

Electric Vehicles and the Energy Transition: Unintended Consequences of Time-of-Use Pricing

Megan R. Bailey David P. Brown Erica Myers
Blake Shaffer Frank A. Wolak

March 2025

Online Appendix

Appendices


A Treatment Messaging

This appendix details the communication and in-app experience for each group at the beginning of the treatment (starting on July 5, 2023).

A.1 Control Group

Email Communications: No email correspondence was initiated with participants in the Control group after initial enrollment.

In-App Experience: The Control group participants continued to experience the baseline features of the application. EV owners in this group could only monitor their charging data and schedule their EV charge start time within the App. The Figure below illustrates their in-app experience.



A.2 Time-of-Use (TOU) Group

Email Communications: Participants in the TOU group were sent an email with the subject line “Action needed – earn additional rewards” and a preheader stating “You now earn an extra 3.5 cents/kWh on off-peak home charging in FortisAlberta’s EV Smart Charging Pilot.” The users observed the information provided in the Figure below upon opening the email.

PROGRAM UPDATE!

Please take action to earn more rewards

Congratulations, you have been selected to receive an additional benefit during our FortisAlberta pilot program!

What You Get

3.5 cents for every kWh charged off-peak at home

When You Get It

Beginning July 5, 2023, off-peak hours are defined as:

10 AM - 2 PM
10 PM - 6 AM

What You Should Do

Schedule your home charging to take place off-peak



Perks of this pilot update:



Lower Electricity Bills

You can now save up to \$28/month by charging your EV off-peak



Smaller Carbon Footprint

By charging during off-peak hours, you're helping reduce CO2 emissions



Grid Support

Your EV is helping to contribute to overall grid stability and reliability



Frequently Asked Questions

How do I earn more rewards?

Starting 7/5/23, you will earn 3.5 cents/kWh on all off-peak home charging, up to 800 kWh per month per vehicle. This is in addition to your \$50 sign-up and \$100 program completion rewards.


How do I schedule my EV to charge off-peak?

You may use all available options, including but not limited to: your EV app, Optiwatt, or other 3rd party tools

[View Your Dashboard](#)

[View the full FAQ](#)

In-App Experience: Within the application, this group encountered new messaging explaining the adjusted TOU rate structure in the program card. Additionally, these participants were granted access to activate or deactivate any Optiwatt scheduling functionalities, a feature unavailable during the initial phase of our field experiment.



A.3 Managed Group

Email Communications: Participants in the Managed group received the following email with the subject “Action needed – earn additional rewards” and a preheader that read “You now earn an extra 3.5 cents/kWh on all home charging in FortisAlberta’s EV Smart Charging Pilot.” The users observed the information provided in the Figure below upon opening the email.



Set Departure Times

Sign Up

Add Home

Add EV

Add Utility

PROGRAM UPDATE!

Please take action to earn more rewards

Congratulations, you have been selected to receive an additional benefit during our FortisAlberta pilot program!

What You Get

3.5 cents for every kWh charged at home

What Will Happen

Beginning July 5, 2023, Optiwatt will occasionally adjust your car's charging to accommodate the needs of the grid

What You Should Do

Set up Scheduled Departure in the Optiwatt app, to ensure your car reaches your battery target by your departure time

To charge urgently, opt out in the Charge Forecast of the Home tab, but lose that day's extra 3.5 cents/kWh reward



Perks of this pilot update:



Lower Electricity Bills

You can now save up to \$28/month by shifting your charging schedule



Smaller Carbon Footprint

By charging during off-peak hours, you're helping reduce CO2 emissions



Grid Support

Your EV is helping to contribute to overall grid stability and reliability



Frequently Asked Questions

How will my charging behaviour change?

If you previously scheduled your charging in another app, or didn't schedule your charging at all, Optiwatt will alter your schedule. But rest assured, Scheduled Departure will ensure your range needs are always met.

How do I set up Scheduled Departure in Optiwatt?

Click the button below, or navigate to the Optiwatt app. In the Charging tab, you will see the option to set a Scheduled Departure time for each day of the week. We ensure your EV will reach your Charge Limit by that time each day.

Set Departure Times

[View the full FAQ](#)



HOME | NEWS | FAQs

1540 Market St San Francisco, California 94102

No longer want to receive these emails? [Unsubscribe](#)

6

In-App Experience: This group was presented with augmented messaging within the application, detailing the new incentive and scheduling parameters. The participants in the Managed group were not permitted to disable the Optiwatt scheduling feature, which was enabled to allow managed charging. They were encouraged, though not mandated, to use the Scheduled Departure functionality through notification. If Scheduled Departure was not enabled for a participant, a default time of 8:30 AM was applied to all weekdays.

Schedule

Tesla
2021 Model 3



Add device



Charge Limit: 86% • 333 km



Allow Optiwatt to Schedule Charging



Setting enabled due to pilot. Opt out daily in the Charge Forecast on the Home tab, but you'll lose that day's reward.

