITC. None of the states besides New York had changes in subsidy rates during the sample period so any state rebate should be accounted for by the state fixed effects. In a later section, I show that the results are robust to adding explict controls for state-level incentives. I also allow utility to vary by installer by including installer fixed effects for each of the 65 installers that placed over 100 bids and made at least one sale. Let M be the total number potential buyers in the sample, then the log likelihood function is:

$$LL(\alpha,\beta,\delta,\lambda) = \sum_{i}^{M}\sum_{j\in\mathcal{K}_i} \mathbb{1}[i\ choose\ j]\cdot ln(Prob_{ij}) \tag{12}$$

$$= \sum_{i}^{M}\sum_{j\in\mathcal{K}_i} \mathbb{1}[i\ choose\ j]\cdot ln\left(\frac{exp\left(\frac{\alpha B_{ij}+\mathbf{x}'_{ij}\beta+\delta_\tau}{1-\lambda}\right)\left(\sum_{k\neq k^0} exp\left(\frac{\alpha B_{ik}+\mathbf{x}'_{ik}\beta+\delta_\tau}{1-\lambda}\right)\right)^{-\lambda}}{1+\left(\sum_{k\neq k^0} exp\left(\frac{\alpha B_{ik}+\mathbf{x}'_{ik}\beta+\delta_\tau}{1-\lambda}\right)\right)^{1-\lambda}}\right)$$

After solving for the maximum likelihood estimates, I sort each of the sellers into quality groups. If a seller submitted less than 100 total bids during the sample, they are placed in the "transient" seller group. If a seller made more than 100 bids during the sample they are considered a "permanent" seller and are placed into one of ten quality groups based on their fixed effects. In particular, I sort each firm into a group based off of the decile of its fixed effect estimate. For instance, if we ordered each firm by the fixed effect estimates and firm j was in the 7th decile, the seller would be placed into "Permanent - Quality Group 7". Let $\hat{\varphi}_j$ denote installer j's fixed effect estimate, then quality group $\mathcal{Q}^l$ is defined as follows:

$$\mathcal{Q}^l \equiv \{j \quad | \quad j \in Permanent \quad \cap \quad Pr[X \leq \hat{\varphi}_j] = l/10\} \tag{13}$$

Each seller's fixed effect will be higher if he systematically wins auctions more often than transient sellers conditional on price and non-price characteristics of his bids. I interpret sellers that consistently perform better (conditional on price bids) as having higher unobserved quality. Unobserved quality will include things like star rating, installation experience, customer service, and warranty services offered.

After sorting the installers into the eleven quality groups, I re-estimate the demand model but now including fixed effects for each seller quality group in the buyers' utility function instead of the full set of firm fixed effects. Figure 4 summarizes the non-price characteristics of each installer's bid (that enter buyer utility).

There are several reasons for sorting the sellers into quality groups. For one, it allows me to group several installers to calculate each seller's expected marginal profit from entering a given auction. Computing expected marginal profits is required to estimate the entry cost parameters. If I did not sort the sellers into groups, then each permanent seller would have a different expected marginal profit from entering a given auction (because they each would have different

Figure 4: Non-Price Characteristics of Bids

$$\begin{array}{c} \text{Quality Group} \\ \begin{bmatrix} \text{Transient} & \text{Quality Grp. 3} \\ \text{Quality Grp. 1} & \ldots \\ \text{Quality Grp. 2} & \text{Quality Grp. 10} \end{bmatrix} \end{array} \times \begin{array}{c} \text{Panel Type} \\ \begin{bmatrix} \text{Standard} \\ \text{Premium} \\ \text{Premium Plus} \end{bmatrix} \end{array} \times \begin{array}{c} \text{Inverter} \\ \begin{bmatrix} \text{String (Standard)} \\ \text{Microinverter} \end{bmatrix} \end{array}$$

Notes: When each seller places a bid in an auction, their quote will also include a vector of seller characteristics that is observable to the buyer. The panel quality and inverter type are observed in the data and the seller's quality group is determined by sorting firms by the first-stage installer fixed effects, The system size is also included in the non-price characteristics of the bid. The microinverter category also includes power optimizers.

probabilities of winning conditional on entry). This is problematic because some sellers may only participate in a few auctions of a particular type (i.e. Installer X may just bid on three large projects in Connecticut in 2016Q1) which means that the estimated marginal profit would have a very high variance due to being calculated using only a handful of observations. Similarly, some of the seller fixed effects may be estimated imprecisely. A seller with just slightly more than 100 bids may have gotten "lucky" and won auctions at a higher rate than their long-term winning probability; this would lead to a large fixed effect estimate and incorrect estimates of markups. In later sections, I discuss the robustness of the demand estimates to changes in the number of quality groups specified.[34] Finally, the quality groups allow for an easier economic interpretation of results compared to listing all of the individual firm fixed effects estimates. In a later section, I show that the demand estimates after grouping the sellers are very similar to the estimates with the full set of fixed effects.

3.4.2 Inferring Markups and Marginal Costs

In the next step, I recover a markup estimate for each bid in the data. To do so, I use the final demand estimates to form each firm's first-order condition for an optimal bid from equation 8. Notice that the FOC does not have a closed-form since it contains two expectations $\frac{\partial \mathcal{P}_{ij}(B_{ij}, \mathbf{x}_{ij}, \mathcal{Q}_j|\tau_i)}{\partial B_{ij}}$ and $\mathcal{P}_{ij}(B_{ij}, \mathbf{x}_{ij}, \mathcal{Q}_j|\tau_i)$. Therefore, we have to integrate the firm's probability of winning over different realizations of competitor sets and competitor bid prices that are unknown to the installer at the time of bidding. To understand the evaluation of each firm's FOC, consider an auction i that is of type τ_i. To evaluate the integrals, I follow the following procedure:

[34] I also discuss different estimation approaches such as using a K-means clustering algorithm to sort the sellers into groups and an iterated estimator that updates the quality groups at each stage until convergence.

1. First, obtain non-parametric estimates of the entry probabilities for each project type and each seller quality group. For example, the probability that "Permanent - Quality Group 7" sellers enter New York auctions in 2015Q2. This estimate is just the ratio of auctions entered divided by total auctions of that type. I assume a seller is a potential entrant for auction i if they entered at least one auction of type τ_i.

2. Next, use the probabilities from the previous step to simulate the entry decisions into auction i for each potential entrant in $\mathcal{N}(\tau_i)$.

3. Draw price bids and non-price characteristics for each of the simulated entrants using the empirical joint distribution of bids and non-price characteristics in the data. For example, if a type $\mathcal{Q}^l$ seller enters a simulated auction of type τ_i; then randomly draw a bid (both bid price and non-price characteristics together) from the pool of all bids placed by type $\mathcal{Q}^l$ sellers in auctions of type τ_i.

4. Evaluate the choice probabilities $Prob_{ij}$ and demand semi-elasticities $\frac{\partial Prob_{ij}}{\partial B_{ij}}$ inside the integrals given the bid prices and the competitors observed characteristics.

5. Repeat this process S times[35] and take the average of all the simulated choice probabilities, and simulated demand semi-elasticities to obtain estimates for the two expectations. Let s denote the simulation iteration, then the expressions are:

$$\widehat{\mathcal{P}_{ij}} = \frac{1}{S}\sum_{s=1}^{S} Prob_{ij}^s, \qquad \widehat{\frac{\partial \mathcal{P}_{ij}}{\partial B_{ij}}} = \frac{1}{S}\sum_{s=1}^{S} \frac{\partial Prob_{ij}^s}{\partial B_{ij}} \tag{14}$$

6. Finally, use the average choice probabilities, and average demand semi-elasticities to calculate the markup portion of each bid. The markup term for firm j in auction i is equal to $-\frac{\widehat{\mathcal{P}_{ij}}}{\widehat{\frac{\partial \mathcal{P}_{ij}}{\partial B_{ij}}}}$. Once we have an estimate of the markup term, the firm's FOC provides a one-to-one mapping that we can use to recover the marginal cost of each project in the data:

$$\widehat{c_{ij}} = B_{ij} + \frac{\widehat{\mathcal{P}_{ij}}}{\widehat{\frac{\partial \mathcal{P}_{ij}}{\partial B_{ij}}}} \tag{15}$$

This process allows me to infer a project-specific marginal cost for every bid in the data. I then use the estimated marginal costs to form a non-parametric cost distribution for each seller-project-type pair (e.g., a cost distribution "Permanent - Quality Group 7" sellers in New York 2015Q2 auctions).

[35] I simulate 1000 iterations of each auction type.

The choice of project type categories is crucial for obtaining credible estimates of markups. Defining project type categories exemplifies a trade-off between bias and variance. On the one hand, defining too few project types can bias markup estimates if projects are heterogeneous. For example, installations in New York will be different ex-ante than projects in Northern California because of differences in labor costs, permitting requirements, and differences in the set of possible competing bidders. Likewise, a New York project in 2015 will be different from a New York project in 2016 because of differences in hardware input costs, differences in consumer preferences, and differences in potential bidders. For this reason, we would not want to use bids placed in Northern California 2015Q1 when simulating a New York 2016Q3 auction because New York installers are not likely to be using these bids to form their expectations. However, if I define too many project categories, (i.e., each state-week has its own category), then markup estimates for each bid will have high variance because there is only a handful of projects to use to simulate realizations of each auction. Figure 5 displays the project type definitions used in the primary analysis. I specify 80 different project types determined by the location, time-period, and size of the project. I allow for projects to vary across the quarter of sample within a state to account for the rapidly changing competitive environment on the platform and evolving input costs over time. In later sections, I discuss the robustness of the results to changes in the definition of auction types.

Figure 5: Project Types

$$\begin{array}{ccccc} \text{Location} & & \text{Time Period} & & \text{Project Size} \\ \begin{bmatrix} \text{CA North} & \text{CT} & \text{NY} \\ \text{CA South} & \text{MA} & \end{bmatrix} & \times & \begin{bmatrix} \text{2015Q1} & \ldots \\ \text{2015Q2} & \text{2016Q4} \end{bmatrix} & \times & \begin{bmatrix} \text{Small } (\leq 6.5 \text{ KW}) \\ \text{Large } (> 6.5 \text{ KW}) \end{bmatrix} \end{array}$$

Notes: Each auction is categorized based of the location, time period, project size. There are $5 \times 8 \times 2 = 80$ project types. For example, {CT, 2016 Q1, Large} defines a single project type.

3.4.3 Entry Cost Parameters

In the final step. I use the estimated marginal costs to form each firms' *pre-entry* expected marginal profit from entering an auction i. For each bid in the data, I can calculate the firms' *post-entry* expected profit (before the buyer makes a choice) using the bid price, marginal cost, and probability of winning. The *post-entry* expected profit for seller j in auction i is equal to $(B_{ij} - \widehat{c_{ij}}) \cdot Prob_{ij}$. To calculate a seller's *pre-entry* expected profit $\widehat{\mathbb{E}[\pi_{ij}]}$ from entering an auction i; I take the average over all of the *post-entry* expected profits that are realized by the seller's quality group $\mathcal{Q}_j$ for projects of type τ_i. Define $N(\tau_i, \mathcal{Q}_j)$ as the total number of bids placed by type $\mathcal{Q}_j$ sellers in auctions of type τ_i, then *pre-entry* expected profit are estimated as:

$$\widehat{\mathbb{E}[\pi_{ij}]} = \frac{1}{N(\tau_i, \mathcal{Q}_j)} \sum_{i \in \tau_i} \sum_{j \in \mathcal{Q}_j} (B_{ij} - \widehat{c_{ij}}) \cdot Prob_{ij} \tag{16}$$

Here, I use $\sum_{i \in \tau_i}$ to mean the sum over all auctions of type τ_i and $\sum_{j \in \mathcal{Q}_j}$ to indicate the sum over all bids submitted by sellers of type $\mathcal{Q}_j$. Next, I use the *pre-entry* expected profits $\widehat{\mathbb{E}[\pi_{ij}]}$ to estimate the entry cost parameters via maximum likelihood. The log-likelihood is given by:

$$EntryLL(\mu, \sigma) = \sum_{i}^{M} \sum_{j \in \mathcal{N}(\tau_i)} \Bigg\{ \mathbb{1}[j\ enters\ i] \cdot ln \left(\Phi \left(\frac{ln \left(\widehat{\mathbb{E}[\pi_{ij}]} - \mu(\tau_i, \mathcal{Q}_j) \right)}{\sigma(\mathcal{Q}_j)} \right) \right) \tag{17}$$

$$+ \Big(1 - \mathbb{1}[j\ enters\ i]\Big) \cdot ln \left(1 - \Phi \left(\frac{ln \left(\widehat{\mathbb{E}[\pi_{ij}]} - \mu(\tau_i, \mathcal{Q}_j) \right)}{\sigma(\mathcal{Q}_j)} \right) \right) \Bigg\}$$

Where $\mathbb{1}[j\ enters\ i]$ is an indicator function that equals one if seller j enters auction i and is zero otherwise. I assume that μ is a linear function of project location (state dummies), seller type (quality group dummies), and across years (2016 dummy). I also allow σ to vary across permanent and transient sellers.

