White Paper

Nest Learning Thermostat Efficiency Simulation: Update Using Data from First Three Months

Nest Labs
April 2012

1. Introduction

The energy savings described herein are primarily gained by changing the thermostat setpoint temperature while occupants are away from the home and when they are asleep. For the vast majority of occupants who have no setback schedule programmed, the greatest savings are realized simply by generating a setback schedule that suits their daily routines. Further energy savings can be gained through (1) Auto-Away, or reducing heating, cooling and air conditioning (HVAC) usage during extended absences from the home (e.g. while on vacation), and (2) making very small changes in the setpoint temperature (e.g. by changing the setpoint by one degree Fahrenheit), perhaps in response to learnings from the Nest Leaf.

The goal of this white paper is to demonstrate that certain features of the Nest® Learning Thermostat™ are expected to result in energy savings, based on the described simulations. Using the simulation findings and weighting them across all U.S. zip codes, it's estimated that the Nest Learning Thermostat can save an average of $173 per year, before applying any of Nest's additional energy-saving features such as Auto-Away™ and the Leaf. Modifications of the simulations will also be undertaken as needed.

It is important to note that the energy savings described herein are expected without resorting to any optimized control of the HVAC system components.

It is also important to note that the strategy of the Nest Learning Thermostat is not solely energy savings. The Nest Learning Thermostat places a high priority on the user's comfort. This white paper makes assumptions of households with mild to moderate energy consciousness in mind.

In April 2012, the simulation model was calibrated using data sampled in January/February 2012 to estimate the energy savings experienced. Devices with a good schedule, as defined in Section 6.2, had an average savings per device of 19.5%, with the highest saving of 36.1%. This is consistent with the EPA's estimate that a properly programmed thermostat can save 20% on heating and cooling costs.

2. List of Features Simulated

Auto-Away. The Nest Learning Thermostat's "auto-away" feature automatically detects non-occupancy events, whether they last for several hours or multiple days (e.g. a vacation). The Auto-Away feature is based on algorithms that interpret occupancy sensor data and provide a confidence determination of whether or not the occupants are away from the home. When the confidence level is high that occupants are away, the Auto-Away feature makes a decision to override the existing schedule to save additional energy. During relatively short periods of non-occupancy, from a few hours to a few days, the setpoint is set to a value where substantial efficiency gains can be realized.

Auto-Schedule. The Nest Learning Thermostat automatically learns a user's preferred temperature as well as schedule. The learning algorithm is based on the user's manual temperature selection on the device. Through the proprietary automatic learning algorithm, the thermostat replays a setback schedule, greatly benefitting the vast majority who train the thermostat wisely. Note that the learning algorithm simulated herein provides energy savings without substantially impacting the user's comfort.

The Nest Leaf. Another way the Nest Learning Thermostat encourages users to select efficient temperatures is the display of a green leaf icon known as the "Nest Leaf." The Nest Leaf displays when the person controlling the thermostat has chosen an energy-efficient setting. The Leaf appears at different temperatures in different households, based on (1) the calculated efficiency of the household, (2) the HVAC system model and (3) the user's prior behavior.

Time to target temperature. There are several ways Nest Learning Thermostat encourages users to select an efficient setpoint temperature. One of these features is called "time to temperature." The time-to-temperature feature calculates and displays, in real-time, an indication of how long it takes to reach the current setpoint temperature. Many people use an exaggerated temperature setting hoping to cool or heat the house more quickly. This behavior is both ineffective and inefficient. The Nest Learning Thermostat's time-to-temperature feature shows the user that an exaggerated setpoint temperature takes much longer to reach, which discourages this wasteful behavior. By providing real-time feedback when selecting manual input temperatures, the user becomes more familiar with the HVAC system, which in turn leads to more economical and environmentally friendly use of energy.

3. Methods

3.1. Simulation Setup

General Setup: The thermostat energy simulation model is a custom-made dynamic model based on principles of heat transfer and HVAC equipment performance incorporating state-of-the-art research on building and equipment performance.