Scheduled Departure

Optiwatt may charge during higher electricity rates to ensure your charge limit is reached by departure time



Departure Time 08:30 AM

S M T W T F S

Everyday

Add another departure time

Precondition car

Automatically warm the battery before departure to reduce the effects of cold weather



Climate

Automatically set your cabin's climate to your desired temperature



Set climate to:

< 21°C >



Save the Planet

Charge during lower CO₂ emissions



Save Money

Charge during cheaper rates



Which plan is for you? [Compare plans](#)

EV Smart Charging Pilot

FORTIS
ALBERTA



Receive a \$50 reward after enrollment and a \$100 reward at the end of the pilot after completing the final customer survey

Optiwatt will adjust your home charging, while meeting any departure time you set, and you earn an extra 3.5 cents/kWh on all home charges.

B Attrition

In this appendix, we describe the degree of attrition and explore whether sample attrition could be driving the conclusions we draw in the paper. At the start of the experiment, we randomly assigned 70 cars to the Managed group, 70 to the TOU group, and 62 to the Control group. However, 20 EVs dropped from the Managed group during the post-treatment period. Approximately half of these vehicles, nine in total, lost connection with the Optiwatt App due to “user password errors,” likely caused by a technical issue accessing Tesla’s API. User password errors only occurred in Tesla cars and disproportionately affected cars in the Managed group. Throughout most of the experiment (July-October 2023), Tesla did not support third-party API connections. Optiwatt described the password connection issue as one that arose when their software couldn’t reach Tesla’s system and then after a certain number of tries, it locked the users out and prompted a password reset. Users then needed to reset their passwords in both systems to re-establish the connection, but some did not complete this extra step.

Attrition in the other two groups included 9 EVs from the TOU group, none due to user password errors, and 3 EVs from the Control group, with one due to a user password error. If we compare attrition rates excluding user password errors, there is no statistically significant difference between the TOU and Managed groups, though both have higher attrition rates than the Control group.

The concern with attrition is that drivers who left the experiment might have had different charging behavior and/or different responses to the treatments than those who stayed, potentially affecting the magnitude of the calculated treatment effects. We assess this issue in several ways.

First, Table B1 compares the pre- and post-treatment charging characteristics of cars that completed the experiment with those that dropped out. For the pre-treatment comparison, we used data from the full pre-treatment period. The table shows that the two groups are statistically indistinguishable. This indicates that drivers who left the experiment did not require more total charging, nor did they differ significantly in how much they charged at home or during peak hours. For the post-treatment comparison, we analyzed charging behavior during the first month of the post-treatment period (July), as this month provides the most post-treatment data for the cars that left the experiment at some time post-treatment. Here too the charging behavior of those who left and those who stayed is quite comparable. The total amount charged, the amount charged at home, and charging at peak times are

all statistically indistinguishable between the two groups. This suggests that drivers who left the experiment were not responding differently to the treatment in terms of the amount and timing of charging relative to those who stayed, for the month of July.

Second, we examine treatment effects over time. If drivers who left the experiment had systematically different driving behavior or were differently responsive to treatment than those who stayed, we would expect to see effects either increasing or decreasing over time. Figure B.1 displays the estimated treatment effects for our Constraint Violations (in kWh) for the TOU and Managed groups in peak and off-peak periods. These treatment effects correspond to a variant of regression equation (3) in the main text, adjusted to estimate month-specific treatment effects which replace the indicator variable Post_d with a vector of month-specific indicator variables. Figure B.1 plots the respective coefficients for these interaction terms for months in the post-treatment period. The treatment effects are statistically indistinguishable over time, implying that the drivers who remained in the experiment responded similarly to the treatment as those who were there at the beginning.

Additionally, in Figure B.2 we show the results of a similar analysis as in Figure B.1, focusing on our Charged kWh dependent variable at the vehicle level. Here too the estimated treatment effects are statistically indistinguishable over time. Figures B.1 and B.2 therefore offer further evidence that attrition due to drivers with certain treatment effects or charging patterns disproportionately dropping out of the experiment is unlikely to be driving our results.

Third, we employ a robustness check where we replace the attritted EVs in our sample with a representative Control EV and rerun our analysis as if they had not left our sample. More specifically, for each vehicle that left the experiment (regardless of treatment group), we randomly assigned the charging behavior of an active Control vehicle for each day following its departure. We then re-estimated the results of our main specification. The results of this analysis are presented in Table B2. This analysis yields results that are consistent with our main findings. As expected, with the larger number of EVs leaving the experiment in the Managed group and the inclusion of unmanaged Control EVs in place of the attritted EVs, the estimated treatment effects for the Managed group are attenuated downward. For example, looking at Column (2), the reduction in peak constraint violations decreased to -0.6188 from -0.7201 in our main specification reported in Table 2 in the main text. These results serve as a lower bound for the estimated treatment effects. Despite the

reduction in the estimated effects, we continue to observe similar conclusions as in our main analysis. The Managed group continues to exhibit a significant difference in off-peak violations when compared to TOU.