3.5 Identification

Identification of the demand parameters follows standard arguments. The utility specification includes quality-group fixed effects, state fixed effects, and time-period fixed effects. Therefore, identification of the coefficient on price, α, comes from variation in the prices quoted to different households by a given seller quality group, while also controlling for common demand shocks across time and across states.[36] Namely, the price coefficient is identified by the extent that choice probabilities change as sellers of a given type $\mathcal{Q}_j$ submit higher prices. Controlling for (unobserved) quality of sellers is vital for identifying the price coefficient because the quality of sellers is likely to be correlated with price and also with the probability that individuals choose the bid. For example, more experienced sellers may have lower costs on average and also be more likely to be selected. Identification of the other demand parameters follows similar arguments.

Once the demand parameters are identified, understanding the identification of marginal costs is also straightforward. With demand parameters in hand, we can compute firms optimal markup, which is not a direct function of marginal cost. Using this optimal markup, we can use Equation 15 to create a one-to-one mapping between bid prices and costs.

Finally, the entry cost parameters are identified by variation in expected marginal prof-

[36]Furthermore, controlling for the non-price characteristics of the bid.

its, holding seller type constant.[37] How much does the probability of entry of type Q_j sellers change as expected marginal profits increase? In theory, we could trace out the entire entry cost distribution non-parametrically if we knew the probability of entry at every level of expected marginal profit. In practice, we require exogenous variation in an observed variable, which does not affect the entry cost distribution but enters the seller's ex-ante payoff (expected marginal profit) before the entry decision. I assume that project's size (large or small) does not affect the firm's entry cost, but does affect expected marginal profit. For this assumption to hold, it must be true that bid preparation time and effort is not different for larger residential projects compared to smaller residential projects on average. I also assume that a firm's average entry cost does not vary within a year.[38] Since expected marginal profits will vary each quarter, I am also using variation in expected profits across quarters but within a year to identify the entry cost parameters.

4 Results

4.1 Demand Parameter Estimates

The first column of Table 4 contains the estimates for the baseline demand specification using the two-step grouped fixed effects estimator. λ, the correlation parameter for the nested logit model is significant, indicating that it is important to account for buyer heterogeneity in the likelihood of picking any bid on the platform. For example, some buyers will be more likely to substitute towards the outside option if some options in their choice set were removed. The price coefficient is negative and statistically significant as expected.

The coefficients on the "Premium" and "Premium Plus" dummies are relatively large compared to the price coefficient indicating that buyers are willing to pay a substantial amount more for higher quality equipment. The table also shows the estimates for three of the quality group dummies, Quality Groups 2, 5, and 9. I only report the three groups because of space, I report the rest of the quality group fixed effects in Figure 14 in the appendix. The omitted quality group is "transient" sellers, so the coefficients should be interpreted as changes in utility relative to choosing a similar bid from a "transient" installer. The magnitude of the coefficients on the quality group dummies are also large relative to the price coefficient. The results suggest that there are substantial differences in the unobserved quality of sellers that is critical to buyers decision. In particular, buyers would be willing to pay 19% more for a seller in the 9th quality decile relative to a firm in the 5th decile.

Columns 2-5 of Table 4 show estimates under alternative specifications of the buyer utility. The second column shows how the coefficients change if I do not include any controls for seller quality. The price coefficient is substantially more negative, indicating that there is negative

[37] While also controlling for common entry cost shocks across states and years.

[38] I allow the entry cost to change between 2015 to 2016 but not every quarter.

Table 4: Demand Estimates

	(1)	(2)	(3)	(4)	(5)
λ	0.294***	0.271***	0.275***	0.281***	0.387***
	(0.058)	(0.067)	(0.063)	(0.058)	(0.059)
Price ($/watt)	−0.794***	−0.990***	−0.811***	−0.827***	−0.834***
	(0.108)	(0.114)	(0.124)	(0.112)	(0.101)
Premium Panel	0.482***	0.643***	0.496***	0.473***	0.568***
	(0.072)	(0.079)	(0.085)	(0.098)	(0.069)
Premium Plus Panel	1.164***	1.743***	1.164***	1.085***	1.391***
	(0.143)	(0.165)	(0.175)	(0.159)	(0.147)
Microinverter	0.250***	0.313***	0.247***	0.227***	0.207***
	(0.078)	(0.077)	(0.091)	(0.081)	(0.066)
Seller Quality Group 2	−0.667***			−0.722***	
	(0.206)			(0.205)	
Seller Quality Group 5	0.153			0.114	
	(0.130)			(0.120)	
Seller Quality Group 9	0.697***			0.653***	
	(0.148)			(0.136)	
Star Rating $<$ 5					−0.285**
					(0.119)
Star Rating = 5					0.385***
					(0.097)
$>$ 6 Months on Platform					0.135**
					(0.063)
$>$ 1000 Total Installs					0.113**
					(0.056)
State FE	Yes	Yes	Yes	Yes	Yes
Time-Period FE	Yes	Yes	Yes	Yes	Yes
Quality Group FE	Yes	No	No	Yes	No
Size Control, Large-Project FE	Yes	Yes	Yes	Yes	Yes
Firm FE	No	No	Yes	No	No
Panel Brand FE	No	No	No	Yes	No
Observations	10,545	10,545	10,545	10,545	10,545

Notes: The first column contains estimates for the baseline model. Time-period fixed effects are included for each quarter-year of the sample. Each specification that contains quality-group fixed effects uses 10 quality groups in estimation, only the fixed effects for Groups 2, 5, and 9 are reported in the table, estimates for the other quality groups can be found in Figure 14. Star ratings and total installation experience were extracted from the EnergySage site in September 2017, 8 months after the sample-period was completed.

correlation between installers' unobserved quality and their cost.[39] The third column shows results for the full installer fixed effects model.[40] All of the estimates are very similar to the baseline estimates; this suggests that the specification of quality groups is not driving the results. The fourth column shows that results are also robust to adding panel brand fixed effects.[41] In Section 4.2.1, I discuss the robustness of the estimates to more alternative specifications including different numbers of quality groups.

The last column of Table 4 does not include quality group fixed effects but instead, adds observable measures of installer quality. The model consists of a dummy for if the seller had a five-star rating and another indicator for if the sellers rating was below five stars.[42] I also include a dummy for if the installer had been active on the platform for at least six months and an indicator for if the firm had completed over 1000 total residential installations (both on and off the platform). These observable quality measures are imperfect because they were scraped from the platform database in September 2017, nine months after the sample period ended. Therefore, these variables do not correspond to the exact information that buyers observed when making purchase decisions in the data. Nonetheless, the data can still be helpful for understanding which factors could be contributing to firms' unobserved quality (firm fixed effects). I find that buyers are much more likely to pick sellers that have five-star ratings, and installers with a sub-five star rating were chosen with a lower probability than installers with no ratings. Buyers also prefer sellers with more experience on the platform and with more overall installation experience. The magnitudes of the coefficient estimates also seem to be of similar magnitudes to the quality group fixed effects estimates in the baseline specification.

4.2 Marginal Cost, Markup, and Entry Cost Estimates

After estimating the demand parameters, I use the estimates to infer markups, and marginal costs for each bid in the data using the procedure explained in the previous section. Figure 6 reports summary statistics of the marginal cost and markup results. In the left panel, I decompose the average bid price for each state into an average marginal cost and an average markup component. The full sample average markup was \$1.42/watt, and the average marginal cost was \$2.19/watt. Therefore, markups accounted for 39% of prices on average. This suggests that markups are a major portion of solar PV prices. However, these are gross markup estimates that do not include installers' overhead or marketing costs. Fu et al. (2016) report that installers' benchmark marketing cost should be \$0.31/watt and other overhead expenses should account for \$0.28/watt. I use the estimates from Fu et al. (2016) to calculate installers' net markup. I estimate that net markups were \$0.83/watt and 23% of total system prices on average. The right panel of Figure 6 shows how gross markups and marginal costs changed over the sample.

[39] In the next, subsection I provide evidence of this negative relationship.

[40] The model includes seller fixed effects for each permanent seller, 65 in total. This model is also the first stage used to sort the sellers into quality groups.

[41] I include panel brand dummies for any company that appeared in over 5% of bids in the data.

[42] Sellers without any reviews are the exclude group.

Figure 6: Decomposition of Solar PV Prices

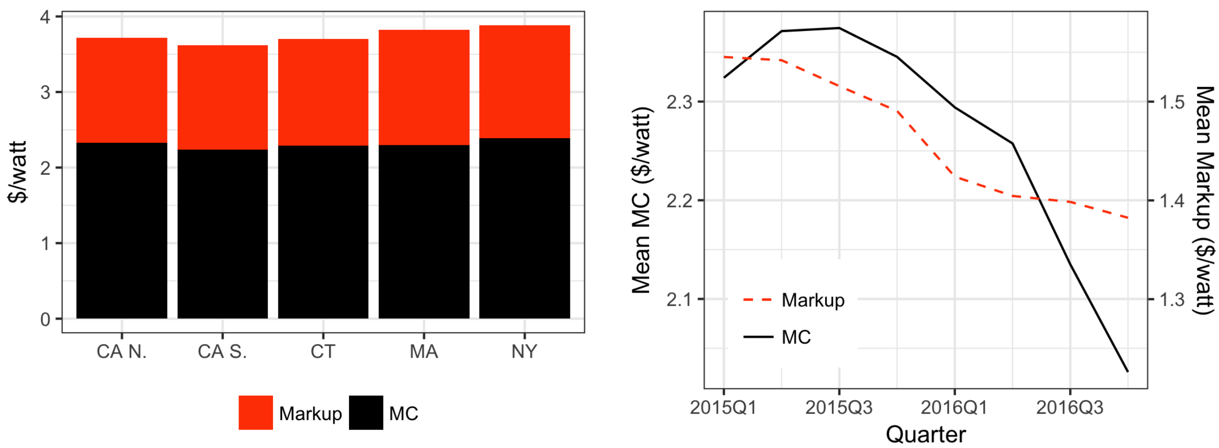

(a) Price Decomposition by State - 2016 Q1

(b) MC & Markups Over Time

Notes: Panel (a) decomposes 2016Q1 price bids into a markups and a marginal cost component for each state on average. The markup component represents the difference between price and marginal cost and does not account for entry costs or other overhead costs. Panel (b) shows how average markups and marginal costs changed over the sample.

Markups fell steadily as the platform became more competitive, accounting for over $0.15/watt decrease bid prices over the sample period. Marginal costs remained steady for the first half of the sample and then sharply declined in the second half of 2016. Figure 17 in the appendix shows that the fall in marginal costs can be accounted for by decreasing hardware costs during the second half of the sample period.[43] In Appendix Figure 16, I plot the average marginal costs and markups for each state over the sample. I also report the state-level estimates in Appendix Table 18.

To gain a better understanding of the main drivers of differences in marginal costs across bids, I run a regression with the implied marginal cost for each bid as the dependent variable and the characteristics of the bid as the explanatory variables. I include three dummy variables for installer quality in the regression: quality groups 1-3 (low quality), quality groups 4-7 (medium quality), and quality groups 8-10 (high quality). The results are shown in Table 7. Note that sellers with higher quality tend to have lower costs. Failing to control for sellers unobserved quality could lead to over-estimates of buyers' price elasticities. If we do not control for the seller's quality, it would appear that customers are picking sellers that have lower prices, but we would also be picking up the negative correlation between price and quality. The estimates from the demand estimates without installer quality group controls (see Column 2 of Table 4) support this conjecture. The price coefficient is indeed lower (more negative). The negative relationship between quality and cost can also be seen in Figure 19 in the appendix. The figure plots the marginal cost distribution for each of the quality groups for one market, it is clear that the cost distribution for the highest quality sellers is shifted to the left.

[43] Hardware Cost estimates were obtained from Bloomberg New Energy Finance. Bloomberg Reports Module Spot Prices and an inverter price index monthly.

The other coefficients in the marginal cost regression also have the expected signs. Systems with high-quality panels and microinverters carry higher marginal costs on average. Also, the negative estimate on project size suggests there are economies of scale within an installation. Economies of scale arise because some parts of the marginal cost of a project such as permitting fees are fixed and do not vary with the size of the installation.

Table 5: Marginal Cost Regressions

	Marginal Cost ($/watt)
Premium Panel	0.115***
	(0.005)
Premium Plus Panel	0.920***
	(0.011)
Microinverter	0.083***
	(0.006)
Size (KW)	−0.029***
	(0.001)
Quality Group (1-3)	0.018**
	(0.007)
Quality Group (4-7)	−0.120***
	(0.007)
Quality Group (8-10)	−0.286***
	(0.008)
State FE	Yes
Time-Period FE	Yes
Observations	40,575
R^2	0.263
Adjusted R^2	0.263
Residual Std. Error	0.427 (df = 40556)
F Statistic	804.506*** (df = 18; 40556)

Notes: 'Transient Sellers' are the omitted quality group. 'Premium' and 'Premium Plus' panel brands are determined by EnergySage based off of panel efficiency, warranty, and other factors.

