Access to Alternatives: Increasing Rooftop Solar Adoption with Online Platforms

Jackson Dorsey*

June 2, 2021

Abstract

Residential solar installation prices, which are set on a case-by-case basis, remain both relatively high and variable despite significant declines in solar panel manufacturing costs. In this paper, I estimate a structural model of the solar installation market to quantify market power and to evaluate the welfare effects of connecting buyers and sellers through an online platform. I find that the platform yields a $1,451 increase in consumer surplus by increasing the number of installation bids each household receives. Households with higher-valued homes attract relatively more bids and reap the largest benefits from the platform. Counterfactual simulations yield two main results: 1) an increase from one to five bids per project causes a 15.5% ($4,000) reduction in gross installation prices and a 33% increase in the number of solar installations, and 2) the solar Investment Tax Credit—which is scheduled to be eliminated in 2022—improves total welfare by mitigating market power in addition to reducing pollution externalities from electricity generation.

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 Business Economics and Public Policy, Kelley School of Business, Indiana University (e-mail: jfdorsey@iu.edu). This paper was previously circulated with the title: "Harnessing the Power of the Sun through Online Platforms: Evidence from the Rooftop Solar Market." I thank Ashley Langer, Derek Lemoine, Gautam Gowrisankaran, Stan Reynolds, Gary Solon, Derek Wolfson, Keith Meyers, Dev Bose, Alvin Murphy, Matt White, Stefan Weiergraeber, Dan Sacks, Emily Baca, Mattijs Wildenbeest, and seminar participants at Arizona State, University of Arizona, University of Maryland (AREC), Cornerstone Research, University of Mannheim, Federal Reserve Bank (Dallas), Charles River Associates, SUNY Albany, Indiana University (Kelley), ISO New England, International Industrial Organization Conference (IIOC), NBER (EEE) Summer Institute, the CU Environmental Economics Workshop, Stanford SITE, AERE Summer Conference, and the American Enterprise Institute for helpful comments, discussions, and suggestions. I also thank Luke Tarbi and Spencer Fields at EnergySage Inc. for providing instrumental help with data access.

Online platforms offer an increasingly common and potentially fruitful way to connect buyers and sellers in markets where services are customized and priced on a case-by-case basis.[1] Examples include business-to-business services—computer programming, legal, and marketing—and consumer services—auto repair, home contracting, and mortgages. In the absence of platforms, buyers in these markets must bear substantial indirect costs of searching for and comparing quotes across potential sellers. These search frictions have implications for market power—when collecting price quotes is costly for buyers, any seller asked to give a quote can expect to be bidding against few or no other sellers, thereby giving that seller incentive to charge a higher markup (Diamond, 1971; Stahl, 1989; Stigler, 1961). Moreover, sellers may charge particularly high prices to consumer groups with high search costs or those that are less price-sensitive (Allen et al., 2019). On the other hand, platforms collect and compile information about each buyer's characteristics, such as a desired project or task, and pass the data to a broad set of sellers that can then submit bids via the platform. Thus, platforms have the potential to provide buyers access to more price quotes than they would otherwise obtain offline through bilateral interactions with sellers. In this way, platforms can increase the number of transactions completed and improve market efficiency through several channels: connecting buyers to better-matched sellers, connecting buyers to lower-cost sellers, or reducing seller markups by increasing competition. In markets with positive externalities such as renewable energy, energy efficiency, and network/communications technologies, an increase in output will also have additional public benefits.

In this paper, I measure the value of an online platform for a product with positive externalities—rooftop solar photovoltaics (PVs)—and document how this value is distributed across consumers and the public. I do this by estimating a structural model that incorporates households' solar installation choice and sellers' strategically bidding on projects. First, I use the estimates to quantify market power in the installation market. Second, I measure the platform's effect on consumers' solar adoption choices and welfare. Third, I use the model to evaluate counterfactuals. In one series of counterfactuals, I measure the effects of changing the the number of installation bids each household obtains on equilibrium prices, the number of solar installations, consumer surplus, and environmental externalities. In another set of policy counterfactuals, I assess the welfare impacts of the largest U.S. solar subsidy program, the Investment Tax Credit (ITC), which is scheduled to be fully eliminated in 2022.

Understanding the effects of platforms on solar installation prices is both important and policy relevant. Although solar prices have fallen, and installations have grown substantially over the past decade, the price of an average solar installation still exceeds $20,000, and fewer than 3% of U.S. single-family homes have a solar system (Barbose et al., 2019). High installation prices cannot be explained by the hardware costs of the solar panels themselves. Figure 1

[1]For instance, Angie's List, Thumbtack, and HomeAdvisor provide platforms for most home and auto services. Upwork, PeoplePerHour, Freelancer Map, and Guru enable online transactions for a wide range of professional services. There are also many specialized platforms for particular services such as Chegg (tutoring), Envato Studio (graphic design), Care.com (child and senior care), or Zebra (auto insurance).

illustrates that hardware materials costs now account for less than 30% of the total price of a U.S. residential solar installation (Barbose et al., 2019; Borenstein, 2017; Feldman, 2014; Fu et al., 2016; O'Shaughnessy et al., 2019).[2] The balance of system cost, often called the "soft cost" in the solar industry, includes installation labor, permitting, inspection, and interconnection costs. Notably, the "soft cost" also includes installers' markups. The fact that soft costs make up a relatively large portion of solar installation prices, and that these costs have been falling relatively slowly over time, suggests that market power could be inhibiting residential solar investments (O'Shaughnessy et al., 2019). There is a host of evidence showing that market power is prevalent in the solar installation market. Pless and van Benthem (2019) develop a test for the presence of market power by measuring the rate at which solar installers pass-through subsidies to consumers; the authors' estimates imply the presence of market power. Borenstein (2017), Gillingham et al. (2016), and Nemet et al. (2016) show evidence that households with high electricity rates and higher expected incomes pay more for solar installations, suggesting that sellers engage in price discrimination. Notably though, many U.S markets now have dozens or even hundreds of solar installers. Thus, imperfectly competitive installation pricing is likely enabled by high search costs. Platforms could therefore reduce installation prices by expanding the number of installers considered by each potential solar buyer.

Figure 1: Residential Solar PV Installation Prices ($/Watt)

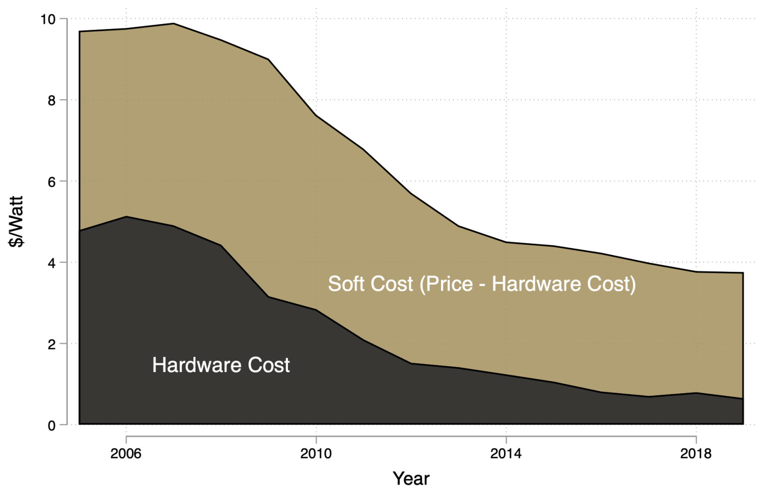

Notes: Figure constructed using data from Barbose et al. (2019).

Imperfect competition also has important implications for the welfare impacts of government subsidy programs. More specifically, the nature of competition in the installation market will influence the share of subsidies that is passed through to end-consumer prices (Weyl and Fabinger, 2013). If monopolistic sellers respond to subsidies by increasing gross prices, then subsidies may only generate small increases in solar installations and associated environmen-

[2] Hardware costs include the panel, inverter, support structures, and electrical wiring.

tal benefits. In this case, subsidies may even reduce overall welfare. On the other hand, it is plausible that subsidies could increase welfare by simultaneously correcting for two prevailing market distortions: 1) unpriced environmental externalities and 2) market power. With these ambiguous predictions in mind, an empirical analysis is needed to evaluate whether government subsidies are welfare improving.

To investigate counterfactual market outcomes and welfare in the solar PV market, I develop a model that includes three stages. In the first stage, a solar installation project is announced. Upon the project announcement, potential sellers learn about the project's characteristics, such as the location of the home seeking the solar installation, the household's monthly electricity use, and the market value of the home. Potential sellers decide whether to participate in the auction by comparing the costs associated with preparing a quote to the expected marginal profits of bidding (Krasnokutskaya and Seim, 2011).[3] The expected profitability from participating and the expected number of bidders depends on the household's characteristics. In the second stage, participating sellers pay the bid preparation cost and learn their marginal cost associated with installing the solar system for that particular household. The participating sellers then submit a price bid to maximize expected profits. An optimal bid is a function of the installer's marginal cost, the household's price elasticity, and expectations about the number of competing bids for that project. In the final stage, the household observes the bids and chooses one of the bidding sellers or the outside option—not installing a solar system or an off-platform installer. Notably, the auction allocation mechanism in this setting differs from a conventional first-price auction. When making a choice, buyers can consider factors other than prices, such as the installers' ratings, experience, and hardware brand. Therefore, I estimate the buyer's choice rule because the allocation mechanism is unobserved. I build on and adapt methods recently developed by Krasnokutskaya et al. (2019) and Yoganarasimhan (2015) to recover parameters in multi-attribute auction settings.[4] The innovation of my model is to account for several key features of the solar market that have not been considered in the previous literature. Namely, I model both strategic bidding by solar installers, as well as imperfectly informed buyers—buyers are only informed about installers that submit bids for their project.

I estimate the model using a novel data set that includes over 37,000 solar installation bids made to over 10,000 households through an online platform. The data is unique relative to most other solar PV data used in the literature in a few ways. First, I can observe non-winning bids and bids made to households that did not purchase a solar system. Observing the number of quotes at the household level allows me to obtain more precise estimates of elasticities and markups.[5] I also observe considerable variation in the number of bids across households and over time. I utilize this variation to identify how prices and buyer installation decisions change

[3]Expected marginal profits of bidding are equal to the probability of being selected, multiplied by the installer's expected markup.

[4]These types of auctions have also been called beauty contest auctions. These auctions are similar to scoring auctions but differ in that the precise scoring rule is not announced to sellers.

[5]Goeree (2008) shows that demand curves are biased toward being too elastic under traditional models that assume the buyers are fully informed about all sellers or products.

with the number of bids per project.

I augment the raw platform data with detailed household-level demographic information. The platform collects and reports information about each household's address and its average monthly electricity expenditure.[6] I supplement that information with marketing-research data on each home's market valuation and demographic variables like the age, race, and educational attainment of the household head. One limitation of the raw platform data is that the sample may not represent the universe of potential solar buyers because households select into using the platform. Consequently, the detailed household information is valuable for a few reasons. First, I use the demographic data to show that platform shoppers appear relatively similar to other solar adopters in the broader market based on observable characteristics. Second, because I observe significant variation in home values—ranging from $100,000 to over $2 million—I am able measure differences in seller participation and pricing across households of differing wealth. Thus, the detailed household information allows me to assess which types of households benefit most from using the platform.

Findings The model estimates imply relatively high markups in solar installation pricing and sizable benefits of using the online platform. I find that markups account for 45% of gross installation prices.[7] On the demand side, I also find that access to the platform increases consumer surplus by an average of $1,451 per solar shopper, or about 8% of the average solar system price. While all households experience some expected welfare gains, these gains differ considerably across markets and household demographics. The platform generates a $1,156 gain in consumer surplus for households in the bottom quartile of home valuation, relative to a $1,476 gain for households in the top quartile. Households with higher valued homes attract more bids through the platform and purchase larger solar systems, which contributes to the disparity in the platform's welfare impact.

The first set of counterfactuals demonstrate that increasing the number of bids obtained by each household has significant impacts on consumer welfare and market outcomes. In the counterfactuals, I vary the number of sellers bidding on each project and sellers update their profit-maximizing bid prices accordingly. The counterfactual exercise reveals that an increase from one to five bids per project causes a 15.5% decline in gross sale prices, a $4,000 reduction on a typical installation. This five-fold increase in bids also promotes a 33% rise in the number of solar installations, a 28% reduction in pollution damages, and a 360% increase in consumer surplus. These findings underscore the potential benefits of reducing search frictions in negotiated price markets.

The second set of policy counterfactuals show the impact of eliminating the federal Investment Tax Credit (ITC) on prices, the number of solar installations, and welfare. The ITC enables solar buyers to deduct 30% of the gross price of a solar installation from their federal taxes; however, policymakers plan to eliminate the ITC in 2022.[8] The counterfactuals reveal

[6] The physical addresses are observed by platform staff but were not disclosed to the author for privacy reasons.

[7] The gross price is the price before subsidies such as the Investment Tax Credit are applied.

[8] The ITC will offer a 26% subsidy in 2020, a 22% credit in 2021, and 0% subsidy in 2022 and beyond.

that removing the ITC increases net sale prices by about 24%. As a result, solar installations and the associated environmental benefits decline by 33% in the absence of the tax credit. The counterfactuals also reveal that eliminating the ITC reduces total welfare by 22%. Surprisingly, I find that the most substantial benefit of the ITC is that it reduces deadweight loss from the exercise of market power. More specifically, each dollar spent on the ITC generates a \$0.77 gain in producer surplus[9] along with a \$0.80 gain in consumer surplus. The subsidy increases market participants' rents by enabling more transactions where the buyer's willingness-to-pay exceeds the seller's marginal cost than would occur without the policy. Moreover, the ITC promotes a modest decrease in pollution externalities—pollution damages fall by \$0.19 per subsidy dollar. Taking all these factors together, I find that the policy increases total net welfare by \$0.69 per subsidy dollar after accounting for the cost of the subsidy.

The results have several important implications for public policy. The structural estimates along with the first set of counterfactuals demonstrate that: 1) installer markups account for a large share of solar installation prices, 2) access to the online platform increases consumer surplus and solar adoption rates, and 3) access to additional bids is a key mechanism that can explain the platform's benefits. However, as of 2016, only 3% of solar buyers used an online platform to purchase a PV system (O'Shaughnessy and Margolis, 2018). Some policymakers have taken steps to steer more buyers and sellers towards platforms. As one example, Connecticut recently introduced a state-sponsored platform in the solar PV market.[10] Policymakers have also used platforms to promote competition in other industries such as healthcare.[11] Therefore, one implication of the results is that policymakers could consider developing their own platforms or encouraging participation on existing platforms to reduce solar PV prices and increase solar PV adoption.

The second set of counterfactuals provides the first welfare analysis of a solar subsidy policy that incorporates consumer welfare, producer welfare, and environmental externalities. Renewable energy subsidies have been a cornerstone of environmental policy over the past two decades. More recently though, policymakers have removed, or considered removing, several existing subsidy policies, such as the ITC. I find that removing the ITC causes a reduction in total welfare by reducing market participants' rents and increasing environmental damages by more than the cost of the subsidy itself. To my knowledge, this is the first paper to account for changes in producer surplus when assessing the welfare effects of renewable energy subsidies. Measuring these producer surplus gains is pivotal in the cost-benefit analysis because consumer surplus gains and environmental benefits alone are not sufficient to cover the subsidy cost.

The results also provide insights on the economics of two-sided platforms more broadly.

[9] Producers' installation marginal profit grows by \$0.80 per subsidy dollar. However, bid preparation costs also increase by \$0.03 per subsidy dollar because the ITC encourages more sellers to bid on projects.

[10] For details see the Connecticut Greenbank's website gosolarct.com operated under Connecticut Legislature's Public Act 11-80.

[11] 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 establishing new individual health insurance marketplaces (platforms) where consumers can shop for, compare, and purchase plans.

Notably, many other industries have shifted transactions towards online platforms in the past decade. In markets that are characterized by significant search costs—-such as home mortgages or customized consumer services—platforms can deliver substantial price reductions and aggregate consumer welfare gains. However, the empirical results also show that magnitude of these welfare gains can differ substantially across consumers when sellers can discriminate on both the extensive margin (selection into bidding) and intensive margin (bid prices). More specifically, the findings indicate that price-sensitive consumers—who stand to gain the most from price savings—–may benefit relatively less from platforms if they attract relatively fewer bids via the platform.

Related Literature This paper builds on multiple literatures. First, it relates to a growing literature on the economics of the residential solar market. The large portion of existing work has focused on estimating the adoption response to price or subsidy changes (Burr, 2014; De Groote and Verboven, 2019; Feger et al., 2017; Gillingham et al., 2016; Gillingham and Tsvetanov, 2019; Hughes and Podolefsky, 2015; Langer and Lemoine, 2017; Reddix II, 2015; Snashall-Woodhams, 2019).[12] I build on this work by estimating solar PV demand using information about household-level bid prices and choice sets. Most previous work focuses on the binary decision of whether or not to install solar and has abstracted away from the choice of installer and competition among installers.[13] Identification of price elasticity of demand in these papers typically exploits variation in subsidy availability across households or changes in mean prices over time. In contrast, I exploit variation across bid prices made to a specific household to identify price elasticities.

Supply-side behavior and incentives are also crucial in determining equilibrium prices and technology adoption decisions. However, most previous literature has not considered, or explicitly modeled, the supply-side of the solar PV market. To fill this gap, I develop a model that accounts for several important features of the market: (1) imperfect competition due to imperfectly informed consumers, (2) selection by sellers into bidding on projects, (3) strategic pricing conditional on each household's characteristics, and (4) sellers that have imperfect information about their competition when placing bids. The model allows me to account for both supply and demand responses to counterfactual environments, such as changes to subsidy policies. One other paper that estimates a supply-side model of the solar PV installation market is Bollinger and Gillingham (2019), who estimate a dynamic supply model to decompose static markups from dynamic pricing incentives, driven by learning-by-doing. In another related paper, O'Shaughnessy and Margolis (2018) compare installation prices paid by solar consumers

[12] Hughes and Podolefsky (2015), Pless and van Benthem (2019), and Gillingham and Tsvetanov (2019) 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 (2019), Snashall-Woodhams (2019), 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.

[13] One exception is Bollinger and Gillingham (2019) who estimate a dynamic discrete choice model using average prices for each installer in each market and assume that buyers are fully informed about all installers. In another related paper Gerarden (2017) formulates a model of competition in the upstream solar panel manufacturing market.

who used an online platform and consumers that did not. They find that platform users pay significantly lower prices for the same solar PV hardware.[14] I build on O'Shaughnessy and Margolis (2018) by developing a structural approach that allows me to assess the effects of the platform on markups, consumer welfare, adoption choices, and pollution externalities.

This paper also pertains to an extensive literature on competition in search markets, the role of platforms, intermediaries, and the internet. Seminal work by Baye and Morgan (2001), Gehrig (1993), Hall and Rust (2003), and Spulber (1996), theoretically investigated the role of intermediaries in search markets. 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), retail gasoline (Luco, 2016), books (Ellison and Ellison, 2009), and freelance computer programming (Krasnokutskaya et al., 2019, 2016). While many previous studies have assessed the impact of online platforms or the internet on average prices, aggregate consumer welfare, or the distribution of producer gains (Goldmanis et al., 2010), there has been little work on which types of consumers are likely to benefit most from online platforms in negotiated price markets. I build on the literature by providing one of the first distributional analyses of consumer welfare gains while accounting for endogenous seller bidding participation and pricing.

