Valuing Solar Subsidies*

Kenneth Gillingham† ⓡ A. Justin Kirkpatrick‡ ⓡ Bryan Bollinger§

October 21, 2026

Abstract

Individuals trade present for future consumption across a range of economic behaviors, and this tradeoff may differ across socioeconomic groups. To assess these tradeoffs, we estimate a dynamic model of residential solar adoption and system sizing using household-level data that offer plausibly exogenous variation in the future benefits from adopting relative to upfront costs. We find implicit discount rates of 17.2%, 15.6%, and 10.9% for low-, medium-, and high-wealth households. This heterogeneity remarkably persists for those with high credit scores. Counterfactual simulations demonstrate opportunities to reduce the regressivity of solar adoption, increase policy cost-effectiveness, and improve welfare for low-wealth households.

Keywords: solar, discount rates, energy policy, distributional impacts, dynamic discrete choice models.

JEL Codes: L94, Q48, H23, D12

*The authors thank Erica Myers, Laura Grant, Soren Anderson, Ryan Kellogg, Steven Haider, Frank Verboven, Rich Sweeney, and the seminar participants at IO+, AEA 2022, AERE 2025, Clemson University, UC Davis, CU Boulder, Marketing for Sustainability at Stanford GSB, NBER EEE Spring Meeting, NBER Distributional Consequences of New Energy Technologies and Policies, the Northeast Workshop on Energy Policy and Environmental Economics, and Marketing Dynamics for their valuable feedback. We also would like to thank an NBER Workshop on Distributional Consequences of Environmental Policies for support. Author ordering was randomized with the AEA randomization tool with confirmation code TzL9Vjl1_wY2.

†Yale University and NBER. Email: kenneth.gillingham@yale.edu.

‡Michigan State University. Email: jkirk@msu.edu.

§Tuck School of Business at Dartmouth College. Email: bryan.k.bollinger@tuck.dartmouth.edu.

1 Introduction

The rate at which individuals appear to trade present consumption for future consumption is important across a range of economic behaviors, including savings, human capital formation, and investment in personal health. Understanding this tradeoff is critical for public policy, which often subjects individual actions and behaviors to incentives or restrictions to achieve long-run objectives that are deemed socially desirable. Laboratory experiments with hypothetical choices to examine intertemporal decision-making are numerous, but credible quasi-experimental estimates of implicit discount rates revealed by market decisions are relatively rare (Hausman, 1979; Lawrance, 1991; Warner and Pleeter, 2001; Bollinger, 2015; de Groote and Verboven, 2019). Evidence on how discount rates vary across policy-relevant sub-populations is even scarcer. As noted by Barsky et al. (1997), econometric estimation of preference parameters, including time preferences, may be "particularly inadequate" when heterogeneity of preferences in the population is important, due to identification issues.

This study examines how consumers of different wealth groups trade off present for future consumption in realized decisions on the adoption of rooftop solar photovoltaic (PV) systems. We find economically meaningful differences in implicit discount rates between low- and high-wealth homeowners: 17.2% versus 10.9%.[1] This implies that high-wealth households implicitly value the same stream of solar benefits over time approximately 50% more than low-wealth households.

These results are important for policy. There has been substantial criticism of past solar policies by economists for primarily benefiting high-wealth households and shifting electricity system fixed costs from relatively wealthy households who are more likely to adopt rooftop solar to relatively poor households who are less likely to adopt (Borenstein and Davis, 2016, 2024). Net Energy Metering (NEM) is a very common policy that has been used at some point by 43 states in the United States. Under the standard NEM policy, solar adopters can export excess solar generation above consumption to the grid and receive credit for it in future months, thus earning compensation at the retail rate of electricity (usually up to the consumer's total annual consumption). The value of this incentive to consumers critically depends on their implicit discount rates, which may vary across wealth groups.[2] Despite the prevalence of NEM policies, we are not aware of any previous research that has evaluated their effectiveness in spurring additional solar PV capacity across wealth levels as a result of differences in discounting, or their effect on consumer welfare.

[1] We refer to the estimated discount rates as 'implicit' because they are the rates that rationalize observed decisions.

[2] Houde and Myers (2021) provide evidence that in the realm of household appliances, higher income consumers weight operating costs more relative to the purchase price.

Our work also relates to questions about the viability of traditional utility models as solar market share increases. If there is more distributed solar generation and the fixed costs of the grid are spread across fewer non-solar customers, this could lead to higher electricity rates and possibly induce further grid defections that could yield a "utility death spiral" (Kind 2013). In response to concerns about equity and the long-run viability of traditional utility business models, many states across the United States are currently reducing their compensation for excess solar electricity fed into the grid. For example, in 2023, California implemented "NEM 3.0," which reduced compensation for excess solar electricity by about 75%, a drastic change in the long-term benefit of rooftop solar. The effects of such a policy could differ across the population if there is heterogeneity in intertemporal tradeoffs.

In this paper, we employ rich and unique micro-data on rooftop solar panel adoption, and on the expected returns from such adoption. Adopting rooftop solar entails an upfront cost to install panels that generate future electricity cost savings. These savings depend on the system characteristics, sunlight, and the compensation rate for excess generation. Using proprietary Google Sunroof data for rooftop-specific expected electricity generation of optimal solar installations for every home in select California zip codes along climate borders, in conjunction with household utility bills, we calculate the value of the flow of benefits of both purchasing and leasing solar as a function of the implicit discount rate for the household. In this calculation, we account for the household's observed electricity consumption and the optimal installation size. We combine data on these variables with observed solar adoptions and household-specific demographic information to estimate our structural model.

In our study, heterogeneous implicit discount rates and marginal utility of income parameters are identified from sizable differences in electricity rates across administratively-determined climate zones and plausibly exogenous household-specific variation in future electricity bill savings from adopting solar. This variation is due to differences in solar irradiance as a function of rooftop characteristics, such as pitch, orientation, and shading from trees and structures. The identification of discount rates is challenging and requires some exogenous source of variation, such as an exclusion restriction (Magnac and Thesmar, 2002). In previous work, the econometric identification of discount rates has often relied heavily on strong functional assumptions.[3] Some papers utilize arguably exogenous changes to the environment, such as the entry of a new choice option (Warner and Pleeter, 2001) or policy change (Bollinger, 2015; de Groote and Verboven, 2019). These approaches using time-series variation typically involve strong assumptions about agent expectations in order

[3]For example, Hausman (1979) assume an exact lifetime for air conditioners, no deterioration over time, no differential in expected inflation rates of appliance prices and electricity, and constant utility over time for air conditioning. Our assumptions are weaker than these and we perform robustness checks over them.

to leverage such variation. In contrast, our approach identifies implicit discount rates from cross-sectional variation at the household level in the relative value of the upfront costs of installing solar and the long-term benefits, leaving us less dependent on assumptions about unobserved consumer expectations for identification. This is more similar to the strategy employed by Koster and Pinchbeck (2022), who leverage differences in property tax rates and house costs for homes similar on all other dimensions to estimate discount rates.

Our estimation approach uses conditional choice probabilities (CCPs) and generally follows Arcidiacono and Miller (2011). However, our utility bill data is anonymized, so we cannot match each non-solar household to a specific utility bill. This leads to an econometric challenge akin to the challenge of incorporating micro-data into BLP estimation (Berry et al., 1995; Conlon and Gortmaker, 2023), due to the fact that household-level likelihoods are no longer independent within zip code.[4] We overcome these challenges by integrating over the observed zip-code empirical distribution of household electricity consumption with replacement.

Based on our estimation results, we run two counterfactuals that change the solar policy while holding the rest of the electricity rate structure fixed. In our first counterfactual, we find that if the NEM compensation is substantially reduced to the average avoided cost of solar generation in California and half of the generated solar electricity is exported to the grid, then solar adoption for both wealth groups would decrease by 9%. Yet the consumer surplus reduction would be over six times higher for high-wealth households, highlighting how much more high-wealth households value NEM benefits. Such a policy would also reduce environmental benefits but would substantially decrease federal tax credit outlays and benefit electric utilities.

In a second counterfactual, we again replace NEM with a scheme that compensates excess solar generation at the avoided cost. We take approximately half of the funding that would have been used for NEM payments and allocate that funding to upfront subsidies (split between a per-W and per-installation subsidy), such that the total number of adoptions for high-wealth households are roughly the same as under the baseline. The result is a very different distribution of installations. Under this counterfactual, low-wealth installations increase by 27% (overall installations increase by 11%, although they are smaller on average, leading to an increase in installed capacity of 18%). This leads to a 30% decrease in consumer surplus for high-wealth households and a 17% increase for low-wealth households. We calculate the marginal benefit of public funds (Hahn et al., 2024) of the additional upfront

[4]To see why, consider that if one specific household is a high-consumption household, this would lower the probability that the other households in a zip code can also be a high-consumption household because the aggregation of all household consumption within a zip code must equal the observed empirical distribution.

subsidies in this counterfactual to be less than one for high- and medium-wealth households, but greater than one (1.23) for low-wealth households. Values greater than one indicate that the willingness to pay of the policy exceeds the government expenditure. This finding implies that subsidies targeted to poorer households not only address distributional equity concerns but are also likely more economically efficient than uniform subsidies.

Our estimates provide credible evidence that individual implicit discount rates that are rationalized by solar adoption decisions exceed common market rates and that these vary by wealth in ways that have important implications for economists and policymakers. This finding of relatively high implicit discount rates accords with results in the 'energy-efficiency paradox' literature of undervaluation of future energy savings in energy efficiency purchases (Metcalf and Hassett, 1999; Allcott and Greenstone, 2012; Gerarden et al., 2017; Gillingham and Palmer, 2014; Gillingham and Myers, 2025).

One novel aspect of our study is that we obtain household-level credit score data to see if credit helps to explain the differences in discount rates; we find households with near-prime or sub-prime credit have higher implicit discount rates (29.4% and 66.2% for medium- and low-wealth households, respectively), demonstrating the likely role of borrowing costs or credit limits in the purchase decision. This is consistent with work showing that credit constraints are important in the adoption of energy-efficient stoves by low-income households in Kenya (Berkouwer and Dean, 2022). However, for those with good credit, the ratio of the implicit discount rates across wealth groups does not change with the inclusion of the household credit score data, suggesting that other mechanisms, such as behavioral anomalies, may also be playing a key role.

Our findings are new to the small but growing literature on residential solar adoption (Kirkpatrick and Bennear, 2014; Hughes and Podolefsky, 2015; Pless and van Benthem, 2016; Gillingham and Tsvetanov, 2017; Feger et al., 2022). One major recent line of solar research is on spatial misallocation of solar (Sexton et al., 2021; Lamp and Samano, 2023; Colas and Saulnier, 2024). Our results highlight a structural inequity in existing solar policy, with a misallocation between upfront subsidies and NEM benefits. We show that standard NEM is a regressive policy and that a switch to upfront subsidies can not only improve cost-effectiveness but also improve distributional equity.

2 Background

2.1 Solar Policy

The adoption of rooftop solar PV has benefited from billions of dollars of subsidies. The most prominent subsidies in the United States over the past two decades have been the many state-level upfront adoption rebates, such as the California Solar Initiative program, the federal investment tax credit (ITC), and NEM. Elsewhere in the world, popular policies include NEM and feed-in tariffs, in which solar adopters send electricity generated from solar to the grid and are compensated at a price that is usually above the retail electricity rate. Net-billing tariffs (NBT) refer to mechanisms in which excess solar generation above consumption that is exported to the grid is compensated at some predetermined rate that is usually below the retail electricity rate. Under NBT, consumers receive a 'net bill' each month and credits for more generation than consumption in each month cannot roll over, as they can under NEM. California's new "NEM 3.0" is a version of NBT, with some additional elements to encourage energy storage. In recent years, most U.S. state rebate programs have been discontinued, but the federal ITC is still in place and NEM is still widely used across the country, despite past and current debates about phasing out both types of incentives.[5]

From a social welfare perspective, policies to encourage adoption of solar can be motivated by innovation market failures, such as learning-by-doing, or as a second-best approach to addressing uninternalized environmental externalities. One key distinction between the different types of solar subsidies is whether they are upfront adoption subsidies, such as tax credits or rebates, or flow subsidies, such as NEM and feed-in tariffs. Typically, a subsidy on the extensive (adoption) margin would sacrifice efficiency on the intensive (generation) margin, as there would be less incentive to improve the productivity of generation and undertake maintenance to ensure generation (Aldy et al., 2023). However, in the case of solar panels, there is very little that can be done to improve efficiency of already-installed rooftop systems and the systems rarely require routine maintenance.[6] Therefore, the decision to adopt solar is made as a "set it and forget it" decision requiring only consideration of the up-front investment and the (constant) flow benefits for the next 25 years, which critically depend on the household implicit discount rate.

The existing literature on the cost-effectiveness of solar subsidies has generally found relatively high costs from upfront subsidies due to free-riders who would have installed solar anyway without the subsidies (Hughes and Podolefsky, 2015; Gillingham and Tsvetanov, 2014). However, it might still be possible to justify at least some rooftop solar subsidies on

[5]The federal "One Big Beautiful Bill Act" of 2025 has provisions to phase out the ITC.

[6]See for instance https://www.energysage.com/solar/101/solar-panel-maintenance/.

economic-efficiency grounds based on uninternalized social costs of greenhouse gas emissions (Rennert et al., 2022), uninternalized local air pollution emissions (Sexton et al., 2021), and learning-by-doing externalities (Bollinger and Gillingham, 2023). Yet without careful policy design, the distributional consequences may be perverse, especially in a state like California that has increasing block pricing for electricity, so that those who consume more (usually wealthier households) pay higher rates and thus benefit more from adding solar.

2.2 Implicit Discount Rates Across Wealth Groups

This paper focuses on estimating implicit discount rates across wealth groups because they are a crucial ingredient for comparing the welfare impacts of subsidy policies that have different time profiles of benefits. However, while implicit discount rates characterize the realized intertemporal tradeoff *implied* by actual decisions, they may capture more than just the pure rate of time preference of different groups. Differences in lending and borrowing costs, the marginal propensity to consume, beliefs, and risk aversion can all also help explain the tradeoffs observed by economists either in the laboratory or in the field (Cohen et al., 2020). Chatterjee and Eyigungor (2024) show that discount factor heterogeneity is needed to explain patterns in credit card usage and default, and Chatterjee et al. (2023) show that such heterogeneity is needed to explain the evolution of credit scores.

Coller and Williams (1999) show that the consumer's real rate of return should equal their discount rate only if it lies between their borrowing and lending costs, and thus differences in borrowing rates or binding credit constraints could lead to higher implicit discount rates for lower-wealth households than others (Allcott et al., 2015). We explore this in our setting by examining how our results differ after accounting for the household's credit score. Wealthier consumers may also have a greater ability to smooth consumption due to greater liquidity, which could also lead to observed intertemporal tradeoffs differing by wealth (Cubitt and Read, 2007). This may be less relevant in our setting, as the flow subsidies we are examining tend to be smooth over time.

It is possible that high-wealth households have different beliefs than low-wealth households about the life of the panel, future electricity costs, how long they will be in their home, and whether the solar panel will be capitalized in the transaction price if they sell. We will discuss each of these and provide evidence suggesting that they are not likely to explain our heterogeneity results. Finally, risk aversion may differ across wealth groups, which have very different levels of consumption. Cohen et al. (2020) show that in a simple two-period tradeoff, risk aversion scales down the pure rate of time preference from the estimated "implicit"

discount rate rationalized by decisions. We allow for this by modeling heterogeneity in the marginal utility of income across wealth groups.

3 Data

In order to identify heterogeneity in implicit discount rates across wealth groups without strict distributional assumptions, it is essential to have rich household-level data on the main drivers of solar adoption. To this end, we assemble a detailed household-level data set on solar adoptions, home characteristics, and household characteristics covering the period 2014 to 2016. Our research design focuses on adjacent zip codes in California that face different electricity rates due to being categorized into different climate zones. We identify 28 zip codes in Pacific Gas & Electric (PG&E) territory that are entirely contained in a single climate zone and have an adjacent zip code in a different climate zone.[7] Figure 1 illustrates the zip codes included in our "boundary zones," all of which are in the greater San Francisco Bay Area of California.

3.1 Data sources

Our first data source contains address-level data on home characteristics from CoreLogic (CoreLogic, Inc., 2018). We focus on all single-family detached owner-occupied non-mobile home residences that were built before 2014 in the 28 zip codes. The data include the year built, heated square footage, and number of stories. This provides the set of potentially adopting households.

Our second data source is publicly-available voter registration data from the California Secretary of State (**?**). For each address, we categorize households to be Democrat, Republican, or mixed.[8] We merge in household characteristics data from InfoUSA (InfoUSA, Inc., 2020), including the number of children, race, home-ownership status, the length of time at the residence, the number of open lines of credit, and the calculated wealth of the household (inclusive of home equity). We also use credit score data from Experian (Experian plc, 2025; Gibbs et al., 2024). If an address has multiple households, we use the data from the household that occupied the home for the plurality of our study period.

[7] Climate zones are administrative boundaries set based on 500' elevation isopleths adjusted to keep zip codes intact; see Appendix A. All included zip codes on both sides of the boundary are in PG&E territory and face similar permitting and interconnection processes.

[8] We identify the household as registered Democrat if and only if both of the two longest-registered 2014 voters are registered Democrats or are registered with the Green Party.

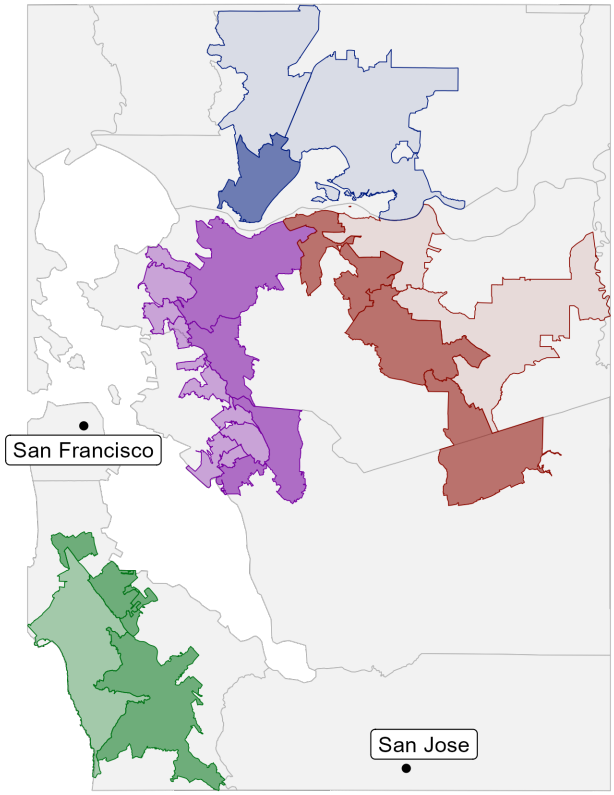

Figure 1: **Sample zip codes**. Zip codes in PG&E service territory that border a CEC climate zone boundary. Climate zones determine the electricity pricing scheme each zip code faces. Adjacent zip codes along a boundary face a similar climate but a different average value of offset electricity. Colors indicate particular boundary zones, with the shading illustrating the climate zone on each side of the boundary.

Our installation-level data set of solar adopters is a restricted-access version of the Lawrence Berkeley National Laboratories "Tracking the Sun" (TTS) database (Lawrence Berkeley National Laboratory, 2018). It contains the address of the installation, application date, installation date, system size (in watts), and total system cost exclusive of subsidies and tax credits. Our study window is after the California Solar Initiative (CSI) subsidy period had ended. We match the solar adopter data to the data on potential adopters with a 96% success rate, and we remove any households that adopted prior to 2014. The final sample contains 7,244 solar adoptions during the study period out of a total of 183,667 potential adopting households. Finally, we bring in data on the amount of sun reaching each solar rooftop from Google Sunroof and data on household electricity consumption from PG&E (**?**). To help provide interpretability for our wealth results, we merge in the credit data from Experian.[9] Data construction details are in Appendix A.

Table 1: Sample Summary Statistics

	All		Adopters		Non-adopters	
	Mean	Std. dev	Mean	Std. dev	Mean	Std. dev
Variables (N=183,667 households)						
Wealth ($1,000's)	2500.09	1017.81	2624.27	909.25	2494.99	1021.70
1(Low wealth bin)	0.33	0.47	0.24	0.43	0.34	0.47
1(Medium wealth bin)	0.33	0.47	0.42	0.49	0.33	0.47
1(High wealth bin)	0.33	0.47	0.34	0.47	0.33	0.47
Lines of credit (count)	0.67	1.41	0.75	1.53	0.66	1.40
1(Children present)	0.32	0.47	0.41	0.49	0.31	0.46
Length of residence (years)	15.56	12.26	13.42	10.65	15.65	12.31
Square Footage (1,000 sq.ft.)	1.79	0.75	2.16	0.81	1.77	0.75
1(Single story)	0.32	0.47	0.45	0.50	0.32	0.47
1(Dem voter registration)	0.43	0.49	0.45	0.50	0.43	0.49
1(Possible non-owner)	0.18	0.39	0.06	0.24	0.19	0.39
Adopted System Characteristics (N=7,244)						
System Size (kW)			5.29	2.36		
Installation Cost ($1,000's)			24.23	12.6		
Low-Wealth Adopter Sys. Size (kW)			4.91	2.23		
Med-Wealth Adopter Sys. Size (kW)			5.51	2.34		
High-Wealth Adopter Sys. Size (kW)			5.29	2.45		
1(Leased system)			0.45	0.5		

[9]Experian performed the matching based on names and addresses, reaching a match rate of 63%. Experian data is further discussed in Appendix J.1.

Table 1 shows summary statistics for the homes and households in the sample by adopter status, where we divide the sample into three wealth groups.[10] Solar adopters tend to be middle- and high-wealth, are more likely to have children, and have larger homes. Adopters are only slightly more likely to be all-Democratic in voter registration, and have a shorter length of residence.

3.2 Solar Irradiance and Electricity Consumption

Some of the relationships we observe in the summary statistics may be capturing heterogeneity in the pecuniary value of adopting solar. To explicitly calculate the monetary benefit from adopting solar, we need to know how much solar electricity would be generated for any given installation on any given house. Thus, for every house in the base data, we leverage the Google Sunroof data on the 'effective' solar irradiance that reaches each rooftop. By definition, solar irradiance is the amount of sunlight that hits the surface of the earth and is a function of climate. It varies considerably across the United States and around the world, and even within states, e.g., across zip codes. Effective irradiance in our data also accounts for the obstruction of solar irradiance by surrounding structures and vegetation, as well as panel orientations and pitches that may fail to capture all irradiance due to rooftop characteristics. Effective irradiance provides micro-level variation in the electricity generation of a unit of solar capacity, namely from household to household within neighborhoods. Moreover, for a given home, effective irradiance varies across the rooftop, with some portions of the rooftop receiving more sunlight than others. These irradiance data are especially useful because installers typically use Google Solar API or similar software for system sizing when developing quotes for potential clients.

Figure 2 shows an example of the effective irradiance of four homes and what it means for solar generation. Panel (a) shows Google Sunroof imagery of the homes as modeled by Google based on satellite imagery. Panel (b) shows how expected solar generation declines with more panels due to reduced irradiance.[11] Google Sunroof predicts electricity bill savings from adopting solar for each rooftop based on the irradiance profile, electricity rate, and the optimal size as a function of household electricity consumption and utility rate. Homes located in deep shade or with a roof profile that does not angle southward need more panels

[10] The threshold between low and middle wealth is $1.96 million and between middle and high wealth is $2.92 million, yielding three equal sized wealth bins. Wealth is inclusive of the value of the home.

[11] While the expected solar generation is downward sloping with the number of solar panels, the relationship is not necessarily monotonic. Following standard installation practice, each roof segment, or partition of a roof with a common pitch and azimuth, is filled before an additional segment is used. Thus, small non-monotonicities can occur when the first panel on a subsequent segment is installed. When non-monotonicities imply multiple equimarginal crossings, we select the smaller number of panels.

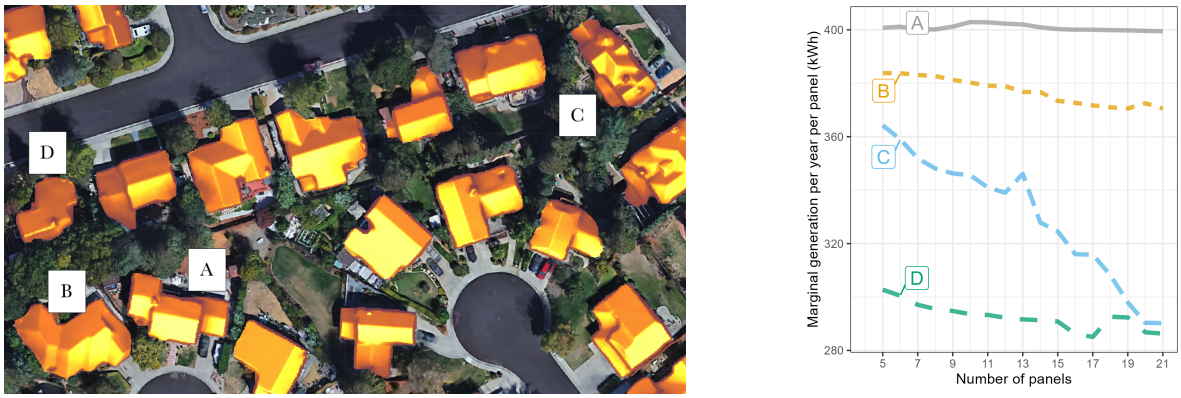

(a) Google Project Sunroof display

(b) Solar generation profiles

Figure 2: **Solar generation profiles** for four selected rooftops on a single block in Solano County, California. Panel (a) shows the Google Sunroof imagry, while panel (b) shows the marginal generation for each household per panel on the vertical axis and the number of panels on the horizontal axis. Contrast dashed green (house D), which exhibits low irradiance (300 kWh/yr) for the first panel and a slow decline up to 15 panels followed by a leveling off, with gray (house A), which exhibits strong irradiance for the first 20 panels (400 kWh/yr). Dashed blue (house C) shows high initial irradiance (365 kWh/yr) but a rapid decline corresponding to the multifaceted roof line.

to generate a given amount of electricity, increasing the cost per kWh to that household. This and the utility rate borders provide our two main sources of identifying variation.

The value of adopting solar depends on household electricity consumption in our setting for two reasons. First, there is minimal compensation for generation above the household's annual consumption.[12] Second, there is an increasing tiered electricity rate structure, so the marginal value of generation to the consumer is greater if it reduces a high-consuming household's electricity imports from the grid. Because of the importance of this rate structure, we obtain data on annual electricity consumption for all customers who reside in any of the 28 PG&E border zip codes in our sample. The timing of solar adoption in the administrative data enable us to match solar adopters with their energy consumption data. For non-adopting households, we only know their zip code. Further details about the consumption data are in Appendix A, as is information about PG&E electricity rates.

We can calculate the marginal levelized cost per kilowatt-hour (kWh) for solar generation by discounting the monetary value of the flow of generation over time.[13] This marginal cost

[12]Compensation for excess generation through Net Surplus Compensation (NSC) was set by the California Public Utilities Commission (CPUC) at roughly two to three cents per kWh, far below the marginal levelized cost of solar for any household in our sample. Thus, we treat NSC as negligible.

