September 2, 2015

MEMO

FROM: Dan Rubado, Evaluation Project Manager, Energy Trust of Oregon;
Phil Degens, Evaluation Manager, Energy Trust of Oregon

TO: Marshall Johnson, Sr. Program Manager, Energy Trust of Oregon;
Paul Sklar, Planning Engineer, Energy Trust of Oregon;
Fred Gordon, Director of Planning & Evaluation, Energy Trust of Oregon

CC: Existing Homes program, CLEAResult

SUBJECT: Follow-up Billing Analysis for the Nest Thermostat Heat Pump Control Pilot

BACKGROUND

In 2013, Energy Trust's Existing Homes program implemented a pilot to test the Nest thermostat as an advanced heat pump control device in single family homes with central, air source heat pumps with electric resistance backup. This pilot was evaluated by Apex Analytics in 2014. The resulting electricity savings were established through billing analysis conducted by Energy Trust and reviewed by Apex, using a partial year of follow-up data. The plan at that time was to provide a preliminary savings estimate for Nest thermostats in heat pump homes and return to the analysis after a full year of data was available to establish a more definitive savings estimate. The pilot ended with 177 successful Nest thermostat installations and a comparison group of 299 similar heat pump homes. The 2014 analysis estimated that the annual electric savings attributable to Nest were 781 kWh (90% CI: 316, 1246), or 4.7% of annual electric use and roughly 12% of the average heating load. The 2014 evaluation report can be found on Energy Trust's website at: http://energytrust.org/library/reports/Nest_Pilot_Study_Evaluation_wSR.pdf. Once a full year of post-installation electric use data were available, we went back and re-analyzed the annual electric savings for pilot homes.

METHODS

Housing and occupant characteristics for pilot participant and comparison homes were retrieved from Energy Trust's project tracking database and from ancillary data collected during the pilot and evaluation project. Monthly electric utility billing data for each home were retrieved from Energy Trust's utility database. Homes that could not be matched to utility data were dropped from the analysis. We computed the raw daily average electric usage for each billing period for each home. Daily usage was the primary unit for the analysis. Weather data from nearby weather stations were retrieved from the

National Climatic Data Center. Daily average temperature was used to calculate heating degree-days (HDDs) and cooling degree-days (CDDs) for each billing period for each home. The HDD and CDD variables were computed for reference temperatures for every degree ranging from 45 to 85 °F. The HDD and CDD values were then divided by the number of days in each billing period to obtain average daily HDD and CDD variables, which could be directly compared with the average daily electric use.

The pre-pilot period of the study was defined as June 2012 through July 2013. Thermostat installations were conducted from August through December 2013, so this period was excluded from the analysis. The post-installation period was defined as January 2014 through February 2015. Monthly electric use readings from each study period were identified and flagged in the analysis dataset.

Using similar methods to the 2014 analysis, we re-ran our analysis of the pilot homes to determine the energy savings attributable to the Nest thermostat in heat pump homes. Participant and comparison homes were screened for potential issues with their billing data, outliers in annual electric usage, large swings in electric usage, solar PV systems, and other Energy Trust-funded efficiency measures. The sample attrition is described in more detail in the results section. We used two analysis methods to compute electricity savings. We first used a very similar multilevel model specification to the 2014 analysis, with the addition of cooling terms. Then, we used a PRISM-like analysis to weather normalize annual usage and computed savings as the difference-in-differences.

Multilevel Model

The first approach was to compare the pre-to-post change in electric use between the study groups using the multilevel mixed effects model. Average daily electric use was modeled as a function of average daily HDDs and CDDs, a study period (pre- vs. post-installation) flag, study group (participant vs. comparison) flag, home square footage, and year built. Interaction terms between the study period flag, study group flag, and HDD and CDD variables were added to model the impact of the Nest thermostat, depending on weather conditions. Additional terms were added to model the relationship between HDDs and CDDs and electric use separately for each home in the sample, which makes this method analogous to the PRISM method. The advantages to this type of model are that it accounts for repeated observations over time within each home and simultaneously computes the effect and variance of the study group and study period. The primary drawback is that it applies the same HDD and CDD reference temperatures to all homes. The following formula describes the resulting linear mixed effects model:

