

NBER WORKING PAPER SERIES

MEASURING THE ENVIRONMENTAL BENEFITS OF ELECTRIC VEHICLES
(RELATIVE TO THE CAR THAT WASN'T BOUGHT)

Erich Muehlegger
David S. Rapson

Working Paper 27197
<http://www.nber.org/papers/w27197>

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
May 2020

We gratefully acknowledge research funding from the State of California Public Transportation Account and the Road Repair and Accountability Act of 2017 (Senate Bill 1) via the University of California Institute of Transportation Studies. We thank Ben Dawson and Shotaro Nakamura for excellent research assistance. The views expressed herein are those of the authors and do not necessarily reflect the views of the National Bureau of Economic Research.

NBER working papers are circulated for discussion and comment purposes. They have not been peer-reviewed or been subject to the review by the NBER Board of Directors that accompanies official NBER publications.

© 2020 by Erich Muehlegger and David S. Rapson. All rights reserved. Short sections of text, not to exceed two paragraphs, may be quoted without explicit permission provided that full credit, including © notice, is given to the source.

Measuring the Environmental Benefits of Electric Vehicles (Relative to the Car that Wasn't Bought)

Erich Muehlegger and David S. Rapson

NBER Working Paper No. 27197

May 2020

JEL No. Q48,Q52,Q58,R4

ABSTRACT

The true net environmental benefit of an electric vehicle (EV) is measured relative to the vehicle that an EV buyer would have bought and driven had they not opted for an EV. This “counterfactual” vehicle cannot be observed, but its fuel economy can be estimated. We use quasi-experimental variation in a generous California EV subsidy program to show that buyers of EVs would have, on average, purchased relative fuel-efficient gasoline-powered cars had they not gone electric. The true incremental pollution abatement arising from the EV is thus substantially smaller when compared to this appropriate reference vehicle, as opposed to, say, the average new passenger car.

Erich Muehlegger
University of California, Davis
Economics Department
One Shields Ave
Davis, CA 95616
and NBER
emuehlegger@ucdavis.edu

David S. Rapson
Department of Economics
University of California, Davis
One Shields Avenue
Davis, CA 95616
dsrapson@ucdavis.edu

1 Introduction

Policy-makers, in the U.S. and around the world, view adoption of electric vehicles (EVs) as central to addressing urban air pollution, lowering carbon emissions and reducing overall petroleum consumption. The long-run vision is one of an electrified transportation sector powered by clean, renewable energy and achievable through a combination of generous subsidies and government mandates. The setting for this paper, California, provides one such example. California is at the forefront of the EV transition in the United States.¹ Still, this transition remains in the early stages. Fewer than 5 percent of the California light-duty vehicle fleet is electric and, according to the California Energy Commission, roughly 27 percent of California's electricity is obtained from renewable sources. Both of these are increasing rapidly. California's Governor Brown issued an Executive Order B-16-12 that calls for 1.5 million zero-emission vehicles statewide by 2025 as part of a goal to reduce petroleum use in cars and trucks by 50 percent by 2030. Moreover, California's Renewable Portfolio Standard calls for 50 percent of electricity to be generated from renewable sources by 2030. Senate Bill 100 will, if passed, mandate a path to a 100 percent renewable grid.

The extent to which electric vehicle adoption helps to achieve environmental policy goals depends on the emissions of electric vehicles *relative* to the alternative vehicle. Intuitively, a household that switches to an electric vehicle will generate a larger environmental benefit if it replaces a gas guzzler than if it replaces a gas sipper. Importantly, it is insufficient to simply observe what car is sold or retired when a household buys an EV. What matters is the fuel efficiency of the car that *would have been purchased* had the household chosen a conventional vehicle instead.²

In this paper, we seek to estimate the fuel economy of the cars that are displaced when California households buy EVs. We propose a transparent, quasi-experimental design that allows us to estimate the distribution of fuel economy of cars that would have been purchased if an EV had not. Our approach builds on previous work that documents a causal increase in EV purchases induced by a major California EV subsidy policy called the Enhanced Fleet Modernization Program (EFMP).³ The pilot program began in July 2015 - and over the first two years

¹In 2020, roughly half of all EVs on the road in the US are in California.

²Although not the focus of this paper, there are also important considerations on the supply side. If EV penetration proceeds far in advance of electric grid transformation, the maximum potential environmental benefits will be capped as a result of powering transportation with fossil fuels. Furthermore, as Jenn et al. (2016) notes, interactions between policies may offset the environmental benefits of electric vehicles by allow automakers more flexibility to sell less efficient vehicles.

³Muehlegger and Rapson (2018)

allocated roughly \$72 million in state funding. The eligibility rules of this program expose some households to thousands of dollars of subsidies for EV purchases, while withholding those subsidies from others in a way that allows for a “treatment/control” comparison. This allows us to compare what the distribution of vehicles purchased in locations in which the EFMP induced people to adopt fuel-efficient electric vehicles relative to locations where consumers were ineligible for the program.

Every estimate of the pollution abatement benefits of EVs either implicitly or explicitly assumes a displaced, or “counterfactual”, vehicle. Some use stated preference data (e.g. Holland, Mansur, Muller and Yates (2016)), while others may assume a particular model pair for comparison (e.g. Archsmith, Kendall and Rapson (2015)). We are aware of only one other working paper that attempts to empirically estimate attributes of the vehicles displaced by EV adoption (Xing, Leard and Li (2019)). Our paper is complementary to Xing et al. (2019) as ours relies on a different source of variation, deploys a different empirical estimation strategy, and focuses on a different sub-population (middle-income and low-income households).

We compare the distribution of vehicle fuel economy purchased by households living in EFMP-eligible zip codes and EFMP-ineligible zip codes and find evidence that: (1) the program does increase the average fuel efficiency of the set of vehicles purchased in a zip code, but that (2) subsidy participants would have purchased relatively fuel-efficient vehicles in the absence of the program. These results suggest that, in the absence of the program, households would have purchased vehicles with an average fuel economy of roughly 35 miles per gallon, substantially higher than the average fuel economy of light-duty passenger vehicles purchased in California over our sample period. This finding substantially moderates the potential air emissions or fuel savings benefits of incentives relative to what might be calculated if the “reference” vehicle were, for example, assumed to have a fuel economy equal to the average new car.

