

Solar Market Frictions: The Role of Platforms and Policies

Jackson Dorsey

October 12, 2024

Abstract

This paper estimates a structural model of the solar photovoltaic (PV) installation market using detailed data from an online platform. The model incorporates households' solar installation choices and sellers' strategic bidding. Counterfactual simulations evaluate how equilibrium outcomes respond to changes in market structure and government subsidies, yielding two main results: (1) an increase from one to five bids through the platform reduces gross installation prices by \$4,000 (15.5%), and (2) the U.S. Solar Investment Tax Credit increases total surplus on the platform by \$1.35 per dollar of subsidy expenditure by mitigating market power and reducing pollution externalities.

JEL Codes: D22, D44, L11, L15, Q40, Q41.

⁰Department of Economics, The University of Texas at Austin (e-mail: jackson.dorsey@austin.utexas.edu). This paper was originally the first chapter of my dissertation and was also previously circulated with the title: "Access to Alternatives: Increasing Rooftop Solar Adoption with Online Platforms." I thank coeditor Daniel (Yi) Xu, Ashley Langer, Derek Lemoine, Gautam Gowrisankaran, Stan Reynolds, Gary Solon, Keith Meyers, Dev Bose, Alvin Murphy, Matt White, Stefan Weiergraeber, Emily Baca, Derek Wolfson, Dan Sacks, Andrew Butters, Mattijs Wildenbeest, Dan Ackerberg, four anonymous referees, and numerous seminar participants helpful comments, discussions, and suggestions. I also thank Spencer Fields at EnergySage Inc. for instrumental help with data access.

Solar photovoltaics (PV) allow homeowners to generate zero-emissions electricity by bolting panels to their rooftops. Consequently, policymakers have offered numerous incentives to encourage household investment in solar PV and reduce electricity sector emissions. Nonetheless, adoption remains sparse, with fewer than 4% of U.S. single-family homes having installed a solar system by 2020 (Barbose et al., 2021).

Like many contracting services, solar installations are customized and priced case by case. Hence, buyers in these markets must bear substantial hassle costs to contact potential sellers. These 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, increasing that seller’s optimal markup. The social cost of market power may be magnified in markets with positive externalities such as renewable energy, energy efficiency, and communications technologies.

Online platforms can alleviate market frictions by linking buyers to sellers. Platforms compile information about a buyer’s characteristics and pass the data to a broad set of sellers who can then submit bids via the platform. Thus, platforms can provide buyers access to more price quotes than they would otherwise obtain offline. In this way, platforms may increase the number of transactions completed and improve market efficiency by connecting buyers to better-matched sellers or reducing prices by increasing competition.

In this paper, I estimate a structural model of the rooftop solar market that incorporates households’ solar installation choices and sellers strategically bidding on projects through an online platform. First, I use the estimates to quantify markups embedded in solar installation bids submitted through the platform. Second, I measure the platform’s effect on consumers’ solar adoption choices and welfare. Third, I evaluate counterfactuals. In one series of counterfactuals, I measure the impact of changing the number of installation bids each household obtains on equilibrium prices, economic surplus, solar installations, and environmental externalities. In the next set of counterfactuals, I assess the welfare impacts of the most extensive U.S. solar subsidy program, the Investment Tax Credit (ITC), for buyers and sellers on the platform.

Understanding how market structure influences solar installation prices is essential for

addressing both affordability concerns and low adoption rates in the solar industry. Despite falling prices, the average solar installation still exceeded \$20,000 in 2020. Figure 1 shows that hardware material costs account for less than 30% of the total price of a U.S. residential solar installation. The balance of system cost, or “soft cost,” includes installation labor, permitting, inspection, interconnection costs, and installers’ markups. High soft costs suggest that market power may be restricting residential solar investments (O’Shaughnessy et al., 2019). Evidence supports the presence of market power within the solar installation market. For instance, Pless and van Benthem (2019) show that solar subsidy pass-through rates imply market power. Similarly, Borenstein (2017) and Gillingham et al. (2016) show that households with high electricity rates and higher expected incomes pay more for solar installations, indicating price discrimination. Many U.S. markets now have numerous installers, suggesting that high hassle costs enable imperfectly competitive pricing. Therefore, platforms could reduce installation prices by increasing competition among installers.

Imperfect competition can also impact the efficiency of government subsidy programs by determining the share of subsidies passed through to end-consumer prices (Weyl and Fabinger, 2013). If monopolistic sellers respond to subsidies by increasing gross prices, subsidies may only generate small increases in solar installations and associated environmental benefits. On the other hand, subsidies may increase welfare by addressing unpriced environmental externalities and market power. Thus, an empirical analysis is needed to evaluate the welfare effects of solar subsidies.

To investigate counterfactual market outcomes and welfare in the solar PV market, I develop a three-stage model. First, a solar installation project is announced, and potential sellers learn about its characteristics, deciding whether to bid based on expected marginal profits. In the second stage, sellers learn their marginal cost and submit bids. The optimal bid depends on the installer’s cost, the household’s price elasticity, and expected competition. In the final stage, the household chooses one of the bidding sellers or the outside option.

The allocation mechanism is a multi-attribute auction (i.e., a beauty contest auction) where buyers consider factors other than prices, such as installers’ ratings and experience. I estimate the model using methods from Krasnokutskaya et al. (2020) and Yoganarasimhan (2015) and a

novel dataset of residential solar auctions from a major online platform. Unlike most solar PV data, I observe non-winning bids and bids to households choosing the outside option, allowing precise estimates of elasticities and markups.

I augment the platform data with detailed household-level demographic information, including average monthly electricity expenditure, home market value, and demographic variables like age, race, and educational attainment. While these data are highly detailed, a limitation is that the platform sample may not fully represent the population of potential solar buyers. Nonetheless, I demonstrate that the observable characteristics of platform shoppers closely resemble those of other solar adopters in the broader market. Notably, the sample exhibits significant variation in home values, allowing me to measure differences in seller participation and pricing across households of differing wealth. Therefore, the data allow for an assessment of which households benefit most from using the platform.

Findings: I find that 45% of gross installation prices are due to markups over marginal cost, higher than previous estimates in the literature (e.g., [Bollinger and Gillingham, 2019](#)). Notably, these markups exclude installers' overhead and marketing costs. Based on [Fu et al. \(2016\)](#), I calculate that installers' average net profit margins in 2016 were 24.8%

Leveraging a full-information assumption,¹ I find that buyers derive substantial consumer surplus gains from using the online platform. However, these consumer surplus gains differ across markets and household demographics. Wealthier households benefit more, as higher-valued homes attract more bids and purchase larger solar systems.

The first set of counterfactuals shows that increasing the number of bids per household significantly impacts consumer welfare and market outcomes. Varying the number of bids per project, I find that an increase from one to five bids lowers gross prices by 15.5%, about \$4,000 per installation, and boosts consumer surplus and solar purchases. This result highlights the potential benefits of reducing search frictions in negotiated price markets.

The second set of counterfactuals measures the impact of eliminating the Solar Investment Tax Credit (ITC) on prices, economic surplus, and pollution externalities. The ITC enables

¹Estimating consumer surplus gains requires an assumption that buyers are fully informed about the outside option, which includes off-platform installers' prices and characteristics.

solar buyers to deduct 30% of the gross installation price from federal taxes. Removing the ITC increases net bid prices by at least 20% and reduces bids per project by over 7%. Each subsidy dollar increases total surplus on the platform by at least \$1.35, including \$0.55 in consumer surplus, \$0.69 in producer surplus, \$0.11 in environmental benefits.² This highlights a surprising potential benefit of solar subsidies—under imperfect competition, subsidies can improve welfare by enabling more transactions where the buyer’s willingness to pay exceeds the seller’s marginal cost. However, the results should be interpreted cautiously because they do not capture all of the welfare impacts of the ITC, such as the impact on off-platform sellers and buyers.

The results have several policy implications. The model estimates and counterfactuals suggest that installer markups are substantial and that increasing competition through online platforms can reduce prices. Yet, as of 2016, only 3% of solar buyers used an online platform (O’Shaughnessy and Margolis, 2018). Some policymakers, such as the state of Connecticut, have begun advertising platforms to potential solar adopters.³ Similar approaches in other industries, like healthcare,⁴ suggest that policymakers could develop or advertise platforms to reduce solar PV prices and boost adoption.

The second set of counterfactuals is relevant to renewable energy policy. In recent years, policymakers have removed or considered removing several existing subsidy policies, such as the ITC. I find that removing the ITC substantially reduces the economic surplus among users of the online platform and modestly reduces associated environmental benefits. These results, therefore, suggest that renewable energy subsidies can—in some cases—help correct market power distortions in addition to reducing pollution externalities.