The remainder of the paper proceeds as follows: in the next section, I discuss the details of the online platform, provide descriptive statistics, and show regression evidence of selective seller participation and discriminatory pricing. 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 in Section 4. Section 5 presents the welfare and counterfactual results, and Section 6 concludes.

2 Background and Data

Shopping for a rooftop solar system is time-intensive because installers generally do not post prices. Therefore, inquiring buyers often need to call installers and schedule site visits to obtain project proposals and price quotes. Because search is costly in this market, many buyers only receive a limited number of price quotes, which increases the incentive for sellers to exercise market power. A 2017 survey conducted by the National Renewable Energy Laboratory found over 80% of solar shoppers obtained two or fewer quotes before making a decision (Sigrin et al., 2017). Recently, online platforms have emerged as an alternative to bilateral negotiations between buyers and sellers.

As one example, the U.S. based quote aggregation platform EnergySage Inc. facilitates connections between potential solar customers and a network of solar PV installers. More specifi-

[14] In ongoing work, Bollinger et al. (2020) exploit experimental variation in the number of active installers in a market to estimate the effect of competition on prices and solar adoption.

cally, the EnergySage platform enables households to conduct multi-attribute auctions to select an installer for their installation projects. Multi-attribute auctions, also called beauty contest auctions, refer to a procurement mechanism in which each bidder submits a multi-dimensional bid that includes a price and a vector of other characteristics such as the panel hardware brand. The buyer then selects the winning bidder based on the multi-dimensional bids. Multi-attribute auctions are related to scoring auctions but differ in that the auctioneer does not explicitly announce the choice rule ex-ante (i.e., weights on each characteristic) as they would in a scoring auction. Each EnergySage auction includes several stages. First, consumers create an account on the platform's website and provide information including: the physical address of the potential installation, a monthly electricity bill, and an indication of whether they have obtained other solar installation bids offline.[15] Second, registered installers[16] receive a notification of the project which includes details such as a Google Maps photo of the buyer's roof (depicted in Figure 6 of the appendix), as well as the monthly electricity usage of the buyer, and the information about whether the buyer has other off-platform quotes. Installers can then submit a project quote to the buyer, which includes the system price, panel brand, inverter brand, and details about the seller, such as a star 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.[17] Figure 7 in the appendix shows an example of the purchaser's comparison tool on the platform.

A distinguishing feature of the multi-attribute auction 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. Solar systems can be payed for as a cash purchase, purchased via loan, or leased from an installation company. Although many installers offer leases or power purchase agreements, 97% of buyers on EnergySage choose to buy a system with cash or a loan. The significant skew towards host ownership instead of leases is likely because EnergySage shows buyers a calculation of the net present value of each contract. It turns out that purchase agreements nearly always offer a higher net present value relative to lease agreements. The trend towards host ownership is not unique to the platform, overall market shares for leased systems have been declining considerably in recent years. Barbose et al. (2019) reports that the market share for leases among residential systems grew dramatically from 2007 to 2012, reaching nearly 60%. However, lease shares have fallen sharply in recent years, with third-party ownership comprising just 38% of installations in 2018. The percentage of solar leases is expected to continue to fall due to the emergence of residential loan products and a move away from the third-party ownership model by SolarCity/Tesla and other major national installers.

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

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

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

2.1 Data and Descriptive Statistics

This study's primary data set includes bid prices and consumer purchase choices for solar auctions on the EnergySage platform. The sample includes auctions originating between 2014 and 2016 within the states of Arizona, California, Colorado, Connecticut, Massachusetts, New York, and Texas.[18] EnergySage collects information on each household's address, the household's average monthly electricity bill, and an indicator for whether or not the buyer obtained other solar installation quotes from off of the platform. EnergySage did not collect consumer demographic information during the sample period of this study, so I also augment the EnergySage data with household demographic data from Axciom's Infobase obtained through Infinite Media, Inc. This consumer marketing database compiles household-level data from various sources such as public records, (e.g., property ownership, professional licenses, or voting registries), and consumer surveys (e.g., magazine subscriber lists, catalogs, and warranty cards.). Infinite Media, Inc. successfully appended demographic information for over 80% of the EnergySage households. I drop all households without a complete set of demographic variables and all projects where a lease agreement was selected.[19] I also drop a handful of other observations that appeared to either be miscoded or outliers.[20] Additionally, I merge each project with environmental benefits estimates from Sexton et al. (2018) by ZIP code.[21]

Table 1 provides descriptive statistics of the 10,253 potential installation projects from the EnergySage platform. The average project received 3.54 bids, however, there is noticeable variation in the number of bids across projects, which I illustrate in Figure 2a. A quarter of projects obtained two or fewer bids, whereas 25% of projects received five or more bids. Figure 2a shows that there is substantial variation in the mean number of bids obtained across the core-based statistical areas. For example, projects in San Diego-Carlsbad, CA remain substantially more competitive than those in Tucson, AZ throughout the sample.[22] We can also see that the mean number of bids per project trends upward throughout the three-year sample period. At the same time, Figure 2b demonstrates that mean bid prices fell consistently. The households in the sample differ from one another in a number of ways. For one, households vary in electricity expenditure. The average household reported an electricity bill of $187/month with a standard

[18]In the analysis, I only consider residential projects from 2014 onward because both buyer and seller participation on the platform was very limited in the platform's first year, 2013. I also drop projects in the final quarter of 2016 because some buyers' choices were still pending when the data set was compiled.

[19]I drop these projects because comparing per-watt prices for leases vs. purchases is not straightforward. Furthermore, these projects compose less than 4% of the choices and thus discluding them is unlikely to have significant effects on the analysis.

[20]I drop projects with system capacities (mean across bid proposals) under 3 kW and over 15 kW. I drop price quotes under $2/watt or over $7/watt. I drop households that reported monthly electric bill under $50 or over $500. Finally, I drop projects with a home market valuation below $100,000.

[21]Sexton et al. (2018) calculate the annual pollution damages avoided in dollars per kilowatt of residential solar capacity for each ZIP code. I scale the pollution damage estimate by the capacity of the system and assume a 20-year system life span with a 5% annual discount rate to determine the net present value of environmental damages avoided.

[22]Projects that are outside a core-based statistical area (CBSA) or are in CBSA with fewer than 100 total projects are placed into a new category by state. For example, a project located in Aspen, Colorado are defined as "Other, CO"

Table 1: Project Summary Statistics

	Mean	SD	25-%tile	50-%tile	75-%tile
Number of Bids	3.54	1.84	2.00	3.00	5.00
Home Market Value ($1,000s)	687.49	508.66	325.00	550.00	900.00
Age of Home	42.24	33.23	17.00	36.00	58.00
Home Size (sqft)	2238.57	863.92	1608.00	2100.00	2692.50
Env. Damage Avoided ($/KW-yr) - ZIP	46.69	18.42	26.12	51.84	53.34
Annual Solar Output (KWh/KW) - ZIP	1428.28	199.17	1198.72	1513.38	1559.34
Electricity Bill ($/month)	187.03	89.38	120.00	167.28	240.00
HH Head Age	51.86	13.86	42.00	50.00	62.00
HH Head Race - Asian (0,1)	0.12	0.32	0.00	0.00	0.00
HH Head Race - Black or Hispanic (0,1)	0.12	0.32	0.00	0.00	0.00
HH Head Holds Bachelor's Degree (0,1)	0.49	0.50	0.00	0.00	1.00
Has Off-Platform Quotes (0,1)	0.17	0.38	0.00	0.00	0.00

Notes: The number of bids, household monthly average electric bill, and an indicator for whether the consumer has off-platform quotes are recorded and reported directly by EnergySage. Annual environmental damages avoided per kW capacity are calculated at the zip code level by Sexton et al. (2018). Additional variables come from Infinite Media's consumer database and were merged with the EnergySage data by property address.

deviation of $89/month. The monthly electricity bill is a key variable because it helps determine the optimally-sized solar system for the household. Additionally, we see differences in search behavior, about 17% of the households reported having other quotes from installers off of the platform. We also see that the households vary in the expected solar generation productivity and environmental benefits depending on their ZIP code.

I acquire the remainder of the consumer data from Infinite Media, Inc. The data contain two types of information about the households: (1) data on the home itself such as the age of the structure, the square footage, and the home's market valuation, (2) information about the head of the household such as age, race, and whether they have a college degree. To verify the precision of the data merge, Figure 10 in the appendix plots a binned scatter plot of the home square footage from the Infinite Media consumer data against the monthly electricity expenditure reported by households directly to the platform. The figure shows that the variables from each data set are very highly correlated in the way that we would expect.

Table 9 in the appendix shows the composition of projects across states and across time. We see that the number of projects initiated through the platform grew dramatically over the sample period. Each year the number of projects roughly triples. We can also see that over 40% of the projects are located in California. This pattern is consistent with the fact that California is home to over 40% of solar installations in the U.S. Figure 9 maps the locations of all of the projects in the sample and we can see that the projects tend to be more concentrated in larger metropolitan area and cities.

A potential disadvantage of the platform data is that platform users may not be representative of the population of solar shoppers. An important question is whether platform users are likely to have systematically different willingness-to-pay relative to solar shoppers off of the platform? The answer to this question has important implications for the interpretation and

Figure 2: Competition and Prices by Core-Based Statistical Area Over Time

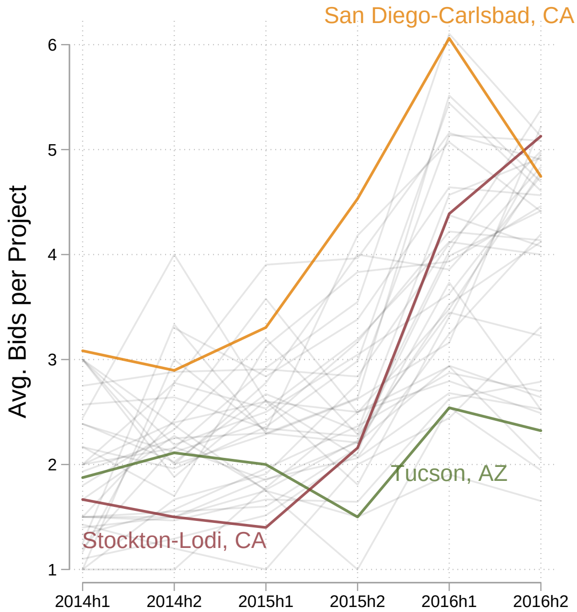

(a) Mean Bids Per Project

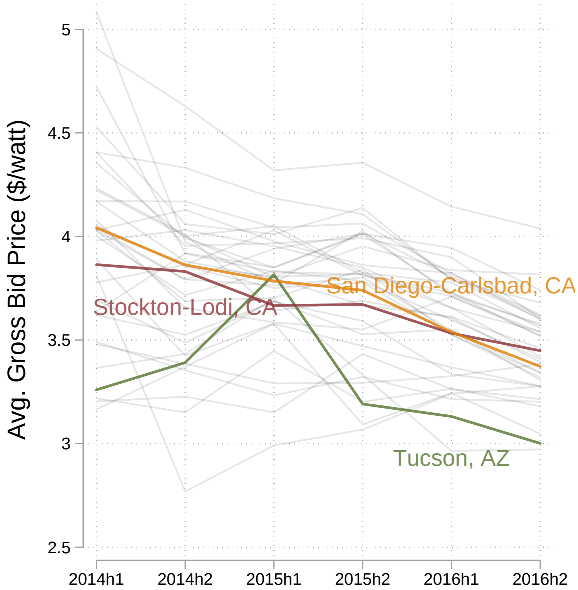

(b) Mean Bid Prices

Notes: Each grey line represents a CBSA. Projects that are outside a CBSA or in CBSA with fewer than 100 total projects are placed into a distinct category by state. For example, a project located in Aspen, Colorado would be defined as "Other, CO".

external validity of the results.

Although, I am not able to test these questions directly, I provide a comparison of the income distribution of households using the platform with the income distribution of the broader population of solar adopters. To do so, I collect information from Barbose et al. (2018) on incomes for U.S. households with solar installations. The household-level data from Barbose et al. (2018) is not publicly available but the authors calculate the income distributions of solar adopters by state and year.[23] The household incomes in Barbose et al. (2018) are reported in $50k-width bins so to make the two data sets comparable, I aggregate my income data for the platform shoppers into three income categories: "<$50k", "$50k-100k", and ">$100k". Figure 11 displays the income distributions for three groups: (1) all solar PV adopters, (2) households that purchased a solar system through the platform, and (3) all households that used the platform including those that did not buy a solar system. We see that the income distribution for population of solar PV appears relatively similar to the income distribution of the subset of households that purchased a solar system through the platform. Solar adopters are relatively affluent overall (Borenstein, 2017), over 55% of adopters have household incomes above $100k and only 13% have an income under $50k. Similarly, 11% of solar adopters in the EnergySage sample had

[23]I only consider the solar PV household income data from Barbose et al. (2018) for the years and states used for the main analyses: AZ, CA, CO, CT, MA, NY, and TX in 2014-2016.

incomes below $50k but were slightly less likely to have incomes above $100k (46%). Thus, platform users appear relatively similar to the broader population of solar buyers in terms of income. If anything, platform users are slightly lower-income relative to the broader population of solar PV adopters. Despite the similarity in income distribution, there could be other differences between platform users and non-users and, therefore, the reader should interpret the results with this caveat in mind. An advantage of the data is that there is a large amount of demographic heterogeneity across households within the sample. Therefore, when estimating the demand model, I focus on estimating heterogeneity in preferences across households with different wealth and energy use behavior. Estimating these heterogeneous preferences allows me assess which types of households benefit most from using the platform.

Characteristics of Solar Bids

Table 2 provides information about the characteristics of the bids submitted for these projects. There are several variables that can differ across each bid: bid price in dollars per watt, proposed capacity of the system to be installed, the brand of the solar panel, and the type of inverter. Buyers often obtain quotes for several different types of panel (module) brands[24] and EnergySage shows buyers a rating classification of each panel brand. For example, premium panels have higher efficiency and better warranties.[25] More efficient panels are attractive because for a given physical system size[26] a more efficient panel will create more electricity.[27] About 34% of bids in the data are for panels designated as "premium". EnergySage also identifies the very highest quality panels as "premium plus",[28] however, these panels are much more expensive and less than 4% of bids offer a "premium plus" panel. Another defining component of a bid is the inverter to be installed. The inverter converts the direct current (DC) output of the PV panel into alternating current (AC). String inverters are the cheapest 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 may fail to achieve optimal 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. I combine microinverters and power optimizers into a single category and I refer to them broadly as systems with microinverters.[29] Bids including microinverter technology are common on the platform comprising 73% of total

[24] 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.

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

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

[27] The platform only allows each seller to place a single bid. For example, a seller cannot place two different bids for different panel qualities.

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

[29] The industry often refers to this class of inverters as module level power electronics.

bids.[30]

Another factor that can influence a buyer's choice between bids is the quality of the seller. EnergySage provides several metrics that allow potential buyers to gauge the quality of each seller including a star rating, information about the installer's experience, as well as customer reviews. In the data, each sellers' name is anonymized, but I observe an unique installer ID that allows me to link each installers bids and performance across projects. I also observe each sellers' rating as of the end of 2016. Additionally, I observe each installers' experience in terms of total number of residential solar installations completed as of late 2016. About 60% of bids came from installers that had a five-star (out of five) rating at the end of the sample, whereas 18% of bids came from installers that had not yet been rated. The average bid came from an installer with nearly 2,400 previous installations completed, but there was notable variation in experience.[31]

Panel B of Table 2 shows the difference in mean characteristics for winning versus non-winning bids for a single market, the Los Angeles core-based statistical area (CBSA) in the first half of 2016. We see that winning bids had lower prices on average, but we also see that winning bids were more likely to be from five-star-rated installers and are more likely to have premium panels and microinverters. These correlations suggest that non-price factors are also important in the buyer's choice between bids.

Table 2: Summary Statistics - Bid Characteristics

Panel A: Full Sample

	Mean	SD
Total Gross Price ($ 1000s)	25.79	(10.22)
Unit Price ($/watt)	3.57	(0.49)
System Capacity - KW	7.30	(2.88)
Premium Panel (0,1)	0.34	(0.47)
Premium Plus Panel (0,1)	0.04	(0.21)
Microinverter (0,1)	0.73	(0.44)
Installer Rating = 5 Star (0,1)	0.60	(0.49)
Installer Rating = 4.5 Star (0,1)	0.08	(0.27)
Installer Rating $\leq$ 4 Star (0,1)	0.18	(0.38)
No Ratings (0,1)	0.18	(0.38)
Experience: # of Installs (1000s)	2.40	(4.55)
Observations	37080	

Panel B: Los Angeles CBSA - 2016H1

	Selected Bid (0,1)			
	0		1	
	Mean	SD	Mean	SD
Total Gross Price ($ 1000s)	23.59	(8.83)	22.74	(7.03)
Unit Price ($/watt)	3.53	(0.29)	3.41	(0.22)
System Capacity - KW	6.72	(2.55)	6.60	(1.74)
Premium Panel (0,1)	0.54	(0.50)	0.64	(0.49)
Premium Plus Panel (0,1)	0.01	(0.10)	0.09	(0.29)
Microinverter (0,1)	0.78	(0.42)	0.86	(0.35)
Installer Rating = 5 Star (0,1)	0.67	(0.47)	0.73	(0.46)
Installer Rating = 4.5 Star (0,1)	0.01	(0.09)	0.00	(0.00)
Installer Rating $\leq$ 4 Star (0,1)	0.24	(0.42)	0.18	(0.39)
No Ratings (0,1)	0.24	(0.42)	0.18	(0.39)
Experience: # of Installs (1000s)	3.19	(3.88)	3.16	(3.83)
Observations	964			

Notes: The installer ratings and experience variables were recorded at the end of 2016 and therefore do not vary across auctions for a given installer.

[30] 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.

[31] There are a several installers that do not report a number of residential installations that they have completed, I set these installer's experience level equal to the median installation experience in the overall sample.

2.2 Descriptive Evidence: Selection and Price Discrimination

Having discussed the data, I present several descriptive results about sellers' participation and bid prices that motivate the model developed in the next section.

There are a few mechanisms by which platforms could facilitate increases in solar PV adoption and improve consumer welfare. First, the availability of the platform could change the number of bids obtained by each household. An increase in the number of bids could increase solar adoption rates if sellers are differentiated in cost or in quality. If sellers have heterogeneous costs, then more bids equates to giving each household more cost draws and consequently a lower expected minimum bid price even if sellers' pricing strategy is held fixed. Similarly, if sellers are vertically or horizontally differentiated, then more bids per project could lead to more solar purchases by linking buyers to higher quality or better-matched sellers. Seller quality could vary due to ratings, reviews, experience, warranties, and other service offerings. Additionally, a change in the number of bidders will change each firms' optimal pricing strategy. Thus, the equilibrium effects of the platform on prices, solar adoption, and consumer welfare hinges not only on how many sellers are registered for the platform, but also on how sellers choose projects to bid on and how sellers set prices.

Figures 3a and 8a (Appendix) illustrate variation in auction participation across two important household characteristics: home market valuation and monthly electricity expenditure. Figure 3a depicts a binned scatter plot with home market valuation on the x-axis and the mean number of bids obtained in each bin on the y-axis, as well as a quadratic fit line. We see that homes with higher market valuations attract more bids through the platform. Sellers bid 30% more frequently on homes valued over \$1 million compared to homes worth under \$300,000. In Figure 8a we see an analogous pattern with electricity expenditure, households with monthly bills below \$100 get fewer than three bids on average, relative to nearly four bids for households with bills above \$300/month.