Finally, as we describe in Section 5.2, in December 2023 we offered the 59 remaining control customers to opt-in to a managed charging program that runs for 6 months. 35 respondents completed the survey and 34 opted into the program. 6 out of 34 (18%) EVs that opted into the Managed program unenrolled. This rate of attrition is comparable to the level observed for those EVs that were automatically enrolled into our initial managed charging treatment. More specifically, 11 EVs actively unenrolled from the initial managed treatment. Removing the 9 EVs that had password errors in the Managed group and left the program (described above), this is an active unenrollment rate of 11 out of 61 (18%).

Table B1. Comparison of Pre-Treatment and Post-Treatment Characteristics: Compliers vs. Non-Compliers

Variable	Pre-Treatment			Post-Treatment		
	Completed	Left	t-test (p-value)	Completed	Left	t-test (p-value)
Home Share (%)	75.79 (23.87)	73.68 (23.84)	0.65	75.79 (28.59)	78.84 (28.96)	0.60
Charge Duration (Minutes)	244.86 (161.65)	260.34 (156.46)	0.61	244.58 (173.14)	229.42 (144.41)	0.62
Energy Charged (kWh)	21.70 (9.43)	25.15 (13.10)	0.16	21.64 (11.00)	21.61 (12.26)	0.99
Max kW Charge (Power)	6.65 (2.36)	7.11 (3.26)	0.45	6.99 (3.14)	8.60 (4.15)	0.06
Off-Peak Share (%)	50.55 (18.52)	47.76 (16.20)	0.39	55.54 (20.75)	54.61 (22.45)	0.84
Off-Peak Share (%) - Home Only	52.36 (21.60)	49.02 (20.45)	0.41	58.89 (24.42)	59.10 (23.12)	0.96
Tesla (%)	85.29 (35.52)	84.38 (36.89)	0.89	85.29 (35.52)	84.38 (36.89)	0.89
Number of EVs	170	32		170	32	


Notes. This table compares pre-treatment and post-treatment (during month of July) average values of various charging variables at the vehicle level separated by EVs that completed the experiment and those that left. Parentheses contain the standard deviations. Home Share represents the percentage of total charging kWh at home, Charge Duration is the daily number of minutes the EV is charged at home, Energy Charged is the kWh charged per day at home, Max KW Charge is maximum power of charge used per day at home, Off-Peak Share is the percentage of kWh charged in the off-peak either at home or away, and Off-Peak Share - Home Only is the percentage of kWh charged in the off-peak at home only. Tesla is the percentage of EVs that are Tesla and Number of EVs is the count of EVs. ANOVA (p-value) reports the p-value from one-way ANOVA tests for differences in means across groups.

Table B2. Estimated Treatment Effects by Group - Attrition Robust


Group	Hours	(1) Charge kWh	(2) Constraint Violations
TOU	Peak	-0.187 (0.039)	-0.749 (0.174)
	Off-Peak	0.227 (0.046)	0.965 (0.276)
Managed	Peak	-0.051 (0.029)	-0.619 (0.195)
	Off-Peak	0.052 (0.040)	-0.083 (0.101)
Treatment Effect Comparison			
TOU - Managed	Peak	-0.136 [0.001]	-0.130 [0.594]
	Off-Peak	0.175 [0.000]	1.047 [0.001]
Observations		1,201,026	127,512

Notes. This table provides the estimated treatment effects for the EV-level dependent variable Charge kWh and the transformer-level dependent variable Constraint Violations (in kWh), using at-home charging only. The estimated treatment effects are separated into Peak and Off-Peak hours. Treatment Effect Comparison compares the treatment effects for TOU and Managed by Peak and Off-Peak, with p-values reported in the brackets for the Wald tests assessing the null hypothesis that the estimated treatment effects for TOU and Managed are equal. All specifications include fixed effects at the day-of-sample, and hour-of-day level. The Column (1) specification includes EV-level fixed effects while the Column (2) includes transformer-level fixed effects. Standard errors (in parentheses) in Column (1) are clustered at the transformer level for vehicles assigned to Managed and the EV-level for vehicles assigned to Control and TOU. Standard errors in Column (2) are clustered at the transformer level.


Figure B.1. Estimated Treatment Effects by Group and Month - Constraint Violations (kWh)




(a) TOU - Peak



(b) TOU - Off Peak




(c) Managed - Peak




(d) Managed - Off Peak

Notes. Upper and lower bars represent the 95% confidence interval for each coefficient.


Figure B.2. Estimated Treatment Effects by Group and Month - Charged kWh




(a) TOU - Peak



(b) TOU - Off Peak



(c) Managed - Peak



(d) Managed - Off Peak

Notes. Upper and lower bars represent the 95% confidence interval.