The results also provide insights into the economics of two-sided platforms more broadly. Many other industries have shifted transactions online in the past decade. In markets characterized by high hassle costs—such as home mortgages or customized services—platforms

²Producer surplus includes gross profits and bid preparation costs. I use exit survey data to estimate incremental adoption and avoided pollution damages from the ITC.

³See Connecticut Greenbank’s website gosolarct.com, under CT Legislature’s Act 11-80.

⁴The Affordable Care Act established health insurance marketplaces to reduce premiums.

can deliver substantial price reductions and consumer welfare gains. However, the empirical results also show that the magnitude of these welfare gains can differ substantially across consumers when sellers can discriminate on both the extensive margin (selection into bidding) and the intensive margin (bid prices). More specifically, the findings indicate that price-sensitive consumers may benefit relatively less from platforms because they attract fewer bids.

Related Literature: This paper builds on multiple literatures. First, it relates to the growing literature on the economics of the residential solar market. Several papers estimate the adoption response to price or subsidy changes. For instance, [Hughes and Podolefsky \(2015\)](#), [Pless and van Benthem \(2019\)](#), and [Gillingham and Tsvetanov \(2019\)](#) use quasi-experimental approaches to estimate the elasticity of demand and quantify the adoption response to subsidy programs. [Feger et al. \(2022\)](#), [De Groot and Verboven \(2019\)](#), and [Langer and Lemoine \(2022\)](#) develop dynamic discrete choice models focusing on the binary decision of when to adopt in order to evaluate subsidy policies. I build on this work by modeling buyers' choice of installer and competition among installers.⁵

Supply-side behavior and incentives are crucial in determining equilibrium prices and technology adoption. Most previous literature has yet to model the supply side of the solar PV market. To fill this gap, I develop a model accounting for (1) imperfect competition due to imperfectly informed consumers, (2) seller selection into bidding, (3) strategic pricing based on household characteristics, and (4) sellers' imperfect information about their competition. This model allows me to account for supply and demand responses to counterfactuals, such as subsidy policy changes. One notable 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. [O'Shaughnessy and Margolis \(2018\)](#) compare installation prices paid by solar consumers who used an online platform and those who did not, finding platform users pay lower prices. I build on [O'Shaughnessy and Margolis \(2018\)](#) by developing a structural approach to measure markups, assess platform effects on consumer welfare, and investigate counterfactuals.

⁵[Gerarden \(2023\)](#) models competition in the upstream solar manufacturing market.

This paper also pertains to the literature on competition in search markets, platforms, intermediaries, and the Internet. Seminal work by [Baye and Morgan \(2001\)](#) theoretically investigated intermediaries in online markets. Recent empirical studies examine the effect of intermediaries or technology increasing price transparency in industries like life insurance ([Brown and Goolsbee, 2002](#)), waste management ([Salz, 2022](#)), health care ([Brown, 2019](#)), and retail gasoline ([Luco, 2019](#)). Many studies assess the impact of online platforms on prices and aggregate consumer welfare, but few examine which consumers benefit most in negotiated price markets. I provide a novel distributional analysis of consumer welfare gains, accounting for endogenous seller bidding participation and pricing.

The remainder of the paper proceeds as follows: the next section details the online platform and provides descriptive statistics. Section 3 develops a model of buyer and seller behavior in the solar PV market, and Section 4 discusses the methods used to estimate the model. 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, often requiring buyers to call installers and schedule site visits. These hassle costs can enable sellers to increase markups. A 2017 National Renewable Energy Laboratory survey found over 80% of solar shoppers contacted two or fewer installers before deciding ([Sigrin et al., 2017](#)). Recently, online platforms have emerged as an alternative to direct buyer-seller interactions.

For example, the U.S.-based platform EnergySage Inc. connects potential solar customers with a network of installers. EnergySage allows households to conduct multi-attribute auctions, where bidders submit proposals that include price and other characteristics like solar panel brand. The buyer selects the winning bidder based on these multi-dimensional bids. In multi-attribute auctions, buyers can choose based on any criteria and are not obligated to select the lowest price. These auctions are similar to scoring auctions, except that the choice rule is not announced to the bidders.

Each EnergySage auction has several stages. First, consumers create an account on the

platform and provide details like the installation address and monthly electricity bill. Second, registered installers receive project notifications, including a Google Maps photo of the buyer's roof (see Appendix Figure A.2) and other details. Installers then submit quotes, which include system price, panel brand, inverter brand, and seller information like star ratings. Buyers can then select a quote or opt not to purchase any offers. Appendix Figure A.2 shows the purchaser's comparison tool on the platform.

Solar PV systems can be paid for with cash, loans, or leases. On EnergySage, 97% of buyers choose to buy with cash or a loan, likely due to the platform's net present value calculations, which favor purchases over leases. This trend reflects a broader market shift away from leases. According to [Barbose et al. \(2021\)](#), the market share for leased systems grew from 2007 to 2012 but fell to 35% by 2020, with further declines expected due to the rise of residential loan products.

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 from 2014 to 2016 in Arizona, California, Colorado, Connecticut, Massachusetts, New York, and Texas. EnergySage collects information on each household's address, average monthly electricity bill, and whether the buyer obtained other off-platform quotes. The auction data is then merged with household-level demographic data from Acxiom's Infobase and ZIP-code level solar environmental benefits estimates from [Sexton et al. \(2021\)](#).⁶ See Appendix B for more details on data processing and sample construction.

Project Characteristics: Appendix Table A.1 provides descriptive statistics of the sample of 10,488 potential installation projects. Households in the sample vary in monthly electricity expenditure, search behavior, solar generation potential, and home valuation. The mean electric bill was \$187/month with a standard deviation of \$89/month, and about 17% of households reported having quotes from off-platform installers. The average project received 3.54 bids, with noticeable variation across projects—a quarter of projects obtained two or fewer bids, while 25% received five or more bids. Appendix Figure A.3 shows that the mean number of

⁶This marketing database compiles household-level data from sources like public records.

bids per project increased while bid prices fell during the sample period in most core-based statistical areas (CBSAs). The figure also illustrates the substantial variation in bids per project and prices across CBSAs. Lastly, Appendix Table A.2 shows the composition of projects across states and time. The number of projects tripled each year, with over 40% in California, reflecting the state's large share of U.S. solar installations.⁷

Although the platform data is highly detailed, a limitation of the data is that platform users may not represent all solar shoppers, potentially affecting external validity. I use income data from Barbose et al. (2018) to compare platform users to solar adopters more broadly. Appendix Figure A.5 shows that the income distribution of platform users is similar to the broader population of solar adopters, though platform users are slightly lower-income. Despite this similarity, there could be other differences, so the reader should interpret my results with this caveat in mind.

An advantage of the data is that there is a large amount of demographic heterogeneity. Therefore, when estimating the demand model, I focus on estimating heterogeneity in preferences across households with different wealth and energy use behaviors. Estimating these heterogeneous preferences allows me to assess which types of households benefit most from using the platform.

Characteristics of Solar Bids: Appendix Table A.3 provides details about the bids submitted. Bids differ in price per watt, system capacity, panel brand, and inverter type. Buyers often get quotes for different panel brands, with EnergySage providing ratings for each make. Premium panels, like those from LG Electronics, have higher efficiency and better warranties, making up 34% of bids. Premium Plus panels, like those from SunPower, are less common and more expensive, comprising less than 4% of bids.

The inverter, converting DC to AC, is another key component. String inverters are the cheapest but can underperform in shade. Microinverters and power optimizers perform better in partial shade but are pricier, comprising 73% of bids.⁸

Seller quality also influences a buyer's choice, and EnergySage provides metrics like star

⁷Appendix Figure A.4 maps the projects, showing a concentration in major metro areas.

⁸Microinverters are identified by the brands Enphase Energy or SolarEdge Technologies.

ratings, experience, and reviews. Although sellers' names are anonymized, each installer is identified in the data by unique IDs. By the end of 2016, 60% of bids came from five-star-rated installers, while 18% came from unrated installers. The average bid was from an installer with nearly 2,400 completed installations, but there is notable variation in experience.⁹

A notable limitation of the data is that sellers using the platform during the sample period were primarily small to mid-sized installers, and therefore, the supply-side estimate may not directly apply to larger national installers.

Panel B of Appendix Table A.3 compares winning and non-winning bids in the Los Angeles area for the first half of 2016. Winning bids had lower prices and were more likely from five-star-rated installers with premium panels and microinverters, indicating the importance of non-price factors.

To compare platform purchases to the broader market, I appended EnergySage data to the Lawrence Berkeley Lab's Tracking the Sun Database. Regressions with zip code and time fixed effects in Appendix Table A.4 reveal that platform prices are about 15% lower, and have similar system sizes to off-platform transactions. Additionally, the regressions show that platform transactions are slightly less likely to include SunPower modules (rated Premium Plus) but more likely to use LG panels (rated Premium) and microinverters. So, while there are some differences in the equipment used, there does not appear to be a systematic bias toward higher-quality equipment or larger systems on the platform.

