Subsidies for Close Substitutes: Evidence from Residential Solar Systems

Abajian, Alexander and Pretnar, Nick

9 August 2023

Subsidies for Close Substitutes: Evidence from Residential Solar Systems

Alexander Abajian[1*] Nick Pretnar[1†]

[1] *University of California, Santa Barbara*

August 9, 2023[‡]

Abstract

Policies promoting residential solar system adoption are designed assuming the associated generation displaces retail electricity purchases on a one-for-one basis. This assumption is not innocuous; electricity from residential solar systems is unlikely to be perfectly substitutable with grid electricity. We estimate a model of U.S. residential electricity demand allowing for spatial heterogeneity and imperfect substitution between forms of electricity to quantify the implications for green energy subsidization. We find subsidies inducing one *kWh* of residential solar electricity demand displace only 0.5 *kWh* of grid consumption. As an emissions reduction policy subsidies had national abatement costs of \$332 per $MTCO_2$ in 2018.

Keywords: Residential PV systems, residential electricity demand, rebound effects, energy subsidies

JEL Classification: H23, Q42, Q48, R23

* alexander_abajian@ucsb.edu

† npretnar@ucsb.edu

‡ We thank peers in the University of California Santa Barbara macroeconomics working group along with seminar participants at the 2022 meetings of the Society for Economic Dynamics in Madison, Wisconsin, the 2022 Southern California Applied Economics Conference at the University of California San Diego, the 2023 AERE Sessions at the Eastern Economic Association, and 2023 AERE Sessions at the Midwest Economic Association International for helpful comments. We also thank Javier Birchenall, Timothy Fitzgerald, Kyle Meng, Esteban Rossi-Hansberg, and two sets of three anonymous referees for helpful comments. A previous version of this paper was circulated under the title *An Aggregate Perspective on the Geo-spatial Distribution of Residential Solar Panels*.

1 Introduction

The United States has implemented a series of tax, regulatory, and trade policies aimed at increasing the adoption of residential solar photovoltaic (PV) systems over the past decade. The Inflation Reduction Act extends the largest of these incentives, a credit equal to 30 percent of the costs of installing PV systems households can claim against their federal tax liability, out until 2034. The stated goal of promoting PV adoption is reducing residential electricity purchases from the grid in order to lower both carbon emissions and strain on local power networks during peak hours.[1] However, the extent to which induced demand actually displaces retail electricity sales is critical for evaluating the efficacy of solar PV subsidization policies. To this end, this paper examines how effective historical government support for residential PV uptake has been in reducing residential demand for electricity from the grid.

From a purely physical perspective it may seem reasonable to assume that the electricity generated by residential solar PV systems should in turn directly lower consumption from the grid. After electricity generated by PV systems is inverted and stepped up into the AC current that is used in homes, it is indistinguishable from electricity drawn from the grid. However, as an empirical question, it is unclear whether generation by residential systems displaces grid demand on a one-for-one basis. There are many reasons, such as intermittency and reliability, that may preclude direct substitution of rooftop solar for electricity drawn from the grid in all instances. Indeed, if the two sources of electricity were perfect substitutes from an economic perspective, households should solely purchase one form when faced with differentiated prices.

The limited uptake of solar PV systems by U.S. households, even after substantial price declines in the past decade, provides a strong signal this substitutability is likely imperfect. Despite subsidization and regulatory support at multiple levels of government residential solar systems generated less than 1% of all electricity consumed by households from 2010-2018, and only 2% at the

[1] The White House recently suspended tariffs on imported silicon photovoltaic cells and modules and invoked the Defense Production Act to promote domestic solar production with goals of "cutting energy costs for families, strengthening our grid, and tackling climate change" by increasing solar adoption (see here).

national level as of 2022. Residential PV system installation rates vary considerably between otherwise similar locations which would offer comparable PV performance. As siting mechanically affects system productivity, some spatial differentiation is to be expected due to variation in the flux of solar radiation per unit area (insolation) across locations. Despite this, we show there is significant dispersion in uptake rates between neighboring Census tracts, despite the fact that within a county one might expect driving factors such as insolation as well as grid and PV system prices to be similar.

These facts lead to the primary question which we explore in this paper. How effective are subsidies for residential PV systems at both inducing additional demand for solar generation and reducing demand for more emissions-intensive retail electricity from the grid? We begin with an overview of data on local electricity and PV system prices taken from various administrative sources and use it to highlight three stylized facts: (1) half of existing variation in solar uptake at the Census-tract level is unexplained by insolation, income, and other observable variation within counties, (2) there is historical dispersion within counties between the costs of electricity generated by residential PV systems and the grid electricity prices faced by residential customers, and (3), while past increases in residentially generated electricity from PV systems probably did concurrently displace some demand from the grid, this crowding out has not been one-for-one.

To rationalize these empirical findings we form a structural model of demand for electricity that allows for imperfect substitutability between electricity from residential PV systems and that drawn from the grid. The model incorporates differentiated prices for each form of electricity at the county level as well as spatial heterogeneity in preferences for the composition of electricity consumed. This lets the price dispersion we see in the data, along with any latent variation in tastes for solar electricity across locations, affect modeled demand. We estimate the model's structural parameters and perform experiments which examine how demand would differ in a counterfactual where existing residential PV subsidies are absent. The counterfactuals illustrate the effects of subsidization policies on demand for both residential solar and retail electricity for the entirety of the lower-48 U.S. states, highlighting both the aggregate effects of subsidies as well as the dispersion of these effects across locations.

Our estimates indicate that existing policies have been effective at increasing demand for residential solar. Counterfactual simulations for 2018 suggest existing subsidies increased household demand for electricity from residential PV systems by 255% relative to an alternative regime absent any government support. The aggregate shift masks fairly heterogenous effects of policy across locations. We find that the costs of inducing an additional $1kWh$ of demand for solar electricity consumption range from \$0.06 per kWh in the 10$^{\text{th}}$ percentile county to over \$0.31 per kWh at the 90$^{\text{th}}$ percentile. These locational differences in the efficacy of subsidies suggest the current uniform federal credit leads to misallocation if the goal is to maximize induced production, similar to results from studies in other countries (Lamp and Samano 2023).

At the national level subsidies lead to a solar rebound effect (Qiu, Kahn, and Xing 2019; Aydın, Brounen, and Ergün 2023), where decreases in residential PV prices cause both solar and aggregate electricity consumption to rise. We find that the effect is small, equal to roughly 0.6% of total residential electricity consumption in 2018. This rebound effect stems from the imperfect crowding out of grid electricity by newly induced solar demand. Existing subsidies increase modeled solar demand by 19.5 billion kWh, but displace only 9.8 billion kWh of demand from the grid. The implied rate of displacement, -0.50, is in line with our reduced-form estimates using state-level panel data for monthly residential electricity consumption between 2014 and 2022.

This lack of one-for-one displacement has important implications for the implied abatement costs associated with PV subsidization policies. The combination of all subsidies aimed at inducing solar demand imply an average national abatement cost of \$332 per metric ton of CO_2 ($MTCO_2$), roughly twice what these costs would be under the assumption of perfect displacement. We conclude by discussing the implications of this imperfect crowding out effect. Given trends in current policy toward subsidies in lieu of pollution quotas or taxes, our findings add to the growing body of evidence on the shortcomings of second-best approaches to mitigation.[2]

[2]This may be especially salient given the Inflation Reduction Act also provides large subsidies for the adoption of heat pumps and electric vehicles, both of which are promoted as green alternatives to existing emissions-intensive technologies (Davis 2023).

2 Background

Residential PV systems have gone from a nascent technology to a widely available option for electricity provision in the U.S. over the past 20 years. The combination of a favorable policy environment and a 75% decline in the average cost (per Watt) of PV systems since 2000 helped drive the installation of almost 2 million small-scale solar systems through 2019 (Barbose et al. 2019).[3] Beyond traditional issues of incidence and distributional consequences (Borenstein and Davis 2016; Borenstein 2017; Pless and Benthem 2019), the use of subsidies raises questions regarding the extent to which present policies are achieving the goals of increasing system uptake and reducing demand for emissions-intensive electricity from the grid. These issues will become increasingly relevant going forward given the trend in current policy toward incentivizing households to purchase green alternatives for existing goods.

In this section we first undergo a brief survey of the existing literature. We contribute to several strands including research on rebound effects, household PV system adoption, and misallocation in the environmental context. We go on to perform a decomposition exercise of residential PV uptake in 2018 that gives a spatial perspective on the distribution of residential systems. The exercise shows that a significant share of within-county variation in residential PV uptake is unexplained by insolation, income, and other observables.

2.1 Existing Literature

Our study contributes foremost to research on the rebound effect, the literature motivated by how changes in the efficiency of using energy resources as inputs affects final demand (Khazzoom 1980; Chang, Wang, and Shieh 2018). In the case of electricity, efficiency improvements lower the amount of generation required to produce a given amount of services that use electricity as an input, but may also induce additional demand due to reducing the relative prices for such services. Empirical estimates for rebound effects in electricity consumption

[3] Small-scale PV, a broader measure capturing all non-utility scale installations by both households and small commercial enterprises, comprised 4.5% of U.S. electricity generating capacity in 2018 (EIA 2020).

vary considerably due to differentiated definitions, methods, and settings (Chan and Gillingham 2015; Borenstein 2015). Papers by Deng and Newton (2017), Qiu, Kahn, and Xing (2019), and Aydın, Brounen, and Ergün (2023) are the first to use quasi-experimental methods to estimate a household solar rebound effect — the extent to which residential PV generation increases electricity consumption. They exploit intraday fluctuations in electricity generation from residential PV systems and find households increase their combined consumption of solar- and grid-drawn electricity on days when household generation is especially high. We contribute to this literature in estimating a solar rebound effect using a structural model that accounts for how the income, price, and substitution effects of subsidies affect residential electricity consumption. This allows us to examine how the spatial variation in subsidies and insolation impacts both local and aggregate rebound effects for the entire United States.

A second literature has focused on the extensive margin of residential solar PV system adoption. Rich household-level data have allowed for the estimation of granular discrete choice models of household installation decisions in the U.S. and international settings (De Groote and Verboven 2019; Gillingham and Tsvetanov 2019; Langer and Lemoine 2022; Feger, Pavanini, and Radulescu 2022). These studies illustrate a range of findings particular to household PV uptake, notably the myopic behavior of households with regard to future payments from renewable energy generation (De Groote and Verboven 2019) and the importance of accounting for forward-looking behavior (Langer and Lemoine 2022) and tariff structure (Feger, Pavanini, and Radulescu 2022) when designing policy incentives to increase PV adoption. While our model of demand is less granular than those above, it allows us to focus on how households' substitution between PV systems and grid electricity affects their total electricity consumption. We can directly examine how existing policies have shaped aggregate electricity demand rather than focusing on discrete system adoption decisions in specific regions.