C Supplementary Tables and Figures

C.1 Extensive Margin

In this section, we focus on data at the vehicle level to evaluate whether EV owners in the TOU or Managed groups differentially adjusted their daily frequency or quantity of charging kWh post-treatment, either at home only or in aggregate (i.e. at home and away charging). By looking at these cases separately, we can evaluate if there was a shift in the location of charging (e.g., from away to home charging) post-treatment for either group.

We estimate the following equation, using all vehicles in our sample, for each day d and vehicle i :

$$Y_{id} = \beta Post_d \times Group_i + \alpha_i + \tau_d + \eta_{id} \quad (1)$$

in which Y_{id} is one of two dependent variables: (1) a charging indicator variable if charging occurred on day d and (2) the Charge kWh is the summation of total charging kWh on day d . $Post_d$ is the post-treatment indicator that equals 1 starting on July 5, 2023, and 0 otherwise, $Group_i$ represents two indicator variables for the TOU and managed treatment groups. α_i is a vehicle-level fixed effect, τ_d is our day-of-sample fixed effect, and η_{id} is the error term. We cluster standard errors at the vehicle level.

We define a “day” between 9:00 AM and 8:59 AM the following day to capture the fact that EV owners systematically make their charging timing decisions in the afternoon/evenings. We consider two specifications where our dependent variables are constructed using at-home charging only and charging both at home and away.

Table C1 presents the results of our extensive margin analysis. We find no evidence of a significant change in charging frequency or charging kWh at the daily level for either treatment group relative to the Control, including at-home-only and both home and away charging. These results indicate that there is no evidence that EV owners responded to either treatment by shifting their charging location and/or aggregate charging patterns at the daily level differentially relative to the Control.

Table C1. Extensive Margin Analysis

	Charging Indicator		Charging kWh	
	Home-Only	Home and Away	Home-Only	Home and Away
TOU $\times$ Post	0.013 (0.027)	0.007 (0.026)	0.378 (0.645)	1.052 (0.815)
Managed $\times$ Post	0.017 (0.026)	0.013 (0.026)	0.218 (0.659)	-0.206 (0.843)
Observations	47,337	47,337	47,337	47,337

Notes. This table provides the estimated vehicle-level treatment effects for equation (1) for the dependent variables Charging Indicator and Charging kWh, using either at-home-only or both home and away charging. All specifications include fixed effects at the vehicle and day-of-sample level. Standard errors are clustered at the vehicle level. Statistical Significance * $p < 0.10$, ** $p < 0.05$, and *** $p < 0.01$.

C.2 Comparison of EVs in Fortis' Territory to EVs in U.S. Metro Areas

In this Appendix, we compare driving and charging patterns in our sample to EVs in the United States. In particular, Optiwatt provided us with charging data from a randomized subsample of EVs in 14 major cities across the United States: Los Angeles, Sacramento, San Diego, San Francisco, San Jose, Las Vegas, Phoenix, Orlando, Miami, Chicago, Houston, San Antonio, Austin, and Seattle. These data cover the pre-treatment period in our sample (April 1, 2023 - July 4, 2023).

First, we focus on EVs in the US sample that were not on a TOU pricing program. This removes the 5 cities in California that are on default TOU programs and EVs in the remaining cities that reported to Optiwatt that they were on a TOU rate. This selection criteria is implemented to compare EVs that are on flat retail rates, as was the case in our Fortis sample pre-treatment.

Table C2 evaluates how our Fortis sample compares to the non-TOU US EVs sample, using data covering our pre-treatment period. This assessment of balance uses the same variables displayed in Table 1 in the main text. While we do observe significant differences for a number of variables, the differences across the two samples are modest. EVs in the US sample charge more at home and a larger percentage in the off-peak. However, they are in comparable ranges. Further, the US sample has a higher proportion of Teslas, but both samples largely consist of Tesla EVs. We observe a comparable amount of daily energy charged at home and max power drawn from the chargers across the two samples. We take these results to demonstrate that while there are differences across the two samples, our Fortis EV sample is not an

outlier compared to EV charging and driving behavior in large US cities.