[13]The levelized cost is a measure computed as the present discounted value of lifetime costs divided by discounted lifetime production. We use the default rate of 4% employed by Google Sunroof. This is a standard default rate in the industry used in system sizing software that solar installers use when recommending

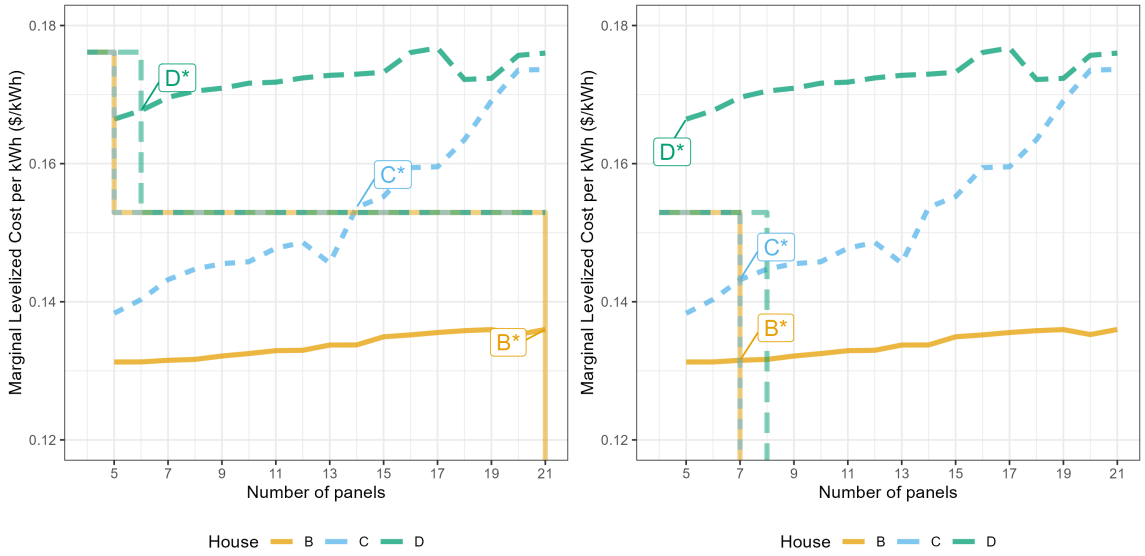

(a) 5th Consumption Bin

(b) 2nd Consumption Bin

Figure 3: **Optimal System Sizes for Two Example Consumption Levels**. Figure shows the marginal per-kWh price of solar generation for the subset of homes from Figure 2b along with the marginal grid price (dashed line) for consumption. Optimal installation sizes are indicated with "*" for each house.

per kWh is increasing with more solar panels due to the declining marginal generation with more panels (from reduced irradiance due to the best locations being taken first). Figure 3 plots the per-kWh levelized cost over the number of panels for three of the four home profiles from Figure 2 (one home is removed to simplify the figure). These levelized cost curves are the upward sloping lines. Figure 3 also shows the weakly decreasing step functions illustrating the cost of the marginal unit of grid electricity for each home. Drops in these step functions occur where the number of panels on each rooftop generates a sufficient amount of electricity for the household to drop down to a lower electricity rate tier.

The intersection of each step function with the upward-sloping levelized cost curve is marked with a letter for the household, and it refers to the optimal solar system size where the marginal benefit of another panel equals the marginal cost. We divide our sample into five equal electricity consumption bins based on average consumption over the sample period. Panel (a) assumes that the households that are in the highest bin, while panel (b) assumes that the households in the 2nd bin.

system size. This rate is lower than the household rates we estimate. However, the optimal sizing decision is generally not made by consumers alone, but rather by the installation contractor who presents the cost and expected payoff for a specific recommended system when generating a quote for homeowner consideration. A lower discount rate tends to increase the savings relative to up-front costs, which favors the installer using the lowest reasonable rate.

To better understand Figure 3, begin with panel (a). Consider home D (green dashed line). The upward-sloping marginal cost per kWh dashed line crosses the electricity rate step function when the system size contains six panels (pointed out by D*). Home D receives much less irradiance than home B, and thus home B has a much lower marginal cost per kWh (yellow upward-sloping line). The electricity rate step function for home B steps down to a lower tier at a solar system size of 21 panels. This system size is also where the step function crosses the marginal cost per kWh line, thus indicating that the optimal system size for home B is 21 panels. Home C falls in between, with an optimal system size of 13 panels.

Panel (b), which shows the same four homes if they had electricity consumption in the second consumption bin, illustrates much smaller optimal system sizes. Homes B and C both have optimal system sizes of seven panels, while home D has such poor irradiance that it has an optimal system size equal to the minimum installation size.[14] In other words, it would not make sense to install a solar system on this rooftop if the electricity consumption was in the second consumption bin due to the lower electricity rates consumers in this consumption bin face. The basic logic described here underpins our optimal sizing model.

3.3 Relationship between Effective Irradiance and Adoption

In Figure 4 we provide model-free evidence that shows the relationship between the probability of adopting solar and the expected panel generation from a fairly typical 15-panel array.[15] As the expected generation increases due to higher solar irradiance, the probability of adoption increases. However, high- and medium-wealth households respond more to expected generation than low-wealth households, both in absolute and relative terms. Since low-wealth households are less responsive to increases in the flow payoffs, as appears to be the case in this figure, they appear to place less value on these future payoffs relative to the present. This provides suggestive evidence on the variation underlying identification of differences in implicit discount rates in our structural model.

4 Model

We develop a dynamic discrete choice model of residential solar installations. The motivation for including dynamics to estimate implicit discount rates is that consumers have a "buy-or-wait" decision about adopting solar, with expectations about future electricity prices and

[14] We set the minimum installation size to five panels which corresponds to the smallest observed installation in our data.

[15] The average optimal sized array is 15 panels. The average optimal-sized array conditional on adoption is 21 panels.

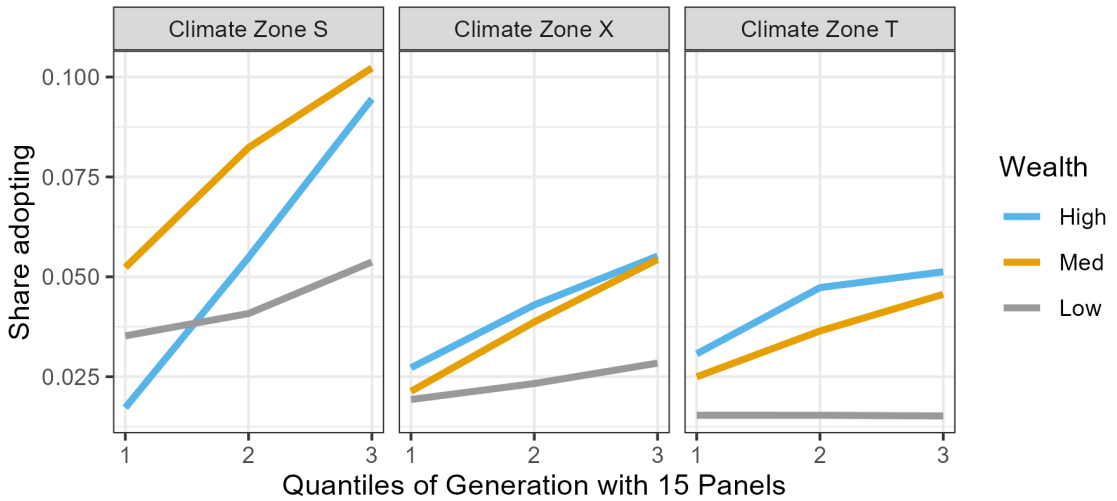

Figure 4: **Model-free evidence**. Observed share adopting solar over tertiles of irradiance by wealth by climate zone. The figures show the change in observed share adopting as the total generation with 15 panels increases.

solar panel prices. A key innovation in our model is to include a static optimal solar system size decision based on setting the marginal benefits of an additional panel from reduced use of grid electricity equal to the marginal costs of installing that panel. Because of the granularity of our data, we can calculate what the optimal installation size is for every potential adopting household as a function of electricity consumption, as is illustrated in Figure 3.

The optimal system sizing model is important in our setting both for correctly calculating the payoffs from installing solar for a given consumer and for running counterfactuals that change upfront versus future payoffs, and thus affect the optimal system size. Once the household and installer determine the optimal solar system size for the home, the household then decides whether to install solar using a discounted net present value calculation. This calculation explicitly accounts for consumer expectations of future electricity prices and solar prices. Details are in Appendix B.

4.1 Solar Installation Decision

Our dynamic discrete choice model of solar installations, conditional on the optimal system size, follows the literature in treating the adoption of solar as an exit action (e.g., de Groote and Verboven (2019)). A distinguishing feature of our model is that it leverages household-level variation in characteristics and climate zones for identification.

The consumer decides whether or not to install solar, $j \in \{0, 1\}$, where $j = 1$ is the decision to install. The flow utility for choice j at time t is given by $u_{jt}(q_0, p_t) + \sigma\epsilon_{jt}$, in

which q_0 is the household's baseline electricity consumption prior to installing solar and p_t is the electricity price. The utility when adopting solar depends on the sizing function, which depends on these two arguments, which will determine q^*, the electricity generated by an optimally-sized solar system. The ϵ_{jt} is a stochastic structural error term with scaling parameter σ.

We normalize the flow utility of not adopting solar to $u_{0t} = 0$; accordingly, we include the savings from using less grid electricity in the value of adopting solar. In what follows, we will derive an expression for the value of both adopting solar and not adopting. Since we treat adopting solar as an exit action, the expression for adopting solar will capture the discounted stream of expected utility for all periods subsequent to the adoption; i.e. $u_{1t} = v_{1t}$, where v_{jt} is the choice-specific value function for choice $j \in \{0, 1\}$ at time t (absent the additive stochastic shock). The value of not adopting will be equal to the continuation value, as there is the option of adopting later.

The total expected utility at time t is given by:

$$V_t(p_t, C_t, F_t) = \sum_{\tau=t}^{\infty} \rho^{\tau-t} \mathbb{E}\left[\max_{j\in\{0,1\}} \left(u_{j\tau}(q_0, p_t) + \sigma\epsilon_{j\tau}\right)\right], \tag{1}$$

in which ρ is the household's discount factor, C_t is the variable cost of installing solar to the consumer, and F_t is the fixed cost of the system. Both C_t and F_t are inclusive of rebates and the tax credits.

By Bellman's optimality principle, we can also define the value function recursively (dropping the t subscript for notational convenience) as:

$$\begin{aligned} V(p, C, F) &= \max_{j\in\{0,1\}} \left(v_j + \sigma\epsilon_j\right), \text{ where} \\ v_j &\equiv \mathbb{E}\left[u_j | p, C, F\right] + \rho\mathbb{E}\left[V(p', C', F' | p, C, F)\right]. \end{aligned}$$

The prime notation indicates the next period's value of the state variables (and we drop the t subscript for notational clarity). We include an expectation term for u_j in addition to the next period's continuation value because solar adoption is treated as a terminal choice whose payoff also depends on the evolution of the state variables (p, C, and F). We follow de Groote and Verboven (2019) and assume rational expectations. We define a period in our model as a quarter.

In order to derive an expression for the choice-specific value function of purchasing solar, v_1, we start with the total expected economic value from the adoption of solar over its lifespan T relative to not adopting (not inclusive of the ϵ_j):

$$\delta_1 \quad = \quad \int_{q_0-q^*}^{q_0} \sum_{\tau=1}^{T} \left(\rho(1+\zeta)(1-\lambda)\right)^{\tau-1} p_\tau(x)dx - C(K^*) - F. \tag{2}$$

The integrand in (2) reflects the present value of future costs of grid electricity avoided given solar generation of q^* over the lifespan of the array, which is given by T.[16] It is determined by the per unit price of grid electricity, $p(x)$, which potentially varies (as a step function) in x, and by the grid electricity consumed, which is the difference in total electricity consumption per period, q_0, and the electricity generated by the optimally sized solar PV system q^*. Consistent with the sizing model and following de Groote and Verboven (2019), ζ is the real trend in electricity prices and λ is the panel depreciation factor.[17] The $C(K^*)$ and F are the variable and fixed costs faced by the consumer of the solar installation of optimal size K^*,[18] and the variable costs decline each period by a factor η such that $C(K^*)' = \eta C(K^*)$.[19]

The expression for the expected economic value of purchasing solar over the life of the solar array can thus be simplified to:

$$\delta_1 = \theta q^* \bar{p} - C(K^*) - F, \tag{3}$$

in which we define:

$$\theta \equiv \sum_{\tau=1}^{T} \left((1+\zeta)(1-\lambda)\rho\right)^{(\tau-1)} \tag{4}$$

The θ term can be interpreted as the number of current period payoffs that the household values as equivalent to the stream of solar net benefits over T periods (similar to an indifference break-even period but accounting for panel depreciation and rate increases). The $\bar{p}$ is the current average cost of grid electricity avoided by adoption of an optimally-sized installation

[16] There is a very small approximation here: we treat the average price of offset grid electricity as not changing as a result of the very small rate of panel depreciation, given by λ, leading to slightly less consumption at the household's lower tier.

[17] We estimate a climate zone and consumption bin-specific ζ according to $p'(x) = (1+\zeta)p(x) + \epsilon^p$, with results in Appendix D. Depreciation is incorporated in our model by further discounting the stream of benefits of adopting generated from a non-depreciated solar array of the optimal size.

[18] We assume that all homeowners have sufficient tax liabilities to qualify for the 30% tax credit offered on solar during this period, which we believe is a reasonable assumption for California homeowners due to the fact that the tax credit can be carried over to future years and would be fully captured by the installing company for leased systems.

[19] Upon examination of the data, the decline in solar PV costs appears to be in the variable costs and F appears to be constant over time. This is consistent with cost declines that relate to either panel costs or increased labor efficiency. We allow the fixed price component to vary by boundary zone.

that produces q^*.[20] Thus, $\theta q^* \bar{p}$ captures the direct value of solar electricity consumption over the installation's life.

Next, we need to make an assumption about what consumers expect to do after the lifetime of their solar array. We formulate the problem as an infinite-time dynamic problem, with the simplifying assumption that once a household makes the decision to adopt solar, the household will re-adopt after the solar installation reaches the end of its working life, i.e., every T quarters. This seems like a reasonable assumption given that the household has already chosen to adopt.[21] An alternative could be to set utility to equal zero after the life of the solar array, but this would not capture either the net benefits of a new solar system or the cost of consuming grid electricity in later periods.[22]

Under these assumptions, the expected value of installing solar is given by the stream of benefits from installing solar again every T quarters:

$$v_1 = \omega \left[\underbrace{\frac{1}{1-(\rho(1+\zeta))^T}\theta q^* \bar{p}}_{p^{EL}} - \underbrace{\frac{1}{1-(\rho\eta)^T}C(K^*) - \frac{1}{1-\rho^T}F}_{p^{INV}} \right] + X\beta. \tag{5}$$

The terms within braces in equation (5) are made up of the value of the stream of offset grid electricity costs (p^{EL}) and the upfront investment price (p^{INV}).[23] These terms are a function of a discount factor ρ, adjusted based on the trends in solar and electricity prices, as well as the depreciation rate. The primary effect of ρ is to change the relative size of the two terms, with a higher discount rate placing more emphasis on the installation costs.[24] The pecuniary

[20] As solar output declines gradually, it is possible for $\bar{p}$ to also change slightly with tiered pricing, but we will account for this in the estimation of how average price changes for different consumption bins over time.

[21] Some consumers might continue to use a very old depreciated system until it no longer works at all, but we believe that using the standard 25-year lifespan is more reasonable than assuming all consumers continue to use very old depreciated systems, many of which would no longer be working properly.

[22] In de Groote and Verboven (2019), the authors argue that they do not have to take a stand on whether the dynamic problem is finite time or infinite time when using CCP estimation. This is technically true, but if consumers view the problem as an infinite-time problem and the problem is treated as a finite-time problem, this implicitly assumes that there is zero utility in period $T+1$ (and later periods), since the expression for the value of not adopting includes the current utility of not adopting plus additional T periods of value when adopting and $T+1$ when not adopting. This inconsistency is removed if the value in period $T+1$ is zero when adopting solar, but this means that the consumer treats the costs of grid electricity needed in the current period as the same as the electricity costs in period $T+1$ (either grid electricity costs or a new solar array).

[23] This is the notation used in de Groote and Verboven (2019), who refer to these as the 'price terms.'

[24] This is an indirect utility specification; we implicitly assume that utility is linear in numeraire non-solar consumption and that agents can borrow or save at an implicit rate of $1/(1-\rho)-1$.

term is pre-multiplied by the marginal utility of income, ω.[25] This changes the relative value placed on the economic benefit of installing solar and the additional non-pecuniary utility derived from solar adoption, captured in the $X\beta$, such as from a warm glow effect (Andreoni, 1990). The $X\beta$ includes time-invariant household-level demographic variables and home characteristics, as well as an area fixed effect interacted with wealth, and a time fixed effect, both of which we assume consumers expect to remain constant in future periods. The inclusion of wealth by area fixed effects captures area-specific heterogeneity in fixed costs, search costs, and other local solar market aspects.

4.2 Solar Leases

Nearly 45% of adopters during our study period used a third-party "leased" system wherein the lessor bears all up-front cost to install the rooftop system, and the lessee agrees to pay a price per unit of consumption during the life of the lease as part of a Power Purchase agreement (PPA). In terms of up-front costs and benefits over time, the adoption decision for a lessee is different from that of a purchaser. Specifically, leasing has the effect of amortizing the cost of installing over the life of the panels. We are the first in the literature modeling solar demand to explicitly model the lease option.

Given standard pricing approaches for leased systems, we show in Appendix F that the value of leasing simplifies to:

$$v_1^l = \omega \left[p^{EL} - \kappa^l p^{INV} \right] + X\beta^l. \tag{6}$$

This is the same expression as for purchasing, except the addition the κ^l; κ^l depends on discount rate and changes the relative tradeoff between upfront costs (which the installer incurs but passes on to the household through the PPA payments, along with a markup) and the long-term benefits of solar when leasing vs. purchasing. The κ^l is determined by the installer's pricing rule for how it amortizes the cost of installing, its expected rate of return over the life of the panels, and the household's discount rate (which determines how they value the future PPA payments). Households with $\kappa^l < 1$ (those with a high enough discount rate) get more economic benefit from leasing and households with $\kappa^l > 1$ get more economic benefit from purchasing.

In our specification, we assume that households have an unobserved, permanent lease-purchase type such that a lease-type household will consider only leasing when considering adopting and vice versa for purchase-types. We view this as a less-restrictive assumption than

[25] Explicitly modeling savings behavior and the possible dependence of the marginal utility of income on solar assets is beyond the scope of this paper and would require additional data we do not have access to.

parameterizing lease-purchase type by observable household characteristics.[26] We discuss integration over the unobserved type in Section 5.2.

4.3 Continuation value when not adopting

To close the model, we need to make assumptions about consumer expectations. We follow Scott (2014) and de Groote and Verboven (2019) by decomposing the expected continuation value shown in equation (1) into a deterministic value plus a short run prediction error, ϵ^V, such that households are correct on average:

$$\rho\mathbb{E}\left[V(p', C', F'|p, C, F)\right] = \rho[V' + \epsilon^V]. \tag{7}$$

We define $\epsilon_0 \equiv \epsilon^0 + \rho\epsilon^V$ such that the error term for non-adoption shown in equation (1) is inclusive of both the stochastic utility of non-adoption and the short-run prediction error. We assume ϵ_0 and ϵ_1 are distributed Type 1 extreme value with scale variance of σ. With these assumptions, we can write the below expression for the value of non-adoption using conditional choice probabilities (Hotz and Miller 1993), treating adoption as an exit state:

$$v_0 \quad = \quad \rho\left(v_1' - \sigma\ln(Pr_1') + \sigma\gamma\right), \tag{8}$$

in which v' is a calculated next period value of adopting, Pr_1' is the next period adoption probability, and γ is the Euler-Mascheroni constant. Without loss of generality, we normalize $\sigma = 1$.

4.4 Final econometric specification

In our main specification, we assume that consumers calculate v' for both purchasing and leasing by advancing prices according to the price trends and that they formulate a belief about next period adoption probabilities, Pr', using current period state variables (conditional on type and consumption), which we estimate using a semi-parametric, flexible logit expression. Under this assumption, the difference in the expressions for v_1 and v_0 in equations (5) and (8) yields the difference in the value of adopting versus not adopting, now including the relevant

[26] We provide a robustness check where we randomly reassign observed adopters to lease vs. purchase type, showing that our assumption is not pivotal to our estimates.

subscripts:

$$\begin{aligned} v_{1ibt} - v_{0ibt} &= \omega_i \left[\frac{1-\rho_i(1+\zeta_b)}{1-(\rho_i(1+\zeta_b))^T} \theta q^*_{ib} \bar{p}_{ibt} - \kappa_i \left(\frac{1-\rho_i \eta}{1-(\rho_i \eta)^T} C_t(K^*_{ib}) + \frac{1-\rho_i}{1-(\rho_i)^T} F_i \right) \right] \\ &+ (1-\rho_i) X_i \beta + \rho_i \left(\log(Pr'_i) - \gamma \right) \end{aligned} \tag{9}$$

in which κ_i is equal to one for purchasing and equal to κ^l when leasing. Under the assumption of time-invariant lessor/purchaser type, Pr' is the semi-parametric estimate of adopting in the next period when of the same type. As a robustness check, we also estimate the model using an alternative assumption for V' and Pr' following de Groote and Verboven (2019) by using realizations of the value of adoption in the next period and constructing the semi-parametric estimate of the next period adoption probabilities with the realizations of next-period state variables.

We allow the marginal utility of income, ω_i, and the discount factor, ρ_i, to depend on wealth.[27] Heterogeneity across households in ζ_b results from differences in rate increases for different consumption bins due to the tiered pricing structure, in K^*_{ib} and q^*_{ib} from differences in the optimally sized system for consumption bin b, in F_i from geographic variation in the fixed costs of solar installation, and in $\bar{p}_{ibt}$ from differences in the optimally sized system and electricity prices. A household's probability of adoption conditional on consumption bin is:

$$Pr^p_{ibt} = \Lambda \left(v_{1ibt} - v_{0ibt} \right), \tag{10}$$

in which Λ is the standard Logit cdf. These probabilities will form the basis of the likelihood estimation procedure.

5 Identification and Estimation

5.1 Identification

5.1.1 Identifying Variation

To identify both the discount rate and the marginal utility of income we need 1) variation that shifts the long-term stream of benefits (the offset costs of grid electricity) relative to the upfront costs, and 2) variation that shifts the total economic value of adopting solar relative to other factors.

[27] We thank Ariel Pakes for noting the importance of heterogeneous marginal utility of income identified separately from discount rates. The parameters of ω_i are estimated with an exponential transformation $\omega_i = e^{\omega_0 + \omega_{med} 1(\text{wealth}_i = \text{med}) + \omega_{low} 1(\text{wealth}_i = \text{low})}$. The parameters of ρ_i are estimated with a Normal cdf transformation $\rho_i = \Phi\left(\rho + \alpha_{low} 1(\text{wealth}_i = \text{low}) + \alpha_{med} 1(\text{wealth}_i = \text{med})\right)$.

Climate zone borders provide variation in the marginal price of electricity, and thus variation in $\bar{p}$ for an installation of a given size. This provides variation in the stream of benefits provided by solar relative to the upfront costs of solar panels, which allows us to identify the discount rate; this is the same strategy as used in other work, such as the concurrent work by Snashall-Woodhams (2024).

To identify the marginal utility of income in conjunction with the discount rate, we further leverage household-level variation in irradiance. Here, the endogenization of installation size is important. If households were to install the same number of panels regardless of the irradiance, then separately identifying both the discount rate and marginal utility of income would not be feasible. However, households respond on the intensive margin to the costs of electricity. For example, many households install enough panels to offset their full consumption. In this situation, higher irradiance does not increase the stream of benefits but rather it lowers the upfront cost required for the same stream of benefits, since households will install fewer panels.[28]

Figure E.1 in the Appendix shows the distribution of the number of panels required for different levels of generation for our observed installations. This variation results from the variation in irradiance. Because a different number of panels is required for a given level of output and there is little incentive under NEM to oversize a system, we have variation in the upfront costs required to generate any given amount of solar electricity. Thus, variation in household-level irradiance provides variation in the total economic value of adopting solar, which allows us to also identify the marginal utility of income.

The previously shown Figure 4 provides a sense of the cross-sectional variation in irradiance that is providing variation in $q^*\bar{p}$ and K^*, along with the variation in electricity consumption and electricity rates (from the climate zones). A key takeaway from the figure is that while higher irradiance on a rooftop leads to more adoption, the rate of increase in adoption is much steeper for high-wealth and medium-wealth households than low-wealth households. In other words, greater payoffs over the life of the solar system increase adoption more for high-wealth and medium-wealth households than low-wealth households.

Our identification strategy differs from previous work estimating discount factors that draws entirely from policy changes that provide temporal shocks to future values (de Groote and Verboven, 2019; Bollinger, 2015). Such an approach is dependent on the variation in the estimated Pr' for identification of the discount rate (see equation (9) and recall that

[28]This raises the natural question of how small mispredictions of the optimal system size might affect the results. If the mispredictions only affect the upfront costs, this would be absorbed in our fixed effects, but effects on flow payoffs could influence the resulting discount rates. We run a robustness check rescaling irradiance by plus or minus 5% and find that for nearly all installations the flow payoff changes less than 2% from the baseline, suggesting that small mispredictions should not be a major concern.

the discount factor enters both θ and the Pr'). Our empirical setting intentionally covers a period with few changes *over time* in regulations or electricity rates that could affect Pr'.

5.1.2 Identification Assumptions

There are three key assumptions in our identification strategy for the implicit discount rate parameters. First, we assume that the profile of irradiance hitting each rooftop is exogenous with respect to the utility of adopting solar, conditional on household/home characteristics, electricity consumption, and a set of geographic fixed effects, including climate boundary zone fixed effects. Second, we assume that homes on one side of a climate zone boundary are similar to homes on the other, after conditioning on this same set of controls. We believe this second assumption is very reasonable given the careful choice of zip codes in our study. Third, we follow the literature in assuming rational expectations for state variables and we assume these expectations are the same for all wealth groups, an assumption we will discuss at length in Section 6.4.

One concern with our first identifying assumption could be endogenous sorting, whereby homeowners intentionally purchase homes that are better suited for solar in a way our controls do not already address.[29] We consider this to be a small concern given the high cost of California real estate in comparison to the value of solar, and the fact that solar installations were still relatively uncommon in the empirical setting. Furthermore, the average length of residence for households in the sample is 13.6 years, indicating that most households in our data purchased homes well before solar could have become a consideration.

Another potential concern with our first identifying assumption is the possibility that households manipulate their local environment to improve solar irradiance reaching their rooftop. For example, households could cut or trim trees. This would lead us to underestimate the fixed costs of adopting solar and would be an issue for us if households with smaller implicit discount rates are more likely to incur the same fixed cost than those with higher implicit discount rates. However, much of our sample has few tall trees in residential neighborhoods due to small lot sizes and relatively dry landscapes in California. In addition, we include wealth fixed effects, which should help address any selection issue in tree trimming/cutting across wealth bins.

[29] E.g. two-story homes have smaller roof areas holding square footage constant. Preferences for two-story homes may be correlated with unobservable tastes for solar. We include an indicator for two-story homes in our controls.

5.2 Estimation

We estimate the household-level adoption model via maximum likelihood estimation. If we could observe the purchase/lease type $e \in \{\text{lease, purchase}\}$ and consumption bin $b \in \{1, ..., 5\}$, then each household's contribution to the likelihood would be written simply as:

$$\mathbb{L}_{ibe} = \prod_t [Pr^e_{ibt}]^{y_{it}} [1 - Pr^e_{ibt}]^{(1-y_{it})}, \tag{11}$$

in which Pr^e_{ibt} is the probability of household i adopting in time t, conditional on being type e and bin b.

In practice, we do not observe the consumption bin b, nor the type e for non-adopting households. Each non-adopter household is one of ten possible consumption $\times$ type combinations $\{1, ..., 5\} \times \{\text{leaser, purchaser}\}$. We specify weights w_{ibe} as the probability that household i consumes in consumption bin b and is of type e. With weights w_{ibe}, we integrate the likelihood function over the unobserved heterogeneity:

$$\mathbb{L} = \prod_i \sum_b w_{ib} \sum_e w_{ie|b} \prod_t [Pr^e_{ibt}]^{y_{it}} [1 - Pr^e_{ibt}]^{(1-y_{it})} \tag{12}$$

which requires evaluating ten conditional likelihoods per observation of household i and time t and then integrating over consumption and lease types. In our main specification, we treat the combination of consumption bin b and type e as permanent, unobserved heterogeneity. The weights are based on ratios of individual- and zip-level likelihoods using the method laid out in Arcidiacono and Miller (2011); estimation details are in Appendix G.

6 Results

6.1 Parameter Estimates

The estimates of our structural model include a large set of coefficients, most of which are controls, such as demographic and voting variables. Table 2 presents our primary estimation results for key parameters and all of the remaining parameter estimates can be found in Appendix Table H.1. First, we present our estimated marginal utility of income (ω) across the wealth groups, as this parameter is very important for our counterfactuals and provides context for our implicit discount rate results. We observe that the high-wealth group has a much lower marginal utility of income, at 0.05, than the low-wealth group, at 0.49. The

medium-wealth group falls in between. The marginal utility of income is an ordinal rather than cardinal parameter, so we interpret relative values across different wealth groups. The low-wealth group has a marginal utility over nine times higher than the high-wealth group.

