Seasonal Savings Impacts in Southern California Gas Territory

Winter 2018/19

In January 2019, Nest launched its Seasonal Savings schedule tune-up algorithm targeted at 292,152 thermostats in the Southern California Gas (SCG) service territory. The table below summarizes the heating impacts of Seasonal Savings on these homes. The analysis followed the same evaluation approach as used in the Winter 2017/18 analysis and included a randomly selected control group of 32,460 thermostats (10% of the target pop) to provide unbiased impact estimates.

The analysis found an estimated 3.7% reduction in heating use per opt-in thermostat. These savings are significantly lower than the 11.7% found for the winter 2017/18 deployment. The change may be due in part to the unusually high savings found last year combined with the inherent variability in heating patterns in mild climates and the later deployment this year (Jan 15th vs Dec 19th) along with a smaller average set point temperature reduction (0.42°F this year vs. 0.54°F last year).

Summary of Impacts	Standard
Participation	292,152 targeted
Qualified on-line, running heating schedule	184,131
% of targeted	63%
Opted In	117,075
% of qualified	64%
Net Change in Set Point Temperatures per opt-in	
Change in Scheduled Set Points	0.59°F ±0.03°F
Change in Actual Set Points	0.42°F ±0.04°F
Heating Runtime Analysis	
% Heating Savings (degree-day regression)	3.7% ±1.9%
therms/ opt-in participant	2.7 ±1.4
therms aggregate	317,292
kWh aggregate (furnace fans and .04% heat pumps)	368,153

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 324,612 thermostats were identified as being in the potential target population -- which was defined as Nest customers with central heating in the SCG service territory. To provide for an unbiased evaluation of impacts, a control group of 32,460 thermostats (10%) was randomly selected from this population. The Seasonal Savings algorithm was deployed to the remaining 292,152 thermostats on January 15, 2019. Overall, 63% (184,131) of the targeted thermostats qualified to run Seasonal Savings -- they were on-line in running a heating schedule during the four week deployment period. Among qualified thermostats, 64% (117,075) of the customers opted to enroll in Seasonal Savings. The qualification rate was the same as the prior winter's deployment and the opt-in rate was a little lower (69% last winter).

Savings Analysis

Seasonal Savings makes changes to customer heating schedules which then leads to more efficient heating set points which then leads to a reduction in heating system runtime hours.

The evaluation analyzed the changes in the customer's heating schedules and set points to demonstrate program impacts. Energy savings were quantified by analyzing heating system runtime. The randomized control group provided a true experimental design for the evaluation -- comparing the impacts for the entire targeted participant group (including those that did not opt-in) to the control group. This evaluation approach is called an Intent-to-Treat or Randomized Encouragement Design (RED).

An RED eliminates self-selection bias because it directly estimates the impact of being in the target participant group -- not the impact of actually participating. To estimate the savings per participant that opted-in, the RED results can be adjusted for 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%).

Analysis of Set Points

The average scheduled thermostat set points for the Seasonal Savings target population and the control group are shown in Figure 1 with a vertical line marking the date of deployment.