Finally, a defining feature of this market is that sellers can adjust their participation and pricing strategies based on the observable characteristics of buyers. For instance, installers may choose to forego submitting bids to low-income households or to bid higher prices to households with higher electricity bills. Appendix Section C shows descriptive evidence of selective entry and price discrimination. These patterns motivate the structural model outlined in the next section.

⁹For several installers lacking residential installation experience data, I assign the median experience level from the overall sample.

3 Model

Having discussed the data, I developed a structural model that incorporates heterogeneous buyer preferences and strategic participation and bidding by sellers. 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 the geographic market where the project is located, the time period, and the characteristics of the household. For each project of type \mathbf{z}_i , there is a set $\mathcal{N}(\mathbf{z}_i)$ 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 $\mathcal{J}_i \subset \mathcal{N}(\mathbf{z}_i)$ be the set of sellers that decide to participate in the auction for project i . Buyer i then chooses between the project bids ($j \in \mathcal{J}_i$) and an outside option ($j = 0$) to maximize their utility. The outside option includes the choice to forgo purchasing solar or purchasing from an off-platform installer. Buyer i 's utility from selecting option j is given by:

$$u_{ij} = B_{ij}\alpha_i + \mathbf{x}'_{ij}\beta + \mathbf{w}'_j\gamma + \delta_i + \zeta_{ig} + (1 - \lambda)\varepsilon_{ij}$$

$$\alpha_i = \mathbf{z}_i^{(1)'}\tilde{\alpha}, \quad \delta_i = \mathbf{z}_i^{(2)'}\tilde{\delta}. \quad (1)$$

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 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 options within a group g . 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 leads to the nested logit model, which accommodates correlation in preferences within pre-specified groups. Here, one group is the outside option, and the other includes all platform bids ("the inside option"). Some households may visit the platform out of curiosity about solar prices, while others are committed to selecting an installation bid. The nested logit model can accommodate these unobserved differences in preferences. As λ approaches zero, there is no correlation in preferences for each inside option, reducing to the standard logit model. As λ approaches one, preferences for each "inside option" become perfectly correlated. The overall level of utility is not identified, so I normalize the utility of the outside option to zero plus an error term ($u_{i0} = 0 + \varepsilon_{i0}$).

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.¹⁰ 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

¹⁰I consider the robustness of the results to a discrete-continuous model in Appendix D.

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.

3.2 Supply

The supply-side model involves firms deciding whether to submit bids to potential buyers and at what price. Sellers lack precise information about the number, characteristics, and prices of competing bids, though they can observe the total number of bids for an auction ex-post and know which firms use the platform in their area. Thus, I assume suppliers know the distribution of potential competition 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 competitors entering the auction, the characteristics of those competitors, and the distribution of possible prices those competitors would submit. Additionally, they know the mean utility of the buyer (but not the random component of utility).¹¹ 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 $\eta_{ij} \sim \text{Lognormal}(\mu(\mathbf{z}_i, \mathbf{w}_j), \sigma^2(\mathbf{w}_j))$.¹² The expected bid preparation cost depends on both the project type and the seller's type. I assume the bid preparation cost is i.i.d. across projects and firms conditional on the seller's type and the project type 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

¹¹I assume that sellers know all of the parameters of the buyer's utility function.

¹²The variance parameter allows random shocks to bid preparation costs across auctions.

Low variance means entry costs for firms of type \mathbf{w}_j are relatively stable across auctions.

marginal cost of completing the project c_{ij} .

To keep the model empirically tractable, I assume that non-price characteristics, \mathbf{x}_{ij} , such as panel quality and inverter type, are not strategic choices by bidders. This implies firms cannot systematically change these characteristics in counterfactuals. This assumption is supported by the data, as installers typically use the same equipment for consecutive projects, with changes in hardware likely predetermined by inventory. Appendix Figure A.6 shows that installers use the same non-price characteristics as their previous bid over 74% of the time, with this probability remaining consistent across household demographics and competition levels.¹³

Analogously to the demand side, I assume that system capacity bids are not strategic choices for sellers. Although installers provide a recommended system capacity, there is little variation in capacity offers for a particular project, as installers typically size systems to match the household's annual energy usage. The second column of Appendix Table A.6 supports this assumption by showing that proposed capacities are not influenced by strategic factors. While the number of expected bids affects prices (Column 1), it is statistically unrelated to proposed system capacities (Column 2).

In formulating the sellers' bidding problem, I follow [Levin and Smith \(1994\)](#) and assume that prior to entering an auction, sellers know the project type and the joint distribution from which their marginal cost, non-price characteristics, and system capacity will be drawn, $F_{C X Q | w_j, z_i}$. Notably, the distribution depends on both the seller's type and the project type. However, firms only learn their exact marginal cost, non-price characteristics, and the capacity of the system after deciding to enter an auction. Intuitively, this means installers must exert some effort to prepare a customized bid and only learn about the exact marginal cost of a project by completing the bid preparation process. Although firms do not perfectly know their marginal costs prior to entry, they do know the distribution from which their marginal cost will be drawn. Therefore, sellers will know, for example, that buyers with large electric bills will be more likely to need a large system, but they will not know the exact size and cost

¹³Installers use the same hardware as their previous bid 75% of the time for homes valued below \$1 million, and 72% for homes valued above \$1 million.

of the installation ex-ante. After the firms make their entry decisions in stage one, each firm's marginal cost and non-price characteristics are drawn from $F_{CXQ|w_j, z_i}$ and the installer then decides on a price bid during the second stage.

The assumption that installers do not know their marginal cost prior to entry is not clear-cut. For instance, [Samuelson \(1985\)](#) proposes an alternative model where marginal costs are known before entry, leading to selective auction participation based on private value realizations. In Appendix E, I investigate these competing models and fail to find evidence of selective entry, supporting the [Levin and Smith \(1994\)](#) model where costs are not perfectly known at the entry stage.¹⁴

Price Bidding: It will be helpful first to 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 :

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

Where q_{ij} is the system capacity, B_{ij} is firm j 's per-unit price bid, and c_{ij} is firm j 's marginal cost. $\mathcal{P}_{ij}(B_{ij}|\mathbf{x}_{ij}, \mathbf{w}_j, \mathbf{z}_i)$ is the equilibrium expected probability of winning the auction conditional on placing a bid price of B_{ij} . The equilibrium expected probability of being selected is also a function of the project type \mathbf{z}_i , the seller's type \mathbf{w}_j , and the non-price characteristics of the bid \mathbf{x}_{ij} . We work with expected probabilities because the seller does not know which exact 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

¹⁴Although the [Levin and Smith \(1994\)](#) model fits the data better than [Samuelson \(1985\)](#), other models might improve the fit further.

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 their bid price and non-price characteristics. 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 A.5, 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 statistically significant. As additional evidence, Table A.5 also shows that the "order of bid" is not associated with a change in the probability of that bid being selected. These regressions suggest firms are not making significant changes in bidding strategy based on the order they submit a bid, justifying the assumption of simultaneous bidding.¹⁵

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

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

¹⁵For tractability, I also assume that bids submitted by sellers are binding. In rare cases, sellers will alter bid characteristics at a later stage (e.g., change the panel brand due to an inventory stock out), but such changes must be approved by the buyer. In such cases, only the final bid characteristics are observed. This is an inherent limitation of the data, and there exists a possibility that some bias exists in a small subset of the choice data due to these unobserved bid changes.

Recall that $Prob_{ij}$ is the probability that buyer i selects firm j 's bid conditional on a 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 .¹⁶ Since each firm's entry draw and marginal cost draw are 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:

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

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 \mathbf{x}_{il} .

Firm j 's first-order condition (FOC) for an optimal bid is given by:

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

Given a vector of non-price characteristics, the optimal bid function $B_{ij}^*(c_{ij}|\mathbf{x}_{ij}, \mathbf{w}_j, \mathbf{z}_i)$ 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

¹⁶ G is only a function of \mathbf{z}_i because a seller's type \mathbf{w}_j and non-price characteristics \mathbf{x}_{ij} are private information at the time of bidding.

expected profits conditional on entering the auction for project i can be expressed as follows:

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

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:

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

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

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

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

3.3 Equilibrium

For each seller j , a strategy consists of a participation function $\mathbf{w} \times \mathbf{z} \times \mathbb{R}_+ \rightarrow 0, 1$ and a bidding function $\mathbf{w} \times \mathbf{z} \times \mathbf{x} \times \mathbb{R}_+ \rightarrow \mathbb{R}_+$. Sellers decide to enter based on project type, seller type, and the entry cost shock, then bid using project type, seller type, marginal cost draw, and non-price characteristics. I follow the literature's focus on type-symmetric pure strategy Bayesian equilibrium (Krasnokutskaya et al., 2020), where sellers of the same type and characteristics use the same strategies. Participation equilibrium occurs when sellers satisfy Inequality 7, and bidding equilibrium when Equation 5 is met. Krasnokutskaya et al. (2020) prove the existence of this equilibrium, though uniqueness in participation is not guaranteed. The next section details the estimation procedure.