The entry cost parameter estimates can be found in Appendix Table 19. In Figure 7, I have also used the estimated parameters to plot the predicted entry probabilities for Massachusetts in 2016 for three seller types as a function of expected marginal profits. Permanent sellers are more likely to enter for any level of expected profits. Also notice that even if expected profits get very large, the entry probabilities do not approach 100%. This finding could be explained by installers having obligations to complete other projects outside the platform and may be capacity constrained and unable to submit a bid regardless of how profitable the project appears to be.

Figure 7: Predicted Probabilities of Entry

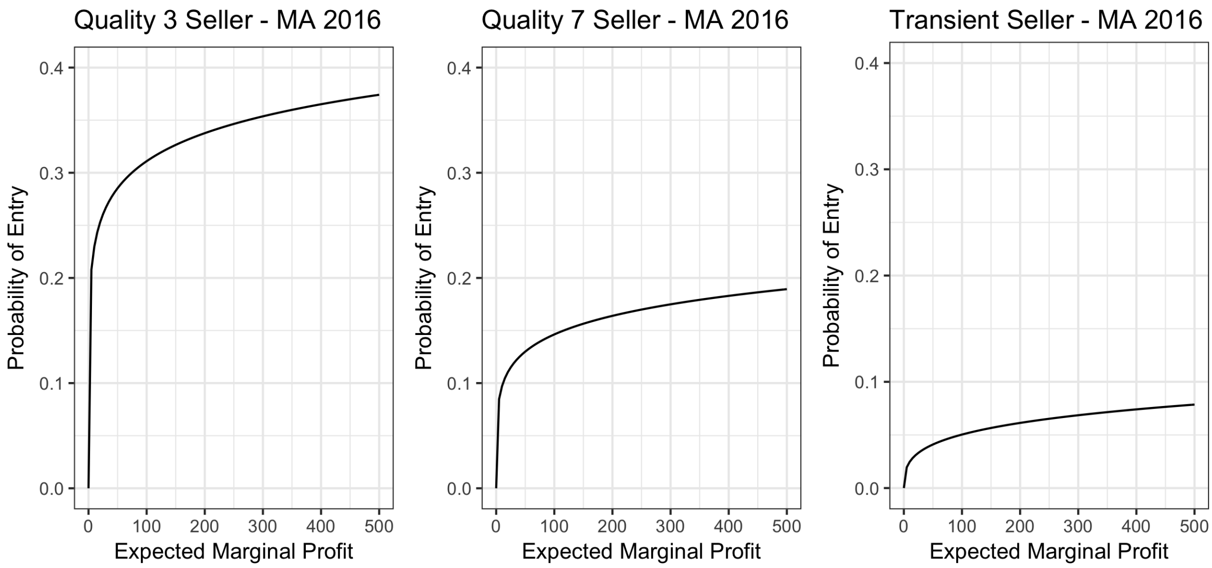

Notes: This figures plots the predicted probabilities of entry for three seller types at varying levels of marginal profits ($). The entry probabilities are determined using the estimated entry cost parameters in Table 19.

4.2.1 Comparison of Estimates to Reported Costs and Robustness Checks

In this section, I compare the marginal cost estimates to existing survey-based estimates and reported costs. I also discuss the robustness of the results to changes in the specification of the model.

Figure 18 in the appendix compares the estimated marginal costs and markups from the model to existing estimates from National Renewable Energy Laboratory[44] (NREL) and reported marginal costs from installers' earnings statements. The estimated marginal costs are similar to the marginal costs reported in quarterly reports by large publicly traded installers, Solar City, Vivint, and Sunrun. In the first quarter of 2016, Solar City reported that installation costs were $1.90/watt and Vivint, and Sunrun reported costs of $2.34/watt and $2.73/watt, respectively. My estimate for the average marginal cost in 2016Q1 falls right between these reports at $2.29/watt. My cost estimate is slightly higher than NREL's benchmark estimates of $1.89/watt. This result is not surprising considering the NREL analysis quantifies a lower bound for costs (i.e., industry best practices). Overall, the model estimates appear to be consistent with existing studies and reports.

Before turning to the counterfactual exercises, I also perform a series of robustness checks to ensure that the demand parameters are not sensitive to small changes in modeling choices.

[44]See Fu et al. (2016).

Table 15 shows how the demand results change after adding additional controls. For example, I add controls for financing offers (loans), state-level incentives, other demand shifters (state-year fixed effects), and time-varying controls for installer experience on the platform.[45] The price coefficient does not change much under any of these alternative specifications.

Table 15 explores alternative methods for calculating the quality-group fixed effects. In column two, instead of sorting the firms into ten equally sized groups, I alternatively define the quality groups by using a k-means clustering algorithm. The k-means algorithm assigns each installer into ten clusters to fit the data best; each cluster can contain a different number of firms.

In the third column, I again use ten equally sized quality groups but I use an iterated estimator for the groupings. At each iteration, the installers are assigned to new groups based off of the firms average residuals from the previous iteration. I continue to update the groups until convergence. The motivation for this estimator is the possibility of misclassifying the quality groups under the two-step estimation approach. Namely, if the installer fixed effects are imprecisely estimated, some installers may be sorted into the "incorrect" group. Fortunately, I find the price coefficient is very similar to the baseline estimate using both alternative estimators. Finally, Appendix Figure 15 shows how the ratio of the price coefficient to the nesting parameter $\alpha/(1-\lambda)$ changes if I use a different number of quality groups.[46] The ratio largely fluctuates if I use fewer than five quality groups in estimation. The estimated fraction $\alpha/(1-\lambda)$ stabilizes at around six quality groups. Any specification with over six groups, including the firm fixed effects model (65 groups) yields similar estimates.

I also test the robustness of the markup estimates to changes in the definition of auction type categories. One concern is that the project size categories do not allow for sufficient heterogeneity across projects of different sizes. Column 2 of Appendix Table 17 shows that markup estimates barely change if I use three project size categories instead of the two categories used in the baseline model. Another possibility is that I am not sufficiently controlling for geographic heterogeneity in competition, costs, and preferences. Appendix Figure 20 shows how the average number of bids changes across counties and also how average bid prices vary across counties. Massachusetts and Connecticut are smaller states, so the variation in prices across counties is very small. However, in New York and California there is more geographic variation in prices. Specifically, prices are higher in major urban areas like New York City and San Francisco. To account for this heterogeneity, I separate New York into two different areas, NYC Metro and Upstate (rest of state). I also divide California into three geographic regions (instead of two), Bay Area, Urban Southern California, and Rural California (rest of state). Column 3 of Appendix Table 17 shows that the markup results are robust to these changes. Furthermore,

[45] Three of the four states did not have state incentives that varied over the sample so any state incentive should be picked up through the state fixed effects. New York is the only state that had subsidies that varied over the sample period. The estimated coefficient on "loan offered" is negative, likely because this variable is endogenous (consumers are allowed to request financing when they create an account).

[46] Under nested logit demand, the term $\alpha/(1-\lambda)$ enters multiplicatively into each firm's optimal markup, so this ratio determines how large the markup estimates will be.

Columns 4 and 5 show that the results are also robust to changes in the length of "time-period" used to sort the projects into categories.

4.3 Counterfactuals

In this section, I use the recovered industry cost structure and buyer preferences to simulate several counterfactual environments. In particular, the counterfactual experiments allow us to understand how access to the platform affects prices, solar system adoption rates, and welfare.

In the first counterfactual, I simulate prices and adoption rates in the case that buyers were not able to use the platform to solicit bids and there is also no federal subsidy available for solar PV. Instead, I assume that buyers only call one installer for a price quote. More specifically, I assume that each buyer randomly samples one seller based on market share.[47] Chosen installers are aware that they are no longer competing against other bidders on the platform (although they still compete against the outside option which could include national installers like SolarCity) so they will adjust their price accordingly.

In the second counterfactual, I consider the case where buyers still cannot use the platform but have access to the 30% federal subsidy. Again, sellers will re-optimize their price to capture some of the subsidies.

Next, I simulate outcomes for a situation where buyers can access the platform but do not have access to the federal subsidy. Sellers in this counterfactual scenario can update their entry strategies to account for the removal of the subsidy and also update bidding strategies accordingly. I can then compare how giving consumers access to the platform equates to giving them the subsidy but no way to cheaply access more potential sellers (the second counterfactual).

Finally, I run a counterfactual experiment where buyers are given access to the platform, but sellers must hold their prices constant at the prices they quoted in the first counterfactual.[48] By holding prices constant at the residual monopolists' prices, I can isolate how much of the change in adoption rate caused by the platform is coming from giving buyers access to more and potentially higher quality installers, as opposed to lower prices.

4.3.1 Solving for Counterfactual Equilibria

For each counterfactual, I numerically solve for the new equilibrium following a policy change. To obtain the counterfactuals, I use a backward solution strategy. In particular, I first derive the equilibrium bidding strategies of sellers given competing sellers entry behavior. Next, I go back and solve for optimal seller entry given the equilibrium bidding strategies. I then combine these two steps to compute the equilibrium outcome for the entire system.

[47] I sample from all potential entrants for that project, the sample is weighted based on the total sales on the platform.

[48] More precisely, I assume that every price that buyers receive on the platform is equal to the average price bid under the no platform counterfactual for the same project type.

To derive a seller's equilibrium bid under a counterfactual scenario, we need information on the equilibrium distribution of bids (to calculate each firm's markup). However, the equilibrium bid distribution will change under the counterfactual, and we need to estimate the new distribution. I begin by using the original distribution of bids (and entry probabilities) observed in the data and then use a root-finding algorithm to calculate a new optimal bid price for each quote in the data. After the new optimal prices have been solved for, we now have a new updated distribution of bids. I then update each firm's equilibrium entry probability given this updated distribution of bids. I repeat the process of updating the bid distribution and entry probabilities until convergence. For any counterfactuals that only involve a single bidder, the firms' strategy will not be a function of other firms' equilibrium bids so we can solve for the new prices in one step.

4.3.2 Model Fit

In Table 6 below, I use the estimated structural parameters to simulate entry, prices, and purchases and compare the model predictions to the observed data. Relative adoption is the percentage of total solar system sales compared to the observed data (.i.e 100% relative adoption would mean that the same number of PV systems were purchased as in the observed data). The model predicts bid prices and the number of entrants very closely to the observed data. The model slightly overpredicts the number of adoptions compared to the observed data. Overall, the model appears to fit the data quite well.

Table 6: Model Fit

	Avg. Bid Price	Takeup Relative to Observed	Avg. Number of Bids
Observed	\$3.61/watt	-	3.95
Model Prediction	\$3.64/watt	110.02%	3.96

4.3.3 Main Counterfactual Results

I report results from the counterfactual exercises in Table 7. The first column shows results for the baseline simulation where each buyer cannot access the platform and instead draws a single installer (based on platform market share) to place a sole bid for the project. In the baseline simulation, there is also no 30% investment tax credit (ITC) for solar purchases. I find that the average bid price is \$3.46/watt (the average net bid price is also \$3.46/watt because there is no subsidy in this case).[49]

In the second column, we see that providing a 30% subsidy for solar results in 17% higher bid prices if consumers receive only one bid. Sellers can capture part of the subsidy by setting

[49]Low Price Bid is the average minimum price bid received by consumers, for this counterfactual each buyer only receives one quote so this is equal to the overall average bid price.

higher prices, while losing no sales (since net costs will be lower for consumers). The average net price from the perspective of consumers is 18% lower under the subsidy policy. In this scenario, adoption increases by 52% relative to the baseline. This finding is similar to other estimates in the literature. For example, Hughes and Podolefsky (2015) find that the California Solar Initiative (CSI) subsidies doubled the number of total installations that occurred. The average subsidy was slightly larger under CSI, with the average household receiving about $1.50/watt, relative to approximately $1.20/watt under the ITC. We also see that both consumer and producer surplus increase by 85% and 123% respectively. However, these gains are mainly a transfer of rents because total welfare only increases by 8% after we subtract the cost of the subsidy. These are estimates of private welfare and do not account for environmental benefits of the solar PV systems.

Table 7: Counterfactual Results

Panel A: Counterfactuals

	(1) Baseline - No Platform, No ITC	(2) No Platform, with ITC	(3) Platform Access, No ITC	(4) Platform Access, with ITC
Low Bid Price (Mean)	$3.46	$4.06	$2.92	$3.32
Net Low Price (Mean)	$3.46	$2.84	$2.92	$2.32
Number of Bids (Mean)	1	1	3.52	3.96
Adoption (% Increase)	-%	52%	104%	255%
CS (% Increase)	-%	85%	156%	273%
PS (% Increase)	-%	123%	69%	319%
Welfare (% Increase) = CS+PS-Subsidy Cost	-%	8%	89%	79%

Panel B: Decomposition of Adoption Increase

	Platform Access, No ITC, with Baseline Prices
Lowest Bid Price (Mean)	$3.47
Number of Bids (Mean)	3.52
Adoption (% Increase)	74%

Notes: In Panel A, the first column shows the results for the baseline simulation where each buyer cannot access the platform and instead draws a single installer (based of platform market share) to place a single bid for the project, there is no 30% investment tax credit (ITC) in the baseline simulation. Columns 2-4 show results of different counterfactual environments. Panel B assumes that the buyer can access the platform but sellers are forced to submit the same price as they would if they were the only bidder.