Figure 3b (and 8b) reveals disparities in the bid prices across households. The figure plots a binned scatter plot with the total installation price (before tax credits) on the y-axes. The total installation price for each observation is linearly adjusted for capacity (kW) and the time that the project originate (half year) so that each observation is more comparable. Therefore, the standardized prices should be interpreted as a bid price for a mean-sized system in the first half of 2016 (H1). We see that households with more expensive homes (and those with higher electricity bills, see 8b) receive higher size-adjusted bid prices. These higher prices are despite the fact that these projects are more competitive on average, as shown in 3a. While these pricing disparities could be linked to systematic differences in costs, this relationship is suggestive that installers are bidding higher prices to households that are likely to be more inelastic.

Figures 3c and 8c further investigate bid prices differences across households. Figure 3c, again, plots the relationship between bid prices and home market valuation, but plots the relationship separately across markets with varying degrees of competition. I define a market as a CBSA-half-year. I then separate households into four categories based on the market in which

Figure 3: Auction Participation and Pricing by Home Market Value

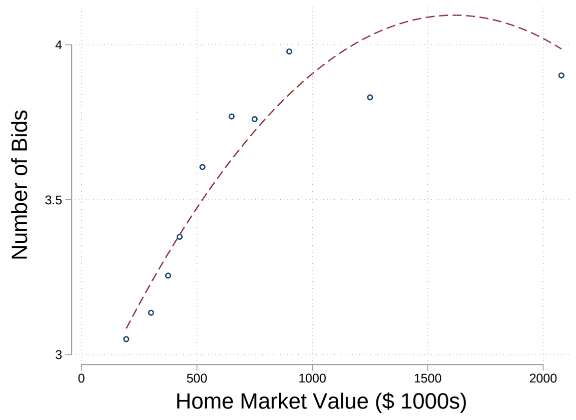

(a) # of Bids by Home Value

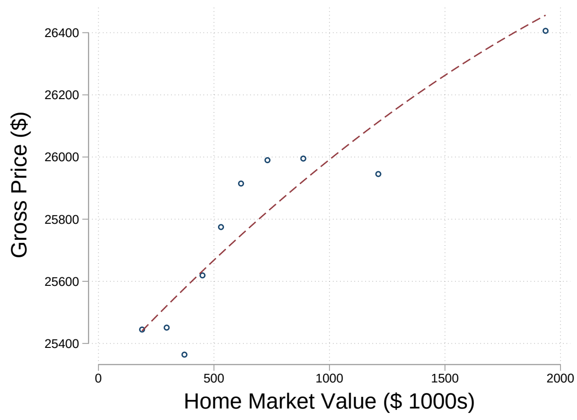

(b) Bid Prices by Home Value

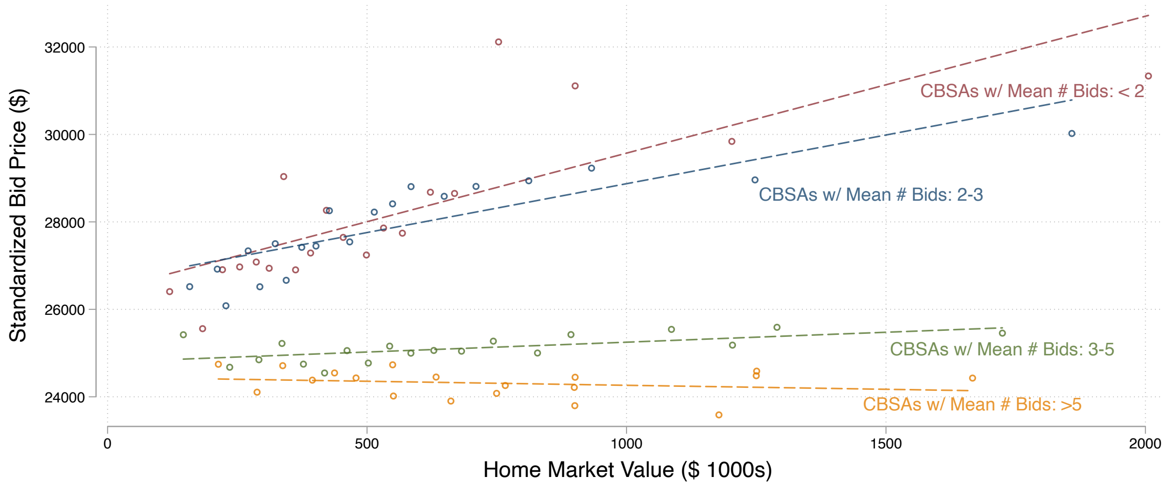

(c) Prices, Home Value, Mean Bids in CBSA

Notes: Panels b,c linearly adjust the prices for the capacity (kW) and the time (half year) the project occurred before plotting. In panel c, the mean bids in the market is the average number of bids across all projects within the same CBSA and the same half year.

they are located: markets with under 2 bids per project on average, markets with 2-3 bids, markets with 3-5 bids, and markets with over 5 bids on average. There are several notable patterns. Unsurprisingly, we see that bid prices made in more competitive markets are systematically lower across all home values. This relationship is also consistent with the correlation shown in Figure 2 that bid prices fall over time as bidding competition increases. However, we see that the relationship between home valuation and bid prices changes with competition. In less competitive markets, those with fewer than three bids per project, there is a large increase in bid prices associated with both higher home values and higher electricity bills. In contrast, we see that bid prices are relatively constant across home valuations and energy bills in more competitive markets with more than three bids per project. An implication of these patterns is that more bidding competition is associated with a large reduction in bid prices for households with expensive homes or high electricity expenditures. These patterns in the data are also supported by the first column of Appendix Table 11, which shows the results of a fixed effects regression of bid prices on the mean number of bids in the household's market, as well as interactions of mean number of bids with household electricity expenditure and home valuation.

A plausible explanation of the heterogeneous relationship between competition and bid prices in Figure 3c is differences in buyers' price-elasticities. For instance, if households with lower home valuations are more price sensitive, then sellers' optimal markups will be relatively lower if they are bidding against few or no competitors. In particular, if households are more price-sensitive then a higher bid by a monopolist or a duopolist will be more likely to be rejected by the buyer.

If home valuation and price elasticity are indeed negatively related, then the relative welfare effects of the platform are ambiguous. On the one hand, wealthier households, in terms of home valuation, attract more bids (Figure 3a) and larger bid price reductions from added competition (Figure 3c). On the other hand, if wealthy households are more price-inelastic then a relatively larger bid price reduction could still lead to a relatively smaller consumer surplus gain.

3 Model

Motivated by these patterns in the data, I develop a structural model that incorporates heterogeneous buyer preferences and strategic participation and bidding by sellers. In the following sections I describe the estimation procedure and use the estimates to evaluate welfare impacts and counterfactuals.

In the model, each buyer i seeks to procure installation services for a single indivisible project using a multi-attribute auction. Throughout the paper, I use i to refer to both a buyer and their respective project. Buyer i's project is distinguished by its project type $\mathbf{z}_i$, which is characterized by geographic market where the project is located, the time period, and the characteristics of the household. For each project of type $\mathbf{z}$, there is a set $\mathcal{N}(\mathbf{z})$ of potential sellers that choose whether or not to submit a bid for the project.

Each seller j is differentiated by their type which is characterized by a vector $\mathbf{w}_j$. A seller's type could be distinguished by a relatively parsimonious measure such as a star-rating category, a relatively higher dimensional variable such as a unique installer ID (i.e., seller fixed effects), or a combination of variables.

Each seller's type is observable to both the buyer and the other potential sellers. If a seller chooses to participate in the auction for project i they then also select a price bid B_{ij}. Each seller is only permitted a single bid for each project. Sellers' bids are characterized by their price in addition to a vector of non-price characteristics $\mathbf{x}_{ij}$ such as panel quality and inverter type. In contrast to the seller's type $\mathbf{w}_j$, $\mathbf{x}_{ij}$ is allowed to vary across projects for a given seller.

3.1 Demand

The allocation rule in a multi-attribute auction comes from the buyer's choice problem. Let $\mathcal{K}_i \subset \mathcal{N}(\mathbf{z}_i)$ 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. Buyer i's utility from selecting option j is given by:

$$\begin{aligned} u_{ij} = B_{ij}\alpha_i + \mathbf{x}'_{ij}\beta + \mathbf{w}'_j\gamma + \delta_i + \zeta_{ig} + (1-\lambda)\varepsilon_{ij} \\ \alpha_i = \mathbf{z}_i^{(1)\prime}\tilde{\alpha}, \quad \delta_i = \mathbf{z}_i^{(2)\prime}\tilde{\delta}. \end{aligned} \tag{1}$$

Here B_{ij} is the bid price for option j, and α_i is the price sensitivity of buyer i. Buyer price sensitivity, α_i, is a function of of an m-dimensional vector household characteristics denoted $\mathbf{z}^{(1)} = \mathbf{z}_{1:m}$, such as the household's home market valuation. Utility is affected by $\mathbf{x}_{ij}$, the non-price characteristics of the bid, such as the panel brand quality, and the solar inverter type. Utility also depends on attributes of each seller that are fixed across bids, $\mathbf{w}_j$, such as installer fixed effects. The δ_i term is a demand shifter for buyer i that allows utility for all of the "inside options" to vary depending on a p-dimensional vector of household characteristics $\mathbf{z}^{(2)} = \mathbf{z}_{m+1:m+p}$ such as the geographic market and time-period. Notice that the variables determining the project type, $\mathbf{z}$, include both sets of household-level variables in $\mathbf{z}^{(1)}$ and $\mathbf{z}^{(2)}$. Choices are also influenced by ε_{ij}, 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 nested logit model (Cardell, 1997). 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, each buyer has no correlation in preferences for each "inside option", and the model reduces to the standard logit model. As λ goes to one, the random component of buyers' preferences for each "inside option" become perfectly correlated. 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.

In modeling the buyer's choice, I assume each buyer chooses the installation option that delivers the highest utility per unit of capacity. This assumption simplifies the demand model because I can model the buyer's decision as a simple discrete choice instead of a discrete-continuous choice. Thus B_{ij} in utility is the bid price in dollars per watt. Current public policies largely dictate each buyer's optimal system capacity. In particular, net-metering rules allow residential solar customers to sell electricity generated by their rooftop system to their utility at the retail electricity rate, as long as the household's annual generation does not exceed their annual consumption. Any solar generation that exceeds the household's annual consumption is compensated far below the retail rate. As a result, the system capacity that will deliver the largest net present benefit to the buyer is the capacity that equates the expected annual solar generation with the expected annual electricity use.[32]

3.2 Supply

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, 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 competitors that offer bids for a particular project. However, they can observe the total number of bids that were submitted to an auction 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.

I model suppliers bidding behavior as a two-stage process. In the first stage, each potential bidder $j \in \mathcal{N}(\mathbf{z}_i)$ 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. Additionally, they know the mean utility of the buyer (but not the random component of utility).[33] Therefore, each firm can form an expectation about their profits, conditional on the decision to enter the auction.

If seller j decides to enter the auction for the project, they incur a bid preparation cost η_{ij}, where $\eta_{ij} \sim Lognormal(\mu(\mathbf{z}_i, \mathbf{w}_j), \sigma^2(\mathbf{w}_j))$. The expected bid preparation cost depends on both the project type and the seller's type. I assume that bid preparation costs are i.i.d. across projects and firms and is private information of each potential bidder. If a firm decides to enter auction i, then the firm learns the non-price characteristics of their bid $\mathbf{x}_{ij}$, the capacity of the system to be installed q_{ij}, and the marginal cost of completing the project c_{ij}.

To make the model empirically tractable, I assume that the non-price characteristics, $\mathbf{x}_{ij}$,

[32] In practice, each installer can propose a different system capacity when bidding through the platform. In a later section, I will show that the demand estimates are robust to controlling for the proposed system capacity as a non-price bid attribute. I also show that the demand elasticities are relatively similar if I use the discrete-continuous choice utility formulation proposed by Hanemann (1984).

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

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, the assumption also finds support in the data. For example, installers 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.

Analogously to the demand-side, I also assume that the system capacity is not a strategic choice for the sellers. Although each installer provides a recommended system capacity in their bid, there is relatively little variation in the capacity offers made for a particular project. In particular, a regression with project ID dummies can explain over 80% of the total variation in system capacity offers in the data. Installers attempt to choose a system capacity that will generate enough to match the household's annual energy use. As further support for this assumption, the second column of Table 11 provides evidence that installers' capacities are not affected by strategic factors. In particular, I run a regression of proposed system capacity on the mean number of bids across projects in the same market (CBSA-half-year). This regression provides a measure of how seller bids change within a market as the number of expected competitors increases. Column 1 shows that the expected number of bids for a project affects the bid prices; however, Column 2 demonstrates that the expected number of competing bids is not related to proposed system capacities.

When firms make their entry decision, they do not know their marginal cost, non-price characteristics, or the exact capacity of the system. However, they do know the joint distribution from which their marginal cost, non-price characteristics, and system capacity will be drawn, $F_{CXQ|\mathbf{w}_j,\mathbf{z}_i}(c, \mathbf{x}, q|\mathbf{w}_j, \mathbf{z}_i)$. The distribution depends on both the seller's type and the project type. Therefore, sellers will know that buyers with large electric bills will be more likely to need a large system but they will not know the exact capacity of each system before deciding to enter in the auction. After the firms make their entry decisions in stage one, each firm's marginal cost and non-price characteristics are drawn from $F_{CXQ|\mathbf{w}_j,\mathbf{z}_i}$ and the installer then decides on a price bid during the second stage.

Price Bidding

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:

$$\max_{B_{ij}} \quad q_{ij}[B_{ij} - c_{ij}] \cdot \mathcal{P}_{ij}(B_{ij}, \mathbf{x}_{ij}, \mathbf{w}_j|\mathbf{z}_i) \tag{2}$$

Where q_{ij} is the system capacity, B_{ij} is firm j's per-unit price bid, c_{ij} is firm j's marginal cost, and $\mathbf{x}_{ij}$ are the firm's non-price characteristics for project i. $\mathcal{P}_{ij}(B_{ij}, \mathbf{x}_{ij}, \mathbf{w}_j|\mathbf{z}_i)$ 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 having type $\mathbf{w}_j$. The equilibrium expected probability of being selected is a function of the type of project $\mathbf{z}_i$. We work with expected probabilities because the seller does not know exactly which competitors they will face nor the bids of those competitors. We note that the solution to the bid pricing problem is not a function of the system capacity realization, q_{ij} enters the expected profit function multiplicatively and, therefore, does not directly influence the optimal per-unit bid price. However, the system capacity can indirectly affect the price bid if system capacity and marginal cost are correlated.

When formulating firms' expectations, I assume that all sellers make entry decisions simultaneously, and submit their bids simultaneously. 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 are 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 Appendix Table 10, I regress bid price on the order that a bid was submitted, controlling for the total number of bids submitted for the project, installer fixed effects, CBSA fixed effects, and half-year fixed effects. The coefficient 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:

$$\begin{aligned}\mathcal{P}_{ij}(B_{ij}, \mathbf{x}_{ij}, \mathbf{w}_j|\mathbf{z}_i) &= \mathbb{E}[Prob_{ij}(B_{ij}, \mathbf{x}_{ij}, \mathbf{w}_j; \mathbf{B}_{i,-j}, \mathbf{X}_{i,-j}, \mathbf{W}_{-j}|\mathbf{z}_i)] = \\ &\int Prob_{ij}(B_{ij}, \mathbf{x}_{ij}, \mathbf{w}_j; \mathbf{B}_{i,-j}, \mathbf{X}_{i,-j}, \mathbf{W}_{-j}|\mathbf{z}_i) \cdot dG(\mathbf{B}_{i,-j}, \mathbf{X}_{i,-j}, \mathbf{W}_{-j}|\mathbf{z}_i)\end{aligned} \tag{3}$$

Recall that $Prob_{ij}$ is the probability that buyer i selects firm j's bid conditional on realized vector of competing price bids $\mathbf{B}_{i,-j}$, having a stacked vector of non-price characteristics $\mathbf{X}_{i,-j}$, and having types $\mathbf{W}_{-j}$. G represents the joint distribution function of $\mathbf{B}_{i,-j}, \mathbf{X}_{i,-j}$, and $\mathbf{W}_{-j}$ occurring in equilibrium, conditional on the project being of type $\mathbf{z}_i$. Since each firm's entry draw and marginal cost draw is assumed to be i.i.d., we can express dG as the product of the probabilities that each competing firm l decides to enter the auction and then bids B_{il} and has non-price characteristics $\mathbf{x}_{il}$.

I define the optimal bid function as $B^*_{il}(c_{il}|\mathbf{x}_{il}, \mathbf{w}_l, \mathbf{z}_i)$ and $H(\mathbf{w}_l, \mathbf{z}_i)$ as the probability that a potential seller l that is of type $\mathbf{w}_l$ enters an auction of type $\mathbf{z}_i$. Then we have:

$$dG(\mathbf{B}_{i,-j}, \mathbf{X}_{i,-j}, \mathbf{W}_{-\boldsymbol{j}}|\mathbf{z}_i) = \prod_{l \in \mathcal{N}(\mathbf{z}_i)\setminus\{j\}} H(\mathbf{w}_l, \mathbf{z}_i) \cdot dF_{CX|\mathbf{w}_l,\mathbf{z}_i}\big(B^{*-1}(B_{il}|\mathbf{x}_{il}, \mathbf{w}_l, \mathbf{z}_i), \mathbf{x}_{il}|\mathbf{w}_l, \mathbf{z}_i\big) \tag{4}$$

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:

$$(B_{ij} - c_{ij})\frac{\partial \mathcal{P}_{ij}(B_{ij}, \mathbf{x}_{ij}, \mathbf{w}_j|\mathbf{z}_i)}{\partial B_{ij}} + \mathcal{P}_{ij}(B_{ij}, \mathbf{x}_{ij}, \mathbf{w}_j|\mathbf{z}_i) = 0 \tag{5}$$

Given a vector of non-price characteristics, the optimal bid function $B^*_{il}(c_{il}|\mathbf{x}_{il}, \mathbf{w}_l, \mathbf{z}_i)$ is defined implicitly by Equation 5.

Seller Participation

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 η_{ij}. When firm j decides to enter auction i, they only know the project type, their own seller type, and their private entry cost draw. Firm j's expected profits conditional on entering the auction for project i can be expressed as follows:

$$\mathbb{E}[\pi_{ij}|\mathbf{w}_j, \mathbf{z}_i] = \int \Big[q_{ij} \cdot (B^*_{ij}(c_{ij}|\mathbf{x}_{ij}, \mathbf{w}_j, \mathbf{z}_i) - c_{ij}) \cdot \mathcal{P}_{ij}(B^*_{ij}, \mathbf{x}_{ij}, \mathbf{w}_j|\mathbf{z}_i)\Big] dF_{CXQ|\mathbf{w}_j,\mathbf{z}_i}(c_{ij}, \mathbf{x}_{ij}, q_{ij}|\mathbf{w}_j, \mathbf{z}_i) \tag{6}$$

Recall that $F_{CXQ|\mathbf{w}_j,\mathbf{z}_i}(c, \mathbf{x}, q|\mathbf{w}_j, \mathbf{z}_i)$ is the joint distribution of non-price characteristics, marginal costs, and system capacity whose realization is not known to the firm at the time of entry. Therefore, the firm will enter the auction as long as:

$$\mathbb{E}[\pi_{ij}|\mathbf{w}_j, \mathbf{z}_i] \geq \eta_{ij} \tag{7}$$

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

$$H(\mathbf{w}_j, \mathbf{z}_i) = \Phi\left(\frac{ln\Big(\mathbb{E}[\pi_{ij}|\mathbf{w}_j, \mathbf{z}_i]\Big) - \mu(\mathbf{z}_i, \mathbf{w}_j)}{\sigma(\mathbf{w}_j)}\right) \tag{8}$$

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 $\mathbf{z}_i$ to all potential entrants $\mathcal{N}(\mathbf{z}_i)$.