Finally, we contribute to a literature examining the spatial heterogeneity in the social costs of renewable energy adoption. Benefits from PV adoption may accrue through the channel of pollution mitigation from displaced grid demand and lower carbon emissions, as well as by relieving grid congestion (Sexton et al. 2021; Lamp and Samano 2023; Dauwalter and Harris 2023). Callaway, Fowlie, and McCormick (2018) show in the U.S. that the displaced emissions of new renewable energy deployment vary substantially across physical location. Sexton et al. (2021) and Dauwalter and Harris (2023) extend this analysis to local pollutants and show that the environmental benefits from new PV capacity are highly dependent on location and the displaced generation. The above studies combine to illustrate how the uniformity of existing subsidies at the federal level, along with the agnostic treatment of spatial variation in the benefits from avoided generation, is in general inefficient. In a related study Lamp and Samano (2023) show for Germany that gains in social welfare may also be possible if existing PV capacity could be moved around to better reflect the spatially-differentiated value of electricity.

Like these papers we focus on the efficacy of existing subsidies in the context of induced demand and to a lesser extent carbon abatement across space. However, to the best of our knowledge, we are the first to explore how the extent that residential generation displaces grid demand determines the implications of solar PV subsidies for aggregate residential electricity consumption. We show that the imperfect crowding out of grid demand by new PV systems affects both the spatial distribution of abatement and aggregate abatement overall. Our pass-through estimates from counterfactuals, supported by evidence from time series data, suggest that assumptions of one-for-one grid displacement by new residential PV demand may be overly optimistic.

2.2 U.S. Residential Solar

The most granular comprehensive data on residential solar panel uptake across the entire continental United States come from Stanford University's *DeepSolar* Project (Yu et al. 2018).[4] *DeepSolar* provides a cross section of the universe of residential solar installations within the contiguous United States at Census-tract levels as of December 2018. Figure 1 shows a choropleth map of county-level average daily solar insolation in native units — kilowatt-hours (kWh) per square meter per day ($kWh/m^2/d$). This measure captures the maximum amount of en-

[4]The *DeepSolar* Project utilizes machine learning techniques to analyze satellite imagery and infer the location and surface area of small-scale solar panel installations across the lower-48 U.S. states. The data are currently available as a static snapshot as of December 2018. This dataset links solar panel location and surface area with average daily solar insolation, measured in kilowatt-hours per square meter per day ($kWh/m^2/d$), at the Census-tract level.

ergy a square meter of PV cells could produce over the course of the average day in a given location.

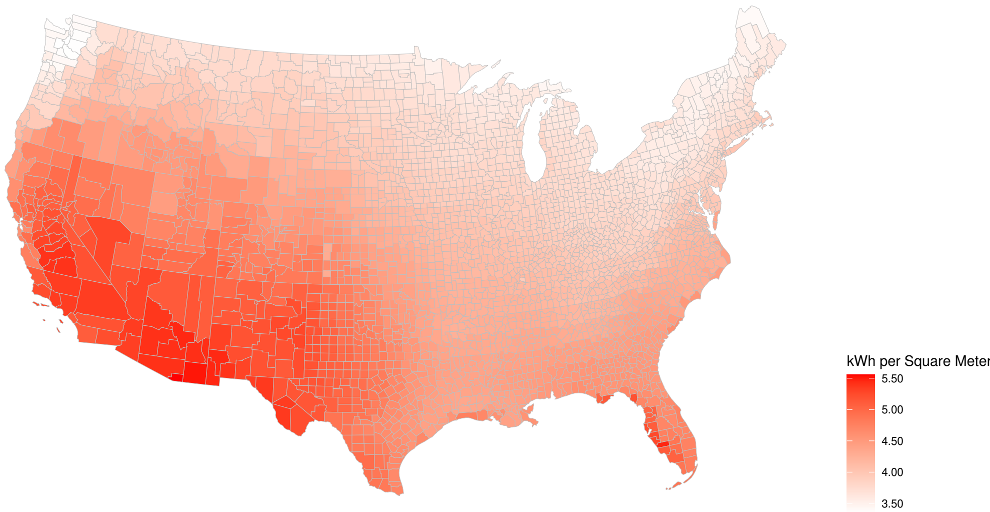

Figure 1: County-level average daily solar insolation ($kWh/m^2/d$) from the December 2018 *DeepSolar* data (Yu et al. 2018).

Figure 2 displays the extent of dispersion in uptake of residential solar installations. The color gradient in this figure is increasing in the percentages of households at the county level that have installed solar capacity (as of December 2018). While solar panel uptake is highly concentrated in the Southwest, wherein some counties have uptake rates of over 10%, household installations are prevalent in a few Northern coastal areas as well. Despite the concentration of installations in urban areas of the Pacific Northwest and Northeastern U.S., uptake rates appear to follow the insolation gradient.

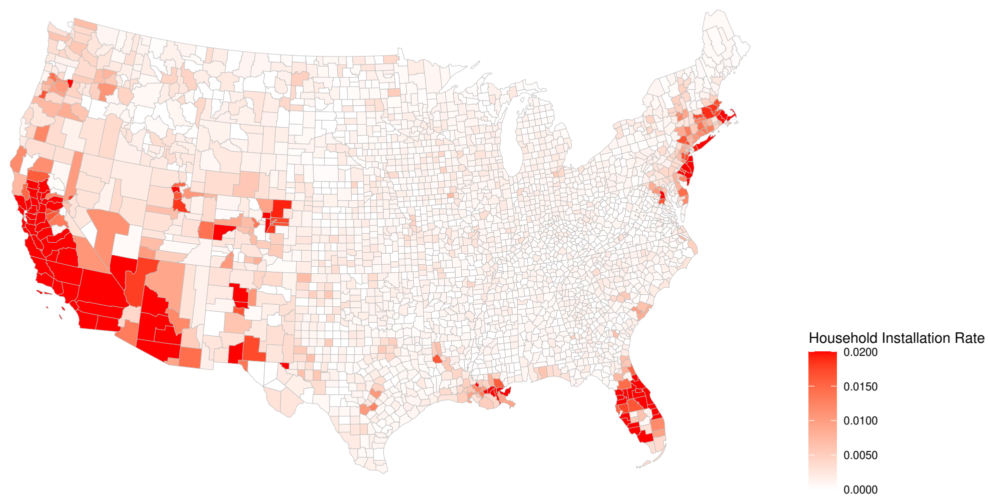

Figure 2: Share of households in a given county with residential solar installations from *DeepSolar* (Yu et al. 2018).

2.3 Decomposing Solar Panel Uptake

A sort of visual regression of Figure 2 on Figure 1 would suggest insolation does explain a large portion of uptake. To formally test this we begin with a reduced-form specification for household demand for PV systems at the Census-tract level. We regress the average panel area per household in Census tract k of county c, PA_{kc}, on income and insolation along with county-level fixed effects, ω_c:

$$PA_{kc} = \alpha_1 \ \ln(income_{kc}) + \alpha_2 \ insolation_{kc} + \omega_c + \epsilon_{kc} \tag{1}$$

Results from variations of the specification in equation (1) are shown in Table 1. Column (1) formalizes the heuristic test available from comparing local insolation in Figure 1 to installation rates in Figure 2; insolation is positively correlated with uptake and explains roughly one-fifth of tract-level variation in adoption. Column (2) confirms existing results that higher-income locations also exhibit greater panel uptake (Borenstein and Davis 2016). Adding county-level fixed effects ω_c along with insolation and income explain half of observed variation in column (5). The Census-tract level granularity is particularly helpful as it allows for decomposition within counties while fixed effects control for any variation in prices and policies across counties.[5]

Table 1: Decomposition of Solar Panel Area per Household

	(1)	(2)	(3)	(4)	(5)
Insolation	0.566	-	0.576	-	0.621
	(0.131)		(0.150)		(0.513)
Log Income ($1000s)	-	0.391	0.413	-	0.368
		(0.183)	(0.152)		(0.188)
County Fixed Effects	No	No	No	Yes	Yes
R^2	0.18	0.05	0.24	0.45	0.49
N	71,651	71,651	71,651	71,651	71,651

Standard errors clustered at the state level are shown in parentheses. Observations are weighted by the number of households in the Census-tract using Stata's *aweights* command. Log income is the Census-tract average reported in the ACS 5-year survey from 2015, measured in thousands of dollars.

Adding a large vector of local climatological and demographic controls to the regression in (5) attenuates the magnitude of coefficients slightly but does not change the qualitative result or model fit.[6] If it is the case where local differences

[5]Variation in the regulatory structure governing utilities and generators, state renewable portfolio standards, and the policy of individual load-serving entities are just some of the factors affecting local electricity markets. While there is substantial heterogeneity in subsidization policies and regulation of the electric grids at the county- and state-levels, we believe the intra-county variation is likely much more limited.

[6]Adding a vector of demographic controls comprised of black and white race shares, the rate of residents with less than a high school education, population density, the rate of households with mortgages, GOP voting percentages in 2016, total land area, median home values and rental payments, and a diversity index and climatological controls, including local heating and cooling degree days, coordinate grid centroid, elevation, heating and cooling design temperatures, frost days, average air temperature, relative humidity, atmospheric pressure, and wind speed raises the R^2 in column (5) to 0.54 while the coefficients on insolation and income remain roughly unchanged.

in system prices and grid electricity are captured by county-level fixed effects, it is somewhat surprising that these level effects, along with variation in insolation and income, explain only half of system uptake in column (5). To attempt to rationalize the large share of uptake not attributable to observables, we form a structural model for residential PV system demand that allows for latent characteristics to contribute for the unexplained variation in our decomposition exercise thus far.

3 Model

This section lays out the theoretical framework we use to characterize U.S. demand for residential PV systems relative to grid electricity and non-electrical consumption. A household i is utility-maximizing and indexed by its county of residence c in year t. Households are price-takers with exogenous income each period. Each household has preferences over an electricity composite good, q^e_{ict}, and an outside commodity, $\bar{c}_{ict}$, which captures all non-electricity consumption. The electricity composite combines electrical consumption from the grid, g_{ict}, with consumption of residentially generated solar electricity, s_{ict}, using a quasi-constant elasticity of substitution (Q-CES) aggregator that allows for non-homotheticity in demand for each form of electricity. The electricity composite is

$$q^e_{ict} = \left(\gamma_j(g_{ict} - \underline{g})^{\frac{\rho-1}{\rho}} + (1-\gamma_j)(s_{ict} - \underline{s})^{\frac{\rho-1}{\rho}}\right)^{\frac{\rho}{\rho-1}}, \quad \rho \geq 0 \quad \text{and} \quad \gamma_j \in (0,1) \tag{2}$$

Where j indexes the state of residence for agent i. The parameters γ_j and ρ respectively govern the weight households put on each form of electricity and the substitutability between the two forms. As $\rho \to \infty$ solar and grid electricity are perfect substitutes. As $\rho \to 0$ the aggregator is Leontief. We allow the weight of grid electricity in the composite, γ_j, to vary across states. Households in states with lower γ_j values will prefer an electricity composite that has a higher share of solar electricity. The degree of non-homotheticity in grid and solar demand is parameterized by a pair of reference levels, $\underline{s}$ and $\underline{g}$.