This project contributes a datapoint on an important policy parameter to the literature, and speaks to policy in several ways. First and foremost, this project provides empirical evidence on the environmental benefits of an electric vehicle incentive program targeting low-income and middle-income households. Although historical adoption of electric vehicles has been heavily tilted toward high-income households, visions of electrifying the transportation sector require substantial adoption by low- and middle-income households. This paper provides some of the first evidence as to the environmental benefits of encouraging adoption amongst these households. Accurately quantifying these benefits is central to achieving California’s

statewide climate change mitigation goals for 2030 and beyond. In addition, the demand side substitution away from conventional vehicles has important fiscal implications. Like many other states, gasoline consumption per capita in California has declined over the past decade. Absent recent state fuel tax increases, tax revenues and fuel consumption are comparable to levels a decade ago. A widespread transition to electric vehicles would accelerate a transition away from fuel taxes as a primary means of support for infrastructure investment.⁴ Our work directly speaks to whether EVs are more likely to displace relatively high (or low) fuel economy vehicles which contribute less (or more) to fiscal coffers on a per-mile basis, and to the speed at which the transition to alternative methods of financing infrastructure might need to occur.

2 Data

We obtained data, covering 2015 to 2017, from a third-party data source that aggregates information about all vehicles purchased and registered within California. We observe total monthly vehicle purchases made by consumers (i.e., excluding fleet or government purchases) in each of California’s roughly 8,000 census tracts at the make-model-model-year level (e.g., 2013 Toyota Camry). We match vehicles by make, model and model-year to fuel economy data from www.fueleconomy.gov⁵ and construct the distribution of fuel economy ratings for purchased vehicles for each census-tract-month during our study period.


As expected, the distribution of fuel economy, plotted in Figure 1 is heavily skewed - the vast majority of vehicles purchased during our study period are those with combined fuel economy ratings between 20 and 30 miles per gallon. This is partially driven by the fact that our vehicle purchase data includes both new and used vehicles, and used vehicles tend to have lower fuel economy.

Figure 2 plots the quarterly trend in both CA vehicle sales (in millions of units on the primary axis) and mean fuel economy (in miles per gallon on the secondary axis). Although sales did not meaningfully change over the study period, average fuel economy rose modestly, as

⁴Davis and Sallee (forthcoming) document that electric vehicles have led to a \$250 million decrease in gasoline tax revenues, and go on to explore the optimal electric vehicle mileage tax.

⁵We string-match our transaction data to make and model makes from the EPA fuel economy rating database in three steps. First, we match models names across the two datasets, conditional on fuel and engine type. For vehicles that do not match, we then remove drive-train information from both datasets (e.g., AWD), again conditioning on fuel and engine type. We hand-code the remaining matches. Where multiple trim lines exist, we assign the fuel economy rating of the most popular trim line for the model. But, in a few instances, due to the coarse nature of the data on engine type (e.g., both plug-in and conventional hybrids are classified as HEVs), there are a few instances where a single model might offer versions with different engine technologies. We examine the robustness of our results to assumptions about the fuel economy of these models in our appendix table and find little evidence that assumptions about this small subset of vehicles meaningfully impacts our conclusions.

Figure 1: **MPG distribution, Light-duty vehicles purchased in CA, 2015-2017**



Note: The figure plots the number of units sold in California between 2015 and 2017, binned by fuel economy (measured in miles per gallon).


consumer preferences shift towards higher fuel economy vehicles and new, more fuel efficiency engine technologies (e.g., battery-electric vehicles) grow in popularity.

3 Program Details

To examine fuel economy of the displaced vehicles, we leverage variation from the EFMP pilot. We have previously examined the impact of the pilot program on the adoption of electric vehicles in Muehlegger and Rapson (2018). We briefly summarize the important details of the program here.⁶ The EFMP offers incentives to low- and middle-income households to replace current vehicles with cleaner and more fuel efficient vehicles. During our study period, EFMP-eligibility was limited to two pilot regions: San Joaquin Valley and the South-Coast Air Quality Management District (i.e., roughly the Los Angeles metropolitan area). Within these locations, eligibility was limited to households with incomes $< 400\%$ of the federal poverty line with lower income households receiving more generous incentives. In addition, households within “disadvantaged communities” (DACs) as defined by the CalEPA EnviroScore, were eligible for more generous incentives.

⁶For more detail, please see Muehlegger and Rapson (2018).

Figure 2: Sales and mean fuel economy of purchased vehicles

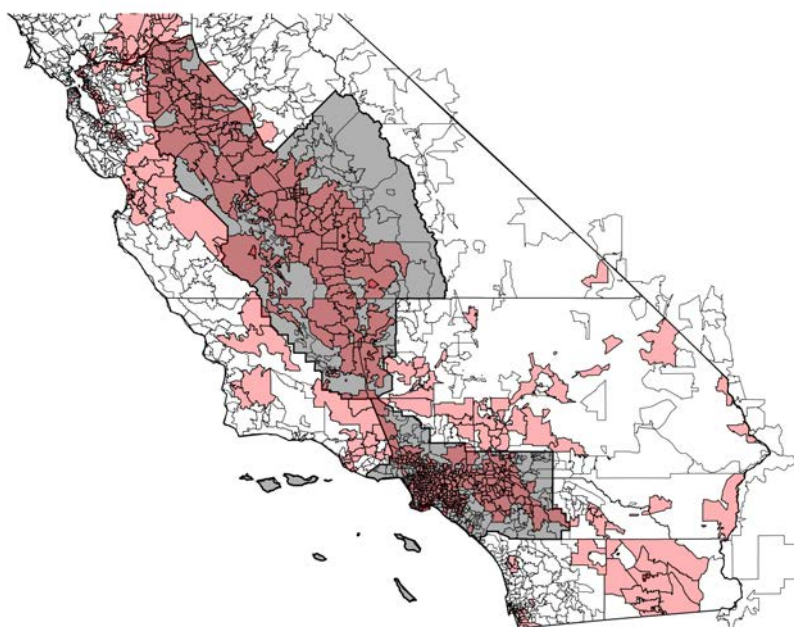


Note: The figure plots quarterly sales in California (on the primary axis, in millions of units) and the mean fuel economy of those vehicles (on the secondary axis, in miles per gallon).

Figure 3 maps zip code boundaries for the Southern two-thirds of California. Regions in grey are the San Joaquin Valley and South Coast AQMDs, the two AQMDs that piloted the EFMP over our study period. The zip codes in pink are those that contain a disadvantaged census tracts. Thus, means-tested households in zip codes that are both grey and pink would be eligible for the subsidy. Outside of the grey and pink boundaries of the two participating AQMDs and disadvantage zip codes, households would be ineligible.

