Cars of the future, today? Estimating the contribution of electric vehicles to California's residential electricity demand

Fiona Burlig, James Bushnell, David Rapson, and Catherine Wolfram*

February 24, 2020

Abstract

California is now home to over 650,000 electric vehicles (EVs), less than 5% of which are charged at home using a meter dedicated to EV use. State policy has thus been forced to rely upon either survey data or approximations based on selected samples to estimate the extent and timing of residential electricity use devoted to EVs. We match a novel dataset comprised of 1.7 billion household electricity meter readings to electric vehicle adoption events at the address-level from 2014-2017 in California. We use these rich data in conjunction with a panel fixed effects approach to estimate the effects of EV adoption on electricity load. In our sample, EVs increase household load by 0.10 to 0.15 kWh per hour, or 17-25 kWh per week, the majority of which is concentrated during evening and nighttime hours. While these estimates are roughly half of the estimates used as an input into state EV-related forecasts and policies, the load impacts are concentrated in the late night and early morning, corresponding to higher marginal emissions factors than if charging had taken place mid-day.

*Burlig: Harris School of Public Policy and Energy Policy Institute, University of Chicago, and NBER. Email: burlig@uchicago.edu. Bushnell: Department of Economics and Davis Energy Economics Program, UC Davis, and NBER. Email: jbbushnell@ucdavis.edu. Rapson: Department of Economics and Davis Energy Economics Program, UC Davis. dsrapson@ucdavis.edu; Wolfram: Haas School of Business and Energy Institute at Haas, UC Berkeley, and NBER. Email: cwolfram@berkeley.edu. This research was partially supported with funds from CA SB1. Jesse Buchsbaum, Benjamin Dawson, Shelly He, and Ellen Lin provided excellent research assistance. All remaining errors are our own.

1 Introduction

Increasingly, plans for future reductions in carbon emissions are working toward a rapid reduction in the carbon intensity of electricity production coupled with a transition of other sectors toward the adoption of electricity as a substitute fro conventional fossil fuels. This process is often referred to as *electrification.* In California, for example, Senate Bills 32 and 100, commit the state to aggressive new greenhouse gas (GHG) reductions through a strategy focusing on the electrification of the residential, industrial, commercial, and, most significantly, transportation sectors. In particular, the state has devoted substantial financial resources to a broad suite of policies aimed at electrifying passenger transportation.

An overarching strategy of transitioning transportation and other energy use applications to electricity has profound implications for the electricity sector. This sector has itself experienced massive changes to the profiles of both end-use energy consumption and production. Residential electricity demand has been flat or declining for a decade, and mid-day electricity demanded from the grid has declined significantly due to the expansion of residential rooftop solar production. Strategies promoting electrification of transportation, home heating, and other applications raise the prospect of additional massive shifts in both the level and timing of electricity demand. This in turn creates large implications for the economics and reliability of electric systems.

Despite this prospect, relatively little empirical evidence is available about the impacts that electrification has had on residential electricity consumption to date. Although California is now home to over 650,000 electric vehicles (EVs), less than 5% of these vehicles are charged at home using a meter dedicated to EV use. Infrastructure planning and the implementation of important policies such as EV incentive programs and the Low Carbon Fuel Standard have had to rely upon either survey data or heuristic approximations to estimate the amount and timing of electricity use devoted to EVs.

In this paper we present what we believe to be the first attempt to rigorously and empirically measure the impacts of EV adoption on household electricity consumption. We apply hourly electricity consumption data from 2014 to 2017 for a purpose-built sample of 10% of the households in California?s three large investor-owned utilities (IOUs): Pacific Gas and Electric (PGE), Southern California Edison (SCE), and San Diego Gas and Electric (SDGE). We combine these data with household-level EV registration data to estimate the impact of EVs on residential consumption.

Our main analysis deploys an event study approach in which we pair household-level data on EV adoption with household-level data on electricity consumption to estimate the change in load resulting from EV adoption. This approach enables us to estimate the relationship between EV adoption and load for the average EV-owning household, something that has been challenging in prior analyses because of the lack of data. We compare our estimates of household usage to load at EV-dedicated meters to demonstrate the potential selection bias involved in current estimates of EV-load that rely exclusively upon directly measured households. Our results indicate that this bias is substantial and significant. Households with dedicated EV-meters consume roughly 2-3 times the amount of electricity per day than our estimates indicate is being consumed at households without dedicated meters. This implies that the expected impact on electricity demand, as well as the amount of usage of EVs, may be substantially inflated due to this selection bias.

We also examine the hourly breakdown, or "load-profile," of EV charging electricity demand. The bulk of EV charging is happening in the early morning hours, when the California electricity system is not capacity constrained and is also clean. We also examine the heterogeneity in charging behavior by vehicle type and by electricity rate class. There is substantial variation across vehicle types, with Tesla owners consuming much more electricity than those of other vehicle makes. There is also evidence that those choosing to enroll in EV-specific rate plans consume more electricity both before and after they purchase their EV.

Section 2 provides further background on the anticipated impact of EVs on the electricity sector. Section 3 summarizes our data sources. In Section 4 we describe our empirical approach and in Section 5 we present our results. Our conclusions are summarized in Section 6.

2 Background

Over the last decade California has experienced a sharp reduction in the GHG intensity of its electricity production. According to the Emissions Inventory maintained by the California Air Resources Board, GHG emissions from the commercial electricity sector (excluding industrial self-generation) have declined from roughly 120 mmTons in 2007 to just over 50 mmTons in 2017 (Borenstein et al. (2019)). At the same time overall electricity consumption, after experiencing a sharp contraction during the financial crises and its aftermath, has been relatively flat. This appears to signal progress on the GHG front in the dimension of reduced intensity *and* per-capita consumption; however, the overall emissions picture is more complicated.

While electricity purchases from the grid have declined slightly over the decade, both natural gas and gasoline consumption have risen. These facts illustrate the challenging fact that the share of energy consumption has been shifting toward the more carbon-intensive fossil fuels. California policy and the planning apparatus are heavily invested in reversing that trend.

As described below, there are a ambitious programs promoting the adoption and use of EVs. The growth in EVs is anticipated to be a contributor to a reversal in the decline of electricity consumption over the next decade. As illustrated in Figure 1, recent forecasts from the California Energy Commission's California Energy Demand (CED) indicate a rise in consumption to over 330 TWh per year by 2030 in the "mid" case.[1] Current charging by EVs is estimated to account for less than 1% of statewide electricity consumption in 2018, but is forecast to grow by up to 10 times over the next decade by the CEC.

Figure 1: Load forecast estimates from the California Energy Demand report by the California Energy Commission.