Households have preferences over the electricity composite, q^e_{ict}, and outside consumption, $\bar{c}_{ict}$, according to a CES utility function:

$$u_j(q^e_{ict}, \bar{c}_{ict}) = \left(\delta_j \left(q^e_{ict}\right)^{\frac{\kappa-1}{\kappa}} + (1-\delta_j) \left(\bar{c}_{ict}\right)^{\frac{\kappa-1}{\kappa}} \right)^{\frac{\kappa}{\kappa-1}}, \quad \kappa \geq 0 \quad \text{and} \quad \delta_j \in (0,1) \tag{3}$$

The parameters δ_j and κ govern the weight and substitutability households assign to the electricity composite and outside consumption. Like above, a larger δ_j will lead to households preferring a higher share of total expenditure dedicated to electricity. We assume that at the state level, all counties in state j have the same weight parameters, γ_j and δ_j. The parameters governing substitutability, κ and ρ, along with reference levels, $\underline{s}$ and $\underline{g}$, are identical for all counties and states.

Within a county households are assumed to be heterogeneous only in exogenous income, y_{ict}. The household's period-t budget constraint is

$$\bar{c}_{ict} + p^g_{ct}\, g_{ict} + p^s_{ct}\, s_{ict} = y_{ict} \tag{4}$$

Households within a county face identical and constant flow prices for solar and grid electricity each year. At the national level households face a uniform price for the outside good, $\bar{c}_{ict}$, which serves as the numéraire. p^s_{ct} is the local average price in county c of period t a household faces in order to consume one kWh of residential-solar electricity.[7] Meanwhile, p^g_{ct} is the local average price a household

[7] As we lack access to household-level data combining panel installations, income, and electricity consumption, uniform market prices (including p^s_{ct}) across households within a county are necessary for estimating demand at the county level using available data. The least innocuous aspect of this price uniformity assumption is that the marginal price of one unit of solar electricity is, on average, the same for all households within a county in a given year. A thorough scrutiny of empirical evidence backing up this assumption is available in Technical Appendix C. In short we find that the portion of system-level variation in levelized costs of energy (LCOE) explained by system capacity is small (less than 10%) relative to the portion of cost that is either constant within locations or explained by other factors (e.g., year of installation, module efficiency, or inverter type). We will discuss in detail how solar prices are constructed in Section 4.

pays for one kWh of electricity consumed from the grid.[8]

Our model treats solar electricity consumption as a continuous choice rather than a discrete choice over adoption choice. We elect to abstract from the discrete aspect of solar adoption because our primary focus is how subsidization has affected total residential electricity consumption in a given location rather than the number of adoptees. The solar flow cost (i.e., the price of one kWh of residential-solar-generated electricity) embeds upfront panel installation costs, regional net-metering policies, regional solar subsidization and taxation policies, and system-level productivity which will vary both over time, due to technological change, and by region, due to insolation and climatological differences.[9] An accurate accounting of the effects of such subsidies will naturally consider how policies impact aggregate consumption flows from different electricity sources.

Lemma 1. *The household's Marshallian demand system features Engel curves that are linear in income. The model admits a county-level representative consumer.*

All proofs are relegated to Technical Appendix A. Assuming prices of grid and solar electricity are constant across all households within a county ensures we can examine the local average demand system for a county-level representative consumer. Households maximize (3) composed with (2) subject to (4). Along with uniform pricing within a county, this nested Q-CES preference structure implies that the demand generated by aggregating across all households within a county coincides with that of a utility-maximizing representative agent. Lemma 1 states that the Marshallian demand functions generated by the Q-CES structure satisfy the conditions for aggregation outlined in Lewbel (1989). From here forward all

[8]Data limitations preclude us from imposing a variable marginal price schedule. As prior structural models in the literature abstract from block pricing, we view the assumption that prices for electricity from the grid are constant across households within counties as fairly benign (De Groote and Verboven 2019; Feger, Pavanini, and Radulescu 2022). While, in practice, instantaneous electricity prices faced by households may vary with intensity and time of utilization in some locations, this variation in marginal pricing is substantially smaller than gaps we observe between solar and grid prices in much of our sample. Borenstein and Bushnell (2022) find that while 58% of residential customers face varying marginal electricity prices, the absolute difference between minimal and maximal marginal prices is less than \$0.02 per kWh on average, between 10-20% of average costs. There is also evidence that household electricity demand is responsive to changes in average rather than marginal prices (Ito 2014).

[9]We describe our process for constructing these prices in the estimation section.

theoretical and empirical analyses are performed on the conditions for optimal allocations of a county-level representative agent.

3.1 Comparative Statics

Our model is useful for examining how subsidies that lower the unit price of residential solar electricity, p^s_{ct}, impact residential solar demand. In this section we establish conditions under which rebound and backfire effects occurs in our model. In our setting a rebound effect occurs when declining composite electricity prices, due to the falling opportunity cost of residential solar, lead to an overall increase in total electrical consumption.[10] Although related, our definition of the rebound effect in the model does not strictly map one-for-one into changes in the real quantity of electricity consumed (as measured in kWh), but rather an increase in demand for the composite, q^e_{ct}, attributable to a decline in the optimal price index, p^e_{ct} via falling p^s_{ct}.[11] As s_{ct} is a normal good, its demand will rise as p^s_{ct} falls. However, the degree to which declining solar prices will affect demand for the composite electricity good and grid electricity is less immediately apparent.

Definition 1. *A subsidization policy **backfires** if households' consumption of grid electricity rises as solar prices fall.*

Definition 1 makes explicit a potential perverse outcome from subsidization. Note that while all policies which backfire due so because of rebound effects, rebound effects are not necessarily associated with backfiring policies (Chan and Gillingham 2015). Under Definition 1, even if both g_{ct} and s_{ct} rise in response to a decline in p^s_{ct} then the policy has backfired as grid consumption has risen. We now characterize the conditions which determine whether subsidization policies will cause rebounding and backfiring effects.

Proposition 1. *Let p^e_{ct} and $\widetilde{p}_{ct}$ be optimal price indices for the electrical and aggregate consumption composite goods. Let $\varepsilon^{e,s}_{ct}$ be the elasticity of p^e_{ct} with respect to p^s_{ct}, and, similarly, let $\widetilde{\varepsilon}^s_{ct}$ be the elasticity of $\widetilde{p}_{ct}$ with respect to p^s_{ct}. Demand for the electricity*

[10] Our notion, while particular to our model, is similar to a more general phenomenon characterized in the literature (Chan and Gillingham 2015; Kulmer and Seebauer 2019; Qiu, Kahn, and Xing 2019).

[11] Note that we are dropping household-level indices, i, as we are now operating in an environment characterized by county-level representative consumers.

composite is strictly decreasing in the price of solar, leading to ***rebound*** *effects, if and only if*

$$\kappa(\widetilde{\varepsilon}^{s}_{ct} - \varepsilon^{e,s}_{ct}) - \widetilde{\varepsilon}^{s}_{ct} < \frac{p^{s}_{ct}\,\underline{s}}{\widetilde{y}_{ct}} \tag{5}$$

where $\widetilde{y}_{ct} = y_{ct} - p^{s}_{ct}\,\underline{s} - p^{g}_{ct}\,\underline{g}$*. This condition is always satisfied if* $\underline{s} \geq 0$*. When* $\underline{s} < 0$*, the existence of rebound effects depends on the sign of* (5)*.*

Proposition 1 establishes that under the Q-CES structure, a decline in the solar price will lead to rebound effects if conditions on κ, $\underline{s}$, and the price-index elasticities with respect to p^{s}_{ct} are met.[12] If households are faced with a non-negative reference level parameter, $\underline{s}$, rebound effects are guaranteed to occur for all combinations of strictly positive solar and grid prices.[13] When $\underline{s} < 0$, however, rebounding effects are not guaranteed unless the price-index elasticities satisfy (5).

Proposition 2. *Subsidization policies reducing* p^{s}_{ct} *will* ***backfire****, causing grid electricity demand to rise, if any of the following conditions hold:*

1. *The structural parameters satisfy* $\underline{s} \geq 0$*, and* $0 < \rho \leq \kappa < 1$*.*

2. $\underline{s} = 0$*, and* ρ*,* κ*, and price-index elasticities satisfy*

$$\frac{\rho - \kappa}{1 - \kappa} < \frac{\widetilde{\varepsilon}^{s}_{ct}}{\varepsilon^{e,s}_{ct}} \quad \text{if} \quad \kappa < 1 \quad \text{or} \quad \frac{\rho - \kappa}{1 - \kappa} > \frac{\widetilde{\varepsilon}^{s}_{ct}}{\varepsilon^{e,s}_{ct}} \quad \text{if} \quad \kappa > 1 \tag{6}$$

3. $\underline{s}$*, and* ρ*,* κ*, and price-index elasticities satisfy*

$$(\rho - \kappa)\varepsilon^{e,s}_{ct} - (1 - \kappa)\widetilde{\varepsilon}^{s}_{ct} < \frac{p^{s}_{ct}\,\underline{s}}{\widetilde{y}_{ct}} \tag{7}$$

Proposition 2 guarantees that backfiring effects will occur independent of local variation in price-index elasticities *only* if condition 1 holds. Note that while this condition is sufficient, it is not necessary for backfiring effects to occur. Indeed,

[12] Note that the price indices, p^{e}_{ct} and $\widetilde{p}_{ct}$ are defined in Lemmas A1 and A2 of Technical Appendix A.

[13] When we estimate the model, indeed we estimate that $\widehat{\underline{s}} > 0$, predicting that residential-solar subsidization policies in all geographic locales within the U.S. lead to rebound effects.

condition 2 says that $\underline{s}$ can be zero with $\kappa < \rho < 1$ and backfiring effects can still be observed as long as the relative price-index elasticities satisfy the left-hand condition in (6). A similar requirement occurs for $\kappa > 1$ at $\underline{s} = 0$. Clearly, it must be that $\rho < \kappa$ for $\frac{\rho-\kappa}{1-\kappa}$ to be positive if $\kappa > 1$, but the relative price-index elasticities will still dictate whether backfiring effects occur. Condition 3 says that even when $\underline{s}$, ρ, κ, and price-index elasticities do not necessarily satisfy the more narrow sufficiency conditions described in 1 and 2, backfiring effects can still be observed. In fact, they are increasingly likely to be observed, regardless of the values of ρ, κ, and price-index elasticities, as income effects get stronger. There are, however, still conditions when $\underline{s}$ is finite or even negative under which backfiring effects will be observed: we simply require that (7) holds.