Table 2: Primary Estimation Results

Wealth	Marginal Utility of Income	Annual Discount Rate	$\bar{\bar{\theta}}$	Ratio
High	0.05 (0.02)	10.9% (0.03%)	48.9	1.5
Med	0.25 (0.19)	15.6% (13.35%)	34.9	1.1
Low	0.49 (0.4)	17.2% (0.25%)	31.7	1
All		14.5%	38.5	–

Heteroskedasticity-robust standard errors in parentheses, calculated by the delta method. The difference between the high-wealth and low-wealth underlying marginal utility of income structural parameters is statistically significant, with a p-value < 0.0001. The difference between the high-wealth and low-wealth underlying discount rate structural parameters is statistically significant, with a p-value < 0.0001.

Some of our most important results are the annual discount rates in Table 2.[30] We find that the high-wealth households have an annual implicit discount rate of 10.9%. Medium-wealth households have a somewhat higher rate of 15.6%. And low-wealth households have an even higher rate of 17.2%. All three of these estimates are significant, and the difference between the low-wealth and high-wealth groups are also significant (as can be seen in the coefficient α_{low} in Appendix Table H.1). The average rate for all groups is 14.5%. This finding is remarkably close to the average finding in de Groote and Verboven (2019) of 15%, estimated based on households in Belgium.

The final two columns in Table 2 provide a way to understand the implications of the implicit discount rate results. Recall from our model that θ converts an initial \$1 payoff from offset electricity costs each period over the panel lifespan (100 quarters) into a present value. This calculation takes into account the decrease in generation from panel degradation λ and the increase in savings from expected electricity price increases ζ_i before applying the implicit discount factor ρ_i. The difference between the high- and low-wealth groups is substantial, with high-wealth households valuing the \$1 flow payoff at \$48.9, while the low-wealth households value it at only \$31.7. The ratio of the value of the two is 1.5. This level of implicit discount rates and the remarkable difference in how consumers of different wealth groups value the future payoffs from solar have very important implications for our counterfactual results.

[30] We calculate these by transforming the structural coefficients, e.g., for high-wealth households $\Phi(\rho_i)^{-4} - 1$, since the time step is a quarter.

In Appendix I, we include the results from a range of robustness checks. These include allowing for rebound effects, alternative trends for power purchase rates for leases, alternative scaling and values for trends in solar and electricity prices, alternative assumptions for T, inclusion of race covariates, randomly re-assigning lease/purchase type among adopters, using a finite time horizon, and use of realizations for V' and Pr' instead of the predicted values. We consistently find that the stream of benefits from offset grid consumption is valued roughly 50% more by high wealth households relative to low wealth households across these specifications.

6.2 Price Elasticity of Demand

The implicit discount rate and the own price elasticity of demand are two of the most important parameters driving our counterfactuals. Our model is flexible enough to yield price elasticities of demand that vary by wealth group. We calculate the elasticities by simulating adoption (incorporating resizing) under a 1% reduction of the system cost over the sample period. Table 3 reports the own-price elasticities of demand by wealth group and lease type.

Table 3: Price Elasticity of Demand

Wealth	Rate	Installations	Purchases	Leases
High	10.9%	-0.5	-0.3	-0.9
Med	15.6%	-1.6	-1.3	-1.9
Low	17.2%	-2.4	-2.4	-2.4
All	14.5%	-1.4	-1.1	-1.8

The price elasticities range between -0.3 and -2.4. What is most notable is that the elasticities of demand for high-wealth households are much smaller than for low-wealth households. These price elasticities should be interpreted as the response over a quarter to a long-lasting price change, affecting both current and future prices. Thus, they would be expected to be slightly larger (in absolute value) than estimates in the existing reduced-form literature, which tend to rely on short-run variation in prices. Indeed, this is the case. Other studies, using different time frames and often other settings, tend to find price elasticities in the range of -0.65 to -1.2 (Hughes and Podolefsky, 2015; Gillingham et al., 2016).

6.3 Model Fit

Although examining price elasticities is one way to gauge the reasonableness of the model, we also examine the model fit. We predict adoption probabilities for each household, time period,

consumption bin, and type, and aggregate the probability of adopting within the study time period by wealth level. We can then compare our model-predicted aggregate probabilities to the data. Overall, we find that our model fits the data well, with total installations only overpredicted by 2.3%. The gap between predicted and data is slightly larger for purchased systems than leased systems, but remains relatively small. For more detail, see Appendix H.

6.4 Why are Implicit Discount Rates Heterogeneous?

In Section 2, we discussed possible explanations for heterogeneous implicit discount rates across wealth groups besides differences in the pure rate of time preference, including differences in hassle costs, borrowing costs, the marginal propensity to consume, beliefs, and risk aversion. We account for possible differences in risk aversion with heterogeneous marginal utility terms, and differences in the marginal propensity to consume are less likely in our setting.

Heterogeneity in hassle costs should be time invariant and would not scale in the size of the installation or irradiance. Thus, we would expect wealth-specific hassle costs to be absorbed by the wealth x boundary zone and wealth x lease/purchase FEs. The same would be true for time-invariant differences in information about upfront costs across wealth groups. In contrast, heterogeneity in borrowing costs might scale with installation size. There is evidence that present-biased individuals are more likely to have substantially more credit card debt (Meier and Sprenger, 2010), which might suggest that borrowing costs or constraints might matter for lower-wealth households, who appear to have a higher implicit discount rate.

Recall that in our primary specification we include the number of credit lines the consumer has access to in our primary specification as a potential control for borrowing costs. To more thoroughly test whether the differences in implicit discount rates are explained by a common proxy for borrowing costs, we leverage the individual-level Experian credit score data (see Appendix J.1 for details). We allow the household discount factor to be given by:

$$\begin{aligned}\rho_i &= \Phi(\rho^0 + \alpha^{med} 1\{\text{wealth}_i\text{=med}\} + \alpha^{low} 1\{\text{wealth}_i\text{=low}\} \\ &+ \alpha^{med,PC} 1\{\text{wealth}_i\text{=med}\}1\{\text{credit}_i\text{=poor}\} + \alpha^{low,PC} 1\{\text{wealth}_i\text{=low}\}1\{\text{credit}_i = \text{poor}\}),\end{aligned}$$

in which we use a dummy to indicate households with poor credit (near-prime or sub-prime).[31] We group all high-wealth households together since there are so few with poor credit scores.

[31] For households with missing credit data, we allow for similar flexible interactions between an indicator for missing data and wealth.

We find that medium- and low-wealth households with poor credit have implicit discount rates of 29.4% and 66.2%, respectively. Because credit scores are indicative of liquidity constraints and influence borrowing costs (Adams et al., 2009), this provides suggestive evidence that borrowing costs or limits matter for the implicit discount rates.

The estimated implicit discount rates for households with good credit scores are 15.9%, 15.9%, and 25.2% across high-, medium-, and low-wealth consumers.[32] The ratio of the low-wealth rate to the high wealth rate is 1.58, which is nearly the same as the ratio in our primary result. This suggests that when households have good credit, there must be another reason for the heterogeneity in implicit discount rates across wealth groups. This does not rule out other borrowing frictions or limits, as borrowing costs or limits are not solely determined by credit score.[33] Indeed, it is quite possible that borrowing limits or other frictions may play some role, along with differences in time preferences (perhaps due to behavioral anomalies).

Any differences across wealth in information frictions (the costs of accessing accurate information) that affect consideration of solar are likely to be time invariant, and would be similar to hassle costs in that they would be absorbed by our wealth x boundary zone and wealth x lease/purchase type fixed effects. Recall that we are leveraging within-wealth group variation in irradiance, which should be unrelated to information frictions. It may be possible that certain installers microtarget parts of neighborhoods in a way that is correlated with wealth, potentially leading to different amounts of information being provided to the different wealth groups. However, such geographic targeting of different wealth groups by installers would be largely picked up with our geography x wealth fixed effects, as the timing of any campaigns would be expected to be plausibly exogenous.

The solar purchase decision depends on consumers' expectations over future electricity rates given their initial consumption and climate zone (ζ), panel degradation (λ), and panel life (T). Under the standard assumption of common rational expectations, biased consumer beliefs would play no role. But systematically different expectations by different wealth households could influence the implicit discount rates. Allcott and Taubinsky (2015) provided evidence showing biases in the expectations of savings from CFL lightbulbs, although they do not examine the link with demographic characteristics.

We view systematically different beliefs for ζ across wealth groups as unlikely for two reasons. First, our sample consists of homeowners in California, who all have at least a reason-

[32]These estimated rates are somewhat higher than our main results. We attribute this difference to the incomplete matching of credit scores and thus a different subsample.

[33]There is evidence that credit scores alone are not sufficient in predicting default risk (Albanesi et al., 2022; Albanesi and Vamossy, 2024). van Rijn et al. (2021) show that income and net worth are predictive of lower loan rates for automobiles.

able degree of wealth, so differences in electricity rate expectations driven by systematically different views on utility pricing seem unlikely. Second, we estimate consumption-specific measures of ζ, yet the distribution of electricity consumption is not particularly different across wealth groups (see Figure A.1b), greatly reducing the likelihood of differing expectations over wealth. Similarly, systematic biases that differ across wealth groups in the panel depreciation rate λ are highly unlikely as these are engineering estimates that come straight from the installers and are easily found online.

There are two possible reasons to think there might be heterogeneity in the T across wealth groups. The first is that consumers in different wealth groups could have different expectations about panel lifespans. However, given that panels tend to have 25-year warranties and that contractors give the same information about panel lifespans to all consumers, one would expect all consumers to have the same expectations for how long panels will continue to produce electricity.

The second possible reason is that some consumers may have different beliefs about how long they will live in their home than others. For example, households with lower wealth may believe they are more likely to move before the 25-year lifespan of the solar panels. However, evidence from a nationwide survey in Bollinger et al. (2025) of 1,673 homeowners who either are considering solar panels or EVs shows that the fraction of homeowners who plan on staying in their home more than 20 years is nearly identical across income groups (Figure J.1). In general the expectations look similar across income groups (see Appendix J.2).[34] Even if such expectations were different, prior work has shown that the capitalization for solar is very high (Dastrup et al., 2012; Qiu et al., 2017; Gillingham and Watten, 2024); since capitalization of other housing attributes (swimming pools, fireplaces, garages, and air conditioning) does not seem to vary with wealth (Sirmans et al., 2006), we view it unlikely that heterogeneous beliefs about capitalization of solar panels explains our effect.[35] Indeed, solar installers generally tell all of their customers that they should expect their home values to rise along with the value of the system. All of this evidence suggests that different beliefs across wealth groups is less likely to be a major factor in explaining the heterogeneity in implicit discount rates.

We further explore systematically different beliefs by performing a bounding exercise. Specifically, we calculate equivalent values of each of the three parameters $\{\zeta, \lambda, T\}$ that

[34]Note that California homeowners are unlikely to be in the lowest income category due to the high cost of homes in the state.

[35]Gillingham and Watten (2024) present evidence that low-income households may have lower capitalization of the residual value of a solar system in the home transaction price than high income households. However, this assumes the same discount rate in calculating the net present value as high-income households and if the discount rate for low-income is higher, then capitalization could be very similar.

would result in our estimated θ but constant ρ over wealth groups.[36] For example, for ζ, we take estimates of the quarterly discount factor $\rho = \{0.974, 0.964, 0.961\}$ and using (4), we solve for the value of $\tilde{\zeta}$ such that $(1+\tilde{\zeta})(1-\lambda)\rho^{low} = (1+\zeta)(1-\lambda)\rho^{high}$. While our estimated ζ implies an annual electricity rate increase of 2.7%, the $\tilde{\zeta}$ required to equate the low-wealth and high-wealth implicit discount rates would be -2.8%, which is clearly unreasonable. An equivalent exercise for $\tilde{\lambda}$ would require that low-wealth households believe that panel generation would decline 6.6% per year, a factor of 8.2 times the engineering estimate of 0.8%. The same exercise for $\tilde{T}$ yields the result that low-wealth households would have to expect an effective panel lifetime $\tilde{T} = 11.2$ to fully explain the differences in implicit discount rates, which is less than half of the 25-year warranty for solar systems. This seems unlikely, as it would require low-wealth households to have dramatically shorter expectations of T than high-wealth households or beliefs of much lower capitalization, neither of which is supported by the evidence available. Thus, we cannot entirely rule out that different beliefs across wealth groups contribute in some way to the heterogeneity in implicit discount rates, but the evidence suggests that different beliefs are not likely to be the primary explanation.

7 Counterfactuals

We explore two counterfactual policies. First, we reduce the compensation for excess solar generation fed back into the grid. Instead of compensation at the retail rate, as under standard NEM, we use an estimated levelized avoided cost of energy of \$0.062/kWh.[37] This first counterfactual roughly approximates some of the changes under California's NEM 3.0 reform. Second, we again reduce the compensation rate for excess solar generation, but accompany it with an upfront subsidy. Under both scenarios, we hold solar prices fixed. The market for solar modules is global, so module prices are unlikely to adjust. Installers could lower their prices in response to the reduced demand, although we would have to assume a model of competition to calculate this response. We discuss this possibility further when presenting the welfare results.

The upfront subsidy in the second counterfactual is chosen to keep the total capacity of solar installed roughly the same in the counterfactual as in the observed data, assuming 100%

[36] That the consumer is correct in expectation over future values of prices is a common assumption in the literature even when identifying heterogeneous preferences. For instance, Gowrisankaran and Rysman (2012) estimate a similar linear autoregressive specification for camcorder prices and incorporate it into a model with persistent consumer heterogeneity. See also Bayer et al. (2016) and Hornbeck et al. (2024).

[37] This estimate is from a typical solar generation profile in the California Public Utilities Commission Avoided Cost Calculator derived by E3 Consulting (E3, 2022)

pass-through.[38] This counterfactual uncovers the potential efficiency and equity improvements possible from changing the policy, while avoiding political fallout (and reduced environmental gains) from diminishing the solar industry.

Before turning to the results, it is worth a brief discussion of how we estimate the outcomes in our counterfactuals. Since state transitions may change under our counterfactuals, we follow similar assumptions as in our parameter estimation, allowing consumers to expect a deterministic evolution of the state variables with a short-run prediction error (see Appendix K). We assume that the policy changes are expected to remain in place for the future. We use the optimal solar system sizing model to incorporate how changes in policy affect system sizing. The baseline that we compare our policy counterfactuals to is the observed data. One limitation of our approach is that leasing is modeled as time-invariant heterogeneity, so it does not change in the counterfactual, although we would expect this to be a minor factor, as leasing is prevalent in all wealth groups and appears to be driven more by other factors outside the relative financial benefits of leasing versus purchasing.

7.1 Lowering Compensation for Exported Solar

Lowering the compensation for excess solar electricity fed into the grid to \$0.062/kWh leads to a leftward shift of the demand curve for solar installations because the benefits of solar are reduced. One challenge in running this counterfactual is that the change in the net present value of a given system depends on how much solar electricity is consumed by the household when the solar generation occurs (and thus offsets buying electricity at the retail rate) versus being exported to the grid. For example, at one extreme, one could imagine a household that uses all of the electricity generated by the solar system, so that none of the electricity is fed into the grid. For this household, a reform that lowered the compensation rate for excess electricity would not change the net present value of adopting solar. At the other extreme, a household that fed nearly all of the electricity into the grid would have a much lower solar net present value after the reform. We report our results for several potential "splits" between solar generation that offsets consumption and solar generation that is exported. Specifically, we report $\{0, 30, 50, 70, 99\}$ percent exported. For each split, we allow for re-sizing, then predict adoption conditional on the resized flow benefits and up-front costs.

Table 4 presents the first results of the counterfactual analysis, showing the percent change in solar adoption for each of the three wealth groups and the entire sample.[39] Each row

[38] Under a specific model of competition we could calculate the optimal pass-through rate and use this instead.

[39] We hold solar prices fixed, but if we were to allow solar prices to adjust downward, these declines would be lower, although we do not expect it would be by much given how large the global market is.

Table 4: Percent Change of Solar Installations in Counterfactual Relative to Baseline

	% solar exported				
Wealth	0	30	50	70	99
High	0	-3	-8	-14	-23
Med	0	-6	-12	-17	-21
Low	0	-5	-7	-10	-13
All	0	-5	-9	-14	-19

shows the percent change in adoptions in the counterfactual relative to the baseline and the columns indicate the percent of solar generation exported to the grid. When none of the solar is exported to the grid, the counterfactual is identical to the baseline.

When 99% of solar generation is exported to the grid (last column in Table 4), we observe that adoption falls by 23% among high-wealth households and by only 13% among low-wealth households (Table 4). The larger decline for high-wealth households reflects their greater sensitivity to reduced NEM benefits, driven by lower implicit discount rates and higher marginal electricity prices under the tiered rate structure. High-wealth households tend to consume more electricity and thus tend to be on a higher electricity rate tier, so their bill offset under current NEM policy is larger.

As we reduce the percent export from 99%, the tiered rate structure for electricity prices becomes even more important. For example, with 50% of the electricity consumed by the household, more higher-wealth households are going to be interested in solar under the counterfactual because they tend to consume more electricity at a higher electricity rate tier. This force explains why the high-wealth group has the same percent reduction in solar installations as the low-wealth group, at approximately 8%. Indeed, when 30% of solar is exported, this second force dominates and the percent reduction for high-wealth households is *lower* than for low-wealth households.

If solar systems are sized to fully or nearly fully offset annual electricity consumption, we would anticipate that approximately 30-50% of electricity being exported for a typical household whose members generally leave the home for work during the day. The actual amount is unobservable and there is seasonality, given that most solar systems overproduce relative to consumption during the summer and underproduce in the winter.

7.2 Replacing NEM Compensation with Upfront Subsidy

In our second counterfactual, we reduce the compensation for excess solar generation as in the first counterfactual but we also add an additional upfront subsidy, swapping a flow

payoff for an upfront benefit. There are two common variations in the design of upfront solar subsidies. One approach is to provide an incentive based on the system size (a per-watt capacity subsidy). Under a capacity subsidy, households installing larger systems–who are often wealthier households–would receive greater subsidies than those installing smaller systems. Alternatively, the government could provide a fixed upfront subsidy that is the same for all solar consumers. Fixed upfront subsidies are less common, but they might have preferable distributional consequences. We explore a per-watt capacity subsidy and a split policy in which 50% of the upfront subsidy is per-watt and 50% is a fixed upfront subsidy that everyone receives.

We calculate the subsidy required (assuming 100% pass-through) to ensure that the total number of adoptions among high-wealth households remains roughly equal to the baseline number under the split policy. One reason for ensuring that solar adoptions remain the same for high-wealth households might be political constraints—since medium- and low-wealth households are more responsive, the subsidy is the smallest necessary to keep adoption rates greater than zero for all wealth groups. For each solar consumer, we estimate the additional subsidy cost to the government and the savings to the utility. Specifically, we calculate the change in the net present value from reducing the NEM compensation and then calculate the net present value of the utility savings (using a 6% implicit discount rate for installers, aligned with our lease model). If we wanted to allow for incomplete pass-through, government expenditures would have to increase by some additional amount (depending on the pass-through rate) which would be directly transferred to the installers.

We find that roughly half of the savings from the reduced NEM compensation would be sufficient to hold the total solar adoptions by high-wealth households constant if it is allocated to upfront subsidies, using both a 30% and 50% grid export rate. If 30% of generation is exported to the grid, then the subsidies required are \$731 per installation and \$0.15 per watt. For 50% of generation exported, the subsidies are \$1,961 and \$0.42 per watt.[40] Table 5 presents the counterfactual results for each wealth group and for the total population.

In Table 5, we present the percent change from the baseline in solar installations and solar capacity for each of the three scenarios: no subsidy (repeating the Table 4 results for reference), savings from reducing NEM compensation used for an upfront capacity-based per-watt upfront subsidy, and savings used for the split capacity and per-installation subsidy. The total aggregate value of the upfront subsidy is identical between the capacity subsidy and the split subsidy. The capacity-based subsidy is \$0.30/W for the 30% export rate case and \$0.84/W for the 50% export case.

[40]The fraction rebated is approximately 50% of the revenues no longer used for NEM, but varies somewhat due to the resizing that occurs in response to different subsidy levels.

Table 5: Percent Change of Solar Installations and Capacity in Second Counterfactual Relative to Baseline

	No subsidy		Per-W		Per-install and per-W	
Wealth	Adopted	Capacity	Adopted	Capacity	Adopted	Capacity
30% exported to grid						
High	-3	-6	0	-2	0	-3
Med	-6	-8	3	2	3	1
Low	-5	-8	6	5	8	5
All	-5	-7	2	2	3	1
50% exported to grid						
High	-8	-18	0	-1	0	-6
Med	-12	-20	10	10	10	5
Low	-7	-20	19	19	27	18
All	-9	-19	9	8	11	4

"Adopted" refers to number of adoptions. "Capacity" refers to kW installed.

We find that the upfront subsidy leads to more solar installations by low-wealth and medium-wealth households than in the baseline, especially at the 50% export rate. A key reason for this finding is the heterogeneity in implicit discount rates, as the upfront subsidies are not discounted while the current NEM policy benefits would be. The heterogeneity in the marginal utility of income (see Table 2) also contributes. The increase in adoptions by low-wealth households is even more notable under the split subsidy. Low-wealth households tend to have smaller systems, so a per-installation subsidy will provide greater benefits to these households than a capacity-based subsidy.

Another notable finding in Table 5 is a difference between the change in the number of installations and solar capacity across scenarios. For example, in the 50% export case, low-wealth households observe a 27% increase in installations but only a 18% increase in solar capacity. The explanation here is that, even after optimal resizing, solar system sizes for low-wealth households are substantially smaller than for higher-wealth households. This finding highlights the value of modeling system sizes in our analysis.

7.3 Welfare and Marginal Value of Public Funds

Changes in welfare-relevant outcomes are crucial inputs to the policy process, but are complicated in the electricity market. There are distortions from raising revenue either through taxes or by increasing electricity rates. Further, the policy changes affecting utility

revenues can lead to adjustments to the electricity rate (Hahn and Metcalfe, 2021), including changes to the entire tiered rate structure. Given the difficulty in ascertaining how such changes might occur, our analysis holds the electricity rate structure fixed (future work could relax this by estimating a new equilibrium). Thus, our analysis should be seen as short-run or medium-run, prior to any new rate cases that would revise electricity rates.

In Table 6, we report changes in welfare-relevant outcomes for each wealth group and for the entire sample for a 50% export rate (see Table L.3 for a 30% rate). The numbers for each wealth group are for the outcomes attributed to the solar adoption decisions of each of the wealth groups, rather than the welfare-relevant effects on that group, as in a standard distributional analysis. In the second column, we calculate the gross post-subsidy changes in consumer surplus by integrating under each consumers' *new* demand curve, inclusive of the subsidies and accounting for the possibility of system resizing. The change in the excess solar generation compensation rate shifts each household's demand curve for adopting solar to the left, and the size of the shift depends on their discount rate, their irradiance curve, their electricity consumption, and the initial tiered pricing structure. We carefully account for heterogeneity by calculating the welfare effect for each household (integrating over the unobserved consumption state with the estimated weights) and then aggregating, rather than assuming a representative consumer.

In column three, we calculate the changes in pre-tax installer surplus. For this calculation, we need an estimate of the installer margin. We draw upon existing data for this estimate. Using a module cost of \$0.67 per watt and inverter cost of \$0.32 per watt (US Department of Energy, 2015) plus labor cost of \$0.33 per watt and permitting cost of \$0.19 per watt (Chung et al., 2015), combined marginal costs in 2014 would be \$1.51. At the average solar price in our data in Q1 2014 of \$5.12 per watt, this implies a 71% margin. This estimated margin also aligns nearly perfectly with the implied margin using our estimated elasticity of -1.4 under the assumption of a single installer (we aggregate our installers) with monopolistic pricing.[41]

In the fourth column Table 6, we show the changes in the pre-tax utility surplus (revenues net of avoided cost of generation) from both the changes in NEM subsidies and the change in solar adoption, integrating again over all households. In column five, we present the aggregate avoided environmental and health damages resulting from changes in solar installations using the monetized values of the externalities in Sexton et al. (2021).[42] Finally, we show the

[41]For a comparison of our elasticity to the literature, Gillingham and Tsvetanov (2014) finds an elasticity of -0.65, Pless and Van Benthem (2019) find an elasticity of -0.85 for purchased systems, and Hughes and Podolefsky (2015) find an estimate of approximately -1.2.

[42]We use the "dollarization" method from this paper and do not split up the calculation by wealth group as this is infeasible. We use a social cost of carbon set at \$175/ton in \$USD2017 (Rennert et al., 2022).

Table 6: Changes in Welfare-Relevant Outcomes (Millions $; % Change)

Wealth	Δ Consumer Surplus	Δ Installer Surplus	Δ Utility Surplus	Δ Avoided Damages	Δ Government Subsidies
CF1: Lowered Compensation for Exported Solar					
High	-$13.9 (-49%)	-$8.2 (-21%)	$19.2 (54%)	-$7.9 (-22%)	-$3.5 (-21%)
Med	-$8.2 (-30%)	-$11.3 (-25%)	$23.6 (59%)	-$11.0 (-26%)	-$4.7 (-25%)
Low	-$2.3 (-17%)	-$5.5 (-23%)	$11.4 (60%)	-$5.5 (-25%)	-$2.3 (-23%)
All	-$24.4 (-35%)	-$25.0 (-23%)	$54.2 (57%)	-$24.3 (-24%)	-$10.5 (-23%)
CF2: Lowered Compensation for Exported Solar plus Split Subsidy					
High	-$9.8 (-30%)	-$6.5 (-16%)	$14.3 (35%)	-$2.5 (-6%)	$7.8 (29%)
Med	-$1.2 (-3%)	-$2.5 (-5%)	$11.9 (23%)	$2.5 (4%)	$13.1 (36%)
Low	$3.1 (16%)	$2.6 (8%)	$2.9 (10%)	$5.0 (15%)	$9.9 (45%)
All	-$7.9 (-9%)	-$6.4 (-5%)	$29.0 (24%)	$5.0 (4%)	$30.9 (36%)
CF2-CF1: Add the Split Subsidy to Lowered Compensation					
High	$4.1 (19%)	$1.7 (5%)	-$4.9 (-19%)	$5.4 (16%)	$11.3 (50%)
Med	$7.0 (27%)	$8.8 (20%)	-$11.7 (-36%)	$13.5 (30%)	$17.8 (61%)
Low	$5.4 (34%)	$8.0 (31%)	-$8.5 (-49%)	$10.5 (40%)	$12.2 (67%)
All	$16.5 (26%)	$18.6 (18%)	-$25.2 (-33%)	$29.4 (28%)	$41.4 (59%)

All estimates assume 50% export to grid.

change in government expenditures due to differences in payouts from the ITC and subsidies. Further details on our calculations can be found in Appendix L.

We find that lowering the compensation for exported solar to the avoided cost reduces consumer surplus (in the sample) by $24.4 million, reduces installer surplus by $20.6 million, increases utility surplus by $54.2 million due to the reduction of bill offsets for net solar generation, reduces avoided damages by $24.3 million, and reduces government tax credit expenditures by $10.5 million. The bulk of the impacts are due to the decisions of the high-wealth group.[43]

Under the split subsidy, aggregate consumer surplus declines by -$7.9 million from baseline, but increases for the low-wealth households by $3.1 million. This demonstrates a clear equity tradeoff. Government subsidies increase by $30.9 million, installer surplus declines by $5.3 million, and utility pre-tax surplus increases by $29.0 million. Avoided environmental and health externalities increase by $5.0 million.[44] If additional subsidies were needed to achieve

[43]If we were to allow solar prices to optimally adjust downward towards zero under this counterfactual, these changes would all be smaller since the reduction in the number of installations would not be as large. Recall that installers would only decrease prices if it increased their profits.

[44]Given that total installations increase in this scenario, if solar prices were allowed to adjust they would do so upwards, attenuating these numbers slightly.

the reduction in post-subsidy prices due to incomplete pass-through, then this would mean government expenditures would have to be higher to achieve the same price change and thus same demand response, an additional transfer from government to installers.

The key takeaway here is that switching from an NEM subsidy to an upfront subsidy can maintain the level of the solar market, maintain the environmental and health benefits, improve the consumer surplus for low-wealth households, and reduce utility expenditures (which could reduce the cost-shift and potentially also benefit low-wealth households).[45]

7.3.1 Marginal Value of Public Funds

One promising approach to gain insight into the overall welfare effects of subsidy counterfactuals without making assumptions about the deadweight loss from raising funds is to calculate the marginal value of public funds (MVPF) (Finkelstein and Hendren, 2020). The MVPF is the ratio of the policy beneficiaries' willingness to pay for a policy to the government expenditures of the policy. It is a measure of the effectiveness of government expenditures. The gains to the beneficiaries include the change in consumer and installer surplus, electric utility expenditures (accounting for taxation, which goes to the government), plus the externality damages avoided. The government expenditures account for the changes in the incentive payments shown in Table 6, and also account for changes in tax revenues. A MVPF of greater than one indicates that the government expenditure is less than willingness to pay, providing one indication that the policy is net beneficial.[46]

We begin by calculating the MVPF from moving from the first counterfactual to the baseline. In this calculation, the only change in government expenditures is the change in the fiscal outlays for the federal solar investment tax credit. We make the calculation using the 50% export to grid assumption with 100% subsidy pass-through from installer. We find the average MVPF is 1.71, with a MVPF of 2.57, 1.42, and 0.96 for high-, medium-, and low-wealth households, respectively.[47] These surprisingly large MVPF estimates come about

[45]Recall that our final sample consists of 183,667 households, which represents 3.2% of all owner-occupied single-family detached homes in California, so the results would have to be scaled up by a factor of 31 to cover all of the state.