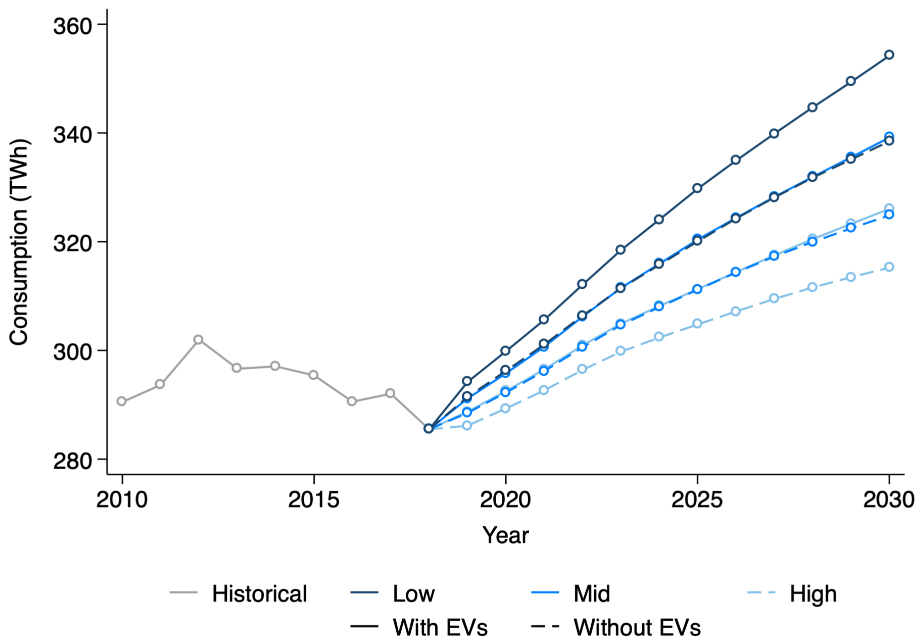

While the contribution of EVs may seem modest at first glance, it is important to note that they account for almost all the expected growth in the electric system over the next decade and that the timing of these charging loads could result in their comprising a much larger share of system net peak consumption (i.e. net of renewable, primarily solar output).

The timing of EV load will be a crucial factor in determining how electricity markets will be affected. The profile of residential load is already changing rapidly as a result of investments in behind-the-meter solar generation. Figure 2 presents net residential electricity demand by

[1] *Electricity and Natural Gas Demand Forecast* (2018).

hour-of-day in PG&E territory from 2014-2017 according to our subsample (which we will describe in Section 3). When considering the impact that a given amount of EV-related demand will have on the system, whether is falls during the mid-night trough or near the evening peak will materially affect the economic value and potentially environmental impact of the energy consumed.

Figure 2: Hourly residential net electricity load, PG&E 2014-17

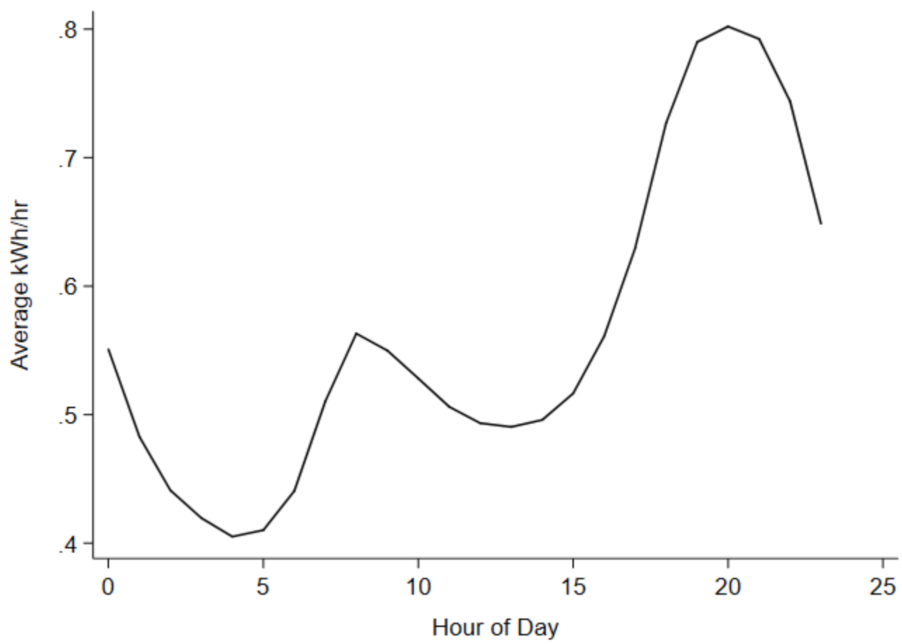

2.1 EV Policy in California

As mentioned above, transportation electrification is a central pillar of California's decarbonization goals. Despite some concerns over the ultimate carbon benefits of electrification (Holland et al. (2018)), EVs do provide some benefits over conventional vehicles (Archsmith et al. (2015). These aspirations were articulated in the form of a 2012 executive order by Governor Brown to have 1.5 million EVs on the road by 2025, and a separate goal of 5.0 million EVs by 2030. Both the state and federal governments have adopted policies that are at least partly intended to promote the supply and demand of EVs. The California Zero Emission Vehicle (ZEV) Mandate generates credits for manufacturers that sell EVs, and requires all manufacturers to either produce or purchase these credits. Similarly, the Corporate Average Fuel Economy standards offer an additional incentive to manufacturers that produce EVs.

On the demand side, there are large federal and state subsidies. As part of the American Clean Energy and Security Act of 2009, up to $1.5 billion in federally-funded tax credits were made available to consumers of each manufacturer. In California, the Clean Vehicle Rebate Project (CVRP) offers new EV buyers between $1,500 and $2,500 for new Plug-in Hybrid Electric Vehicles (PHEV) and Bettery Electric Vehicle (BEV) purchase, respectively. These are often augmented by an array of other state and local incentives such as high-occupancy vehicle lane access and/or free or subsidized charging. Despite some concerns about the distributional impacts of these subsidies (Borenstein and Davis (2016)), the California and Federal incentives are clearly having an impact on EV adoption (Muehlegger and Rapson (2018)). Through 2017, roughly 700,000 EVs claimed a total of an estimated $4.7 billion in federal subsidies have been paid to EV buyers nationwide. In California, 340,000 EVs have been purchased under the CVRP for a total of over $770 million in subsidies as of October 25, 2019.

Figure 3 displays heat maps of where EV purchases are concentrated in California in 2014 and 2017. Most EV purchase activity occurs in cities along the cost, with major concentrations in the Bay Area, Los Angeles and San Diego. At the end of our sample in December 2017, there are 423,297 plug-in EVs registered in California. This represents a 2.9% share of the 14.6 million passenger vehicles that were registered in California that year. Table 1 shows the number of EVs we observe in our sample (see Section 3 below for a description of our sampling approach) relative to the total number of EVs in the state.

While the stated goals for the adoption of EVs in the state are straightforward, understanding the translation of EVs on the road to electricity demand is a more complicated task. The challenges involved are described in more detail in the following section.

Table 1: EV Counts Within the Study by IOU, and the California Total.

Utility	EV Count
PGE	74,468
SCE	64,378
SDGE	3,125
Study Total	141,971
California Total	423,297

2.2 Measuring EV Electricity Consumption

By far the largest challenge in evaluating the impact of EV growth on the electric system is the lack of directly-measured consumption data for residential home charging. Home charging

Figure 3: EV distribution and growth between 2014 and 2017 throughout California.