Table C2. Balance on Observable Characteristics by Sample Using Pre-Treatment Data


Variable	Fortis Sample	U.S. Sample	T-Test (p-value)
Home Share	75.46 (23.82)	83.74 (24.83)	.00
Charge Duration (Minutes)	247.31 (160.56)	279.11 (189.17)	.01
Energy Charged (kWh)	22.25 (10.14)	23.18 (8.93)	.21
Max KW Charge (Power)	6.72 (2.52)	6.80 (3.17)	.67
Off-Peak Share (%)	50.11 (18.17)	54.66 (20.13)	.00
Off-Peak Share (%) - Home Only	51.83 (21.41)	56.43 (22.05)	.00
Tesla (%)	85.15 (35.65)	98.49 (12.20)	.00
Number of EVs	202	1,985	

Notes. This table compares average values of various charging variables at the vehicle level between EVs in the Fortis' territory and 9 metropolitan areas across the United States: Las Vegas, Phoenix, Orlando, Miami, Chicago, Houston, San Antonio, Austin, and Seattle. The U.S. sample excludes EVs on Time-of-Use (TOU) plans and those that never charge at home, focusing on data from April 1, 2023, to July 4, 2023, the same period as the pre-treatment period for the Fortis sample. Parentheses contain the standard deviations. Home Share represents the percentage of total charging kWh at home, Charge Duration is the daily number of minutes the EV is charged at home, Energy Charged is the kWh charged per day at home, Max KW Charge is the maximum power of charge used per day at home, Off-Peak Share is the percentage of kWh charged in the off-peak either at home or away, and Off-Peak Share - Home Only is the percentage of kWh charged in the off-peak at home only. Tesla is the percentage of EVs that are Tesla and Number of EVs is the count of EVs. T-Test (p-value) reports the p-value from t-tests on the equality of means between the two groups.

Second, we are interested in evaluating how the US EVs that report being on a TOU rate in the Optiwatt sample charge their cars at home. Figure C.1 shows the average hourly charging kWh at home on days where charging occurs using all 14 major US cities provided by Optiwatt. For consistency, we focus on our pre-treatment sample period. These descriptive results are consistent with our main findings. EVs on TOU rates in the US sample have the highest average charged kWh arising in the evening off-peak period, with reduced charging in the evening peak. The largest charging kWh occur at midnight in the US sample. This is likely driven by the fact

that many TOU rate structures have the lowest prices starting at midnight, as is the case in California’s EV2 rate (Valdberg et al., 2022).

Figure C.1. Average Charged kWh by Hour in 14 U.S. cities



Notes. This figure presents the mean hourly charging kWh for EVs, including only days on which EVs incurred a positive charge at home, from 14 major cities across the United States: Los Angeles, Sacramento, San Diego, San Francisco, San Jose, Las Vegas, Phoenix, Orlando, Miami, Chicago, Houston, San Antonio, Austin, and Seattle. Data includes charging sessions between April 1, 2023, to July 4, 2023, the same period as the pre-treatment period for the Fortis sample. It only includes vehicles that self-report being on a TOU rate on the Optiwatt app.

C.3 Cluster Robustness

In this Appendix, we consider alternative levels of clustering for our standard errors. Table C3 presents the results for our Charge kWh dependent variable that is at the EV level. Column (1) clusters standard errors at the vehicle level, while column (2) clusters standard errors at the transformer level for vehicles assigned to Managed and the EV level for vehicles assigned to Control and TOU. The latter reflects the level of clustering from our main analysis. These results demonstrate that the standard errors and resulting statistical inference are highly robust to the level of clustering.

Table C4 presents results for our transformer-level dependent variable Constraint Violations (in kWh) with and without wild bootstrap robust standard errors. In our main analysis, for our transformer-level dependent variable Constraint Violations (in kWh), we cluster the standard errors at the transformer level resulting in 21

clusters. We implement the wild bootstrap approach detailed in [Roodman et al. \(2019\)](#) to evaluate the robustness of our results to having a relatively small number of clusters. Table C4 presents the results of our regression analysis with and without wild bootstrap clustered standard errors. Column (1) presents the results from our main analysis, while column (2) implements the wild bootstrap results. Comparing across the two columns, we observe a slight increase in the p-values (in parentheses) that evaluate if the individual treatment effects are statistically significantly different from zero and a corresponding widening of the 95% confidence intervals (in brackets). However, the statistical inference that we draw from our analysis is robust.

Table C3. Estimated Treatment Effects by Group - Charge kWh with Alternative Clustering

Group	Hours	(1)	(2)
		Charge kWh	Charge kWh
TOU	Peak	-0.203 (0.041)	-0.203 (0.041)
	Off-Peak	0.229 (0.047)	0.229 (0.047)
Managed	Peak	-0.048 (0.034)	-0.048 (0.029)
	Off-Peak	0.062 (0.040)	0.062 (0.040)
Treatment Effect Comparison			
TOU - Managed	Peak	-0.155 [0.000]	-0.155 [0.000]
	Off-Peak	0.167 [0.001]	0.167 [0.001]
Observations		1,131,426	1,131,426

Notes. This table provides the estimated treatment effects for the EV-level dependent variable Charge kWh, using at-home charging only. Column (1) clusters standard errors at the vehicle level, while column (2) clusters standard errors at the transformer level for vehicles assigned to Managed and the EV level for vehicles assigned to Control and TOU. The estimated treatment effects are separated into Peak and Off-Peak hours. Treatment Effect Comparison compares the treatment effects for TOU and Managed by Peak and Off-Peak, with p-values reported in the brackets for the Wald tests assessing the null hypothesis that the estimated treatment effects for TOU and Managed are equal. All specifications include fixed effects at the vehicle, day-of-sample, and hour-of-day level.