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 bidder's 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.

4.1 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{x} , variables that characterize the non-price attributes of each participating installer's bid; and (4) $\mathbf{z}^{(2)}$, variables that shift the buyer's preference for all of the installation bids (i.e., shifts the likelihood of picking the outside option).

The bid price that enters buyers' utility for option j is measured in dollars per watt and is scaled to 70% of each installer's gross bid price to account for the 30% Investment Tax Credit (ITC). I refer to this after-incentive price as the "net price".¹⁷ 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 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, sellers are unlikely to observe household income. 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.

The utility for each option is also a function of several non-price characteristics: dummies

¹⁷This specification assumes buyers value a \$1 reduction in gross price the same as a \$1 increase in tax credit, consistent with the literature (e.g., [Langer and Lemoine, 2022](#)).

for premium and premium plus panel categories, a microinverter dummy, 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 the bids in the sample. In Appendix D, I discuss the robustness of the results to changes in the utility specification.

Outside Option: The outside option in this setting encompasses the possibility of either purchasing through an installer outside of the platform or not installing solar PV. As such, the net utility that the buyer obtains from choosing the outside option may also be affected by any hassle costs of contacting sellers directly. Moreover, the value of the outside option can also depend on local electricity prices, off-platform installation prices, and policies such as net energy metering, which all dictate the value of purchasing solar PV through the platform. Thus, 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 in the utility function.¹⁸ More specifically, I interact each of these fixed effects with an indicator variable for whether the choice represents a bid made through the platform (e.g., the inside options). Thus, the associated coefficients allow the relative value of each bid (relative to the outside option) to vary flexibly across households. Importantly, the CBSA fixed effects also control for any state or local incentives that could change the value of solar PV, including net metering rules that affect households’ electricity bills or up-front government rebates.¹⁹

In general, I cannot directly observe solar installation purchases made off of the platform. However, the platform did conduct exit surveys that asked potential buyers about off-platform purchase choices. I leverage these survey data to estimate probabilities of installing solar conditional on selecting the outside option for some of the counterfactual analyses.

¹⁸Projects outside a CBSA or in CBSAs with fewer than 100 projects are grouped into a new state-specific category. For instance, a project in Yuma, AZ, would be classified as “Other, AZ.”

¹⁹Most local incentives remain constant over the sample period, but those in New York and Connecticut vary. In Section 4.5, I exclude these states as a robustness check.

4.2 Markup and Marginal Cost Estimation

After estimating demand, I calculate the markup embedded in each bid using the utility function estimates to form each firm's FOC for an optimal bid from Equation 5. Since the FOC involves expectations that cannot be expressed in a closed form, I evaluate these expectations by simulating the firm's probability of winning across different realizations of competitor sets and bid prices unknown to the installer at the time of bidding. The procedure is as follows:

1. *Estimate Entry Probabilities*: Calculate non-parametric entry probabilities for each project-seller type pair, $\mathbf{z}_i, \mathbf{w}_j$, as the ratio of auctions entered to total auctions for that type pair. Assume a seller is a potential entrant for auction i if they entered at least one auction of type \mathbf{z}_i .
2. *Simulate Entry Decisions*: Use the probabilities from (1) to simulate entry decisions for auction i by potential entrants in $\mathcal{N}(\mathbf{z}_i)$.
3. *Draw bids*: Draw price bids and non-price characteristics for each simulated entrant using the empirical distribution from similar auctions. If a type \mathbf{w}_j seller enters a simulated auction of type \mathbf{z}_i ; then draw a bid vector $\{B_{ij}, \mathbf{x}_{ij}\}$ from the pool of all bids placed by type \mathbf{w}_j sellers in auctions of type \mathbf{z}_i .
4. *Evaluate Choice Probabilities*: Calculate the buyer choice probabilities $Prob_{ij}$ and demand semi-elasticities $\frac{\partial Prob_{ij}}{\partial B_{ij}}$ using the simulated bids and competitor characteristics.
5. *Average Simulations*: Repeat the previous steps S times,²⁰ averaging the choice probabilities and demand semi-elasticities to obtain estimates for the two expectations in the FOC. Let s denote the simulation iteration, then the expressions are:

$$\widehat{\mathcal{P}}_{ij} = \frac{1}{S} \sum_{s=1}^S Prob_{ij}^s, \quad \frac{\partial \widehat{\mathcal{P}}_{ij}}{\partial B_{ij}} = \frac{1}{S} \sum_{s=1}^S \frac{\partial Prob_{ij}^s}{\partial B_{ij}}. \quad (9)$$

²⁰I simulate 100 iterations of each auction type.

6. *Calculate Marginal Costs:* Use the averaged choice probabilities and averaged demand semi-elasticities to compute the markup and marginal costs for each bid via the FOC:

$$\widehat{c}_{ij} = B_{ij} + \frac{\widehat{\mathcal{P}}_{ij}}{\frac{\partial \mathcal{P}_{ij}}{\partial B_{ij}}}. \quad (10)$$

This algorithm provides a unique marginal cost estimate for every bid, which I then use to form a non-parametric cost distribution for each seller-project type pair.

Defining project categories is crucial for credible markup estimates, balancing bias and variance. Too few categories can bias results if projects are heterogeneous, while too many can increase variance due to limited data. To avoid bias, I chose a high-dimensional project-type definition based on five variables: the CBSA, half-year, home market value quartile, household electricity expenditure quartile, and whether the buyer has off-platform quotes.

4.3 Entry Cost Estimation

In the final step, I use the estimated marginal costs to form each firm's *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:

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

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:

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

Where $\mathbb{1}[j \text{ enters } i]$ 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 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.

4.4 Identification and Modeling Assumptions

To consistently estimate the demand parameters, it is necessary that all variables in the utility function, including bid price, are uncorrelated with the preference shocks ε_{ij} and ζ_{ig} . Additionally, the number of bids each buyer receives must be uncorrelated with these shocks. My identification strategy leverages household-level demographic data, including flexible fixed effects to control for household heterogeneity and a rich set of fixed effects to control for bid quality. Specifically, I use panel brand quality controls, inverter type controls, installer star-rating dummies, installer experience controls, and installer fixed effects. To account for buyer heterogeneity, I include electricity bill quartile dummies, a control for off-platform bids, CBSA fixed effects, and half-year fixed effects.

The price coefficients, α_i , are identified by variation in quoted prices after controlling for seller quality, hardware type, and common demand shocks across markets and time. This variation arises from differences in marginal cost conditional on buyer and seller type, such as input cost shocks due to roofing material differences. Despite the rich set of controls, there could be omitted variables correlated with both prices and unobserved preference shocks, which I address in detail in Appendix Section D.

With demand parameters identified, marginal costs are identified by Equation 10, which maps bid prices to costs. Entry cost parameters (μ and σ) are identified by variation in expected marginal profits, holding seller type constant. Identification of entry costs requires an exclusion restriction—an observed variable that affects the seller’s ex-ante expected marginal profit before the entry decision but does not affect the entry cost distribution (Krasnokutskaya and Seim, 2011). 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. This assumption is plausible if installers use their history of past bids to streamline new bids, making preparation costs constant across household demographics.

Having discussed the identification of the model parameters, I next highlight several critical assumptions that underlie my modeling approach. First, my model employs a static demand framework. Although dynamic models are commonly used for the solar PV market, there are several reasons why my *static demand assumption* may be reasonable in this context. Previous studies (e.g. Bollinger and Gillingham, 2019; De Groot and Verboven, 2019; Langer and Lemoine, 2022) used data when solar subsidies were decreasing, making dynamic incentives crucial. During my sample period, most US states had stable subsidies, with exceptions in Connecticut and New York, which I exclude in robustness tests. Additionally, Appendix Figure A.7 shows that pre-incentive solar prices had stabilized by my sample period, suggesting a relatively low incentive for consumers to delay investment for lower future prices.

Second, the model assumes that cost shocks are independent and identically distributed (i.i.d.) conditional on project and installer type. If unobserved factors, such as roofing type, induce correlation in cost shocks among installers, the *i.i.d. cost assumption* would be violated, leading to upwardly biased markup estimates. Appendix G provides evidence of the plausibility of the i.i.d. assumption.