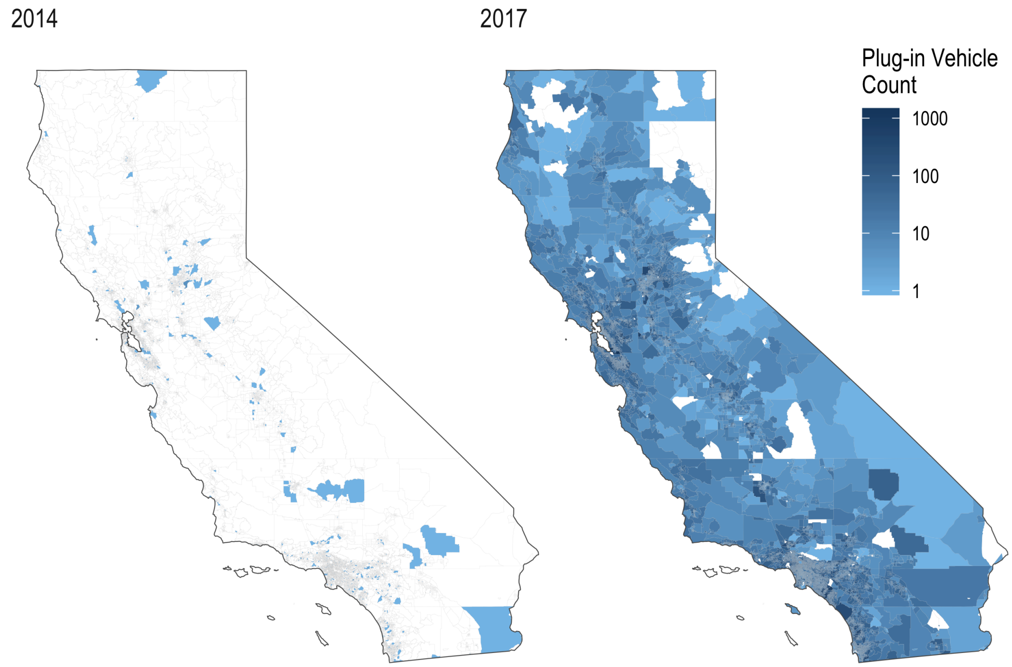

of EVs does not require a separate meter or even separate equipment for low-voltage charging. Consequently, less than 5% of EVs are directly metered when charging at home (Electric (2019)). While charging at networks operated either by commercial charging businesses or vehicle manufactures such as Tesla is directly metered, the ARB estimates that upwards of 80% of EV households charge at home some or all of the time. Thus the vast majority of EV charging is currently unmeasured.[2]

Absent any detailed data on the un-metered customers, California planning and policy has come to rely upon projections based on the small share of EVs that are directly metered. This is problematic because these meters were not deployed randomly. They were chosen to be installed, at potentially high cost, by individual customers. This creates a significant potential 'selection bias' that could cause projections based solely upon this non-random sample to be inaccurate and unreliable. The extent of this bias is, as of yet, untested.

[2]The best data on EV charging use is probably within the vehicles themselves. Most Original Equipment Manufactures (OEMs) collect charging data from the cars they have sold, but these data are held closely due to strategic business interests and privacy concerns.

2.2.1 EV Rate Options

The interaction of electricity rates, vehicle adoption, and energy use is an area that deserves considerable attention. Despite the fact that many customers do not fully respond to complex rate structures (Ito (2014)), the disparity between marginal electricity prices and social marginal cost in California is large enough that it could be a significant impediment to electrification (Borenstein and Bushnell (2018)). Electricity prices may very well be influencing the decision to enroll in an EV rate and whether or not to install a dedicated EV meter. All three investor owned utilities in California offer generally two rates: first an EV rate for the whole house, whereby the entire house is on the time-of-use (TOU) rate, or the option to submeter the EV itself. All EV-specific TOU rates are time-varying by season (summer and winter), and weekends and holidays. When the EV is submetered, only the EV meter is on the TOU rate; the rest of the household remains on their current tariff schedule. Over time, the EV rates at each IOU have changed names and structures. However, they generally include either a whole-house rate that is TOU or an EV-specific TOU rate that leaves the house itself on its existing tariff.

EV rates are typically only offered to individuals with battery electric or plug-in electric vehicles, not hybrid electric vehicles.[3] Thus, a household wishing to make the transition an EV rate need only demonstrate proof of EV ownership. In some cases, the distribution system may require an upgrade in order to support the increase in load. However, these are limited (see CPUC proceeding 19-IEPR-04). To obtain a designated EV (submetered) rate requires the purchase and installation of the meter itself. This can cost between a few hundred and a few thousand dollars.

However the California Public Utilities Commission (CPUC) has required the IOUs to run pilot programs for designated meters. In this case the IOUs have subsidized the meters, and with the case of Pacific Gas and Electric, their most recent pilot program subsidized thee meters by $210, with an additional $17.50 per month. In the case the customer is not satisfied with their EV rate, they are always welcome to opt-out of the tariff structure and revert back to their old tariff. In addition to the ability to opt-out of the EV tariff during the pilot program, PGE also is only keeping participants in the pilot program on the tariff for 12 billing cycles unless the participant decides to remain on the program. Program enrollment for dedicated metering of the EV has been a small fraction of total EV ownership.

[3] The household TOU rate that is currently offered by SCE is an exception. It is open to all households, irrespective of EV ownership.

3 Data description

We have assembled a novel and extremely rich dataset on household electricity consumption and billing and electric vehicle ownership. These data come from two main administrative sources: California's investor-owned utilities and the California Department of Motor Vehicles (DMV). We describe each source in turn below.

3.1 Investor-owned utilities

We obtained electricity consumption and billing data from the three investor-owned utilities in California: Pacific Gas and Electric (PGE), Southern California Edison (SCE), and San Diego Gas and Electric (SDGE). In this draft, we focus our analysis on PGE, but we are working to expand our approach to SCE and SDGE as well.

3.1.1 Sampling

In an ideal world, we would obtain the universe of meters for PGE for all years. In order to reduce the burden on PGE, our sample is limited to 10 percent of the service territory, from 2014 to 2017.[4] We designed a two-part sampling strategy to allow us to capture a large share of the IOU's EVs, while also providing coverage of the utility territory as a whole.

Main analysis sample For our main analysis sample, we aimed to select ZIP codes to over-sample both regions with a high number of electric vehicles and areas with low power supply reliability.[5] For these ZIP codes, we obtained the universe of residential meters, which allows us to include non-EV owners as controls in our analysis. In particular, the sampling occurred as follows:

- We ranked ZIP codes based on 2016 customer-hours of power outages based on a list provided by PGE. We selected ZIP codes on this list in descending order until we had accumulated 4% of the total population of the service territory. We sampled 100% of the residential meters in each of these ZIP codes.
- We ranked ZIP codes based on EV penetration[6], and selected ZIP codes on this list in descending order until we had accumulated 4% of the total population of the service territory. We sampled 100% of the population in each of these ZIP codes.

[4]We are in the process of initiating a data request that would allow us to extend the sample through the end of 2019.
[5]This sampling strategy was designed to allow us to study both EVs and blackouts; we focus on EVs in this paper.
[6]This original dataset on EV penetration by ZIP code was derived from Experian Automotive.

This approach yields a sample of 8% of the PGE service territory, weighted towards high EV penetration and low reliability ZIP codes.

Random sample In addition to our main analysis sample, we also sampled an additional 2% of the service territory to use as a comparison group. For this sample, we simply took a 2% random sample of all households in the PGE service territory, excluding the ZIP codes in the main analysis sample.

3.1.2 Electricity data

We obtained three types of data from PGE: monthly billing information, hourly electricity consumption data, and customer details. In addition to the consumption and billing data, we observe each customer's street address, latitude and longitude, rate class, and a solar panel interconnection date, where applicable. Our sample consists of 362,945 households, and over 1.7 billion hourly electricity consumption observations. Table 2 presents summary statistics from our sample, divided between EV owners and non-EV owners. We observe that EV households are much more likely to have solar, additional meters, and consume more electricity per hour. They also have higher bill consumption and bill amounts than their non-EV-owning counterparts.