4 Estimating the Model

We estimate demand parameters from the model presented in Section 3 to test whether Propositions 1 and 2 hold in the data. We begin with presenting our main estimating equations and briefly outlining our data sources, namely how we compute the price of residential solar electricity. We then estimate the structural parameters on a set of 601 counties between 2010 and 2018 for which we have sufficient data. After examining the in-sample model fit, we extrapolate our estimates to the full set of counties in the contiguous United States. The estimated parameters are then used for simulations in Section 5.

4.1 Estimating Equations

We use generalized method of moments (GMM) to estimate the structural demand parameters governing the model. Our estimating equations are comprised of the representative agents' structural demand functions for each commodity:

$$g_{ct} = \underline{g} + \left(\frac{\gamma_j}{p^g_{ct}}\right)^{\rho} p^e_{ct}(p^g_{ct}, p^s_{ct})^{\rho-1} X^e_{ct}(g_{ct}, s_{ct}, p^g_{ct}, p^s_{ct}) + \eta^g_{ct} \tag{8}$$

$$s_{ct} = \underline{s} + \left(\frac{1-\gamma_j}{p^s_{ct}}\right)^{\rho} p^e_{ct}(p^g_{ct}, p^s_{ct})^{\rho-1} X^e_{ct}(g_{ct}, s_{ct}, p^g_{ct}, p^s_{ct}) + \eta^s_{ct} \tag{9}$$

$$\bar{c}_{ct} = (1-\delta_j)^{\kappa} \widetilde{p}_{ct}(p^g_{ct}, p^s_{ct})^{\kappa-1} \left(y_{ct} - p^s_{ct}\,\underline{s} - p^g_{ct}\,\underline{g}\right) + \eta^{\bar{c}}_{ct} \tag{10}$$

For each county/year pair the targeted moments are given explicitly by equations (8), (9), and (10) below. Demand is decomposed into theoretical predictions based on observed income and prices along with an error term, $\boldsymbol{\eta}_{ct}$.[14,15] X^e_{ct} is county-level optimal electricity expenditure, which can be written as a function of demand for grid and solar, prices, and reference parameters, as demonstrated in Lemma A1 of Technical Appendix A. Because X^e_{ct} is a function of left-hand side variables g_{ct} and s_{ct}, the demand system given by (8), (9), and (10) is implicit.[16] The term y_{ct} in (10) is average household income for a county/year pair taken from the American Community Survey (ACS). The sample we use for estimation contains 601 counties covering 16 states over the 2010-2018 period. Annual averages for residential grid electricity prices and consumption at the county-level come from the EIA.[17,18]

4.2 Residential PV Electricity Pricing and Consumption

Unlike with electricity purchased from the grid, flow prices and quantities for households' consumption of residential PV-generated electricity are not readily observable. In this section we describe how we measure average residential PV generation and flow prices at the system level. We use these system-level measurements to form averages at the county/year frequency for solar consumption,

[14] See Lewbel (2001).

[15] We use bold-face font to denote vectors. In this case $\boldsymbol{\eta}_{ct}$ is a three-dimensional vector comprising the error terms in the GMM-targeted demand system for a given county/year pair (i.e., the difference between observed demand and the values that would be generated by the true set of population parameters governing agents' maximization problems in a correctly specified model). We denote demand-specific errors with superscripts, so that, for example, η^g_{ct} is the error for county/year pair c and t in (8).

[16] When minimizing the GMM objective function, we pass data for g_{ct}, s_{ct}, p^g_{ct}, and p^s_{ct} to the expression for electrical expenditure, while the minimizing routine updates the structural parameters $\underline{g}$ and $\underline{s}$ at each iteration.

[17] The EIA's form 861 collects annual data on the universe of U.S. utility companies' electricity sales from 1990 onward. Respondents submit data on total electricity deliveries, total revenue, and total number of customers by state and end-use sector. These data provide total residential electricity expenditure, consumption, and customers served for each state/utility pair. We then aggregate all consumption and customers served at the county level using the service territory data provided in form EIA-861 to construct average prices and consumption and use the EIA's utility/county crosswalk to map these data to the counties in our sample.

[18] Non-electrical consumption, $\bar{c}_{ct}$, is computed by subtracting both forms of electricity expenditure from county-level average household income taken from the American Community Survey estimates for each year.

s_{ct}, and average solar prices, p^s_{ct}.

Data for constructing prices and quantities are drawn from the Lawrence Berkley National Lab's (LNBL) *Tracking the Sun* (TTS) dataset, which catalogues the majority of small-scale PV installations between 1998 and 2018 (Barbose et al. 2019). Each observation contains the size of the system (in Watts), the location of the installation, and the date the system was installed. We use the LBNL's *PVWatts* API to estimate annual solar generation at the system level for the approximately 700,000 residential installations in the TTS dataset for which we have sufficient location data. Then, for each county/year in which we observe installations, we sum our total measurement for generation across all systems installed either prior to or in that year. We take this total and divide it by annual household counts from the ACS 5-year dataset to form our panel of average household solar consumption at the county/year level.

On a unit basis measuring flow prices, p^s_{ct}, directly in the data is infeasible. The TTS contains detailed data on system-level installation costs as well as local rebates, taxes and transfers, and other pecuniary incentives. Moving forward, we assume that households price electricity generated by PV systems using a user-cost approach. This approach takes into account various factors, such as the fraction of total system costs attributed to each unit of electricity produced, along with taxes, local subsidies, and fixed installation costs. In addition, this approach factors in the ongoing expenses associated with maintenance and depreciation of solar panels. By implementing this approach, the cost (or price) of generating one *kWh* of electricity from a household PV system in county c during period t will be expressed in the same units as the flow price of grid electricity. This ensures consistency in measuring the cost of electricity regardless of the source, allowing for accurate comparisons between the two.[19] We construct this PV user cost using a levelized cost of energy (LCOE) approach common in the engineering literature for each system-level observation in the TTS data.[20] From these residential system-level prices, p^s_{ict}, we then construct size-weighted average prices for all

[19] In the data we observe both households that purchase their own systems and those that buy electricity from a third party-owned (TPO) system. System size and installation costs are similar across both forms of ownership. As electricity generated by TPO systems must be priced so as to recoup the user costs for their owners, we find it reasonable to assume that households value electricity generated by household-owned systems in a similar manner.

[20] See Flowers et al. (2016) for a review of LCOE measures for PV systems.

county/year pairs to get the values of p^s_{ct} used in our primary quantitative analyses. Technical Appendix B.1 provides a complete description of how we construct our estimates for average county/year residential solar generation, s_{ct}. Technical Appendix B.2 provides a detailed description of the system-level LCOE measure we use to get prices, as well as our county/year aggregation procedure.

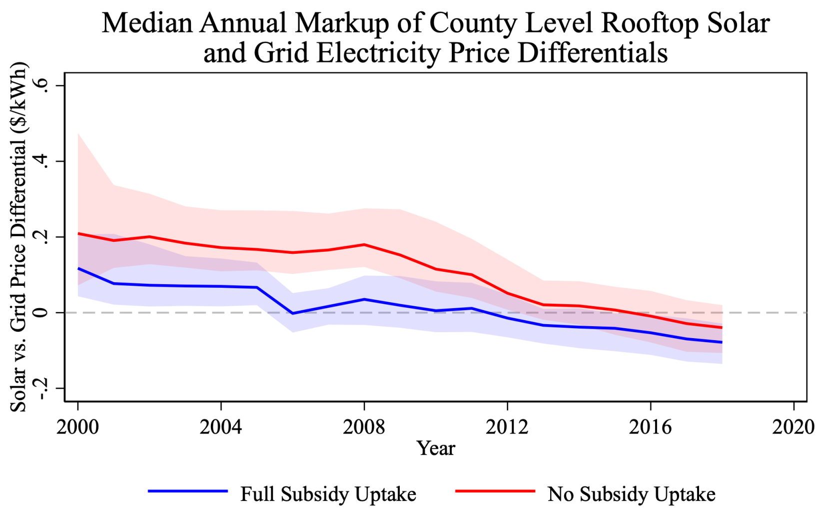

Note: Solid lines denote the annual median difference at the system level between the LCOE we estimate and the grid electricity price ($/kwh) in that county. The dashed grey line indicates the breakeven point where the prices are equal. Shaded regions denote the 10% and 90% percentiles in each year.

Figure 3: Time series of median county-level price wedges, $p^s_{ict} - p^g_{ct}$.

Figure 3 shows how system-level solar markups over local retail electricity prices have changed since 2000. This markup is just $p^s_{ict} - p^g_{ct}$ where $p^s_{ict} \equiv LCOE_{ict}$. A markup value above the dashed gray line indicates a given system-level solar price in its year of installation exceeded that of electricity drawn from the grid. We present two time series, each featuring different assumptions regarding the degree to which households took advantage of available tax credits and other solar-installation incentives during the period of installation.[21] The blue line contains a time series of the median system-level $LCOE_{ict}$ each year constructed under the assumption that households take full advantage of all available subsidies, including federal investment tax credits (ITC), which when con-

[21] In Technical Appendix B.2.2 we engage in additional comparisons of after-subsidy solar prices under alternative subsidization assumptions.

sidered together offset well-over 30% of aggregate residential PV system costs over our sample period.[22,23] The red line, on the other hand, features the time series of median $LCOE_{ict}$ constructed under the assumption that no subsidies, either federal ITC or local, are taken up by households.

Through 2010, we estimate that solar prices exceeded those from grid electricity for both the median and average systems even under maximal assumptions concerning subsidization. Prices excluding subsidies for the median system did not decline well below grid prices until 2016, and even in 2018 many systems provided electricity at a cost well-above that drawn from the grid. This large variation in prices over space and time provides us the variation necessary for identification of structural parameters when estimating the model.

4.3 Identification

Our identification strategy is based on several primary assumptions. First, as solar and grid prices may be endogenous we require a set of valid instruments on which to project the residuals (De Groote and Verboven 2019; Gillingham and Tsvetanov 2019; Lyu 2023). Second, we assume that within each state the county-level econometric errors are mean zero. This amounts to assuming that if the utility-maximization problem of the representative agent in county c corresponds to the actual data generating process, we expect the average representative agent within each state to also be solving the utility-maximization problem we have specified such that these errors are zero.

Let z_{ct} be a period-t vector of instrumental variables for county c. For a given county/year pair this object is three-dimensional and contains the following instruments: 1) a constant term, 2) the national average price per Watt for U.S.

[22] The ITC is effectively a credit worth 30% of total system costs that can be claimed against federal taxes for any residential solar system installed between 2006 and 2019. The Internal Revenue Service (IRS) line-item data suggest households claimed $2.25 billion in federal solar investment tax credits for fiscal year 2018 for the full universe of payees. Our maximal estimate for credits that could be claimed against installations sufficient data in our sample was $1.65 billion for 2018, supporting the assumption that most households successfully claimed the credit. Similar levels of aggregate claims relative to installation costs are present in other tax years (IRS 2019).