Although the EFMP offered large incentives (up to \$9,500 for low income households living in disadvantaged zip codes to purchase battery electric and plug-in hybrid vehicles), smaller incentives were also offered if qualifying households purchased a “conventional” hybrid vehicle. From publicly available data on program participation, we observe 2,474 participants in the program. Table 1 reports the number and average fuel economy (in gallons per mile or GPM-equivalent) of the battery-electric vehicles (BEV), a plug-in hybrid vehicles (PHEV) and hybrid vehicles (HEV) purchased through the program. Hybrid vehicles were the most popular category purchased by program participants - roughly 45% of participants choose a hybrid vehicle. Roughly 38% chose a plug-in hybrid - the remainder chose battery electric vehicles. The average fuel economy (in gallons per mile) of the different vehicles lines up as expected.

Figure 3: Map of EFMP pilot regions and Disadvantaged Zip Codes



Note: Figure maps zip code boundaries in California. Regions in grey are the two EFMP pilot regions (i.e., San Joaquin Valley and South Coast Air Quality Management District). Zip codes in pink are zip codes that partially or wholly contain a census tract classified as disadvantaged by the CalEPA EnviroScore 2.0.

Hybrid vehicles are the “least” fuel efficient of the three (at 0.0269 GPM or roughly 37 miles per gallon), although they are still substantially more fuel efficient than the average conventional vehicle. Plug-in hybrids are slightly more fuel efficient at close to 40 miles per gallon equivalent while battery electric vehicles are the most fuel efficient at over 100 miles per gallon equivalent. For reference, we also report the average fuel economy of all 18.6 million vehicles purchased by consumers in California between 2015 and 2017. All of the vehicles purchased under the program are substantially more fuel efficient than the average vehicle, with a fuel economy of roughly 22 miles per gallon.

Table 1: Vehicle Fuel Efficiency (2015-2017)

	Number of Vehicles	Mean GPM
<u>EFMP Replacement Vehicles</u>		
BEV	431	.0093
HEV	1106	.0269
PHEV	937	.0249
All EFMP Vehicles	2474	.0231
All CA Vehicle Transactions	18.6 M	.0452

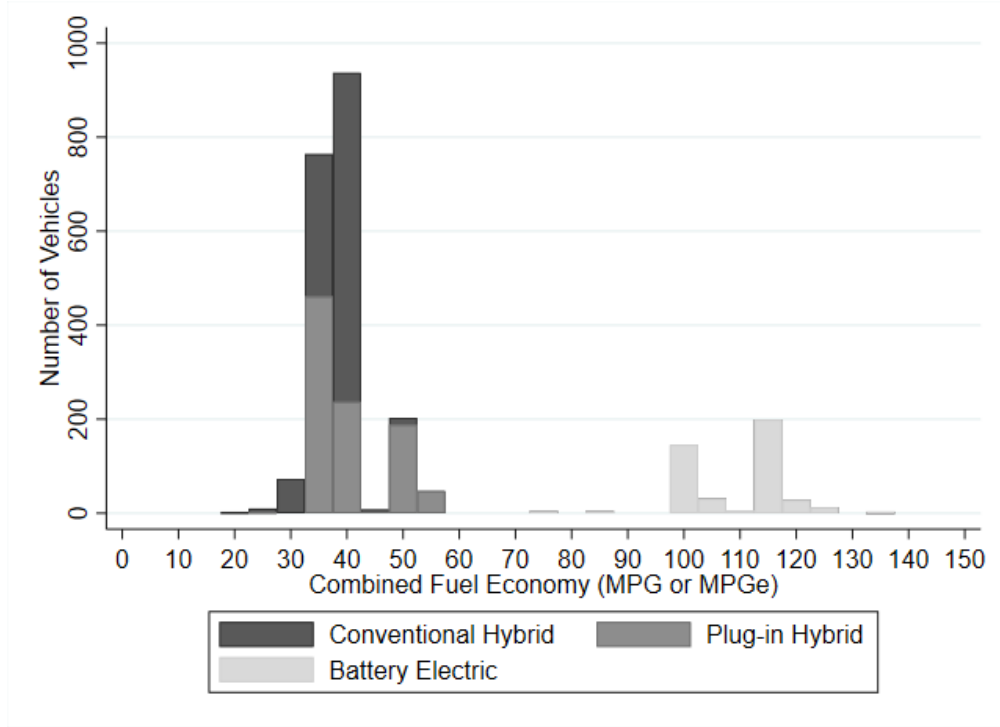
Notes: calculated from EFMP participant records, fleet data and the fueleconomy.gov website.

Average fuel economy by engine type mask substantial variation within group. We plot the distribution of vehicle fuel economy for the three groups of vehicles in Figure 4. For ease of interpretation, we group vehicles by their combined miles-per-gallon rating (or MPGe rating for battery electric vehicles). While plug-in hybrid vehicles are more fuel efficient on average than conventional hybrid vehicles, the fuel economy distributions of the two groups of vehicles overlap substantially. In contrast, battery electric vehicles are substantially more fuel efficient than all plug-in and conventional hybrid vehicles.

Since EFMP program eligibility is set at the zip-code level, we aggregate our census tract-level sales up to the zip-quarter level, and follow the rules of the EFMP program to determine whether households in a particular zip code are eligible for the most generous incentives. Specifically, if a zip code: (a) contains partially or wholly a census tract classified as disadvantaged by the CalEPA EnviroScore 2.0, and (b) is located in either the San Joaquin Valley Aid District or the South Coast Air Quality Management District, we classified the zip code as eligible for the EFMP Plus-Up program. According to program rules, zip codes that do not contain a disadvantaged census tract or are outside of the pilot region are classified as not eligible for the program.

The gradual shift in fuel economy noted statewide in Figure 2 is present in both the EFMP-

Figure 4: Fuel Economy of EFMP Vehicles

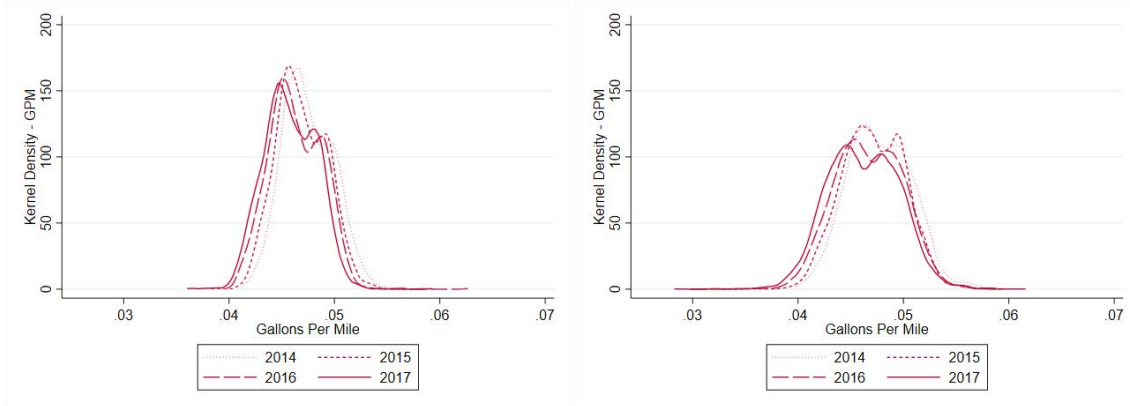


Note: Figure plots the number of vehicles purchased using a Plus-Up incentives offered through the EFMP pilot program through the end of 2017, classified by fuel economy rating (in MPG or MPGe) and by the type of engine technology.