3.2 DMV

In addition to our electricity consumption data, we obtained California DMV registration records for the period 2008 to 2019. Our main dataset contains the universe of EVs registered in the state during this time period, selected using 7-digit VIN stems.[7] This is a uniquely detailed dataset: for each EV, we observe the address, make, model, year, and VIN stem, as well as a series of other vehicle characteristics. We also observe an anonymized unique vehicle identifier, which allows us to track vehicles over time, as well as the registration date. We observe 423,297 unique vehicles in the state of California during this period, 74,468 of which are in ZIP codes belonging to the sample of the PGE service territory that matches our analysis sample. 63,765 of these are in the PGE service territory between 2014 and 2017, the time period of our electricity use information. Figure 4 presents summary statistics on the EVs in our sample. We observe that incremental EV additions are increasing over time: the increase in EVs is happening at an increasing rate in the PGE service territory. In addition, we see that Chevrolet, Nissan,

[7] We are in the process of obtaining similar data for ICE vehicles to use as placebos.

Table 2: Summary statistics on electricity consumption

	Non-EV households	EV-A	EV-B	Other EV households
Panel A: AMI data				
Number of meters	327,582	7,979	69	14,389
Portion of meters with solar	.05578	.2243	.1009	.2111
Average kWh per hour	.7439	1.412	.3363	.7836
Median kWh per hour	.567	.8938	.2771	.6421
Minimum kWh per hour	-21.8	-5.729	-2.65	-50.73
5th Percentile kWh per hour	.003988	.01024	0	.002381
95th Percentile kWh per hour	1.865	4.468	.8643	1.896
Observations	68,071,384	1,661,695	9,347	3,022,370
Panel B: Billing data				
Number of billing accounts	594,998	7,871	60	38,338
Average bill consumption	559.1	1,389	247.3	629.5
Average bill amount	111.1	243.9	34.93	139.7
Median bill consumption	439	798.2	214	499
Median bill amount	69.27	142.4	28.17	81.91
Observations	15,666,187	182,913	1,419	804,499

Notes: This table provides basic summary statistics on households in our sample, throughout our sample period. Observations in Panel A have collapsed hourly data to the week-hour level to ease computational burden.

Tesla, and Toyota are the majority manufacturers.

Matching We use a string matching algorithm to assign EVs to PGE households. We begin by cleaning the data so that common words are represented in the same way in both datasets (*e.g.* "ave" vs. "avenue"; "st" vs "street"; *etc*). Next, we perform an exact match on address. Finally, we use a fuzzy string match to finalize our merge. Out of the more than 63,000 vehicles registered in ZIP codes in our main PGE analysis sample, we matched 57,290 cars to PGE addresses: a match rate of 89.8 percent.[8]

With access to this unique dataset on both electricity use and EV registration, we are able to empirically estimate the effects of EV ownership on energy use among PGE households.

[8]Some of the remaining addresses belong to municipal and other local utilities which share ZIP codes with PGE, so we would not expect them to match to PGE addresses.

Figure 4: Incremental EV additions over time

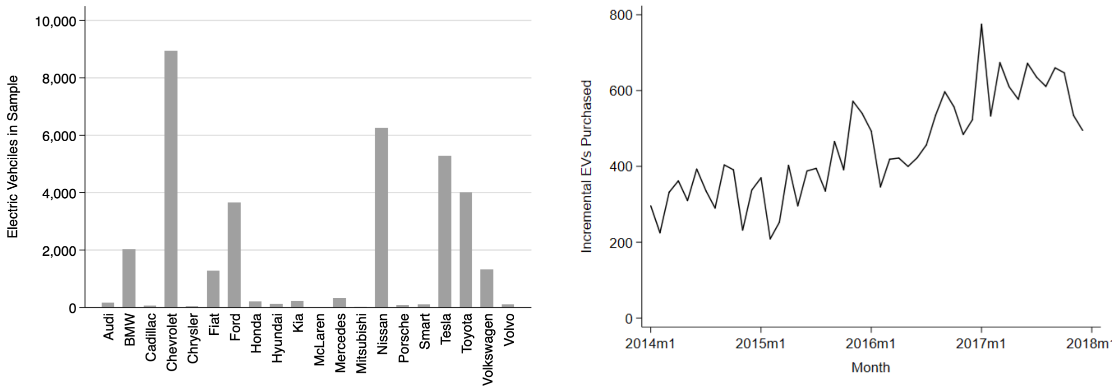

Notes: This figure plots data on our sample of EVs. The left panel displays EV registrations by vehicle manufacturer in the PGE service territory during our sample period. Chevrolet, Nissan, Tesla, Toyota, and to a lesser extent, Ford, are the largest manufacturers by volume. The right panel plots the number of incremental EV additions during our sample over time. The number of EV additions each month is clearly trending upwards.

4 Empirical Approach

In order to estimate the effect of EV charging on residential load, we leverage our high-frequency consumption data in conjunction with our address-level data on EV registration dates, using a panel fixed effects research design.

In our baseline analysis, we simply estimate the effect of an EV being registered in a difference-in-differences specification:

$$Y_{ith} = \beta EV_{it} + \gamma Solar_{it} + \alpha_i + \delta_t + \varepsilon_{ith} \tag{1}$$

where Y_{ith} is electricity consumption in household i during week t in hour h.[9] EV_{it} is an integer-valued variable, equal to the number of EVs registered at household i by week t, and zero for households without EVs. β, our variable of interest, describes the effect of EV registration on electricity consumption, scaled to be in units of kWh per hour. $Solar_{it}$ is an indicator, equal to one if week t is after household i's solar interconnection date, and zero otherwise. We include this control because many households that purchase EVs also install solar panels, which substantially reduce their net load. Without this control, we run the risk of substantially attenuated estimates. α_i are household fixed effects, and δ_t are week-of-sample fixed effects. This is a representation of our baseline specification. We also present variants, including no con-

[9]In some specifications, we collapse the data to the week-by-hour level to speed computation time. As we show in **?**, collapsing high-resolution data to an aggregated level has a limited impact on the estimates, while speeding computation time dramatically. In a future draft, we intend to estimate this and other regressions using the full hourly sample.

trols whatsoever; household-by-year fixed effects; household-by-month-of-year fixed effects; week-of-sample fixed effects; and various combinations thereof. ε_{ith} is an error term, which we two-way cluster at the Census block group and week-of-sample levels.

In addition to this specification, we also produce event study estimates, using the following equation:

$$Y_{ith} = \beta^s \sum_{s=-S}^{s=S} EV_{it} \cdot \mathbf{1}[t = s] + \gamma^s \sum_{s=-S}^{s=S} Solar_{it} \cdot \mathbf{1}[t = s] + \alpha_i + \delta_t + \varepsilon_{ith} \tag{2}$$

This equation is the same as Equation 1 above, except now we plot coefficients for weeks $s \in \{-S, S\}$ separately. This allows us to provide evidence on pre-treatment trends, as well as to trace the effects of EV adoption through time.