2. Each potential seller $j \in \mathcal{N}(\mathbf{z}_i)$ receives a private entry cost η_{ij}. Each potential entrant then compares their entry cost η_{ij} to the expected marginal profit conditional on entering the auction $\mathbb{E}[\pi_{ij}|\mathbf{z}_i, \mathbf{w}_j]$. 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}, learns their non-price characteristics $\mathbf{x}_{ij}$, and learns the exact capacity of the system q_{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 $\mathbf{w} \times \mathbf{z} \times \mathbb{R}_+ \rightarrow \{0, 1\}$, and a bidding strategy $\mathbf{w} \times \mathbf{z} \times \mathbf{x} \times \mathbb{R}_+ \rightarrow \mathbb{R}_+$. Specifically, sellers use information about the project type, their seller type, 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 seller type, 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., 2019). That is, all sellers of the same type use the same participation strategy in equilibrium, and all sellers of the same type 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 Inequality 7, given the strategies of other firms. An equilibrium in the bidding stage requires that all firms satisfy Equation 5 given the other installer's strategies. Krasnokutskaya et al. (2019) 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. The next section describes the estimation procedure in detail.

4 Estimation

I estimate the structural parameters in three steps. First, I solve for the demand parameters via maximum likelihood. Second, I use the estimated demand parameters to simulate firms' 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 bidders' expected marginal profit from entering each auction and estimate the entry cost parameters using observed entry decisions. I discuss the details of each step in the following subsections.

Demand Estimation

From Equation 1 we can see that there are four sets of variables that affect a buyer's utility: (1) B, the bid price (2) $\mathbf{z}^{(1)}$, household demographics that determine price-sensitivity (3) $\mathbf{z}^{(2)}$, variables that shift the buyer's preference for all of the installation bids (i.e., shifts the likelihood of picking the outside option), and (4) $\mathbf{x}$, variables that characterize the non-price attributes of each participating installer's bid. I choose the baseline utility specification to produce estimates that are robust while also keeping the non-linear optimization computationally tractable. Every household has a different choice set so the choice probabilities cannot be aggregated to the market share level as in Berry et al. (1995). Thus, I estimate the demand parameters via maximum likelihood based off the individual choice data, which involves a non-linear search over all the utility parameters, as opposed to only the non-linear parameters as in Berry et al. (1995). In my baseline demand model which controls for CBSA fixed effects as well installer fixed effects (for permanent installers), there are around 100 utility parameters to estimate. Therefore, the nested logit utility specification, which has a closed-form expression for choice probabilities, is more computationally tractable in this context relative to a random coefficients utility model.

The price that enters buyers' utility for option j is measured in \$/watt and is scaled to 70% of each installer's gross bid price to account for the 30% Investment Tax Credit (ITC). I allow for price sensitivity to vary across households with different home market valuations. In particular, I separate the sample into quartiles based on home market valuation and allow for each quartile to have different price coefficients. I use home market valuation to proxy for wealth for a few reasons. First, home valuation is likely to be observable to the installers when placing bids. On the other hand, household income is unlikely observable to the sellers. Second, Caceres (2019) shows that housing wealth is more closely correlated with total wealth than annual income among a large sample of U.S. households. Finally, there is more variation in the home market valuation due to the greater number of binned categories reported in the data.[34]

In the vector of demand shifters, $\mathbf{z}^{(2)}$, I include fixed effects for each CBSA, fixed effects for each half year of the sample, fixed effects for each quartile of household monthly electricity expenditure, and a dummy variable for whether the household obtained off-platform quotes.[35]

[34]The data report the household annual incomes in 9 bins ranging from "\$ 0-14k" to " above \$125k". The data identify the home market valuation of each home in 17 bins ranging from "below \$50k" to "above \$ 2 million". Over 40% of the sample have annual household incomes above \$100k. Therefore, a large share of them fall within the highest reported income category. On the other hand, less than 4% of the sample is in the highest home market valuation category.

[35]Projects that are outside a CBSA or in a CBSA with fewer than 100 total projects are placed into a new category by state. For example, a project located in Aspen, Colorado would be defined as "Other, CO"

These coefficients allow for differences in willingness-to-pay across households located in CBSAs with different expected solar production, for example.

Finally, the utility for each option is also a function of several non-price characteristics: fixed effects for premium and premium plus module categories, a microinverter fixed effect, fixed effects for the installer's star rating category, and a set of fixed effects that measure the installer's installation experience. I also allow for additional heterogeneity in seller quality by including "permanent" installer fixed effects for each installer that placed over 300 total bids through the platform during the sample. These permanent sellers account for over 60% of bids in the sample. In later sections, I discuss the robustness of the results to changes in the utility specification.

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 5. Notice that the FOC does not have a closed form since it contains two expectations $\frac{\partial \mathcal{P}_{ij}(B_{ij},\mathbf{x}_{ij},\mathbf{w}_j|\mathbf{z}_i)}{\partial B_{ij}}$ and $\mathcal{P}_{ij}(B_{ij}, \mathbf{x}_{ij}, \mathbf{w}_j|\mathbf{z}_i)$. 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 evaluation the FOC for each bid, I follow the following procedure:

1. First, obtain non-parametric estimates of the entry probabilities for each project-seller type pair, $\{\mathbf{z}_i, \mathbf{w}_j\}$. This estimate is just the ratio of auctions entered divided by total auctions of that project-seller type. I assume a seller is a potential entrant for auction i if they entered at least one auction of type $\mathbf{z}_i$.

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

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 $\mathbf{w}_j$ seller enters a simulated auction of type $\mathbf{z}_i$; then randomly draw a bid (both bid price and non-price characteristics together) from the pool of all bids placed by type $\mathbf{w}_j$ sellers in auctions of type $\mathbf{z}_i$.

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 competitor's observed characteristics.

5. Repeat the second through fourth step S times each[36] 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:

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

$$\widehat{\mathcal{P}_{ij}} = \frac{1}{S}\sum_{s=1}^{S} Prob_{ij}^{s}, \qquad \widehat{\frac{\partial \mathcal{P}_{ij}}{\partial B_{ij}}} = \frac{1}{S}\sum_{s=1}^{S} \frac{\partial Prob_{ij}^{s}}{\partial B_{ij}} \tag{9}$$

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:

$$\widehat{c_{ij}} = B_{ij} + \frac{\widehat{\mathcal{P}_{ij}}}{\widehat{\frac{\partial \mathcal{P}_{ij}}{\partial B_{ij}}}} \tag{10}$$

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.

The choice of project categories is critical 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, an installation for a household in Denver with high energy use will be different ex-ante than an installation for a relatively lower energy using household in Los Angeles because of differences in labor costs, permitting requirements, and differences in the set of possible competing bidders. In addition, these households will need different sized solar arrays so the price per watt will vary due to within-project economies of scale. Economies of scale arise because some parts of the installation cost are fixed regardless of the system capacity such permitting or inspection costs. Also, if labor costs do not scale linearly with system capacity then we would expect to see decreasing unit costs as system capacity increases. Similarly, a Denver project in 2015 will be different from a Los Angeles 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 Los Angeles 2015H1 when simulating a Denver 2016H2 auction. However, if I define too many project categories, (i.e., each CBSA-week has its own category), then markup estimates for each bid will have higher variance because there will be only a handful of projects to use to simulate realizations of each auction. In my main specification, I employ a relatively high dimensional project type definition to avoid bias. The project types are determined by five variables: the CBSA in which the project is located, the half-year when the project was originated, home market value quartile, the quartile of the household's home electricity expenditure, and an indicator for whether the buyer has off-platform quotes. In a later section, I discuss the robustness of the results to changes in the definition of project types.

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 $q_{ij} \cdot (B_{ij} - \widehat{c_{ij}}) \cdot Prob_{ij}$. To calculate a seller's *pre-entry* expected profit $\widehat{\mathbb{E}[\pi_{ij}]}$ from entering an auction i, I take the average of all the realized *post-entry* expected profits for that seller's type $\mathbf{w}_j$ for projects of type $\mathbf{z}_i$. More precisely, the *pre-entry* expected profit is estimated as:

$$\widehat{\mathbb{E}[\pi_{ij}]} = \frac{1}{N(\mathbf{z}_i, \mathbf{w}_j)} \sum_{i \in \mathbf{z}_i} \sum_{j \in \mathbf{w}_j} q_{ij} \cdot (B_{ij} - \widehat{c_{ij}}) \cdot Prob_{ij}. \tag{11}$$

Where $N(\mathbf{z}_i, \mathbf{w}_j)$ is the total number of bids placed by type $\mathbf{w}_j$ sellers in auctions of type $\mathbf{z}_i$. I use $\sum_{i \in \mathbf{z}_i}$ to indicate the sum over all auctions of project type $\mathbf{z}_i$ and $\sum_{j \in \mathbf{w}_j}$ to indicate the sum over all bids submitted by sellers of type $\mathbf{w}_j$.

Next, I use the *pre-entry* expected profits $\widehat{\mathbb{E}[\pi_{ij}]}$ to maximize the following pseudo log likelihood function:

$$EntryPseudoLL(\mu, \sigma) = \sum_{i}^{M} \sum_{j \in \mathcal{N}(\mathbf{z}_i)} \left\{ \mathbb{1}[j\ enters\ i] \cdot ln\left(\Phi\left(\frac{ln\left(\widehat{\mathbb{E}[\pi_{ij}]}\right) - \mu(\mathbf{z}_i, \mathbf{w}_j)}{\sigma(\mathbf{w}_j)} \right) \right) \right. \tag{12}$$

$$\left. + \Big(1 - \mathbb{1}[j\ enters\ i]\Big) \cdot ln\left(1 - \Phi\left(\frac{ln\left(\widehat{\mathbb{E}[\pi_{ij}]}\right) - \mu(\mathbf{z}_i, \mathbf{w}_j)}{\sigma(\mathbf{w}_j)} \right) \right) \right\}$$

Where $\mathbb{1}[j\ enters\ i]$ is an indicator function that equals one if seller j enters auction i and is zero otherwise. I specify that μ is a linear function of a CBSA fixed effects, half-year fixed effects, installer star-rating fixed effects, installer experience fixed effects, and individual installer fixed effects for permanent sellers. I also allow σ to vary across installer star rating categories.

4.1 Identification and Modeling Assumptions

To identify the buyer utility parameters, we require that all variables entering the utility function, including bid price, are not correlated with the preference shocks ε_{ij} and ζ_{ig}. Also, we need the number of bids made to each buyer to be uncorrelated with the unobserved preference shocks.

My identification strategy exploits the rich household-level demographic information obtained in the marketing-research data. Namely, I employ a utility specification that includes flexible fixed effects to control for household heterogeneity and a rich set of fixed effects to control for bid quality. I include panel brand quality controls, inverter type controls, installer star-rating dummies, installer experience controls, and installer fixed effects for permanent sellers to control for quality differences across bids. To control for buyer heterogeneity, I include electricity bill quartile dummies, a control for whether the household has off-platform bids, CBSA fixed effects, and half year fixed effects.

Identification of the price coefficients for each home value quartile, α_i, comes from variation in the prices quoted adjusting for seller quality, hardware type, common demand shocks across time, demand shocks across CBSAs, and demand shocks to specific household demographic groups (determined by the set of household-level controls such as electric bill quartiles). Intuitively, the price variation that identifies α_i comes from differences in marginal cost conditional on buyer type and seller type—such as price changes driven by input cost shocks. Despite the relatively rich set of buyer and seller controls, there still could be omitted variables correlated with both prices and the unobserved preference shocks. In Section 4.2, I discuss these concerns in more detail and present a series of robustness checks.

The λ parameter is identified by exploiting variation in the number of bids that households receive after controlling for the CBSA, the time-period, home energy expenditure, home market valuation, and whether the household has off-platform quotes. For example, in 2015 H1, some buyers in Hartford, Connecticut, receive two bids while others may receive three or more. To consistently estimate λ, the variation in the number of bids should be coming from supply-side factors. For example, say Household A solicits bids the week after Household B and receives one fewer bid because one of the suppliers is now busy installing other rooftop systems that week. On the other hand, the estimate of λ will be biased if variation in bids is driven endogenously by demand-side factors. For instance, if some buyers are ex-ante more likely to buy solar (i.e., more educated households), and therefore more suppliers choose to bid on their project. If this is the case, it will appear that the additional bids are causing more solar purchases, but in fact, the likelihood of buying solar is causing an increased number of bids. In Section 4.2, I outline and test a series of confounding factors that could correlate with both unobserved buyer tastes and bid quantity.

After the demand parameters are identified, understanding the identification of marginal costs is straightforward. With demand parameters in hand, we can compute firms' optimal markups. Then using the markups, we can employ Equation 10 to create a one-to-one mapping between bid prices and costs. The entry cost parameters are identified by variation in expected marginal profits, holding seller type constant.[37] In particular, μ and σ are pinned down by the extent the probability of entry of type $\mathbf{w}_j$ sellers changes as expected marginal profits increase. In theory, we could trace out the entire entry cost distribution for each seller type non-parametrically if we observed the entry probability at every possible level of expected marginal profit. To generate variation in expected marginal profit across auctions, we need ex-

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

ogenous 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 the household-level demographics—home energy expenditure, home market valuation, and whether the household has off-platform bids—do not affect the firms' entry costs. In contrast, the household-level demographics do affect firms' expected marginal profits. For this assumption to hold, it must be true that bid preparation time and effort is not different depending on household characteristics within a market. Figure 12 provides a useful visualization of the variation that identifies the entry cost parameters. For a single market, the figure shows that there is substantial variation in the expected marginal profit due to variation in the household demographics. We see that, sellers are more likely to enter projects with higher expected marginal profits.

Having discussed the identification of the model, it is also important to acknowledge a potential limitation of my modeling approach: Namely, I employ a static model of demand. While there is a large literature that has used static discrete choice models to estimate demand for durable goods (i.e, Berry et al. (1995)), several related papers have implemented dynamic demand models of the solar PV market (Bollinger and Gillingham, 2019; Burr, 2014; De Groote and Verboven, 2019; Feger et al., 2017; Langer and Lemoine, 2017; Reddix II, 2015; Snashall-Woodhams, 2019). My static demand estimates will be biased if consumers are forward-looking.

There are a few reasons why a static model may be more appropriate in my setting relative to other studies in the literature. First, several studies (Bollinger and Gillingham, 2019; Burr, 2014; De Groote and Verboven, 2019; Langer and Lemoine, 2017; Reddix II, 2015) estimate models using data from a period when solar subsidies were decreasing substantially over time, making dynamic incentives especially crucial in the solar installation decision. In contrast, solar subsidies in most US states did not change substantially during my sample period and were not scheduled to change in the near future. Two exceptions are Connecticut and New York, both which offered solar incentives that changed during the sample.[38] Therefore, I test the robustness of the demand estimates to the exclusion of those states from the analysis. Finally, the majority of existing papers in the literature use data from before 2013, when pre-incentive solar installation prices were declining rapidly over time. Figure 13 in the appendix shows that the median U.S. solar installation price was declining at around 15% per year in the early 2010s. However, installation prices later stabilized, and by the start of my sample, the median installation price fell by only 2% between 2014 and 2015. The relatively stable installation price trend suggests that the households in my sample have relatively low option value in delaying their solar investment to get a lower price next year.[39]

In summary, the absence of declining subsidies, together with relatively stable installation

[38]In Connecticut, the expected performance based buy-down program led to changes in rebate levels for solar installations during the sample period. Similarly, the NY-Sun Incentive featured changing incentive levels over time.

[39]De Groote and Verboven (2019) estimate that consumers use a relatively high discount rate when making solar investments (real implicit interest rate around 15%), suggesting that consumers are unlikely to delay an investment this year to get a price reduction of less than 5% next year.

prices, indicates that a static model of demand is likely to deliver a more reasonable approximation of consumer behavior in my setting. Nonetheless, I make several efforts to mitigate the possibility of bias in the demand parameter estimation. For one, I avoid using variation in prices over time by including time fixed effects in utility; the time fixed effects allow the value of the outside option to change over time. Hence, if consumers expected prices to drop in 2016, the value of the outside option for consumers shopping in 2015 can shift to account for those beliefs. Therefore, time fixed effects provide a reduced-form way of controlling for potential forward-looking behavior in the static model. With this caveat in mind, I now transition to discussing the model estimates in the next subsection.

4.2 Estimation Results

Demand

The top portion of Table 3 contains the estimates for the baseline utility specification. The estimate of the nesting correlation parameter, λ, is 0.37. The nesting parameter is greater than zero, so we can reject the standard logit model. However, the nesting parameter is also lower than one which indicates that there is imperfect correlation in the preference shocks. Consequently, an increase in the number of bids will increase the likelihood of a household adopting solar through improving buyer-specific match value.

As expected, the base price coefficient ($/watt) is negative and relatively large in magnitude. We also see that households in the top three home value quartiles are relatively less price-sensitive, although the interaction terms are not quite statistically significant (p-values of 0.14, 0.15, and 0.14, respectively). The middle portion of the table shows the mean own-price elasticity across all bids made to households in each home value quartile; we see that bottom quartile is the most elastic (mean own-price elasticity $= -1.68$) despite receiving the lowest number of bids on average.

The parameters on the electricity bill quartile dummies reveal that households with higher monthly electricity expenditure obtain a lower utility (per unit capacity) from a solar installation. However, households with larger electricity expenditure typically install higher capacity systems. We will see in Section 5 that the consumers in the highest electric bill quartiles do obtain higher total consumer surplus despite attaining lower consumer surplus per watt.

We also see that the coefficient associated with whether the household has off-platform quotes is also positive. This suggests that although households with offline quotes have more outside options, these households also appear to have higher willingness-to-pay for a solar installation relative to households that did not engage in any offline search.