The following simulations assume a 1,800 sq. ft. single-family home with an average efficiency level. The building envelope heat transfer model begins with a standard $\boldsymbol{U}$*$\boldsymbol{A}$*$\boldsymbol{dT}$ model, where $\boldsymbol{U}$ is the heat transfer coefficient; $\boldsymbol{A}$ is the surface area of the house, and $\boldsymbol{dT}$ is the difference between the indoor and outdoor temperatures. The model also accounts for thermal mass that exchanges directly with the inside and outside environments.

Weather: A typical year data set of hourly values of solar radiation and meteorological elements is used. The third and latest typical meteorological year ("TMY3") data is used, which was developed by the National Renewable Energy Laboratory for the climates simulated. Solar gain through windows is modeled based on hourly solar data from TMY3.

Other Aspects of the Model. A number of details are employed in the simulation in an effort to account for important system dynamics that could have an impact on various thermostat control strategies. Some of these details include:

- Air infiltration is based on a detailed infiltration model that includes wind and stack effects using hourly wind speeds and indoor and outdoor temperatures.
- Heating and cooling equipment is modeled to include transient start-up effects, and interactions with thermal mass and distribution systems.
- The heating equipment is assumed to be forced-air gas furnace. For both heating and cooling, single-stage systems are assumed. In future simulations, we plan to add heat pump systems, non-forced-air systems, as well as multi-stage heating and cooling systems into the model.
- Air conditioner capacity and power draw are modeled as a function of indoor and outdoor conditions with latent (i.e., humidity-related) capacity modeled using an indoor air humidity model that balances the effects of moisture loads from occupants and air infiltration with the dynamic moisture removal capacity of the air conditioner.

Other Parameters. Following are a list of conditions used in the simulation:

- The heating maintenance band is 1.4°F wide, with 1°F below and 0.4°F above the setpoint temperature. (E.g. if the setpoint is 70°F, the system maintains the room temperature between 69.0°F and 70.4°F).
- The cooling maintenance band is 2.0F wide, with 0.7°F below and 1.3°F above the setpoint temperature. (E.g. if the setpoint is 80°F, the system maintains the room temperature between 81.3°F and 79.3°F).
- Heating and cooling temperature is monitored for overshoot. If the temperature in the room (without solar effect) overshoots, then the maintenance band is adjusted to keep the room temp fluctuation as close to the maintenance band as possible.
- The air conditioner system has a minimum cycle time (both on and off) of 5 minutes (unless manually commanded), and the forced-air gas heating system has a minimum cycle time of 3 minutes.

- A simple model of user's window usage is assumed. If the room temperature suggests that the HVAC system should turn on, but the outside temperature is at least 5°F in the favorable direction, the user is assumed to open the windows and not turn on the HVAC system. For example, if the air conditioner set point is 80°F, and the house internal temperature is 83°F then normally the air conditioning system should activate. However, if the outside temperature is 72°F (cooler than 80°F – 5°F = 75°F), then the HVAC system is not activated and instead windows are opened.

3.2. Experimental Conditions

Geographic Locations: We simulated a sample of 12 cities in the continental U.S. with at least one city from each of nine U.S. climate zones seen in Figure 1. The cities included in the simulation are: Atlanta, Boston, Chicago, Dallas, Denver, Houston, Miami, Minneapolis, Phoenix, San Diego, San Francisco, and Spokane.