$$Usage_{ij} = \beta_0 + \beta_1 HDD_{ij} + \beta_2 CDD_{ij} + \beta_3 Group_{ij} + \beta_4 Post_{ij} + \beta_5 SqFt_{ij} + \beta_6 YearBuilt_{ij} + \beta_7 Group_{ij} * Post_{ij} + \beta_8 Group_{ij} * HDD_{ij} + \beta_9 Post_{ij} * HDD_{ij} + \beta_{10} Group_{ij} * Post_{ij} * HDD_{ij} + \beta_{11} Group_{ij} * CDD_{ij} + \beta_{12} Post_{ij} * CDD_{ij} + \beta_{13} Group_{ij} * Post_{ij} * CDD_{ij} + u_{0i} + u_{1i} HDD_{ij} + u_{2i} CDD_{ij} + \epsilon_{ij}$$

Where:

$Usage_{ij}$ = the average daily electric usage for home i during billing month j,
β = the coefficients for each variable in the model,
β_0 = the fixed intercept for all homes,

HDD_{ij} = Heating Degree-Days for home *i* during month *j*,
CDD_{ij} = Cooling Degree-Days for home *i* during month *j*,
$Group_{ij}$ $\{0,1\}$ = dummy variable where 1 indicates that home *i* is part of the participant study group, which is static across all *j* billing months,
$Post_{ij}$ $\{0,1\}$ = dummy variable where 1 indicates that home *i* during billing month *j* is in the post-installation study period,
$SqFt_{ij}$ = square footage of home *i*, which is static across all *j* billing months,
$YearBuilt_{ij}$ = year of construction of home *i*, which is static across all *j* billing months,
u_{0i} = random intercept for home *i* which is independent from ϵ_{ij},
u_{1i} = random slope coefficient of HDD for home *i* which is independent from ϵ_{ij},
u_{2i} = random slope coefficient of CDD for home *i* which is independent from ϵ_{ij}, and,
ϵ_{ij} = model error for home *i* during billing month *j*.

As noted above, HDD and CDD variables with different reference temperatures were tested in the model using all possible combinations from 45 to 85°F. The reference temperatures that resulted in the model with the best fit was selected as the final model, based on the fit statistics (AIC and BIC). A HDD reference temperature of 55°F and CDD reference temperature of 70 °F proved to have the best fit for this sample of homes.

The model provided three key parameter estimates for computing energy savings: the interaction coefficients β_7, β_{10} and β_{13}. Together, these coefficients describe the difference between participant and comparison group homes change in pre- to post-installation average daily usage for a given number of HDDs and CDDs. The model factors out the influence of any differences in square footage or year built between homes on the interaction coefficients. So, the sum of these coefficients is the average daily electric savings. A linear combination of these three coefficients was computed to estimate the weather normalized annual electric savings in kWh per home, as described below. We also computed the pre-pilot average annual electric use and heating usage for the treatment group from the parameter estimates in kWh per home, so that we could calculate energy savings as a percent of annual electric and annual heating loads.

Average Annual Savings = $365 * \beta_7 + LRHDD * \beta_{10} + LRCDD * \beta_{13}$
Average Annual Usage = $365 * (\beta_0 + \beta_3 + AvgSqFt * \beta_5 + AvgYearBuilt * \beta_6) + LRHDD * (\beta_1 + \beta_8) + LRCDD * (\beta_2 + \beta_{11})$
Average Annual Heating Usage = $LRHDD * (\beta_1 + \beta_8)$
Average Annual Cooling Usage = $LRCDD * (\beta_2 + \beta_{11})$

Where:
AvgSqFt = average square feet across all homes in the sample,
AvgYearBuilt = average year of construction across all homes in the sample,
LRHDD = long-run average annual HDDs for each weather station, averaged across the homes in the sample, derived from the Typical Meteorological Year 3 (TMY3) dataset, and,

LRCDD = long-run average annual CDDs for each weather station, averaged across the homes in the sample.