The third column shows how the market prices and adoption rates change if customers can use the platform to solicit bids but no 30% subsidy is available. The average (lowest) bid price that buyers receive would be slightly higher than the case with the 30% subsidy but only one bid. The average low bid is 15% smaller than the baseline case. On the other hand, adoption rates increase by 104% when consumers can use the platform. Additionally, there are substantial surplus gains, particularly for consumers who enjoy a 156% increase in consumer surplus. Despite lower equilibrium bid prices, sellers also experience a 63% increase in overall profits relative to the baseline case. Firm profits are higher because the decrease in prices is more than offset by the increase in total sales. Seller profits are still lower than the single bidder case with the subsidy.

There are two potential channels through which the platform could increase overall adoption. First, bid prices decline when buyers have access to the platform as the market becomes more competitive. Second, the platform could increase adoption by connecting buyers with higher quality sellers. Reduced prices do not appear to be the primary driving factor behind the significant increase in adoption rates resulting from platform access. To see this, notice that adoption is actually higher with the platform and no subsidy even though prices are slightly higher than the "no platform with ITC" simulation. This suggests that the primary driver through which the platform increases adoption could be giving buyers access to higher-quality installers. To explore this possibility, I simulate another counterfactual where I hold sellers entry probabilities the same as in Panel A, Column 3 but I force the sellers to quote prices that are the same as the baseline case. In this counterfactual buyers are offered the same prices on average as they would in the baseline case. Panel B reports the results. Despite prices remaining constant, adoption still increases by 74% relative to the baseline case. This result indicates that the majority (71%) of the increase in adoption can be attributed to linking buyers to higher quality or better-matched installers.

4.3.4 Alternate Assumptions for the "No Platform" Counterfactual

The baseline counterfactual simulation assumes that households would receive only one price quote if they were unable to access the platform. While this is a reasonable starting point, many buyers may collect more than one quote in the absence of the platform. If buyers receive multiple quotes without the platform, the results in the last section would overstate the increase in adoption caused by the platform. Table 8 shows the predicted increase in adoption changes under alternative assumptions about the "no platform" counterfactual. In the second column, I assume that (in expectation) each buyer receives half as many quotes as the under the platform case (Column 3 of Table 7). I implement this counterfactual by reducing each firm's entry probability by half relative to the platform case. Sellers are allowed to update their equilibrium bids accordingly. Under these assumptions, buyers would receive about 1.7 bids on average, and the platform would instead cause a 56% increase in adoption relative to this new baseline. In the third column, I assume that each consumer would get two price quotes without the platform. Under this assumption, some buyers receive more quotes in the "no platform" case than in the "platform" case, over 15% of buyers only received one bid in the "platform" simulation.

Nonetheless, I still find that the platform would increase overall adoption by 26% overall in this case.

Table 8: No Platform Counterfactual - Alternate Asssumptions

	(1) Baseline - Buyers Receive One Bid	(2) Buyers Receive 1/2 as Many Bids	(3) Buyers Receive Two Bids
Effect of Platform on Adoption (% Incr.)	104%	56%	26%

Notes: This table shows the simulated effects of providing consumers with access to the platform under alternative assumption about the no platform counterfactual. In the baseline case, each buyer draws one seller weighted by market share. In column two, I assume each buyer recieves a number of bids that is half as large as what they would recieve on the platform in expectation. More concretely, buyers received an average of just over 3 bids on the platform, so they receive an average around 1.7 bids in this counterfactual. In the third column, I assume that each buyer would receive exactly two bids in absence of the platform.

5 Discussion and Conclusions

Increasing renewable energy investment is viewed as a pertinent part of a larger goal to curb carbon emissions from the electricity sector worldwide. In particular, many national and local governments have emphasized expanding adoption of residential solar PV systems. Despite falling hardware costs and increased public interest in solar energy, installation prices remain prohibitively high for many households to adopt the technology.

In this paper, I estimate a structural model of consumer choice and seller bidding in the residential solar PV market. The model allows me to quantify the importance of price, hardware quality, and installer quality in consumers' solar adoption decision. I also use the model to measure the extent of market power in the solar industry. In particular, I use high-resolution data on firm bidding behavior to identify marginal costs, markups, and bid preparation costs.

My estimates provide evidence that both hardware and installer quality are critical factors in buyers' solar purchase decision. Namely, the average buyer is willing to pay substantially more for a high-quality firm to install their PV system. I also show that supplier markups compose a substantial portion of overall PV system prices. Market power is of particular concern in this setting because solar PV provides positive environmental benefits.

Most current public policies have used subsidies, tax credits, and rebates to encourage adoption of residential solar PV systems. These policies have been shown to be an effective method for increasing solar PV purchases. However, these programs also require large public expenditures and risk being cut if government funds are limited. This paper provides evidence that using platforms to increase competition and expand buyers choice sets can also serve as an effective way to increase adoption. Platforms can reduce system prices by making the market more competitive. Additionally, increasing buyer participation on platforms can also lead to more purchases by linking consumers to higher quality installers.

Both state and federal governments could channel more effort to informing the public about existing platforms or by developing their own platforms to link buyers and sellers. Since platforms have already been designed to connect buyers and sellers in this market, it is likely that increases in adoption from expanding platform participation could come at a relatively low cost and lead to significant welfare gains for both consumers and producers.

References

Barbose, G. L. and Darghouth, N. R. (2016). Tracking the Sun IX: The Installed Price of Residential and Non-residential Photovalic Systems in the United States (Washington: US Department of Energy). *available at emp. lbl. gov/publications/trackin.*

Berry, S., Levinsohn, J., and Pakes, A. (1995). Automobile prices in market equilibrium. *Econometrica*, pages 841–890.

Bollinger, B. K. and Gillingham, K. (2014). Learning-by-doing in solar photovoltaic installations.

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

Brown, J. R. and Goolsbee, A. (2002). Does the Internet make markets more competitive? Evidence from the life insurance industry. *Journal of Political Economy*, 110(3):481–507.

Brown, Z. Y. (2017). An Empirical Model of Price Transparency and Markups in Health Care. Technical report, Working Paper, Columbia University.

Burr, C. (2014). Subsidies and Investments in the Solar Power Market. Technical report, Working Paper, University of Colorado.

Callaway, D., Fowlie, M., and McCormick, G. (2017). Location, location, location: The variable value of renewable energy and demand-side efficiency resources. *Journal of the Association of Environmental and Resource Economists.* (Forthcoming).

Cullen, J. (2013). Measuring the environmental benefits of wind-generated electricity. *American Economic Journal: Economic Policy*, 5(4):107–133.

De Groote, O. and Verboven, F. (2016). Subsidies and myopia in technology adoption: Evidence from solar photovoltaic systems. Technical report, Working Paper, KU Lueven.

Diamond, P. A. (1971). A model of price adjustment. *Journal of Economic Theory*, 3(2):156–168.

Fu, R., Chung, D., Lowder, T., Feldman, D., Ardani, K., and Margolis, R. (2016). *U.S. Solar Photovoltaic System Cost Benchmark: Q1 2016.*

Gehrig, T. (1993). Intermediation in search markets. *Journal of Economics & Management Strategy*, 2(1):97–120.

Gillingham, K., Deng, H., Wiser, R., Darghouth, N., Nemet, G., Barbose, G., Rai, V., Dong, C. G., and others (2016). Deconstructing solar photovoltaic pricing. *The Energy Journal*, 37(3).

Gillingham, K. and Tsvetanov, T. (2014). Hurdles and steps: Estimating demand for solar photovoltaics. Technical report, Working Paper, Yale University, Working Paper, Yale University.

Gowrisankaran, G., Reynolds, S. S., and Samano, M. (2016). Intermittency and the value of renewable energy. *Journal of Political Economy*, 124(4):1187–1234.

Guerre, E., Perrigne, I., and Vuong, Q. (2000). Optimal Nonparametric Estimation of First-Price Auctions. *Econometrica*, 68(3):525–574.

Hall, G. and Rust, J. (2003). Middlemen versus market makers: A theory of competitive exchange. *Journal of Political Economy*, 111(2):353–403.

Hughes, J. E. and Podolefsky, M. (2015). Getting green with solar subsidies: Evidence from the California solar initiative. *Journal of the Association of Environmental and Resource Economists*, 2(2):235–275.

Jensen, R. (2007). The digital provide: Information (technology), market performance, and welfare in the South Indian fisheries sector. *The Quarterly Journal of Economics*, 122(3):879–924.

Krasnokutskaya, E. and Seim, K. (2011). Bid preference programs and participation in highway procurement auctions. *The American Economic Review*, 101(6):2653–2686.

Krasnokutskaya, E., Song, K., and Tang, X. (2017). The Role of Quality in On-Line Service Markets. Technical report, Working paper, Johns Hopkins.

Langer, A. and Lemoine, D. (2017). Designing Dynamic Subsidies to Spur Adoption of New Technologies. Technical report, Working Paper, University of Arizona.

Luco, F. (2016). Who Benefits from Information Disclosure? The Case of Retail Gasoline. Technical report, Working Paper, Texas A&M.

Novan, K. (2015). Valuing the wind: Renewable energy policies and air pollution avoided. *American Economic Journal: Economic Policy*, 7(3):291–326.

Pless, J. and van Benthem, A. (2017). The Surprising Pass-Through of Solar Subsidies. Technical report, National Bureau of Economic Research.

Reddix II, K. D. (2015). *Powering Demand: Solar Photovoltaic Subsidies in California*. PhD thesis, The University of North Carolina at Chapel Hill.

Salz, T. (2017). Intermediation and competition in search markets: An empirical case study. Technical report, Working Paper, Columbia University.

Spulber, D. F. (1996). Market making by price-setting firms. *The Review of Economic Studies*, 63(4):559–580.

Stahl, D. O. (1989). Oligopolistic pricing with sequential consumer search. *The American Economic Review*, pages 700–712.

Stigler, G. J. (1961). The economics of information. *Journal of Political Economy*, 69(3):213–225.

Yoganarasimhan, H. (2015). Estimation of beauty contest auctions. *Marketing Science*, 35(1):27–54.

Appendices

A Tables and Figures

Figure 8: Google Maps Photo of the Rooftop for a Potential Project

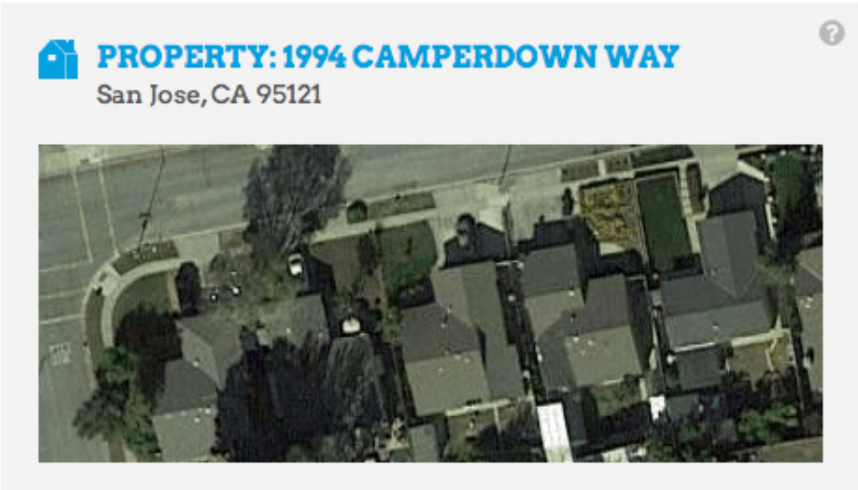

Figure 9: EnergySage Dashboard for Comparing Submitted Quotes

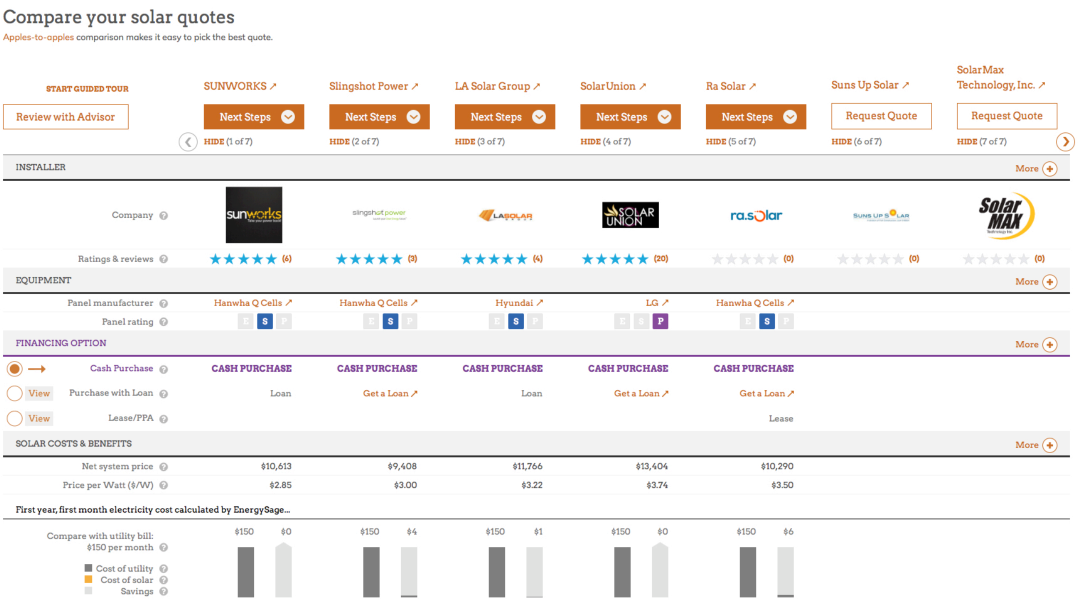

Table 9: Project Summary Statistics by State (Means)