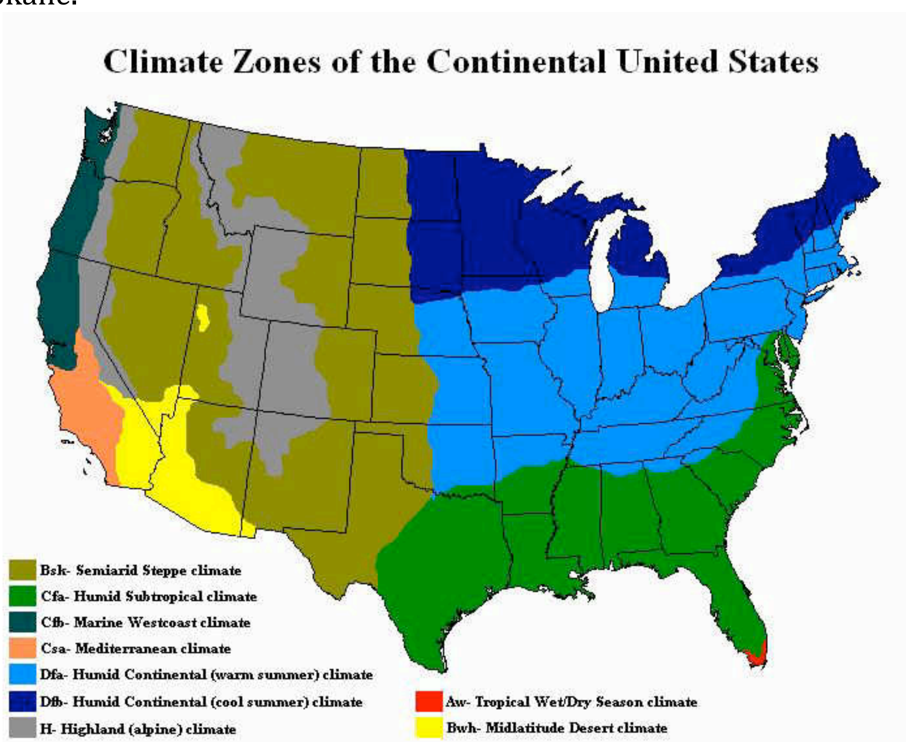

Figure 1: Climate Zones of the Continental United States

Among the list of cities, we have included four cities in which the Nest Learning Thermostat should be available in store at launch: Miami, Dallas, Houston, and San Francisco (Bay Area). Accordingly, there are some climate zones in Figure 1 that are represented with more than one city in the simulation. The simulation assumes all houses to have heating as well as cooling HVAC systems. This means that it ignores the fact that air-conditioning is rarely needed (or equipped in homes) in San Francisco and Spokane, and heating is used infrequently in Phoenix and Miami.

Occupant Types: For purposes of the simulation, the household occupants are classified into one of two types:

- **Type 1 (75%)**: Do not have long and regular mid-day non-occupied periods. Examples of Type 1 occupants include: a family with young children; a retired couple; and households in which at least one member works at home.
- **Type 2 (25%)**: Have predictable, long periods of non-occupancy in the middle of the day. Examples of Type 2 occupants include: a working couple having no children; a working couple having children in all-day daycare; and a family with older children who are out of home during the day.

Program Schedules: Assumptions are made about the schedules and setpoint temperatures that the Nest Learning Thermostat automatically learns for the Type 1 and Type 2 occupants. It is also assumed that the same program schedule is used on weekdays and weekends. Temperature ranges are selected to match mild to moderate energy conscious households. The heating and cooling schedules and setpoint temperatures simulated are as follows:

- **Type 1 heating**: 6am to 10pm, setpoint = 70°F
  10pm to 6am, setpoint = 62°F

- **Type 1 cooling**: 8am to 6pm, setpoint = 79°F
  6pm to 8am, setpoint = 76°F

- **Type 2 heating**: 6am to 8am, setpoint = 70°F
  8am to 6pm, setpoint = 62°F
  6pm to 10pm, setpoint = 70°F
  10pm to 6am, setpoint = 62°F

- **Type 2 cooling**: 8am to 6pm, setpoint = 82°F
  6pm to 8am, setpoint = 76°F

- **"Hold" (no program)**: Heating setpoint = 70°F
  Cooling setpoint = 76°F

The "Hold" setpoint temperatures for heating and cooling represent those who either (1) do not own a programmable thermostat; or (2) own a programmable thermostat but do not program the thermostat. As discussed above, approximately 92% of the population falls into one of these two categories, unless they adjust the temperature manually constantly. Thus, it is assumed that due to the Nest Learning Thermostat features such as learning, time-to-temp, and the Nest Leaf, about 92% of households move from no program schedule to either the Type 1 or Type 2 schedules shown. More accurate values of behavior observed in field trials will be used in future simulations for both baseline condition and learned condition. For example, baseline condition will be changed from "hold" to actual recorded temperature/behavior prior to Nest Learning Thermostat installation. In addition, the learned schedule times and setpoint temperatures will vary from city to city, and schedules may be different on weekends versus weekdays. However, we believe that the simulation using the values shown provides a good starting point.