PRISM-like Analysis

Next, we used a PRISM-like (PRInceton Score-keeping Method[1]) weather normalized annual usage, differences-in-differences approach. We fitted separate weather regression models for each home for both the pre- and post-installation study periods, using HDD and CDD variables. All combinations of HDD and CDD reference temperatures were run for all home-level regression models, from 45° to 85°F. The model results with the highest R-squared for each home and study period were selected to calculate the weather normalized annual usage, using the TMY3 long-run HDDs and CDDs. However, if the model R-squared was less than 0.5 or the HDD coefficient was negative, then we assumed the home was insensitive to weather and used the raw annual usage for the analysis. The primary advantage of this method is that the models are specified for each individual home so the reference temperatures that best fit the data are used. Unfortunately, since the weather normalization is not done simultaneously with the difference-in-differences computation, the error terms are not carried through from the regression models to the savings estimate. Thus, this method tends to understate the statistical significance of the results. The model specifications for weather normalization were:

Average daily usage$_i$ = $\beta_0 + \beta_1 HDD_i(\tau_h) + \beta_2 CDD_i(\tau_c) + \varepsilon_i$
Normalized annual usage$_i$ = $365*\beta_0 + \beta_1 LRHDD_i(\tau_h) + \beta_2 LRCDD_i(\tau_c)$
Normalized heating usage$_i$ = $\beta_1 LRHDD_i(\tau_h)$

Where:
i = home indicator,
β_0 = Estimated average daily "base load" usage for home *i,*
β_1 = Model predicted heating slope,
$HDD_i(\tau_h)$ = Average daily HDDs at reference temperature τ_h,
β_2 = Model predicted cooling slope,
$CDD_i(\tau_c)$ = Average daily CDDs at reference temperature τ_c,
ε_i = Unexplained error term,
$LRHDD_i(\tau_h)$ = Long-run average annual HDDs at reference temperature τ_h, and,
$LRCDD_i(\tau_c)$= Long-run average annual CDDs at reference temperature τ_c.

Next, the difference was taken between the pre- and post-pilot normalized annual electric usage for each home. To determine electric savings while controlling for square footage and year built, we created another regression model where study group predicted the delta in annual usage. The coefficient of the study group variable was the annual electric savings.

[1] Fels, M. (1986). PRISM: An Introduction. Energy and Buildings, 9, 5-18. Retrieved from http://www.marean.mycpanel.princeton.edu/~marean/images/prism_intro.pdf

RESULTS

Pilot homes were removed from the analysis sample for a variety of reasons, including: not matching to billing data, insufficient billing data for analysis (less than six months of billing records in either study period), missing or invalid information on square footage or year built, known solar PV systems funded by Energy Trust, other Energy Trust-funded efficiency measures, outliers in annual electric use (<5,000 kWh/year or >40,000 kWh/year), homes with large swings in annual electric use (more than 100% increase or 50% decrease in kWh/year), and participant homes where Nest thermostats were known to have been uninstalled mid-pilot. After the attrition steps, 60% of participants homes and 67% of comparison homes remained for analysis. Compared to the 2014 evaluation, there was some additional attrition in this analysis. **Table 1** summarizes the sample attrition.

Table 1: Attrition analysis for Nest heat pump control pilot participant and comparison homes.

	Participant Group				Comparison Group			
Phase of Analysis	N Removed	N Homes	% Homes	2012 kWh Usage	N Removed	N Homes	% Homes	2012 kWh Usage
All Nest pilot homes	0	177	100%	--	0	299	100%	--
Homes matched to billing data	-13	164	93%	17,315	-40	259	87%	16,583
Homes removed with solar PV	-5	159	90%	17,379	-3	256	86%	16,677
Homes with sufficient valid billing data	-8	151	85%	17,557	-24	232	78%	16,477
Homes removed with Energy Trust projects	-35	116	66%	17,362	-20	212	71%	16,290
Homes with valid sq.ft. and year built	-5	111	63%	17,316	-5	207	69%	16,294
Outliers removed with low annual usage	-1	110	62%	17,476	-1	206	69%	16,373
Outliers removed with high annual usage	-1	109	62%	17,246	-4	202	68%	15,642
Outliers removed with large changes in annual usage	0	109	62%	17,246	-3	199	67%	15,653
Homes removed where Nest uninstalled	-3	106	60%	17,150	0	199	67%	15,653
Total homes available for analysis		**106**	**60%**	**17,150**		**199**	**67%**	**15,653**

The basic characteristics of the pilot homes in the final sample are summarized in **Table 2** and the geographic distribution of homes is summarize in **Table 3**. The participant and comparison group homes are clearly very similar on all of the dimensions analyzed here. The mean square footage and year of construction are nearly identical between the groups. The percent of site built versus manufactured homes and the geographic distribution were also similar for participant and comparison homes.