	CA North	CA South	CT	MA	NY	Total
Number of Bids	3.947	4.649	3.509	3.669	3.096	3.953
	(1.811)	(1.950)	(1.808)	(1.776)	(1.660)	(1.903)
Project Size (Watts)	6312.3	6630.4	8288.0	7498.0	8887.3	7155.9
	(2902.9)	(2827.0)	(3101.4)	(2950.4)	(3546.9)	(3138.6)
Distinct Panel Brands	3.051	2.816	2.457	2.662	2.620	2.795
	(1.281)	(1.228)	(1.162)	(1.235)	(1.398)	(1.277)
# of Bids w/ Premium Panel	1.027	2.250	1.353	1.468	0.886	1.486
	(0.994)	(1.628)	(1.293)	(1.336)	(0.876)	(1.398)
# Bids w/ Premium Plus Panel	0.214	0.0727	0.266	0.0945	0.235	0.159
	(0.471)	(0.281)	(0.538)	(0.311)	(0.463)	(0.410)
# of Bids w/ Microinverter	2.564	3.619	2.196	3.364	2.156	2.929
	(1.613)	(1.973)	(1.812)	(1.825)	(1.363)	(1.854)
# of Bids from Permanent Sellers	3.421	4.070	2.721	3.280	2.142	3.351
	(1.676)	(1.841)	(1.344)	(1.757)	(1.310)	(1.783)
Observations	10545					

Standard deviations are listed in parentheses.

Table 10: Mean Bid Characteristics by State

	CA North	CA South	CT	MA	NY	Total
Price ($/watt)	3.590	3.476	3.739	3.773	3.765	3.611
	(0.482)	(0.416)	(0.452)	(0.486)	(0.650)	(0.494)
Premium Panel (0,1)	0.266	0.502	0.396	0.408	0.292	0.386
	(0.442)	(0.500)	(0.489)	(0.492)	(0.455)	(0.487)
Premium Plus Panel (0,1)	0.0554	0.0162	0.0778	0.0263	0.0774	0.0412
	(0.229)	(0.126)	(0.268)	(0.160)	(0.267)	(0.199)
Microinverter (0,1)	0.665	0.807	0.642	0.936	0.710	0.761
	(0.472)	(0.395)	(0.479)	(0.245)	(0.454)	(0.426)
Permanent Seller (0,1)	0.887	0.907	0.796	0.913	0.705	0.871
	(0.317)	(0.290)	(0.403)	(0.282)	(0.456)	(0.335)
Observations	40575					

Standard deviations are listed in parentheses.

Table 11: Seller Summary Statistics

	Mean	SD	10-%tile	50-%tile	90-%tile
Total Bids Submitted	289.05	488.14	9.00	87.50	830.00
Avg. Star Rating (if Rated)	4.80	0.64	4.00	5.00	5.00
5 Star Rating	0.84	0.36	0.00	1.00	1.00
Below 5 Star Rating (if Rated)	0.16	0.36	0.00	0.00	1.00
Number of Ratings	6.98	10.66	1.00	3.50	16.00
Total Residential Installs	1579.18	3487.32	0.00	410.00	3500.00
Years Installing Solar	9.02	6.01	3.00	8.00	15.00
N	168				

Table 12: Regressions of Bid Prices ($/watt) on Expected Number of Bids

	(1)	(2)	(3)	(4)
Expected Number of Bids	-0.148***	0.0173***	-0.0186***	0.00765
	(0.00428)	(0.00578)	(0.00420)	(0.00532)
Controls	Yes	Yes	Yes	Yes
State FE	No	Yes	No	Yes
Time FE	Yes	Yes	Yes	Yes
Installer FE	No	No	Yes	Yes
N	40575	40575	40565	40565
R^2	0.312	0.347	0.530	0.535

Each regression includes controls for system size, panel quality (rated by EnergySage), and a dummy for if the system includes a microinverter or power optimizer. The expected number of bids is calculated by taking the mean number of bids for projects in the same state and same quarter as project i. All standard errors are listed in parenthesis and are clustered by project ID. * indicates $p < 0.10$, ** indicates $p < 0.05$, and *** indicates $p < 0.01$

Table 13: Regressions of Bid Prices ($/watt) on Realized Number of Bids

	(1)	(2)	(3)	(4)
Realized Number of Bids	-0.0328***	-0.0132***	-0.0168***	-0.0151***
	(0.00153)	(0.00137)	(0.00124)	(0.00125)
Controls	Yes	Yes	Yes	Yes
State FE	No	Yes	No	Yes
Time FE	Yes	Yes	Yes	Yes
Installer FE	No	No	Yes	Yes
N	40575	40575	40565	40565
R^2	0.294	0.349	0.533	0.537

Each regression includes controls for system size, panel quality (rated by EnergySage), and a dummy for if the system includes a microinverter or power optimizer. All standard errors are listed in parenthesis and are clustered by project ID. * indicates $p < 0.10$, ** indicates $p < 0.05$, and *** indicates $p < 0.01$