Identification The identifying assumption for both of these models is that, conditional on our choice of fixed effects, households that did and did not adopt EVs would have had electricity consumption that was, and would have continued to be, trending similarly in the absence of EV adoption. Figure 5, which estimates a version of Equation 2 supports this assumption: even with no controls, we see that there is no trend in electricity consumption prior to the arrival of an EV. We do see some evidence of noise in our EV dates. It appears that our DMV registration data lags the actual arrival dates somewhat, generating a negative (level) pre-period estimate. However, because there is no trend, this supports our identification assumption.[10]

4.1 Heterogeneity

We are interested in going beyond average treatment effects to understand how EVs affect load shapes, and to compute treatment effects for households on different electricity tariffs. To do this, we extend Equation 2 to allow for heterogeneous treatment effects.

Hourly effects To estimate effects by hour of day, we fully interact all of the terms in Equation 2 with 24 hour-of-day dummies:

$$Y_{ith} = \beta^{s,h} \sum_{s=-S}^{s=S} \sum_{r=0}^{23} EV_{it} \cdot \mathbf{1}[t = s, hour = h] + \gamma^{s,r} \sum_{s=-S}^{s=S} \sum_{r=0}^{23} Solar_{it} \cdot \mathbf{1}[t = s, hour = h] + \alpha_i + \delta_t + \varepsilon_{ith} \tag{3}$$

[10]We are in the process of running specifications where we exclude a few weeks before and after an EV arrives to deal with measurement error.

Figure 5: Event study of EV registration on electricity consumption (limited controls)

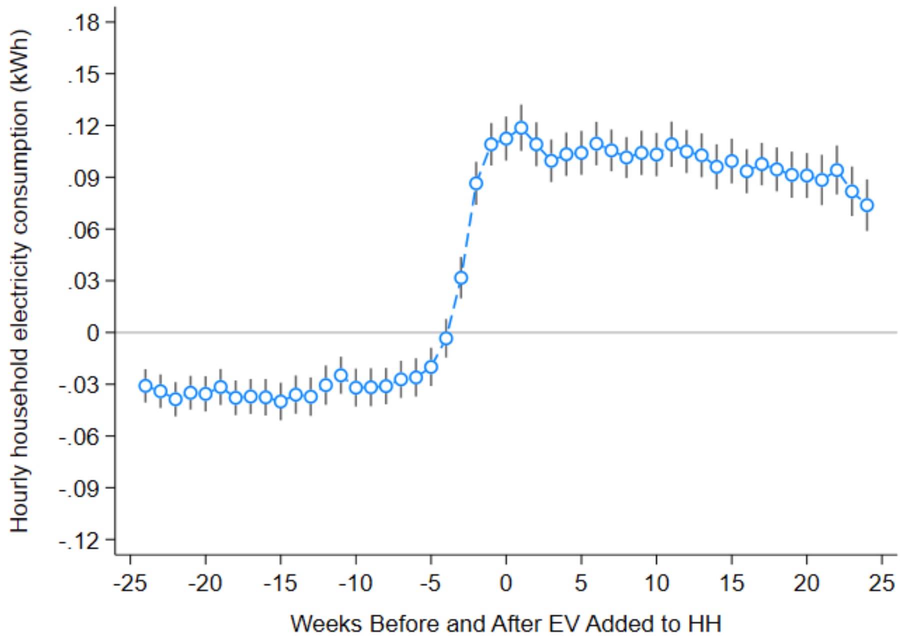

Notes: Each dot reflects the estimated coefficient on an event-time dummy variable, where the EV registration date corresponds to event time -1. The specification includes fixed effects for household. Standard errors clustered at the Census block group level and week-of-sample.

Rate-specific estimates Because households on different electricity tariffs face substantially different pricing schedules, including the potential for time-of-use pricing, we estimate rate-specific effects by, fully interacting all of the terms in Equation 2 with dummies for rate:

$$Y_{ith} = \beta^{s,r} \sum_{s=-S}^{s=S} \sum_{r \in \text{rates}} EV_{it} \cdot \mathbf{1}[t = s, rate = r] + \gamma^{s,r} \sum_{s=-S}^{s=S} \sum_{r \in rates} Solar_{it} \cdot \mathbf{1}[t = s, rate = r] + \alpha_i + \delta_t + \varepsilon_{ith} \tag{4}$$

We interpret these estimates with caution, as they combine selection into these rates with heterogeneous treatment effects. That said, these estimates are potentially informative about policy-relevant quantities, so we include them here.

EV-model-specific estimates Different vehicles have different battery capacities, and are likely to be charged at different times. We explore this through an additional heterogeneous effects approach:

$$
\begin{aligned}
Y_{ith} =& \beta^{s,v} \sum_{s=-S}^{s=S} \sum_{v \in \text{vehicle type}} EV_{it} \cdot \mathbf{1}[t = s, manufacturer = v] \\
&+ \gamma^{s,v} \sum_{s=-S}^{s=S} \sum_{v \in manufacturer} Solar_{it} \cdot \mathbf{1}[manufacturer = v, hour = s] + \alpha_i + \delta_t + \varepsilon_{ith} \qquad (5)
\end{aligned}
$$

Finally, we estimate heterogeneous effects by both rate and hour. To do this, we take our fully interacted samples (Equations3 and 4) and interact them with one another, such that we estimate hour-specific effects for each rate class:

$$
\begin{aligned}
Y_{ith} = & \beta^{s,r,h} \sum_{s=-S}^{s=S} \sum_{r \in \text{rates}} \sum_{k=0}^{23} EV_{it} \cdot \mathbf{1}[t = s, rate = r, hour = k] \\
&+\gamma^{s,r,k} \sum_{s=-S}^{s=S} \sum_{r \in rates} \sum_{k=0}^{23} Solar_{it} \cdot \mathbf{1}[t = s, rate = r, hour = k] + \alpha_i + \delta_t + \varepsilon_{ith} \qquad (6)
\end{aligned}
$$

5 Results

In this section, we present empirical results corresponding to the approaches described in Section 4. We first estimate the average household-level effects of EV adoption on energy consumption, both using a difference-in-differences approach and an event study. We then isolate the load effects by hour of day, and examine heterogeneity in charging patterns on a number of observable dimensions. Finally, by better understanding how load is actually distributed across time, we can take a fresh look at the implications of EV load profiles on marginal damages.

All of the results that follow reflect the overall effect at the household level, and include both changes from EV charging as well as any indirect changes in (non-EV) load patterns that are caused by the introduction of an EV. This interpretation is desirable in some contexts and less so in others. For example, those wishing to forecast the effect of EV-related load on the grid will be interested in the overall effect that includes both direct and indirect components. On the other hand, some may aspire to measure the amount of electricity used for EV charging, which is relevant for (among other things) assigning credits under the Low Carbon Fuel Standards. To use our estimates for the latter purpose requires an assumption that there are no indirect effects of EV adoption on non-EV load, which is an assumption that we are not yet able to test.

There are (at least) two ways that we will attempt to isolate EV load from other contemporaneous factors. First, in what we present here we include controls for solar installation, for which we observe the timing at the household level. While we don't focus on the solar coefficients in this paper, we present them and they may be of independent interest to some readers. Controlling for behind-the-meter solar production is important insofar as there is complementarity between solar and EVs, which is at least one popular narrative (which we will test). Second, our vehicle dataset includes a sample of comparison vehicles that are gasoline powered. In future drafts we will include estimates of the effect of buying non-EV cars on overall household load. While we are optimistic that those two approaches will together help to support the claim that these results primarily reflect direct EV load, we cannot yet rule out the possibility that our results reflect some indirect effects as well.