[23] The assumption that all households are able to fully monetize the ITC may overstate its effects on costs (Borenstein 2017; Pless and Benthem 2019) as some households which purchased PV systems may not have sufficient tax liabilities against which to claim the credit. As such we construct both prices as a bounding exercise.

photovoltaic cell and module imports that year interacted with the longitudinal centroid of county c, and 3) the national average price of retail electricity across all sectors. We assume the annual average price of PV cell and module imports is correlated with county-level flow prices, p^s_{ct}, as these are key hardware components that determine system costs, but orthogonal to model errors in solar demand, η^s_{ct}.[24] Interacting this with county-level longitude exploits the fact that almost all silicon photovoltaic cells and modules used in U.S. systems are imported from East Asian countries, leading to freight costs for PV imports to be increasing in longitude.[25] Further, the national average grid electricity price is strongly correlated with local grid prices, p^g_{ct}, but unlikely to violate exclusion restrictions given each county's residential demand is a very small fraction of total, national electricity consumption across all sectors.

Note that, for the next couple of paragraphs only, we denote the county indices' dependence on state j via notation $c(j)$. Our within-state identifying assumption is as follows:

$$\mathbb{E}\Big[\boldsymbol{\eta}_{c(j),t} \otimes \boldsymbol{z}_{ct} \;\Big|\; c(j) \in j\Big] = \mathbf{0}_{9\times 1} \qquad \forall j \tag{11}$$

where $\otimes$ is the Kronecker product. The logic behind this identifying assumption is straightforward: if the model is correctly specified, agents' optimization should result in the residual terms being mean-zero in expectation regardless of any geographical conditioning. This generates a sufficient number of moment conditions to identify the state-level preference weights γ_j and δ_j.

Now, let $n \in \{g, s, \bar{c}\}$ index the moment conditions associated with implicit demand conditions in (8), (9), and (10), respectively, for each county/year vector of residuals. Let $\boldsymbol{\eta}^n_{jt}$ be a column vector of the n^{th} component of each of $\boldsymbol{\eta}_{c(j),t}$ for a fixed state j in year t. Let $\mathbb{C}_{jt}$ be the number of counties in state j in period t that are included in our sample, let $\mathbb{C}_j = \sum_t \mathbb{C}_{jt}$ be the total number of county/year level observations within state j, and let $\mathbb{C} = \sum_j \sum_t \mathbb{C}_{jt}$ be the total number of

[24] This is reasonable given that residential systems comprise a small share of U.S. end use of PV cells and modules and an even smaller share of the global market; local residential demand in a single U.S. county should have negligible effects on national import prices.

[25] Imports have accounted for about 90% of all U.S. solar PV shipments between 2013 and 2019, among which the largest partners were Vietnam, Malaysia, South Korea, and Thailand. Prior to this period the Chinese share of U.S. imports was high enough to motivate the U.S. to bring an anti-dumping dispute to the WTO.

county/year observations in our dataset.[26] In a slight abuse of notation, let $\boldsymbol{\eta}_j^n$ be the $\mathbb{C}_j \times 1$ dimensional column vector of within-state demand residuals for either grid, solar, or non-electrical consumption across all county/year observations and $\boldsymbol{z}_j$ be the stacked column vector of instruments $\boldsymbol{z}_{c(j)t}$ over our entire sample period for each county in state j. Letting $\boldsymbol{\Theta} = \left(\rho, \kappa, \underline{s}, \underline{g}, \boldsymbol{\gamma}^\top, \boldsymbol{\delta}^\top\right)$ be the column vector comprised of the 36 structural parameters constituting the demand system, define $\boldsymbol{\mu}_j^n$ as the 3-dimensional vector of state-level moments given by:[27]

$$\boldsymbol{\mu}_j^n\left(\boldsymbol{\eta}_j^n, \boldsymbol{z}_j\,;\,\boldsymbol{\Theta}\right) = \frac{1}{\mathbb{C}_j} \sum_t \sum_{c(j)} z_{ct} \cdot \eta_{c(j),t}^n \tag{12}$$

The vector $\boldsymbol{\mu}_j^n\left(\boldsymbol{\eta}_j^n, \boldsymbol{z}_j\,;\,\boldsymbol{\Theta}\right)$ gives the sample analogue of state-j population moments in equation (11) for good n. We stack $\boldsymbol{\mu}_j^g$, $\boldsymbol{\mu}_j^s$, and $\boldsymbol{\mu}_j^{\bar{c}}$ into a 9-dimensional vector $\boldsymbol{\mu}_j$ (without the superscript), which contains all of the orthogonally projected moments for state j. Let $\boldsymbol{\mu}$ be the 144-element column vector stacking the $\boldsymbol{\mu}_j$ from each of the 16 states in our sample. This brings us to our primary exclusion restriction, which is as follows:

$$\mathbb{E}\Big[\boldsymbol{\mu}\left(\boldsymbol{\eta}, \boldsymbol{z}\,;\,\boldsymbol{\Theta}\right)\Big] = \mathbf{0}_{144\times 1} \tag{13}$$

which takes $\boldsymbol{\eta}$, the stack of $\boldsymbol{\eta}_j^g$, $\boldsymbol{\eta}_j^s$, and $\boldsymbol{\eta}_j^{\bar{c}}$ for each state j, the set of instruments $\boldsymbol{z}$ which stacks $\boldsymbol{z}_j$ from each state, and structural parameters $\boldsymbol{\Theta}$ as arguments.

4.4 GMM Estimator and Structural Parameter Estimates

Our GMM estimator operates on the identifying assumption in (13). We use Stata's two-step GMM routine where the initial weight matrix is the 144×144 identity matrix. In the second-stage we compute a robust weight matrix, allowing for heteroskedasticity in errors across each moment condition. Let $\widehat{\boldsymbol{W}}$ be the

[26] Note that our panel is unbalanced; a few counties in our dataset do not have sufficient data in the early years of our sample period to be included.

[27] This object is 3-dimensional because we use 3 instrumental variables.

estimated second-stage, robust optimal-weighting matrix. The estimated vector of structural parameters solves

$$\widehat{\boldsymbol{\Theta}} = \underset{\boldsymbol{\Theta}}{\operatorname{argmin}}\, \boldsymbol{\mu}\big(\boldsymbol{\eta}, \boldsymbol{z}\,;\, \boldsymbol{\Theta}\big)^{\top}\, \widehat{\boldsymbol{W}}\, \boldsymbol{\mu}\big(\boldsymbol{\eta}, \boldsymbol{z}\,;\, \boldsymbol{\Theta}\big) \tag{14}$$

Our approach gives 144 moments which over-identify the 36 structural parameters of interest.

Table 2: Structural Parameter Estimates

Parameter	Mean [Min : Median : Max]
κ	1.281
	(0.084) [0.252]
ρ	4.385
	(0.501) [0.417]
$\underline{g}$	4,501
	(0.055) [14,837]
$\underline{s}$	0.602
	(0.055) [0.095]
γ_j	[0.66 : 0.81 : 0.95]
δ_j	[0.01 0.02 : 0.03]
N	5,110

Asymptotic robust standard errors are shown in parentheses and bootstrapped standard errors are in brackets.

Table 2 shows our point estimates along with the 10-50-90 percentile ranges for the state-level structural weight parameters, γ_j and δ_j. The top part of Table 2 shows our estimates of the two structural parameters, κ and ρ, governing substitution elasticities as well as the two reference levels for each form of electricity, $\underline{g}$ and $\underline{s}$. Substitution elasticities κ and ρ are precisely estimated and are such that $\widehat{\rho} > \widehat{\kappa} > 1$. Our estimate of the reference level for solar electricity is a small positive value, and we can reject the null hypothesis that $\underline{s} < 0$. At face value the reference level for grid consumption, $\underline{g}$, corresponds to about 37% of average household electricity consumption in our sample.

Note that because $\underline{s} \geq 0$, Proposition 1 tells us that rebound effects are *guaranteed* regardless of heterogeneity in the elasticities of county-level price indices. Thus, in all counties over all periods, our quantitative model will predict that subsidizing residential solar leads to higher levels of composite electricity consumption, q^e_{ct}. While it is clear that reducing the unit price of residential solar electricity leads to rebound effects, whether backfiring effects (whereby grid consumption, g_{ct} also increases as a response to a fall in p^s_{ct}) occur is ambiguous. Turning to Proposition 2 note that condition 3 is the relevant one for our parameter estimates, since neither conditions 1 or 2 are satisfied. The local values of the price-index elasticities, as well as local average income, y_{ct}, will determine whether grid consumption is associated with a positive cross-price elasticity, which is indicative of the presence of backfiring effects.

Our estimated substitution parameter between grid and solar electricity, $\widehat{\rho}$, substantially exceeds one, indicating that lowering solar prices should increase the share of electricity derived from PV systems. Further, we cannot rule out backfiring entirely, though in our quantitative analyses we find that it occurs in only a very small number of counties. While the structural parameters are precisely estimated, a Sargan-Hansen test of the over-identifying restrictions fails to hold at standard levels. This suggests the assumption of mean-zero econometric errors for all states and all goods may fail. For this reason, we evaluate in-sample model fit below. We also discuss the results of various out-of-sample extrapolation exercises in Technical Appendix B.3.

4.5 Model Fit

Before moving on to our counterfactual exercises, we assess model fit. Table 3 displays a selection of four fitted moments generated by our model along with their sample analogs, two of which are not targeted directly by our estimation procedure (solar percentage of grid and electricity expenditure share). The columns "Sample Mean" and "Sample Median" contain the observed data moments from our sample of counties over time, while the columns "Model Mean" and "Model Median" contain the analogous moments over the fitted values. The final column displays the portion of observed variance in the data explained by the model as measured by an R^2 statistic. Note that the values in Table 3 are summary statistics over the entire 2010-2018 sample period in *only* those county/year pairs for which we have sufficient data to include them in the estimating sample.

Table 3: Model Fit for County/Year Level Moments, 2010-2018

Moment	Sample Mean	Model Mean	Sample Median	Model Median	Model Fit (R^2)
Average Solar Consumption (kWh/yr)	66.6	52.7	4.94	5.38	0.49
Average Grid Consumption (kWh/yr)	10,153	10,369	9,220	9,326	0.68
Solar as Percent of Grid (%)	0.93	0.77	0.05	0.05	0.57
Elec. Expenditure Share (%)	1.98	1.97	1.89	1.89	0.75
N	5,085	5,085	5,085	5,085	5,085

Table 3 displays four sample moments against their fitted values. We winsorize the set of county/year pairs at the 0.5% level (dropping 25 observations from the right tail of the distribution over modeled solar consumption among the 5,110 observations) prior to comparing model fit. Note that, as a sensitivity analysis, we also winsorized the right rail at 1% and 2% of observations: the fitness results did not change. The R^2 measure displays an unadjusted coefficient of determination relative to a simple mean. Our procedure for computing the fitted values, which cover consumption of solar and grid electricity over the entire set of counties used in our structural estimation, is described in Technical Appendix B.3.1

The modeled moments match the data well. The mean fitted values are near the observed sample means for both types of electricity. Simple R^2 values for the solar-grid ratio and electricity expenditure share suggest the model is able to explain about half of county level variation during our sample period. The model underestimates mean solar electricity consumption slightly while overestimating median consumption. The fitted values for solar consumption are also right skewed, but cannot match the thickness of tails in the data. More formal statistical tests for how the fitted aggregates compare to their sample counterparts also fair well. Student's-t tests allowing for unequal variances fail to reject the equivalence of means for the model and sample values for grid consumption and electricity expenditure share values at the 5% level. However, these tests narrowly reject the equivalence of solar electricity consumption and the solar-grid ratio with sample means at the 5% level. We conclude that the structural model is able to account for approximately half of observed variation in county-level electricity consumption in our sample.