Figure 10: Histogram - Number of Bids per Project

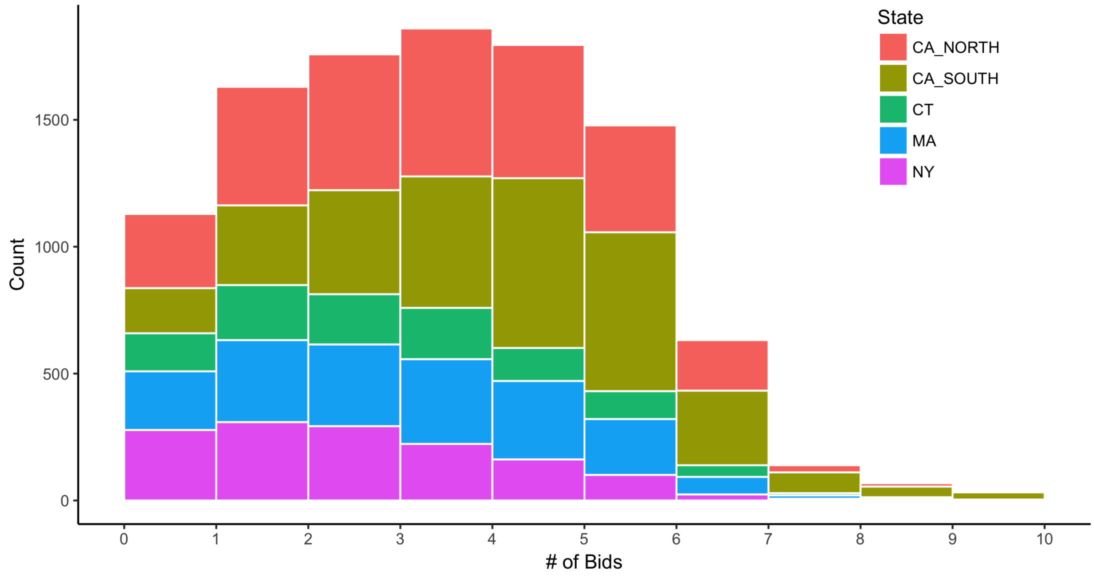

Figure 11: Number of Distinct Installers Submitting Bids on the Platform

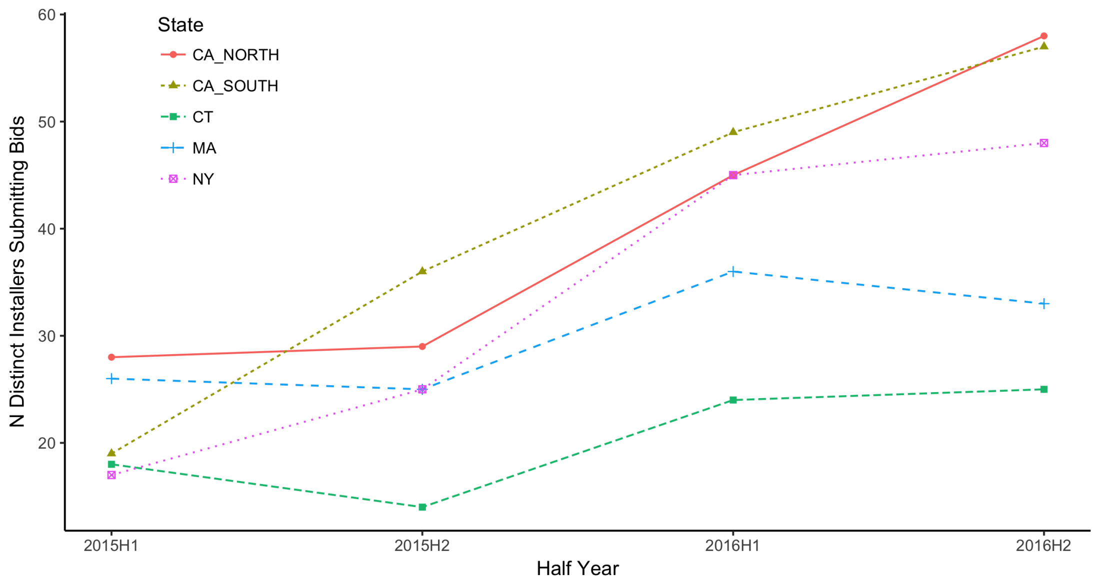

Figure 12: 2016 Q1 Price Distributions by State

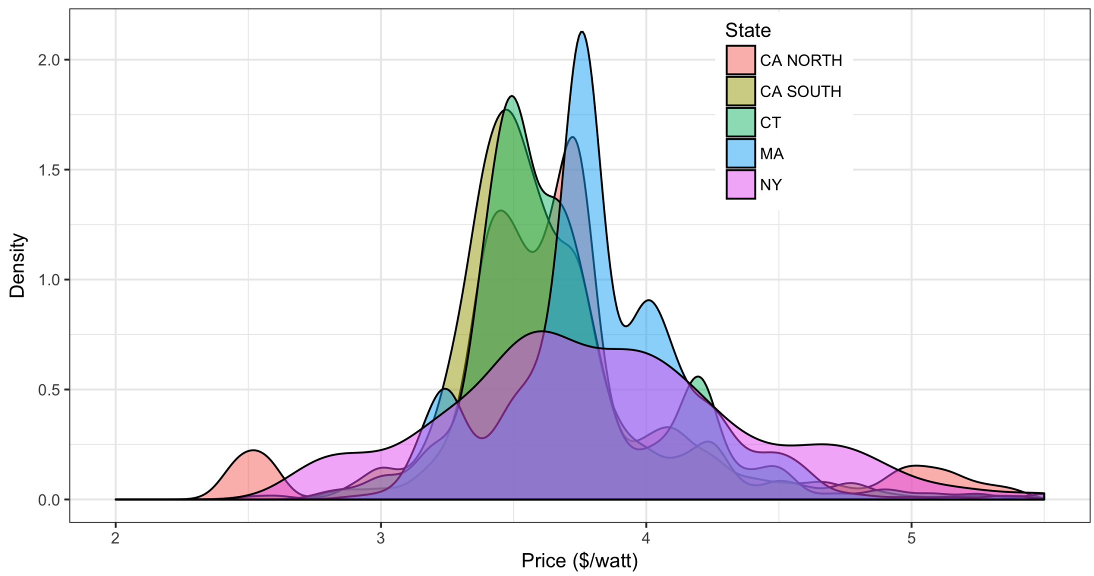

Figure 13: Imputed Installer Locations

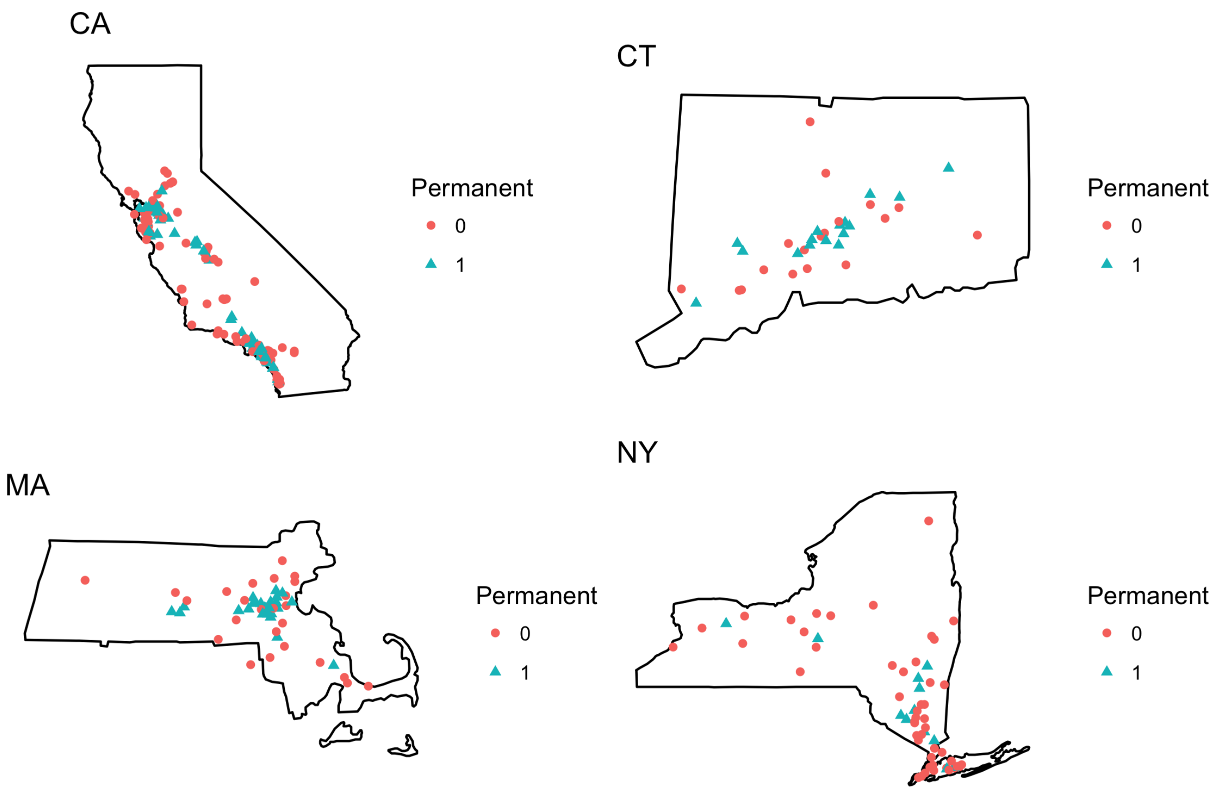

Notes: Estimated installer locations are the centroid of all project locations that the installer submitted bids for throughout the sample. The red circles represent "transient" sellers that submitted fewer than 100 total bids, and the blue triangles represent permanent sellers.

Table 14: Regressions of Bid Price ($/watt) on Order of Bid

	(1) Price ($/watt)	(2) Price ($/watt)
Order of Bid	0.00169	-0.000182
	(0.00133)	(0.00122)
Panel Quality Controls	No	Yes
NBids Control	Yes	Yes
System Size Control	Yes	Yes
State FE	Yes	Yes
Time FE	Yes	Yes
Installer FE	Yes	Yes
N	40565	40565
R^2	0.449	0.537

Panel quality controls include panel quality (rated by EnergySage), and a dummy for if the system includes a microinverter or power optimizer. All standard errors are listed in parenthesis. * indicates $p < 0.10$, ** indicates $p < 0.05$, and *** indicates $p < 0.01$

Figure 14: Quality Group Fixed Effects Estimates

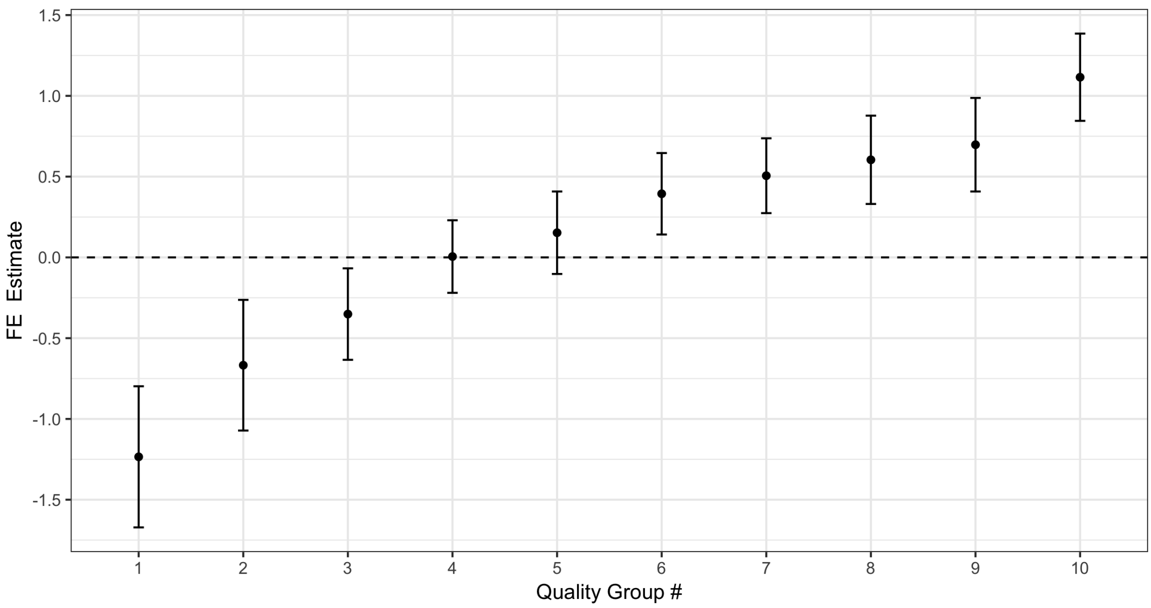

Notes: This figure plots the point estimates for each of the quality group fixed effects from the preferred demand specification (column 1 of Table 4). The bars indicate 95% confidence intervals.

Table 15: Robustness Checks: Alternative Demand Specifications

	(1)	(2)	(3)	(4)
λ	0.294^{***}	0.285^{***}	0.279^{***}	0.292^{***}
	(0.058)	(0.058)	(0.060)	(0.059)
Price ($/watt)	-0.794^{***}	-0.750^{***}	-0.781^{***}	-0.783^{***}
	(0.108)	(0.110)	(0.109)	(0.110)
State Incentive ($/watt)		0.445^{***}		
		(0.172)		
Loan Offered		-0.194^{***}		
		(0.060)		
$>$ 6 Months on Platform				-0.021
				(0.081)
Equipment Quality Controls	Yes	Yes	Yes	Yes
State FE	Yes	Yes	No	Yes
Time-Period FE	Yes	Yes	Yes	No
Quality Group FE	Yes	Yes	Yes	Yes
State-Year FE	No	No	Yes	No
Size Control, Large Project FE	Yes	Yes	Yes	Yes
# Quality Groups	10	10	10	10
Observations	10,545	10,545	10,545	10,545

Notes: Equipment quality controls include premium panel, premium plus panel, and microinverter dummies. State incentives are reported by EnergySage and include both rebates and tax credits. NY was the only state that had an incentive policy that varied over the sample period. '>6 months on the Platform' is dummy determined by installer experience on EnergySage.

Table 16: Robustness Checks: Alternative Demand Specifications 2

	Baseline	K-Means	Iterated
	(1)	(2)	(3)
λ	0.294***	0.275***	0.254***
	(0.058)	(0.059)	(0.061)
Price ($/watt)	−0.794***	−0.820***	−0.835***
	(0.108)	(0.110)	(0.112)
Equipment Quality Controls	Yes	Yes	Yes
State FE	Yes	Yes	Yes
Time-Period FE	Yes	Yes	Yes
Size Control, Large Project FE	Yes	Yes	Yes
# Quality Groups	10	10	10
Observations	10,545	10,545	10,545

Notes: Equipment quality controls include premium panel, premium plus panel, and microinverter dummies. The 'K-means' specifiactions uses, a k-means clustering algorithm to sort the sellers into groups based off the estimated firm FEs. The 'Iterated' specification updates the quality grouping after each estimation round based off of each firm's average residuals, and continues to update the groupings until convergence.

Figure 15: Robustness: Alternative Quality Group Specifications

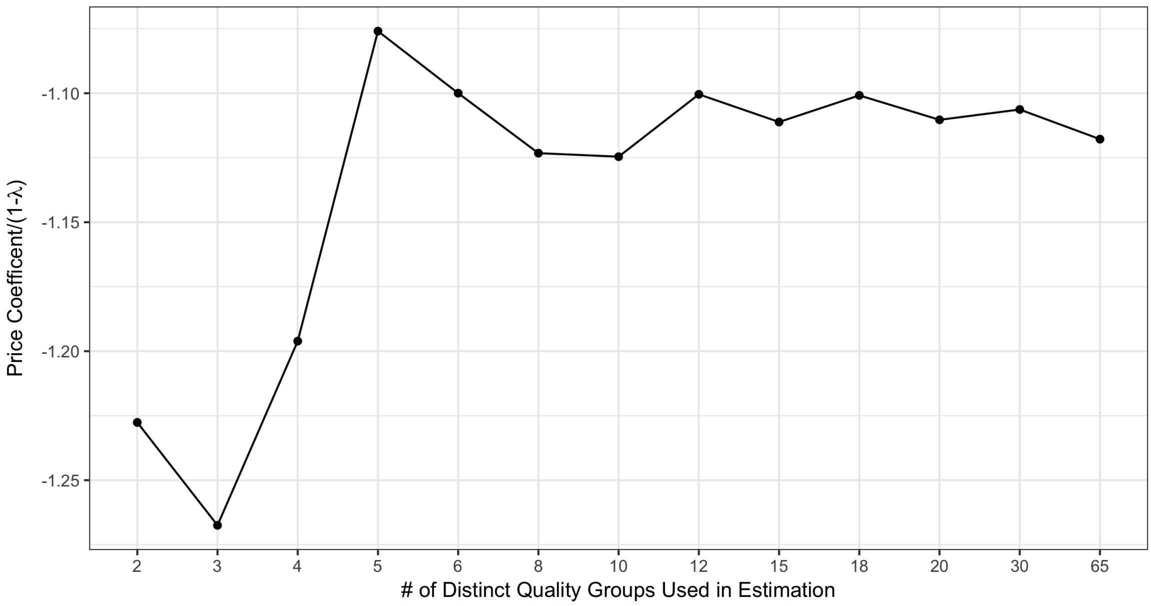

Notes: This figure shows how the ratio of the estimated price coefficient over $1-\lambda$ changes if different numbers of quality groups are used in estimation. The case with 65 quality groups corresponds to the firm fixed effects specification.

Figure 16: Estimated Marginal Costs and Markups by State over Time

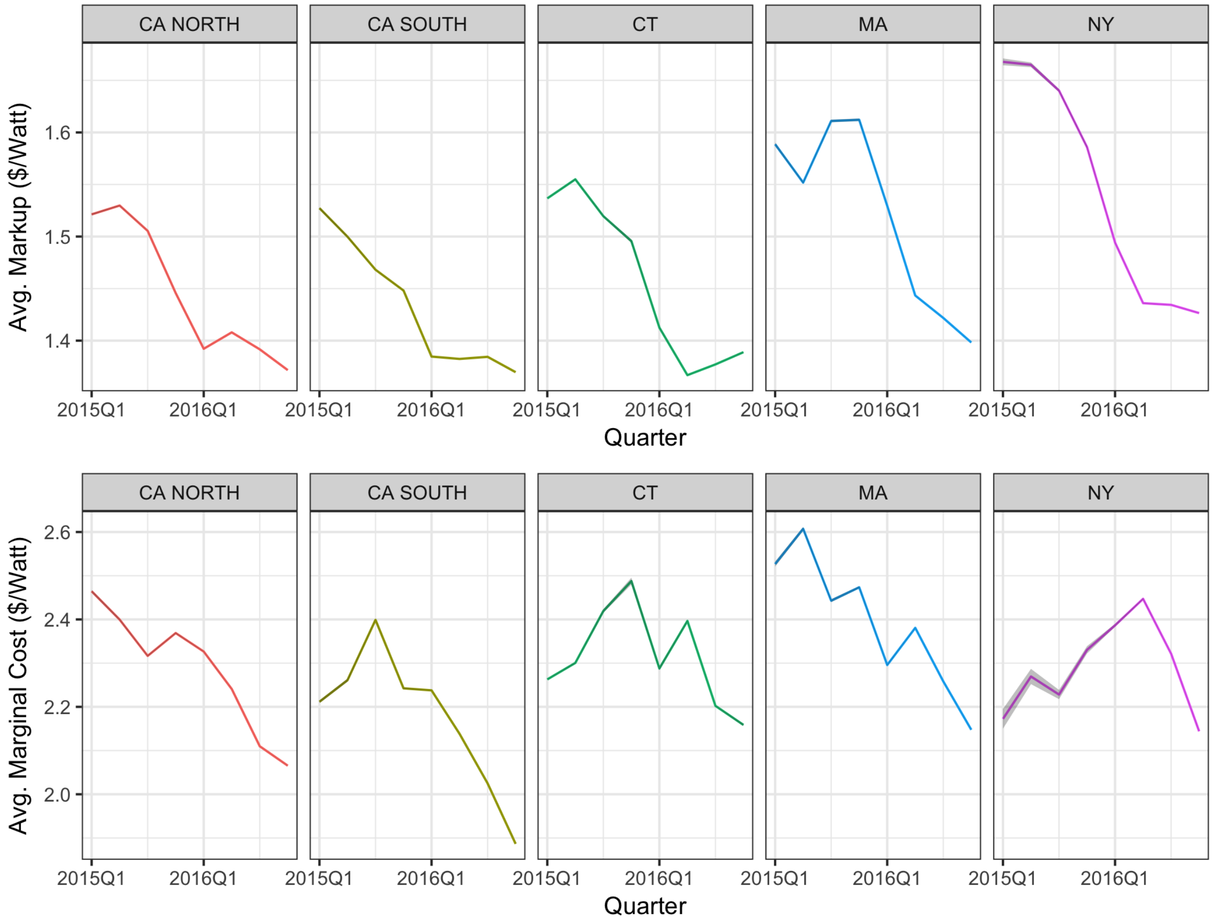

Table 17: Markup and MC Estimates - Alternative Project Classifications

	(1)	(2)	(3)	(4)	(5)
Mean Markup	1.42	1.43	1.42	1.48	1.43
Mean Marginal Cost	2.19	2.19	2.19	2.14	2.19
Mean Own-Price Elasticity	-1.79	-1.78	-1.79	-1.72	-1.78