Table 2: Summary of Nest pilot home characteristics.

Group	N	Mean Square Footage	Mean Year Built	% Site Built
Participants	106	1,681	1978	80%
Comparison	199	1,666	1978	74%
Total	305	1,672	1978	76%

Table 3: Geographic distribution of Nest pilot homes.

	Portland Metro		Willamette Valley		Southern Oregon	
Group	N	% of Homes	N	% of Homes	N	% of Homes
Participants	104	59%	28	16%	45	25%
Comparison	150	50%	48	16%	101	34%
Total	**254**	**53%**	**76**	**16%**	**146**	**31%**

Multilevel Model

We specified multilevel mixed effects models with all combinations of HDD and CDD reference temperatures, with CDD reference temperatures greater than or equal to HDD, between 45°F and 80°F. The model we selected with the best fit had the same reference temperatures as in the 2014 analysis: 55°F for HDD and 70°F for CDD.

The annual electric savings estimate was 645 kWh (**Table 4**), or 3.8% of annual electric use and 14% of heating usage (**Table 5**). This savings estimate is slightly but not significantly lower than the preliminary savings estimate calculated in the 2014 analysis of 781 kWh (which was itself lower than the original savings estimate for heat pump advanced controls), yielding a realization rate of 83%. The savings percentages were based on the participant group's pre-pilot normalized annual electric use of 16,935 kWh and annual heating usage of 4,542 kWh (27% of annual usage), as computed from the model coefficients. We also calculated that normalized annual cooling usage in pilot homes from the model coefficients and found that it was very low, on average, at an estimated 200 kWh per year. As a result, there were no detectable cooling savings.

Table 4: Multilevel model annual electric savings for the Nest thermostat in heat pump homes.

Annual kWh Savings	Std. Err.	90% Conf. Interval	p-value
645	152	376, 914	<0.001*

* Highly statistically significant at the 90% confidence level.

Table 5: Multilevel model annual electric savings for the Nest thermostat in heat pump homes, as percentages of total electric use and heating usage.

% Savings	% Heating Savings	Annual kWh Usage	Annual Heating kWh Usage	% Heating Usage
3.8%	14%	16,935	4,542	27%

PRISM-like Analysis

The annual electric savings estimate using the PRISM-like method was 544 kWh (**Table 6**), or 3.2% of annual electric use and 11% of heating usage (**Table 7**). **Figure 1** displays the average pre- and post-pilot weather normalized annual electric use for each study group to illustrate how the savings were computed as the difference-in-differences. The savings were substantially lower than those estimated from the 2014 analysis, yielding a realization rate of 70%. As noted above, the error of the savings estimate is overstated and the statistical significance is understated due to the two-stage nature of the method. The savings percentages were based on the participant group's pre-pilot normalized annual electric use of 16,876 kWh and heating usage of 5,127 kWh (30% of annual usage), as computed from the model coefficients. These usage estimates are similar to the usage estimates from the multilevel model. The average best fit HDD reference temperatures used for the pre- and post-installation regression models were, 56 and 55°F, respectively, essentially the same as the multilevel model. The average CDD reference temperatures used were 68 and 67°F, respectively, close to the 70°F temperature used in the multilevel model.

Table 6: PRISM-like analysis annual electric savings for the Nest thermostat in heat pump homes.

Annual kWh Savings	Std. Err.	90% Conf. Interval	p-value
544	384	-91, 1178	0.158*

* Borderline statistically significant at the 90% confidence level. However, this method tends to understate statistical significance.