Third, the demand model assumes consumers have full information about all alternatives, including the outside option. This implies that consumers know the prices and characteristics of off-platform installers. Under the *full information assumption*, the value of the outside option can incorporate additional hassle costs of off-platform procurement (e.g., time spent on

contacting sellers), but it requires that consumers are aware of the utility they would receive from choosing an off-platform installer net of hassle costs. This assumption would be violated if, for example, consumer demand is characterized by a more complex search model in which consumers choose whether to undertake a costly effort to learn about offline installers. In such a model, consumer search effort may change endogenously in counterfactuals and bias the counterfactual predictions.

Lastly, the model treats sellers' entry and bidding decisions independently, ignoring joint decisions for simultaneously posted projects. Since only 13% of project postings occur in markets with more than 20 other projects weekly, this simplification is reasonable. However, if simultaneous postings influence seller entry, then my model will overestimate the variance of the marginal cost distribution. With these assumptions in mind, I now transition to discussing the model estimates in the following subsection.

4.5 Model Estimation Results

Demand Estimates: The top portion of Table 1 Panel A contains the estimates for the baseline utility specification. As expected, the base price coefficient is negative and relatively large in magnitude. We also see that households in the top three home value quartiles are relatively less price-sensitive. However, 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 the bottom quartile is the most elastic (mean own-price elasticity of -2.43) 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 expenditures 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 from the platform 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.²¹ This result is consistent with off-platform quotes being positively correlated with preferences for rooftop solar and also that quotes on the platform tend to dominate off-platform installation options.

On the top right side of the table, we see that households value hardware characteristics, preferring higher-rated panel brands and systems with microinverters. We also see that buyers value installer characteristics. All else equal, sellers with a 4.5-star rating and those with ratings below 4.5 stars are less likely to be selected relative to an installer without ratings. On the other hand, installers with a five-star rating are preferred to all other 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 who has completed over 1000 installations compared to an installer who has completed less than 100 installations.

As discussed in Section 4.4, there are several threats to the identification of the utility parameters. Appendix Section D investigates the robustness of the coefficient estimates to adding a series of potential confounding variables into the utility specification. The robustness checks include adding controls for installer experience, panel brand fixed effects, household demographics, house characteristics, a discrete-continuous choice specification, and excluding auctions from states with time-varying subsidies. The estimates are relatively stable across these alternative specifications.

Cost Estimates: Table 1 Panel B 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.

Although the estimates indicate that markups represent a substantial portion of installation prices, it is important to recognize that the estimated markups are gross of installers' fixed operating expenditures. These fixed expenditures include marketing and overhead for

²¹In the raw data, buyers with off-platform quotes are 33% more likely to choose a platform bid than those without off-platform quotes.

facilities and administration. [Fu et al. \(2016\)](#) calculate that installers should pay an average of \$0.33 per installed watt for marketing expenses and \$0.37 per watt for other overhead costs in 2016. Taking these estimates as given would imply that installers on the platform earn an average net profit margin of \$0.85 per watt—or 24.8% of installation prices in 2016.

Appendix Figure A.8a demonstrates that the estimated marginal cost reductions are explained mainly by the fall in wholesale solar PV hardware prices reported by Bloomberg Inc. Figure A.8a 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%. Appendix Figure A.8b 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 than the cost reported by SolarCity/Tesla but marginally lower than the costs reported by Sunrun and Vivint ([Fu et al., 2016](#)). 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.

I use the entry cost parameters, reported in Appendix Table A.7, to calculate the expected bid preparation cost (entry cost) conditional on participating. Table 1 Panel B shows that bid preparation costs fell from \$40 in 2014 H1 to under \$10 in 2016 H2. Appendix Figure A.7 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. Appendix F further investigates cost heterogeneity across installer, hardware, and household characteristics.

5 Welfare and Counterfactuals

In this section, I use the model estimates to evaluate consumers' welfare gains from using the platform. I then investigate the mechanisms driving the welfare results by simulating

counterfactuals that measure the effect of increasing the number of bids per project on market outcomes. Finally, I solve a second set of counterfactuals to evaluate the impacts of solar subsidies.

5.1 Consumer Surplus Gains From Platform Access

I evaluate the extent to which the online platform improves consumers' expected welfare relative to their outside options using the following expression:

$$\mathbb{E}[CS] = \frac{1}{M} \sum_i \frac{1}{\alpha_i} \mathbb{E}[\max_{j \in \mathcal{J}_i \cup 0} u_{ij} - u_{i0}] = \frac{1}{M} \sum_i \frac{1}{\alpha_i} \log \left(1 + \left(\sum_{j \in \mathcal{J}_i} \exp\left(\frac{v_{ij}}{1-\lambda}\right) \right)^{1-\lambda} \right) (\bar{q}_i). \quad (13)$$

Here, $\mathbb{E}[CS]$ is the expected consumer surplus that buyers obtain from their menu of options on the platform (\mathcal{J}_i) relative to their outside option. The last equality derives from the nested logit utility specification, where v_{ij} is the utility buyer i obtains from option j net of the taste shock, $v_{ij} = u_{ij} - \varepsilon_{ij}$. Lastly, recall that utility is measured in dollars per watt, so I scale the consumer surplus measure by each buyer's expected system size, \bar{q}_i , to obtain the expected *total* consumer surplus.

The above expression measures the value that consumers glean from the platform *relative to their outside option*. Thus, under the *full-information assumption*, the above estimate captures the added value the platform provides to buyers as an alternative means of procuring solar PV.

There are a number of caveats to keep in mind when interpreting the consumer surplus estimates. First, I estimate preferences among a selected sample of buyers who opted to shop for solar through an online website. Although we saw in Section 2 that these buyers appear relatively similar to the broader population of solar adopters— it's possible that these estimates relate to a sophisticated consumer segment that can elicit more value from using the online platform. Second, the consumer surplus measure is implicitly a function of the quality and

pricing offered by installers off of the platform.²² Even though the platform composed a small fraction of the overall solar installation market during the sample period, it is still plausible that offline installation price and quality offerings are enhanced in response to the competition from the platform. If so, these consumer surplus estimates would tend to understate the actual gains created by the platform because the calculation assumes that the value of the outside option remains fixed when the platform is removed (instead of decreasing). Third, it is possible that in a scenario where the platform was not available, buyers could contact the exact sellers that made a bid to them through the platform, therefore increasing the value of the outside option. However, most markets in the sample have dozens or even hundreds of operating solar installers, so adding a few additional sellers to the pool of possible off-platform sellers is unlikely to have a substantial impact on a buyer's value from the outside option.

Given the above caveats, Table 2 Panel A shows that the platform increased consumer surplus by \$1,016 for the average household that shopped on the platform. For reference, this welfare gain is equivalent to making a payment of 6% of the mean bid price to each individual who used the platform. 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 \$754 in benefits from the platform relative to \$1,080, \$1,164, and \$1,075 for the top three quartiles, respectively. A similar pattern emerges in the distribution of gains across home values. The lowest-valued quartile of homes sees gains of \$809 relative to \$1,030 for the most expensive quartile of homes.

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 1) 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

²²Without the platform, sellers might increase door-to-door sales or other marketing.

more distinct installers. The distribution of welfare gains in Table 2 Panel B and C, therefore, suggests a high value of obtaining additional installer choices.

5.2 Counterfactuals: Effects of Expanding Buyers' Choice Sets

The results presented in the previous section suggest that platforms can improve consumer outcomes in the residential solar market. However, the success of platforms going forward will hinge on attracting both sellers and buyers to participate. A platform will only reduce margins and prices if there are a sufficient number of sellers bidding through the platform. This raises several additional questions. First, how much of the platform welfare benefit can be explained by the buyers obtaining additional bids? Relatedly, how much does each incremental bid per project reduce solar installation prices? How much does solar adoption and associated environmental benefits increase if buyers obtain more bids? These questions are important for either policymakers or private companies that seek design platforms that reduce solar installation prices.

To investigate the above questions, I measure the effect of expanding buyers' choice sets while 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 on 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 provide additional details about the algorithm for solving counterfactual equilibria in Appendix H.1.

Table 3 displays the outcomes of changing the number of bids from one to 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 a 16.4% reduction in mean bid prices, a 25.2% reduction in the lowest bid, and a 15.5% fall in purchase prices (selected bids). These results imply a \$4,000 gross 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 bid price reduction (12% from the baseline) relative to adding

a fifth bid (0.8% from the baseline).

Incremental bids can reduce prices through two channels: (1) reducing markups through increased competition and (2) reducing installation costs by connecting buyers to sellers with lower installation costs. The counterfactual results indicate that the majority of the price decline from incremental bids comes from reduced markups. More specifically, markups fall by over 40% when the number of bids increases from one to five, whereas marginal costs fall by only 2.5%.