Notes: All markups and marginal costs are recorded in dollars per watt. The first column reports summary statistics for the estimates under the baseline specification of project types. Column two allows for three projects size categories instead of two. The third column uses shorter time-periods, specifically, each time period is two months long instead of three. The fourth column uses 6-month time periods. Finally, the fifth column allows for more geographic heterogeneity by splitting NY and CA into rural and urban areas.

Figure 17: Estimated Average Marginal Cost and Bloomberg Hardware Cost Index

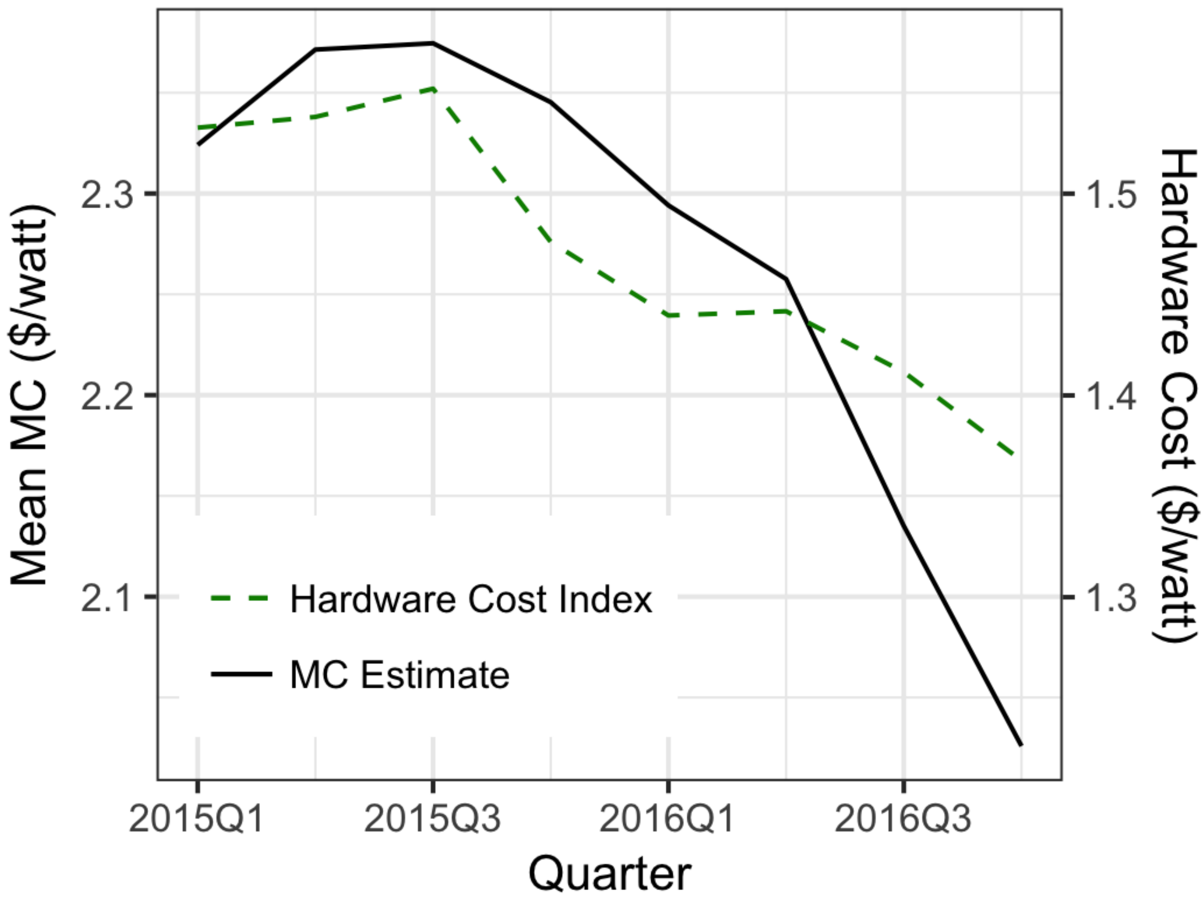

Notes: Hardware Costs were obtained from Bloomberg New Energy Finance. The Hardware Index plotted above is calculated by summing Bloomberg's multicrystaline silicon module overall avg spot price and the PV Insights residential inverter price index (also collected from Bloomberg).

Figure 18: Comparison of Marginal Costs to Existing Estimates

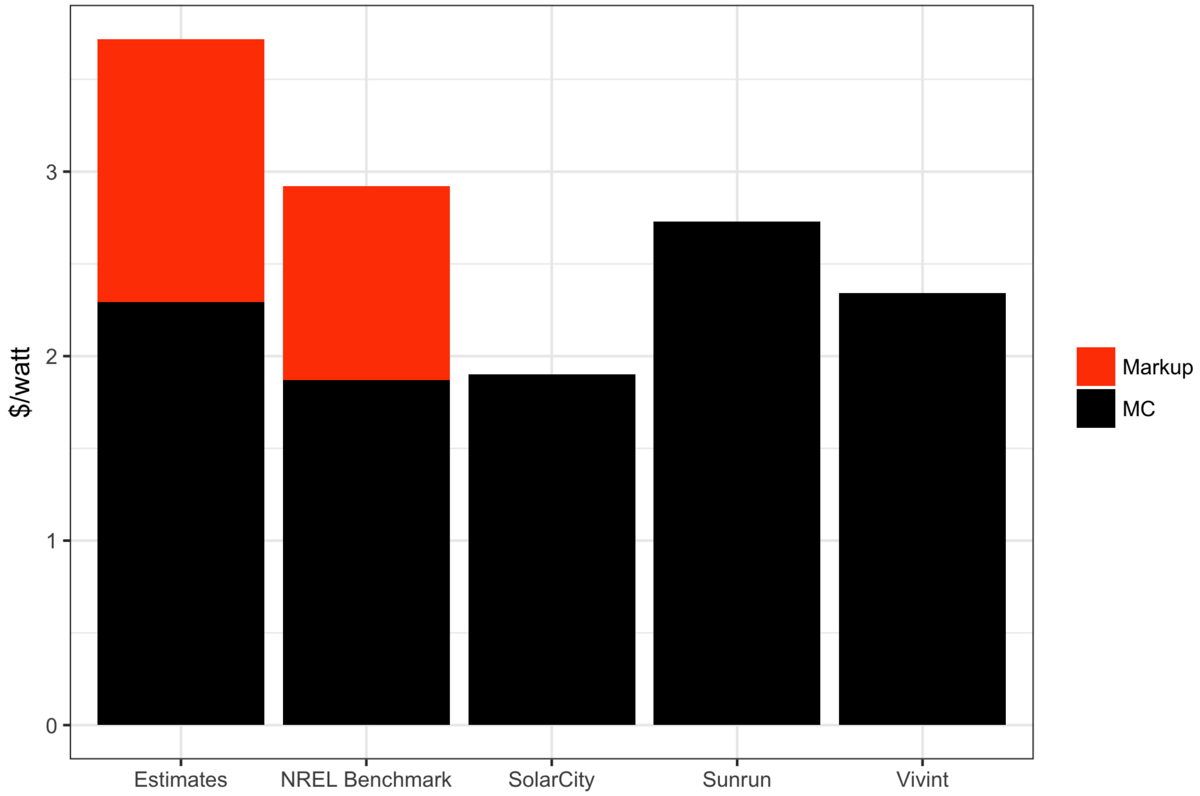

Notes: The first bar plots the average marginal cost and markup estimates from the structural model for 2016Q1. The second bar plots estimates from the National Renewable Energy Laboratory (NREL) 2016 Q1 benchmark analysis. The last three bars plot marginal costs reported in quarterly reports by large publicly traded installers, Solar City, Vivint, and Sunrun.

Table 18: Summary of Estimated Markups and Marginal Cost by Location and Time-Period

Panel A: Marginal Costs ($/watt)

State	Time-Period	Mean	SD
CA NORTH	2015 Q1	2.46	0.42
CA NORTH	2015 Q2	2.4	0.45
CA NORTH	2015 Q3	2.32	0.51
CA NORTH	2015 Q4	2.37	0.47
CA NORTH	2016 Q1	2.33	0.58
CA NORTH	2016 Q2	2.24	0.45
CA NORTH	2016 Q3	2.11	0.44
CA NORTH	2016 Q4	2.07	0.46
CA SOUTH	2015 Q1	2.21	0.42
CA SOUTH	2015 Q2	2.26	0.6
CA SOUTH	2015 Q3	2.4	0.48
CA SOUTH	2015 Q4	2.24	0.4
CA SOUTH	2016 Q1	2.24	0.39
CA SOUTH	2016 Q2	2.14	0.36
CA SOUTH	2016 Q3	2.02	0.44
CA SOUTH	2016 Q4	1.89	0.35
CT	2015 Q1	2.26	0.37
CT	2015 Q2	2.3	0.39
CT	2015 Q3	2.42	0.57
CT	2015 Q4	2.49	0.61
CT	2016 Q1	2.29	0.43
CT	2016 Q2	2.4	0.49
CT	2016 Q3	2.2	0.41
CT	2016 Q4	2.16	0.46
MA	2015 Q1	2.53	0.76
MA	2015 Q2	2.61	0.66
MA	2015 Q3	2.44	0.51
MA	2015 Q4	2.47	0.45
MA	2016 Q1	2.3	0.44
MA	2016 Q2	2.38	0.46
MA	2016 Q3	2.26	0.48
MA	2016 Q4	2.15	0.43
NY	2015 Q1	2.17	0.71
NY	2015 Q2	2.27	0.67
NY	2015 Q3	2.23	0.7
NY	2015 Q4	2.33	0.73
NY	2016 Q1	2.39	0.66
NY	2016 Q2	2.45	0.68
NY	2016 Q3	2.32	0.65
NY	2016 Q4	2.14	0.62

Panel B: Markups ($/watt)

State	Time-Period	Mean	SD
CA NORTH	2015 Q1	1.52	0.1
CA NORTH	2015 Q2	1.53	0.09
CA NORTH	2015 Q3	1.51	0.08
CA NORTH	2015 Q4	1.45	0.07
CA NORTH	2016 Q1	1.39	0.06
CA NORTH	2016 Q2	1.41	0.07
CA NORTH	2016 Q3	1.39	0.06
CA NORTH	2016 Q4	1.37	0.05
CA SOUTH	2015 Q1	1.53	0.14
CA SOUTH	2015 Q2	1.5	0.14
CA SOUTH	2015 Q3	1.47	0.11
CA SOUTH	2015 Q4	1.45	0.12
CA SOUTH	2016 Q1	1.38	0.08
CA SOUTH	2016 Q2	1.38	0.09
CA SOUTH	2016 Q3	1.38	0.08
CA SOUTH	2016 Q4	1.37	0.07
CT	2015 Q1	1.54	0.13
CT	2015 Q2	1.55	0.13
CT	2015 Q3	1.52	0.13
CT	2015 Q4	1.5	0.13
CT	2016 Q1	1.41	0.09
CT	2016 Q2	1.37	0.08
CT	2016 Q3	1.38	0.08
CT	2016 Q4	1.39	0.08
MA	2015 Q1	1.59	0.21
MA	2015 Q2	1.55	0.16
MA	2015 Q3	1.61	0.19
MA	2015 Q4	1.61	0.23
MA	2016 Q1	1.53	0.16
MA	2016 Q2	1.44	0.1
MA	2016 Q3	1.42	0.1
MA	2016 Q4	1.4	0.09
NY	2015 Q1	1.67	0.1
NY	2015 Q2	1.66	0.09
NY	2015 Q3	1.64	0.1
NY	2015 Q4	1.59	0.12
NY	2016 Q1	1.49	0.1
NY	2016 Q2	1.44	0.09
NY	2016 Q3	1.43	0.09
NY	2016 Q4	1.43	0.08