[46]We follow Hahn et al. (2024) in how we handle the installers' surplus, electricity markup, corporate taxes for installers and the electric utility, and government expenditures, using the same assumptions (a 10% corporate tax rate and the 8% economy-wide markup found by De Loecker et al. (2020)). Our calculations would be the same with government-mandated transfers of the post-tax utility savings back to consumers – pre-tax transfers would lead to small changes in the effect on government revenues. Another possible transfer mechanism is a small downward adjustment in utility rates, which would have a secondary effect of reducing installations to some degree since the value of the offset utility bills would be reduced. We anticipate that the MVPF calculations would not change appreciably.

[47]Table M.1 shows the adjusted surplus numbers assuming a 50% export rate (Table M.2 does the same with 30%) for the installers, utility, and government after accounting for the fiscal externalities (the contributions of the added subsidy relative to the economy wide effect and for corporate taxes).

largely because the change in utility surplus is not counted as government expenditure and the change in government expenditure is a secondary effect of the policy change. So this comparison is an unusual MVPF calculation. But it indicates that the marginal social value is positive for the solar subsidies, and that it is over twice as large for subsidies that go to the high wealth households, due to the fact that higher-wealth households would be receiving greater utility from the NEM subsidy.

The findings are slightly different when comparing the first counterfactual to the second counterfactual, which pairs the lower compensation with an upfront subsidy. The average MVPF of adding the upfront subsidy is 0.94, with values of 0.57, 0.97, and 1.23 for high-, medium-, and low- wealth households, respectively.[48] The utility surplus, which is distributed across all utility customers, offsets much of the consumer and installer surplus, so most of the surplus gains are from the avoided externality damages. The most notable finding is that the relative MVPF across wealth bins flips relative to the MVPF of subsidies at the higher compensation rate. Additional subsidies that go to low-wealth households at the lower compensation rate provide more than twice the social value as those that go to high-wealth households.

As robustness checks, we also considered how incomplete or overshifting pass-through would influence the MVPF when adding the additional subsidy in the second counterfactual. We use estimates of 78.7% and 133.3% from Pless and Van Benthem (2019). We also repeat the calculation using a 30% export rate. Results are in Appendix M.3. Because our subsidy is set to keep installation counts approximately equal to baseline, reducing the pass-through rate attenuates MVPF toward one as it represents only a transfer from government to installer surplus with some fiscal externalities, which we calculate. The effect is small, with the 78.7% pass-through changing MVPF for low-wealth households from 1.23 to 1.20, and pass-through of 133.3% increasing low-wealth households MVPF to 1.34. High- and medium-wealth households all remain under 1.0.

Dynamic benefits on the supply side such as learning-by-doing (LBD) would also affect these calculations.[49] Given that Bollinger and Gillingham (2023) find only small LBD in the installation of solar during the early expansion of solar PV in California, we do not expect LBD to affect these numbers too substantially, although it would serve to increase the MVPF. Subsidies could also incentivize innovation in manufacturing, but since the United States

[48]Table M.4 shows the adjusted surplus numbers assuming a 50% export rate and 100% pass-through for the installers, utility, and government after accounting for the fiscal externalities.

[49]Hahn et al. (2024) examine LBD specifically in their Appendix B. The dynamic price reduction on the supply side would lead to higher demand and more spending on incentives. They show that since dynamic pricing incentives from LBD decline with installations under an isoelastic cost specification, the additional terms in the calculation would decline over time.

has not historically seen large investments in domestic solar panel production relative to the overall size of the market (Gerarden et al., 2025), it is unlikely that an additional subsidy of the magnitude we propose here would change these numbers notably.

8 Conclusion

In this paper, we use a novel identification strategy that leverages plausibly exogenous variation in irradiance to identify implicit discount rates used by households in California in their solar installation decisions. We provide model-free evidence that high-wealth and medium-wealth households are more responsive to the value of the solar flow payoff obtained by investing in solar panels than low-wealth households, highlighting the features of the data that drive our results. We then estimate a structural dynamic model that pins down an average implicit discount rate for residential solar adopters of 14.5% and demonstrates clear heterogeneity in the implicit discount rate by wealth, with high-wealth households making decisions consistent with a 10.9% rate and low-wealth households a 17.2% rate. Our results help explain disparities in solar adoption rates between low-wealth and high-wealth households.

We also explore why these differences might occur with credit score data. We find that households with poor credit have higher implicit discount rates, indicating that credit scores matter. But the same 1.5 ratio of the high-wealth to low-wealth rates remains even for households with solid credit, suggesting that the *disparity* in rates between groups is not due solely to credit scores. This is important because if borrowing costs are the only explanation for the heterogeneity in discount rates, then attention would primarily turn to credit-related policies instead of subsidies.[50] We find that the best evidence available renders other potential explanations for the heterogeneity in implicit discount rates less likely, suggesting that credit and differences in the pure rate of time preference are the more likely explanations.

Our counterfactual results indicate that reducing NEM compensation to the avoided cost would reduce the size of the solar market substantially. It would reduce installations across all wealth groups, but would reduce consumer surplus of the high-wealth households much more than low-wealth households. If roughly half of the savings from reducing NEM compensation are allocated to an upfront subsidy (half per installation and half per capacity installed), this would maintain the size of the solar market, reduce government and utility expenditures, and even lead to an increase in consumer surplus for low-wealth households. In short, by carefully structuring the upfront incentives, there is the potential to improve equity and reduce any

[50] Of course, our findings do not rule out the possibility that borrowing may matter or that policies focused on credit, such as low-interest loan programs for low-wealth households, could also increase solar uptake.

potential cost-shift. Such a subsidy results in a MVPF of greater than one for low-wealth households, but less than one for high- and medium-wealth.

California and many other states are reforming NEM policies to better reflect the actual value of solar generation, while the federal ITC is in the process of being phased out shortly. Given these major changes to the financial benefits of adopting solar, it is useful for policymakers to understand how consumers' valuation of different solar subsidies vary by household wealth. Many other green durables have similar characteristics with upfront costs and future fuel savings, and future work may show that our results generalize further. Our study aims to provide a contribution to the economics literature through well-identified estimates of the implicit discount rate across wealth groups, insights into the mechanisms, and guidance to policymakers on the distributional effects of different policy approaches.

References

Adams, William, Liran Einav, and Jonathan Levin, "Liquidity constraints and imperfect information in subprime lending," *American Economic Review*, 2009, *99* (1), 49–84.

Albanesi, Stefania and Domonkos F Vamossy, "Credit scores: Performance and equity," Technical Report, National Bureau of Economic Research 2024.

_ , Giacomo DeGiorgi, and Jaromir Nosal, "Credit growth and the financial crisis: A new narrative," *Journal of Monetary Economics*, 2022, *132*, 118–139.

Aldy, Joseph E, Todd D Gerarden, and Richard L Sweeney, "Investment versus output subsidies: Implications of alternative incentives for wind energy," *Journal of the Association of Environmental and Resource Economists*, 2023, *10* (4), 981–1018.

Allcott, Hunt and Dmitry Taubinsky, "Evaluating Behaviorally Motivated Policy: Experimental Evidence from the Lightbulb Market," *American Economic Review*, August 2015, *105* (8), 2501–2538.

_ and Michael Greenstone, "Is there an energy efficiency gap?," *Journal of Economic Perspectives*, 2012, *26* (1), 3–28.

_ , Christopher Knittel, and Dmitry Taubinsky, "Tagging and targeting of energy efficiency subsidies," *The American Economic Review*, 2015, *105* (5), 187–191.

Andreoni, James, "Impure altruism and donations to public goods: A theory of warm-glow giving," *The Economic Journal*, 1990, *100* (401), 464–477.

Arcidiacono, Peter and Paul B. Ellickson, "Practical Methods for Estimation of Dynamic Discrete Choice Models," 2011, *3* (1), 363–394.

_ **and Robert A Miller**, "Conditional choice probability estimation of dynamic discrete choice models with unobserved heterogeneity," *Econometrica*, 2011, *79* (6), 1823–1867.

_ **and Robert A. Miller**, "Identifying Dynamic Discrete Choice Models off Short Panels," *Journal of Econometrics*, April 2020, *215* (2), 473–485.

Barsky, Robert B, F Thomas Juster, Miles S Kimball, and Matthew D Shapiro, "Preference parameters and behavioral heterogeneity: An experimental approach in the health and retirement study," *The Quarterly Journal of Economics*, 1997, *112* (2), 537–579.

Bayer, Patrick, Robert McMillan, Alvin Murphy, and Christopher Timmins, "A Dynamic Model of Demand for Houses and Neighborhoods," *Econometrica*, 2016, *84* (3), 893–942.

Berkouwer, Susanna B and Joshua T Dean, "Credit, attention, and externalities in the adoption of energy efficient technologies by low-income households," *American Economic Review*, 2022, *112* (10), 3291–3330.

Berry, Steven, James Levinsohn, and Ariel Pakes, "Automobile Prices in Market Equilibrium," *Econometrica*, 1995, *63* (4), 841–890.

Bollinger, Bryan, "Green technology adoption: An empirical study of the Southern California garment cleaning industry," *Quantitative Marketing and Economics*, 2015, *13*, 319–358.

_ **and Kenneth Gillingham**, "Learning-by-Doing in Solar Photovoltaic Installations," 2023.

_ **,** _ **, Sonia Gupta, Hal Hershfield, and Christopher Pike**, "Renewable Energy Technology Increases Place Attachment," *Yale University Working Paper*, 2025.

Borenstein, Severin and Lucas W Davis, "The distributional effects of US clean energy tax credits," *Tax Policy and the Economy*, 2016, *30* (1), 191–234.

_ **and** _ , "The distributional effects of US tax credits for heat pumps, solar panels, and electric vehicles," Technical Report, National Bureau of Economic Research 2024.

Buchsbaum, Jesse, "Long-run price elasticities and mechanisms: Empirical evidence from residential electricity consumers," *Energy Institute Working Paper*, 2022, *331.*

Chatterjee, Satyajit and Burcu Eyigungor, "Understanding Contract Hetergeneity in the Credit Card Market," Technical Report, Discussion paper 2024.

_ **, Dean Corbae, Kyle Dempsey, and José-Víctor Ríos-Rull**, "A quantitative theory of the credit score," *Econometrica*, 2023, *91* (5), 1803–1840.

Chung, Donald, Carolyn Davidson, Ran Fu, Kristen Ardani, and Robert Margolis, "US photovoltaic prices and cost breakdowns. Q1 2015 benchmarks for residential, commercial, and utility-scale systems," Technical Report, National Renewable Energy Lab.(NREL), Golden, CO (United States) 2015.

Cohen, Jonathan, Keith Marzilli Ericson, David Laibson, and John Myles White, "Measuring time preferences," *Journal of Economic Literature*, 2020, *58* (2), 299–347.

Colas, Mark and Emmett Saulnier, "Optimal Subsidies for Residential Solar," 2024.

Coller, Maribeth and Melonie B Williams, "Eliciting individual discount rates," *Experimental Economics*, 1999, *2*, 107–127.

Conlon, Christopher and Jeff Gortmaker, "Incorporating Micro Data into Differentiated Products Demand Estimation with PyBLP," Technical Report, NYU working paper 2023.

CoreLogic, Inc., "Property Tax and Deeds Dataset 1990-2018," `https://corelogic.com` 2018. Accessed: January 2018; Compiled from county recorder offices.

Cubitt, Robin P and Daniel Read, "Can intertemporal choice experiments elicit time preferences for consumption?," *Experimental Economics*, 2007, *10*, 369–389.

Dastrup, Samuel R, Joshua Graff Zivin, Dora L Costa, and Matthew E Kahn, "Understanding the Solar Home price premium: Electricity generation and "Green"" social status," *European Economic Review*, 2012, *56* (5), 961–973.

de Groote, Olivier and Frank Verboven, "Subsidies and time discounting in new technology adoption: Evidence from solar photovoltaic systems," *American Economic Review*, 2019, *109* (6), 2137–72.
E3 Consulting (E3)

E3 Consulting (E3), "Avoided Cost Calculator v1a," Technical Report, E3 Consulting, San Francisco, CA 2022. Online at `https://e3.sharefile.com/share/view/s3fdd4ff8b9db4e95904726427ae54e81/fo5e3775-f03c-4008-be45-9e8d3e765955`. Last visited 2022-07-01.

Experian plc, "Experian Individual-Level Consumer Credit Panel: VantageScore Microdata," `https://experian.com` 2025. Accessed: August 1, 2025; Anonymized individual-level household credit records.

Feger, Fabian, Nicola Pavanini, and Doina Radulescu, "Welfare and redistribution in residential electricity markets with solar power," *The Review of Economic Studies*, 2022, *89* (6), 3267–3302.

Finkelstein, Amy and Nathaniel Hendren, "Welfare analysis meets causal inference," *Journal of Economic Perspectives*, 2020, *34* (4), 146–167.

Gerarden, Todd D, Bryan K Bollinger, Kenneth Gillingham, and Daniel Xu, "Strategic Avoidance and the Welfare Impacts of US Solar Panel Tariffs," Technical Report, National Bureau of Economic Research 2025.

_ , Richard G Newell, and Robert N Stavins, "Assessing the energy-efficiency gap," *Journal of Economic Literature*, 2017, *55* (4), 1486–1525.

Gibbs, Christa N, Benedict Guttman-Kenney, Donghoon Lee, Scott T Nelson, Wilbert H van der Klaauw, and Jialan Wang, "Consumer credit reporting data," Technical Report, National Bureau of Economic Research 2024.

Gillingham, Kenneth and Erica Myers, "Economics of Energy Efficiency," *Handbook on Electricity Regulation*, 2025.

_ **and Karen Palmer**, "Bridging the energy efficiency gap: Policy insights from economic theory and empirical evidence," *Review of Environmental Economics and Policy*, 2014.

_ **and Tsvetan Tsvetanov**, "Hurdles and steps: Estimating demand for solar photovoltaics," Technical Report, Working Paper, Yale University 2014.

_ **and** _ , "Estimating Demand for Solar Photovoltaics," 2017.

_ **, Hao Deng, Ryan Wiser, Naïm Darghouth, Gregory Nemet, Galen Barbose, Varun Rai, and CG Dong**, "Deconstructing solar photovoltaic pricing," *The Energy Journal*, 2016, *37* (3).

Gillingham, Kenneth, Kirkpatrick, A Justin and Bollinger, Bryan, "Data and Code for *Valuing Solar Subsidies*," ICPSR Data Repository 2026. Data and Code Repository for this paper. DOI pending.

Gillingham, Kenneth T. and Asa Watten, "How is rooftop solar capitalized in home prices?," *Regional Science and Urban Economics*, 2024, *107*, 104006. Urban Economics and the Environment.

Google, Inc., "Project Sunroof," `https://sunroof.withgoogle.com` 2018. Accessed: July 1, 2018; Anonymized individual-level household credit records.

Gowrisankaran, Gautam and Marc Rysman, "Dynamics of consumer demand for new durable goods," *Journal of Political Economy*, 2012, *120* (6), 1173–1219.

Hahn, Robert W and Robert D Metcalfe, "Efficiency and equity impacts of energy subsidies," *American Economic Review*, 2021, *111* (5), 1658–1688.

_ **, Nathaniel Hendren, Robert D Metcalfe, and Ben Sprung-Keyser**, "A Welfare Analysis of Policies Impacting Climate Change," Technical Report, National Bureau of Economic Research 2024.

Hausman, Jerry A, "Individual discount rates and the purchase and utilization of energy-using durables," *The Bell Journal of Economics*, 1979, pp. 33–54.

Hornbeck, Richard, Shanon Hsuan-Ming Hsu, Anders Humlum, and Martin Rotemberg, "Gaining Steam: Incumbent Lock-in and Entrant Leapfrogging," Technical Report, National Bureau of Economic Research 2024.

Hotz, V Joseph and Robert A Miller, "Conditional choice probabilities and the estimation of dynamic models," *The Review of Economic Studies*, 1993, *60* (3), 497–529.

Houde, Sébastien and Erica Myers, "Are consumers attentive to local energy costs? Evidence from the appliance market," *Journal of Public Economics*, 2021, *201*, 104480.

Hughes, Jonathan E and Molly Podolefsky, "Getting Green with Solar Subsidies: Evidence from the California Solar Initiative," *Journal of the Association of Environmental and Resource Economists*, 2015, *2* (2), 235–275.

InfoUSA, Inc., "InfoUSA Historical Business and Consumer Database," `https://data-axle.com` 2020. Accessed: January 2020.

Kind, Peter, "Disruptive Challenges: Financial Implications and Strategic Responses to a Changing Retail Electric Business," Technical Report, Edison Electric Institute 2013.

Kirkpatrick, A. Justin and Lori S. Bennear, "Promoting clean energy investment: An empirical analysis of property assessed clean energy," *Journal of Environmental Economics and Management*, 2014, *68* (2), 357 – 375.

Koster, Hans RA and Edward W Pinchbeck, "How do households value the future? Evidence from property taxes," *American Economic Journal: Economic Policy*, 2022, *14* (1), 207–239.

Lamp, Stefan and Mario Samano, "(Mis) allocation of renewable energy sources," *Journal of the Association of Environmental and Resource Economists*, 2023, *10* (1), 195–229.

Lawrance, Emily C, "Poverty and the rate of time preference: evidence from panel data," *Journal of Political Economy*, 1991, *99* (1), 54–77.

Lawrence Berkeley National Laboratory, "Tracking the Sun: Distributed Solar and Storage Restricted Access Data File 2000-2018," Open Energy Data Initiative (OEDI) 2018. Accessed: January 1, 2018; U.S. Department of Energy Contract No. DE-AC02-05CH11231.

Loecker, Jan De, Jan Eeckhout, and Gabriel Unger, "The rise of market power and the macroeconomic implications," *The Quarterly Journal of Economics*, 2020, *135* (2), 561–644.

Magnac, Thierry and David Thesmar, "Identifying dynamic discrete decision processes," *Econometrica*, 2002, *70* (2), 801–816.

Meier, Stephan and Charles Sprenger, "Present-biased preferences and credit card borrowing," *American Economic Journal: Applied Economics*, 2010, *2* (1), 193–210.

Metcalf, Gilbert E and Kevin A Hassett, "Measuring the energy savings from home improvement investments: evidence from monthly billing data," *Review of Economics and Statistics*, 1999, *81* (3), 516–528.

Oliver, Matthew E, Juan Moreno-Cruz, and Kenneth Gillingham, "Microeconomics of the solar rebound under net metering," 2024.

Pless, Jacquelyn and Arthur A. van Benthem, "The Surprising Pass-Through of Solar Subsidies," NBER Working Paper 23260, National Bureau of Economic Research, Cambridge, MA March 2016.

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

Qiu, Yueming, Yi David Wang, and Jianfeng Wang, "Soak up the sun: Impact of solar energy systems on residential home values in Arizona," *Energy Economics*, 2017, *66*, 328–336.

Rennert, Kevin, Frank Errickson, Brian C Prest, Lisa Rennels, Richard G Newell, William Pizer, Cora Kingdon, Jordan Wingenroth, Roger Cooke, Bryan Parthum et al., "Comprehensive evidence implies a higher social cost of CO2," *Nature*, 2022, *610* (7933), 687–692.

Scott, Paul, "Dynamic discrete choice estimation of agricultural land use," 2014.

Sexton, Steven, A Justin Kirkpatrick, Robert I Harris, and Nicholas Z Muller, "Heterogeneous solar capacity benefits, appropriability, and the costs of suboptimal siting," *Journal of the Association of Environmental and Resource Economists*, 2021, *8* (6), 1209–1244.

Sirmans, G Stacy, Lynn MacDonald, David A Macpherson, and Emily Norman Zietz, "The value of housing characteristics: a meta analysis," *The Journal of Real Estate Finance and Economics*, 2006, *33*, 215–240.

Snashall-Woodhams, Nicholas, "Targeting solar subsidies," Technical Report, Working Paper 2024.

US Department of Energy, "Quadrennial Technology Review 2015: An Assessment of Energy Technologies and Research Opportunities," 2015.

van Rijn, Jordan, Shuwei Zeng, and Paul Hellman, "Financial institution objectives and auto loan pricing: Evidence from the survey of consumer finances," *Journal of Consumer Affairs*, 2021, *55* (3), 995–1039.

Warner, John T and Saul Pleeter, "The personal discount rate: Evidence from military downsizing programs," *American Economic Review*, 2001, *91* (1), 33–53.

Online Appendix

A Data Assembly

A.1 Sample Selection

We assemble the study sample by first identifying zip codes in California that are located within the Pacific Gas and Electric service territory. We identify the climate zone associated with each zip code in the Pacific Gas and Electric service territory and extract from the sample homes in those zip codes that lie on a climate zone boundary.[51] The 28 zip codes contained in the sample are shown in Figure 1, with color representing the boundary zone and shading representing the CEC climate zone.

Our base dataset of potential solar adopters is the set of all single-family detached non-mobile home residences that were built before 2014 in the 28 zip codes we identified.

A.2 Google Project Sunroof Matching

We match each household in our data to the nearest Google Project Sunroof (Google, Inc., 2018) record with a greedy spatial matching algorithm. We drop the 5% of households with the largest distance between the geocoded home address and the Google Sunroof record. For each matched household, our data contains the panel-by-panel expected generation. Generation is largely decreasing as panels are ordered by generation, though contiguity requirements may result in some increases. Figure 2b shows the previous four randomly selected rooftops' generation profile.

A.3 Household Consumption Profiles

The PG&E data are anonymized and do not include addresses (**?**). They do include the 5-digit zip code, tariff ID, as well as grid interconnection IDs that identify solar adopters. Thus, for solar adopters, we observe the annual consumption prior to the installation of solar and can designate the closest consumption bin for each adopter. For non-adopters, we see only the full distribution of consumption across the zip code. For each zip code, we remove known adopters and bin all consumption into five equal-sized consumption bins. We also remove a very small number of households (less than 100) enrolled in an early version of the Electric Vehicle (EV) tariff. In estimation, we will integrate over the observed empirical distribution of consumption within each zip code (sampling without replacement).

The study period is 10 quarters, from Q1 2014 through Q2 2016, so we treat consumption as fixed over the study period. We calculate the mean consumption for each bin within each zip code. A density plot of zip-level consumption is shown in Figure A.1a. Figure A.1b shows the consumption levels of adopters and non-adopters by wealth. Across all three wealth bins, adopters have higher consumption on average. As wealth increases from low to medium, the separation between average consumption of non-adopters and average consumption of adopters becomes larger, consistent with lower implicit discount rates and higher sensitivity to the flow of benefits from installing solar.

[51] During the study period, PG&E designated CEC climate zones as "baseline territories". We use PG&E published rates by territory, but note that these territories follow the CEC climate zone boundaries.

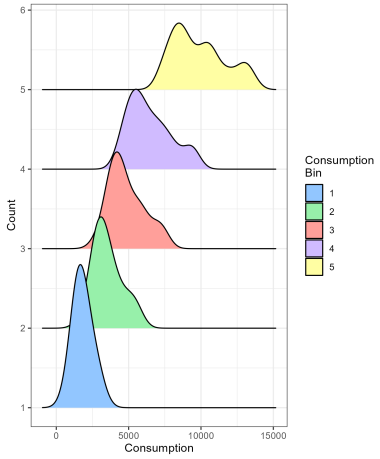

(a) Distribution of zip-level consumption bins calculated from the full distribution of consumption by zip code. For each zip code, consumption is grouped into one of five bins (1 is lowest consumption, 5 is highest) and the mean consumption is calculated for each bin. Plot shows the distribution across zip codes of the five bin's mean consumption.

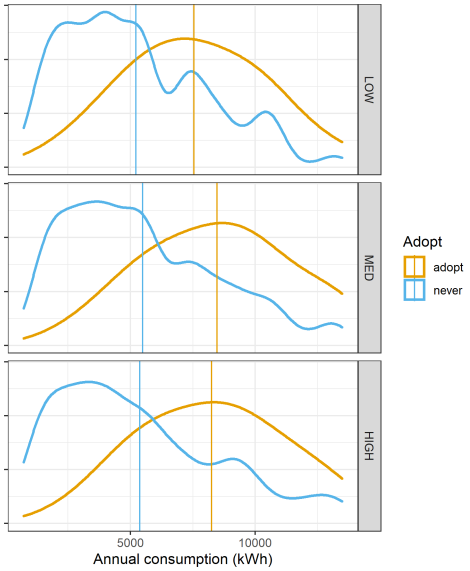

(b) Density of consumption by adopter status and wealth. To calculate density of non-adopters, we assign equal weight to each of the five consumption bins for the household's zip code.

Figure A.1: **Consumption Summary Statistics**

A.4 PG&E Rates and CEC Climate Zones

While all households in the PG&E service territory share the same block pricing steps at any given time (Gillingham, Kenneth, Kirkpatrick, A Justin and Bollinger, Bryan, 2026), California Energy Commission (CEC) climate zones vary in the *width* of a block tier pricing step – hotter inland zones are allowed more baseline consumption before stepping up to the next higher marginal rate relative to cooler coastal zip codes. Our sample includes homes in 28 zip codes that are wholly contained in one of three unique CEC climate zones (PG&E territories S, T, and X). While all homes initially face the same price per kilowatt-hour for their first unit of consumption, homes in warmer CEC climate zones will, as consumption increases, face lower retail rates than will homes in cooler CEC climate zones due to the higher threshold for stepping to the next block tier price. Average retail price per kilowatt-hour is weakly lower in warmer CEC climate zones with higher baseline thresholds.

Climate zones were established by the California Energy Commission (CEC) in 1985, with revisions in 1990 and have largely remained constant through 2019 (Buchsbaum, 2022). They are in general determined by elevation with cuts at 1,000m, 1,500m, 3,000m, and 6,000m, with some exceptions for contiguity of neighborhoods. For billing purposes, PG&E allocated zip codes to each territory. In most cases, when zip codes crossed elevation lines, the zip codes were allocated to the territory of largest coverage with exceptions for very large changes in elevation within a zip code, in which case a zip code was separated into two territories. Our sample consists only of zip codes in which the zip code line coincides with the boundary line.

During the study period, baseline consumption thresholds were adjusted once. Thresholds are shown over time in Figure A.3. The sample contains 34,796 households in zip codes in Territory S, 57,790 households in zip codes in Territory T, and 91,081 households in zip codes in Territory X, for

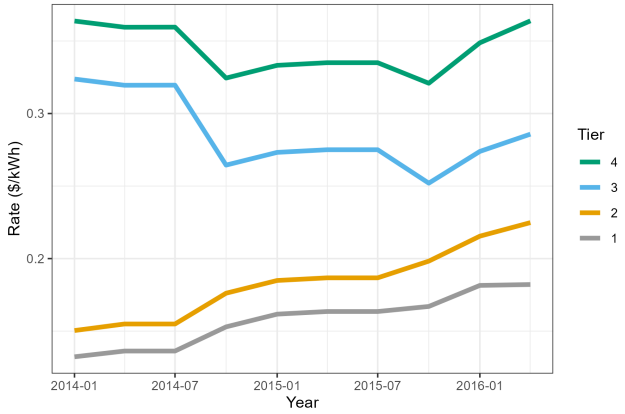

Figure A.2: **Retail Rates**. Retail rates by tier over time in dollars per kilowatt-hour.

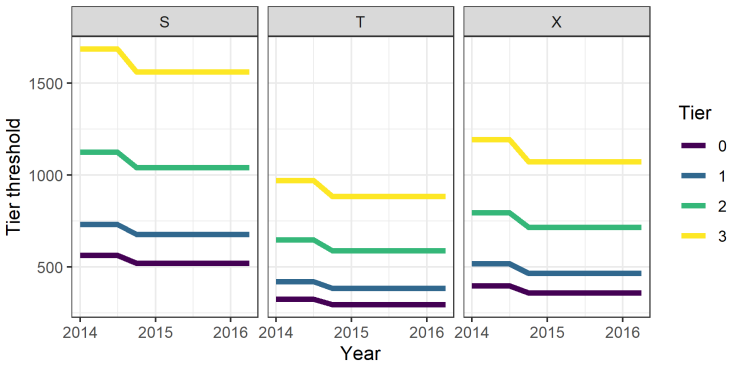

Figure A.3: **Rate tier climate zone thresholds**. CEC climate zone thresholds over time in kilowatt-hours per month. Corresponding rates are shown in Figure A.2

a total sample size of 183,667 representing approximately 3.2% of all owner-occupied single-family detached homes in California.