Table C4. Estimated Treatment Effects by Group - Constraint Violation with Wild Bootstrap Cluster Robust Standard Errors

Group	Hours	(1)	(2)
		Constraint Violations	Constraint Violations
TOU	Peak	-0.772	-0.772
		(0.000)	(0.002)
	Off-Peak	[-1.134, -0.410]	[-1.182, -0.349]
		0.961	0.961
Managed	Peak	(0.002)	(0.007)
		[0.384, 1.538]	[0.305, 1.620]
	Off-Peak	-0.720	-0.720
		(0.001)	(0.001)
Treatment Effect Comparison	TOU - Managed	[-1.087, -0.353]	[-1.147, -0.307]
		-0.146	-0.146
	Off-Peak	(0.145)	(0.176)
		[-0.348, 0.055]	[-0.370, 0.076]
TOU - Managed	Peak	-0.052	-0.052
		(0.826)	(0.824)
	Off-Peak	[-0.534, 0.431]	[-0.567, 0.470]
		1.108	1.108
Observations		(0.001)	(0.003)
		[0.543, 1.673]	[0.469, 1.777]
		127,512	127,512


Notes. This table provides the estimated treatment effects for the transformer-level dependent variable Constraint Violations (in kWh), using at-home charging only. The estimated treatment effects are separated into Peak and Off-Peak hours. Column (1) presents the results from our main analysis with clustered standard errors at the transformer level. Column (2) clusters standard errors at the transformer level, but implements the wild cluster bootstrap detailed in [Roodman et al. \(2019\)](#). p-values are reported in parentheses and confidence intervals are provided in brackets. Treatment Effect Comparison compares the treatment effects for TOU and Managed by Peak and Off-Peak, with p-values reported in the parentheses and confidence intervals are provided in brackets for the Wald tests assessing the null hypothesis that the estimated treatment effects for TOU and Managed are equal. All specifications include fixed effects at the transformer, day-of-sample, and hour-of-day level.

C.4 Constraint Violation Estimates by Constraint Level


In this Appendix, we estimate the hour-specific treatment effects for the Constraint Violation dependent variable, allowing for differential treatment effects by the level of the randomly assigned transformer constraint. We consider three constraint categories: (i) Slack with a capacity limit between 20 - 24 kW, (ii) Moderate with capacity limits 16 - 19 kW, and (iii) Tight with capacity limits 12 - 15 kW. We run our hour-specific constraint violation regression described in Section 5.1 separately for each constraint Category.

Figures C.2 and C.3 present the estimated treatment effects for the TOU and Managed groups, respectively. The results are consistent with those in our main specification shown in Figure 3b in the main text. For the TOU group with slack constraints, while more muted, the pattern of reduced peak period violations and elevated evening off-peak constraint violations persists. These effects increase in magnitude as the virtual transformer constraint becomes tighter leading to more severe violations in the off-peak hours. For the Managed group, the point estimates show a systematic reduction in constraint violations in the evening hours without a corresponding increase in the off-peak period. The magnitude of the reduction in constraint violations increases as the constraints become tighter. This is consistent with the expectation that the pressure on distribution transformers increases as the capacity constraint becomes tighter relative to the underlying non-EV electricity demand.


Figure C.2. Estimated Hourly Constraint Violation Treatment Effects by Constraint Category - TOU Group



(a) Slack




(b) Moderate




(c) Tight

Notes. Upper and lower bars represent the 95% confidence interval.


Figure C.3. Estimated Hourly Constraint Violation Treatment Effects by Constraint Category - Managed Group



(a) Slack



(b) Moderate



(c) Tight

Notes. Upper and lower bars represent the 95% confidence interval.

C.5 Alternative Transformer Grouping

This Appendix considers the effects of alternative transformer groupings for the Control and TOU EVs on our outcome of interest. Specifically, it explores aggregating vehicles based on similar charging habits and EV characteristics, in contrast to the random aggregation approach employed in the main analysis. In reality, EVs on a given transformer may be more homogeneous in their characteristics (and driving/charging habits) because of the observed geographical concentration of EVs (e.g., in high-income neighborhoods, commuting distances, etc.).

It is important to first note that the allocation of EVs in the TOU and Control groups to a specific transformer does not influence their charging behavior, as charging decisions are solely determined by the EV owners. This allows us to arbitrarily regroup these vehicles to examine the effects of different owner groupings on constraint violations. This is in contrast to the Managed group, in which the transformer headroom and the charging behavior of other EVs on the transformer affect charging times post-treatment through the managed charging algorithm.