Figure 19: Non-Parametric Marginal Cost Distributions

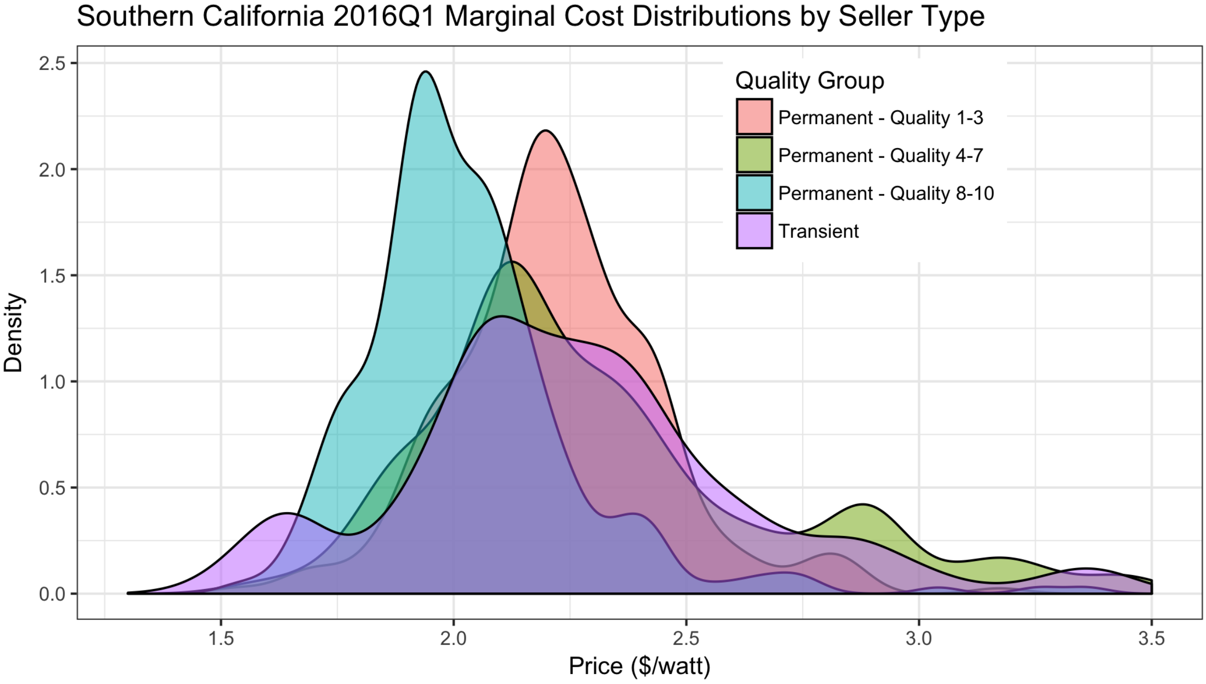

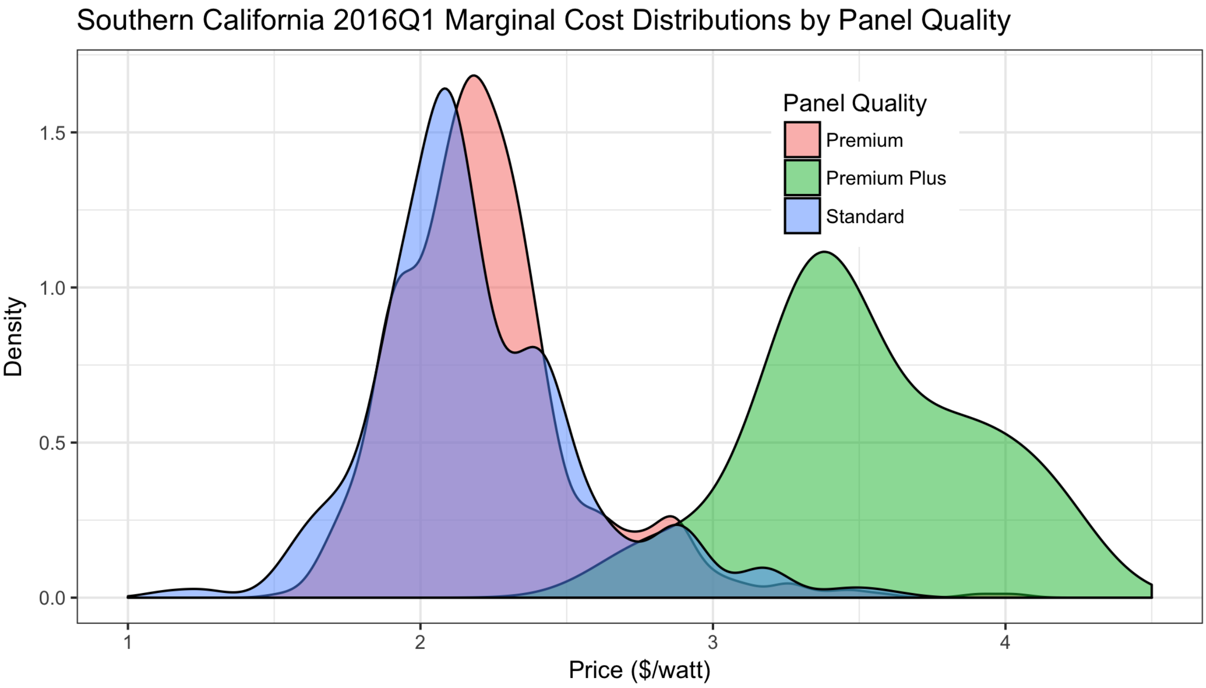

Table 19: Entry Cost Estimates

	Estimates
μ - Transient	13.196***
	(0.591)
μ - Quality Group 1	7.729***
	(0.329)
μ - Quality Group 2	10.574***
	(0.475)
μ - Quality Group 3	6.149***
	(0.183)
μ - Quality Group 4	9.047***
	(0.294)
μ - Quality Group 5	10.522***
	(0.366)
μ - Quality Group 6	8.812***
	(0.261)
μ - Quality Group 7	11.379***
	(0.400)
μ - Quality Group 8	12.886***
	(0.497)
μ - Quality Group 9	12.815***
	(0.477)
μ - Quality Group 10	10.169***
	(0.256)
σ - Transient	7.102***
	(0.482)
σ - Permanent	9.349***
	(0.640)
State FE	Yes
Year FE	Yes
Observations	280,670

Notes: 'Permanent Sellers' include all sellers in quality groups 1-10.

Figure 20: Average Number of Bids Received and Bid Prices ($/watt) by County

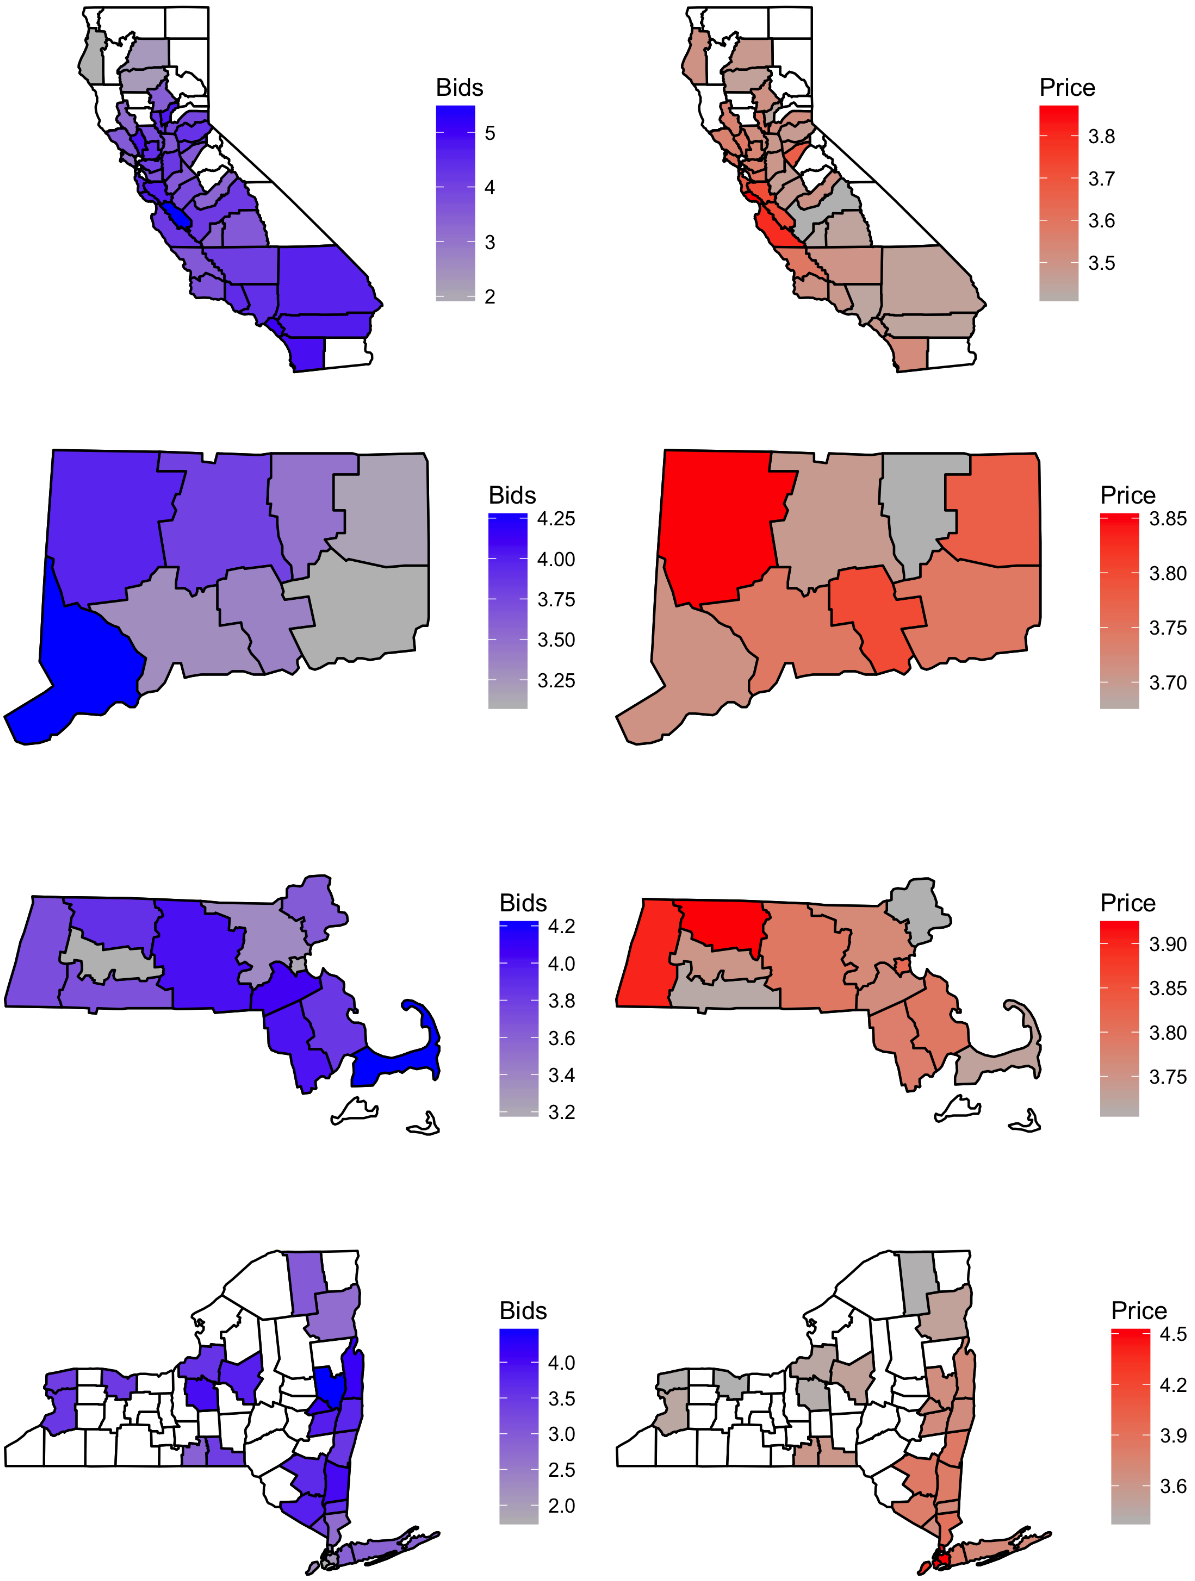

Notes: Maps only include counties that had at least 10 potential projects during the sample.