The simulation results are calculated for both Type 1 and Type 2 occupants, for each of three cases:

- **Learned schedule and setpoint**: Type 1 and Type 2 schedules as shown above. These results show expected energy and costs savings simply by adopting the Nest Learning Thermostat without any away time (e.g. vacations).

- **1°F carving**: Each of the Type 1 and Type 2 occupant types are simulated assuming the user changed his/her behavior based on Nest Leaf display to use 1°F less (for heat) or more (for cool) as their setpoint from the learned schedule. In other words, the setpoint temperatures are all adjusted by 1°F to increase efficiency.

- **Auto-away**: Auto-away that occurs a few hours at a time on daily basis is not simulated here. A wide range of daily auto-away usage is expected, and field trial data will provide an anchoring point for future simulation. For this white paper, both Type 1 and Type 2 occupant types are simulated with two periods of non-occupancy for two weeks at a time (i.e. four weeks of non-occupancy annually). Temperature ranges are selected to match mild to moderate energy conscious households. One two-week absence is simulated in December with safety temperature for heating set to 45°F, and the other two-week absence is simulated in August with safety temperature for cooling set to 95°F. While results shown below reflect these assumptions, simulations with significantly milder temperature choices (62°F for December time away, and 79°F for August time away) are also simulated and commented where appropriate.

Energy Costs: The cost of the energy varies from city to city. The average or most recent energy cost from US Energy Information Administration was used. The values used are in Table below.

City	kWh cost (dollars)	Therms cost (dollars)
Atlanta	0.10	1.56
Boston	0.15	1.47
Chicago	0.10	0.94
Dallas	0.11	1.01
Denver	0.10	0.81
Houston	0.11	1.01
Miami	0.12	1.81
Minneapolis	0.10	0.87
Phoenix	0.10	1.59
San Diego	0.15	0.98
San Francisco	0.15	0.98
Spokane	0.08	1.20