In the top right side of the table we see that hardware characteristics are important to households, and that households prefer higher-rated panel brands and systems with microinverters, holding prices fixed. We also see that buyers value installer characteristics. All else equal, sellers with a rating under five stars are less likely to be selected relative to an installer with-

Table 3: Demand Estimates

Nesting Parameter		**Non-Price Bid Attributes**	
λ	0.372 (0.061)	Premium Panel	0.562 (0.078)
Price Coefficients		Premium Plus Panel	1.429 (0.152)
Constant	-0.717 (0.109)	Microinverter	0.370 (0.083)
$\times$ Home Mkt. Value - Quartile 2	0.068 (0.045)	**Installer Attributes**	
$\times$ Home Mkt. Value - Quartile 3	0.079 (0.053)	Star Rating ≤ 4	-0.260 (0.261)
$\times$ Home Mkt. Value - Quartile 4	0.089 (0.060)	Star Rating = 4.5	-0.410 (0.291)
Household Attributes ($\times$ **Inside Option**)		Star Rating = 5	0.281 (0.109)
Electric Bill - Quartile 2	-0.067 (0.102)	Installs Completed: 100-1000	0.697 (0.251)
Electric Bill - Quartile 3	-0.241 (0.106)	Installs Completed: >1000	0.749 (0.257)
Electric Bill - Quartile 4	-0.584 (0.117)		
Has Off-Platform Quotes	0.368 (0.096)		

Home Mkt. Value Quartile	Mean Own-Price Elasticity	Avg Number of Bids
1	-1.68	3.1
2	-1.56	3.47
3	-1.57	3.84
4	-1.57	3.86

Fixed Effects		Log Likelihood
CBSA Fixed Effects	Yes	-3823.542
Half-Year Fixed Effects	Yes	
Permanent Installer Fixed Effects	Yes	

Notes: The utility specifications CBSA, half-year, and permanent seller fixed effects. Permanent sellers are those that submitted over 300 total bids. All of the project attributes, non-price attributes, and installer attributes are dummy variables. The star rating coefficients should be interpreted relative to installers with no rating. The mean own-price elasticity are calculated based of the realized choice sets and do not account for ex-ante uncertainty in seller participation. Standard errors are in parenthesis.

out ratings. On the other hand, installers with a five-star rating are preferred over non-rated sellers. We also see that an installer's residential installation experience is highly valued by buyers; buyers are willing to pay over 20% more for an installer that has completed over 1000 installations compared to an installer that has completed less than 100 installations.

Demand Robustness Checks

As discussed in Section 4.1, there are several threats to the identification of the utility parameters. The price coefficients will be biased if bid prices are correlated with unobserved installer quality, unobserved hardware quality, or unobserved buyer characteristics. The λ parameter could also be biased if the number of bids is correlated with unobserved buyer characteristics. Table 12 investigates the robustness of the key coefficient estimates to adding a series of potential confounding variables into the utility specification.

The baseline specification controls for star rating and installer experience, as well as seller fixed effects for permanent sellers. A potential concern is that the star rating, experience, and even the seller fixed effects are imperfect controls if quality varies over time. For instance, star ratings can change as installers bid more and complete more projects through the platform.

Column 2 of Table 12 shows that both the price and nesting coefficients are robust to adding controls for the amount of time each installer has been active on the platform. In Column 3, I consider a utility specification with a full set of panel brand fixed effects in lieu of the more parsimonious controls for "premium" and "premium plus" rated panels. The key coefficients remain nearly unchanged after adding these additional hardware brand controls. In Column 4, I include additional household demographic variables that could be correlated with both prices and willingness-to-pay for solar. Namely, I include a control for whether the household head has a college degree, fixed effects for the race of the household head, and a dummy for if the household head is over 65 years old. The price coefficients are very similar after adding these demographic controls. Finally, the bid prices and number of bids could both be correlated with factors that influence the viability of the specific rooftop. The age of the house itself, as well as the size of the house may affect the difficulty and cost of an installation, so I include a set of controls for the physical size of the house (square feet) and the age of the house (years since built). In particular, I add fixed effects for each home size quartile and each house age quartile. The parameters and the implied price elasticities are again robust to these changes. For transparency, Table 14 also shows how the estimates change if some of the installer attributes and household demographic variables are omitted from the utility specification.

It is conceivable that prices are correlated with other unobserved household characteristics that are also correlated with preferences; although these characteristics would need to be observable by the sellers through the platform but at the same time not collected and reported in Axciom InfoBase's household marketing research database.

Another key assumption of the demand model is that the buyer's choice can be expressed as a function of utility per unit capacity. Effectively, this assumption abstracts away from the buyer's system capacity choice. This assumption could be problematic if buyers' choice is influenced by the proposed system capacity that each installer offers. For example, if buyers prefer bids for smaller systems and system capacity is correlated with price per watt due to installation economies of scale, then the price coefficients will be biased. In the second column of Table 13, I add an explicit control for each installer's proposed system capacity as a non-price attribute in utility. The price coefficients are almost invariant to this change, suggesting the system capacity assumption is not consequential for the demand estimates.[40] In the third column, I consider an alternate functional form for buyers' utility function. Hanemann (1984) shows that if a buyer makes a discrete choice followed by a continuous choice, then the utility can be written as a function of the logarithm of price. The intuition is that the buyer should only care about the ratio of price to quality when choosing a product. Taking a log transformation leads to a linear utility function with a logged price term. I find the log-price specification leads to demand estimates that are slightly more elastic than the baseline model. However, I also find that the baseline model with linear prices better fits the data than the log-price model.

[40]In other specifications not shown here, I find that adding the system capacity variable to utility has a large effect on the price coefficient if I also omit controls for household electricity expenditure.

Figure 4: Marginal Cost and Markup Distributions

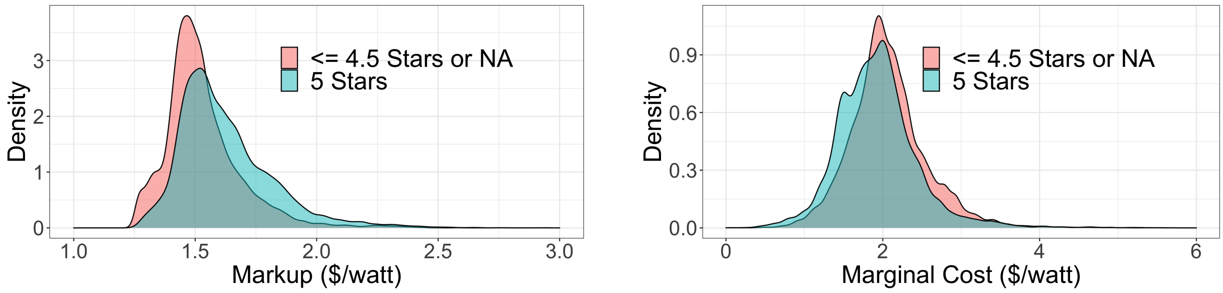

(a) Markup by Seller Rating

(b) MC by Seller Rating

Notes: Kernel densities of the model implied marginal costs and markups.

As a final robustness check, I re-estimate the demand model after removing auctions located in New York and Connecticut, the two states that offered time-varying subsidies during the sample period. Any state or municipal subsidies that are fixed over time are controlled for with the CBSA fixed effects in the baseline utility specification, however, time-varying subsidies could bias the estimated price elasticities if consumers are forward-looking. Table 15 shows that removing Connecticut and New York does not substantially change the demand estimates. Removing these two states leads to slightly higher implied elasticities and a correspondingly lower mean markup of $1.45/watt compared to $1.59/watt estimated from the full sample.

Cost Estimates

With the demand estimates in hand, I solve for markups and marginal cost using the firms' first-order-condition in Equation 5. Figure 4 shows the distribution of both markups and marginal costs across across sellers with differing ratings. Figures 4a and 4b illustrate that five-star rated sellers are more likely to charge higher markups, however these high-rated sellers are also more likely to have lower costs. The figures also illustrate substantial heterogeneity in both costs and markups across projects, and thus it is important to control for both seller-level and household-level heterogeneity.

To further investigate which variables are linked to higher costs and to higher markups, I run regressions with both marginal costs and markups as dependent variables, including all of the variables in buyers' utility (besides price) as regressors. Table 4 shows that the average installation costs are lower for households in the top 75% of home valuation, but these more expensive homes are also subject to much higher markups. We also see that households' electricity expenditure is both negatively correlated with marginal costs and markups. Additionally, the regressions show that higher quality hardware (premium panels and microinverters) increases marginal costs but also earn higher margins. Consistent with Figure 4, we also note that five-star sellers charge higher markups, but their prices are actually lower after accounting for their lower marginal costs. Sellers with more experience charge higher markups but do not have substantially lower costs after controlling for seller rating.

Table 4: Marginal Cost and Markup Regressions

	Dependent variable:		
	MC ($/Watt)	Markup ($/Watt)	Gross Price ($/Watt)
Household Attributes			
Home Mkt. Value - Quartile 2	−0.130 (0.006)	0.140 (0.002)	0.010 (0.006)
Home Mkt. Value - Quartile 3	−0.129 (0.008)	0.159 (0.002)	0.031 (0.007)
Home Mkt. Value - Quartile 4	−0.151 (0.009)	0.191 (0.002)	0.040 (0.008)
Electric Bill - Quartile 2	−0.037 (0.006)	−0.037 (0.002)	−0.074 (0.005)
Electric Bill - Quartile 3	−0.063 (0.006)	−0.049 (0.002)	−0.112 (0.005)
Electric Bill - Quartile 4	−0.075 (0.006)	−0.068 (0.002)	−0.143 (0.005)
Has Off-Platform Quotes	0.0001 (0.005)	−0.001 (0.001)	−0.001 (0.005)
Non-Price Bid Attributes			
Premium Panel	0.066 (0.005)	0.084 (0.001)	0.150 (0.005)
Premium Plus Panel	0.744 (0.011)	0.256 (0.003)	1.000 (0.010)
Microinverter	0.086 (0.006)	0.053 (0.002)	0.139 (0.005)
Installer Attributes			
Star Rating $\leq$ 4	0.017 (0.022)	−0.068 (0.006)	−0.051 (0.020)
Star Rating = 4.5	−0.034 (0.016)	−0.061 (0.004)	−0.095 (0.014)
Star Rating = 5	−0.144 (0.008)	0.066 (0.002)	−0.077 (0.007)
Installs Completed: 100-1000	−0.035 (0.012)	0.083 (0.003)	0.048 (0.011)
Installs Completed: >1000	0.025 (0.013)	0.095 (0.004)	0.120 (0.012)
Permanent Installer FE	Yes	Yes	Yes
CBSA FE	Yes	Yes	Yes
Half-Year FE	Yes	Yes	Yes
Observations	37,080	37,080	37,080
R^2	0.402	0.676	0.494

Notes: Robust standard errors in parentheses. The dependent variables are the model implied marginal cost ($/watt), the model implied markup ($/watt), the observed bid price before tax credits ($/watt). All variables that enter consumer utility are included as covariates including the Home Market Value Quartile dummies which are interacted with price in consumer utility.

Table 5 shows how average markups and marginal costs progressed throughout the sample period. Between the start of 2014 and the end of 2016, marginal costs fell by $0.31 per watt (over $2,200 on an average-sized system), and markups fell by $0.21 per watt (over $1,500 on an average-sized system). Costs and margins fell proportionally, and therefore the (before-subsidy) Lerner index—calculated using all bids—remained relatively stable at about 0.45.

Figure 14a demonstrates that the estimated marginal cost reductions are mostly explained by the fall in wholesale solar PV hardware prices reported by Bloomberg Inc. Figure 14a also indicates that hardware costs account for about half of the total marginal cost. In particular, the Bloomberg hardware cost index, together with the marginal cost estimates, imply that in 2016 H1 hardware costs made up about 28% of the solar installation prices (before subsidies), non-hardware costs such as installation labor and permitting made up 28%, and markups made up 44%. Figure 14b shows how the marginal cost estimates compare to stated costs reported by publicly-traded solar installation firms; the mean estimated marginal cost was slightly higher

Table 5: Summary of Cost Estimates

Panel A: Mean Marginal Cost, Markups, Entry Costs

Half Year	# of Bids	MC ($/watt)	Markup ($/watt)	Lerner Index	Bid Prep Cost ($)
2014 H1	2.07	2.19	1.76	0.46	39.61
2014 H2	2.34	2.14	1.74	0.46	62.39
2015 H1	2.54	2.15	1.69	0.45	45.22
2015 H2	2.72	2.1	1.68	0.45	33.66
2016 H1	3.98	2.03	1.56	0.44	16.04
2016 H2	4.03	1.88	1.55	0.46	9.84

Notes: The table shows the mean number of bids, mean marginal cost, mean markup , mean Lerner Index (markup/ gross price) and the expected bid preparation cost, conditional on submitting a bid for each half-year of the sample.

than the cost reported by SolarCity/Tesla but slightly lower than the costs reported by Sunrun and Vivint. Finally, we compare the estimated markups to other available estimates in the literature. Bollinger and Gillingham (2019) estimate static markups of about $1.20/watt to $1.45/watt using California data from before 2012, whereas the mean markup estimates in the current paper are slightly larger ranging from $1.55-$1.76/watt across 2014-2016.[41]

I use the entry cost parameters, reported in Appendix Table 16, to calculate the expected bid preparation cost (entry cost) conditional on participating. Table 5 shows that bid preparation costs fell from $40 in 2014 H1 to under $10 in 2016 H2. Figure 5 shows that the bid preparation also becomes less-variable over time, with the majority of bid costs in the 0-$20 range by the second half of the sample. Both Figure 5b and Table 16 also show disparity in entry costs across sellers. Sellers with five-star ratings have higher bid preparation costs of $20.97 compared to $13.28 for sellers with a 4.5-star rating. Higher bid preparation costs for higher quality sellers are consistent with higher opportunity costs. For instance, higher-rated sellers may have more project leads off of the platform.

5 Welfare and Counterfactuals

In this section, I use the model estimates to evaluate the consumer welfare effects of the platform. I then investigate the mechanisms driving the welfare results by simulating counterfactuals that measure the causal effect of increasing the number of bids per project. Finally, I solve a second set of counterfactuals to evaluate the welfare effects of government solar subsidies.

[41] One potential reason for the difference in estimated markups in Bollinger and Gillingham (2019), is that the authors assume that buyers are perfectly informed about all sellers' prices in their local market, if buyer's are imperfectly informed about prices, the estimated demand elasticities will be biased towards zero, thus leading to lower estimated markups. The difference in markup estimates could also be due to differences in the studies' sample compositions or changes in market structure over time.

Figure 5: Expected Entry Cost Distributions Conditional On Bidding

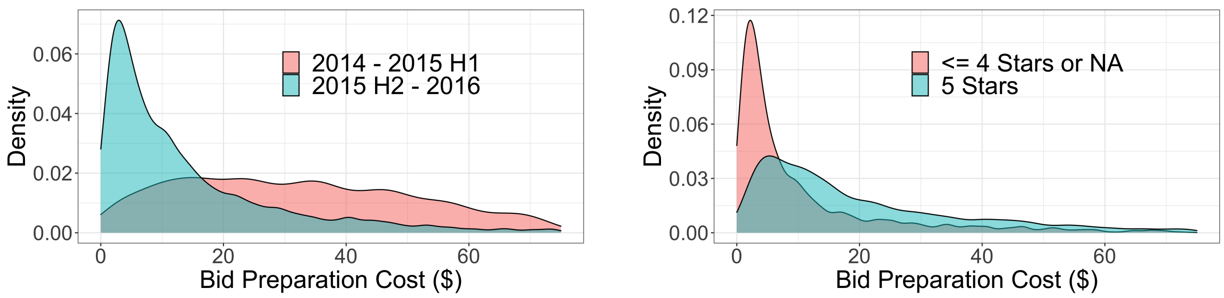

(a) Entry Cost Over Time

(b) Entry Cost by Seller Rating

Notes: 5b shows density plots of the entry cost distribution across seller ratings. 5a shows entry cost density plots for the first half vs the last half of the sample. The expected cost distributions are conditional on the seller deciding to enter a bid.

5.1 Consumer Surplus Gains From Platform Access

In the model, each buyer chooses between a set of installation bids on the platform and an unspecified outside option. The outside option could include either not buying a solar system or buying from an offline installer, therefore the expected consumer surplus associated with each buyers' platform choice set can be interpreted as households' consumer surplus gain from the platform. Recall that consumer utility is measured in dollars per watt, so I need to scale each household's utility to obtain total consumer surplus. To calculate total consumer surplus, I use the standard log-sum formula to calculate consumer welfare per unit of capacity and then I scale each household's per-unit utility by the mean system capacity that was offered to that household.

Table 6: Consumer Surplus Gains from Access to the Platform

Panel A: Consumer Surplus Summary - Full Sample

Statistic	Mean	St. Dev.	Pctl(25)	Pctl(75)
Total Consumer Surplus Per Household ($)	1,451.24	1,319.34	590.31	1,855.41
Consumer Surplus Per Unit ($/watt)	0.20	0.16	0.09	0.27
(Consumer Surplus)/(Mean Bid Price)	0.08	0.06	0.04	0.11

Panel B: Mean CS by Elec. Bill Quartile

Quartile	Bids	Total CS ($)	CS ($/watt)
1	3.12	1077.45	0.21
2	3.54	1542.51	0.24
3	3.64	1663.69	0.21
4	3.85	1535.79	0.16

Panel C: Mean CS by Home Value Quartile

Quartile	Bids	Total CS ($)	CS ($/watt)
1	3.1	1155.85	0.15
2	3.47	1553.43	0.21
3	3.84	1621.06	0.24
4	3.86	1476.33	0.22

Notes: Each column in Panel B and Panel C reports the means value of the variable for each quartile.

Table 6 Panel A shows that the platform increased consumer surplus by $1,451 for the average household. For reference, this welfare gain is equivalent to making a payment of 8% of the mean bid price to each individual that used the platform. Noteably, there is substantial heterogeneity in benefits across households. The top forth of households attain over $1,800 in welfare gains, whereas the bottom 25% experienced less than $600 in gains. Panel B and Panel C demonstrate how consumer welfare varies across electricity expenditure quartiles and home valuation quartiles, respectively. Households in the top electricity expenditure quartile realize smaller per-unit surplus gains relative to households in the lowest quartile. However, these households purchase bigger systems and consequently attain larger total welfare gains. More specifically, households in the bottom electricity bill quartile reap $1,077 in benefits from the platform relative to $1,543, $1,664, and $1,536 for the top three quartiles respectively. A similar pattern emerges in the distribution of gains across home values. The lowest 25% of homes see gains of $1,156 relative to $1,476 for the most expensive forth of homes.

There is also substantial variation in consumer benefits across CBSAs. According to Table 18, the top ten CBSA are all located in the Northeast such as: Norwich-New-London, CT; Worcester, MA-CT; and Cambridge-Newton-Farmington, MA. On the other hand, CBSAs with the lowest welfare gains are located mostly in Texas, Arizona, and non-coastal California including: Fort-Worth-Arlington, TX; Fresno, CA; and Austin-Round-Rock, TX. Table 17 breaks down marginal costs by CBSA and indicates that CBSAs with large consumer surplus gains are also more likely to have high installation marginal costs, whereas CBSAs with smaller consumer gains are more likely to have lower installation costs.

These results indicate that the platform facilitated substantial gains for most consumers but that wealthier households appear to be the largest beneficiaries. This distributional result was not clear ex-ante because on the one hand, wealthier households are less price-sensitive (Table 3) and therefore stand to benefit relatively less from a price reduction. On the other hand, wealthier households attract more bids, see bigger price reductions, and gain access to more distinct installers. The distribution of welfare gains in Table 6 Panel B and C, therefore suggests a high value of obtaining additional installer choices.

These findings motivate several additional questions: First, how much of the platform welfare benefit can be explained by the buyers obtaining additional bids? Relatedly, how much does increasing the number of bids per project reduce solar installation prices? How much does solar adoption and associated environmental benefits increase if buyers obtain more bids? How does imperfect competition in the installation market influence the performance of government subsidies for solar PV?