In the second section of the table, I document changes in consumer and producer surplus as the number of bids per project expands. The estimates confirm that increasing the number of bids can deliver large consumer benefits. More specifically, providing a buyer with five bids causes a 360% increase in buyers' surplus from using the platform relative to the single bid case. Intuitively, buyers collect a relatively small benefit from using the platform if they receive only a single bid. Conversely, the platform can deliver substantial value for consumers when several sellers submit unique bids.

We also see that total producer surplus for on-platform installers increases over twofold if the number of bids increases 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 of bids increases. Despite the increase in total surplus among platform sellers, the expected profit for each individual seller falls by over 50% as the number of bidders increases from one to five. Notably, 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 and do not account for changes that would accrue to installers operating off of the platform.

The third part of the table shows the impact of bid quantity on the number of solar PV installations. Evaluating counterfactual changes in solar adoption is especially challenging because I do not directly observe or model off-platform solar purchases. 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. Therefore, I use data from platform exit surveys to predict the probability that buyers will adopt so-

lar PV, conditional on selecting the outside option. Appendix Table A.8 shows that 27% of survey respondents elect to hire an off-platform installer conditional on choosing the outside option, and 73% chose not to install solar.²³ Additionally, buyers who collected quotes from off-platform sellers prior to arriving on the platform are much more likely to ultimately make off-platform purchases. As such, I specify the conditional probability of hiring an off-platform installer separately for each group of buyers—51.7% for households that already had off-platform quotes prior to using the platform and 19.7% for those who did not previously hold external quotes.

There are a few potential complications of using the exit survey data to predict off-platform purchase choices. For one, the survey represents a selected subset of the sample that may not be representative of the full platform buyer population. To characterize selective participation in the exit survey, I regress an indicator variable for whether each buyer completed the survey on a set of household characteristics in Appendix Table A.9. The regressions indicate that households with higher home values, high (or moderate) electricity usage, and those who receive more bids through the platform are the most likely to complete the survey. These selection patterns suggest that using the exit survey may provide an upper-bound estimate of the full sample probability that potential buyers hire off-platform installers conditionally on choosing the outside option. A second possible limitation of using the off-platform purchase probabilities estimates above is that off-platform purchase behavior may vary with the number of bids that buyers obtain. For instance, buyers who receive fewer bids through the platform may search more off of the platform. To ease this concern, Appendix Figure A.10 illustrates that the off-platform purchase probability does not change systematically with the number of (platform) bids. Moreover, Appendix Table A.10 shows that there is no statistically significant relationship between the number of bids and buyers' off-platform purchases across several regression specifications.

Acknowledging the potential challenges and additional assumptions required to estimate solar adoption effects, I find that expanding the bid set from one to five bids leads to at least a 21% increase in the number of households that adopt solar. I also use estimates from [Sexton](#)

²³The exit survey was completed by 11% of households that selected the outside option.

et al. (2021) to calculate the net present value from avoided pollution damages resulting from these additional solar purchases. I find that expanding buyers' choice sets from one to five bids leads to at least an 18% increase in external benefits via pollution reductions.

5.3 Counterfactuals: Market Impacts of Solar Subsidies

In the next set of counterfactuals, I use the model to assess the impact of government subsidies 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 Solar 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. The history of the ITC has been marked with many contended extensions, planned step-downs, and other proposed policy alterations. The solar industry, environmentalists, and politicians have argued that renewable energy subsidies are critical to addressing climate change. On the other hand, opponents have countered that the tax credits are not cost-effective and place too large a burden on taxpayers.

To understand the effects of removing the ITC, I simulate counterfactual market outcomes both with and without the 30% ITC subsidy. In the counterfactual with the ITC, winning sellers earn 100% of the gross bid price, but consumers pay only 70% of the gross price. For the counterfactual without the ITC, buyers pay the entire gross bid price to selected sellers. In both counterfactual simulations, sellers update both participation and bidding strategies to account for the subsidy availability (or removal).²⁴ Importantly, the counterfactual removal of the ITC may also affect the value that buyers obtain from choosing the outside option through changes in off-platform prices. To account for these changes, I solve two counterfactual specifications that remove the ITC. In the first specification, I hold the normalized value of the outside option fixed ($v_{i0} = 0$). This specification corresponds to a 0% subsidy pass-through rate among installation prices outside the platform. In the second specification, I adjust the value of the outside option to characterize a 100% subsidy pass-through rate to off-platform

²⁴The demand model includes CBSA fixed effects to control for local incentives, so the counterfactuals reflect changes from removing the ITC while holding local policies fixed.

prices. Specifically, I modify the utility of the outside option as $v_{i0} = 0 - 0.3 \alpha_i * \bar{p}_{it}$, where \bar{p}_{it} is the mean gross installation price among *all* installations in buyer i 's local CBSA market at time t using Lawrence Berkeley National Lab data which includes off-platform transactions.²⁵ I discuss the algorithm used to solve the counterfactual equilibria in Appendix H.2.

Panel A of Table 4 shows average market outcomes for each of the counterfactuals. The first column shows results for the case where the ITC is available. In contrast, the second and third columns contain the results with the ITC removed—assuming a 0% off-platform pass-through rate and full pass-through, respectively. Notably, all of the outcomes in both the second and third columns are very similar, suggesting that the counterfactual results are relatively robust to altering the assumption about how removing the ITC affects the outside option. Focusing on the results in the first and third columns, we see that eliminating the ITC results in an 8% reduction in the number of bids per project from 3.49 to 3.24. 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 25.6% increase in the net price of solar bids from \$2.47/watt to \$3.10/watt. This net price increase of \$0.63/watt implies a \$4,600 increase on a typical-sized system. In the same fashion, the mean price of selected bids also increases by 23.7% with the elimination of the ITC. 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%.²⁶ All together, these price changes, along with reductions in seller participation, cause a 32% reduction in the number of purchases made through the platform absent the ITC.

In the lower part of Table 4—Panel B—I calculate the average change in various welfare outcomes per dollar of subsidy expenditures under the ITC. In particular, I estimate these values as $(\bar{Y}_{ITC} - \bar{Y}_{No}) / \overline{Subsidy}_{ITC}$, where the numerator is the change in mean outcomes

²⁵I use data from Barbose et al. (2021) to calculate average per-watt prices for each CBSA by half-year. For CBSAs missing transaction prices, I set \bar{p}_{it} as the statewide average.

²⁶The ITC covers 30% of the gross price, so removing it raises the net price by $10/7$ (42.8%).

between the “ITC” and “No ITC” counterfactuals and the denominator is the mean subsidy expenditure (per potential buyer) when the subsidy is available. I calculate the above measures separately for each of the two counterfactual specifications removing the ITC. Again, focusing on the last column, we see that the ITC fosters substantial welfare increases for both producers and consumers using the platform. The ITC increases average consumer surplus by \$0.55 per subsidy dollar and also raises sellers’ gross marginal profits by \$0.72 per subsidy dollar. On the other hand, the ITC slightly increases bid preparation costs by \$0.03 per subsidy dollar by encouraging more bids. As an additional benefit, the ITC reduces pollution damage by at least 11.7¢ for each dollar of subsidy expenditure.²⁷ Summing all of these effects together and subtracting the \$1.00 subsidy cost, we obtain a net benefit of \$0.35 per subsidy-dollar ($\$0.55 + \$0.72 - \$0.03 + \$0.11 - \$1.00 = \0.35). Notably, the ITC provides sizable welfare gains for producers and consumers on the platform by reducing deadweight loss caused by market power in addition to correcting for environmental externalities. In particular, we have seen that solar installations are priced well above marginal cost. Therefore, the results suggest that a subsidy can promote efficiency gains by encouraging more transactions on the platform where buyers’ willingness to pay exceeds the seller’s marginal cost.

In addition to understanding the average welfare impacts of residential solar subsidies, an important question for policymakers is to determine where subsidies are likely to deliver the largest welfare gains. Table A.11 in the Appendix shows the five CBSAs in which platform users obtain the largest welfare gains and smallest welfare gains from the ITC per dollar of subsidy expenditure. The top CBSAs are located in Arizona, Texas, and Colorado. Platform users in these states obtain large gains from subsidies for two main reasons: First, these areas all have relatively few competing installers, which enables high markups. Additionally, Arizona, Colorado, and Texas all have abundant solar resources, so adding solar installations will generate more electricity and provide more external benefits from pollution reductions. On the other hand, the ITC subsidy provides more modest welfare gains in Connecticut and Massachusetts because there is stronger competition among installers and fewer solar resources

²⁷Pollution damages depend on the probability that buyers choosing the outside option hire off-platform installers. I calibrate these probabilities based on exit survey data.

in these areas. In sum, the results suggest that residential solar subsidies may deliver higher welfare gains in markets where external environmental benefits are larger but also in markets with higher markups.