In addition to spatial CEC climate zone variation in block tier step width, retail electricity rates for each step of block tier pricing vary over time, as shown in Figure A.2. We gather rates and step thresholds from PG&E rate filings to construct a complete panel of electricity rates for our study period (Gillingham, Kenneth, Kirkpatrick, A Justin and Bollinger, Bryan, 2026). Although rates at lower consumption levels (Tiers 1 and 2) steadily increased during the study period shown (2014-2016), higher tier prices decreased slowly over the same time period. Thus, higher-consumption households experienced smaller total increases in the average cost per kilowatt hour during this time period.

B Sizing Model

In this section, we describe the optimal solar system sizing decision. In practice, the system size is calculated by installers based the particular details of the home and household using standard websites, including Google Solar API. The recommended system size is then provided to customers. Customers commonly accept the recommended system size, but have the opportunity to negotiate. Our modeling follows the process installers follow, and we calculate the optimal system size as a function of the system costs (inclusive of upfront subsidies), system generation (as a function of size), and the price of grid electricity, which depends on a household's electricity consumption due to the tiered pricing structure. As a robustness check, we also incorporate the possibility of a solar rebound that increases electricity consumption after the solar installation (Oliver et al., 2024) and vary the assumed panel life T.[52]

Let $TC(K)$ denote the total cost to a household of a solar installation of size K. It is comprised of the fixed cost, FC and a cost per panel, V, charged by the installer. A per-kWh of solar capacity rebate subsidy, S can reduce the upfront cost of the panels. The total cost of the system net of rebates is reduced by a fraction I equal to the investment tax credit.[53] Thus, $TC(K) = (K \cdot (V - S) + FC)(1 - I) = K \cdot C + F$. We define $C = (V - S)(1 - I)$ as the cost per kilowatt of solar capacity faced by consumers (i.e., the post-subsidy variable cost of adding further panels) and $F = FC(1 - I)$ as the fixed cost faced by consumers.

If the levelized cost of electricity generation from the marginal unit capacity is less than or equal to the lowest electricity rate, p_0, we assume that the household that purchases solar will optimally install a system large enough to exactly offset all consumption.[54] The household installs a minimum installation of 5 panels (1.25 kW) if the levelized cost is greater than the highest grid electricity rate, where 5 is the smallest observed installation size in the data. And if levelized costs of the marginal solar capacity are less than the highest grid electricity rate but higher than the lowest grid electricity rate, then the household optimally sizes the solar array to offset just the fraction of grid electricity consumed at the highest rate.

Leased systems in the data are much more likely to be sized equal to current consumption, and indeed the firm incentives are such that installers would prefer recommending the largest system size up to fully offset consumption. Thus, we size purchased systems according to the optimal sizing model, and leased systems according to full offset, which minimizes the absolute error in sizing. Fit statistics are shown in Appendix Table H.2.

Let $Q(K)$ be the quarterly electricity generated by a new solar system of size (installation capacity) K (in units of kWh), and let λ be the quarterly depreciation of panels each quarter. We define Q^* as the annual electricity produced by a new system of optimal size K^*. Let $q(K) = \frac{dQ}{dK}$ be the marginal generation at size K, i.e., the electricity produced by an incremental kilowatt of capacity. Following our descriptive evidence, we assume $q(K)$ is weakly decreasing in K because solar panels are installed first on the best (i.e., highest irradiance) ares of a rooftop. Thus, $q(K) \geq 0$ and $\frac{dq(K)}{dK} \leq 0$. Let the lifetime of the solar panel be given by T quarters.

We define a household's initial grid consumption as Q_0, such that for an installations producing Q kW of electricity with $Q(K) < Q_0$, the remaining grid consumption is $Q_0 - Q(K)$. The interior

[52] For our main results we do not include a solar rebound due to the fact that our model fit is better without including the rebound effect. We find results are robust to shorter panel life assumptions, but the overall model fit is better under $T = 100$ (25 years), the industry standard.

[53] We assume that all homeowners have sufficient tax liabilities to qualify for the 30% tax credit offered on solar during this period.

[54] Excess generation under NEM policies not credited against a month with net withdrawal is compensated at a negligible level after 12 months.

solution to the sizing problem equates the present value of the lifetime stream of offset utility payments that results from an additional unit of capacity to the cost of that capacity. :

$$C(K) = \sum_{t=1}^{T} \left(\frac{(1+\zeta)(1-\lambda)}{1+r} \right)^{t-1} p(Q_0 - Q(K))q(K), \tag{13}$$

in which $p(Q_0 - Q(K))$ is the quarterly cost of the monthly marginal unit of electricity consumption, r is the quarterly discount rate used for sizing, and ζ is the expected increase in real electricity costs.[55]

We define:

$$c(K) = \frac{C(K)}{\left(\sum_{t=1}^{T} \left(\frac{(1+\zeta)(1-\lambda)}{1+r} \right)^{t-1} q(K) \right)}, \tag{14}$$

which is effectively an average, discounted price of solar electricity that consumers compare to the price of grid electricity, $p(Q_0 - Q(K))$, which accounts for the expected increases in both solar and grid prices over time.

Because q is decreasing in K and $p(Q_0 - Q(K))$ is decreasing in K with constant and increasing block rate prices, there will be a unique solution to the first order condition (if any). With increasing block prices, the first units of displaced grid electricity bear the highest marginal costs, and marginal costs of displaced electricity decline discretely as solar capacity increases.[56] The optimal size is a function of upfront incentives and the feed-in-tariff rate, is increasing in q, p, S, and I, and is decreasing in r and V.[57]

Because of the tiered nature of grid prices, the optimal sizing function is not always an interior solution; it is a piece-wise function defined according to the number of tiers in the tariff structure. We illustrate this in the context of a tariff with two distinct tiers of volumetric charges. We abstract from consideration of fixed charges because we assume no households prefer to disconnect from the grid. Let τ be the monthly grid consumption threshold at which rates change from p_0 to p_1 for $p_0 < p_1$, and let q_0 denote monthly household consumption.

[55] We follow de Groote and Verboven (2019) by using a constant and deterministic value for ζ, which we estimate using an AR(1) model with fixed effects. This can also be interpreted as assuming rational expectations, discussed further in Section 5.2.

[56] For decreasing block prices, the opposite is true, and the interior optimum may not be unique.

[57] For the sizing model, we use the Google Sunroof default rate of 4%. We test values around 4% and find that 4% best fits observed installation sizes. This rate is considerably lower than the household rates we estimate. However, the optimal sizing decision is generally not made by consumers alone, but rather by the installation contractor who presents the cost and expected payoff when generating a quote for homeowner consideration. Using a discount rate of 10% in sizing would result in more than half of all installations receiving an optimal size of 5 panels, the smallest size reported by Google Sunroof, a figure that is grossly incongrous with observed data. Optimal sizing also must assume an expected rate of increase in electricity prices, which we estimate from the data, as described in Appendix D and apply to the sizing decision. We find that an alternative sizing that uses an average 2.2% increase in electricity prices suggested by Google Sunroof rather than our estimated rates does not improve model fit. Furthermore, we include a robustness check using a 2.2% increase in electricity prices for sizing and adoption, and find no significant difference in results (see I.3).

The marginal price of grid electricity p depends upon the residual grid demand, i.e., consumption net of solar generation, such that:

$$p = \begin{cases} 0 & Q_0 - Q \leq 0 \\ p_0 & 0 < Q_0 - Q < \tau \\ NSOp_1 & Q_0 - Q \geq \tau \end{cases}.$$

We define the minimum size installation considered as $\underline{K} = 1.25$ kW. The piece-wise-defined solution to the optimal sizing problem for increasing block rates and a household consuming at the highest rate is defined as:

$$K^* = \begin{cases} \underline{K} & c(0) > p_1 \\ q^{-1}\left(\frac{C}{\sum_{t=1}^{T}\left(\frac{(1+\zeta)(1-\lambda)}{1+r}\right)^t p_1}\right) & c(0) \leq p_1, c(Q_0 - \tau) \geq p_1 \\ Q^{-1}(Q_0 - \tau) & p_0 \leq c(Q_0 - \tau) < p_1 \\ q^{-1}\left(\frac{C}{\sum_{t=1}^{T}\left(\frac{(1+\zeta)(1-\lambda)}{1+r}\right)^t p_0}\right) & c(Q - \tau) < p_0, c(Q_0) \geq p_0 \\ Q^{-1}(Q_0) & c(Q_0) \leq p_0 \end{cases}.$$

Figure B.1 depicts these cases, moving in order from the top curve to the bottom, for given alternative q's that define c. The weakly decreasing step function (indicated in red) is the levelized cost of the marginal unit of grid electricity as a function of grid consumption, and the intersection points with each curve indicates the generation from an optimally size installation. Solar generation $q(K)$ is increasing left to right, and remaining grid consumption is increasing right to left.

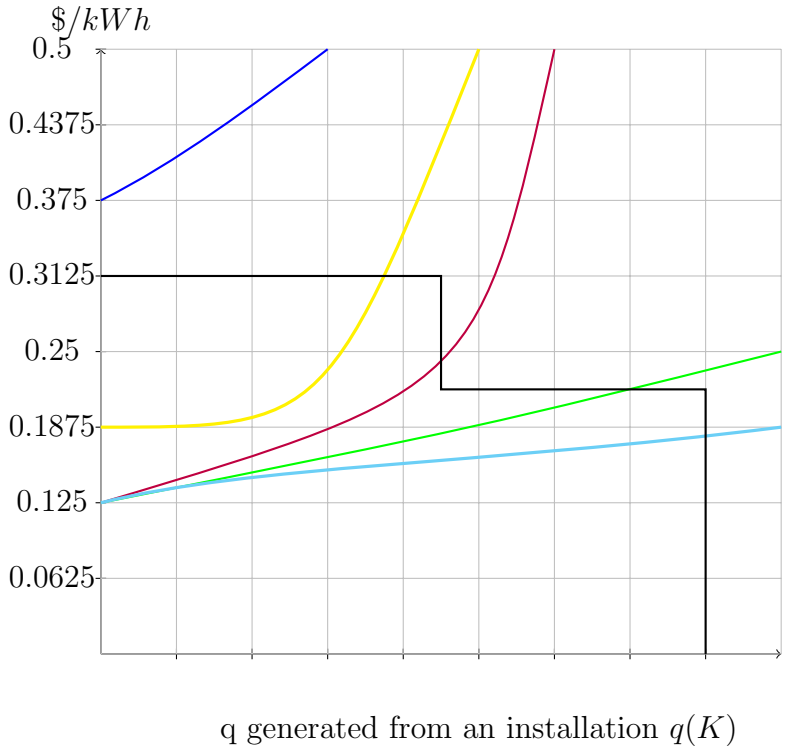

Figure B.1: **Optimal sizing**. Optimal sizing of solar PV arrays is depicted as a function of alternative modeled q' functions, i.e., the electricity generation of marginal units of solar capacity. Depicted are the costs per kWh for alternative q's and a tariff with two tiers of volumetric charges. Grid consumption increases from right to left. Solar capacity increases left to right.

Figure B.2 shows the fit of the optimal sizing model fit over purchased (left) and leased (right) systems, which use the two different assumptions. For purchasers, we use the optimal sizing model, and for lessees, we assume full offset. For non-adopters, we observe only the zip-level distribution of consumption, and thus size each household conditional on consumption at one of five zip-level consumption bins. In Figure B.3, we graphically present the sizing model fit using average annual consumption (left) and the same sizing model fit using the *closest* consumption bin level for the adopting household's zip code. The horizontal striation on the right is due to discretization to the nearest consumption bin. The right panel, despite not using actual consumption, provides a close fit to observed installation sizes, despite the discretization.

Table B shows a low average error in sizing expressed in number of panels. The average error of 1-2 panels is 5-10% of the average installation size (20 panels). This is consistent with installers using similar tools (electricity bills and roof profiles) and assumptions to size a system.

Our preferred sizing model does not account for rebound effects. As a robustness check, we allow for a solar rebound effect, which may either increase the amount for grid consumption, or it may lead to larger panel sizes. For consumers not sizing their array to offset full consumption, any rebound will be through increased grid consumption, if the marginal cost of electricity is reduced due to the solar installation by placing the household on a lower rate tier. For consumers sizing their array to offset full consumption, the levelized cost of added solar generation will be weakly less than the cost of grid electricity, and there will be some scope to increase the size of the installation due to rebound. To do so, we calculate the consumer surplus using a constant price elasticity of

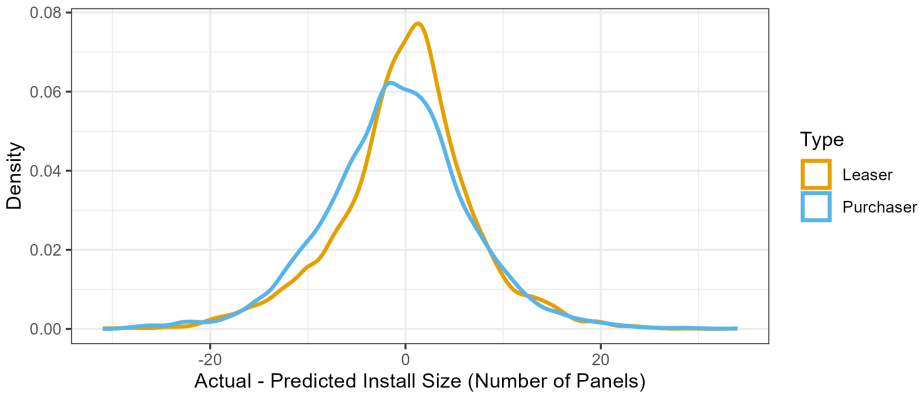

Figure B.2: **Sizing model accuracy**. Density plot of Google Sunroof-method actual-predicted installed system capacity for (orange) leased systems and (blue) purchased systems.

Table B.1: Optimal sizing model fit by type. This table reports the difference between the predicted and actual observed installations. Positive numbers indicate sizing greater than observed. Does not include reported installations with greater than 44 panels (95th percentile).

	Mean Sizing Error		Mean Absolute Sizing Error	
Type	count panels	%	count panels	%
Purchaser	2.01	10%	5.72	29%
Leaser	1.19	6%	4.97	24%
All	1.64	8%	5.38	27%

$\varepsilon = 0.1$ and allow the system size to increase so long as the cost of the increase is lower than the increased surplus. Details are in Appendix C.[58]

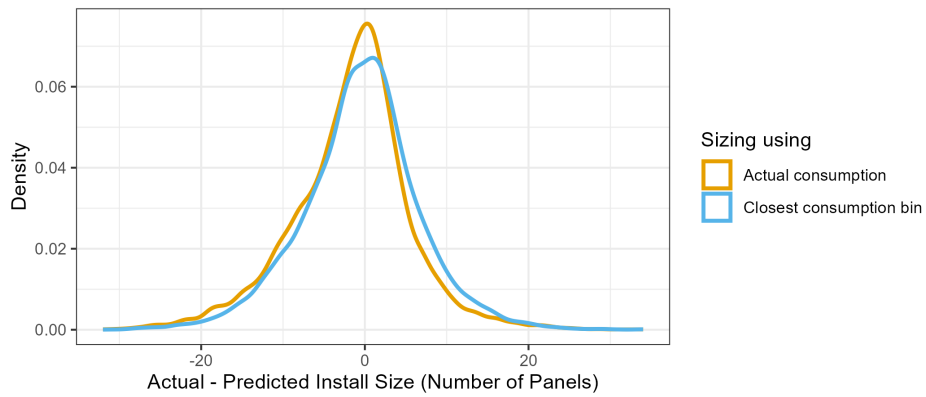

Figure B.3: **Sizing model accuracy**. Density plot of Google Sunroof-method actual-predicted installed system capacity sized using (orange) actual consumption and (blue) using assigned closest consumption bin.

C Sizing with Rebound

We account for rebound effects using a constant demand elasticity for electricity, ε. Let:

$$\log(q) = \varepsilon \log(p) + \log\left(\frac{q_0}{(p_H)^{\varepsilon}}\right), \tag{15}$$

in which the second term is an intercept that sets observed pre-installation consumption to q_0 at the observed pre-installation marginal price of electricity, p^H. We can expect demand to increase by Δq, which is the maximum value of r such that:

$$r \leq -\varepsilon \left(\frac{p_H - p_L(r)}{p_H}\right) q_0, \tag{16}$$

in which p_L is the effective marginal electricity rate with rebound r using the discount rate used to size the installation:

$$p_L(r) = \min_{r^S}\left(\min\left(p(q_0 - q^* + r - r^S), \frac{C/T}{q'\left(K(q^* + r^S)\right)\left(\sum_{t=1}^{T} \frac{1}{(1+\delta)^t}\right)}\right)\right), \tag{17}$$

in which $r^S \leq r$ is the amount of the rebound electricity generated from increasing the installation size and q^* is still defined as the optimal solar generation ignoring rebound. Let Δq^S be the optimal amount of this rebound solar electricity. For households not offsetting full consumption with their installation, the entirety of the rebound effect is through higher grid consumption at

[58]When including rebound with a price elasticity of $\varepsilon = 0.1$, the mean error for purchasers increases to 16%, while the mean absolute error increases to 32%. Thus, we use the no-rebound sizing model for our main estimation due to the high degree of fit and little evidence of the systematic under-sizing that would indicate sizing for rebound. We report results using $\varepsilon = 0.1$ in Appendix I.1.

the new marginal rate, $p(q_0 - q^* + \Delta q)$, and equation (16) holds with equality except at a point in which the rebound pushes the household to the next price tier. For those households offsetting full consumption prior to rebound, the cheapest increment of additional electricity may be (and will almost always start with) the levelized marginal rate of additional solar electricity, but as the installation size increases and the levelized cost or marginal generation increases, it may switch to additional grid consumption.

For households whose rebound consumption is grid electricity (those who don't offset their full consumption), the value of the generated electricity in each period is equal to the offset cost of q^* plus the added surplus from the rebound effect:

$$\begin{aligned}
\Delta S^R &= \int_{p_L}^{p_H} q(p)dp - (p_H - p_L)q_0 + \int_{q_0}^{q_0+\Delta q-\Delta q^S} (p_L - p(q-q^*))dq \\
&= \int_{p_L}^{p_H} p^{\varepsilon}\frac{q_0}{(p_H)^{\varepsilon}}dp - (p_H - p_L)q_0 + \int_{q_0}^{q_0+\Delta q-\Delta q^S} (p_L - p(q-q^*))dq \\
&= \left(\left(\frac{1}{1+\varepsilon}\right)\left(\frac{(p_H)^{1+\varepsilon} - (p_L)^{1+\varepsilon}}{(p_H)^{\varepsilon}}\right) - (p_H - p_L)\right) q_0 \\
&+ \int_{q_0}^{q_0+\Delta q-\Delta q^S} (p_L - p(q-q^*))dq. \qquad (18)
\end{aligned}$$

The first two terms are the additional surplus of consumption when the cost of this electricity is p_L. The third term is equal to zero unless the rebound consumption is at multiple price levels (i.e. at multiple price tiers), in which case we need to account for the fact that some of the grid rebound consumption is at a price lower than p_L.

Households who currently offset full consumption may consume from the grid, and they may not. They will also increase the size of their installation, in which case we need to adjust this surplus term by:

$$p_L(q_0 + \Delta q^S) - \left(TC(K(q^* + \Delta q^S)) - TC(K(q^*))\right). \qquad (19)$$

This term is the levelized value of the solar rebound minus the cost. It is equal to zero if the discount rate used in sizing is the same as that when making installation decisions. Since most sizing decisions are heavily influenced by the installer (using a discount rate of 4%), we allow for these to be different.

As a robustness check, we re-size the optimal sizing model using an elasticity of $\varepsilon = 0.1$, motivated by the literature. We apply the rebound only to purchased systems under the assumption that leased systems are sized to full consumption, as in the main body.

Relative to table B, sizing with rebound over-sizes purchased systems, leading to greater mean sizing error (16%) and greater mean absolute error (32%). Leased systems are sized to offset full consumption, and thus have the same error as in table B.

Nevertheless, we re-estimate our preferred specification and calculate the counterfactual-implied elasticity as in 7. Results are in Appendix I

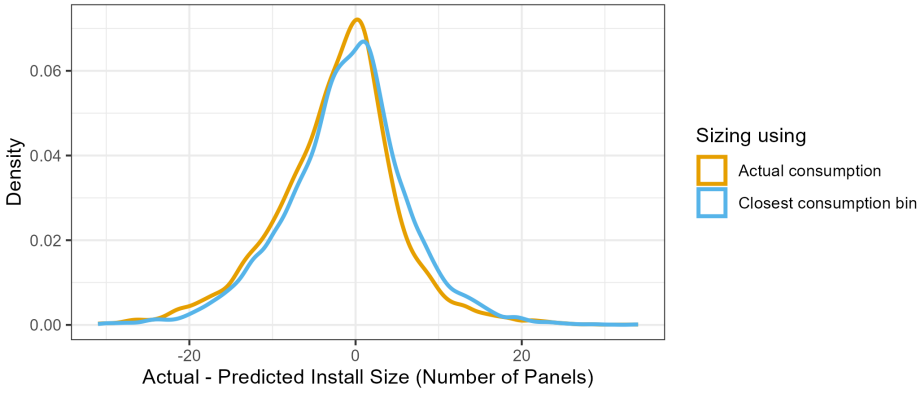

Figure C.1: **Sizing model with rebound accuracy**. Density plot of Google Sunroof-method actual-predicted installed system capacity sized using (orange) actual consumption and (blue) using assigned closest consumption bin, both inclusive of rebound.

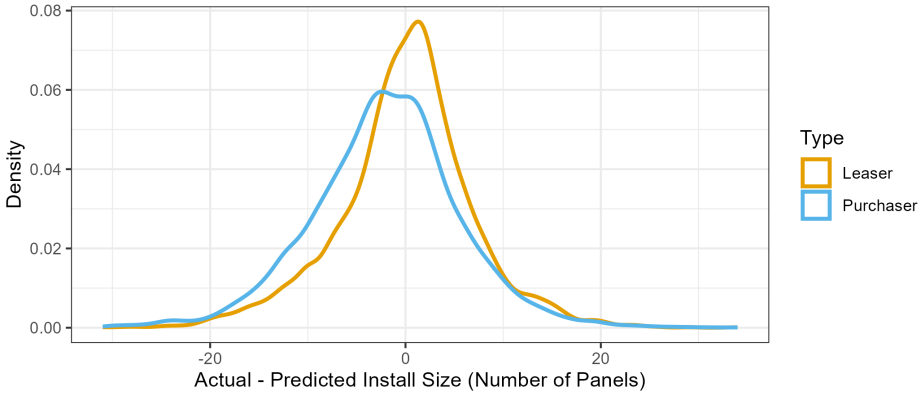

Figure C.2: **Sizing model with rebound accuracy**. Density plot of Google Sunroof-method actual-predicted installed system capacity for (orange) leased systems and (blue) purchased systems, both inclusive of rebound.

Table C.1: Optimal sizing model fit by type. This table reports the difference between the predicted and actual observed installations. Positive numbers indicate sizing greater than observed. Does not include reported installations with greater than 44 panels (95th percentile).

	Mean Sizing Error		Mean Absolute Sizing Error	
Type	count panels	%	count panels	%
Purchaser	3.13	16%	6.21	32%
Leaser	1.19	6%	4.97	24%
All	2.25	11%	5.65	28%

D Estimation of Electricity and Solar Price Trends

D.1 Cost of Solar Installation

We need an estimate of the cost of adopting solar for all potential adoptions, which may be a function of installation size, time, and location. To do so, we estimate the fixed and variable (per-watt) components of price using TTS data for each boundary zone and quarter in our sample. We assume that the fixed price component will vary by boundary zone, and that a common variable price per watt will decline over time at a constant rate η, to be estimated. The decline of panel prices on a per-watt basis is well-documented in the literature. Heterogeneous fixed costs reflect relative wealth and cost of labor across the study areas.

We estimate the solar price model on a sample that includes our study period, as well as the year prior (2013). We deflate costs using the Bureau of Labor Statistics quarterly CPI using the fourth quarter of 2016 as the base period. Our decomposition of total reported price into fixed costs, variable costs, and a common η uses the following specification, which we estimate by nonlinear least squares (NLS):

$$SystemCost_{it} = \sum_{b} \kappa_b \mathbb{1}_{b=b(i)} + \beta_W SystemW_{it} \cdot e^{\beta_T (t-t_0)} + \varepsilon_{it}, \tag{20}$$

in which *SystemCost* is the cost exclusive of subsidies or tax credits, *SystemW* is the size of the system, $t - t_0$ is the elapsed time since the beginning of the cost model estimation sample, Q1 of 2013.

Estimated fixed costs range from \$998 to \$1,622. Per-watt panel estimates start at \$5.12 in the first quarter of 2013, and decline at a rate of $e^{-.014t}$. A value of $\beta_T = -.01403$ yields an estimated quarterly $\eta = 0.986$. We use the estimates in table D.1 to predict fixed and variable installation costs for adopters and non-adopters.

Table D.1: **Cost Model Estimation**: Results from estimation of equation 20 using the sample of all installations in the study area during the study period and the previous year. Results show boundary zone-specific fixed cost intercepts κ, per-watt costs β_W in first quarter 2013, and quarterly rate of decline in per-watt costs, β_T, where $\eta = 1 - \beta_T = .986$. All costs are deflated using BLS CPI-U.

	(1)
κ_A	1622.238
	(154.807)
κ_B	1527.982
	(244.818)
κ_C	1124.275
	(185.183)
κ_D	1119.103
	(164.510)
κ_E	998.059
	(283.709)
κ_F	1212.440
	(148.897)
β_W	5.122
	(0.034)
β_T	-.014
	(0.000)
Num.Obs.	9163
Log.Lik.	-89988.246

D.1.1 Solar Costs By Wealth

Installation costs are assumed to be constant across wealth and boundary zones, implying a value of η (the quarterly expected decline in per-W solar panel costs) that is the same for each wealth group. However, if low- or high-wealth households have systematically higher (or lower) costs, or follow a different decline η, then our estimates may be biased. For instance, if low-wealth solar-interested households receive fewer quotes from installers and thus face weakly higher prices, then we may understate the up-front cost of solar for low-wealth households and incorrectly attribute non-adoption to a higher discount rate. While our main specification includes fixed effects for boundary zone by wealth, which would absorb the average effect, as well as any effect present in the fixed costs, we further check for robustness of estimates for β_T and β_W.

We re-estimate our cost model in 20 but allow β_W and β_T to vary by wealth bin. We estimate the model on the entire sample, as well as each boudnary zone individually. Table D.2 shows the results. No significant difference is present for either β_W, the per-W price of a solar installation, though β_W is significant at the 10% level in the pooled sample. The rate of decline of the per-W price is significantly different for medium-wealth households in Zone A, but is not significant in the pooled sample.

Table D.2: Solar installation cost model estimated including interactions between price per W and rate of decline with household wealth.

	All	Zone A	Zone B	Zone C	Zone D
β_T	−0.014***	−0.015***	−0.013***	−0.013***	−0.014***
	(0.001)	(0.001)	(0.001)	(0.001)	(0.001)
β_T^{Low}	−0.001	−0.002	0.010	−0.003	0.001
	(0.002)	(0.003)	(0.007)	(0.003)	(0.003)
β_T^{Med}	−0.001	−0.005*	−0.004	−0.001	0.001
	(0.001)	(0.002)	(0.004)	(0.002)	(0.002)
κ_{zone}	1239.536***	1866.265***	310.137	1352.965***	650.883**
	(139.263)	(250.792)	(398.022)	(291.438)	(244.629)
β_W	5.118***	5.102***	5.212***	5.087***	5.191***
	(0.043)	(0.082)	(0.102)	(0.098)	(0.072)
β_W^{Low}	0.182+	0.190	−0.443	0.210	0.211
	(0.095)	(0.168)	(0.408)	(0.171)	(0.183)
β_W^{Med}	0.004	0.140	0.262	−0.051	0.028
	(0.058)	(0.134)	(0.259)	(0.106)	(0.116)
Num.Obs.	8381	2989	1015	2256	2121
Log.Lik.	−82 255.065	−29 412.288	−9954.212	−22 240.219	−20 592.283

+ p <0.1, * p <0.05, ** p <0.01, *** p <0.001

D.2 Electricity Prices

We use an AR(1) model with fixed effects to estimate the expectation of the rate of increase in average electricity rates for each consumption bin and in each zipcode's CEC climate zone, $\zeta_{b,z}$, using deflated quarterly data. To account for CPUC orders requiring that the block-tiered price system reduce the gap between high- and low-consumers that resulted in a significant one-time decreases in tiers 3 and 4 and increases in tiers 1 and 2 which took effect in Q3 of 2014, we include a break term in the AR(1) specification as follows:

$$avgprice_{t,bz} = \zeta_{bz} avgprice_{t-1,bz} + \zeta^{b}_{bz} avgprice_{t-1,bz} \mathbb{1}(t = 2014Q3) + \varepsilon_{t,bz} \tag{21}$$

Table D.3 shows results from equation 21. ζ_{bz} estimates range from 2.67 to 5.62 percent annually (.66 to 1.37 percent quarterly).