4. Results

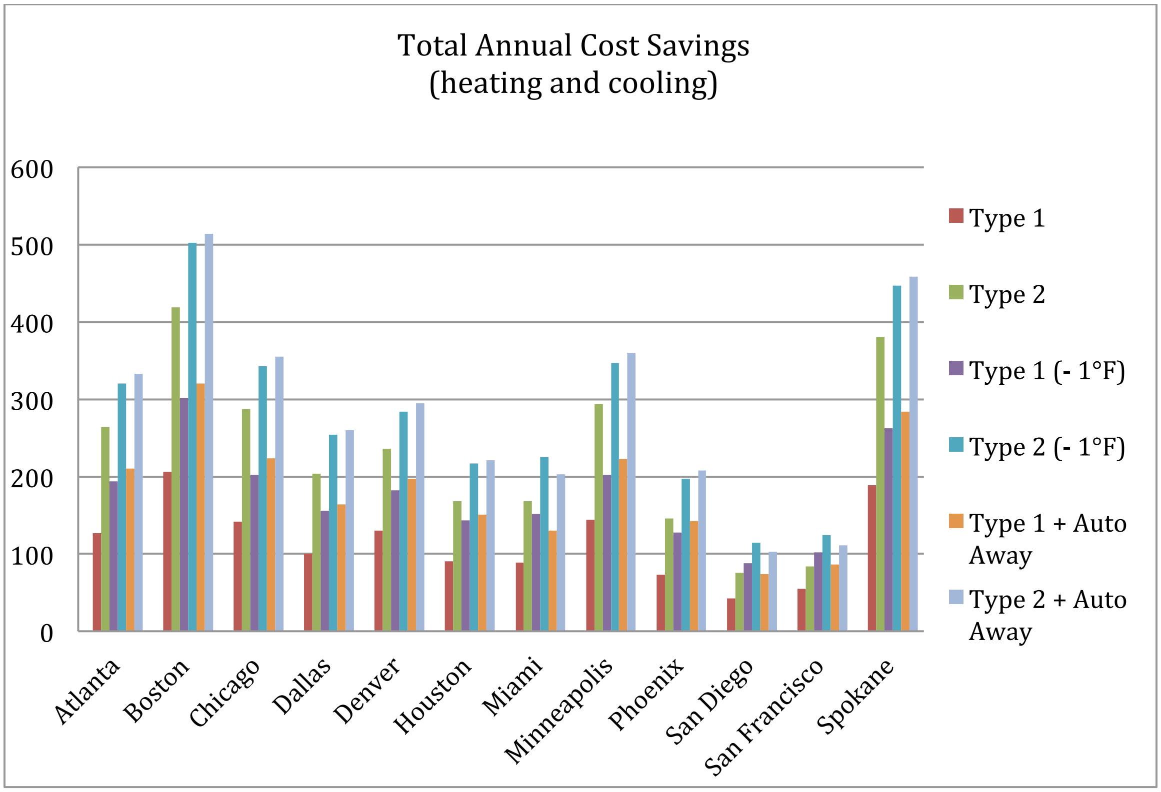

Figure 2: Total Annual Cost Savings (heating and cooling)

Figure 2 shows a summary of simulated annual cost savings for Type 1 and Type 2 occupants with the setpoints as described in Section 3.2, along with savings for the 1°F carving and for the auto-away periods. The costs reflect the calculated savings when compared to the non-programmed "hold" set points. As can be seen, some cities such as Boston and Spokane have much greater savings than others, such as San Diego and San Francisco. This is largely due to the fact that the overall usage and associated costs vary widely by geographic location. As expected, the Type 2 occupants, who are absent from the house for long predictable periods during the day, benefit the most from the Nest Learning Thermostat. Additionally, the 1°F carving in set point temperatures leads to substantial savings in all cases. Finally, the Auto-away feature allows for even greater cost savings when the occupants are away during the two two-week periods each year.