Figure 5: Distribution of GPM for “Treated” and “Untreated” zip codes

(a) Treated Zip Codes

(b) Untreated Zip Codes



Note: The figures plot the annual distributions of fuel economy ratings, measured in gallons per mile (or gpm equivalent), for census tracts eligible for the EFMP Plus Up Incentive (left) and the census tracts ineligible for the EFMP Plus Up incentive (right).

eligible and ineligible zip codes. Figure 5 plots the annual distributions of vehicle fuel efficiency (measured in gallons per mile or the equivalent thereof) for the “treated” locations (i.e.,

disadvantaged zip codes in the pilot regions) and “control” zip codes. The left panel of Figure 5 shows how the distribution of GPMs of the transacted vehicles has shifted downward from 2014 to 2017 in the treated zip codes (DAC = 1, AQMD = 1). The right panel of Figure 5 shows the same thing but for untreated zip codes (DAC = 0 or AQMD = 0). This highlights one advantage of our empirical approach. As average fuel economy is changing throughout California over this period if we were to simply compare mean fuel economy in EFMP-eligible zip codes before and after the introduction of the program, we might misattribute the general trend towards more fuel efficient vehicles to the EFMP program. Rather, by comparing similar locations inside and outside of the pilot regions, we can control for state-wide trends in purchases and better isolate effect of the program as distinct from shifting preferences for Californians.

4 Empirical Approach

Our empirical strategy is based on the implementation rules for the EFMP pilot program. To be eligible for the most generous subsidies, applicants must live within DACs within the pilot program regions. The eligibility rules allow us to compare locations that are similar with respect to demographic characteristics, pre-program adoption patterns and their level of “disadvantage” as measured by the Cal EPA EnviroScore, but which differ in eligibility by whether or not they are located within the two pilot air districts.

We compare the average fuel economy of vehicles purchased in disadvantaged zip codes inside and outside of the two participating AQMDs, before and after the start of the EFMP pilot program. We regress the average fleet fuel economy in a zip-month on λ_{zt} , the fraction of purchases that receive an EFMP subsidy relative to all vehicle transactions in the zip-quarter. We include zip-code fixed effects to control for time-invariant zip code preferences related to fuel economy and region-time fixed effects to control for unobservable changes in the preferences for fuel economy inside and outside the pilot regions.

Formally, we estimate

$$Y_{zt} = \beta_1 \lambda_{zt} + \nu_{tA} + \gamma_z + \epsilon_{zt} \quad (1)$$

where our dependent variable Y_{zt} is the average fuel economy of vehicles purchased by households in zip z in month t . In the tables below we scale up λ_{zt} from fractions to percentage points, and interpret the coefficient β_1 as changes in GPM by a one-percentage-point increase in the share of transactions that received the EFMP plus-up subsidy.

Our primary parameter of interest is the fuel efficiency of the marginal vehicle *replaced* as a result of the incentive. From our estimate of the impact of the incentive on average fuel efficiency, we can back out an estimate of the fuel efficiency of the vehicle which would have been purchased, absent the incentive. To illustrate the intuition, suppose we observe two zip codes that are identical but for their EFMP eligibility status. If the average fuel efficiency in the EFMP-eligible zip code increases significantly when an electric vehicle is purchased through the program relative to the fuel efficiency in a zip code in which households are not eligible, it suggests that the electric vehicle is substantially different than the vehicle that would have been purchased absent the subsidy. In contrast, if the average fuel efficiency does not differ between the two locations when electric vehicles are purchased through the EFMP program, it suggests that the subsidized vehicle was similar in fuel efficiency to the vehicle that would have otherwise been purchased absent the subsidy. To estimate the fuel efficiency of the marginal replacement vehicle, we subtract our estimated coefficient from the average fuel efficiency of the vehicles chosen under the program.⁷ For each regression specification, we report the fuel economy rating (in MPG) of the marginal replacement vehicle in at the bottom of each result column.

4.1 Alternative specifications

Whether the empirical strategy above provides an accurate estimate of the impact of the EFMP program on the average fuel economy in a location depends on an assumption that that fuel economy in locations that avail themselves of the program would have followed a similar path absent the program to locations that either did not participate or were not eligible. Evidence in favor of this assumption can be found in parallel trends in fuel economy across treated and control zip codes before the introduction of the program. If treated and control locations differed along some dimensions that affect trends in fuel economy, we might misattribute ex post differences in fuel economy to the program. Thus, we consider three additional specifications: (1) a “matched” specification, (2) a triple-differences specification, and (3) a within-zip code specification.

In the “matched” specification, we use nearest-neighbor matching (Abadie and Imbens (2006)) to pair disadvantaged zip codes in participating AQMDs with “control” disadvantaged zip codes in non-participating AQMDs, where we match zip codes by pre-period trends in GPM. Although all disadvantaged zip codes have somewhat different socio-demographics

⁷We derive this relationship in the appendix.

than non-disadvantaged zip codes, the eligible set of DACs might differ along dimensions that are correlated with trends in fuel economy preferences. The “matched” specification places additional weight on zip-code that have more similar trends in fuel economy in the pre-program period.

The “triple-differenced” specification includes as an additional set of controls the non-DACs in the pilot and non-pilot regions. To the extent that fuel economy in pilot and non-pilot regions were trending differentially over time, the additional set of non-DACs in and out of the pilot regions allow us to control for those differences. Formally, the specification includes a full-set of interaction fixed effects, ν_{tA} , ϕ_{tD} and γ_z , capturing shocks common to the pilot region, shocks common to DACs and time-invariant zip-level differences.

$$Y_{zt} = \beta_1 \lambda_{zt} + \nu_{tA} + \phi_{tD} + \gamma_z + \epsilon_{zt} \quad (2)$$

Finally, we consider a “within-zip” specification that focuses only on the treated zip codes and identifies the coefficients from quarter-by-quarter variation the number of people who receive subsidies within each “treated” zip code. As above, we include zip code fixed effects, the coefficient estimate identified from a comparison of a given zip code in a quarter in which many people purchase a vehicle through the EFMP program to that same zip code in a quarter in which few people receive a subsidy.