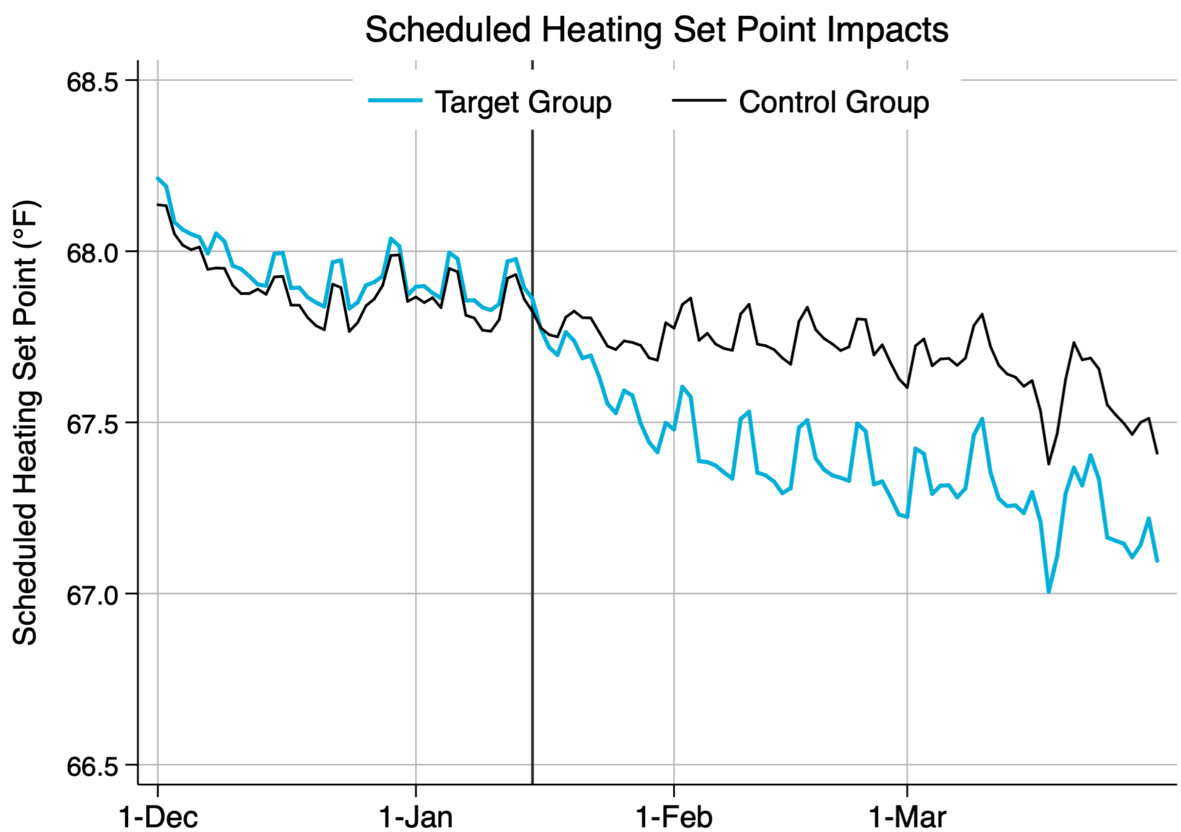

Figure 1. Scheduled Heating Set Points over the season

The graph clearly shows a net reduction in scheduled set points as the algorithm deploys over three weeks. It's worth noting that the plot shows the impacts diluted by the fact that about half (54%) of the target group did not actually opt-in to Seasonal Savings.

Figure 2 directly plots the difference between the two lines in Figure 1 -- providing a better illustration of the schedule impacts. It also plots the differences using the actual executed set points. Actual set points can differ from the schedule due to manual adjustments (via dial or app or web) or to the auto-away feature based on occupancy detection.

