Seasonal Savings Impacts for Marin Clean Energy: Winter 2018/19

January 2019 through April 2019

Prepared for Marin Clean Energy

Executive Summary

On January 14, 2019, Nest launched its winter Seasonal Savings algorithm targeted at 31,141 thermostats in the Marin Clean Energy service territory. This report summarizes the estimated heating impacts of Seasonal Savings on these home, The key findings are summarized in the table below.

Summary of Impacts	**Standard**
Participation	31,141 targeted
Qualified on-line, running heating schedule	24,723
% of targeted	79%
Opted In	16,917
% of qualified	68%
Heating Runtime Analysis	
% Heating Savings (2-way fixed effects regression)	7.0% ±2.7%
therms/ opt-in participant	7.2 ±2.8
therms aggregate	118,903
kWh aggregate	149,338

note: ± values are 95% confidence intervals

Intro

Seasonal Savings is a software algorithm that offers customers an opportunity to make their heating schedules more efficient through a series of very small adjustments to the scheduled temperatures over a three week period. The algorithm results in more energy efficient heating schedules going forward. Customers are offered the program on their thermostat and through the Nest phone app and must opt-in to participate.

Participation

A total of 34,601 Nest thermostats were identified in the target population of Marin Clean Energy customers. A control group of 10% of the target population was randomly selected to not receive treatment so that the evaluation could employ a true experimental design, leaving 31,141 thermostats in the target treatment group.

The 31,141 thermostats were located in 24,702 homes -- 22% of customers had more than one thermostat and there were an average of 1.29 thermostats per home

The algorithm was deployed on January 14, 2019 and 79% of the target thermostats qualified to participate (i.e., were online and running a heating schedule) and 68% of the qualified thermostats had the customer opt to participate. In total, 16,917 thermostats participated in the Seasonal Savings event, equal to 54% of the target group.

Methodology

The evaluation employed a true experimental design by including a randomized control group. Because Seasonal Savings is offered as an opt-in service, the analysis followed a Randomized Encouragement Design (RED) -- which compares the entire targeted participant group (including those that did not opt-in) to the entire control group to assess the net impact of being in the target group. An RED eliminates self-selection bias because it directly estimates the impact of being in the target participant group -- not the impact of opting to participate. To estimate the savings per opt-in participant, the RED results can be adjusted by the opt-in rate. For example, if the RED analysis found 2% savings from being in the target group and there was a 50% participation rate then the estimated savings per opt-in customer would be 4% (2% / 50% = 4%).

Runtime Impact Analysis

The heating runtime recorded by the thermostats can also be used to more directly assess the impacts produced by Seasonal Savings. Figure 1 shows the daily heating run time averaged by week for the target participant group and the matched controls.