Finally, we note our sample contains county-level observations in only 16 states. To perform counterfactual exercises we ideally would like to examine how demand responds to policy changes nationally, not just in the 16 states for which we have full data. Thus, since counties across states only differ in terms of their state-level amenity weights, γ_j and δ_j, we extrapolate our estimates for these weight parameters to all counties in states not featured in our estimating sample, matching on a vector of county-level demographic and climatological covariates taken from Stanford's *DeepSolar* dataset. Technical Appendix B.3.2 explains our amenity-weight extrapolation procedure in detail.

5 Quantitative Exercises

We now engage in several quantitative exercises, referring to simulations involving the estimated model parameters as well as reduced-form evidence using state-level data. First, we use the structural model to simulate counterfactual scenarios that vary prices and preferences to understand how both of these channels affect demand for electricity. Second, given our counterfactual results demonstrate how displacement of grid consumption by new solar demand is not one-to-one, we provide reduced-form evidence from state-level time series that suggest this crowding out is indeed imperfect and grid and residential-solar electricity appear to be imperfect substitutes. Third, we consider how demand patterns vary geographically and what such heterogeneity implies for the efficacy of existing subsidies. Finally, we use the model to estimate the implied local and national cost of carbon abatement associated with residential PV subsidization policies.

5.1 Counterfactual Experiments

We engage in two counterfactual simulations to help us understand how solar panel subsidies interact with preference heterogeneity and price dispersion to affect demand. Our experiments compare a baseline calibrated model to alternative regimes which change solar electricity prices and preferences. We examine how policy and preference changes affect electricity consumption and the solar-grid ratio. We also look at subsidy efficacy, price elasticities, and marginal abatement costs.

All of our counterfactual analyses operate on the 2018 demand environment over the entire contiguous United States. We form a baseline scenario by solving the model for 3,074 U.S. counties in the lower-48 states for which we have sufficient data on household income and electricity prices in 2018. We weight each county by the number of households and aggregate it to the national level. Columns (2) through (4) of Table 4 display our baseline model along with results from our counterfactual experiments.

In aggregate we can judge our model performance against values in column (1), which displays national levels of residential grid and solar consumption taken directly from administrative data for 2018, as reported by the EIA. For the United States as a whole the baseline model over-predicts aggregate residential solar consumption relative to the EIA's estimates. While this undoubtedly reflects a degree of limitation in the explanatory power of the model, we believe this may stem from the EIA having different information on the total level of systems installed through the entirety of the 2018 calendar year than what is catalogued in the TTS dataset which we use to estimate the model.

Table 4: National Electricity Aggregates

	Data Aggregates	Model Baseline	No Subsidies	Median Gammas
	(1)	(2)	(3)	(4)
Solar Consumption (Billion kWh)	17.11	27.19	7.60	4.47
Grid Consumption (Billion kWh)	1,469	1,803	1,813	1,819
Solar-Grid Ratio	0.012	0.015	0.004	0.002

Table 4 shows electrical aggregates over simulated demand for 3,074 counties. Units for electricity are in billions of kilowatt-hours for solar and grid consumption to aid readability and allow for meaningful comparisons of magnitudes across scenarios. The solar-grid ratio is the ratio of the two electricity sources aggregated over all households at the national level. Aggregate residential solar generation for the U.S. is taken from the EIA's annual electricity module here.

Columns (3) and (4) present aggregate outcomes under two different counterfactual experiments. Our main result is featured in column (3). In this experiment we set all incentives and flat transfers to zero in all counties. The lion's share of the resulting price change stems from removing the 30% federal ITC. We then re-solve the model in each county to calculate hypothetical quantities of 2018 demand for s_c, g_c, and $\bar{c}_c$ without any government subsidization of residential solar. The simulations in (3) thus predict how much solar and grid electricity U.S. households would consume in a world where no subsidies for residential solar installations were present. Government price support increases residential solar demand by 19.5 billion kWh — a 255% increase relative to our counterfactual where there are no subsidies. We conclude from this that historical subsidies have been a dominant force driving residential PV uptake; the model suggests well over two-thirds of existing demand would not exist absent government support.

We run an additional simulation to illustrate how much of the variation in solar demand is determined by the geographic dispersion in solar-specific preferences (heterogeneous γ_j) suggested by the estimated model. Column (4) sets each county's grid electricity weight equal to the median U.S. level ($\gamma = 0.815$). This experiment turns off preference heterogeneity across counties while maintaining observed dispersion in county-level incomes and prices. Preference heterogeneity drives the vast majority of variation in solar panel uptake. As $\hat{\rho}$ is substantially larger than unity, local values of electricity preference weights along with solar-grid price differentials play a very important role in determining how households construct their optimal electricity bundle.

5.2 Rebound Effects

In accordance with Proposition 1 subsidies increase aggregate electricity consumption when going from the counterfactual in column (3) to our baseline model in (2). The magnitude is small, with subsidies increasing consumption by 9.7 billion kWh, or about 0.6% of U.S. residential electricity consumption. The increase in solar consumption of 19.5 billion kWh is offset only partially by a decrease in grid consumption of 9.8 billion kWh. The implied aggregate rate of displacement is -0.50: households purchased 0.50 fewer kWh from the grid for each additional unit of solar electricity consumed, implying a rebound effect of 50 percent.

Our estimated solar rebound effect is admittedly higher than existing estimates in the literature derived from quasi-experimental estimates on household panel data. Deng and Newton (2017) and Qiu, Kahn, and Xing (2019) respectively find rebound effects between 17 and 21 percent in Australia and 18 percent in Phoenix, Arizona. Aydın, Brounen, and Ergün (2023) find smaller effects in the

Netherlands of 7.7 percent after controlling for demand shifting between periods. However, these estimates exclusively capture an intensive rebound margin. A regression of household level total electricity consumption on concurrent household PV system generation identifies the marginal effect of higher generation on consumption for households with PV systems installed. This will not account for how the extensive effects of subsidies that induce new PV system adoption will change baseline household electricity consumption after uptake, a channel we believe our model captures on aggregate.

Finally, the model suggests that backfiring, as described in Proposition 2, is unlikely to be a concern. This is not to say that local backfiring effects do not occur (they do in a few places), but rather that across the entire economy isolated instances of backfiring are not large enough to cause an aggregate backfiring effect. If the ultimate goal of policymakers is to subsidize solar in order to reduce carbon-intensive grid consumption, it works: relative to the no-subsidy scenario, aggregate grid consumption falls by more than total electrical consumption rises.

5.3 Empirical Evidence of Imperfect Substitution

Our counterfactual simulations imply that increases in solar uptake due to solar subsidies do not directly offset an equal amount of grid consumption. This lack of one-for-one displacement of grid demand by induced solar generation in our counterfactuals is testable. If new residential solar generation does not fully displace grid consumption, we should see a similar pattern of imperfect crowding out in aggregate time series for residential grid demand as residential solar consumption has risen over the past decade. We test this hypothesis using panel data at the monthly frequency across U.S. states. If households on average were exactly substituting their consumption of grid electricity with that from residential PV systems, we would expect to see measured average grid consumption per household fall on a one-for-one basis with measured average residential solar production. This effect is captured in equation (15) below:

$$g_{jmt} = \underbrace{\beta\, s_{jmt}}_{H_0:\ \beta=-1} + \boldsymbol{\theta}^{\top}\boldsymbol{X}_{jmt} + \upsilon_j + \nu_m + \epsilon_{jmt} \tag{15}$$

where the variables g_{jmt} and s_{jmt} are average customer-level consumption of grid and residential solar electricity in state j in month m of year t. The vector $\boldsymbol{X}_{jmt}$ contains a time series of state-level cooling and heating degree days (CDDs and HDDs) to account for the fluctuations in demand owing to variation in weather, while υ_j and ν_m are state and month fixed effects.[28] This empirical strategy is similar to those used in Cullen (2013), Callaway, Fowlie, and McCormick (2018), Sexton et al. (2021), and Dauwalter and Harris (2023) who examine how additions of renewable capacity affect generation by carbon-intensive existing producers. We estimate (15) in levels as a Dickey-Fuller test rejects the null hypothesis for the presence of unit roots for in time series for per-customer grid consumption in all states.

Since demand for residential PV systems is unlikely to be independent of grid prices (which in-turn determine the quantities of electricity demanded) estimation of equation (15) using OLS may suffer from omitted variable bias. To account for potential endogeneity we use a one-year lagged value of cumulative U.S. solar panel imports (measured in kilowatts) as an instrument for average solar consumption. This variable is highly correlated with state-level residential solar generation (first-stage F statistics are over 100 in all specifications below). We use cumulative shipments as our baseline instrument due to a better first-stage fit and a closer intrinsic connection between cumulative wattage and generation. We believe the exclusion restriction is reasonable as local grid demand is unlikely to be influenced directly by the cumulative imports of solar panels to the U.S. in the prior year.

Table 5 displays results from iterating over different specifications of the regression in Equation (15) above. Point estimates for the coefficient on solar generation are negative in all five specifications and become significant at traditional levels once controls for CDDs and HDDs are added. This negative and significant association remains even when an alternative variable (lagged cumulative U.S. imports of solar panels in dollar terms) is used to instrument for solar generation. These results suggest that there is indeed a "crowding out" effect as residential generation displaces electricity consumed from the grid.

[28] The data we use for this regression comes from the EIA for years 2014-2022. The EIA tracks monthly consumption of grid electricity by the residential sector in each of the 50 states along with the District of Columbia. Monthly estimates of residential solar electricity generation at the state level are also available. We use these values along with the data on monthly end-use residential customers to create a time series of average per-customer grid and solar consumption.