We take two alternative approaches to group EVs into virtual transformers. First, we compute the average daily charged kWh at-home in off-peak hours pre-treatment. This measure was chosen because it captures both the intensity of at-home charging (in kWhs) and the timing of when a household tends to charge at home. We then rank EVs in the Control and TOU groups separately by this measure and allocate EVs into 10-EV virtual transformers starting with the EVs with the highest value on this charging measure down to those with the lowest. This approach groups the high, medium, and low off-peak charging EVs together in separate virtual transformers. This grouping will be referred to as “Alternative 1”.

Second, we use an adapted kmeans clustering approach using pre-treatment data to group EVs into 10-EV transformers based on several characteristics, including the average daily charged kWh at-home in off-peak hours, the percentage of charging kWhs at-home, the average daily duration of charging (in minutes), estimated EV battery range, and the maximum power of charge used per day at home. This grouping will be referred to as “Alternative 2”.

Kmeans clustering is effective at partitioning the EVs in our sample into groups to minimize the within-cluster differences. However, it does not ensure that the groups are of equal size. As a result, we develop an iterative approach to capture features of kmeans clustering, while allocating EVs into groups of 10. We start the algorithm by clustering EVs in TOU and Control separately using k-means clustering with 7

groups (representing our target number of transformers). For clusters with 10 or more EVs, we calculate the sum of squared errors (SSE) for all EVs in the cluster. For each of these clusters, we select the 10 EVs with the lowest SSE to achieve the most similarity of EVs within the cluster. For the remaining EVs that are not allocated into these 10-EV transformer groups, we rerun the kmeans clustering algorithm and follow the same process of finding groups of 10 EVs. This process continues until we have 10-EV transformer groups for all EVs in the TOU and Control groups.

For both alternative transformer groupings, we assigned the randomized transformer-by-day capacity limits that were used in our experiment. These capacity limits are used to compute the hourly transformer constraint violations with these alternative transformer groups.

Figure C.4 presents the average hourly Constraint Violations (in kWh) for the TOU and Control groups using our baseline transformer grouping from our experiment and Alternative 1 transformers. For the Alternative 1 groups, we observe a modest increase in the average constraint violations for both the Control and TOU groups pre- and post-treatment, compared to our baseline grouping. This is consistent with the fact that, for a subset of the alternative transformers, we observe an increase in coincidental charging leading to more constraint violations, while for others we have groups of EVs with lower overall charging levels leading to lower constraint violations. On average, we observe a small increase in violations for both groups compared to our baseline grouping. We continue to observe similar Constraint Violations for the Control pre- and post-treatment, but a sizable increase for TOU post-treatment.

Figure C.5 presents the corresponding results using Alternative 2 transformer grouping. We do not observe a systematic increase in average constraint violations for either group. This is likely driven by the fact that we are using a wider array of characteristics to group EVs, leading to less homogeneity in coincidental charge timing within a transformer.


Table C5 presents regression results using the specification from our main analysis for the Constraint Violations dependent variable. We adjust our main specification by including only TOU and Control transformer groups to evaluate the impacts of the alternative groupings. Standard errors continue to be clustered at the transformer level. We use the wild bootstrap approach detailed in Roodman et al. (2019) for statistical inference to address concerns over having a small number of clusters. P-values are reported in parentheses and 95% confidence intervals are reported in brackets.

Column (1) provides estimated treatment effects for the TOU group using the

transformer groupings from our experiment to serve as a baseline. Columns (2) and (3) present the estimated results with the two alternative transformer groupings. Despite having allocated EVs based on charging characteristics, we do not observe a large increase in the estimated treatment effects. In fact, the effects are reduced with the Alternative 2 transformers. These results could be driven by the fact that while we observe an increase in Constraint Violations for a subset of the alternative transformers due to more coincidental charging, for others we observe a reduction as we have a group of EVs that charge less often and/or have less charged kWh. In addition, for Alternative 1 we also observe a corresponding increase in the average Constraint Violations for the Control group both pre- and post-treatment, offsetting the increase in average TOU Constraint Violations.


These results suggest that while having more homogeneous EVs on a virtual transformer may increase Constraint Violations for certain transformers, it can reduce violations for others. More broadly, these results demonstrate that the key conclusions drawn from our analysis and sample are robust to alternative transformer groupings for the TOU and Control EVs.

Figure C.4. Average Transformer Violations by Group and Hour - Baseline and Alternative Transformers 1



Notes. Average Transformer Violations represents the average magnitude of hourly distribution transformer constraint violations (in kWh) across all virtual transformers for the pre- and post-treatment periods. Baseline represents the virtual transformers from our main analysis, while Alternative 1 is alternative transformer grouping 1 which groups EVs based on their daily average charge kWhs in off-peak hours. The shaded areas represent our off-peak hours.