Table 3 displays the average treatment effect of EV adoption on post-adoption electric load under various specifications. Moving from left to right across the columns reflects increasingly rich controls. The relatively large coefficient in the first column, which has no fixed-effect controls, reflects the fact that EV-adopting households consume more electricity than the average household in our sample. Examining within-household changes via the addition of household fixed effects brings the coefficient into the range of 0.10 to 0.15 kWh per hour across all remaining specifications.

The coefficients can be converted into daily or weekly kWh by multiplying by 24 or 168, respectively. Thus, the coefficient range of of 0.10 and 0.15 translates to between 17 and 25 kWh of net EV charging load per week. To put this into perspective, a 2016 Nissan Leaf (a fully battery-electric vehicle) has a battery capacity of between 24 and 30 kWh and can support a travel range of 84 to 107 miles (source: Edmunds).

Table 3: Difference-in-differences estimates of EV registration effect on electricity consumption

	kWh/hr	kWh/hr	kWh/hr	kWh/hr	kWh/hr	kWh/hr
EV Post	0.25***	0.12***	0.12***	0.10***	0.15***	0.15***
	(0.02)	(0.02)	(0.02)	(0.02)	(0.01)	(0.03)
Solar Post	-0.32***	-0.48***	-0.43***	-0.53***	-0.36***	-0.41***
	(0.03)	(0.04)	(0.03)	(0.03)	(0.03)	(0.03)
Constant	0.80***	0.81***	0.80***	0.81***	0.80***	0.80***
	(0.02)	(0.01)	(0.01)	(0.00)	(0.00)	(0.00)
HH FEs	No	Yes	No	No	No	No
HHxYear FEs	No	No	Yes	No	Yes	Yes
HHxMofY FEs	No	No	Yes	Yes	No	Yes
Week-of-Sample FEs	No	No	No	Yes	Yes	Yes
Observations	70,051,899	70,051,861	70,044,099	70,044,209	70,051,762	70,044,099
Within R^2	0.00	0.04	0.23	0.19	0.08	0.23

Notes: Standard errors clustered by Census block group are shown in parentheses. Significance at the 10%, 5%, and 1% levels are denoted by *,**,***, respectively.

One might wish to rule out the possibility that our difference-in-differences estimates are picking up pre-existing trends in electricity usage that may differ between EV adopters and other households in the sample. To examine this possibility we estimate an event study, the results of which are presented in Figure 6. The event study specification aligns most closely with the controls in the right-most column of Table 3, and the difference between the pre- and post-adoption usage is roughly the same – in the range of 0.12 to 0.15.

Figure 6: Event study of EV registration on electricity consumption

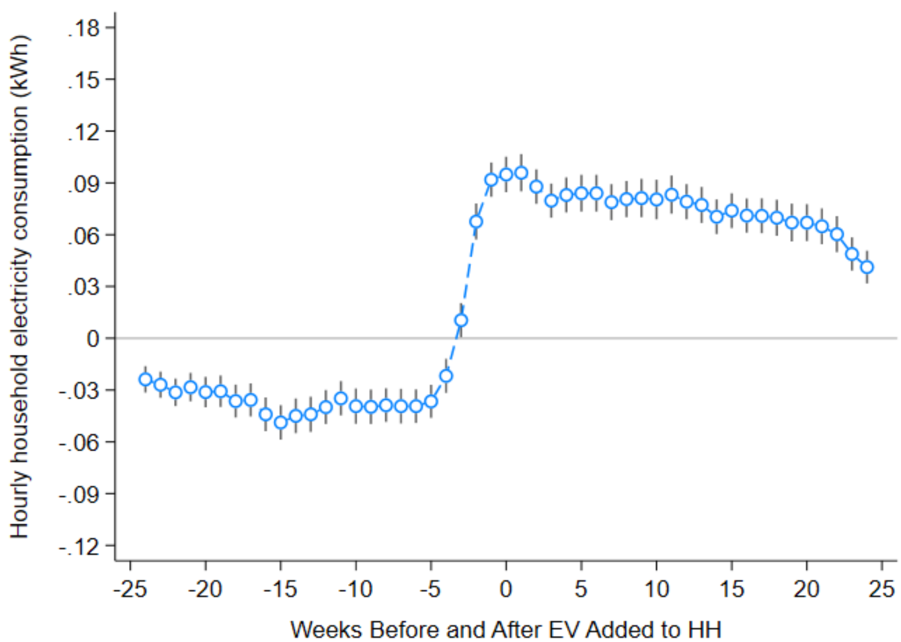

Notes: Each dot reflects the estimated coefficient on an event-time dummy variable, where the EV registration date corresponds to event time -1. The specification includes fixed effects for household-by-month-of-year, household-by-year and week-of-sample. Standard errors clustered at the Census block group level.

Reassuringly, results from this preferred specification are qualitatively and quantitatively similar to those retrieved with limited controls that were shown in Section 4 above. Three features of the event study are worth noting. First, the pre-adoption usage in EV households appears flat, providing evidence that the difference-in-differences estimates are not misinterpreting pre-existing trends. Second, there is measurement error in the registration date variable. The event study uses event-time -1 as the registration week, and yet it is clear that the treatment effect of EV adoption begins to arise 2-4 weeks before then. Finally, there appears to be attenuation of the EV load effect in the post adoption period. There are two potential explanations. Either EV owners are changing the amount or source of their EV charging load, which could manifest either as reduced usage of the EVs over time or the substitution of charging away from the home towards commercial or workplace charging; or there could be some EV-induced change in non-EV household electricity use in these households.

5.1 Placebos on ICE registration events

For the next draft we will run the difference-in-differences and event study specifications on adoption of gasoline-powered cars. This will evidence about the extent to which purchasing *any* car leads to changes in the overall household electricity load profile.

5.2 Timing of EV load & implication for emissions

Figure A1 presents our estimates of how home EV charging is distributed across hours of the day, and how that maps onto the marginal emissions externalities during those hours. EV charging is concentrated during nighttime hours, reaching its maximum at 1am. Throughout the night, EV charging drops each hour before virtually disappearing between 10am and 6pm. Some households appear to plug in their EVs after returning home from work in the early evening, which is ill-timed with respect to scarcity on the California grid. According to California's Independent System Operator, in 2019 the system reached peak load at 5:50pm (on August 15). Any amount of load during or near peak hours will have a disproportionately large effect on the grid.

The concentration of EV load in nighttime hours is advantageous from that perspective. The nighttime charging pattern is primarily determined by pre-programming of the EVs themselves, which are set to begin charging around midnight. Interestingly, and perhaps not surprisingly based on when people are at work, the timing of home EV load is more-or-less non-overlapping with solar production.

The emissions externalities on California's grid also vary by time-of-day. The highest damages occur from electricity produced in the late evening and early hours of the morning. Damages in these hours reflect the possibility that coal and/or gas are the marginal sources of generation.[11] It is quite clear from Figure A1 that the timing of residential EV charging in our sample is highly correlated with the intensity of emissions externalities.

[11] While there is no coal generation in California, the same is not true of the western grid to which California is connected.