5.2 Counterfactuals

Effects of Expanding Buyers' Choice Sets

To investigate the mechanisms driving the platform's consumer benefits. I measure the effect of expanding buyers' choice sets, holding preferences fixed. In particular, I simulate counterfactuals in which I alter the number of installers that bid on each buyer's project. In the counterfactuals, installers are informed about the number of other competing sellers (if any), and each installer updates their optimal bid price based the number of competitors they will face. In the counterfactuals, I assume that sellers know the number of competing bidders but not the identities of the other bidders. I draw the identities of the bidders randomly with the probabilities weighted by seller entry probabilities observed in the data.[42] After I draw the installers for each project, buyers choose from competing bids and the outside option.[43] I provide additional details about the algorithm for solving counterfactual equilibria in the appendix.

Table 7 displays the outcomes of changing the number of bids from one and five. I report all outcomes relative to the single-bid case (i.e., $\frac{y|bids=N}{y|bids=1}$). The top of the table shows that an increase from one to five bids leads to 15.5% reduction in mean bid prices, a 25.2% reduction in the lowest bid, and a 15.3% fall in purchase prices (selected bids). These results imply a $4,000 gross bid price decline (before subsidies) for a typical-sized system. The table also reveals a substantial decrease in the marginal effect of competition, namely, adding a second bid causes a much greater marginal price reduction (9.7% from the baseline) relative to adding a fifth bid (1.2% from the baseline).

Table 7: Effects of Number of Bids Per Project on Market Outcomes

	Counterfactual # of Bids Per Project				
	1 - Baseline	2	3	4	5
Relative Prices					
Mean Price Per Watt (All Bids)	1	0.880	0.855	0.842	0.834
Mean Price Per Watt (Lowest Bid)	1	0.834	0.788	0.762	0.745
Mean Price Per Watt (Selected)	1	0.903	0.874	0.857	0.845
Relative Quantities, Externalities, & Consumer Surplus					
# of Solar Installations	1	1.114	1.198	1.272	1.333
Solar Installed Capacity (KW)	1	1.112	1.194	1.267	1.327
Annual Solar Output (KWh)	1	1.108	1.187	1.257	1.314
Pollution Damages Avoided ($)	1	1.097	1.167	1.230	1.282
Consumer Surplus ($)	1	1.859	2.502	3.089	3.597
Producer Surplus (On Platform)					
Total Producer Surplus ($)	1	1.501	1.843	2.140	2.379
Producer Surplus Per Bidding Seller ($/bid)	1	0.750	0.614	0.535	0.476

Notes: Table summarizes counterfactual simulations varying the number of bids that each project receives. All outcomes are reported relative to the one-bid case.

The second part of the table shows the impact of bid quantity on the share of households that purchase a solar system. Measuring changes in the solar adoption is more challenging than measuring changes in prices because I do not observe possible off-platform solar purchases by households in the sample. As a consequence, when a buyer chooses the outside option in the counterfactuals, this could either mean that they choose an off-platform installer or that they do not purchase a solar system. Fortunately, I do observe whether each buyer had any bids from offline installers. I assume that all buyers *do* purchase a solar system off of the platform if

[42]The seller entry probabilities are calculated separately for each project type. I also assume that installers know the entry probabilities of each competitor.

[43]I simulate 100 iterations of each project and measure the average outcome across all iterations to reduce simulation noise.

they have an off-platform quote and also choose the outside option. This assumption provides a lower bound of the effect of an additional bid on solar adoption. The assumption yields a lower bound because as the number bids increase in the counterfactuals, any measured increase in solar purchases only comes from a subset of buyers that do not have any off-platform quotes.[44] We see that expanding the bid set from one to five bids leads to at least a 33% increase in the number of households that adopt solar. Relatedly, a five-fold increase in bids leads to a 32% increase in total installed capacity and a 31% increase in expected solar output accounting for the expected capacity and location of the installations. I also use estimates from Sexton et al. (2018) to calculate the net present value from avoided pollution damages resulting from these additional solar purchases.[45] I find that expanding buyers' choice sets from one to five bids leads to a 28% increase in external benefits via pollution reductions from fossil-fueled power plants. The final row of the table confirms that increasing the number of bids can deliver large consumer benefits. More specifically, providing a buyer with four additional solar bids causes a 360% increase in consumer surplus.

In the third section of the table, I document changes in producer surplus as the number of bids per project expands. The producer surplus estimates should be interpreted with some caution because the estimates only quantify changes in producer surplus for the sellers that are participating on the platform. The estimates do not account for changes in producer surplus that would accrue to installers operating off of the platform. With this caveat in mind, we see that total producer surplus for on-platform installers increases over twofold if the number of bids increase from one to five. The increase in total producer surplus is driven primarily by the rise in the number of transactions that occur as the number bids increases. Despite the increase in *aggregate* producer surplus, the expected profit for each individual seller falls by about 50% as the number of bidders increases from one to five.

Market Impacts of Solar Subsidies

In the next set of counterfactuals, I use the model to assess government subsidies' impact on market outcomes. Over the past few decades, subsidy programs have been a cornerstone of policymakers' efforts to expand renewable energy investment. The largest U.S. program promoting rooftop solar adoption has been the federal Investment Tax Credit (ITC). The ITC was originally established by the Energy Policy Act of 2005, and allowed solar buyers to deduct 30% of the solar installation price from their federal taxes through 2019. In 2020, the incentive dropped to 26% and is scheduled to be eliminated by 2022. The expiration of the ITC has been controversial. The solar industry, environmentalists, and many politicians have argued that renewable energy subsidies are critical to addressing climate change. On the other hand, op-

[44] The lower bound interpretation also relies on an implicit assumption that buyers report truthfully whether they have off-platform bids.

[45] Sexton et al. (2018) calculate the annual pollution damages avoided per unit of residential solar capacity for each ZIP code. To calculate the life-time net present value of a solar installation, I assume that each solar system has a 20-year life span and a discount rate of 5%.

ponents have countered that the tax credits are not cost-effective and place too large a burden on taxpayers.

Table 8: Effects of Government Subsidies on Competition, Prices, and Welfare

Panel A: Market Outcomes

	Counterfactual	
	ITC	No ITC
Competition		
Bids Per Project	3.35	3.06
Prices		
Gross Price Per Watt (All Bids)	3.52	3.10
Net Price Per Watt (All Bids)	2.47	3.10
Net Price Per Watt (Lowest)	2.25	2.82
Net Price Per Watt (Selected)	2.40	2.97
Output and Welfare Effects		
Solar Installations	-	-33%
Pollution Damages Avoided	-	-33%
Consumer Surplus	-	-35%
Producer Surplus	-	-56%
Bid Preparation Costs	-	-42%
Total Welfare incl. Subsidy Cost	-	-22%

Panel B: Average Effects of ITC Subsidies

Δ Pollution Damages Avoided/ Subsidy Cost	0.19
Δ Consumer Surplus/ Subsidy Cost	0.80
Δ Producer Surplus/ Subsidy Cost	0.70
Δ Bid Preparation Costs/ Subsidy Cost	0.03
Δ Welfare/ Subsidy Cost	0.69

Notes: Table summarizes counterfactual market simulations with and without the federal investment tax credit which compensates solar buyers for 30% of the purchase price on a solar system.

To understand the effects of removing the ITC, I simulate counterfactual market outcomes with the 30% subsidy and without the 30% subsidy. In the counterfactual simulations, sellers update both participation and bidding strategies to account for the subsidy availability (or removal). I discuss the algorithm used to solve the counterfactual equilibria in the appendix.

Panel A of Table 8 shows that eliminating the ITC results in an 9% reduction in the number of bids per project. Without the ITC, fewer installers find it worthwhile to submit bids because, for some projects, expected profits no longer cover their bid preparation costs. Removing the ITC also results in a 25.5% increase in the net price of solar bids obtained by buyers. This net price increase of $0.63/watt, implies a $4,600 increase on a typical-sized system. The consumer burden of removing the ITC is partly mitigated by sellers reducing their gross bid prices by $0.42/watt on average. Sellers find it optimal to reduce offer prices after the subsidy removal to recoup some sale quantity losses. If sellers did not change their bids in reaction to the ITC removal, the net bid price would increase by 42.8%.[46] The mean lowest bid price and the mean price of selected bids also increase by 25.3% and 23.7% when the ITC is eliminated.

The bottom of Panel A shows that removing the ITC also has significant effects on solar installation decisions, consumer welfare, producer welfare, and environmental outcomes. More specifically, the number of solar installations falls by 33% without the subsidy. Similarly, the expected pollution reductions caused by solar system investments decline by 33%. Removing

[46]The ITC provides a credit for 30% of the gross price, therefore, with the ITC, the net price equals 7/10 of the gross price, so removing the ITC without changing the gross bid price leads to a 10/7 (42.8%) increase in net price.

the ITC also results in a 35% reduction in consumer surplus attained by solar shoppers. The withdrawal of the ITC causes even more significant losses for producers. In particular, producer surplus falls by 56% due to the combination of lower gross margins and reduced sales volume. Sellers do benefit from a 42% reduction in bid preparation cost due to lower auction participation. Notably, the participation decline comes mostly from installers with relatively high bid preparation costs. However, as we saw in previous sections, bid preparation costs are relatively small compared to the overall installation cost. The bottom row of Panel A, we see that withdrawing the ITC subsidies causes a 22% reduction in total welfare. That is to say, we see that the ITC subsidies cause an increase in total welfare.

In Panel B, I investigate the mechanisms driving the welfare result. The top row shows that the ITC reduces pollution damages by \$0.19 for each dollar of subsidy expenditure. This number indicates that the subsidy is a relatively inefficient way to reduce pollution damages by itself. However, the next row shows that the ITC fosters substantial welfare increases for producers and consumers in the solar market. The ITC increases average consumer surplus by \$0.80 per subsidy-dollar. Also, the subsidy expands producer marginal profits by \$0.70 per subsidy-dollar. We also see that the ITC increases bid preparation costs by \$0.03 per subsidy-dollar by encouraging more bids. Summing all of these effects together and subtracting the $\$1.00$ subsidy cost, we obtain the net welfare benefit of $\$0.69$ per subsidy-dollar ($\$0.19+\$0.80+\$0.70-\$0.03-\$1.00=\$0.69$). Notably, the subsidy's external environmental benefits only explain a portion of the welfare gain, whereas the market participants' increased rents explain most of the welfare gain. The ITC increases welfare by correcting for unpriced environmental externalities and reducing deadweight loss caused by market power. In particular, we have seen that solar installations are priced well above marginal cost. Therefore, the subsidy encourages more transactions where buyers' willingness-to-pay exceeds the seller's marginal cost. As a result, the ITC causes a net welfare gain.

There are several limitations and caveats to keep in mind when interpreting the results of the welfare analysis. First, I do not consider several potentially important externalities of solar installations such a electricity transmission and distribution costs (Feger et al., 2017), learning-by-doing (Bollinger and Gillingham, 2019), peer effects (Bollinger and Gillingham, 2012), or generation reserve costs (Gowrisankaran et al., 2016) in the welfare calculation. Second, the ITC policy covers all residential solar installations in the U.S., whereas the model is estimated using only solar installation projects that originated through the online platform. The intensity of bidding competition is likely to be substantially different for projects originating off-platform relative to on-platform. For this reason, we would expect off-platform bids to respond differently to the ITC availability, yielding different quantitative welfare effects. Thus, we should be cautious to extrapolated the exact welfare estimates out of sample. Nonetheless, if we believe that competition for off-platform solar projects is less intense than competition for projects on the platforms–and therefore subject to higher markups—we would also expect substantial net benefits of subsidizing off-platform projects to constrain market power. Finally, the US

ITC policy also provides a tax incentive for utility-scale solar projects. Market structure in the utility-scale market is likely to differ from the residential market, and therefore, the welfare impacts of removing the ITC will also be likely to vary across these two markets.

6 Conclusion

Online platforms offer a convenient means for buyers and sellers to connect in markets with negotiated prices. Nonetheless, many such markets still widely operate via bilateral negotiations. In the solar PV market context, we have seen that online platforms can facilitate sizable price reductions and welfare gains for consumers. A primary mechanism through which the platform helps consumers is by providing them access to additional bids. These price reductions are particularly relevant in the solar PV market because (non-hardware) soft costs, which include installer markups, now account for over 70% of the price of a typical solar installation. Thus, policymakers aimed at reducing solar prices could consider informing the public about existing platforms or develop their own platforms to link buyers and sellers. Platforms could also yields similar benefits in other markets characterized by high search costs such as home mortgages or building energy efficiency retrofits.

Notably, though, we have seen that welfare gains will not be evenly distributed across consumers when sellers can bid on projects selectively. In particular, households with higher home values obtain the most significant benefits because these wealthier households attract relatively more bids through the platform. This finding implies that other complementary policies, such as progressive subsidies, may be required if policymakers desire more equitable solar investment across income groups. These distributional effects are also important to keep in mind as other markets move more towards online platforms to link market participants.

The counterfactuals also provide new evidence on the welfare impacts of investment subsidies. Policymakers most commonly justify renewable energy subsidies based on positive environmental externalities. However, government subsidies are often introduced into markets with multiple existing distortions such as imperfect competition, imperfect information, in addition to environmental externalities. Robinson (1933) was the first to note that subsidies could improve welfare in imperfectly-competitive markets, and Judd (2002) later argued that government subsidies could have particularly large welfare benefits in markets for capital-intensive goods. This paper empirically quantifies the welfare benefits of a prominent investment subsidy in the solar PV market, and highlights that subsidies can improve total welfare by constraining market power in addition to reducing pollution damages from electricity generation. These results have implications for tax and subsidy policy in other imperfectly-competitive markets.

References

Allen, J., Clark, R., and Houde, J.-F. (2019). Search frictions and market power in negotiated-price markets. *Journal of Political Economy*, 127(4):1550–1598.

Barbose, G., Darghouth, N., Elmallah, S., Forrester, S., Kristina SH, K., Millstein, D., Rand, J., Cotton, W., Sherwood, S., and O'Shaughnessy, E. (2019). Tracking the sun: Pricing and design trends for distributed photovoltaic systems in the united states - 2019 edition.

Barbose, G., Darghouth, N., Hoen, B., and Wiser, R. (2018). Income trends of residential pv adopters: An analysis of household-level income estimates.

Baye, M. R. and Morgan, J. (2001). Information gatekeepers on the internet and the competitiveness of homogeneous product markets. *American Economic Review*, 91(3):454–474.

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

Bollinger, B. and Gillingham, K. (2012). Peer effects in the diffusion of solar photovoltaic panels. *Marketing Science*, 31(6):900–912.

Bollinger, B. and Gillingham, K. (2019). Learning-by-doing in solar photovoltaic installations. *Available at SSRN 2342406*.

Bollinger, B., Gillingham, K., and Lamp, S. (2020). Equilibrium effects of competition on solar photovoltaic demand and pricing. *Working Paper*.

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

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.

Caceres, C. (2019). Analyzing the effects of financial and housing wealth on consumption using micro data. Technical report.

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

De Groote, O. and Verboven, F. (2019). Subsidies and time discounting in new technology adoption: Evidence from solar photovoltaic systems. *American Economic Review*, 109(6):2137–72.

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

Ellison, G. and Ellison, S. F. (2009). Search, obfuscation, and price elasticities on the internet. *Econometrica*, 77(2):427–452.

Feger, F., Pavanini, N., and Radulescu, D. (2017). Welfare and redistribution in residential electricity markets with solar power.

Feldman, D. (2014). Photovoltaic (PV) pricing trends: Historical, recent, and near-term projections.

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.

Gerarden, T. (2017). Demanding innovation: The impact of consumer subsidies on solar panel production costs. Technical report, Working paper, Harvard University.

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

Goeree, M. S. (2008). Limited information and advertising in the us personal computer industry. *Econometrica*, 76(5):1017–1074.

Goldmanis, M., Hortaçsu, A., Syverson, C., and Emre, Ö. (2010). E-commerce and the market structure of retail industries. *The Economic Journal*, 120(545):651–682.

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.

Hanemann, W. M. (1984). Discrete/continuous models of consumer demand. *Econometrica: Journal of the Econometric Society*, pages 541–561.

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.

Judd, K. L. (2002). Capital-income taxation with imperfect competition. *American Economic Review*, 92(2):417–421.

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. (2019). The Role of Quality in On-Line Service Markets (forthcoming). *Journal of Political Economy.*

Krasnokutskaya, E., Terwiesch, C., and Tiererova, L. (2016). Trading across Borders in Online Auctions. *Report, Johns Hopkins University.*

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.

Nemet, G. F., O'Shaughnessy, E., Wiser, R., Darghouth, N., Barbose, G., Gillingham, K., and Rai, V. (2016). Characteristics of Low-Priced Solar Photovoltaic Systems in the United States. *Berkeley, CA: Lawrence Berkeley National Laboratory.*

O'Shaughnessy, E., Nemet, G. F., Pless, J., and Margolis, R. (2019). Addressing the soft cost challenge in us small-scale solar pv system pricing. *Energy Policy*, 134:110956.

O'Shaughnessy, E. and Margolis, R. (2018). The value of price transparency in residential solar photovoltaic markets. *Energy Policy*, 117:406–412.

Pless, J. and van Benthem, A. A. (2019). Pass-through as a test for market power: An application to solar subsidies. *American Economic Journal: Applied Economics*, 11(4):367–401.

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

Robinson, J. (1933). *The economics of imperfect competition.* Springer.

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

Sexton, S. E., Kirkpatrick, A. J., Harris, R., and Muller, N. Z. (2018). Heterogeneous environmental and grid benefits from rooftop solar and the costs of inefficient siting decisions. Technical report, National Bureau of Economic Research.

Sigrin, B., Dietz, T., Henry, A., Ingle, A., Lutzenhiser, L., Moezzi, M., Spielman, S., Stern, P., Todd, A., Tong, J., et al. (2017). Understanding the evolution of customer motivations and adoption barriers in residential solar markets: survey data. Technical report, National Renewable Energy Laboratory-Data (NREL-DATA), Golden, CO (United ....

Snashall-Woodhams, N. (2019). Targeting solar subsidies. *Working Paper*.

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.

Weyl, E. G. and Fabinger, M. (2013). Pass-through as an economic tool: Principles of incidence under imperfect competition. *Journal of Political Economy*, 121(3):528–583.

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

Appendix

Algorithm for Solving Counterfactuals with Fixed # of Bids

1. For each project i, start with a vector of all bids submitted for projects of that type ($\mathbf{B_0}$), a fixed number of bids (N), and entry probability weights for each potential seller for that auction type ($\mathbf{E_0}$).

2. Calculate each firm's optimal price given the current distribution of prices and entry probabilities from step one. Store the new vector of bids $\mathbf{B_1}$.

   - Equation 5 is the first order condition for each firm's optimal price. The first order condition does not have a closed form, so simulate S=100 iterations of each auction type to approximate the integrals numerically.

3. Measure the difference between each of the original prices and the updated prices. Stop if $||abs(\mathbf{B_1} - \mathbf{B_0})||_\infty < \delta_b$. Otherwise replace $\mathbf{B_0}$ with $\mathbf{B_1}$ and then start over at Step 1.

   - I set $\delta_b = .00001$

Algorithm for Solving ITC Counterfactuals with Endogenous Bidding

1. For each auction type, start with a vector of all bids submitted for projects of that type ($\mathbf{B_1}$) and start with an entry probability for each potential entrant for that auction type ($\mathbf{E_0}$).

2. Draw S=100 vectors of non-price characteristics for potential entrant. Draw each vector of non-price characteristics at random from the list of all bids made by that project-seller type pair.