There are several limitations and caveats to keep in mind when interpreting the results of the ITC subsidy analysis. First, the analysis above does not capture changes in producer surplus that accrue to off-platform sellers. For example, if buyers substitute from on-platform sellers to off-platform sellers in the counterfactuals without the ITC, the producer surplus estimates above may overstate the aggregate producer surplus benefits of the ITC. On the other hand, if removing the ITC leads some buyers who previously chose off-platform installers to not adopt solar PV, then the estimates may understate the aggregate producer surplus benefits of the ITC. Second, I do not consider several potentially important externalities of solar installations such as electricity transmission and distribution costs (Feger et al., 2022), learning-by-doing (Bollinger and Gillingham, 2019), peer effects (Bollinger and Gillingham, 2012), or generation reserve costs (Gowrisankaran et al., 2016) in the welfare calculations. Third, 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. Buyer preferences and the intensity of bidding competition may be substantially different for projects that are remote from the platform. For this reason, we would expect equilibrium outcomes to respond differently to the ITC availability, yielding different quantitative welfare effects. Thus, we should be cautious about extrapolating the exact welfare estimates out of sample. Nonetheless, if competition for off-platform solar projects is less intense than the competition for projects on the platform—and therefore subject to higher markups—we would also expect substantial net benefits of subsidizing off-platform projects to constrain market power. Finally, the ITC also provides a tax incentive for utility-scale solar projects. Market structure in the utility-scale market is likely to differ from that of the residential market. 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 welfare gains for consumers. A primary mechanism through which a platform can help consumers is by providing them access to additional bids, thereby lowering bid prices.

Reducing installation prices is 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 developing their own platforms to link buyers and sellers. Platforms could also yield similar benefits in other markets characterized by large hassle 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. 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, and 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 investigates the removal of a prominent investment subsidy in the solar PV market and highlights that subsidies have the potential to improve welfare by constraining market power in addition to reducing pollution damages. These results have implications for tax and subsidy policy in other imperfectly-competitive markets.

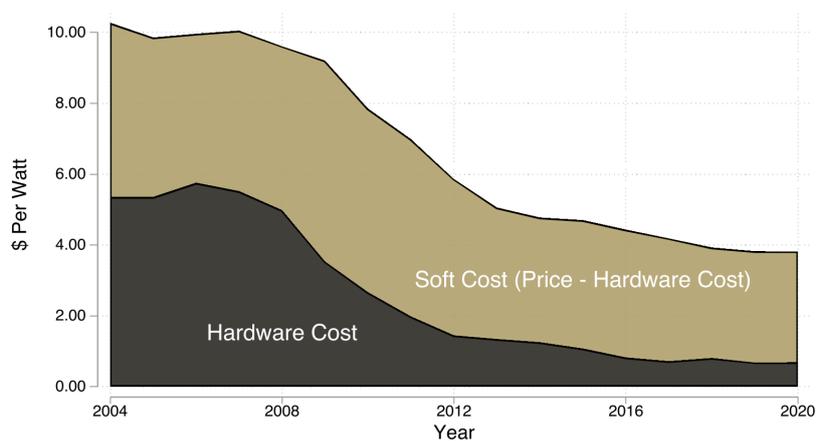
References

- G. Barbose, N. Darghouth, B. Hoen, and R. Wiser. Income trends of residential pv adopters: An analysis of household-level income estimates. Technical report, Lawrence Berkeley National Lab (LBNL), 2018.
- G. L. Barbose, N. R. Darghouth, E. O’Shaughnessy, and S. Forrester. Tracking the sun: Pricing and design trends for distributed photovoltaic systems in the united states (2021 edition). Technical report, Lawrence Berkeley National Lab (LBNL), 2021.
- M. R. Baye and J. Morgan. Information gatekeepers on the internet and the competitiveness of homogeneous product markets. *American Economic Review*, 91(3):454–474, 2001.
- B. Bollinger and K. Gillingham. Peer effects in the diffusion of solar photovoltaic panels. *Marketing Science*, 31(6):900–912, 2012.
- B. Bollinger and K. Gillingham. Learning-by-doing in solar photovoltaic installations. *Available at SSRN 2342406*, 2019.
- S. Borenstein. 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, 2017.
- J. R. Brown and A. Goolsbee. Does the Internet make markets more competitive? Evidence from the life insurance industry. *Journal of Political Economy*, 110(3):481–507, 2002.
- Z. Y. Brown. An Empirical Model of Price Transparency and Markups in Health Care. Technical report, Working Paper, 2019.
- C. Caceres. Analyzing the effects of financial and housing wealth on consumption using micro data. Technical report, International Monetary Fund, 2019.
- O. De Groote and F. Verboven. Subsidies and time discounting in new technology adoption: Evidence from solar photovoltaic systems. *American Economic Review*, 109(6):2137–72, 2019.

- F. Feger, N. Pavanini, and D. Radulescu. Welfare and redistribution in residential electricity markets with solar power. *The Review of Economic Studies*, 89(6):3267–3302, 2022.
- R. Fu, D. Chung, T. Lowder, D. Feldman, K. Ardani, and R. Margolis. Us solar photovoltaic system cost benchmark: Q1 2016. Technical report, National Renewable Energy Lab.(NREL), 2016.
- T. D. Gerarden. Demanding innovation: The impact of consumer subsidies on solar panel production costs. *Management Science*, 2023.
- K. Gillingham and T. Tsvetanov. Hurdles and steps: Estimating demand for solar photovoltaics. *Quantitative Economics*, 10(1):275–310, 2019.
- K. Gillingham, H. Deng, R. Wiser, N. Darghouth, G. Nemet, G. Barbose, V. Rai, C. G. Dong, and others. Deconstructing solar photovoltaic pricing. *The Energy Journal*, 37(3), 2016.
- G. Gowrisankaran, S. S. Reynolds, and M. Samano. Intermittency and the value of renewable energy. *Journal of Political Economy*, 124(4):1187–1234, 2016.
- J. E. Hughes and M. Podolefsky. Getting green with solar subsidies: Evidence from the California solar initiative. *Journal of the Association of Environmental and Resource Economists*, 2(2):235–275, 2015.
- K. L. Judd. Capital-income taxation with imperfect competition. *American Economic Review*, 92(2):417–421, 2002.
- E. Krasnokutskaya and K. Seim. Bid preference programs and participation in highway procurement auctions. *The American Economic Review*, 101(6):2653–2686, 2011.
- E. Krasnokutskaya, K. Song, and X. Tang. The role of quality in internet service markets. *Journal of Political Economy*, 128(1):75–117, 2020.
- A. Langer and D. Lemoine. Designing dynamic subsidies to spur adoption of new technologies. *Journal of the Association of Environmental and Resource Economists*, 9(6):1197–1234, 2022.

- D. Levin and J. L. Smith. Equilibrium in auctions with entry. *The American Economic Review*, pages 585–599, 1994.
- F. Luco. Who benefits from information disclosure? the case of retail gasoline. *American Economic Journal: Microeconomics*, 11(2):277–305, 2019.
- E. O’Shaughnessy, G. F. Nemet, J. Pless, and R. Margolis. Addressing the soft cost challenge in us small-scale solar pv system pricing. *Energy Policy*, 134:110956, 2019.
- E. O’Shaughnessy and R. Margolis. The value of price transparency in residential solar photovoltaic markets. *Energy Policy*, 117:406–412, 2018.
- J. Pless and A. A. van Benthem. Pass-through as a test for market power: An application to solar subsidies. *American Economic Journal: Applied Economics*, 11(4):367–401, 2019.
- J. Robinson. *The economics of imperfect competition*. Springer, 1933.
- D. B. Rubin. The bayesian bootstrap. *The annals of statistics*, pages 130–134, 1981.
- T. Salz. Intermediation and competition in search markets: An empirical case study. *Journal of Political Economy*, 130(2):310–345, 2022.
- W. F. Samuelson. Competitive bidding with entry costs. *Economics letters*, 17(1-2):53–57, 1985.
- S. Sexton, A. J. Kirkpatrick, R. I. Harris, and N. Z. Muller. Heterogeneous solar capacity benefits, appropriability, and the costs of suboptimal siting. *Journal of the Association of Environmental and Resource Economists*, 8(6):1209–1244, 2021.
- B. Sigrin et al. Understanding the evolution of customer motivations and adoption barriers in residential solar markets: survey data. Technical report, National Renewable Energy Laboratory (NREL), 2017.
- E. G. Weyl and M. Fabinger. Pass-through as an economic tool: Principles of incidence under imperfect competition. *Journal of Political Economy*, 121(3):528–583, 2013.
- H. Yoganarasimhan. Estimation of beauty contest auctions. *Marketing Science*, 35(1):27–54, 2015.

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



Notes: Figure constructed using data from [Barbose et al. \(2021\)](#). Each line shows the median cost in dollars per watt for residential solar installations in the U.S. Hardware costs include the panel, inverter, support structures, and electrical wiring.