Figure 7: Timing of EV load vs emissions factors

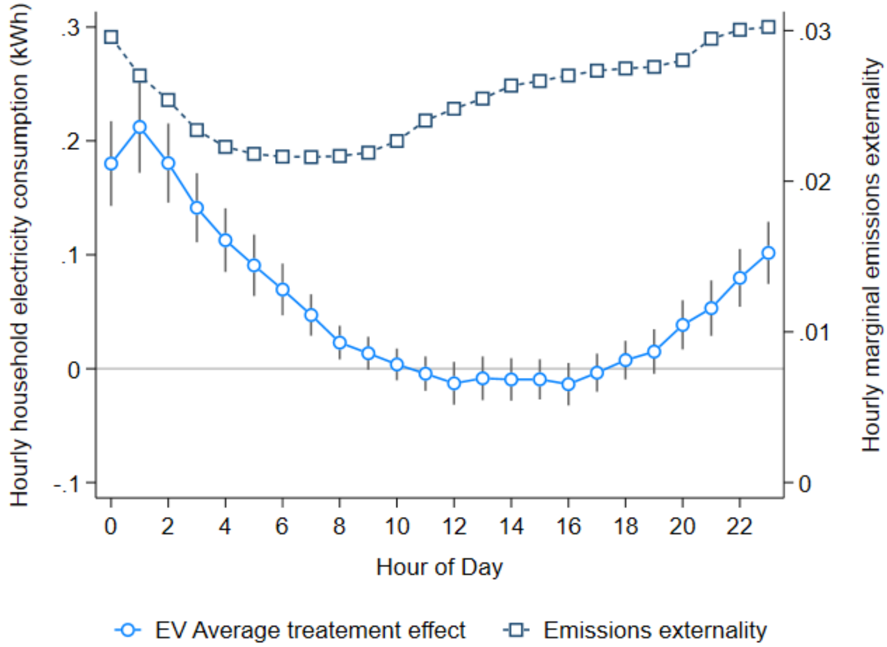

Notes: EV average treatment effects from difference-in-differences specification with household-by-month-of-year, household-by-year and week-of-sample fixed effects, and standard errors clustered at the Census block group level. Emissions externalities estimates sourced from Borenstein and Bushnell (2018).

5.3 Heterogeneity

There are many reasons why one may wish to understanding the heterogeneity of charging intensity and temporal patterns. Following the over-arching motivation of this paper, for example, one may wish to forecast the effect of EV load on the electric grid. However, current EV owners represent a small fraction of total passenger car ownership. To the extent that certain types of EV owners exhibit different charging patterns than others, or that charging patterns differ across electricity rate classes and vehicle models, these may be helpful for forecasting and planning. In this section we present estimates of EV charging by hour-of-day across various dimensions of heterogeneity.

Figure 8 shows how charging patterns differ across electricity rate types. The solid line represents estimated EV load for customers on the EV tariff. This tariff is time-varying by hour-of-day, but is not subject to the increasing block tiers that are built into the non-EV tariff. It is apparent that the level of charging under the EV rate is significantly higher than that of charging under the flat rate, but that the proportion of this charging in a given hour of the day is roughly the same.

Figure 8: EV charging patterns by electricity tariff type:
EV vs non-EV rates

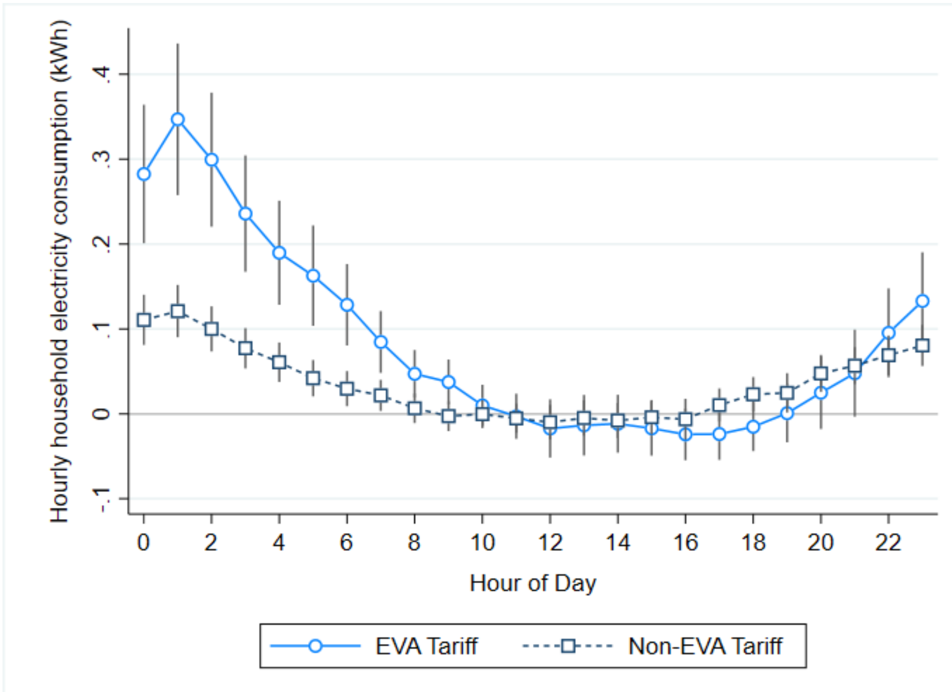

Notes: EV average treatment effects from difference-in-differences specification with household-by-month-of-year, household-by-year and week-of-sample fixed effects, and standard errors clustered at the Census block group level.

Figure 9 breaks out load effects by popular EV types. Perhaps not surprisingly, Tesla charging represents significantly more electricity load than the Nissan Leaf, Chevrolet Volt, or all other EVs. Teslas consume roughly two times the average baseline effect estimated in our difference-in-differences, which translates to an average of roughly 50 kWh per week in home charging.

Overall, these load estimates are roughly half the size of those used by the state for forecasting and policy-making. A 2018 California Energy Commission report (**??**) shows residential EV demand to be between 0.25 and 0.41 kWh per hour. This is 1.6 to 2.7 times our average treatment effect estimates. Moreover, the aggregate level of charging assumed by CEC also appears to be high. Separate CEC forecasts assume that EVs consume roughly 80 kWh per week in electricity from all charging sources, or 0.48 kWh per hour.[12] While some of this will be served by non-residential charging, this number still appears high relative to our findings. There are several potential explanations. It is possible that EVs in our sample are not being driven as

[12]Interestingly, total EV electricity consumption is assumed to vary based only on the number of EVs on the road, and not as a function of vehicle fuel economy or intensity of vehicle use. This is likely due to a paucity of empirical estimates of the price elasticity of demand for eVMT (vehicle-miles traveled in an EV).