5 Results

We present our results in Table 2. Column (1) corresponds to our primary “difference-in-differences” specification from equation (1). Columns (2) through (4) correspond to our alternative “matched”, “triple-differenced” and “within-zip code specifications. In panel A, the explanatory variable of interest is the fraction of transactions (in percentage points) in a zip*quarter purchased via an EFMP subsidy. Across all columns, we estimate negative coefficients that are within a narrow range. These coefficients are consistent with the subsidized vehicles being more fuel-efficient than the average vehicle purchased in a zip-quarter. Interpreting the coefficient in column (1), if 10 percent of vehicles in a zip-quarter receive an EFMP subsidy, the mean fuel efficiency (measured in gallons per mile) decreases by -0.00055 gallons pre mile. This is equivalent to increasing the average fuel economy of purchased vehicles (with a mean fuel economy of roughly 23 percent) to an average of roughly 23.3 miles per gallon.

The relatively modest increase is reflective of two features of the EFMP program. The first is

that roughly 80 percent of the vehicles purchased under the EFMP program during this period were plug-in or conventional hybrid vehicles. Although both are more fuel efficient than the fleet average, they are significantly less fuel efficient than battery electric vehicles, which were purchased much less frequently under the program.

The second relates to what vehicle would have been purchased absent the incentive. By combining our estimates in Panel A with information about the average fuel efficiency of vehicles purchased under the EFMP program, we can obtain an estimate the fuel economy of the counterfactual vehicles that would have been purchased absent the incentive. To be clear, we don't observe this hypothetical purchase in our data; rather, we observe the purchases and average fuel economy in locations that are similar to the EFMP-eligible zip codes (i.e., also classified as "disadvantaged" by CalEPA), but were not eligible for the EFMP incentive by virtue of being located outside the pilot regions. Under the assumption that the similar, but ineligible location is reflective of what might have happened absent the EFMP program, we can back-out an estimate of the fuel economy of the vehicle that would have been purchased absent the subsidy. We report the fuel economy of the "marginal replacement" vehicle at the bottom of panel A.

Absent the subsidy, our results suggest that individuals who purchased vehicles during the study period would have tended to purchase relatively fuel efficient vehicles, with fuel economy averaging roughly 35 miles per gallon as alternatives. This moderates the impact of the EFMP program on environmental outcomes and fuel consumption, as it suggests program participants would have purchased "gas-sippers" rather than "gas-guzzlers" in the absence of the program.

5.1 Estimates by vehicle type

In this section we explore the possibility that the alternative vehicle that would have been chosen absent the incentive might vary by type of EV purchase. That is, the (unobserved) counterfactual vehicle may be different for households that purchase a BEV, households purchasing a PHEV, and households purchasing a conventional HEV. We directly test this by estimating separate coefficients for each engine type. Specifically, we estimate the average fuel efficiency in a zip-month as a function of the share of transactions receiving EFMP plus up, by each engine type (BEV, PHEV and HEV). As we did above, we then subtract the coefficients from the average fuel efficiency of the each engine type to estimate the fuel efficiency of the replaced vehicle.

Table 2: **Effect of EFMP Incentives on fleet GPM, main results**

Panel A: Impact of fraction of vehicles subsidized on average GPM				
	(1) DinD	(2) Matched DinD	(3) Triple Diff	(4) Within Treatment Zip
% point EFMP Transactions	-0.000055** (0.000024)	-0.000045* (0.000025)	-0.000050** (0.000025)	-0.000055** (0.000024)
Observations	9230	12246	21294	6643
R-Squared	0.88	0.92	0.89	0.89
Marg. repl. MPG	35.0	36.3	35.7	35.0

Panel B: Impact of average subsidy on average GPM				
	(1) DinD	(2) Matched DinD	(3) Triple Diff	(4) Within Treatment Zip
Avg. PU Subsidy	-0.0016*** (0.00055)	-0.0013** (0.00057)	-0.0015** (0.00058)	-0.0016*** (0.00055)
Observations	9230	12246	21294	6643
R-Squared	0.88	0.92	0.89	0.89

Dependent variable is average GPM in a zip*quarter. Standard errors are clustered by zip code. Columns 1, 2, 3 and 4 reported regression results for the unmatched Differences-in-Differences, the matched Difference-in-Differences, the Triple-difference and the within-zip code specifications, respectively. In panel A, the explanatory variable of interest is the fraction of transactions (in percentage points) in a zip*quarter purchased with an EFMP subsidy. In Panel B, the explanatory variable of interest is the average plus-up subsidy (in thousands of dollars) received by buyers across all vehicles purchased in a zip*quarter.

We present the estimates of the specification in 3. Unsurprisingly, we find that battery electric vehicles purchased under the EFMP program have the largest impact on fleet fuel economy. This is consistent with Figure 4, which shows that BEVs are substantially more fuel-efficient than either plug-in or conventional hybrid vehicles. Despite these considerations, we do not find strong differences when we estimate the marginal vehicles that would have been purchased in the absence of the program.

5.2 Implications for Emissions Abatement

The true emissions abatement contribution of an EV purchase is a function of three main factors: 1) the difference between the fuel economy of the EV and the car that otherwise would have been purchased, 2) the amount the car will be driven (vehicle-miles traveled, or VMT), and 3) the upstream emissions profile of the electricity generators that power a plug-in EV. In this section, we use our estimates of counterfactual vehicle fuel economy to calculate the implied CO2 abatement under two different electricity grid generation profiles. Throughout, we hold VMT constant at the California state average (14,453 miles per driver per year).⁸

Roughly speaking, the marginal source of electricity generation in much of the US is an effi-

⁸U.S. Federal Highway Administration

Table 3: **Effect of EFMP Incentives on fleet GPM , by vehicle type**

Panel A: Impact of fraction of vehicles subsidized on average GPM				
	(1) DinD	(2) Matched DinD	(3) Triple Diff	(4) Within Treatment Zip
% point EFMP BEV Transactions	-0.00018** (0.000085)	-0.00011 (0.000072)	-0.00018** (0.000085)	-0.00018** (0.000085)
% point EFMP PHEV Transactions	-0.000044 (0.000048)	-0.000035 (0.000037)	-0.000029 (0.000049)	-0.000044 (0.000048)
% point EFMP HEV Transactions	-0.000011 (0.000042)	-0.000023 (0.000042)	-0.000011 (0.000042)	-0.000011 (0.000042)
Observations	9230	12246	21294	6643
R-Squared	0.88	0.92	0.89	0.89
MR MPG BEV	37.1	48.9	37.1	37.1
MR MPG PHEV	34.2	35.2	36.0	34.2
MR MPG HEV	35.7	34.2	35.7	35.7