Table 1: Demand and Cost Estimates

Panel A: Demand Estimates					
λ		0.372 (0.063)	Premium Panel		0.562 (0.076)
Price Coefficients			Premium Plus Panel		1.429 (0.150)
Constant		-0.717 (0.111)	Microinverter		0.370 (0.079)
× Home Mkt. Value - Quartile 2		0.068 (0.044)	Installer Attributes		
× Home Mkt. Value - Quartile 3		0.079 (0.051)	Star Rating ≤ 4		-0.260 (0.256)
× Home Mkt. Value - Quartile 4		0.089 (0.059)	Star Rating = 4.5		-0.409 (0.246)
Household Attributes (× Inside Option)			Star Rating = 5		0.281 (0.096)
Electric Bill - Quartile 2		-0.067 (0.098)	Installs Completed: 100-1000		0.697 (0.237)
Electric Bill - Quartile 3		-0.241 (0.102)	Installs Completed: >1000		0.749 (0.244)
Electric Bill - Quartile 4		-0.584 (0.111)			
Has Off-Platform Quotes		0.368 (0.092)			
Home Mkt. Value Quartile		Mean Own-Price Elasticity			Avg Number of Bids
1		-2.43			3.1
2		-2.25			3.47
3		-2.26			3.84
4		-2.26			3.86
Fixed Effects			Log Likelihood		
CBSA Fixed Effects	Yes		-3823.542		
Half-Year Fixed Effects	Yes				
Permanent Installer Fixed Effects	Yes				
Panel B: Mean Marginal Cost, Markups, Entry Costs					
Half Year	# of Bids	MC (\$/watt)	Markup (\$/watt)	Lerner Index	Bid Prep Cost (\$)
2014 H1	2.07 (0.07)	2.19 (0.34)	1.76 (0.34)	0.46 (0.09)	39.57 (10.32)
2014 H2	2.34 (0.06)	2.14 (0.34)	1.74 (0.34)	0.46 (0.09)	62.38 (13.15)
2015 H1	2.54 (0.04)	2.15 (0.33)	1.69 (0.33)	0.45 (0.09)	45.2 (8.84)
2015 H2	2.72 (0.04)	2.1 (0.34)	1.68 (0.34)	0.45 (0.09)	33.66 (6.99)
2016 H1	3.98 (0.04)	2.03 (0.32)	1.56 (0.32)	0.44 (0.09)	16.04 (3.58)
2016 H2	4.03 (0.02)	1.88 (0.32)	1.55 (0.32)	0.46 (0.09)	9.84 (2.17)

Notes: Panel A includes demand estimates with CBSA, half-year, and permanent seller fixed effects. Permanent sellers are those who submitted over 300 total bids. All project attributes, non-price attributes, and installer attributes are dummy variables. The star rating coefficients are relative to installers with no rating. The mean own-price elasticity is calculated based on the realized choice sets and does not account for ex-ante uncertainty in seller participation. Standard errors are in parentheses. Panel B 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. Bayesian Bootstrap standard errors (Rubin, 1981) in parentheses. Bootstrap weights for each auction are drawn according to a Dirichlet distribution with $\alpha = 1$ across 100 bootstrap samples.

Table 2: Consumer Surplus Gains from Access to the Platform

Panel A: Consumer Surplus Summary - Full Sample

	Mean	Std. Dev.
Total Consumer Surplus Per Household (\$)	1015.62 (190.29)	923.27 (181.47)
Consumer Surplus Per Unit (\$/watt)	0.14 (0.03)	0.11 (0.02)
(Consumer Surplus)/(Mean Bid Price)	0.06 (0.01)	0.04 (0.01)

Panel B: Mean CS by Electricity Bill Quartile

Quartile	Bids	Total CS (\$)	CS (\$/watt)
1	3.12 (0.03)	754 (203.55)	0.15 (0.04)
2	3.54 (0.03)	1079.53 (315.04)	0.17 (0.05)
3	3.64 (0.03)	1164.29 (304.08)	0.15 (0.04)
4	3.85 (0.04)	1074.81 (325.84)	0.11 (0.03)

Panel C: Mean CS by Home Value Quartile

Quartile	Bids	Total CS (\$)	CS (\$/watt)
1	3.1 (0.03)	808.88 (229.36)	0.11 (0.03)
2	3.47 (0.03)	1087.17 (314.25)	0.15 (0.04)
3	3.84 (0.04)	1134.51 (363.7)	0.17 (0.05)
4	3.86 (0.04)	1033.1 (354.79)	0.16 (0.05)

Notes: Panel A reports the mean and standard deviation of each variable across all households in the sample. Each column in Panel B and Panel C reports the means value of the variable separately for each quartile. Bayesian Bootstrap standard errors (Rubin, 1981) in parentheses. Bootstrap weights for each auction are drawn according to a Dirichlet distribution with $\alpha = 1$ across 100 bootstrap samples.

Table 3: 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 (.)	0.88 (0.022)	0.855 (0.026)	0.842 (0.028)	0.834 (0.028)
Mean Price Per Watt (Lowest Bid)	1 (.)	0.834 (0.021)	0.788 (0.025)	0.762 (0.027)	0.745 (0.027)
Mean Price Per Watt (Selected)	1 (.)	0.904 (0.022)	0.874 (0.026)	0.857 (0.028)	0.845 (0.029)
Consumer & Producer Surplus (On Platform)					
Consumer Surplus (\$)	1 (.)	1.859 (0.019)	2.502 (0.043)	3.089 (0.086)	3.597 (0.131)
Total Producer Surplus (\$)	1 (.)	1.501 (0.063)	1.843 (0.126)	2.14 (0.178)	2.379 (0.224)
Producer Surplus Per Bid (\$/bid)	1 (.)	0.75 (0.031)	0.614 (0.042)	0.535 (0.044)	0.476 (0.045)
Relative Quantities & Externalities					
# of Solar Installations	1 (.)	1.075 (0.006)	1.13 (0.007)	1.178 (0.01)	1.218 (0.012)
Pollution Damages Avoided (\$)	1 (.)	1.064 (0.007)	1.111 (0.008)	1.152 (0.009)	1.186 (0.011)

Notes: Summary of counterfactual simulations varying the number of bids that each project receives. All outcomes are reported relative to the one-bid case. Bayesian Bootstrap standard errors (Rubin, 1981) in parentheses. Bootstrap weights for each auction are drawn from a Dirichlet distribution ($\alpha = 1$) across ten bootstrap samples.

Table 4: Effects of Government Subsidies on Competition, Prices, and Welfare on Platform

	Counterfactual		
	ITC	No ITC (0)	No ITC (1)
Panel A: Market Outcomes			
Bids Per Project	3.491 (0.039)	3.240 (0.058)	3.237 (0.058)
Gross Price Per Watt (All Bids)	\$3.523 (0.004)	\$3.123 (0.054)	\$3.098 (0.053)
Net Price Per Watt (All Bids)	\$2.466 (0.003)	\$3.123 (0.054)	\$3.098 (0.053)
Net Price Per Watt (Lowest)	\$2.252 (0.005)	\$2.842 (0.053)	\$2.819 (0.054)
Net Price Per Watt (Selected)	\$2.396 (0.008)	\$2.987 (0.04)	\$2.966 (0.042)
Purchases on Platform (Normalized)	1 (.)	0.669 (0.038)	0.679 (0.037)
Panel B: Change in Welfare per ITC Subsidy Expenditure (On Platform)			
		(0)	(1)
Δ Consumer Surplus		\$0.570 (0.031)	\$0.552 (0.031)
Δ Gross Producer Surplus		\$0.711 (0.053)	\$0.718 (0.054)
Δ Bid Preparation Costs		-\$0.034 (0.005)	-\$0.033 (0.005)
Δ Pollution Damages Avoided		\$0.117 (0.013)	\$0.113 (0.013)
Δ Subsidy Cost		-\$1.000 (0)	-\$1.000 (0)
Δ Welfare (On Platform)		\$0.364 (0.014)	\$0.350 (0.013)

Notes: Panel A summarizes average outcomes across counterfactual market simulations with and without the federal investment tax credit. The second column, “No ITC (0)”, assumes a 0% pass-through rate of the ITC removal to prices off of the platform. The third column “No ITC (1)”, assumes a 100% pass-through rate of the ITC removal for prices off of the platform. Panel B calculates the relative change in welfare measures between the “ITC” and “No ITC” cases per dollar change in subsidy expenditure. Welfare estimates in Panel B do not include off-platform transactions. Bayesian Bootstrap standard errors (Rubin, 1981) in parentheses. Bootstrap weights for each auction are drawn from a Dirichlet distribution ($\alpha = 1$) across ten bootstrap samples.