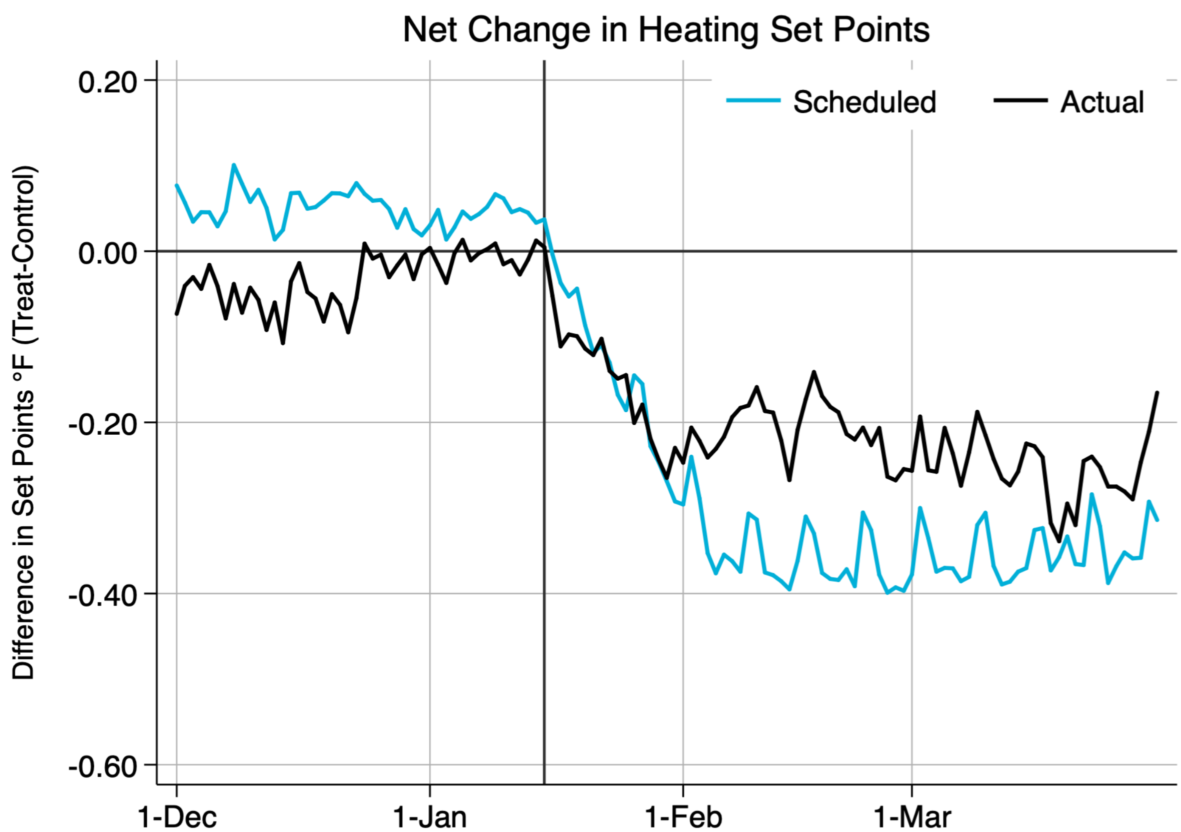

Figure 2. Difference between Target and Control Group set points: Scheduled and Actual

The figure shows a clear impact on set points. The net change in set points during the season was quantified using a regression analysis that included both thermostat and date fixed effects (see Appendix A for details) which found a net reduction in scheduled set points of 0.34°F per targeted thermostat, equal to 0.59°F per opt-in (57% opt-in rate for customers in the analysis). The average reduction in actual set points was 0.42°F (±0.04°) per opt-in participant. This reduction is a little smaller than the 0.54°F reduction that was found in the winter 2017/18 deployment.

Heating Runtime Analysis

The heating runtime recorded by the thermostats can be used to assess savings. Figure 3 shows the daily heating run time averaged by week for the target group and control group over the season.