Panel B: Impact of average subsidy on average GPM				
	(1) DinD	(2) Matched DinD	(3) Triple Diff	(4) Within Treatment Zip
Avg. PU BEV Subsidy	-0.0037** (0.0018)	-0.0024 (0.0015)	-0.0037** (0.0017)	-0.0037** (0.0018)
Avg. PU PHEV Subsidy	-0.0011 (0.00097)	-0.00091 (0.00072)	-0.00078 (0.00099)	-0.0011 (0.00097)
Avg. PU HEV Subsidy	-0.00032 (0.0017)	-0.00080 (0.0017)	-0.00035 (0.0017)	-0.00032 (0.0017)
Observations	9230	12246	21294	6643
R-Squared	0.88	0.92	0.89	0.89

Dependent variable is average GPM in a zip*quarter. Standard errors are clustered by zip code. Columns 1, 2, 3 and 4 reported regression results for the unmatched Differences-in-Differences, the matched Difference-in-Differences, the Triple-difference and the within-zip code specifications, respectively. In panel A, the explanatory variable of interest is the fraction of transactions, by engine type, (in percentage points) in a zip*quarter purchased with an EFMP subsidy. In Panel B, the explanatory variable of interest is the average plus-up subsidy, by engine type, (in thousands of dollars) received by buyers across all vehicles purchased in a zip*quarter.

cient combined-cycle natural gas generator. It may be dirtier in some places (e.g. the Midwest) and cleaner in others (California and regions with a substantial endowment of hydroelectricity). While several recent papers have estimated the marginal emissions profile of EVs (e.g. Holland, Mansur, Muller and Yates (2016) and Archsmith, Kendall and Rapson (2015)), here we acknowledge that the grid composition is changing and readers with interests across varying jurisdictions may wish to apply different grid emissions profiles. We thus use two reference points: one in which EVs are charged by a natural gas plant (0.418 kg of CO₂ per kWh⁹), and a second that uses 50 percent natural gas and 50 percent carbon-free electricity (0.209 kg of CO₂ per kWh). These assumptions conform to actual policy and other contributions to the literature.¹⁰

Table 4 reports estimates of annual CO₂ savings per EV. Comparing results across columns,

⁹U.S. Energy Information Administration

¹⁰The California cap-and-trade program assumes a generic electricity source emits 0.428 kg of CO₂ per kWh (Borenstein, Bushnell, Wolak and Zaragoza-Watkins (2019)).

Table 4: **Implied CO₂ savings (kg per EV per year)**

Electricity generation source	Counterfactual	Reference vehicle	
		Average new car purchased	Average new or used car purchased
Natural gas	198	1,202	1,692
50% natural gas, 50% renewable	1,934	2,938	3,429

Notes: Table reports kg of CO₂ per year abated by EV relative to reference vehicle. Natural gas generation is assumed to emit 0.418 kg of CO₂ per kWh. Purchased EV MPGe is the average of purchased EVs under the EFMP program in our sample. PHEVs are assumed to operate in electric mode for 55 percent of miles traveled (source: U.S. Department of Energy).

one can see the inverse relationship between EV CO₂ savings and fuel economy of the comparison vehicle. Abatement relative to the (fuel-efficient) counterfactual vehicle that we estimate in this paper is substantially lower than when the reference vehicle is less efficient, as would be the case with a reference vehicle reflective of new (or new and used) vehicle purchases over the same period. To the extent EV subsidies are intended to promote CO₂ abatement, the policy implications are profound. Suppose a policy-maker were to assume the reference vehicle is an average new car purchased over our sample period. The true CO₂ savings would be less than one-sixth what may be expected if the grid is powered by natural gas. For cleaner grids, the error would be smaller but still substantial – roughly 33 percent less emissions reductions than anticipated when the grid emissions profile is halfway between natural gas and zero-carbon.

6 Conclusion

In this project we seek to understand what cars are displaced when California households buy EVs. We implement multiple variants on a simple, quasi-experimental design that allows us to estimate the distribution of fuel economy of cars that would have been purchased if an EV had not. The eligibility rules of this program expose some households to thousands of dollars of subsidies for EV purchases, while withholding those subsidies from others. We exploit the pilot nature of the EFMP program to understand what cars are displaced when California households purchase EVs. To do so, we compare the distribution of vehicle fuel economy purchased by households living in EFMP-eligible zip codes and EFMP-ineligible zip codes. Our results are remarkably robust to changes in specifications, with estimates of the replaced vehicle fuel economy falling within a very narrow range.

We find evidence that: (1) the program increases the average fuel efficiency of the set of

vehicles purchased in a zip code, but (2) subsidy participants were most likely to purchase relatively fuel efficient vehicles in the absence of the program. This second point moderates the potential air emissions or fuel savings benefits of EV subsidies. If the “reference” vehicle were assumed to have a fuel economy equal to the average new car purchased, a policy-maker would dramatically overstate the true greenhouse gas savings. The extent of overstatement depends on the composition of electricity generation sources powering the EV. If the EV is charged by an efficient natural gas generator, the implied CO₂ emissions savings would be overstated by roughly six times; if the EV is charged by a 50/50 mix of natural gas and renewables, CO₂ savings would be overstated by fifty percent.

References

- Abadie, Alberto and Guido Imbens**, “Large Sample Properties of Matching Estimators for Average Treatment Effects,” *Econometrica*, 2006, 74, 235–267.
- Archsmith, James, Alissa Kendall, and David 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, James Bushnell, Frank Wolak, and Matthew Zaragoza-Watkins**, “Expecting the Unexpected: Policy Choice and Emissions Market Design,” *American Economic Review*, 2019, 109 (1), 3953–3977.
- Davis, Lucas W and James M Sallee**, “Should Electric Vehicle Drivers Pay a Mileage Tax?,” *NBER Environmental and Energy Policy and the Economy*, forthcoming.
- Holland, Stephen, Erin Mansur, Nicholas Muller, and Andrew Yates**, “Are There Environmental Benefits from Driving Electric Vehicles? The Importance of Local Factors,” *American Economic Review*, 2016, 106 (12), 3700–3729.
- Jenn, Alan, Inês ML Azevedo, and Jeremy J Michalek**, “Alternative fuel vehicle adoption increases fleet gasoline consumption and greenhouse gas emissions under United States corporate average fuel economy policy and greenhouse gas emissions standards,” *Environmental science & technology*, 2016, 50 (5), 2165–2174.