Figure C.5. Average Transformer Violations by Group and Hour - Baseline and Alternative Transformers 2



Notes. Average Transformer Violations represents the average magnitude of hourly distribution transformer constraint violations (in kWh) across all virtual transformers for the pre- and post-treatment periods. Baseline represents the virtual transformers from our main analysis, while Alternative 2 is alternative transformer grouping 2 which clusters EVs into groups of 10 based on their similarity across several charging and EV characteristics. The shaded areas represent our off-peak hours.

Table C5. Estimated Treatment Effects for the Alternative Transformer Grouping - Constraint Violation with Wild Bootstrap Cluster Robust Standard Errors

Group	Hours	(1) Baseline Transformers	(2) Alt. Transformers 1	(3) Alt. Transformers 2
TOU	Peak	-0.772	-0.772	-0.695
		(0.002)	(0.000)	(0.000)
		[-1.184, -0.355]	[-1.085, -0.452]	[-1.077, -0.333]
	Off-Peak	0.961	0.977	0.933
		(0.006)	(0.001)	(0.004)
		[0.302, 1.617]	[0.406, 1.559]	[0.338, 1.524]
Observations		85,008	85,008	85,008

Notes. This table provides the estimated treatment effects for the TOU group for the transformer-level dependent variable Constraint Violations (in kWh), using at-home charging only. The estimated treatment effects are separated into Peak and Off-Peak hours. The regression analysis only includes the TOU and Control transformers. Column (1) presents the results using the transformer groupings used in our main analysis, column (2) considers alternative transformer grouping 1 which groups EVs based on their daily average charge kWhs in off-peak hours, and column (3) considers alternative transformer grouping 2 which clusters EVs into groups of 10 based on their similarity across several charging and EV characteristics. Standard errors are clustered at the transformer level with the wild cluster bootstrap procedure detailed in [Roodman et al. \(2019\)](#). p-values are reported in parentheses and confidence intervals are provided in brackets. All specifications include fixed effects at the transformer, day-of-sample, and hour-of-day level.

C.6 Maximum Transformer Demand as the Number of EVs Vary

As the number of EVs on a transformer (N) increase, the maximum demand on the transformer is expected to increase due to periodic coincidental EV charging. However, heterogeneity in charging times across EV owners suggests that the growth in maximum demand should diminish as N increases. To quantify this relationship in our setting, we exploit the fact that EVs in the TOU and Control groups made their charging decisions independently of the virtual transformer assignments in our experiment. This independence enables us to construct counterfactual virtual transformers with varying EV allocations, allowing us to systematically analyze how the number of EVs on a transformer affects its maximum demand—an important input for transformer sizing decisions.


For both the Control and TOU groups, we randomly allocate EVs into virtual transformers of size N , where N varies from 2 to 10.¹ For each transformer, we calculate the maximum hourly demand that arises in the post-treatment period, where hourly demand consists of at-home EV charging from the N EVs on the transformer plus the representative non-EV residential load provided by Fortis. Then, for each treatment group, we compute the average of the maximum hourly demand for the post-treatment period across all of the virtual transformers. We repeat this process 100 times with different random allocations of EVs to new counterfactual virtual transformers of size N . We take the average across all 100 iterations to form an average post-treatment maximum demand on a virtual transformer with N EVs for the Control and TOU groups separately.

Figure C.6 presents the average maximum demand on the virtual transformers in the post-treatment period for the Control and TOU groups using the process outlined above. As expected, we observe an increase in the maximum transformer demand as the number of EVs increase, but it increases at a decreasing rate. The maximum demand on the virtual transformers for the TOU group lies above the Control group. This is consistent with our main findings that EVs in the TOU group systematically shifted their charging to the off-peak hours leading to a rise in the degree of coincidental charging post-treatment. Further, we see that while both groups' maximum demands are increasing at a decreasing rate, there is a divergence between TOU and Control as the number of EVs on the transformer increases. This

¹In our experiment, we have 62 Control and 70 TOU EVs. For a certain N values, it was not possible to allocate all EVs to virtual transformers of size N . For example, when $N = 3$, we could only create 20 and 23 3-EV virtual transformer groups for the Control and TOU groups, respectively. The remaining EVs were dropped.

is also consistent with the higher degree of coincidental charging in the TOU group in the post-treatment period, leading the rate of the increase in the maximum demand to decline more slowly as N increases.

Figure C.6. Average Maximum Demand on Virtual Transformer by Number of EVs



References

- Roodman, D., M. Ø. Nielsen, J. G. MacKinnon, and M. D. Webb (2019). Fast and wild: Bootstrap inference in stata using boottest. *The Stata Journal* 19(1), 4–60.
- Valdberg, A., D. Gomez, E. G. Barnes, and B. Ellis (2022, March). Compliance filing of southern california edison company, san diego gas & electric company, and pacific gas and electric company pursuant to ordering paragraph 2 of decision 16-06-011. Technical report, Public Utilities Commission of the State of California, San Francisco, CA.