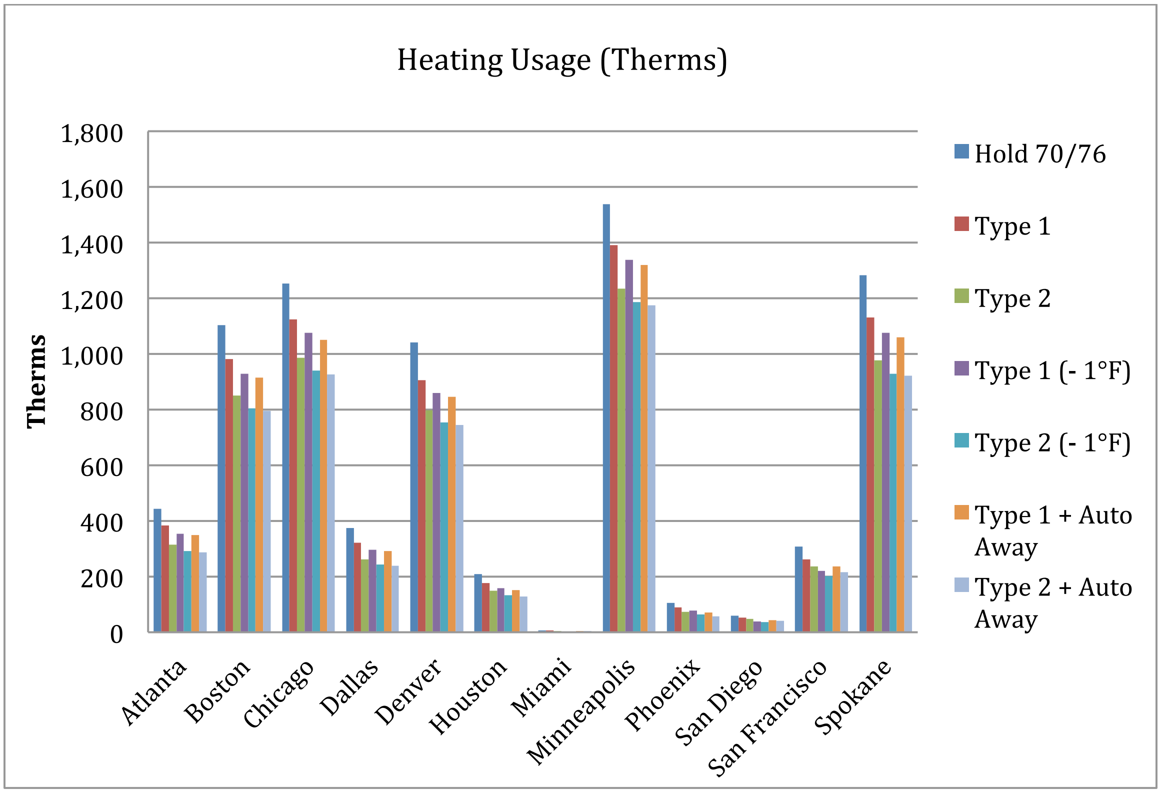

Figure 3: Annual Heating Usage in Therms

Figure 3 shows the simulated annual energy usage for heating in therms for Type 1 and Type 2 occupants having a basic schedule, 1°F carving, and with auto-away absences. Also shown for comparison is the "Hold" baseline households that have no program. Less energy is used by the Type 2 households in general, and for households using 1°F less as well as those with an auto-away detected absence simulated. It is also notable that Miami, Phoenix and San Diego had very low annual energy usage for heating.

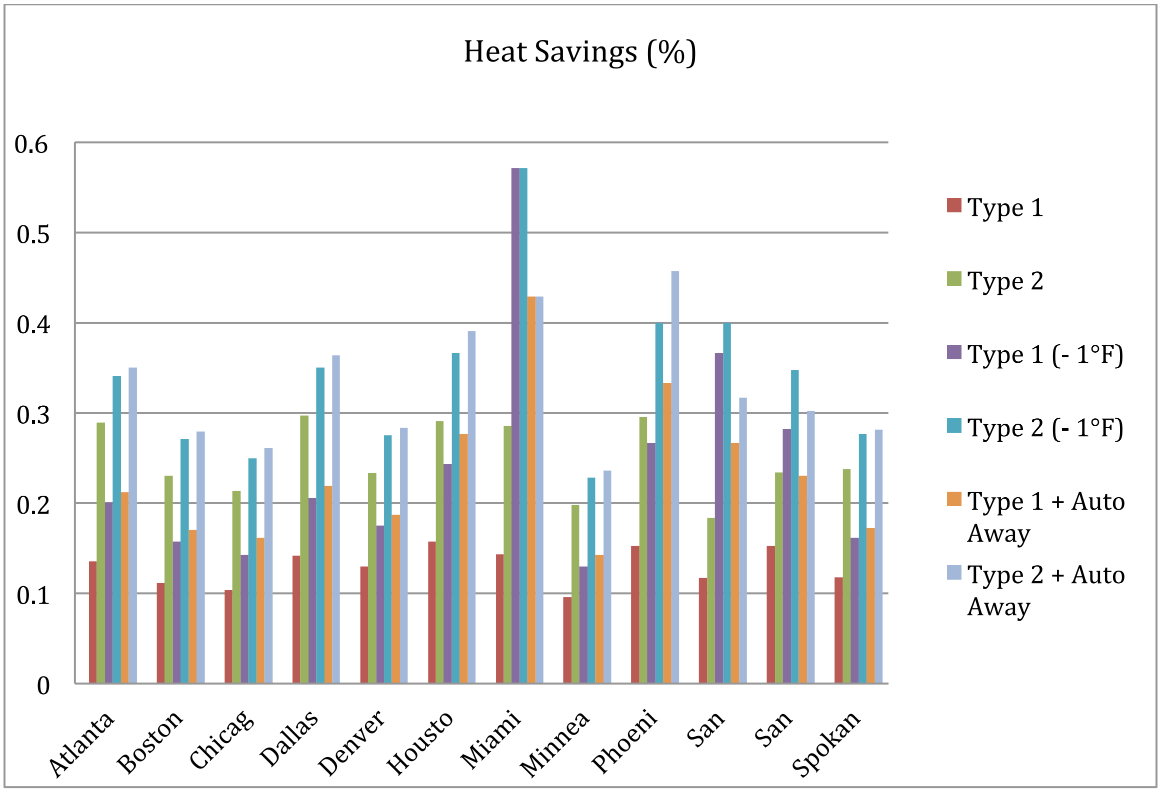

Figure 4: Annual Heating Savings in Percent

Figure 4 shows the simulated annual heating energy savings in percentage when compared to the "hold" baseline, for Type 1 and Type 2 occupants having a basic schedule, 1°F carving, and with auto-away absences. As expected greater energy savings are realized by the Type 2 households in general, and for households using 1°F less as well as those with an auto-away detected absence simulated. Note that the values for Miami are not very significant since that city has very little heating usage in any case.