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

Xing, Jianwei, Benjamin Leard, and Shanjun Li, “What Does an Electric Vehicle Replace?,” Technical Report, National Bureau of Economic Research 2019.

A Appendix

A.1 Estimating the marginal replacement vehicles

We report rough estimates of the fuel efficiency of the marginal replacement vehicle, using the potential outcome framework with never-takers (NT), compliers (C), and always-takers (AT). We consider this framework in the context of the intent-to-treat in with the difference-in-difference specification. Defining z_i as the intent to treat indicator, and D_i to be the treatment “take-up”, i.e. the adoption of EVs. Given that the observation i is at the zip-year level, we interpret that within each observation, a fraction of it has $D_i = 1$, and 0 for the rest. (Does this even make sense?) The fractions of populations within zip-quarter that fall into each of the treatment- and assignment blocks are shown as follows:

Table A.1: Potential treatment framework: Share of population within each treatment-assignment (ITT) status

	$z_i = 0$	$z_i = 1$
$D_i = 0$	$\pi^{NT} + \pi^C$	π^{NT}
$D_i = 1$	π^{AT}	$\pi^{AT} + \pi^C$

Our outcome of interest is the fuel efficiency measure (average-gallons per mile), which will be expressed as a function of D_i .

$$E[Y_i|z_i = 1] = \pi^{NT}E[Y(0)|NT] + \pi^CE[Y(1)|C] + \pi^{AT}E[Y(1)|AT] \quad (3)$$

$$E[Y_i|z_i = 0] = \pi^{NT}E[Y(0)|NT] + \pi^CE[Y(0)|C] + \pi^{AT}E[Y(1)|AT] \quad (4)$$

The ITT estimate should therefore give us

$$E[Y_i|z_i = 1] - E[Y_i|z_i = 0] = \pi^C(E[Y(1)|C] - E[Y(0)|C]) \quad (5)$$

In our specifications, we scale the treatment by the zip-year level realizations of π_i^C . Given that the share of compliers are small and has minimal effect to the shares of always-takers and never-takers, and assuming that there are no heterogeneous treatment effects by the realized complier shares, z_i is 0 if in the non-assigned zip-years and π_i^C for assigned zip-years

$$E[Y_i|z_i = \pi_i^C] - E[Y_i|z_i = 0] = \pi_i^C(E[Y(1)|C] - E[Y(0)|C]) \quad (6)$$

So the coefficient on the share variable estimates $(E[Y(1)|C] - E[Y(1)|C])$. To estimate the replacement vehicle fuel efficiency, i.e. $E[Y(0)|C]$, we subtract the replaced vehicles' average fuel efficiency $(E[Y(1)|C])$ from the coefficient of π_i^C .

A.2 Endogeneity and Instrumental Variables

Given that the denominator of the main regressor λ_{zt} is all vehicle transacted including ICEs transacted in the zip-quarter, rather than EVs only, we believe that the endogeneity issue is less of a concern. We still address this issue by constructing an instrument for the EFMP-share (and an analogous instrument for average subsidies). We instrument for EFMP-share (over all types of vehicles transacted) using the count of EFMP transactions in zip z at time t , normalized by average total (including ICE) number of transactions in that zip in all quarters except the current one, scaled by the ratio of sales in all other zips in AQMD*DAC in the current time period relative to others.

Formally, denoting the number of post-period quarters as T , the quarter in which the EFMP program becomes active as t^* and average number of transactions in zip z in quarter t as $Q_{zt} = \sum_i \mathbf{1}(\text{zip} = z, \text{time} = t)$, we construct the instrument for EFMP-share as follows, where Q is now total transactions including ICEs:

$$IV_{zt} = \frac{\sum_i \mathbf{1}(\text{Subsidy}_{izt} > 0, \text{zip} = z, \text{time} = t)}{\frac{\sum_{r \neq t, r \geq t^*} Q_{zr}}{T-1} \frac{\sum_{x \neq z} Q_{xt}}{\sum_{r \geq t^*} \sum_{x \neq z} Q_{xr} / T - 1}} \quad (7)$$

Table A.2: Effect of EFMP Incentives on fleet GPM, main results

	Panel A: Impact of fraction of vehicles subsidized on average GPM					
	OLS			IV		
	(1) DinD	(2) Matched DinD	(3) Triple Diff	(4) DinD	(5) Matched DinD	(6) Triple Diff
% point EFMP Transactions	-0.000055** (0.000024)	-0.000045* (0.000025)	-0.000050** (0.000025)	-0.000065*** (0.000024)	-0.000046* (0.000025)	0.00029 (0.00046)
Observations	9230	12246	21294	9230	12246	21294
R-Squared	0.88	0.92	0.89	166824.7	160295.8	128.7
F				33.8	36.2	-171.7
Marg. repl. MPG	35.0	36.3	35.7			

	Panel B: Impact of average subsidy on average GPM					
	OLS			IV		
	(1) DinD	(2) Matched DinD	(3) Triple Diff	(4) DinD	(5) Matched DinD	(6) Triple Diff
Avg. PU Subsidy	-0.0016*** (0.00055)	-0.0013** (0.00057)	-0.0015** (0.00058)	-0.0019*** (0.00058)	-0.0013** (0.00059)	0.0070 (0.011)
Observations	9230	12246	21294	9230	12246	21294
R-Squared	0.88	0.92	0.89	13635.1	13344.5	253.1
F						

Dependent variable is average GPM in a zip*quarter. Standard errors are clustered by zip code. Columns 1, 2 and 3 are OLS regressions for the unmatched Differences-in-Differences, the matched Difference-in-Differences and the Triple-differenced specifications, respectively. Columns 4, 5 and 6 present IV estimates of columns 1 through 3 using the instrument.

A.3 Robustness checks

Due to the level of aggregation of our data, we need to make several assumptions when mapping make-model-model-years to fuel economy ratings.

In our data, we cannot definitively identify from our data whether BMW I3s, that are technically classified in our data as battery-electric vehicles, have "range-extender options" that allow for a gasoline engine to recharge the battery. The range extender option provides functionality very similar to a plug-in hybrid vehicle (with commensurately lower fuel economy).