Table 5: Residential Customer Solar/Grid Tradeoff

Dep. Var: g_{jmt}	(1)	(2)	(3)	(4)	(5)
Solar per Customer (s_{jmt})	-0.815	-0.594	-0.658	-0.899	-0.529
	(0.957)	(0.388)	(0.259)	(0.269)	(0.257)
HDDs	-	-	0.242	0.245	0.236
			(0.013)	(0.013)	(0.013)
CDDs	-	-	1.074	1.069	1.079
			(0.033)	(0.034)	(0.033)
State and Month Fixed Effects	No	Yes	Yes	Yes	Yes
Instrument	Imports	Imports	Imports	Value ($)	Overid.
$p\,(H_0 : \beta = -1)$	0.85	0.29	0.19	0.71	0.07

Observations are weighted by the number of customers each month at the state level using Stata's *aweights* command. Robust standard errors are shown in parenthesis. State-level population-weighted time series for CDDs and HDDs are taken from NOAA's Climate Prediction Center.

However, in all specifications examined we see a coefficient on solar consumption smaller than unity in magnitude. Test statistics of the null hypothesis that $\beta = -1$ are shown in the bottom row of Table 5. While the estimates are not precise enough to reject the null entirely, they are fairly low in two of the three specifications we examine that allow for both climatological controls and fixed effects. The relatively low p-value for the null hypothesis under the over-identified specification in column (5) remains present regardless of the choice of IV estimator.[29] Results from the panel data indicate that there is likely substantial but incomplete displacement of grid electricity consumption by household PV systems, which our model rationalizes. Our preferred estimation of $\widehat{\beta} = -0.66$ in column (3) is qualitatively similar to the aggregate displacement rate of -0.50 we find in our structural simulations, providing evidence that our structural results are not merely by-products of model selection.

[29] This result is also robust to the sign of measurement error we may expect in household solar generation; if some household solar generation goes unobserved by utilities or unreported to the EIA we would expect $\widehat{\beta}$ to be biased upward.

5.4 Geographical Heterogeneity in Induced Demand

How are counties associated with demand profiles that may respond differently to subsidies? Figure 4 displays the modeled increase in solar electricity consumption due to subsidization across all counties in the lower-48 states. Subsidies raise aggregate solar electricity consumption by 255 percent relative to our modeled counterfactual. There is substantial regional variation in household responsiveness to solar subsidization. The modeled county-level responsiveness ranges from an 86% increase at the 10^{th} responsiveness percentile to an over ten-fold increase in demand at the 90^{th} percentile. This counterfactual also allows us to calculate both local and aggregate own-price elasticities of solar demand. At the national level we find a population-weighted price elasticity of demand for solar of -3.20, larger than existing estimates for the extensive margin of adoption in the literature (Pless and Benthem 2019; Gillingham and Tsvetanov 2019). At the county level the mean (median) elasticity is closer to -3.45 (-2.92).

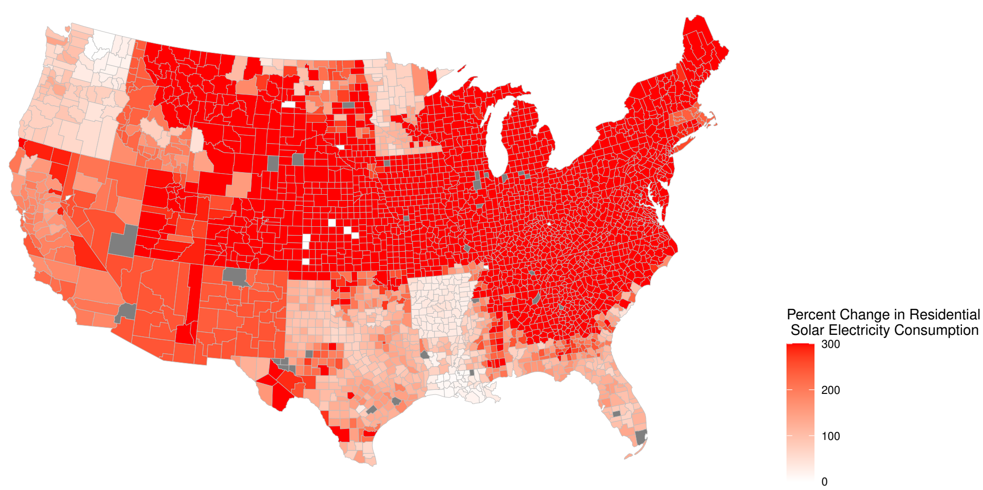

Figure 4: Increase in 2018 residential solar electrical induced by modeled subsidies.

We find that policies are cost-effective in terms of the cost of induced demand relative to observed electricity prices. The modeled subsidies induce 19.5 billion

kWh of demand at the cost of $1.44 billion in outlays by all levels of government. This translates to an average cost of induced demand of slightly under $0.07 per kWh. Figure 5 shows this measure of efficiency at the county level, mapping the changes in solar demand between our counterfactual and baseline scenarios divided by total government subsidy expenditure in the baseline model in each county. This ratio captures the average expenditure by a combined local, state, and national government, in dollars per kilowatt-hour, per unit of residential solar generation induced through subsidies. We find a wide range for the effectiveness on a per-dollar basis across counties. Our estimates suggest a range of costs between $0.06 per kWh in the 10^{th} percentile county to over $0.31 per kWh at the 90^{th} percentile. While the government cost of induced demand is lower than the private costs of solar electricity in over 2,500 of 3,074 counties, our estimates show that misallocation is likely if the singular goal of subsidies is to increase solar demand.

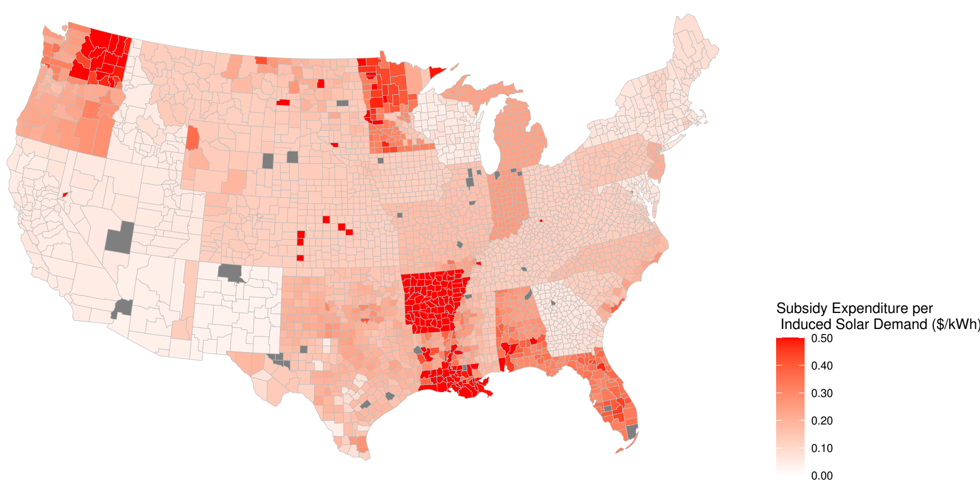

Figure 5: Subsidies per induced unit of solar demand.

5.5 Carbon Abatement Costs

Finally, we examine the relative efficacy of solar subsidies in the context of social marginal costs associated with changes in demand for electricity from the grid. This follows a recent strain of literature examining the spatial heterogeneity in externalities associated with the marginal generation of electricity which in turn lead to the benefits of displacing demand to change across locations (Callaway, Fowlie, and McCormick 2018; Sexton et al. 2021; Borenstein and Bushnell 2022; Dauwalter and Harris 2023). These studies show that the marginal social costs of electricity generation, both in terms of local pollutants and carbon emissions, vary substantially across space due to variation in the composition of marginal electricity generation as well as the distance between customers and generators.

Our model allows us to examine how changes in grid electricity demand due to residential solar subsidies translate to emissions reductions at the county level. As we also solve for the value of subsidies in the baseline scenario, we can combine these two figures to calculate the implied carbon abatement costs associated with solar subsidies at the county level. To translate our estimates for displaced grid consumption into carbon abatement, we take each county-level change in grid demand and multiply it by a local emission factor estimated by Borenstein and Bushnell (2022). These emissions factors give the carbon emissions associated with a marginal *kWh* of electricity consumed in each county due to the additional fuels used by generators. Figure 6 displays the geographic distribution of county-level marginal abatement costs. The missing gray counties are include those which experience backfiring as shown in Figure 6 leading to negative abatement costs.

We find that the implied local costs of abatement vary wildly across geographies. As this estimate combines dispersion in baseline grid prices, local subsidies, and emissions factors associated with marginal grid demand, our county level estimates range from \$150 per metric ton of carbon ($MTCO_2$) at the 10^{th} percentile to over \$500 per $MTCO_2$ at the 90^{th} percentile. At the national level, our estimates indicate the total reduction in grid demand from PV subsidies reduced emissions from the electricity sector by 4.3 million metric tons of CO_2 in

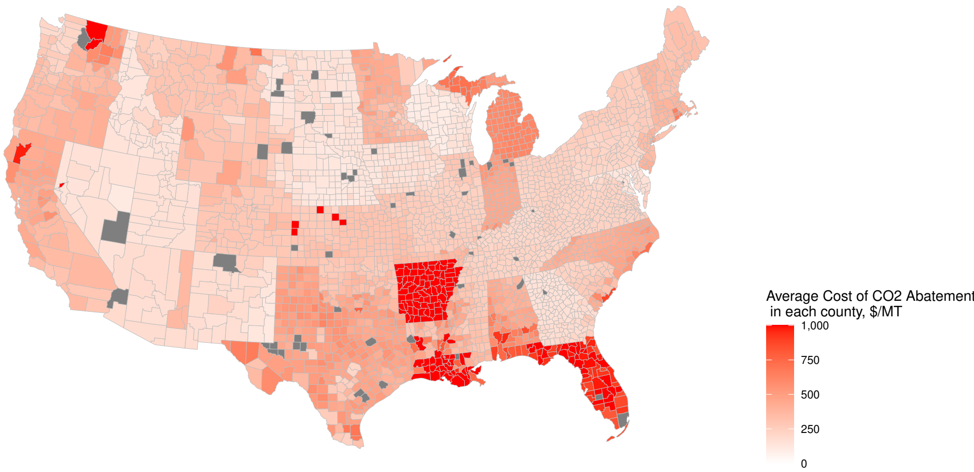

Figure 6: County-level implied carbon abatement costs from solar subsidies in units of $ per $MTCO_2$.

2018. This translates to an aggregate abatement cost of $332 per $MTCO_2$.[30] The highly disperse nature of abatement costs implied by our model reflect findings in the literature on PV deployment (Lamp and Samano 2023) and environmental policy more broadly (Cruz and Rossi-Hansberg 2022) that illustrate the importance of spatial heterogeneity when considering the effects of uniform taxes and subsidies at large scales.