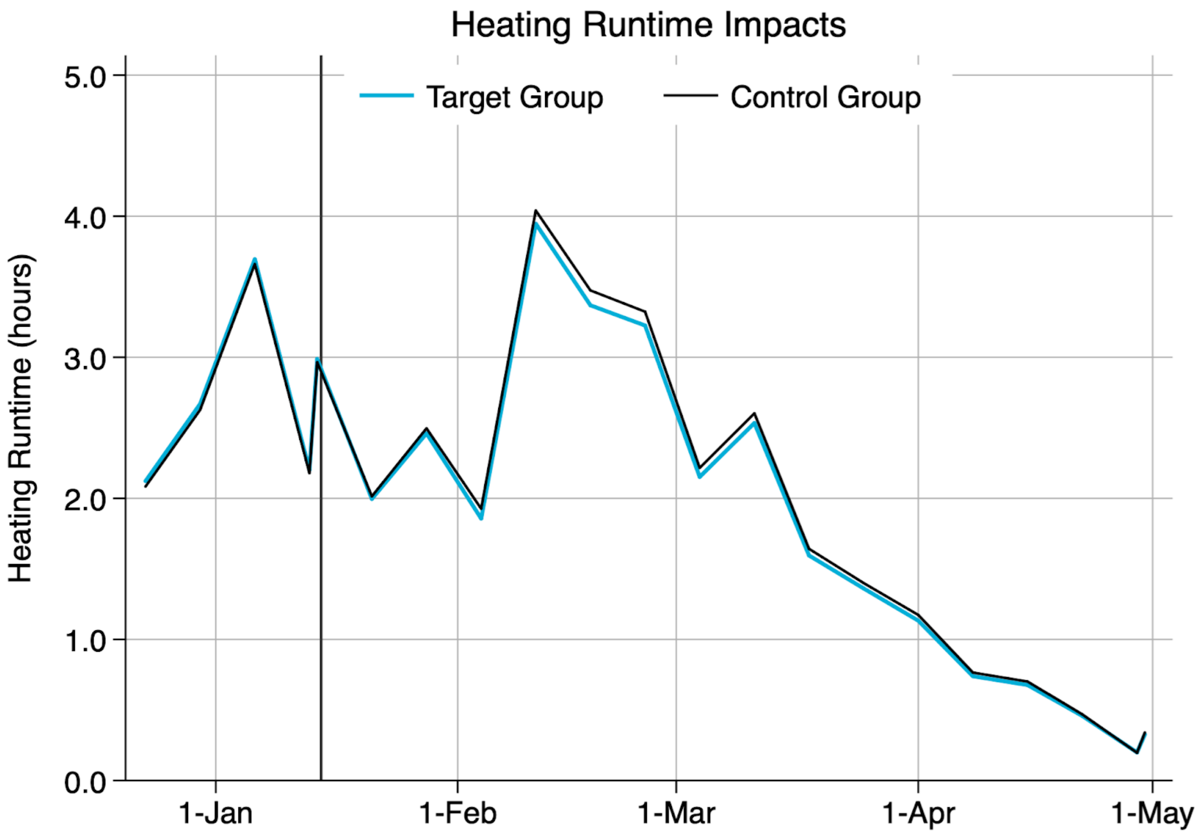

Figure 3. Average Daily Heating Runtime by week: Target and Control Groups

The heating runtime is very similar for the two groups in the weeks prior to the deployment and then the participant runtime becomes noticeably lower than the controls -- the blue line drops below the black line. The difference is most clear during the weeks with the most runtime.

The impact of Seasonal Savings on heating runtime can be quantified using regression analysis similar to methods employed to analyze utility meter data. The analysis dataset included more than 4.9 million daily runtime observations on 30,011 target participants and 3,321 controls.

Two statistical modeling approaches were used to analyze the impacts - a degree day based model that includes thermostat specific fixed effects (similar to billing data analysis methods) and a difference-in-differences (DiD) model that included both thermostat and date fixed effects. The model with the smaller standard error for percent savings is selected as the better model.

In this case, the HDD model provided the smaller estimated uncertainty with net savings per opt-in estimated at 6.97% ±2.69% (the two way fixed effects model estimate was 6.59% ±3.94%). The confidence interval is fairly wide due to highly variable nature of daily heating runtime and the small control group. These results are somewhat larger than the net savings of

5.3% ±1.8% found in the prior winter's deployment but the differences between the two winters are not statistically significant.

Energy Savings

The estimated percent heating savings were converted to energy savings based on the actual heating run time of the participants during the period (January 15, 2019 through April 30, 2019) and using estimated heating system input rates. For gas furnaces, input rates were based on climate zone average values in DEER which averaged 48.5 kBtu/hr. The electric power draw of these furnaces was estimated at 553 Watts (primarily air handler power). There were 265 heat pumps in the opt-in group (1.6% of total opt-ins) and we estimated the average power draw at 2.98 kW using a sizing estimation algorithm.

The results of the savings calculations are shown in the table below.

Table 1. Energy Savings Summary

		Runtime Hours		Energy Savings per opt-in		Aggregate Energy Savings	
Heat Type	# opt-in	Post Period baseline	Saved	therms	kwh	therms	kWh
Gas	16,575	212	15	7.2	8.2	118,903	135,573
Heat Pump	265	250	17	0	52	0	13,764
Total	16,840					118,903	149,338

Gas heated homes that opted to participate in Seasonal Savings are estimated to have saved 15 hours of furnace runtime which equals 7.2 therms of natural gas and 8.2 kWh of electricity (primarily from reduced air handler power use). The estimated total post-deployment gas heating use was just 103 therms. The 265 opt-in participants with heat pump are estimated to have reduced runtime by 17 hours on average which equals 52 kWh over the season. Aggregate savings were 118,903 therms of natural gas and 149,338 kWh of electricity.