3. Draw S=100 uniform draws for each potential entrant to determine random entry for each simulation iteration.

   - Choose entrants for each simulation iteration by determining if the random uniform draw is less than $\mathbf{E_0}$
   - *Note: To ensure convergence, I hold the initial S=100 sets of simulated entrants fixed throughout the algorithm even though the entry probabilities will change in the counterfactuals. I use an importance sampling approach similar to Guerre et al. (2000) to adjust for the fact that I do not update the sets of entrants at each step.*

4. Set $\mathbf{E_0}$=$\mathbf{E_1}$

5. Calculate each firm's optimal price given the current distribution of prices $\mathbf{B_1}$ and entry probabilities $\mathbf{E_1}$. Store the new vector of bids $\mathbf{B_2}$.

- Equation 5 is the first order condition for each firm's optimal price. The first order condition does not have a closed form, so use the S=100 simulation iterations of each auction type to approximate the integrals numerically. When calculating the averages, I use importance weights to adjust for that fact that the competitors were drawn according to $\mathbf{E_0}$ instead of $\mathbf{E_1}$.

6. Use the updated prices (and conditional winning probabilities) from Step 2 to calculate each potential entrant's expected marginal profit of entering the auction. Then use the new expected profits to update each firm's entry probability. Store the new entry probabilities $\mathbf{E_2}$.

7. Measure the difference between each of the original prices and the updated prices and measure the difference between the original and updated entry probabilities. Stop if $||abs(\mathbf{B_2} - \mathbf{B_1})||_\infty < \delta_b$ and $||abs(\mathbf{E_2} - \mathbf{E_1})||_\infty < \delta_e$. Otherwise replace $\mathbf{B_1}$ with $\mathbf{B_2}$ and $\mathbf{E_1}$ with $\mathbf{E_2}$ and then start over at Step 4.

   - I set $\delta_b = 0.00001$ and $\delta_e = .0.00001$.

Figure 6: Google Maps Photo of the Rooftop for a Potential Project

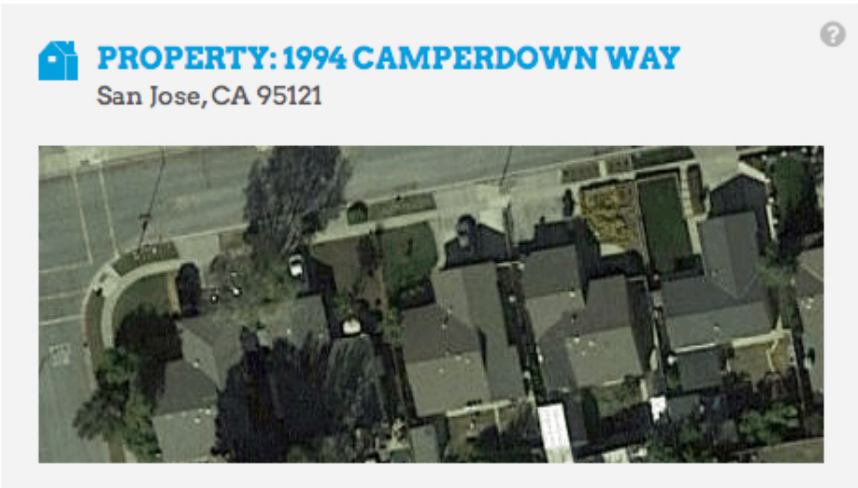

Figure 7: EnergySage Dashboard for Comparing Submitted Quotes

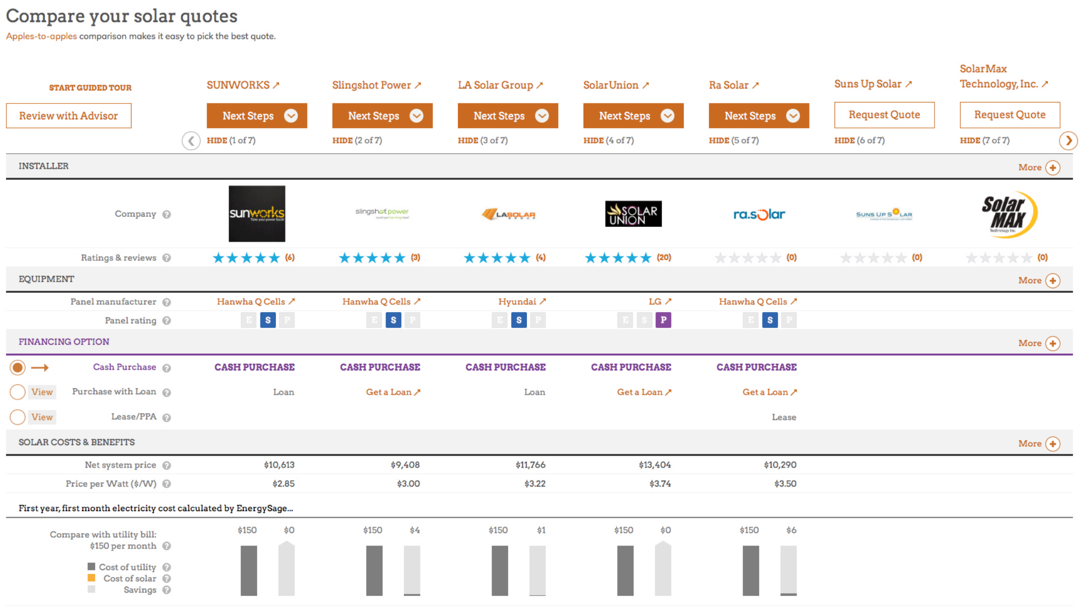

Notes: EnergySage quote comparison page in late 2016.

Table 9: Project Count by Location and Year

	2014	2015	2016	Total
AZ	57	110	563	730
CA	353	1053	3417	4823
CO	30	140	365	535
CT	132	559	330	1021
MA	129	400	1025	1554
NY	62	173	698	933
TX	49	156	687	892
Total	812	2591	7085	10488

Table 10: Regressions of Bid Price ($/watt) on Order of Bid

	(1) Gross Price ($/watt)	(2) Gross Price ($/watt)
Order of Bid	0.00164	0.00125
	(0.00147)	(0.00141)
Total Bids Control	Yes	Yes
System Size Control	Yes	Yes
CBSA FE	Yes	Yes
Half-Year FE	Yes	Yes
Installer FE	No	Yes
N	37080	37080
R^2	0.299	0.363

Notes: All standard errors listed in parenthesis are clustered by project id.

Table 11: Effects of Competition on Bid Prices and Proposed System Size

	Gross Price ($)	System Capacity (W)
Electric Bill ($/month)	6.066	22.40
	(0.783)	(0.803)
Home Value ($ 1000s)	0.391	0.162
	(0.153)	(0.146)
Mean # of Bids in Market	-364.9	15.11
	(32.39)	(31.68)
Mean # of Bids in Market × Electric Bill ($/month)	-0.938	-0.243
	(0.181)	(0.192)
Mean # of Bids in Market × Home Value ($ 1000s)	-0.0596	0.00773
	(0.0341)	(0.0340)
Capacity, Capacity^2 Controls	Yes	-
CBSA FE	Yes	Yes
Half-Year FE	Yes	Yes
Panel Brand FE	Yes	Yes
Inverter Brand FE	Yes	Yes
Installer FE	Yes	Yes
N	37080	37080
R^2	0.948	0.597

Notes: The mean number of bids in a market is defined as the average number of bids for all projects within the same CBSA and the same half-year. Household electric bill and home value variables are demeaned before running the regressions. Standard errors clustered by project are in parentheses.

Figure 8: Auction Participation and Pricing By Electricity Expenditure

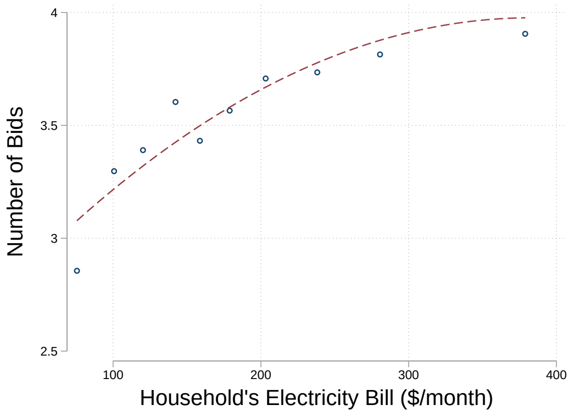

(a) # of Bids by Elec. Bill

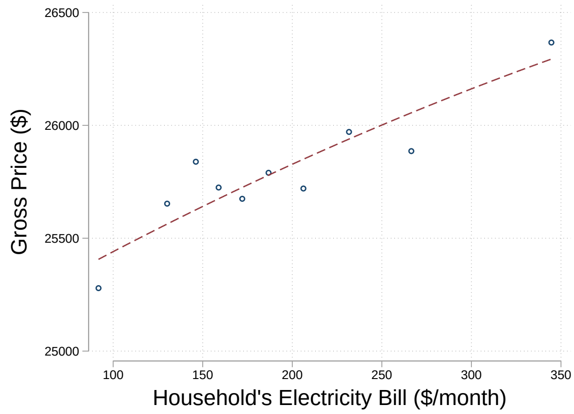

(b) Bid Prices by Elec. Bill

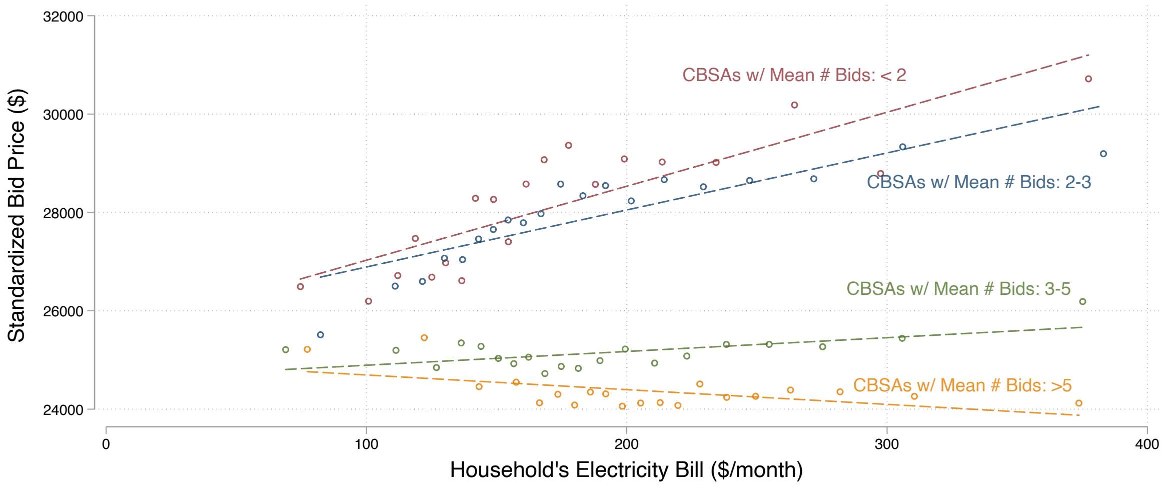

(c) Prices, Elec. Bill, Mean Bids in CBSA

Notes: Panels b,c linearly adjust the prices for the size (kW) of the system size and the time (half year) the project occurred before plotting. In panel c, the mean bids in the market is the average number of bids across all projects within the same CBSA and the same half year.

Figure 9: Potential Project Locations

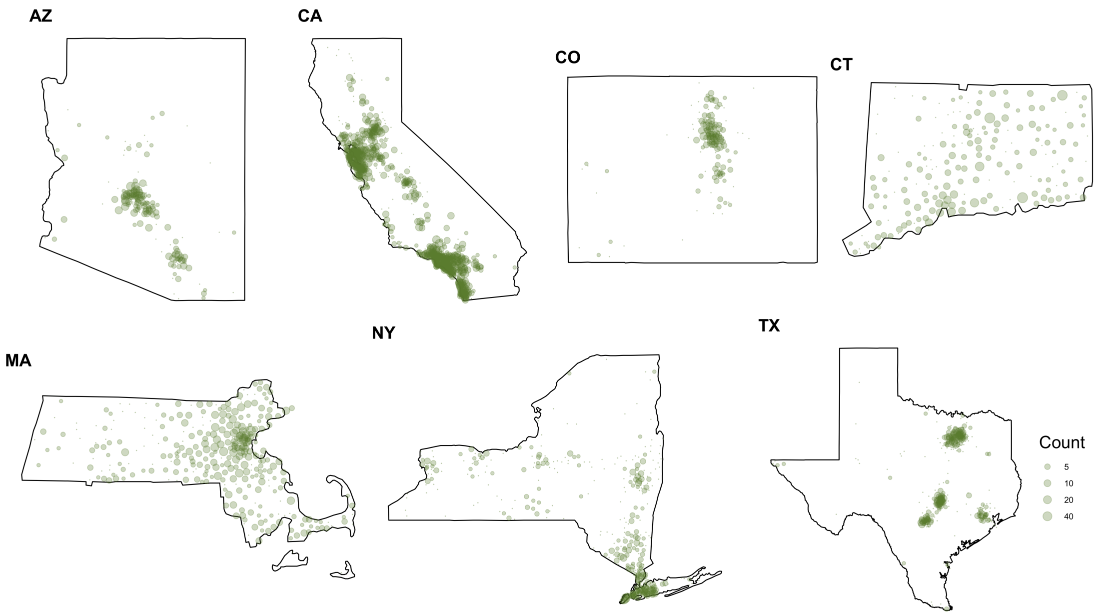

Notes: Count is the total number of potential projects within a ZIP code during the full sample.

Figure 10: Correlation of Variables from Different Data Sources

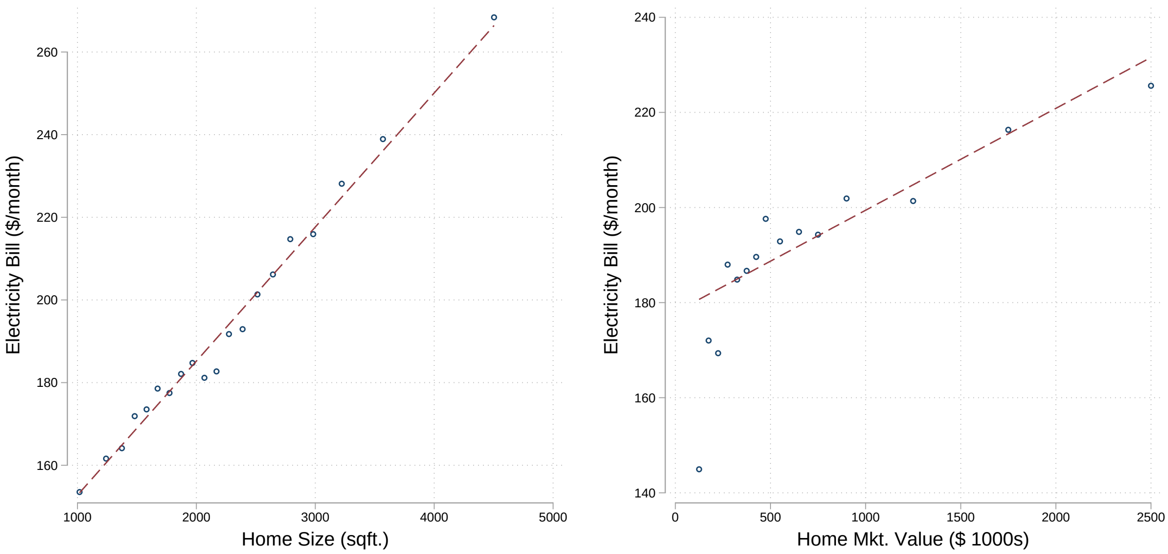

Notes: The binned scatter plots project the average monthly electricity bill as reported by consumers directly through the EnergySage website on the y-axis. Home square footage and home market value from Infinite Media, Inc. are plotted on the x-axis. Infinite Media, Inc. reports home market value as a range of values for each household and the middle value of the range is used.

Figure 11: Household Income Distributions - Platform Users vs. All Solar PV Adopters

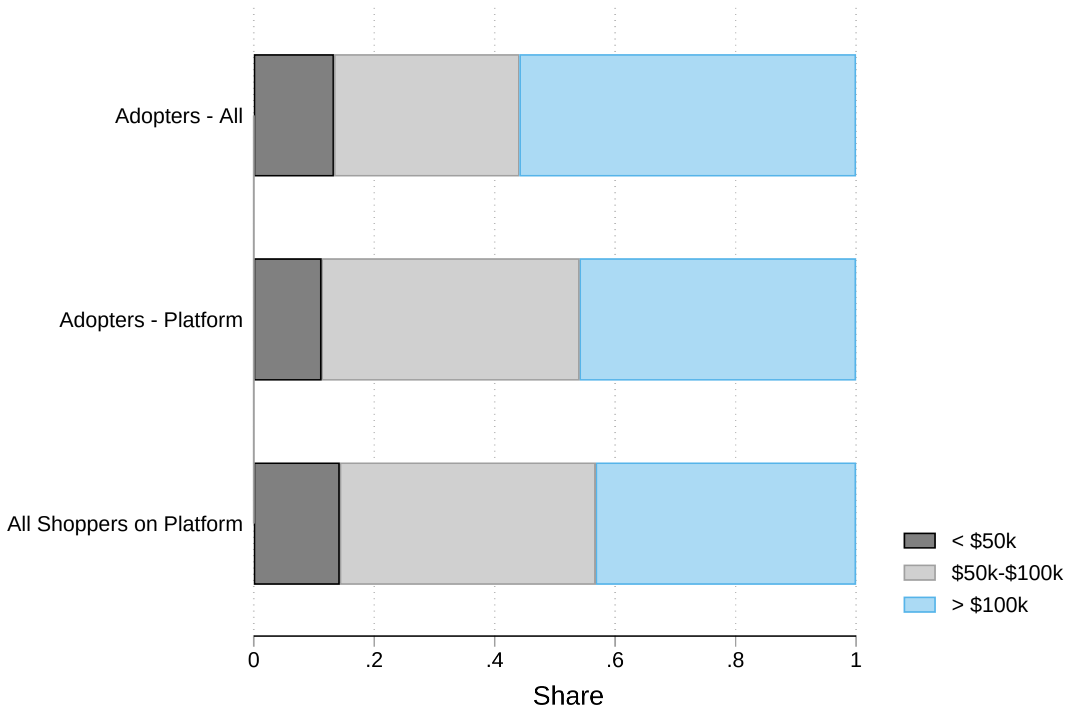

Notes: The data on income distributions for off-platform solar adopters were obtained from Barbose et al. (2018).