Table 7: PRISM-like analysis annual electric savings for the Nest thermostat in heat pump homes, as percentages of total electric use and heating usage.

% Savings	% Heating Savings	Annual kWh Usage	Annual Heating kWh Usage	% Heating Usage
3.2%	11%	16,876	5,127	30%

Figure 1: Average normalized annual electric use for participant and comparison homes in the pre- and post-pilot study periods.

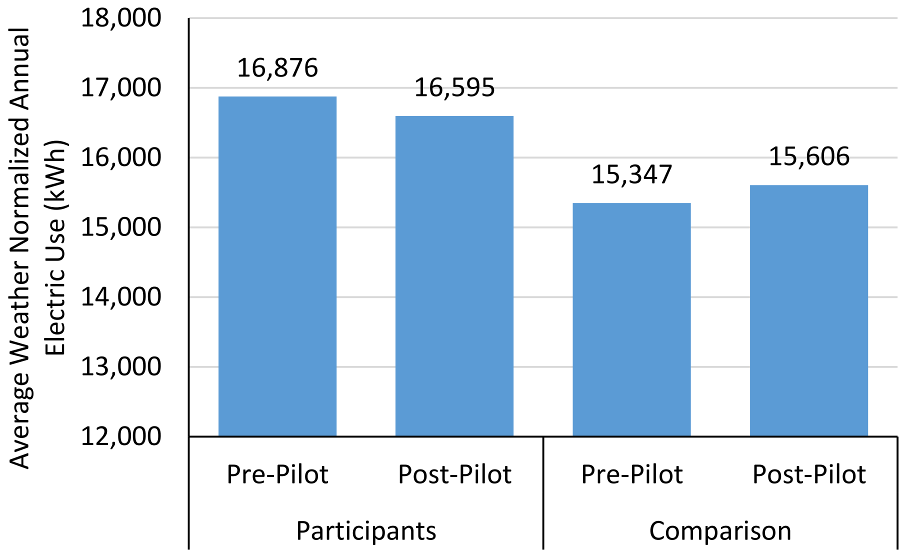

We created graphs of the changes in weather normalized annual usage from the pre-to-post installation periods for individual homes to illustrate the distributions used to compute the average savings. **Figure 2** shows the distribution of changes in usage with the kernel density. **Figure 3** shows a scatter plot of the changes in usage as a function of the pre-period annual usage. The large amount of scatter and substantial overlap between the distributions of the participant and comparison homes demonstrates that there is a large amount of variability in the results for individual homes. However, there were significant electric savings on average.

Figure 2: Distribution of changes in normalized annual electric use per dwelling unit, by study group.

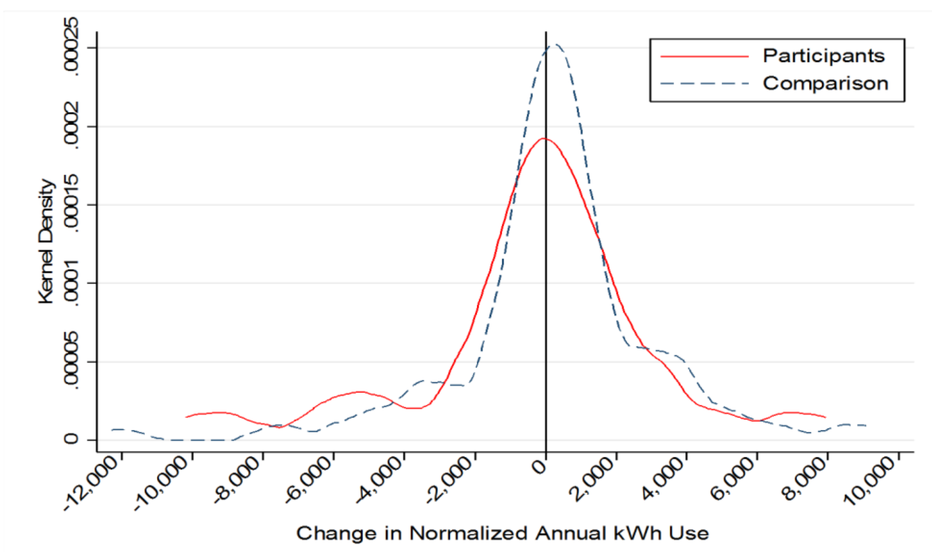

Figure 3: Scatterplot of changes in normalized annual electric use per dwelling unit versus pre-installation normalized annual electric use, by study group.