Table D.3: **Electricity Rate Model Estimation**: To estimate equation 21, we calculate the average electricity rate for each quarter from Q1 2012 to Q3 2016, which captures the two years prior to the study period. We calculate this average rate using annual consumption that reflects the distribution of consumption in each zipcode and deflate using the BLS CPI-U. For each level of consumption and CEC climate zone, designated by $\{S, T, X\}$, we calculate the total variable cost portion of the bill and divide by the annual consumption. The unit of observation is defined by the combinations of zipcode, consumption level, and CEC climate zone. Consumption bin 5 is the highest quintile of consumption.

	(1)
$\zeta_{1,S}$	1.012
	(0.003)
$\zeta_{2,S}$	1.012
	(0.003)
$\zeta_{3,S}$	1.013
	(0.003)
$\zeta_{4,S}$	1.014
	(0.003)
$\zeta_{5,S}$	1.007
	(0.003)
$\zeta_{1,T}$	1.012
	(0.002)
$\zeta_{2,T}$	1.012
	(0.002)
$\zeta_{3,T}$	1.014
	(0.002)
$\zeta_{4,T}$	1.011
	(0.002)
$\zeta_{5,T}$	1.007
	(0.002)
$\zeta_{1,X}$	1.012
	(0.002)
$\zeta_{2,X}$	1.013
	(0.002)
$\zeta_{3,X}$	1.012
	(0.002)
$\zeta_{4,X}$	1.010
	(0.002)
$\zeta_{5,X}$	1.008
	(0.002)
Num.Obs.	3420
R2	0.999
R2 Adj.	0.999
Std.Errors	HC Robust
Time step	Quarterly

E Number of Panels for a Given Level of Production

Figure E.1 shows the number of panels required (for observed installations) to produce 20 different levels of solar output, q^*. This provides us with variation in the upfront costs, as installers recommend the installation size that makes sense for the household given its irradiance, roof profile, and electricity costs.

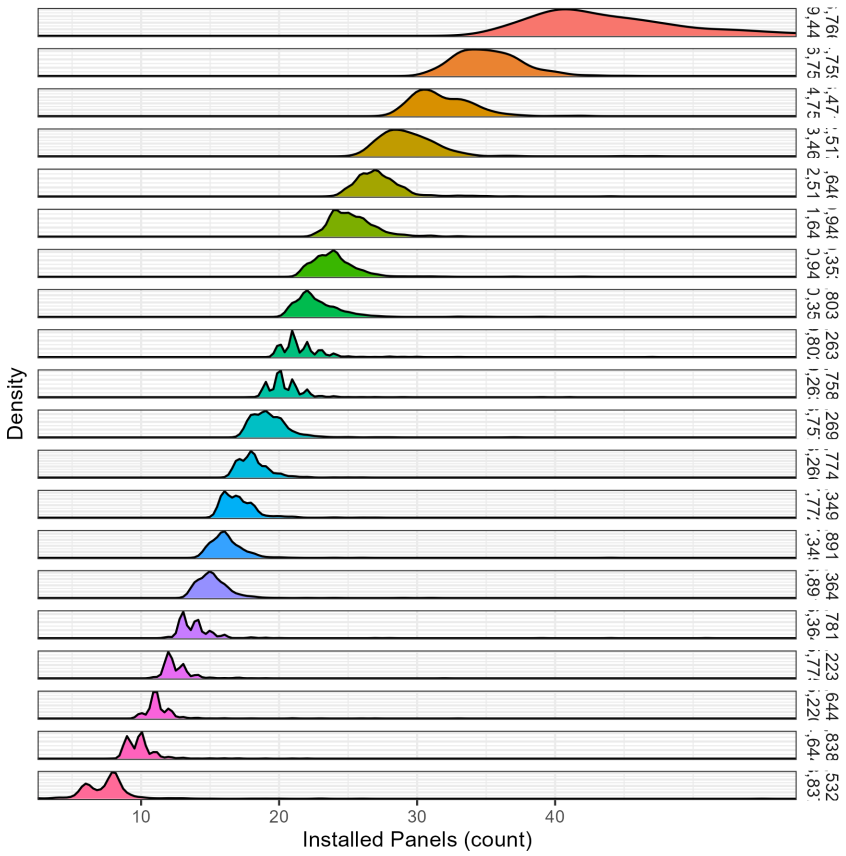

Figure E.1: The distribution of optimal K^* (in panel counts) conditional on the quantity of per-period offset electricity q^*. Data for K^* is from observed installations while the corresponding annual generation q^* is from Google Project Sunroof. K^* enters the adoption decision as a variable up-front cost. Therefore, this figure represents the variation in up-front cost associated with very narrow bands of per-period generation. This variation is solely from roof orientation and irradiance.

F Leasing Model

Nearly 45% of adopters during our study period used a third-party "leased" system wherein the lessor bears all up-front cost to install the rooftop system, and the lessee agrees to pay a price per unit of consumption during the life of the lease. In terms of up-front costs and benefits over time, the adoption decision for a lessee is distinctly different from that of a purchaser. Specifically, leasing has the effect of amortizing the cost of installing over the life of the panels. We are the first in the literature modeling solar demand to explicitly model the lease option.

For leased systems, the economic value of the lease over its lifespan is:

$$\delta_1^l = \omega\theta q^* \bar{p} - \theta^{ppa} q^* p^{ppa}, \tag{22}$$

in which we define:

$$\theta^{ppa} \quad = \quad \sum_{\tau=1}^{T} \left((1+\zeta^{ppa})(1-\lambda)\rho\right)^{(\tau-1)}, \tag{23}$$

and in which p^{ppa} is the starting price of solar electricity and ζ^{ppa} captures the annual change in the electricity price guaranteed by the PPA, which we set at 0.04 based on industry standards (we show the results from a robustness check using 0.03 in Appendix I.2).

We assume that the price the installer charges reflects the levelized cost of installing the system:

$$p^{ppa} = \frac{C(K^*) + F}{\theta^I q^*} + markup, \tag{24}$$

in which the θ^I captures the discounted sum of benefits per kWh using the installer's implicit discount rate, set at 6%: $\theta^I = \sum_{t=1}^{T} \left((1+\zeta^{ppa})(1-\lambda)\rho^I\right)^{(t-1)}$. We also allow for a transfer of *markup* to the installer for fronting the installation costs. We parameterize this transfer by letting κ^{TC} denote the profit margin based on the cost of providing the capital for the installation, and κ^p a per-kWh markup. We can write (24) as:

$$p^{ppa} = (C(K^*) + F)(1+\kappa^{TC})\frac{1}{\theta^I q^*} + \kappa^p. \tag{25}$$

The first multiplicative factor in (25) applies the markup to the installation cost and then divides by θ^I, which amortizes the marked-up cost over 25 years at the installer's implicit discount rate. Finally, the amortized cost is divided by the per-period generation q^*, plus a per-kWh markup.

By substitution, for leased systems, the economic value of the lease over its lifespan can thus be written:

$$\delta_1^l = \theta q^* \bar{p} - \theta^{ppa} q^* \kappa^p - \frac{\theta^{ppa}}{\theta^I}(1+\kappa^{TC})(C(K^*) + F). \tag{26}$$

We assume all of the markup is through the κ^{TC}, which captures the full resolution of p^{PPA} and set $\kappa^{TC} = 0.7$ based on the observed markup reported by a prominent solar leasing company's annual reports during our study period;[59] correspondingly, we assume κ^P is equal to zero. These assumptions together yield prices p^{ppa} that are consistent with observed lease terms, approximately \$0.18-0.19 per kWh (we have tested and found results that are robust to other combinations of κ^{TC} and κ^{PPA} that together yield similar total markups).

We can write the value of leasing then as:

$$v_1^l = \omega\left[p^{EL} - \kappa^l p^{INV}\right] + X\beta^l. \tag{27}$$

[59] See `https://ir.tesla.com/press-release/solarcity-announces-first-quarter-2014-financial-results`, which states a markup of 50-55%. We assume markup is pre-ITC, and add 15% for the ITC to total 70%.

with $\kappa^l \equiv \frac{\theta^{ppa}}{\theta^l}(1+\kappa^{TC})$ capturing the relative tradeoff between upfront costs and long-term benefits of solar when leasing vs. purchasing. Households with $\kappa^l < 1$ get more economic benefit from leasing and households with $\kappa^l > 1$ get more economic benefit from purchasing.

G Estimation Details

G.1 Likelihood

Each non-adopter household is one of ten possible consumption $\times$ type combinations $\{1, ..., 5\} \times$ {lessee, purchaser}. We specify weights w_{ibe} as the probability that household i consumes in consumption bin b and is of type e. With weights w_{ibe}, we integrate the likelihood function over the unobserved heterogeneity. Restating Equation 12:

$$\mathbb{L} = \prod_i \sum_b w_{ib} \sum_e w_{ie|b} \prod_t [Pr^e_{ibt}]^{y_{it}} [1 - Pr^e_{ibt}]^{(1-y_{it})},$$

which requires evaluating ten conditional likelihoods per observation of household i and time t.

We posit the weights w_{ibe} as $w_{ib} \times w_{ie}$, where the following adding up constraints apply:

$$\begin{aligned} &\textstyle\sum_{e'=1}^{2} w_{ie'} = 1 \\ &\textstyle\sum_{b'=1}^{5} w_{ib'} = 1 \\ &\textstyle\sum_{i \in z} w_{ib} = \frac{N_z}{5} \end{aligned} \tag{28}$$

The first two constraints require that each household i have weights that sum across consumption bin b and type e to equal 1. The constraint in equation (28) requires that, within a zipcode z, the weights must sum to 1/5th of the number of households in the sample for that zipcode (due to the use of quintiles). This guarantees that the moments of the consumption distribution match the empirical distribution. Type e weights have no analogous restriction, only that they sum to unity.

The constraint in equation (28) makes w_{ibe} dependent on w_{jbe} for any $i \in z, j \in z$. Methods of calculating these weights such as those in Arcidiacono and Miller (2011), though applicable for type-weights in our context, are not appropriate for dependent weights, and an alternative is employed here. To account for dependence, we integrate over randomly drawn allocations of b that comport with the empirical distribution, which we describe below.

G.2 Estimation Procedure

Here, we describe the estimation of parameters $\theta = \{\rho, \omega, \beta\}$, weights $w_{ie|b}$ and w_{ib}, and conditional choice probabilities (CCP), the next-period probability of adopting. Following Arcidiacono and Miller (2011) and Arcidiacono and Ellickson (2011), we also estimate $\pi(e)$, the probability of being of type e.[60] We implement the expectation maximization (EM) algorithm which proceeds in two stages, and update CCP's in a third stage. First, the expectation stage in which weights $w_{ie|b}$ and w_{ib} and $\pi(e)$ are calculated from the data conditional on an estimate of θ, and the maximization

[60] Note that no analogous π is estimated for consumption bin as, by construction, the probability of being in bin b is 0.2.

step in which θ is estimated conditional on taking weights and $\pi(e)$ as fixed. Finally, CCP's are updated using updated type and bin weights.

G.2.1 Step 1: Initialization

1. Estimate CCP's using a weighted flexible logit described in the following section. Doing so requires weights w_{ibe} over type e and consumption b, which are initialized to .1 for non-adopters and 1 for adopters for their observed type and consumption and zero otherwise. Then, use the flexible logit to predict the next-period probability of adopting, $Pr_{ibe}^{e'}$ (see below).

2. Initialize parameters to be estimated in the maximization step $\hat{\theta}^{(0)}$, $\hat{\pi}^{(0)}(e) = 0.475$ for $e =$ lessee (implying $\hat{\pi}(purchaser) = 1 - \hat{\pi}(leaser)$). This initialization reflects the proportion of observed adoptions that are leases in the data. Set $m = 0$ for notation.

3. Draw $R = 1000$ random allocations of b that place $\frac{1}{5}$ of the non-adopting households into each bin, guaranteeing that $b^{(r)}$, the rth allocation, satisfies the constraint in 28 for each allocation r.

4. Evaluate the conditional likelihood L_{ibe} for every $\{b, e\}$ combination for every household i using $\hat{\theta}^{(0)}$ and equation 11.

5. Calculate $w_{ie|b}$, the type-weight, as $w_{ie|b}^{(0)} = \frac{\hat{\pi}^{(0)}(e)L_{ibe}}{\sum_{e'} \hat{\pi}^{(0)}(e')L_{ibe'}}$.

G.2.2 Step 2: Maximization

Take type-weights and CCP's as fixed within the maximization step. For each evaluation of θ, the likelihood is calculated in the following steps:

1. Calculate $L_{ib} = \sum_{e'} w_{ie|b} L_{ibe}$ using type-weights from step 1.

2. Turning to the b weights, for each r, calculate the likelihood of observing r conditional on the parameters and $w_{ie|b}$. To do so, take the product of $L_{ib^{(r)}}$ where $L_{ib^{(r)}}$ is the likelihood for household i conditional on being in the bin drawn in allocation r. Formally, $L_r = \Pi_i L_{ib^{(r)}}$.

3. Calculate allocation weights $w_z^{(r)} = \frac{L_r}{\sum_{r'} L_{r'}}$.

   (a) Note that allocations r where households are allocated to consumption bins b that better explain the observed outcome are weighted higher than allocations that poorly explain the observed outcome. All weights sum to 1 within household i, and all weights satisfy Equation (28) by definition.

4. Calculate $L_z = \sum_r w_z^{(r)} L_r$.

5. Maximize the log-likelihood $L = \sum_z log(L_z)$.

Upon convergence, the maximization step yields updated $\hat{\theta}^{(m+1)}$.

G.2.3 Step 3: Expectation

Weights $w_{ie|b}^{(m)}$ and $\hat{\pi}^{(m)}(e)$ are then updated following Arcidiacono and Ellickson (2011) and using $L_{ibe}(\hat{\theta}^{(m+1)})$.

1. Calculate $w_{ie|b}^{(m+1)} = \frac{\hat{\pi}^{(m)}(e)L_{ibe}}{\sum_{e'} \hat{\pi}^{(m)}(e')L_{ibe'}}$
2. Calculate $w_{ib}^{(m+1)} = \sum_{r'} 1(b_i^{(r')} = b)w_z^{(r')}$
3. Calculate $w_{ibe}^{(m+1)} = w_{ie|b}^{(m+1)} w_{ib}^{(m+1)}$
4. Calculate $\hat{\pi}(e)^{(m+1)} = \frac{\sum_i 1(e'=e)w_{ibe'}^{(m+1)}}{\sum_i \sum_e w_{ibe}^{(m+1)}}$, the mean type-weight across non-adopting households, weighted by the household's bin weight.

G.2.4 Iterate Steps 2-3

Re-estimate the weighted flexible logit using updated $w_{ibe}^{(m+1)}$ and update CCP's. Repeat steps 2-3 until convergence both in the maximization step and the expectation step, as in Arcidiacono and Ellickson (2011). Convergence overall requires convergence in (i) $\hat{\theta}$, (ii) $\hat{\pi}(e)$, and (iii) CCPs.

G.3 Conditional Choice Probabilities

In estimation we utilize the state transition estimates shown in Appendix D. The final component of equation (12) is the conditional choice probability (CCP), which is the predicted probabilities of adopting in the next period. We again follow Arcidiacono and Miller (2011) and use a flexible logit to estimate the probability of adopting for household h at time t.[61] Weights w_{ibe} are used in the flexible logit and a separate probability of adopting is estimated for leasers and purchasers. In estimation, the next-period probability of adopting $Pr_{ibe}^{e'}$ is updated in a third step using updates of w_{ibe} to re-weight the logit estimation and following the update of the weights.

The flexible logit is specified using all interactions of square footage, length of residence, voter affiliation, presence of children, race, quarterly savings $q^*\bar{p}$, net system cost $C(K^*)$, consumption bin b, and a second degree polynomial in time. We also include a 3rd degree polynomial in time and boundary zone fixed effects. The CCP's (predicted probabilities of adopting in the next period) are generated by advancing the time by 1 period (including advancing $\bar{p}$ by ζ and $C(K^*)$ by η) and predicting the logit response. As a robustness check, we also generate CCP's using the next-period state of $\{\bar{p}, C(K^*)\}$ (see Appendix I.8).

G.4 Estimation Transformations

The parameters of ω_i are estimated with an exponential transformation

$$\omega_i = e^{\omega_0 + \omega_{med} 1(\text{wealth}_i = \text{med}) + \omega_{low} 1(\text{wealth}_i = \text{low})}.$$

[61] A bin estimator would be feasible, but components of Equation 9 and the leaser analogue are continuous

The parameters of ρ_i, the household discount factor, are estimated with a Normal cdf transformation:

$$\rho_i = \Phi(\rho^0 + \alpha^{med} 1\{\text{wealth}_i\text{=med}\} + \alpha^{low} 1\{\text{wealth}_i\text{=low}\}).$$

Similar functional forms are used when credit data is incorporated into ρ_i (see Appendix J.1). All dollar amounts are scaled to be in thousands of dollars.

H Structural Parameter Results

The full set of results are shown in Table H.1. In addition to the results discussed in the main text, we find that length of residence and the number of children are negative but insignificant. While previous literature has indicated these demographics correlate with adoption, these other papers did not condiiton on the economic value of adopting solar.[62] We condition on the house being listed as owner-occupied in the Corelogic data, but not all of these are coded as owner-occupied in InfoUSA; the indicator for being listed as non-owner-occupied in the InfoUSA data has a large and significant negative intercept. Square footage is positive in the linear term and close to zero in the quadratic, while 2+ stories is negative indicating single-story homes receive more utility from adopting relative to equal-sized two-story homes, possibly due to greater roof area for optimal solar panel siting.

The ι terms capture differential effects across boundary zones A through D (A is the omitted category) across wealth and type, which is of particular concern if solar installers use wealth for lead generation or targeting. Intercept shifts for consumption bins show a U-shaped form: the lowest consumption bin, ϕ_{bin1} is closest to the omitted category, bin 5, while ϕ_{bin2} is the lowest. This is consistent with households that are highly energy-cognizant investing in energy efficiency *and* solar adoption. Consumption bin intercept shifts for leases are positive, but only one is significant, and then only at the 10% level.

[62]The lack of significance (and negative sign) may also be the result of conditioning on wealth and consumption since households with children tend to have higher consumption. Households with longer length of residence tend to be higher wealth as well.

Table H.1: Structural parameter results (untransformed)

Param. Grp	Parameter	Estimate	se	t	pval	
ω	ω_0	-2.932	0.439	-6.678	0.000	***
	ω_{med}	1.548	0.413	3.744	0.000	***
	ω_{low}	2.223	0.577	3.856	0.000	***
ρ	ρ_0	1.951	0.001	1863.326	0.000	***
	α_{med}	-0.147	0.355	-0.415	0.678	
	α_{low}	-0.189	0.006	-29.325	0.000	***
β	β_0	25.094	0.526	47.693	0.000	***
	Wealth: middle 1/3rd	-5.609	1.901	-2.950	0.003	**
	Wealth: lowest 1/3rd	-4.621	0.889	-5.201	0.000	***
	Lease	-14.330	5.313	-2.697	0.007	**
	Lease x Wealth: middle 1/3rd	5.106	0.070	72.641	0.000	***
	Lease x Wealth: lowest 1/3rd	4.921	0.341	14.451	0.000	***
	Voter affiliation: D	-1.725	0.087	-19.912	0.000	***
	Voter affiliation: D + lease	2.111	0.964	2.190	0.029	*
	Has 1 or more lines of credit	1.829	0.084	21.722	0.000	***
	Lines of credit	0.776	0.399	1.945	0.052	.
	Length of residence	-0.349	0.184	-1.893	0.058	.
	Has children	-0.307	1.429	-0.215	0.830	
	Listed as renter-occupied in InfoUSA	-35.760	0.158	-226.558	0.000	***
	Sqft (1,000's)	3.777	0.145	26.108	0.000	***
	Sqft^2 (1,000's)	-0.311	0.092	-3.388	0.001	***
	Stories	-4.430	1.829	-2.422	0.015	*
ι_{area}	ι_B	-6.709	0.490	-13.702	0.000	***
	ι_C	4.615	0.687	6.720	0.000	***
	ι_D	1.829	0.667	2.743	0.006	**
	Wealth: lowest 1/3rd: ι_B	10.489	0.182	57.783	0.000	***
	Wealth: middle 1/3rd: ι_B	6.259	0.489	12.809	0.000	***
	Wealth: lowest 1/3rd: ι_C	0.996	0.684	1.456	0.145	
	Wealth: middle 1/3rd: ι_C	-5.615	0.657	-8.553	0.000	***
	Wealth: lowest 1/3rd: ι_D	1.910	3.072	0.622	0.534	
	Wealth: middle 1/3rd: ι_D	-4.093	0.543	-7.545	0.000	***
	$\iota_{B,lease}$	1.636	0.459	3.567	0.000	***
	$\iota_{C,lease}$	-1.336	0.688	-1.942	0.052	.
	$\iota_{D,lease}$	-1.667	0.203	-8.194	0.000	***
$\phi_{consumption}$	ϕ_{bin1}	-2.816	0.667	-4.224	0.000	***
	ϕ_{bin2}	-8.599	0.367	-23.440	0.000	***
	ϕ_{bin3}	-5.839	0.263	-22.181	0.000	***
	ϕ_{bin4}	-3.858	0.901	-4.283	0.000	***
	$\phi_{bin1,lease}$	2.128	0.349	6.101	0.000	***
	$\phi_{bin2,lease}$	2.799	0.363	7.709	0.000	***
	$\phi_{bin3,lease}$	2.886	1.068	2.701	0.007	**
	$\phi_{bin4,lease}$	1.430	1.576	0.908	0.364	

0 '***' 0.001 '**' 0.01 '*' 0.05 '.' 0.1 ' ' 1

Includes time fixed effects. Robust std. errors from Arcidiacono and Miller (2011)

Table H.2: Model predicted and actual adoption shares

		Total			Purchases			Leases		
Wealth	Households	Predicted	Actual	Predicted:Actual	Predicted	Actual	Predicted:Actual	Predicted	Actual	Predicted:Actual
High	61,229	2,545	2,445	1.041	1,682	1,553	1.083	863	892	0.967
Med	61,223	3,091	3,052	1.013	1,555	1,540	1.010	1,536	1,512	1.016
Low	61,215	1,777	1,747	1.017	686	711	0.965	1,091	1,036	1.053
All	183,667	7,413	7,244	1.023	3,923	3,804	1.031	3,490	3,440	1.015

To assess model fit, we use parameter estimates in Table H.1 to predict adoption probabilities for each household, time period, consumption bin, and type. We then calculate the probability of adoption within the study period as one minus the product of the probability of non-adoption in each time period. We then aggregate across consumption bin and type for each household using the weights from estimation. As in estimation, for adopters we observe the type (leaser vs. purchaser) and consumption bin with certainty, which is reflected in the weights. Finally, we aggregate by wealth bins and type to summarize total predicted adoption during the study period, which can be directly compared to observed adoptions. While our consumption bin weights are restricted to reflect the zip-level distribution of consumption, our type weights (leaser vs. purchaser) are unrestricted. We find little separation, with a mean lease-type weight of 0.486 and a standard deviation of 0.036.

We further assess model fit by examining the ratio of predicted to actual installations (as described above) by zipcode. Results reveal no systematic bias in wealth across zipcodes. Zipcodes with lower numbers of households in wealth bins (left side of plot H.1) are less accurate, but zipcodes with larger numbers of households are more accurate.

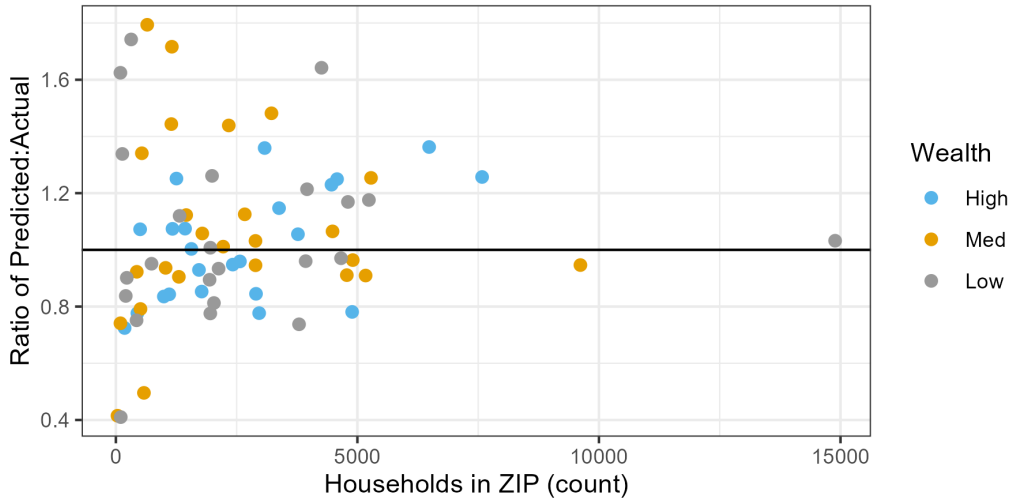

Figure H.1: **Model fit**. Ratio of predicted to observed installations by zip and wealth. Points near to 1 represent zip x wealth combinations with accurate predictions. Note that 7 zip x wealth bins with zero adoptions are omitted here.

I Robustness Checks

I.1 Model Estimates with Rebound

We resize optimal installations assuming a rebound of $\varepsilon = .1$ for purchased systems, holding leased systems sized to full offset as before (see Appendix C for sizing results). Using updated sizes, we re-estimate the model. Results are largely robust to sizing with rebound and are qualitatively identical to results presented above. With resizing, the estimated implicit discount rates across wealth are 10.2%, 14.7%, and 16.2%, which can be compared to results in Table 2 of 10.9%, 15.6%, and 17.2%. Differences between high- and low-wealth household implicit discount rates remain significant. While estimated implicit discount rates are lower at all wealth levels, the ratio of θ for high to low wealth remains identical at 1.5.

Table I.1: Implicit Discount Rates: Rebound elasticity for $\varepsilon = 0.1$

Wealth	Marginal Utility of Income	Annual Discount Rate	$\bar{\bar{\theta}}$	Ratio
High	0.07 (0.03)	10.2% (0.03%)	51.7	1.5
Med	0.28 (0.19)	14.7% (14.14%)	37.1	1.1
Low	0.51 (0.34)	16.2% (0.14%)	33.8	1
All		13.7%	40.9	–

Heteroskedasticity-robust standard errors in parentheses calculated by delta method. The difference between the high-wealth and low-wealth underlying marginal utility of income structural parameters is statistically significant, with a p-value of < 0.0001. The difference between the high-wealth and low-wealth underlying implicit discount rate structural parameters is statistically significant, with a p-value < 0.0001.

I.2 Model Estimates with Alternative Increase in Power Purchase Rates

We assume the rate of increase in power purchase rates for leased systems is contracted at 4% (see Equation 23), and in this section, we sensitivity test this assumption using $\zeta^{PPA} \in \{.03, .05\}$. Based on the data, we assume lease customers are assumed to size to full offset. Thus, we do not re-size under the alternative values of ζ^{PPA}.

Results are qualitatively similar on implicit discount rates, with $\zeta^{PPA} = .03$ resulting in estimated implicit discount rates of 10.8%, 15.5%, and 17.1% for high, medium, and low-wealth households. At $\zeta^{PPA} = .05$, estimated rates are 8.5%, 12.4%, and 13.7%, slightly lower than main estimates of 10.9%, 15.6%, and 17.2%.