6 Conclusion

This paper provides a structural model to analyze the effects of PV subsidies in the presence of geographic heterogeneity in income, prices, and preferences. Avoiding the assumption that electricity from residential solar panels is perfectly

[30]The range of existing estimates of the abatement costs associated with solar PV subsidies is quite broad (Gillingham and Stock 2018). For residential systems, Gillingham and Tsvetanov (2019) estimate solar subsidies in Connecticut were associated with abatement costs of $364 per $MTCO_2$, while back of the envelope calculations in Crago and Chernyakhovskiy (2017) find costs of $184.

substitutable with retail purchases from the grid allows us to more carefully examine implications for aggregate electricity consumption from all sources. We leverage rich data from multiple academic and administrative sources to illustrate that uptake is spatially diffuse, imperfectly correlated with factors such as insolation and income, and often remains high in locales where solar prices exceed prices of electricity from the grid. With aggregate time series we show how the apparent imperfect substitutability suggested by observed price dispersion emerges in the form of imperfect crowding out of residential grid consumption by households' PV production.

We then form and estimate a structural model for preferences over a composite electricity good comprising solar- and grid-drawn electricity along with outside consumption. The model allows for flexible amenity parameters at the state level and performs reasonably well at fitting uptake for counties across the contiguous United States. When estimated, our findings suggest that solar subsidies have increased aggregate consumption at the national level by 255%. Subsidies succeeded at displacing demand from the grid at a relatively high cost, in part due to imperfect crowding out of grid consumption by induced solar demand. At the national level we estimate on average that $0.07 of subsidy expenditure was required to induce an additional kWh of demand for electricity from PV systems. This point estimate masks substantial heterogeneity in spatial responsiveness to subsidies; targeted policies would be much more cost-effective.

More importantly, our results suggest that subsidies inducing an additional one kWh of residential PV generation do not perfectly crowd out one unit of demand from the grid. Our reduced-form evidence suggests this displacement is between 0.6 and 0.8, and our model suggests more substantial rebound effects of 50%. This imperfect crowding out is quantitatively important; our point estimate of a national abatement cost of $332 per $MTCO_2$ is twice as high as it would be absent the rebound effect described above. Existing policies promoting residential PV adoption are costly in terms of reducing emissions from electricity consumption.

Recent policies such as the ten-year extension of the 30% residential Investment Tax Credit for PV systems or the solar mandate in Title 24 of California's Building Energy Efficiency Standards indicate continued intentions of policymakers to adopt second-best solutions for mitigating pollution associated with electricity demand. The high costs of abatement and evidence of imperfect crowding out suggest policymakers should exercise caution when viewing residential solar subsidies as alternatives to traditional abatement. This is exacerbated by the inherent regressive aspect of green consumption subsidies both implied in our model and born out in the data (Borenstein and Davis 2016; Borenstein 2017). Our findings add to the growing body of evidence that second-best policies, even if more politically palatable, are unlikely to be substitutes for addressing externalities at the source.

Our work shows that amenity preferences play an important role in both the uptake of environmental goods and the efficacy of government policy. Future work could attempt to examine more closely the underlying structures of environmental preferences and adopt a dynamic setting for PV system uptake that allows for heterogeneity in household-level investment decisions across other forms of green durables.[31] Crucial to these (and other) investigations will be determining the rate at which households substitute between existing goods and the emerging "green" alternatives.[32] If environmentally friendly goods prove insufficiently adequate substitutes for current products, subsidies may be ineffective at abating the use of existing products.

[31] Buchsbaum (2023) examines how long-run changes in electricity prices affect households' uptake of energy-efficient durable goods and PV systems in the Californian setting.

[32] Lyu (2023) finds that PV system adoption itself has a positive spillover in terms of inducing additional electric vehicle demand, suggesting the potential for latent *complementarity* to further complicate this calculation.

References

Aydın, Erdal, Dirk Brounen, and Ahmet Ergün (2023). "The rebound effect of solar panel adoption: Evidence from Dutch households". *Energy Economics* 120, p. 106645 (cit. on pp. 3, 5, 27).

Barbose, Galen et al. (2019). "Tracking the Sun: Pricing and Design Trends for Distributed Photovoltaic Systems in the United States-2019 Edition" (cit. on pp. 4, 17).

Borenstein, Severin (2015). "A microeconomic framework for evaluating energy efficiency rebound and some implications". *The Energy Journal* 36.1 (cit. on p. 5).

— (2017). "Private net benefits of residential solar PV: The role of electricity tariffs, tax incentives, and rebates". *Journal of the Association of Environmental and Resource Economists* 4.S1, S85–S122 (cit. on pp. 4, 19, 36).

Borenstein, Severin and James B Bushnell (2022). "Do two electricity pricing wrongs make a right? Cost recovery, externalities, and efficiency". *American Economic Journal: Economic Policy* 14.4, pp. 80–110 (cit. on pp. 12, 33).

Borenstein, Severin and Lucas W Davis (2016). "The distributional effects of US clean energy tax credits". *Tax Policy and the Economy* 30.1, pp. 191–234 (cit. on pp. 4, 8, 36).

Buchsbaum, Jesse (2023). "Long-Run Price Elasticities and Mechanisms: Empirical Evidence from Residential Electricity Consumers" (cit. on p. 36).

Callaway, Duncan S, Meredith Fowlie, and Gavin McCormick (2018). "Location, location, location: The variable value of renewable energy and demand-side efficiency resources". *Journal of the Association of Environmental and Resource Economists* 5.1, pp. 39–75 (cit. on pp. 5, 29, 33).

Chan, Nathan W and Kenneth Gillingham (2015). "The microeconomic theory of the rebound effect and its welfare implications". *Journal of the Association of Environmental and Resource Economists* 2.1, pp. 133–159 (cit. on pp. 5, 13).

Chang, Juin-Jen, Wei-Neng Wang, and Jhy-Yuan Shieh (2018). "Environmental rebounds/backfires: Macroeconomic implications for the promotion of environmentally-friendly products". *Journal of Environmental Economics and Management* 88, pp. 35–68 (cit. on p. 4).

Crago, Christine Lasco and Ilya Chernyakhovskiy (2017). “Are policy incentives for solar power effective? Evidence from residential installations in the Northeast”. *Journal of Environmental Economics and Management* 81, pp. 132–151 (cit. on p. 34).

Cruz, José-Luis and Esteban Rossi-Hansberg (2022). *Local Carbon Policy*. Tech. rep. National Bureau of Economic Research (cit. on p. 34).

Cullen, Joseph (2013). “Measuring the environmental benefits of wind-generated electricity”. *American Economic Journal: Economic Policy* 5.4, pp. 107–133 (cit. on p. 29).

Dauwalter, Travis E and Robert I Harris (2023). “Distributional Benefits of Rooftop Solar Capacity”. *Journal of the Association of Environmental and Resource Economists* 10.2, pp. 487–523 (cit. on pp. 5, 6, 29, 33).

Davis, Lucas W (2023). *The Economic Determinants of Heat Pump Adoption*. Tech. rep. National Bureau of Economic Research (cit. on p. 3).

De Groote, Olivier and Frank Verboven (2019). “Subsidies and time discounting in new technology adoption: Evidence from solar photovoltaic systems”. *American Economic Review* 109.6, pp. 2137–72 (cit. on pp. 5, 12, 19).

Deng, Gary and Peter Newton (2017). “Assessing the impact of solar PV on domestic electricity consumption: Exploring the prospect of rebound effects”. *Energy Policy* 110, pp. 313–324 (cit. on pp. 5, 27).

EIA (2020). “Monthly energy review”. *DOE/EIA-0035 (2020/9). Office of Energy Statistics, US Department of Energy Washington* (cit. on p. 4).

Feger, Fabian, Nicola Pavanini, and Doina Radulescu (2022). “Welfare and redistribution in residential electricity markets with solar power”. *The Review of Economic Studies* 89.6, pp. 3267–3302 (cit. on pp. 5, 12).

Flowers, Mallory E et al. (2016). “Climate impacts on the cost of solar energy”. *Energy Policy* 94, pp. 264–273 (cit. on p. 17).

Gillingham, Kenneth and James H Stock (2018). “The cost of reducing greenhouse gas emissions”. *Journal of Economic Perspectives* 32.4, pp. 53–72 (cit. on p. 34).

Gillingham, Kenneth and Tsvetan Tsvetanov (2019). “Hurdles and steps: Estimating demand for solar photovoltaics”. *Quantitative Economics* 10.1, pp. 275–310 (cit. on pp. 5, 19, 31, 34).

IRS, Internal Revenue Service (2019). “Statistics of Income - Individual Income Tax Returns, Line Item Estimates” (cit. on p. 19).

Ito, Koichiro (2014). "Do consumers respond to marginal or average price? Evidence from nonlinear electricity pricing". *American Economic Review* 104.2, pp. 537–63 (cit. on p. 12).

Khazzoom, J Daniel (1980). "Economic implications of mandated efficiency in standards for household appliances". *The energy journal* 1.4 (cit. on p. 4).

Kulmer, Veronika and Sebastian Seebauer (2019). "How robust are estimates of the rebound effect of energy efficiency improvements? A sensitivity analysis of consumer heterogeneity and elasticities". *Energy Policy* 132, pp. 1–14 (cit. on p. 13).

Lamp, Stefan and Mario Samano (2023). "(Mis) allocation of Renewable Energy Sources". *Journal of the Association of Environmental and Resource Economists* 10.1, pp. 195–229 (cit. on pp. 3, 5, 6, 34).

Langer, Ashley and Derek Lemoine (2022). "Designing dynamic subsidies to spur adoption of new technologies". *Journal of the Association of Environmental and Resource Economists* 9.6, pp. 1197–1234 (cit. on p. 5).

Lewbel, Arthur (1989). "Exact aggregation and a representative consumer". *The Quarterly Journal of Economics* 104.3, pp. 621–633 (cit. on p. 12).

— (2001). "Demand Systems with and without Errors". *American Economic Review* 91.3, pp. 611–618 (cit. on p. 16).

Lyu, Xueying (2023). "Are Electric Cars and Solar Panels Complements?" *Journal of the Association of Environmental and Resource Economists* 10.4, pp. 1019–1057 (cit. on pp. 19, 36).

Pless, Jacquelyn and Arthur A van Benthem (2019). "Pass-through as a test for market power: An application to solar subsidies". *American Economic Journal: Applied Economics* 11.4, pp. 367–401 (cit. on pp. 4, 19, 31).

Qiu, Yueming, Matthew E Kahn, and Bo Xing (2019). "Quantifying the rebound effects of residential solar panel adoption". *Journal of Environmental Economics and Management* 96, pp. 310–341 (cit. on pp. 3, 5, 13, 27).

Sexton, Steven et al. (2021). "Heterogeneous solar capacity benefits, appropriability, and the costs of suboptimal siting". *Journal of the Association of Environmental and Resource Economists* 8.6, pp. 1209–1244 (cit. on pp. 5, 6, 29, 33).

Yu, Jiafan et al. (2018). "DeepSolar: A machine learning framework to efficiently construct a solar deployment database in the United States". *Joule* 2.12, pp. 2605–2617 (cit. on pp. 6–8).