These savings results do not include any savings achieved after April 2019.

*Note: This study is specific to the Seasonal Savings program deployed for eligible, participating MCE customers during the 2018/19 heating season. The results found herein do not necessarily represent expected results from the Seasonal Savings program under different conditions.

Appendix A: Statistical Methods

The HVAC runtime recorded by the thermostats can be used to directly assess the impacts produced by Seasonal Savings. A variety of model specifications can be used and will tend to produce similar estimates, especially in large scale RCT/RED designs. Nest employed two different model specifications in this study:

1) a standard billing data analysis style fixed effects model that analyzes pre and post deployment runtime data for the treatment and control groups and includes degree day terms to account for weather;
2) a two way fixed effects model similar to that used for the set points analysis -- including thermostat and date specific fixed effects

The two models tend to provide similar results for large sample RCT/RED studies with substantial pre and post deployment runtime data available. Some potential advantages of the degree day model include: the ability to estimate savings for periods other than the exact post deployment period analyzed; familiarity in the energy program evaluation community with standard specifications available; the ability to account for differences in weather patterns between the treatment and control groups that could result from any imbalance in the randomization or in the matching process for synthetic control groups, especially if geographically dispersed. But the two way fixed effects model accounts for date-specific factors that may not be accounted for by degree days alone and does not rely on any strong relationship between weather and runtime.

The degree day model specification is:

(Eq. 2) $Runtime_{it} = \beta_1 * DD_{it} + \beta_2 * Treat_i * DD_{it} + \beta_3 * Post_t + \beta_4 * Post_t * DD_{it} + \beta_5 * Post_t * Treat_i + \beta_6 * Post_t * Treat_i * DD_{it} + Tstat_i + \varepsilon_{it}$

where:

$Runtime_{it}$ is the hours of HVAC runtime for thermostat i on day t

DD_{it} is the heating (cooling) degree days (base 60° heating, 65° cooling) for thermostat i on day t

$Post_t$ is an indicator variable equal to 1 if date t is after the deployment date, otherwise 0

$Treat_i$ is an indicator variable equal to 1 if the thermostat is in the target treatment group and 0 if it is in the control group

$Tstat_i$ is the thermostat specific fixed effect for thermostat i

β_1 is the estimated change in HVAC runtime per degree day

β_2 is the estimated additional change in HVAC runtime per HDD for thermostats in the target treatment group

β_3 is the estimated change in HVAC runtime in the post deployment period

β_4 is the estimated added change in HVAC runtime per degree day in the post period

β_5 is the estimated additional change in HVAC runtime in the post deployment period for thermostats in the target treatment group – which is an estimate of runtime savings that are constant per day for the treatment in the post period

β_6 is the estimated change in HVAC runtime per degree day in the post deployment period for thermostats in the target treatment group – which is an estimate of runtime savings per degree day for the treatment

ε_{it} is the random error term for thermostat i on date t. The variance is calculated accounting for clustering within thermostat

The savings during the post deployment period are then calculated based on the coefficients β_5 and β_6 and the number of days and degree days in the post period. Percent savings are calculated based on post deployment runtime. The estimated standard errors of the savings are calculated using the variance-covariance matrix from the regression that accounts for clustering within thermostat.

The second regression model is simpler in that it models daily runtime data as a function of thermostat and date-specific fixed effects and then includes a simple postXtreatment dummy variable to estimate the net impact.   The model specification is:

(Eq. 3) $Runtime_{it} = \beta_1 * PostTreat_{it} + Tstat_i + Date_t + \varepsilon_{it}$

where:

$PostTreat_{it}$ is a dummy variable equal to 1 if thermostat i is in the target treatment group and date t is in the post-deployment period, otherwise it is 0

$Tstat_i$ is the thermostat specific fixed effect for thermostat i

$Date_t$ is the date-specific fixed effect for date t

β_1 is the net impact of Seasonal Savings on runtime estimated by the regression model

ε_{it} is the random error term for thermostat i on date t. The variance is calculated accounting for clustering within thermostat

The savings from this model are directly shown from the coefficient β_1. The estimated standard errors of the savings are calculated using the variance covariance matrix from the regression that accounts for clustering within thermostat.