Table I.2: Implicit Discount Rates: $\zeta^{PPA} = 0.03$

Wealth	Marginal Utility of Income	Annual Discount Rate	$\bar{\bar{\theta}}$	Ratio
High	0.06 (0.03)	10.8% (0.03%)	49.2	1.5
Med	0.25 (0.18)	15.5% (13.37%)	35.2	1.1
Low	0.47 (0.36)	17.1% (0.22%)	32.0	1
All		14.4%	38.8	–

Heteroskedasticity-robust standard errors in parentheses calculated by delta method. The difference between the high-wealth and low-wealth underlying marginal utility of income structural parameters is statistically significant, with a p-value of < 0.0001. The difference between the high-wealth and low-wealth underlying discount rate structural parameters is statistically significant, with a p-value of < 0.0001.

Table I.3: Implicit Discount Rates: $\zeta^{PPA} = 0.05$

Wealth	Marginal Utility of Income	Annual Discount Rate	$\bar{\bar{\theta}}$	Ratio
High	0.06 (0.03)	8.5% (0.02%)	60.1	1.5
Med	0.35 (0.24)	12.4% (10.17%)	43.6	1.1
Low	0.53 (0.36)	13.7% (0.07%)	39.9	1
All		11.5%	47.8	–

Heteroskedasticity-robust standard errors in parentheses calculated by delta method. The difference between the high-wealth and low-wealth underlying marginal utility of income structural parameters is statistically significant, with a p-value of < 0.0001. The difference between the high-wealth and low-wealth underlying discount rate structural parameters is statistically significant, with a p-value of < 0.0001.

I.3 Model Estimates with Alternative Scaling of η, ζ

In our main results, we estimate a value of η, the rate of decline of variable panel costs, from the data. We also estimate a zip code x consumption bin-specific value of ζ, the rate of increase in average electricity price. These estimates form the consumer expectations of future panel and avoided electricity rates associated with sizing and adopting. We sensitivity test around these results by scaling both η and ζ_{bz} by $\pm$ 10%, and by setting ζ_{bz} equal to a rate of 2.2% per year, a factor commonly used by solar installers for sizing calculations and used by Google Project Sunroof in calculating savings.

Estimated implicit discount rates for all three are little changed. Scaling η, ζ_{bz} by a factor of 1.1 results in estimated rates of 10.7%, 15.2%, and 16.8%, nearly identical to the main results. Scaling by a factor of .9 results in estimated rates of 10.6%, 15.3%, and 16.9% for high, medium, and low-wealth households, also nearly identical to the main results.

Table I.4: Implicit Discount Rates with scaled-up ζ, η

Wealth	Marginal Utility of Income	Annual Discount Rate	$\bar{\bar{\theta}}$	Ratio
High	0.06 (0.03)	10.7% (0.03%)	51.8	1.5
Med	0.25 (0.18)	15.2% (13.20%)	37.0	1.1
Low	0.48 (0.34)	16.8% (0.19%)	33.5	1
All		14.2%	40.8	–

Heteroskedasticity-robust standard errors in parentheses calculated by delta method. The difference between the high-wealth and low-wealth underlying marginal utility of income structural parameters is statistically significant, with a p-value of < 0.0001. The difference between the high-wealth and low-wealth underlying discount rate structural parameters is statistically significant, with a p-value of < 0.0001.

Table I.5: Implicit Discount Rates with scaled-down ζ, η

Wealth	Marginal Utility of Income	Annual Discount Rate	$\bar{\bar{\theta}}$	Ratio
High	0.11 (0.04)	10.6% (0.02%)	48.0	1.5
Med	0.31 (0.13)	15.3% (14.71%)	34.5	1.1
Low	0.55 (0.33)	16.9% (0.16%)	31.5	1
All		14.2%	38.0	–

Heteroskedasticity-robust standard errors in parentheses calculated by delta method. The difference between the high-wealth and low-wealth underlying marginal utility of income structural parameters is not statistically significant, with a p-value of 0.754. The difference between the high-wealth and low-wealth underlying discount rate structural parameters is statistically significant, with a p-value of < 0.0001.

Setting $\zeta_{bz} = 2.2\%$ for all consumption bins and areas primarily removes cross-sectional variation in energy price trends across consumption bins. Estimated implicit discount rates are nearly unchanged at 10.7%, 15.4%, and 17.1% for high-, medium-, and low-wealth households relative to our main results of 10.9%, 15.6%, and 17.2%, respectively.

Table I.6: Implicit Discount Rates with $\zeta = 2.2\%$

Wealth	Marginal Utility of Income	Annual Discount Rate	$\bar{\bar{\theta}}$	Ratio
High	0.11 (0.04)	10.7% (0.03%)	41.0	1.5
Med	0.35 (0.17)	15.4% (13.63%)	30.1	1.1
Low	0.56 (0.26)	17.1% (0.10%)	27.5	1
All		14.4%	32.9	–

Heteroskedasticity-robust standard errors in parentheses calculated by delta method. The difference between the high-wealth and low-wealth underlying marginal utility of income structural parameters is statistically significant, with a p-value of < 0.0001. The difference between the high-wealth and low-wealth underlying discount rate structural parameters is statistically significant, with a p-value of < 0.0001.

I.4 Model Estimates assuming $T = 22$

To test sensitivity of our results to the assumption on panel life, we set assumed panel life to $T = 22$. Panel life enters both the sizing model and the estimation of θ; here we set $T = 22$ in both to maintain consistency.

Results are highly similar to main results. Estimated implicit discount rates across wealth are 11.2%, 16.1%, and 17.9%, slightly higher than main results. The ratio of average θ between household wealth bins of 1.5:1 and 1.1:1 is nearly identical to the main results.

Table I.7: Implicit Discount Rates Assuming Panel Life $T = 22$.

Wealth	Marginal Utility of Income	Annual Discount Rate	$\bar{\bar{\theta}}$	Ratio
High	0.11 (0.02)	11.2% (0.02%)	45.5	1.5
Med	0.33 (0.09)	16.1% (13.62%)	33.0	1.1
Low	0.55 (0.22)	17.9% (0.14%)	30.0	1
All		15%	36.2	–

Heteroskedasticity-robust standard errors in parentheses calculated by delta method. The difference between the high-wealth and low-wealth underlying marginal utility of income structural parameters is statistically significant, with a p-value of < 0.0001. The difference between the high-wealth and low-wealth underlying discount rate structural parameters is statistically significant, with a p-value of < 0.0001.

I.5 Model Estimates including Race Covariates

Data from InfoUSA allow us to extract race and ethnicity for each household head. We group all households into four groups: Hispanic, White (non-Hispanic), Asian, and Black/Mixed and include indicators for this race variable in estimation. We allow indirect utility for adopting to vary by race interacted with wealth bins, but do not interact it with implicit discount rate or marginal utility of income. Race is not used in the sizing model, and enters only in the adoption model.

Results are similar to main results, though slightly higher across all wealth bins. Estimated implicit discount rates across wealth are 13.0%, 18.4%, and 20.4%, while the ratio of average θ between household wealth bins of 1.6:1 and 1.1:1 remains nearly identical to the main results.

Table I.8: Implicit Discount Rates Including Race Covariates

Wealth	Marginal Utility of Income	Annual Discount Rate	$\bar{\bar{\theta}}$	Ratio
High	0.02 (0.03)	13.0% (0.11%)	41.7	1.6
Med	0.23 (0.16)	18.4% (2.13%)	29.7	1.1
Low	0.49 (0.95)	20.4% (0.89%)	26.8	1
All		17.2%	32.7	–

Heteroskedasticity-robust standard errors in parentheses, calculated by delta method. The difference between the high-wealth and low-wealth underlying marginal utility of income structural parameters is statistically significant, with a p-value < 0.0001. The difference between the high-wealth and low-wealth underlying implicit discount rate structural parameters is statistically significant, with a p-value of < 0.0001.

I.6 Model Estimates when Randomizing Lease vs. Purchase

Our main specification uses the estimation method of Arcidiacono and Miller (2011) to integrate over permanent unobserved heterogeneity in the type of adoption a household would consider – lease vs. purchase. To address concerns that type is not permanently assigned but rather a function of discount rates and costs, we re-estimate the model after randomly reassigning observed adopters to a purchase or lease type. We interpret unchanged estimates as evidence that unobserved type is not pivotal in explaining adoption decisions. To preserve the wealth and consumption bin adoption distributions, we reassign adoption amongst adopters within strata defined by these key observed characteristics, keeping the counts and types of adoptions constant within strata.

Results in Table I.9 are largely unchanged from the main results with slightly lower estimated discount rates for all wealth levels. The ratio of high-wealth to low-wealth θ remains nearly identical to the main specification at 1.5:1.

Table I.9: Implicit Discount Rates Estimated when Randomizing Lease vs. Purchase

Wealth	Marginal Utility of Income	Annual Discount Rate	$\bar{\bar{\theta}}$	Ratio
High	0.003 (0.001)	10.8% (0.02%)	49.5	1.5
Med	0.316 (0.138)	15.7% (11.83%)	34.8	1.1
Low	0.608 (0.323)	16.6% (0.22%)	32.9	1
All		14.3%	39.0	–

Heteroskedasticity-robust standard errors in parentheses calculated by delta method. The difference between the high-wealth and low-wealth underlying marginal utility of income structural parameters is statistically significant, with a p-value of < 0.0001. The difference between the high-wealth and low-wealth underlying discount rate structural parameters is statistically significant, with a p-value of < 0.0001.

I.7 Model Estimates Assuming Finite Time Horizon

We re-estimate the adoption model as a finite time model as in de Groote and Verboven (2019), which implicitly constrains the value function to include zero utility and zero continuation value at the end of the 25-year lifespan of the panels. In contrast, our main specification assumes re-adoption at $T = 100$ quarters. Results in Table I.10 imply higher implicit discount rates than our main results. The ratio between high- and low-wealth households valuation of the flow of NEM benefits, captured by θ, are slightly higher (1.6:1) than the main specification results of 1.5:1.

Table I.10: Implicit Discount Rates Estimated Assuming Finite Time Model

Wealth	Marginal Utility of Income	Annual Discount Rate	$\bar{\bar{\theta}}$	Ratio
High	0.01 (0)	13.0% (0.04%)	41.7	1.6
Med	0.17 (0.14)	18.7% (13.13%)	29.2	1.1
Low	0.43 (0.62)	20.6% (0.60%)	26.5	1
All		17.3%	32.5	–

Heteroskedasticity-robust standard errors in parentheses calculated by delta method. The difference between the high-wealth and low-wealth underlying marginal utility of income structural parameters is statistically significant, with a p-value of < 0.0001. The difference between the high-wealth and low-wealth underlying discount rate structural parameters is statistically significant, with a p-value of < 0.0001.

I.8 Model Estimates Using Next-Period Realized State Values

In our main specification, we estimate CCP's from a first-stage flexible logit and generate Pr' terms (next-period probability of adoption) by advancing all state variables by the expected ζ, η, t. As an alternative, we generate Pr' terms instead by taking the one-period lead of the model's predicted probability of adopting. This assumes that individuals know and anticipate the next-period realizations of $\bar{p}$ and $C(K^*)$.

Results in I.11 indicate slightly higher estimated discount rates of 12.5% to 20.5%. The ratio of high-wealth to low-wealth θ remains close to the main specification at 1.6:1.

Table I.11: Implicit Discount Rates Estimated with Next-Period Realized State Values for Conditional Choice Probabilities

Wealth	Marginal Utility of Income	Annual Discount Rate	$\bar{\bar{\theta}}$	Ratio
High	0.003 (0.003)	12.5% (0.06%)	43.1	1.6
Med	0.221 (0.22)	18.3% (16.58%)	29.8	1.1
Low	0.485 (2.56)	20.5% (2.59%)	26.7	1
All		17%	33.2	–

Heteroskedasticity-robust standard errors in parentheses calculated by delta method. The difference between the high-wealth and low-wealth underlying marginal utility of income structural parameters is not statistically significant, with a p-value of 0.256. The difference between the high-wealth and low-wealth underlying discount rate structural parameters is statistically significant, with a p-value of < 0.0001.

J Details on Borrowing Costs and Expectations

J.1 Potential Role of Borrowing Costs

One explanation for the heterogeneous implicit rates we estimate may come from credit constraints that induce higher borrowing costs. Although our sample households are all homeowners who would be expected to have access to home equity lines of credit (which is one of our controls), credit constraints may still drive our results. To explore this possibility, we merge our sample with household-level credit data from Experian based on the household address.

We obtain a 63% match rate with Experian data. Of those matching, only 6.7% have credit scores in the sub-prime (less than 680) or near-prime (less than 780) range, which we consider to be "poor" credit. Adopters have a slightly smaller share of poor credit at 5.9%. Poor credit is more prevalent in low-wealth households (12.4%) and medium-wealth households (6.0%) relative to high-wealth households (2.7%). Thus, we construct a binary indicator for poor credit, and another binary indicator for un-matched households.

We estimate our model incorporating credit data and show that heterogeneity across wealth is persistent, even conditional on credit. The intent here is to flexibly control for the credit-induced cost of borrowing, allowing for a comparison across wealth conditional on prime-or-better credit. We augment the specification shown in Section 6.4 with interactions between poor credit and wealth, as well as interactions between un-matched status and wealth. Due to the small share of high-wealth households that have poor credit, we restrict the effect of poor credit to be zero for high-wealth households. We include intercept, lease intercept, and marginal utility of income (ω) interactions for poor credit as well, all interacted with low- or medium-wealth. We fully interact 'missing' status with wealth in intercept, lease intercept, and discount rate.

The results in J.1 suggest that credit constraints are not likely to fully explain heterogeneity in discount rates. We estimate conditional (on prime or better credit) implicit discount rates of 15.9% to 25.2% (high-wealth and low-wealth, respectively). Even when we control for credit-by-wealth interactions, the difference in implicit discount rates between low-wealth and high-wealth households is still statistically significant, though substantially higher. Importantly, the $\bar{\theta}$ ratio is nearly identical to our main results at 1.6:1, showing that even for homeowners with excellent credit (and thus would presumably have access to ample credit), there are notable differences in implicit discount rates.

Table J.1: Implicit Discount Rates Estimated with Interacted Controls for Credit

Wealth	Credit	Marginal Utility of Income	Annual Discount Rate	$\bar{\bar{\theta}}$	Ratio
High	All	0.0012 (8e-04)	15.9% (0.28%)	34.2	1.6
Med	$\geq$ Prime	6e-04 (0.0021)	15.9% (1.69%)	34.2	1.6
Low	$\geq$ Prime	0.167 (0.1079)	25.2% (0.46%)	21.6	1
Med	Poor	0.209 (1.1268)	29.4% (10.88%)	18.5	2.1
Low	Poor	0.1146 (0.2899)	66.2% (55.06%)	9.0	1
All			20.1%	29.6	–

All values are conditional on matching household credit (63% of sample). Heteroskedasticity-robust standard errors in parentheses, calculated by delta method. The difference between the high-wealth and low-wealth underlying marginal utility of income structural parameters is statistically significant, with a p-value of < 0.0001. Though the point estimates are non-monotonic in wealth between high- and medium-wealth underlying marginal utility of income structural parameters, the difference between them is not statistically significant, with a p-value of 0.087. The difference between the high-wealth and low-wealth underlying discount rate structural parameters for prime or better credit is statistically significant, with a p-value of < 0.0001.

J.2 Survey Evidence on Home Tenure Expectations

To better understand consumer expectations about the length of time they will live in their current home, we draw upon evidence from Bollinger et al. (2025). That study runs a nationwide representative survey using the vendor Dynata. There was a screening requirement for participants to be considering purchasing (or have recently purchased) either solar panels or an electric vehicle. There are also results restricted to homeowners. The final sample of 3,305 respondents, of which 1,673 are homeowners. All but one homeowner answered the income question (non-homeowners were more likely to not respond to the income question). Figure J.1 presents the fraction of responses in an income bin by the expected home tenure bin (there are five possible bins they could choose from). In other words, on the far left we see that roughly 32% of households making over $150,000 per year expect to stay in their home more than 20 years. A key takeaway from this figure is that a similar fraction of households expect to stay in their homes for over 20 years across all income groups. Broadly, the figure shows similar patterns, with some increase in very short tenures from the lowest-income group, possibly due to retirees who do not expect to stay long in their home. Note that few homeowners earn under $50,000 a year in California, so we are not very concerned about this result.

There are also similar expectations across income groups for future electricity price growth, as can be seen in Figure J.2 below. There is an extremely similar figure for expectations of future solar prices as well (Figure J.3). In short, the evidence suggests that there are negligible differences across wealth groups in expectations of home tenure, future electricity prices, and future electricity prices.

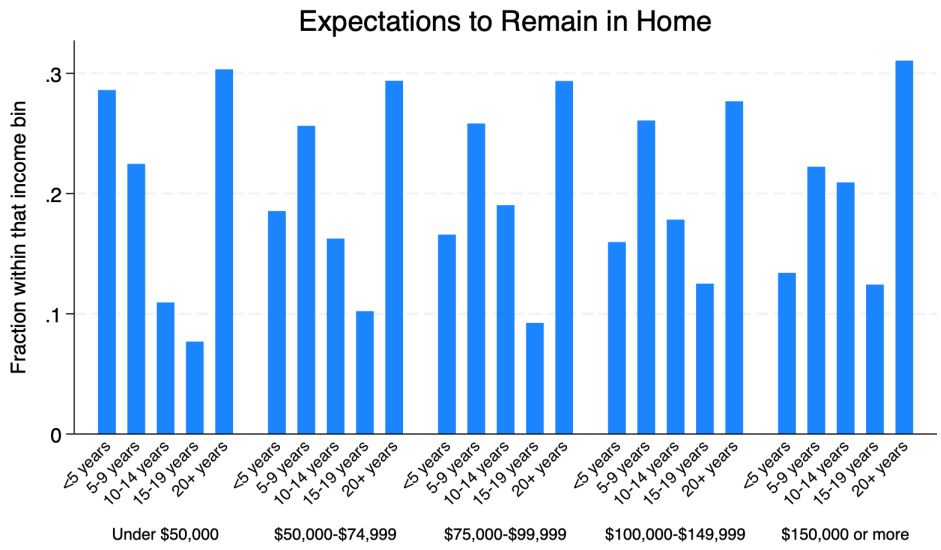

Figure J.1: **National Survey**. Expectation to stay in current residence by income.

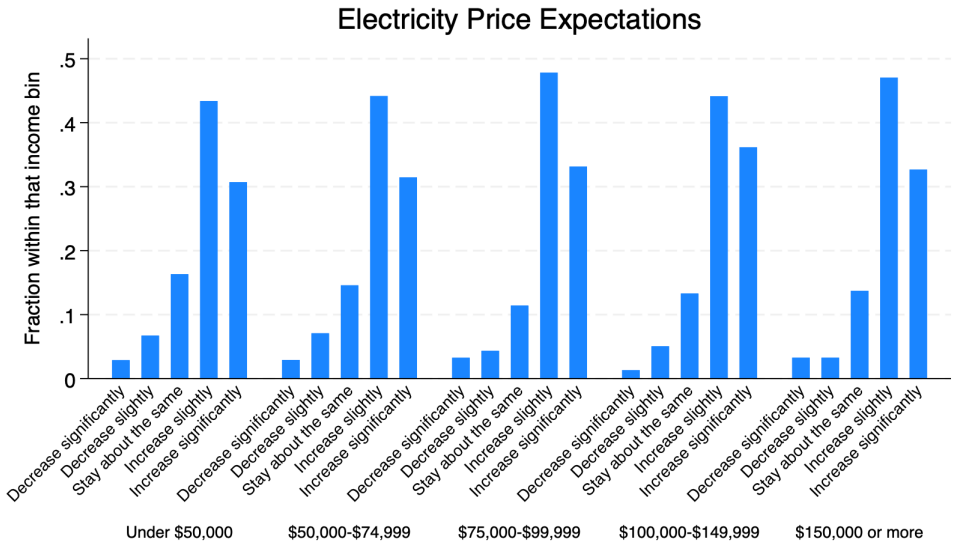

Figure J.2: **National Survey**. Expected electricity prices by income.

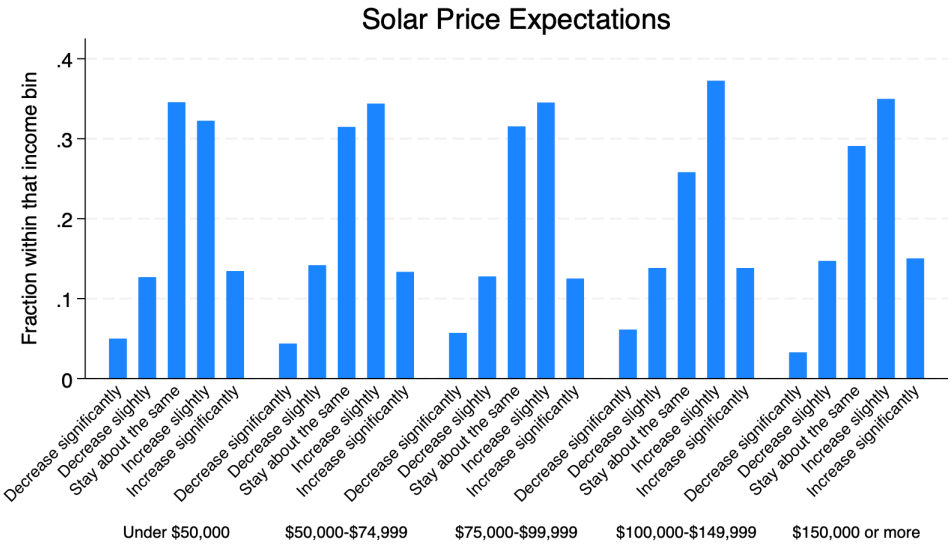

Figure J.3: **National Survey**. Expected solar prices by income.

K Counterfactual Methodology

K.1 Counterfactual Methods

All counterfactual scenarios report the ratio of predicted counterfactual adoptions to baseline predicted adoptions occurring within the study sample and during the study period. In all cases, we calculate adoption as one minus the joint probability of choosing "do not adopt" in each of the 10 time periods in the study window. This allows us to report the total number of adoptions predicted over the period, rather than examining across time periods. In the following, we dive into the details of the counterfactual methodology.

The value of adoption, δ_1, is a terminal state and so there is no continuation value. To calculate δ_0, we need to calculate the value function. We do this by taking the infinite sum given in 1. In practice, we forward simulate the state variables many periods in the future (S periods), assume $V(p_S, VC_S, FC_S) = \delta_1(p_S, VC_S, FC_S)$ (i.e. treat the value as if everyone who had not yet adopted solar at period S purchases solar), which is a benign assumption if S is large enough. We then then backward iterate to calculate δ_0 in the current period using the below expression for the inclusive value of adoption in each period:

$$\begin{aligned} & E\left[\max(\delta_0 + \sigma\epsilon_0, \delta_1 + \sigma\epsilon_1)\right] \\ &= \sigma E\left[\max(\delta_0/\sigma + \epsilon_0, \delta_1/\sigma + \epsilon_1)\right] \\ &= \sigma\left[\log\left(\exp(\delta_0/\sigma) + \exp(\delta_1/\sigma)\right) + \gamma\right]. \end{aligned} \tag{29}$$

Under counterfactual regimes, we must account for the fact that the probability of adopting in the next period will also change. I.e., counterfactuals require counterfactual CCP's – as the utility payoff of adopting changes (in one or all periods), the probability of adopting in the next period also changes. Here, we use model estimates to update both $\delta_1 - \delta_0$ and $\log(Pr')$ as well.

Our strategy for calculating counterfactual $\log(Pr'^{+})$ is to calculate the *difference* in adoption utility implied by the estimated parameters of the model. Arcidiacono and Miller (2020) discuss identification of counterfactual CCPs and show conditions under which counterfactual CCPs are identified. Single action finite dependence (which holds if the distribution of states does not depend on the initial choice for a particular action) is one such assumption. Under our counterfactual, we continue with our assumption that households' expectations are correct excepting a short term prediction error (as in Scott (2014) and de Groote and Verboven (2019)) using the same solar and electricity price trends that existed in the current environment.

We write future changes in the probability of adopting as an infinite series of (known) flow utility payoffs. Let A capture the scaled (by $1/\sigma$) utility of adopting today relative to the discounted utility of adopting tomorrow, $A = (u - \rho u')/\sigma$, and B capture the next period adoption probability and Euler-Mascheroni constant, so that we can write the value of adopting as:

$$\begin{aligned} \frac{1}{\sigma}(\delta_1 - \delta_0) &= \overbrace{\frac{1}{\sigma}(1-\rho(1+\zeta))\theta q^* \bar{p} - \frac{1}{\sigma}(1-\rho)FC - \frac{1}{\sigma}(1-\rho\eta)VC + \frac{1}{\sigma}(1-\rho)X\beta}^{A} \\ &+ \underbrace{\rho\left(\log(Pr') - \gamma\right)}_{B}. \end{aligned}$$

Under each counterfactual scenario, we change the utility of adopting today by Δu and the expected utility of adopting tomorrow by $\Delta u'$, so that $\Delta A = (\Delta u - \rho\Delta u')/\sigma$.

The counterfactual change in the probability of adopting, Pr^+, is determined by the change in A plus the change in B:

$$Pr^+ = \Lambda\left(\frac{1}{\sigma}(\delta_1 - \delta_0) + \Delta A + \Delta B\right). \tag{30}$$

Note we can write:

$$\Lambda^{-1}(Pr) \quad = \quad \log\left(\frac{Pr}{1-Pr}\right),$$

and thus

$$\begin{aligned}\log(Pr') &= \Lambda^{-1}(Pr') + \log(1-Pr') \\ &= \frac{1}{\sigma}\left(\delta_1' - \delta_0'\right) + \log(1-Pr'),\end{aligned} \tag{31}$$

where δ_1' is the next-period utility of adopting and δ_0' is the next period utility of non-adoption.

Using this expression, we can write B as:

$$B = \rho(\log(Pr') - \gamma) = \rho\left(\frac{1}{\sigma}(\delta_1' - \delta_0') + \log(1-Pr')\right) - \rho\gamma,$$

and we can write B^+ (the value of B under the counterfactual) as:

$$B^+ = \rho\left(\log(Pr'^+) - \gamma\right) = \rho\left(\frac{1}{\sigma}(\delta_1'^+ - \delta_0'^+) + \log(1-Pr'^+)\right) - \rho\gamma.$$

By plugging in the values for δ and δ^+, we can write the change in B as:

$$\begin{aligned}\Delta B &= B^+ - B \\ &= \rho\left(\frac{1}{\sigma}(\delta_1' - \delta_0') + \Delta A' + \Delta B' + \log(1-Pr'^+)\right) - \rho\gamma - B \\ &= \rho\left(\frac{1}{\sigma}(\delta_1' - \delta_0') + \Delta A' + \rho(\log(Pr''^+) - \log(Pr'')) + \log(1-Pr'^+)\right) - \rho\gamma \\ &\quad -\rho\left(\frac{1}{\sigma}(\delta_1' - \delta_0') + \log(1-Pr')\right) - \rho\gamma \\ &= \rho\Delta A' + \rho^2(\log(Pr''^+) - \log(Pr'')) + \rho(\log(1-Pr'^+) - \log(1-Pr')) \\ &= \rho\Delta A' + \rho\Delta B' + \rho(\log(1-Pr'^+) - \log(1-Pr')),\end{aligned} \tag{32}$$

where $\Delta A'$ is the next-period difference between counterfactual and actual A. ΔB can be rewritten using the following recursion:

$$\begin{aligned}\Delta B &= \rho\Delta A' + \rho(\log(1-Pr'^+) - \log(1-Pr')) + \rho\Delta B' \\ &= \rho\Delta A' + \rho^2\Delta A'' + \rho\left(\log(1-Pr'^+) - \log(1-Pr')\right) \\ &+ \rho^2\left(\log(1-Pr''^+) - \log(1-Pr'')\right) + \Delta B'',\end{aligned} \tag{33}$$

where we denote the next period's next-period logged probability of adopting as $\log(Pr'')$ and $\log(Pr''^+)$, and additional future probabilities of adopting as $\log(Pr_s^{'})$ and $\log(Pr_s^{'+})$.

The expression for ΔB in (33) is a recursive sum that can rewritten as:

$$\Delta B = \sum_{s=1}^{\infty} \rho^s \Delta A_s' + \sum_{s=1}^{\infty} \rho^s \left(\log \left(\frac{1 - Pr_s'^{+}}{1 - Pr_s'} \right) \right). \tag{34}$$

The s subscript on A_s' and Pr_s' indicates that A, Pr, and Pr'^{+} change in future periods based on the evolution of the state variables and are affected by the counterfactual changes (electricity prices, solar prices, etc.). Each recursion produces $\Delta A'$ that is equal to ΔA scaled by ρ,[63] as well as a discounted logged ratio of the probability of not adopting in the next period under the counterfactual to the probability of not adopting in the next period under the baseline. While this ratio is close to unity, a positive ΔA, for example, will result in a higher probability of adopting in subsequent periods, decreasing the probability of not adopting and decreasing the logged ratio.