Figure 12: Model Fitted Entry Probabilities - 2016 H1, Nassau County-Suffolk County, NY

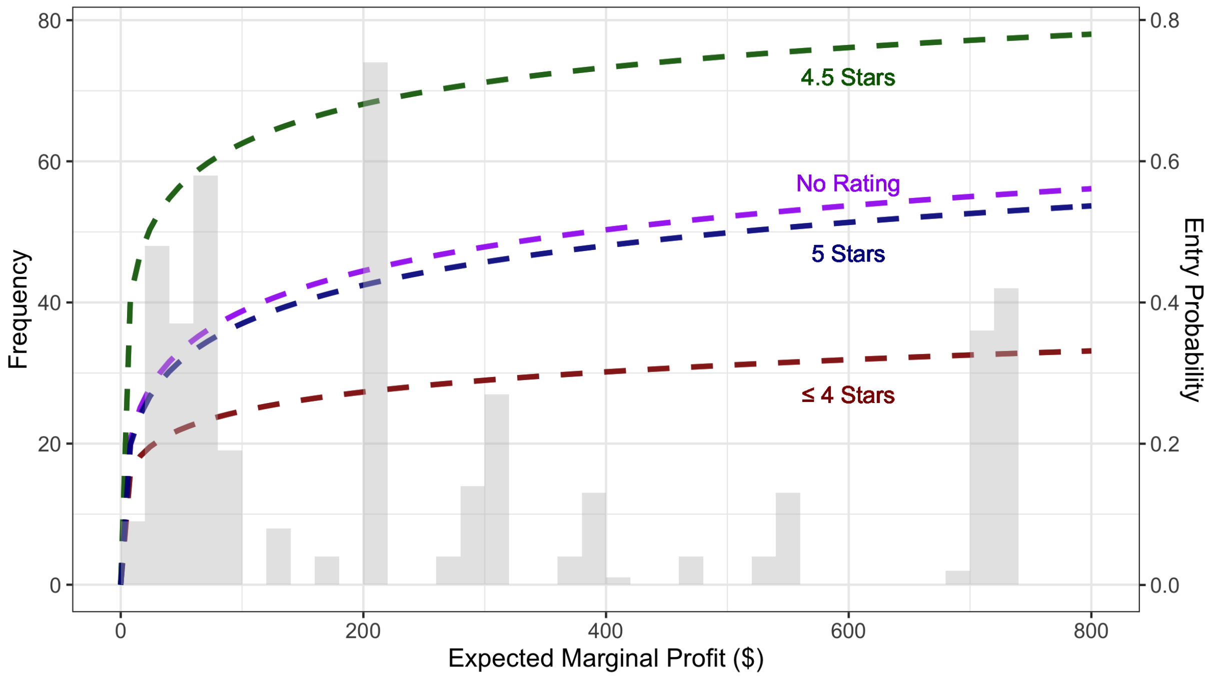

Notes: For a single market, the lines plot the change in entry probability as expected marginal profits increase for sellers with different ratings. The entry probabilities are implied by the estimated entry cost distribution. The grey bars show the frequency distribution of expected marginal profit for projects in Nassau County-Suffolk County, NY in 2016 H1.

Figure 13: US Solar Installation Price Dynamics

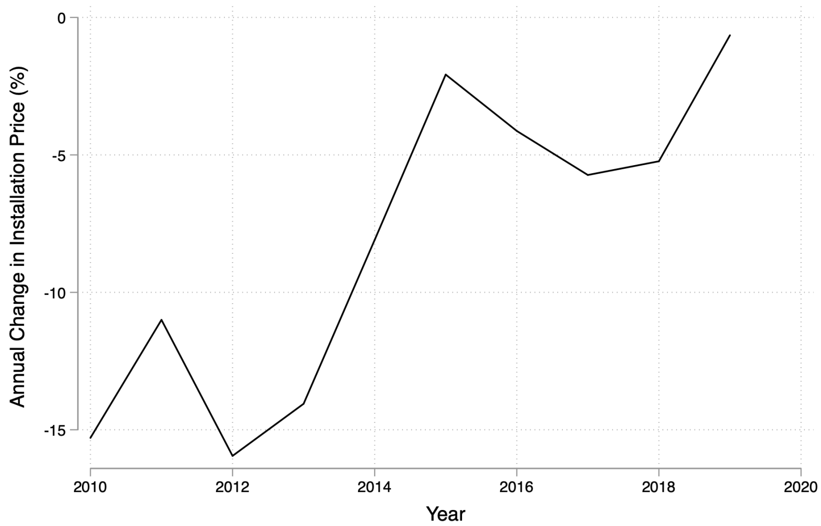

Notes: The line represents the annual percentage change in the median residential solar installation price (per watt) . Figure constructed using data from Barbose et al. (2019).

Figure 14: Assessing Marginal Cost Estimates

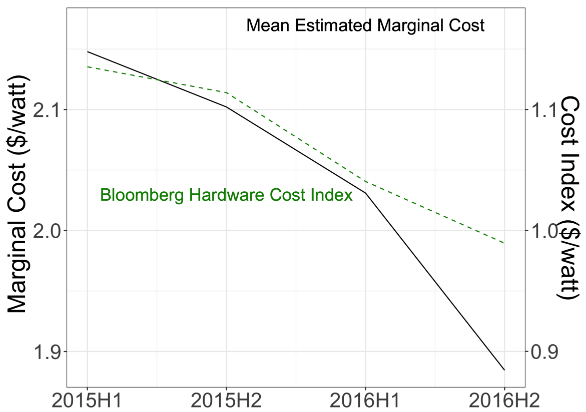

(a) Hardware Cost and MC Estimate

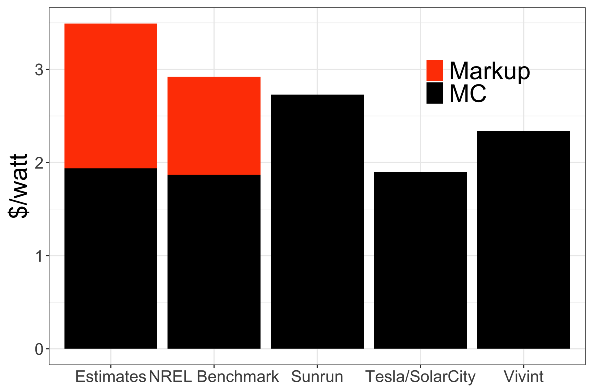

(b) MC Estimates vs. Stated Costs

Notes: Panel 14a compares the evolution of estimated marginal costs (mean) to Bloomberg's solar PV cost index for the final two years of the sample. The Bloomberg cost index is the sum of the Bloomberg's polysilicon panel cost index and Bloomberg's inverter cost index. Panel 14b compares estimated marginal cost (mean) to NREL's 2016 cost benchmark and stated costs reported by three large publicly-trade installers.

Table 12: Alternate Model Specifications - Adding Controls

	(1) - Base	(2)	(3)	(4)	(5)
λ	0.37 (0.06)	0.37 (0.06)	0.37 (0.06)	0.37 (0.06)	0.37 (0.06)
β - Price	-0.72 (0.11)	-0.72 (0.11)	-0.71 (0.11)	-0.71 (0.11)	-0.71 (0.11)
β - Price $\times$ Home Mkt. Value Q2	0.07 (0.04)	0.07 (0.04)	0.07 (0.05)	0.06 (0.05)	0.06 (0.05)
β - Price $\times$ Home Mkt. Value Q3	0.08 (0.05)	0.08 (0.05)	0.08 (0.05)	0.07 (0.05)	0.06 (0.06)
β - Price $\times$ Home Mkt. Value Q4	0.09 (0.06)	0.09 (0.06)	0.09 (0.06)	0.07 (0.06)	0.07 (0.07)
Mean Own-Price Elasticity	-1.59	-1.59	-1.58	-1.60	-1.60
Mean Markup ($/watt)	1.59	1.59	1.60	1.58	1.58
Log Likelihood	-3823.54	-3821.79	-3819.75	-3820.01	-3821.43
Installer Attributes					
Fixed Effects for Permanent Installers	Y	Y	Y	Y	Y
Installer Rating and Exper. Controls	Y	Y	Y	Y	Y
Experience/Time on Platform Controls	N	Y	N	N	N
Non-Price Bid Attributes					
Hardware Quality Controls	Y	Y	N	Y	Y
Panel Brand Fixed Effects	N	N	Y	N	N
Project Attributes $\times$ Inside Good					
Electric Bill Quartile Fixed Effects	Y	Y	Y	Y	Y
CBSA Fixed Effects	Y	Y	Y	Y	Y
Half-Year Fixed Effects	Y	Y	Y	Y	Y
Off-Platform Quote Fixed Effect	Y	Y	Y	Y	Y
Bachelor's Degree Fixed Effect	N	N	N	Y	N
Race Fixed Effects	N	N	N	Y	N
Over Age 65 Fixed Effect	N	N	N	Y	N
Home Sq. Footage Quartile Fixed Effects	N	N	N	N	Y
Home Age Quartile Fixed Effects	N	N	N	N	Y

Notes: The top panel displays the nesting parameter and price coefficients for the base model and alternate demand specifications. Standard errors are in parentheses. The middle panel shows the mean own-price elasticity and mean implied markup across all bids in the data. The bottom panel indicates which additional controls are included in utility. All models include fixed effects for all "permanent" installers, defined as any installer that submitted over 300 bids during the sample. All models include dummies for installers' star rating in 2016 and overall residential installation experience. The second model includes two dummies that indicate if the installer had 1) been bidding on EnergySage for at least 6 months, and 2) been bidding on EnergySage for over a year at the time the bid was submitted. All models include a microinverter dummy to control for hardware quality. Models 1,2,4, and 5 include dummies for "premium" and "premium plus" panel brands. The third model includes panel-brand dummies for the largest 7 panel manufacturers. The fourth model includes additional demographic control variables such as the education, race, and age of the household head. The fourth model also controls for the home square footage quartile. All models include electric bill quartile fixed effects, CBSA fixed effects, and half-year fixed effects. All variables listed under *Project Attributes* are interacted with the "inside good". The fifth model interacts price with a dummy for whether the potential buyer stated that the they already had quotes from another installer off of the platform. Standard errors are in parentheses.

Table 13: Alternate Model Specifications - Functional Form and System Capacity Choice

	(1) - Base	(2)	(3)
λ	0.37 (0.06)	0.37 (0.06)	0.37 (0.06)
β - Price	-0.72 (0.11)	-0.72 (0.12)	-1.91 (0.3)
β - Price $\times$ Home Mkt. Value Q2	0.07 (0.04)	0.07 (0.05)	0.18 (0.12)
β - Price $\times$ Home Mkt. Value Q3	0.08 (0.05)	0.08 (0.05)	0.2 (0.15)
β - Price $\times$ Home Mkt. Value Q4	0.09 (0.06)	0.09 (0.06)	0.23 (0.16)
Mean Own-Price Elasticity	-1.59	-1.61	-1.70
Mean Markup ($/watt)	1.59	1.57	1.49
Log Likelihood	-3823.54	-3823.50	-3823.49
Price Variable	$/watt	$/watt	ln($/watt)
Seller Proposed System Capacity Control (W)	N	Y	N
Installer Attributes			
Fixed Effects for Permanent Installers	Y	Y	Y
Installer Rating and Exper. Controls	Y	Y	Y
Non-Price Bid Attributes			
Hardware Quality Controls	Y	Y	Y
Project Attributes $\times$ Inside Good			
Electric Bill Quartile Fixed Effects	Y	Y	Y
CBSA Fixed Effects	Y	Y	Y
Half-Year Fixed Effects	Y	Y	Y
Off-Platform Quote Fixed Effect	Y	Y	Y

Notes: The first column presents the baseline utility estimates with the price variable measured linearly in $/watt. The second column allows each seller's proposed system capacity (W) to enter as a non-price attribute in the buyer utility function. The final column estimates a demand with the natural logarithm of unit price entering utility. Standard errors are in parentheses.

Table 14: Alternate Model Specifications - Removing Controls

	(1) - Base	(2)	(3)	(4)	(5)
λ	0.37 (0.06)	0.34 (0.06)	0.39 (0.06)	0.41 (0.06)	0.44 (0.06)
β - Price	-0.72 (0.11)	-0.7 (0.11)	-0.83 (0.1)	-0.59 (0.1)	-0.57 (0.09)
β - Price $\times$ Home Mkt. Value Q2	0.07 (0.04)	0.07 (0.04)	0.07 (0.04)		
β - Price $\times$ Home Mkt. Value Q3	0.08 (0.05)	0.09 (0.05)	0.08 (0.05)		
β - Price $\times$ Home Mkt. Value Q4	0.09 (0.06)	0.1 (0.06)	0.1 (0.06)		
Mean Own-Price Elasticity	-1.59	-1.50	-1.90	-1.51	-1.51
Mean Markup ($/watt)	1.59	1.68	1.33	1.67	1.68
Log Likelihood	-3823.54	-3840.94	-3887.81	-3849.38	-3867.63
Installer Attributes					
Fixed Effects for Permanent Installers	Y	Y	N	Y	Y
Installer Rating and Exper. Controls	Y	N	Y	Y	Y
Non-Price Bid Attributes					
Hardware Quality Controls	Y	Y	Y	Y	Y
Project Attributes $\times$ Inside Good					
Electric Bill Quartile Fixed Effects	Y	Y	Y	N	N
CBSA Fixed Effects	Y	Y	Y	Y	N
State Fixed Effects	N	N	N	N	Y
Half-Year Fixed Effects	Y	Y	Y	Y	Y
Off-Platform Quote Fixed Effect	Y	Y	Y	Y	Y

Notes: The top panel displays the nesting parameter and price coefficients for the base model and alternate demand specifications. Standard errors are in parentheses. The middle panel shows the mean own-price elasticity and mean implied markup across all bids in the data. The bottom panel indicates which additional controls are included in utility. Models 1,4, and 5 include dummies for installers' star rating in 2016 and overall residential installation experience: one dummy indicating the installer has completed over 100 installs , and another dummy indicating over 1000 installs completed. Models 1,2,4, and 5 include fixed effects for all "permanent" installers, defined as any installer that submitted over 300 bids during the sample. All models include dummies for microinverter, "premium panel brand" and "premium plus panel brand" to control for hardware quality. Some models include electric bill quartile fixed effects, CBSA fixed effects, state fixed effects, and half-year fixed effects. All variables listed under *Project Attributes* are interacted with the "inside good". Standard errors are in parentheses.

Table 15: Alternate Model Specifications - Drop States with Time-Varying Subsidies

	(1) - Base	(2)
λ	0.37 (0.06)	0.38 (0.07)
β - Price	-0.72 (0.11)	-0.78 (0.13)
β - Price $\times$ Home Mkt. Value Q2	0.07 (0.04)	0.06 (0.06)
β - Price $\times$ Home Mkt. Value Q3	0.08 (0.05)	0.08 (0.07)
β - Price $\times$ Home Mkt. Value Q4	0.09 (0.06)	0.1 (0.07)
Mean Own-Price Elasticity	-1.59	-1.73
Mean Markup ($/watt)	1.59	1.45
Log Likelihood	-3823.54	-3009.82
Sample	Full	Drop NY & CT
Installer Attributes		
Fixed Effects for Permanent Installers	Y	Y
Installer Rating and Exper. Controls	Y	Y
Non-Price Bid Attributes		
Hardware Quality Controls	Y	Y
Project Attributes $\times$ Inside Good		
Electric Bill Quartile Fixed Effects	Y	Y
CBSA Fixed Effects	Y	Y
Half-Year Fixed Effects	Y	Y
Off-Platform Quote Fixed Effect	Y	Y

Notes: The first column presents the baseline utility estimates. The second column shows estimates of the equivalent model, but drops all auctions in New York and Connecticut. New York and Connecticut had changes in solar incentives during the sample period.

Table 16: Entry Cost Estimates

$\boldsymbol{\mu}$		$\boldsymbol{\sigma}$	
Star Rating ≤ 4	4.349 (0.597)	Constant	4.733 (0.234)
Star Rating = 4.5	-2.818 (0.195)	Star Rating ≤ 4	3.567 (0.530)
Star Rating = 5	0.275 (0.124)	Star Rating = 4.5	-0.143 (0.262)
Installs Completed: 100-1000	0.953 (0.103)	Star Rating = 5	0.173 (0.171)
Installs Completed: $>$1000	1.101 (0.114)		

Installer Rating	**Mean Bid Preparation Cost**	**Share of Total Bids**
<= 4 Stars	$ 5.33	0.14
4.5 Stars	$ 13.28	0.08
5 Stars	$ 20.97	0.61
No Ratings	$ 18.62	0.16

Fixed Effects		**Pseudo Log Likelihood**
CBSA Fixed Effects in $\boldsymbol{\mu}$	Yes	-62192.04
Half-Year Fixed Effects in $\boldsymbol{\mu}$	Yes	
Permanent Installer Fixed Effects in $\boldsymbol{\mu}$	Yes	

Notes: The top panel shows several of the parameter estimates from the entry cost model, as well as standard errors in parentheses. Coefficients for the CBSA fixed effects, half-year fixed effects, and permanent installer fixed effects in μ are not shown. The middle panel summarizes the expected bid preparation costs conditional on bidding.

Table 17: Marginal Costs Across CBSAs

Panel A: Lowest Marginal Cost CBSAs in 2016 H1

CBSA	MC (Mean)	MC (SD)
Phoenix-Mesa-Scottsdale, AZ	1.46	0.58
Tucson, AZ	1.58	0.36
Other, TX	1.62	0.47
Austin-Round Rock, TX	1.68	0.41
Other, CO	1.68	0.35
Other, AZ	1.7	0.4
Fort Worth-Arlington, TX	1.72	0.36
Dallas-Plano-Irving, TX	1.74	0.43
Denver-Aurora-Lakewood, CO	1.74	0.34
San Antonio-New Braunfels, TX	1.86	0.34

Panel B: Highest Marginal Cost CBSAs in 2016 H1

CBSA	MC (Mean)	MC (SD)
New York, NY	2.42	0.7
Providence-Warwick, RI-MA	2.4	0.43
Other, CT	2.36	0.68
Hartford, CT	2.28	0.46
New Haven-Milford, CT	2.28	0.43
Worcester, MA-CT	2.28	0.4
Other, NY	2.26	0.62
Other, MA	2.25	0.52
San Jose-Sunnyvale-Santa Clara, CA	2.25	0.5
Bridgeport-Stamford-Norwalk, CT	2.23	0.45

Notes: The top panel shows the mean and standard deviation of marginal cost for the ten lowest cost CBSAs in 2016 H1. The lower panel shows the mean and standard deviation of marginal cost for the ten highest cost CBSAs in 2016 H1.

Table 18: Consumer Surplus Gains from Access to the Platform by CBSA

Panel A: CBSAs with Smallest Consumer Surplus Gain ($)

CBSA	CS (Mean)	CS (SD)
Other, TX	201.92	192.54
Fort Worth-Arlington, TX	302.19	196.11
Fresno, CA	459.66	278.33
Austin-Round Rock, TX	513.07	344.91
Other, AZ	520.15	338.13
Dallas-Plano-Irving, TX	528.81	327.74
Phoenix-Mesa-Scottsdale, AZ	551.27	505.39
Denver-Aurora-Lakewood, CO	706.89	522.41
Bridgeport-Stamford-Norwalk, CT	732.01	457.52
San Antonio-New Braunfels, TX	780.45	434.97

Panel B: CBSAs with Highest Consumer Surplus Gain ($)

CBSA	CS (Mean)	CS (SD)
Norwich-New London, CT	3212.59	2019.84
Worcester, MA-CT	3182.82	2059.99
Cambridge-Newton-Framingham, MA	3012.63	1887.74
Hartford, CT	2878.94	1781.09
Boston, MA	2653.86	1526.9
New Haven-Milford, CT	1960.25	1183.3
Other, MA	1952.9	1160.51
Other, NY	1785.88	1125.69
Other, CT	1737.28	1149.16
Providence-Warwick, RI-MA	1726	906.67

Notes: The top panel shows the mean and standard deviation of consumer surplus for the bottom ten CBSAs. The lower panel shows the mean and standard deviation of consumer surplus for the top ten CBSAs.