Figure 9: EV charging patterns by popular EV model

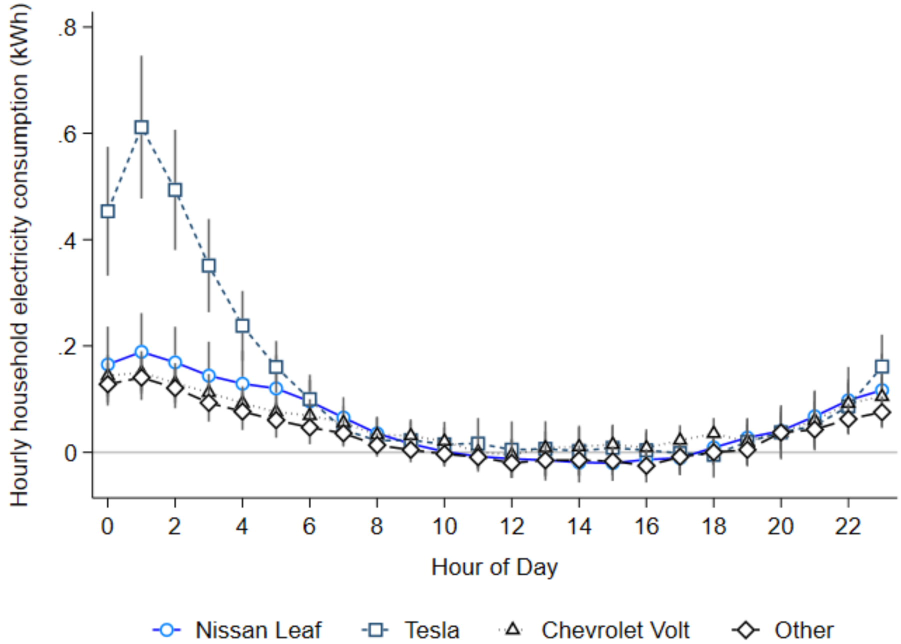

Notes: EV average treatment effects from difference-in-differences specification with household-by-month-of-year, household-by-year and week-of-sample fixed effects, and standard errors clustered at the Census block group level.

much as the average car assumed in the CEC forecast. It may also be the case that a higher proportion of charging for the cars in our sample is being met by workplace or commercial chargers. While evidence that we hope to be able to cite soon implies that the latter is unlikely, at this point we cannot rule out or measure the extent of either of these factors.

6 Conclusion

In this paper, we combine extremely rich data on household electricity consumption – over 1.7 billion observations of hourly electricity use in PGE's service territory – with spatially resolved data on electric vehicle ownership – address level DMV registration information for every EV in California – from 2014 to 2017 to estimate the effects of household EV adoption on electricity load.

We document several key findings. First, on average, EV adoption increases household load by 0.15 kWh per hour, an increase of approximately 20 percent over the mean consumption in a non-EV-owning household. However, these treatment effects are substantially smaller than the expected future change in load used by state agencies: our treatment effects are roughly half the size of the forecasted increase in load in California Energy Commission projections.

Second, these increases in load are not uniform throughout the day. In particular, EV charging occurs largely during the late night and early morning hours, with the bulk of charging taking place between 10 PM and 4 AM. This is important for two reasons: the shape of load has an important role to play in determining future investment into the grid (*e.g.* an all-solar grid will be ineffective at supplying a midnight peak); and the marginal emissions on the grid vary across times of day. Though the emissions-time gradient in California is less steep than other places in the United States, EV charging is still occurring during what are by far the dirtiest hours of the day.

Third, we find substantial heterogeneity in EV charging by electricity rate and model of car. All customers exhibit similar hourly patterns, but the level of our estimated treatment effects differ substantially across groups. Customers on the most popular EV tariff, EV-A, have much larger treatment effects than other households not on an EV tariff. This is likely a combination of both selection and treatment effects; in ongoing work, we aim to use variation in the timing of tariff changes to isolate these. In addition, we see heterogeneity in levels of charging by vehicle: as expected, Tesla owners use substantially more electricity than all other vehicles.

This paper is, to our knowledge, the first to empirically evaluate the relationship between electric vehicle adoption and electricity use *in situ*. These early results suggest that EVs are likely to have a large impact on the future of California electricity consumption, though this impact may be smaller than previously thought. Our average estimates also mask substantial heterogeneity. In ongoing work, we are studying the role of policy in shaping the effects of EV adoption on electricity consumption.

References

Archsmith, James, Alissa Kendall, and David S Rapson, "From Cradle to Junkyard: Assessing the Life Cycle Greenhouse Gas Benefits of Electric Vehicles," *Research in Transportation Economics*, 2015, *63* (3), 397–421.

Borenstein, Severin and James B Bushnell, "Do Two Electricity Pricing Wrongs Make a Right? Cost Recovery, Externalities, and Efficiency," Technical Report, National Bureau of Economic Research 2018.

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

— **, James Bushnell, Frank Wolak, and Matthew Zaragoza-Watkins**, "Expecting the Unexpected: Uncertainty and Environmental Market Design," *American Economic Review*, 2019, (Forthcoming), 191–234.

Electric, Pacific Gas &, "Joint IOU Electric Vehicle Load Research Report," Technical Report, California Energy Commission 2019.

Electricity and Natural Gas Demand Forecast

***Electricity and Natural Gas Demand Forecast**, Technical Report, California Energy Commission 2018.*

Holland, Stephen, Erin Mansur, Nicholas Muller, and Andrew Yates, *"Decompositions and Policy Consequences of an Extraordinary Decline in Air Pollution from Electricity Generation," Technical Report, National Bureau of Economic Research 2018.*

Ito, Koichiro, *"Do consumers respond to marginal or average price? Evidence from nonlinear electricity pricing," American Economic Review, 2014, 104 (2), 537–63.*

Muehlegger, Erich and David S Rapson, *"Subsidizing Mass Adoption of Electric Vehicles: Quasi-Experimental Evidence from California," Technical Report, National Bureau of Economic Research 2018.*

A Appendix

Figure A1: EV contribution to total load

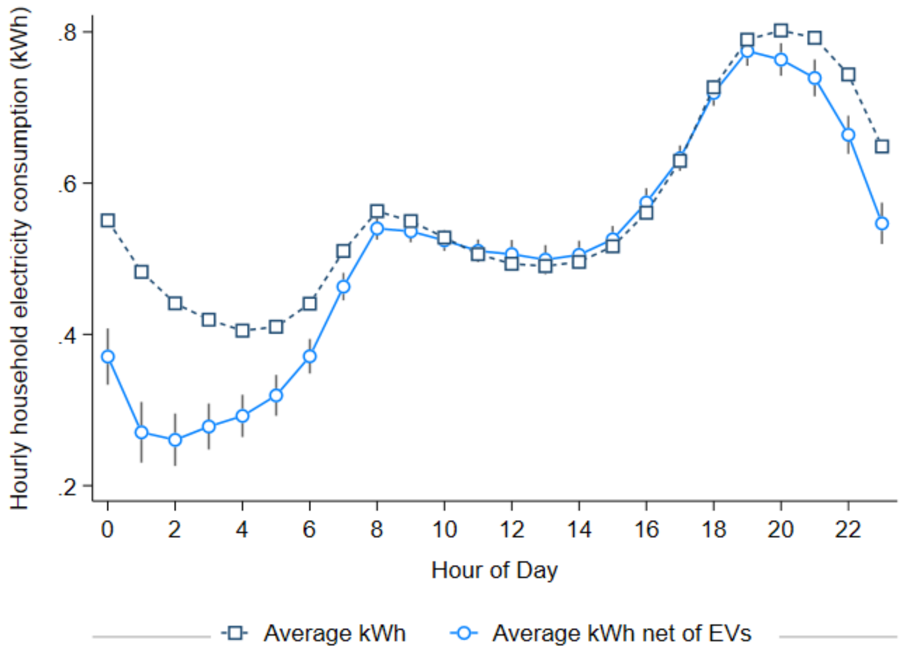

Notes: EV average treatment effects from difference-in-differences specification with household-by-month-of-year, household-by-year and week-of-sample fixed effects, and standard errors clustered at the Census block group level.