To calculate the second term in equation (34), we forward-simulate the deterministic progression of C and $q^*\bar{p}$ to the period $s = S = 151$ using the trends given by η and $\zeta_{b,z}$. We assume that all households will make a terminal choice at this time period, and calculate the expected value of that choice as $\log(\exp(\frac{\delta_1^{(S)}}{\sigma}) + exp(0)) + \gamma$. Setting $s = 150$, we substitute the expected value of not adopting as the previous quantity scaled by ρ and calculate the probability of adopting at $s = 150$. We record the log of one minus that probability, scaled by ρ^s. We then iterate from $s = 150$ to $s = 0$, taking a cumulative sum of $log(1 - Pr)$ at each value of s. At $s = 0$, we discard the continuation value in favor of using Equation 34, and substitute the cumulative sum of $log(1 - Pr)$ for the denominator of the second term. We then calculate δ_1 under the counterfactual – a straightforward task for any counterfactual that re-scales C and/or $q^*\bar{p}$, and repeat the process to calculate the numerator of the second term in 34.

The first term in equation (34) is straightforward. Adding $\Delta A'$ to 34, we see:

$$\Delta A + \Delta B = \sum_{s=0}^{\infty} \rho^s \Delta A_s' + \sum_{s=1}^{\infty} \rho^s \left(\log \left(\frac{1 - Pr_s'^{+}}{1 - Pr_s'} \right) \right),$$

which allows for closed form calculation of the infinite sum of the first term, including cases where $\Delta A' = \eta \Delta A$, as is the case when counterfactual changes apply to the variable cost term.

Finally, with $\Delta A + \Delta B$ in hand, we calculate the counterfactual Pr'^{+} using Equation 30.

K.2 Counterfactual Subsidy Amounts

Exported to grid	Subsidy Title	Per-W	Per-install
30%	Per-W only	\$0.30	\$0
30%	Per-W and per-install	\$0.15	\$731
50%	Per-W only	\$0.84	\$0
50%	Per-W and per-install	\$0.42	\$1,961

We set the total subsidy amount equal to half of the reduction in utility expenditures from the NEM reform before allowing for re-sizing, splitting this in half again between per-W and

[63] For example, under a counterfactual price change that scales the variable cost of solar by some amount in every period, ΔA is equal to that difference in the variable installation costs and we have that $\Delta A' = \eta \Delta A$. A counterfactual change in the fixed costs of solar leads to $\Delta A' = \Delta A$.

per-install subsidies. For each possible share exported to grid, we then adjust these subsidies upwards proportionally until the net change in high-wealth households is equal to zero, an increase of approximately 4%. Actual reductions in the utility expenditures depend on both the re-sizing conditional on export rates and subsidies, and the adoption decision conditional on export rates, subsidies, and optimal size. We report changes in adoption inclusive of these adjustments in Table 5.

L Counterfactual Welfare-Relevant Outcomes

L.1 Calculation Methodology

We formulate demand as the probabilistic quantity of capacity installed as a function of the fixed and variable costs of installation:

$$d_{ibet}(C_t) = K^*(C_t, p_t; \gamma_i, e_i, b_i)\tilde{Pr}^{adopt}(C_t, FC_i, p_t, e_i, b_i, X_i, K^*(C_t, p_t; \gamma_i, e_i, b_i)), \tag{35}$$

where K^* is the optimal sizing function, C_t is the per-watt cost of a panel at time t, p_t are the retail electricity rates, γ_i represents the roof profile of household i, and b_i and e_i are household consumption bin b and type e, respectively. This represents quantity (in watts of solar capacity) demanded conditional on adopting. F_i is the fixed cost of installation which does not enter the optimal sizing K^*. $\tilde{Pr}^{adopt}$ is the unconditional probability of adopting conditional on K^*, household demographics X_i, time t, type, consumption bin, F, C_t, and p_t. We can use model predictions to estimate a conditional probability of adopting Pr^{adopt}. Thus, we calculate $\tilde{Pr}^{adopt}$ using the following:

$$\tilde{Pr}_t^{adopt} = \Pi_{s=1}^{t}\left(1 - \mathbb{1}(s \geq 2)Pr_{t-s+1}^{adopt}\right)Pr_t^{adopt}.$$

For each household, type, consumption bin, time period, and value of C_t, we calculate K^* as a function of price, generate $q^*\bar{p}$ and TC, and use the adoption model to calculate the probability of adopting.

L.1.1 Consumer surplus

Household surplus is calculated by integrating equation (35) over average cost per watt, defined as:

$$CS_{ibe} = \sum_{t=1}^{10}\int_0^m \left(\tilde{C}_t + x\right) d_{ibet}(\tilde{C}_t + x)dx,$$

where the term $\tilde{C}_t$ is given as:

$$\tilde{C}_t = \frac{F + K^*C_t}{K^*}.$$

The upper limit of integration m is chosen to approximate the intercept of the demand curve. Because Type I extreme value shocks are unbounded, the probability of adopting never reaches 0. In practice, we numerically integrate over the range of $[0, 5.60]$, and assume that the probability of adopting when $\tilde{C}_t$ is \$5.60 per watt larger than what is observed in the data is negligible.[64] The consumer surplus, weighted by the unconditional probability of adopting, is summed over the 10 time periods in the sample.

Because K^*, q^*, and other important quantities vary over time within a household and type, we calculate the expected capacity installed, generation, government expenditures including counterfactual subsidies and expenditures on the federal investment tax credit (ITC), utility expenditures, and avoided damages using a similar equation, which appropriately weights each time period by the probability of not adopting prior to that period and the probability of adopting:

[64] the maximum value of $\tilde{C}_t$ in our data is \$4.17/W.

$$K^*_{ibe} = \sum_{t=1}^{10} \tilde{Pr}_t^{adopt} K^*_{ibet}. \tag{36}$$

Finally, we take the weighted sum across type and consumption bin using w_{ibe}, resulting in a total for each household in our sample.

L.1.2 Government expenditures

Counterfactual 3 explores the replacement of the NEM embedded subsidy with an explicit subsidy paid per-watt and per-installation. In addition to explicit subsidies, payment of the federal ITC will change under the counterfactuals (e.g., fewer households might adopt smaller systems). We report these effects as $\Delta GovernmentExpenditures$. Changes in private expenditures are accounted for in CS_i.

L.1.3 Utility profits and expenditures

We evaluate the change in utility expenditures by considering the subsidies embedded in the NEM policy. Because current NEM policy provides electric bill savings in excess of the avoided cost savings, the difference between the avoided cost and $\bar{p}$ is an embedded subsidy. Furthermore, critics of NEM policy note that the embedded NEM subsidy is substantial, is primarily paid to high-wealth solar adopters, and increases electricity prices on lower-wealth non-adopters. Our second counterfactual is based on a reform of NEM policy that does not entirely eliminate the embedded subsidy – under this counterfactual, households can no longer "bank" excess solar generation (and are instead paid at the avoided cost rate for export to the grid). Thus, we calculate the change in utility expenditures as the change in the embedded NEM subsidy, calculated as:

$$\Delta UtilityExpenditures_{ibet} = \theta_{6\%} \left((1-s) \times q^{*'} (\bar{p'} - 0.062) - q^*(\bar{p} - 0.062) \right),$$

in which $\theta_{6\%}$ is the value of θ with a 6% discount rate (an appropriate rate for an investor-owned utility), s is the export share $s \in \{0.3, 0.5\}$, $Subs^W$ and $Subs^I$ are the explicit subsidies, $q^{*'}$ and $\bar{p'}$ are counterfactual generation and average value of offset electricity after resizing. $\bar{p'} \geq \bar{p}$ as the average value of non-grid-exported electricity increases as more is exported to the grid. We calculate total $\Delta Subsidy_i$ using equation (36) and then summing over type and consumption using the weights w_{ibe}.

L.1.4 Installer surplus

The installer margin for any given installation is given by $p(\frac{m^{solar}}{1+m^{solar}})$, where p is the installation price inclusive of the ITC. We assume a 71% margin (main manuscript Section 7.3), which corresponds to a markup of 2.5. We can replace price with expected revenue to calculate installer surplus (pre-tax). We report aggregated installer surplus in relevant welfare tables. For contribution to the MVPF, we calculate the fiscal externality on government expenditures in Appendix M.1.

L.1.5 Environmental damages avoided

A primary goal of NEM policy is to encourage the adoption of renewable energy in order to decrease negative externalities associated with fossil-fuel based electricity generation. Counterfactual welfare calculations include a "dollarization" of environmental and health damages avoided based on Sexton

et al. (2021), which calculates the hourly marginal emissions factors for all dispatchable generators capable of serving, *inter alia*, the relevant portion of the grid operated by the California Independent System Operator (CAISO) during the period including 2014-2016. These marginal emissions factors for each potentially responding plant are calculated as dollars of damages per unit of emissions of three criteria pollutants: NOx, SOx, and PM2.5, plus CO2 emissions.[65] Yearly damages are calculated for each potentially responding plant based on emissions location, local dispersion, and local population density and composition. Dollarized yearly damages are then summed over all potentially responding plants based on estimated responsiveness. Hourly generation profiles for each zip code in the sample are used to convert one watt of capacity K to avoided emissions, and yearly damages are summed to the 25-year lifetime of the panel using a 2% social discount rate, a value of $\lambda = 0.8\%$, and an expected decline in grid emissions equal to 1.75% per year. Finally, we sum expected environmental damages avoided over sample time and type as noted before.

L.2 Counterfactual Welfare Results

Counterfactual per-W and per-install subsidies are as follows:

Table L.1: Total Welfare for a 50% export rate.

Wealth	Consumer Surplus $1,000's	Installer Surplus $1,000's	Gov Subsidy Expenditures $1,000's	Gov ITC Expenditures $1,000's	Utility Surplus $1,000's	Avoided Damages $1,000's	Adoptions	Total Capacity kW	Total Generation MWh/y	Total Generation kWhy/kW
Baseline										
High	$42.2	$46.6	$0.0	$19.6	-$55.0	$44.0	2545	14,014	21,046	1,502
Med	$35.4	$56.5	$0.0	$23.7	-$63.5	$53.7	3091	17,007	26,898	1,582
Low	$15.7	$29.4	$0.0	$12.3	-$30.5	$27.7	1777	8,769	13,797	1,573
All	$93.3	$132.5	$0.0	$55.7	-$149.0	$125.4	7413	39,791	61,741	1,552
Lowered Compensation for Exported Solar										
High	$28.3	$38.4	$0.0	$16.1	-$35.8	$36.1	2347	11,494	17,530	1,525
Med	$27.2	$45.2	$0.0	$19.0	-$40.0	$42.7	2730	13,535	21,618	1,597
Low	$13.4	$23.9	$0.0	$10.0	-$19.1	$22.3	1654	7,038	11,216	1,594
All	$68.9	$107.5	$0.0	$45.2	-$94.8	$101.1	6731	32,068	50,365	1,571
Lowered Compensation for Exported Solar plus per-W Subsidy										
High	$33.6	$37.8	$11.6	$15.9	-$42.3	$43.4	2545	13,826	20,831	1,507
Med	$34.3	$51.2	$15.7	$21.5	-$54.0	$59.1	3403	18,727	29,598	1,580
Low	$17.5	$28.9	$8.8	$12.1	-$27.2	$33.1	2107	10,464	16,422	1,569
All	$85.4	$117.8	$36.1	$49.5	-$123.5	$135.6	8054	43,017	66,851	1,554
Lowered Compensation for Exported Solar plus Split Subsidy										
High	$32.4	$40.1	$10.5	$16.9	-$40.7	$41.5	2535	13,229	20,041	1,515
Med	$34.3	$54.1	$14.1	$22.7	-$51.7	$56.2	3395	17,800	28,256	1,587
Low	$18.8	$31.9	$8.9	$13.4	-$27.6	$32.8	2264	10,356	16,368	1,581
All	$85.4	$126.1	$33.5	$53.0	-$120.0	$130.5	8194	41,385	64,665	1,563

[65] a Social Cost of Carbon of $175/ton (in $USD2017) is used (Rennert et al., 2022).

Table L.2: Total Welfare for a 30% export rate.

Wealth	Consumer Surplus $1,000's	Installer Surplus $1,000's	Gov Subsidy Expenditures $1,000's	Gov ITC Expenditures $1,000's	Utility Surplus $1,000's	Avoided Damages $1,000's	Adoptions	Total Capacity kW	Total Generation MWh/y	Total Generation kWhy/kW
Baseline										
High	$42.2	$46.6	$0.0	$19.6	-$55.0	$44.0	2545	14,014	21,046	1,502
Med	$35.4	$56.5	$0.0	$23.7	-$63.5	$53.7	3091	17,007	26,898	1,582
Low	$15.7	$29.4	$0.0	$12.3	-$30.5	$27.7	1777	8,769	13,797	1,573
All	$93.3	$132.5	$0.0	$55.7	-$149.0	$125.4	7413	39,791	61,741	1,552
Lowered Compensation for Exported Solar										
High	$34.7	$43.8	$0.0	$18.4	-$47.5	$41.2	2462	13,136	19,845	1,511
Med	$31.9	$52.1	$0.0	$21.9	-$53.4	$49.4	2911	15,647	24,821	1,586
Low	$15.1	$27.1	$0.0	$11.4	-$25.4	$25.4	1695	8,044	12,697	1,578
All	$81.7	$122.9	$0.0	$51.6	-$126.4	$116.1	7069	36,828	57,362	1,558
Lowered Compensation for Exported Solar plus per-W Subsidy										
High	$36.6	$43.0	$4.1	$18.0	-$49.6	$43.3	2536	13,792	20,773	1,506
Med	$34.3	$54.2	$5.2	$22.8	-$58.9	$55.0	3176	17,403	27,532	1,582
Low	$16.5	$28.9	$2.8	$12.1	-$28.7	$29.1	1878	9,209	14,482	1,573
All	$87.5	$126.1	$12.1	$53.0	-$137.1	$127.4	7591	40,404	62,787	1,554
Lowered Compensation for Exported Solar plus Split Subsidy										
High	$36.2	$44.0	$3.9	$18.5	-$49.2	$42.9	2534	13,660	20,603	1,508
Med	$34.4	$55.4	$4.9	$23.3	-$58.3	$54.2	3174	17,179	27,206	1,584
Low	$17.0	$29.9	$2.8	$12.6	-$28.9	$29.1	1917	9,202	14,494	1,575
All	$87.6	$129.3	$11.6	$54.3	-$136.3	$126.2	7626	40,040	62,303	1,556

Table L.3: Changes in Welfare-Relevant Outcomes (Millions $; % Change)

Wealth	Δ Consumer Surplus	Δ Installer Surplus	Δ Utility Surplus	Δ Avoided Damages	Δ Government Subsidies
CF1: Lowered Compensation for Exported Solar					
High	-$7.5 (-22%)	-$2.9 (-7%)	$7.5 (16%)	-$2.8 (-7%)	-$1.2 (-7%)
Med	-$3.5 (-11%)	-$4.4 (-8%)	$10.1 (19%)	-$4.3 (-9%)	-$1.9 (-8%)
Low	-$0.5 (-4%)	-$2.3 (-8%)	$5.0 (20%)	-$2.3 (-9%)	-$1.0 (-8%)
All	-$11.6 (-14%)	-$9.6 (-8%)	$22.6 (18%)	-$9.3 (-8%)	-$4.0 (-8%)
CF2: Lowered Compensation for Exported Solar plus Split Subsidy					
High	-$6.0 (-17%)	-$2.6 (-6%)	$5.8 (12%)	-$1.1 (-3%)	$2.8 (13%)
Med	-$1.0 (-3%)	-$1.2 (-2%)	$5.3 (9%)	$0.5 (1%)	$4.4 (16%)
Low	$1.3 (8%)	$0.6 (2%)	$1.6 (6%)	$1.4 (5%)	$3.0 (20%)
All	-$5.7 (-6%)	-$3.2 (-2%)	$12.7 (9%)	$0.8 (1%)	$10.3 (16%)
CF2-CF1: Add the Split Subsidy to Lowered Compensation					
High	$1.5 (5%)	$0.3 (1%)	-$1.7 (-4%)	$1.6 (4%)	$4.0 (19%)
Med	$2.6 (8%)	$3.3 (6%)	-$4.9 (-10%)	$4.8 (10%)	$6.3 (24%)
Low	$1.8 (11%)	$2.9 (10%)	-$3.4 (-14%)	$3.7 (14%)	$4.0 (28%)
All	$5.9 (8%)	$6.4 (5%)	-$10.0 (-9%)	$10.1 (9%)	$14.3 (23%)

All estimates assume 30% export to grid.

M Marginal Value of Public Funds (MVPF)

M.1 MVPF Methodology

The MVPF is the ratio of the policy beneficiaries willingness to pay for a policy to the government expenditures of the policy (Finkelstein and Hendren, 2020). We calculate the MVPF for counterfactual 2 relative to counterfactual 1 as the addition of up-front subsidies along with the NEM reform provides an appropriate setting to evaluate those expenditures. We construct the MVPF for each wealth bin to separately examine the effectiveness of counterfactual 2 spending across the wealth distribution.

The MVPF takes the form of

$$\left(\Delta' CS + \Delta' InstallerSurplus + \Delta' UtilityExpenditures + \Delta' Damages\right) / \Delta' GovtExpenditures,$$

in which Δ' is the difference between counterfactual 2 and counterfactual 1 values. Under the 50% export assumption, direct changes in government expenditures via the ITC and subsidies increase from -\$10.5M to +\$30.9M (see Table 6). However, the MVPF includes the fiscal externality of the change in government expenditures resulting from changes in tax revenues on utility and installer profits. Damages and consumer surplus are unchanged in MVPF calculations.

M.1.1 Utility MVPF effects

Following Hahn et al. (2024) and noting that the utility in our case is an investor-owned utility (so α in their study is zero), we define the government fiscal externality from utility tax revenues as:

$$(P - [(LCOE + c_t) \cdot (1 + m)]) \cdot \tau,$$

in which LCOE is the levelized avoided cost of solar energy.

Further noting that $(LCOE + c_t)$ is equal to the avoided cost calculated in E3 (2022), we calculate the government fiscal externality as the sum across observations of $(P_i - 0.062)(1 + m)\tau$, weighted by the probability of adopting. We use a corporate tax rate of $\tau = 0.10$ and an economy-wide markup of $m = 0.08$. As an illustrative calculation and to roughly match real-world short-run marginal borrowing costs, we use $\theta_{4\%}$, a government discount rate of 4%, to calculate the present value of the fiscal externality which reflects the borrowing cost associated with changes in tax revenue. Our results would scale with different assumptions of the government borrowing rate.

In the MVPF numerator, we take the change in Consumer Surplus, $\Delta' CS$ and the calculated change in damages $\Delta' Damages$ directly from Table 6. We account for the change in utility expenditure as in Table 6, but adjust for markup and the corporate tax rate using the same assumptions as before. We calculate utility expenditures for each potential adoption and sum over all observations, weighed by probability of adopting. In order to adjust for markup m and tax rate τ, we calculate the change in utility expenditures for any possible installation as:

$$(P - [(LCOE + c_t) \cdot (1 + m)]) \cdot (1 - \tau).$$

To calculate the total expected effect, we multiply this number by the probability of the installation occurring and aggregate across all possible installations. We calculate the present value of this flow using $\theta_{6\%}$ as in Appendix L.

M.1.2 Installer MVPF effects

Installer profits change due to the different number and sizes of installations under the counterfactual. To assess the impact on profits, we again assume a markup of $m^{solar} = 2.5$ (Appendix L). The MVPF calculation in (Hahn et al., 2024) accounts for industry average markup $m = 0.08$ and divides installer surplus between fiscal externality and installers using corporate tax rate $\tau = 0.10$, as in utility effects above. We calculate analogous fiscal externalities for installers, using the implied marginal cost of $\frac{p}{m^{solar}+1}$. The surplus generated by a given installation, relative to the economy-wide average is:

$$p - \left(\frac{p}{m^{solar}+1}\right)(1+m) = p\left(\frac{m^{solar}-m}{1+m^{solar}}\right).$$

We can replace price with revenue to get an expression for the total additional surplus relative to the economy average, and then multiply by τ to calculate the additional government surplus and by $1-\tau$ to calculate the installers' surplus.

Pass-through and Installer MVPF Alternative pass-through assumptions would affect installer surplus, government expenditures, and the fiscal externalities from installer profits. Pless and Van Benthem (2019) find a range of pass-through between 78.7% for purchased systems and 133.3% for leased systems (indicating over-shifting), with a central estimate of 104%, close to our main assumption of 100%.

Under alternative pass-through assumptions, we maintain the scenario goal of maintaining installation quantities for high-wealth households using a per-W and per-Installation subsidy set such that the total predicted high-wealth installations are held constant. If pass-through is below 100%, then those subsidies must increase by $\frac{1-\iota}{\iota}$ where ι is the pass-through. We do so by increasing government expenditures by this excess subsidy amount, then allow installers to capture the excess subsidy as pure surplus. Since the effective subsidy remains constant, adoptions and their sizes remain constant. Installer surplus increases, and MVPF calculations for taxation are applied, subtracting from installer surplus but offsetting government expenditures. Finally, the MVPF is calculated. Note that consumer surplus, utility surplus, and avoided damages are not affected.

M.2 MVPF Results: CF1 to Baseline

We present MVPF calculations and results for a policy that moves from Counterfactual 1 (lowered export compensation rate) to baseline (NEM policy), assuming a 50% and 30% export from the grid and pass-through of 100% here.[66] Figures for Installer Surplus, Utility Surplus, and Government Expenditures are adjusted for the fiscal externality using the method above.

[66] We do not apply alternative pass-through to these calculations for three reasons. First, it is unclear on how pass-through would work for the baseline policy in which the NEM incentive, rather than a subsidy payment, is being passed. Second, California policy explicitly moved away from baseline (NEM subsidies) and returning to the baseline is precluded by statute. Finally, pass-through of approximately 100% is well-documented in the literature. We focus our results on the MVPF of CF1 to CF2.

M.2.1 MVPF Results for 50% export to grid

Table M.1: Changes in Outcomes (Millions $) and Marginal Value of Public Funds from Comparing the Two Counterfactual Scenarios. Assumes 50% export to grid and pass-through of 100%.

Wealth	Consumer Surplus	Installer Surplus	Utility Surplus	Avoided Damages	Government Expenditures	MVPF
High	$13.9	$7.2	-$16.4	$7.9	$4.9	2.57
Med	$8.2	$9.8	-$19.9	$11.0	$6.4	1.42
Low	$2.3	$4.8	-$9.6	$5.5	$3.1	0.96
All	$24.4	$21.8	-$45.8	$24.3	$14.4	1.71

M.2.2 MVPF Results for 30% export to grid

Table M.2: Changes in Outcomes (Millions $) and Marginal Value of Public Funds from Comparing the Two Counterfactual Scenarios. Assumes 30% export to grid and pass-through of 100%.

Wealth	Consumer Surplus	Installer Surplus	Utility Surplus	Avoided Damages	Government Expenditures	MVPF
High	$7.5	$2.5	-$6.2	$2.8	$1.8	3.71
Med	$3.5	$3.9	-$8.3	$4.3	$2.6	1.29
Low	$0.5	$2.0	-$4.1	$2.3	$1.3	0.52
All	$11.6	$8.4	-$18.7	$9.3	$5.7	1.87

M.3 MVPF Results: CF1 to CF2

In this Appendix section, we present MVPF calculations and results for Counterfactual 2 relative to Counterfactual 1, assuming a 50% export from the grid and pass-through rates of 78.7%, 100%, and 133.3%. Amounts for Consumer Surplus, Utility Surplus, and Avoided Damages are unchanged from the primary welfare results. Installer Surplus, Utility Surplus, and Government Expenditures are adjusted for the fiscal externality using the method above. We vary pass-through based on high and low estimates from Pless and Van Benthem (2019). Because our counterfactual subsidies applied are targeted at maintaining the number of adopters, incomplete pass-through results in increasing subsidies to an amount such that post-pass-through subsidies are unchanged. Thus, the changes in pass-through only affect installer surplus and government expenditures.

M.3.1 MVPF Results for 50% export to grid

Table M.3: Changes in Outcomes (Millions $) and Marginal Value of Public Funds from Comparing the Two Counterfactual Scenarios. Assumes 50% export to grid and pass-through of 78.7%.

Wealth	Consumer Surplus	Installer Surplus	Utility Surplus	Avoided Damages	Government Expenditures	MVPF
High	$4.1	$4.1	-$4.3	$5.4	$14.3	0.65
Med	$7.0	$11.1	-$10.3	$13.5	$21.9	0.98
Low	$5.4	$9.1	-$7.5	$10.5	$14.6	1.20
All	$16.5	$24.4	-$22.1	$29.4	$50.8	0.95

Table M.4: Changes in Outcomes (Millions $) and Marginal Value of Public Funds from Comparing the Two Counterfactual Scenarios. Assumes 50% export to grid and pass-through of 100%.

Wealth	Consumer Surplus	Installer Surplus	Utility Surplus	Avoided Damages	Government Expenditures	MVPF
High	$4.1	$1.5	-$4.3	$5.4	$11.7	0.57
Med	$7.0	$7.7	-$10.3	$13.5	$18.4	0.97
Low	$5.4	$7.0	-$7.5	$10.5	$12.5	1.23
All	$16.5	$16.2	-$22.1	$29.4	$42.6	0.94

Table M.5: Changes in Outcomes (Millions $) and Marginal Value of Public Funds from Comparing the Two Counterfactual Scenarios. Assumes 50% export to grid and pass-through of 133.3%.

Wealth	Consumer Surplus	Installer Surplus	Utility Surplus	Avoided Damages	Government Expenditures	MVPF
High	$4.1	-$0.9	-$4.3	$5.4	$8.8	0.49
Med	$7.0	$4.5	-$10.3	$13.5	$14.5	1.01
Low	$5.4	$5.0	-$7.5	$10.5	$10.0	1.34
All	$16.5	$8.7	-$22.1	$29.4	$33.3	0.97

MVPF is little changed qualitatively under alternative pass-through scenarios. In all cases, MVPF is greater than 1 for low-wealth. MVPF surpasses 1 for medium-wealth under the assumption that pass-through is 133.3%.

M.3.2 MVPF Results for 30% export to grid

We report MVPF results for the 30% export to grid assumption here, assuming pass-through of 78.7%, 100%, and 133.3%.

Table M.6: Changes in Outcomes (Millions $) and Marginal Value of Public Funds from Comparing the Two Counterfactual Scenarios. Assumes 30% export to grid and pass-through of 78.7%.

Wealth	Consumer Surplus	Installer Surplus	Utility Surplus	Avoided Damages	Government Expenditures	MVPF
High	$1.5	$1.2	-$1.5	$1.6	$5.2	0.56
Med	$2.6	$4.0	-$4.2	$4.8	$7.7	0.93
Low	$1.8	$3.2	-$3.0	$3.7	$4.8	1.19
All	$5.9	$8.4	-$8.7	$10.1	$17.7	0.89

Table M.7: Changes in Outcomes (Millions $) and Marginal Value of Public Funds from Comparing the Two Counterfactual Scenarios. Assumes 30% export to grid and pass-through of 100%.

Wealth	Consumer Surplus	Installer Surplus	Utility Surplus	Avoided Damages	Government Expenditures	MVPF
High	$1.5	$0.2	-$1.5	$1.6	$4.2	0.46
Med	$2.6	$2.8	-$4.2	$4.8	$6.5	0.92
Low	$1.8	$2.5	-$3.0	$3.7	$4.1	1.22
All	$5.9	$5.6	-$8.7	$10.1	$14.9	0.87

Table M.8: Changes in Outcomes (Millions $) and Marginal Value of Public Funds from Comparing the Two Counterfactual Scenarios. Assumes 30% export to grid and pass-through of 133.3%.

Wealth	Consumer Surplus	Installer Surplus	Utility Surplus	Avoided Damages	Government Expenditures	MVPF
High	$1.5	-$0.6	-$1.5	$1.6	$3.1	0.33
Med	$2.6	$1.7	-$4.2	$4.8	$5.2	0.94
Low	$1.8	$1.9	-$3.0	$3.7	$3.4	1.31
All	$5.9	$3.0	-$8.7	$10.1	$11.7	0.89

The qualitative results remain the same over all pass-through assumptions. Under a 30% export rate, all pass-through assumptions indicate MVPF of greater than 1 for low-wealth households, and less than 1 for all high and medium-wealth households.