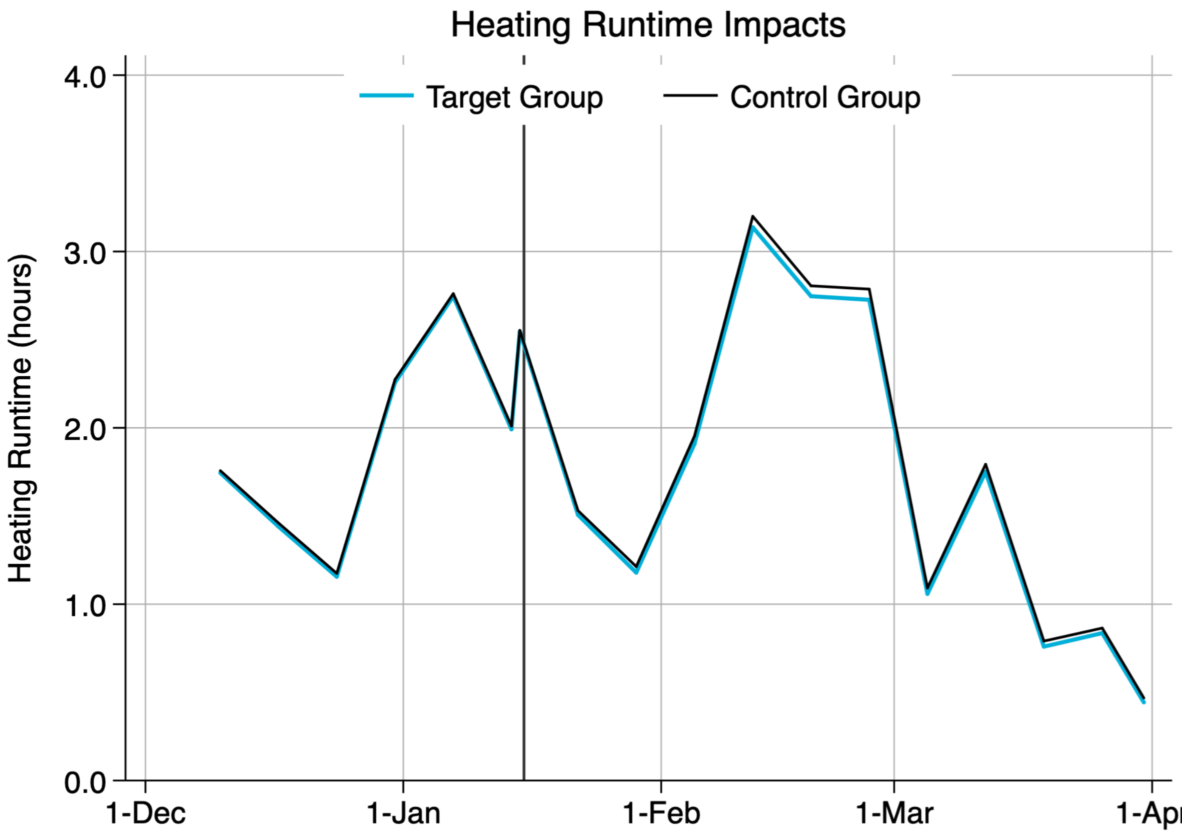

Figure 3. Average Daily Heating Runtime by week: Target and Control Groups

The graph shows that the two groups are nearly indistinguishable although a close examination shows a small decline in the target group compared to the control group after the deployment -- especially clear during the peak heating days in February.

The impact of Seasonal Savings on heating runtime was quantified using regression analysis similar to methods employed to analyze utility meter data. The analysis dataset included 45 million daily runtime observations from 277,829 targeted participants (116,320 opt-ins) and 30,789 controls. HVAC runtime was modeled as a function of degree days with interaction terms for the post period and treatment group and including thermostat specific fixed effects. Appendix A provides more details on the modeling specifications.

The regression analysis estimated that Seasonal Savings reduced heating runtime by 3.7% (±1.9%) per opt-in participant. These savings are much lower than the 11.7% (±1.9%) found in the winter 2017/18 evaluation. The prior year percent savings were unusual -- the largest ever found for Seasonal Savings impacts (although the energy savings of 10 therms/opt-in savings were lower than has been found for many deployments in colder climates). So the reduction in percent savings for this year, although large, was not completely unexpected.

Energy Savings

The estimated percent heating savings were converted to therms and kilowatt-hours based on the actual heating run time of the participants' heating systems during the post-deployment period and using estimated heating system natural gas and electric input rates.

We used the same estimated average gas furnace input rate as used in the 2017/18 evaluation (which had been based on DEER averages by climate region) -- which was 50.1 kBtu/hr. A reduction in heating hours will also result in a reduction in furnace electric usage -- primarily from the air handler fan but also other ancillary uses (e.g., direct vent fan). We used the same 580W estimated power draw as used in the 2017/18 evaluation. In addition, additional electric savings will be provided for the 0.04% of the opt-in thermostats that control heat pumps (heat pump customers were specifically excluded from the target sample but Seasonal Savings is deployed by structure and a few structures had both gas heat and a heat pump). We used the prior year estimate of 3.63 kW as the average power draw for heat pumps. The table below shows the resulting savings per opt-in participant and in aggregate.

Energy Savings Summary

	Impacts per Opt-In Participant		
	Gas Heat	Heat Pump	Aggregate
# Opt-in Participants	116,277	43	116.320
Hours Baseline	148	144	
Hours Saved	5.5	5.3	
Gas Savings (therms)	2.7 ±1.4	0	317,292
Electricity Savings (kWh)	3.2 ±1.6	19 ±10	368,153

These savings results do not include any savings achieved after April 2019 (i.e., the end of the winter and persistence into the following heating season are both assumed to be zero) -- a full accounting of savings would likely result in a larger total.

*Note: This study is specific to the Seasonal Savings program deployed by Nest for eligible, participating Southern California Gas 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

Analysis of Net Changes in Set Points

The net impact of Seasonal Savings on customer set points is estimated by analyzed the daily average set points for all thermostats in the target participant group and the control group across the pre and post deployment data for the season. The analysis accounts for both thermostat-specific and date-specific effects using a two way fixed effects model and excludes days with no heating (cooling) runtime. The net impact of Seasonal Savings can then be estimated using a single explanatory variable – an indicator for post deployment period in the treatment group. This same model is used to analyze the scheduled set points and the actual executed set points. The statistical model is:

(Eq. 1) $Tset_{it} = \beta_1 * PostTreat_{it} + Tstat_i + Date_t + \varepsilon_{it}$

where:

> $Tset_{it}$ is the average set point for thermostat i on day t
>
> $PostTreat_{it}$ is a dummy variable equal to 1 if thermostat i is in the target treatment group and day j 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 set points 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

Analysis of Runtime Savings from Daily Data

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 produce similar estimates, especially in large scale RCT/RED designs. In this study, Nest employed 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. An alternative model that simply employs a difference-in-differences type specification with both thermostat and date fixed effects is also used in some studies. But that model tends to work best when the climate doesn't vary much across the population and so was not employed in this study.

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 period

β_4 is the estimated additional 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.