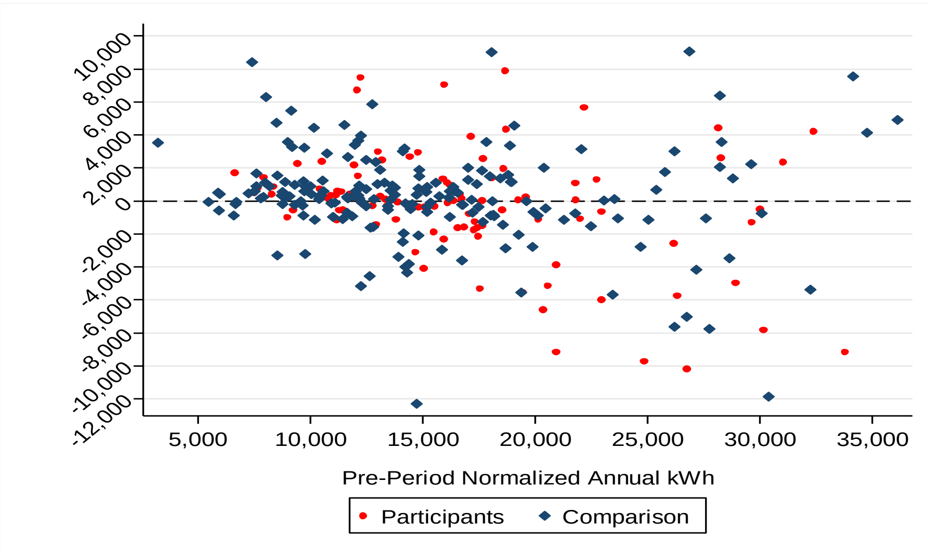

CONCLUSIONS & RECOMMENDATIONS

This billing analysis is a follow-up to Energy Trust's preliminary savings results for the Nest thermostat in homes with central heat pump systems, presented in a 2014 evaluation report. Using a full year of post-pilot electricity billing data, we were able to re-analyze pilot homes to update the electric savings estimate. Unfortunately, after a year elapsed, there was 40% attrition of Nest participant homes (compared to 36% in the original analysis), reducing the final sample for the analysis. Most of these homes were removed from the analysis because they received Energy Trust incentives for additional electric efficiency measures in the follow-up period. Since more participant homes were removed than comparison homes (40% versus 33%), this source of attrition could bias the results. The differential attrition combined with high variability in the results and a participant group that may not represent typical smart thermostat customers means that the generalizability of this study may be limited. As a result, the savings estimate may be subject to additional changes in the future, given additional participants and data to analyze. Ideally, to nail down the electric savings for Nest thermostats in heat pump homes, we would need a larger sample of participants who purchased their thermostats through typical market channels. However, this pilot provides the only data we currently have to assess savings for the Nest thermostat in heat pump homes.

The estimated annual electric savings for the Nest thermostat differed somewhat, depending on the analysis method used. The multilevel model approach yielded a savings estimate of 645 kWh per year and the PRISM-like approach resulted in an estimate of 544 kWh per year. These two approaches were equally valid and the model outputs and results are similar. We believe the average of the two results provides a reasonable estimate of savings for Nest. Thus, the average annual electric savings due to the Nest thermostat in heat pump homes was 594 kWh per year, which is equivalent to 3.5% of annual electric use and 12% of heating usage (**Table 8**). This represents a 24% decrease from the preliminary savings estimate of 781 kWh per year from the 2014 analysis (realization rate of 76%), although the difference was not statistically significant. The updated savings estimate is based on a full year of post-installation data and should replace the preliminary estimate as the electric savings claimed by Energy Trust for Nest thermostats installed in heat pump homes.

Table 8: Final annual electric savings for the Nest thermostat in heat pump homes.

Annual kWh Savings	% Savings	% Heating Savings
594	3.5%	12%