In addition, our data contains information on fuel type, (gasoline, diesel, hybrid, and electric), but does not offer distinctions between PHEVs and HEVs. This is an issue for the following models (model years in parentheses):

- Ford C-Max (2013-17)
- Ford Fusion (2013-18)
- Honda Accord (2014)
- Hyundai Sonata (2016-17)
- Toyota Prius (2012-15)

We test the robustness of our analysis to each of these groups of vehicles for which we cannot definitively assign fuel economy ratings. For the former (the BMW I3), we run specifications excluding the I3 from the fleet and find little impact on our results, consistent with the modest market share of the vehicle. For the latter group of vehicles, for which we do not know whether a given vehicle is a plug-in or conventional hybrid, we perform a bounding analysis, first under the assumption that all of the vehicles have the most "optimistic" fuel economy rating (as if all the vehicles sold were PHEVs) and then under the most "pessimistic" assumptions (that all the vehicles sold were conventional HEVs). We present our main specification under all three assumptions below.

Table A.3: Effect of EFMP Incentives on fleet GPM, dropping BMW I3

Panel A: Impact of fraction of vehicles subsidized on average GPM				
	(1) DinD	(2) Matched DinD	(3) Triple Diff	(4) Within Treatment Zip
% point EFMP Transactions	-0.000057** (0.000023)	-0.000052** (0.000022)	-0.000052** (0.000024)	-0.000057** (0.000023)
Observations	9230	12194	21294	6643
R-Squared	0.88	0.93	0.89	0.89
Marg. repl. MPG	34.7	35.4	35.4	34.7

Panel B: Impact of average subsidy on average GPM				
	(1) DinD	(2) Matched DinD	(3) Triple Diff	(4) Within Treatment Zip
Avg. PU Subsidy	-0.0017*** (0.00056)	-0.0015*** (0.00050)	-0.0015*** (0.00058)	-0.0017*** (0.00056)
Observations	9230	12194	21294	6643
R-Squared	0.88	0.93	0.89	0.89

Dependent variable is average GPM in a zip*quarter. Standard errors are clustered by zip code. Columns 1, 2 and 3 are OLS regressions for the unmatched Differences-in-Differences, the matched Difference-in-Differences and the Triple-differenced specifications, respectively. Column 4 presents the results restricting the sample to eligible zip codes (i.e., DACs within the relevant pilot regions). We dropped BMW I3 from the sample, because we could not be confident as to which of the transactions in our data had range extenders (which would make it a PHEV, and has significantly worse implied fuel efficiency than the pure BEV version).

Table A.4: Effect of EFMP Incentives on fleet GPM, with “optimistic” fuel type assumption

Panel A: Impact of fraction of vehicles subsidized on average GPM				
	(1) DinD	(2) Matched DinD	(3) Triple Diff	(4) Within Treatment Zip
% point EFMP Transactions	-0.000055** (0.000024)	-0.000044* (0.000025)	-0.000050** (0.000025)	-0.000055** (0.000024)
Observations	9230	12246	21294	6643
R-Squared	0.88	0.92	0.89	0.89
Marg. repl. MPG	35.1	36.3	35.7	35.1

Panel B: Impact of average subsidy on average GPM				
	(1) DinD	(2) Matched DinD	(3) Triple Diff	(4) Within Treatment Zip
Avg. PU Subsidy	-0.0016*** (0.00055)	-0.0013** (0.00057)	-0.0015** (0.00058)	-0.0016*** (0.00055)
Observations	9230	12246	21294	6643
R-Squared	0.88	0.92	0.89	0.89

Dependent variable is average GPM in a zip*quarter. Standard errors are clustered by zip code. Columns 1, 2 and 3 are OLS regressions for the unmatched Differences-in-Differences, the matched Difference-in-Differences and the Triple-differenced specifications, respectively. Column 4 presents the results restricting the sample to eligible zip codes (i.e., DACs within the relevant pilot regions). The dependent variable for the treated group is constructed such that for make-model years for which we are uncertain about the engine type (PHEV or HEV, BEV or PHEV) we take the more fuel efficient type.

Table A.5: Effect of EFMP Incentives on fleet GPM with “pessimistic” fuel type assumption

Panel A: Impact of fraction of vehicles subsidized on average GPM				
	(1) DinD	(2) Matched DinD	(3) Triple Diff	(4) Within Treatment Zip
% point EFMP Transactions	-0.000055** (0.000024)	-0.000045* (0.000025)	-0.000049** (0.000025)	-0.000055** (0.000024)
Observations	9230	12246	21294	6643
R-Squared	0.88	0.92	0.89	0.89
Marg. repl. MPG	35.0	36.3	35.7	35.0

Panel B: Impact of average subsidy on average GPM				
	(1) DinD	(2) Matched DinD	(3) Triple Diff	(4) Within Treatment Zip
Avg. PU Subsidy	-0.0016*** (0.00055)	-0.0013** (0.00057)	-0.0015** (0.00058)	-0.0016*** (0.00055)
Observations	9230	12246	21294	6643
R-Squared	0.88	0.92	0.89	0.89

Dependent variable is average GPM in a zip*quarter. Standard errors are clustered by zip code. Columns 1, 2 and 3 are OLS regressions for the unmatched Differences-in-Differences, the matched Difference-in-Differences and the Triple-differenced specifications, respectively. Columns 4 presents the results restricting the sample to eligible zip codes (i.e., DACs within the relevant pilot regions). The dependent variable for the treated group is constructed such that for make-model years for which we are uncertain about the engine type (PHEV or HEV, BEV or PHEV) we take the less fuel efficient type.

A.4 Geographical units

One challenge in constructing the dataset is that policy assignment into the EFMP is done at the zip-code level, while the transaction data is at the Census tract level. In order to reconcile the discrepancies in the geographical units, we construct a merged dataset in which the unit of observation is zip-quarter. This is done by weighting the tract level observations by the share of geographical overlap with a given zipcode. We show the cross tabulation of treatment status at the tract level and that at the zipcode level. We see that there is one tract-zip overlap in which the tract was treated but the zip was not. This is fixed in the analysis data.

Table A.6: Cross-tabulation of tracts to zipcodes

DAC tract	0		DAC zip 1		Total	
	No.	%	No.	%	No.	%
0	7,727.0	55.6	6,178.0	44.4	13,905.0	100.0
1	1.0	0.0	3,834.0	100.0	3,835.0	100.0
Total	7,728.0	43.6	10,012.0	56.4	17,740.0	100.0

Figure A.1 shows the number of zipcodes that a given Census tract overlaps.

Figure A.1: **Number of zip codes within census tracts**