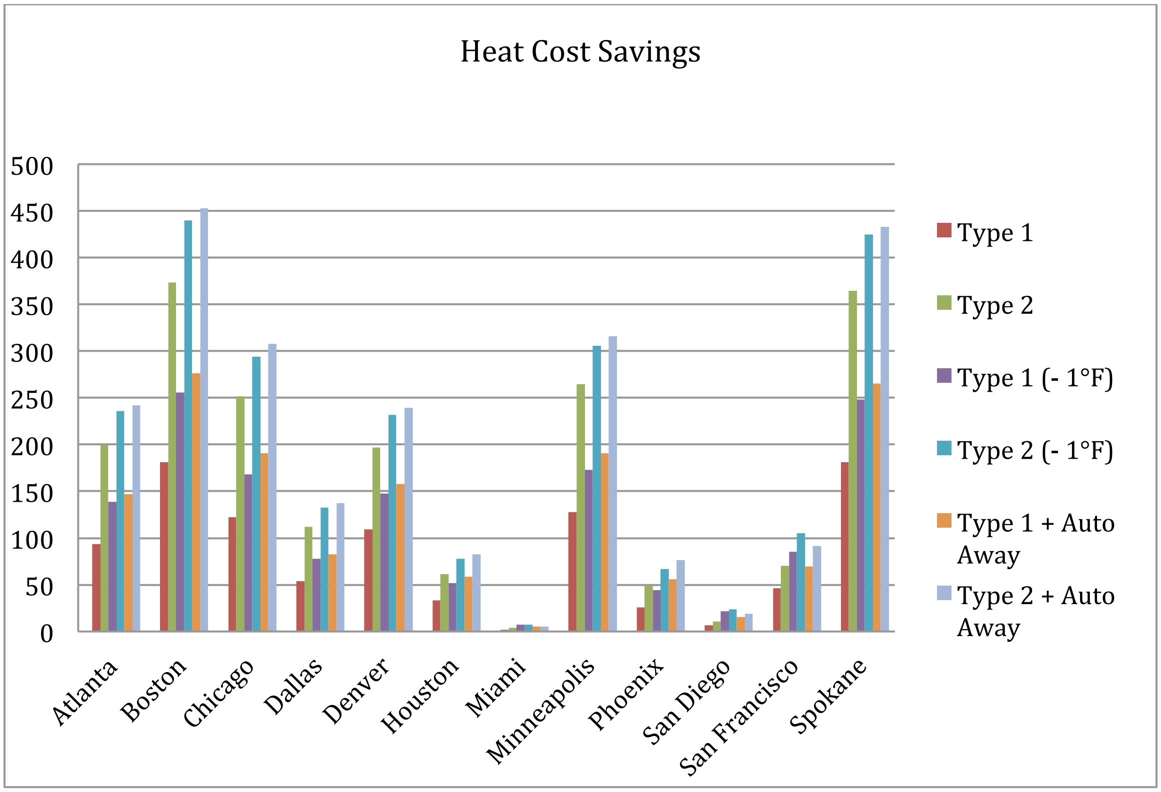

Figure 5: Annual Heating Cost Savings

Figure 5 shows the simulated annual heating cost savings when compared to the “Hold” baseline households, for Type 1 and Type 2 occupants having a basic schedule, 1°F carving, and with auto-away absences. The costs for energy vary by city and range from \$0.81 per therm in Denver, to \$1.59 per therm in Phoenix. As expected, greater savings are realized by the Type 2 households in general, and for households with 1°F less as well as those with an auto-away detected absence simulated.

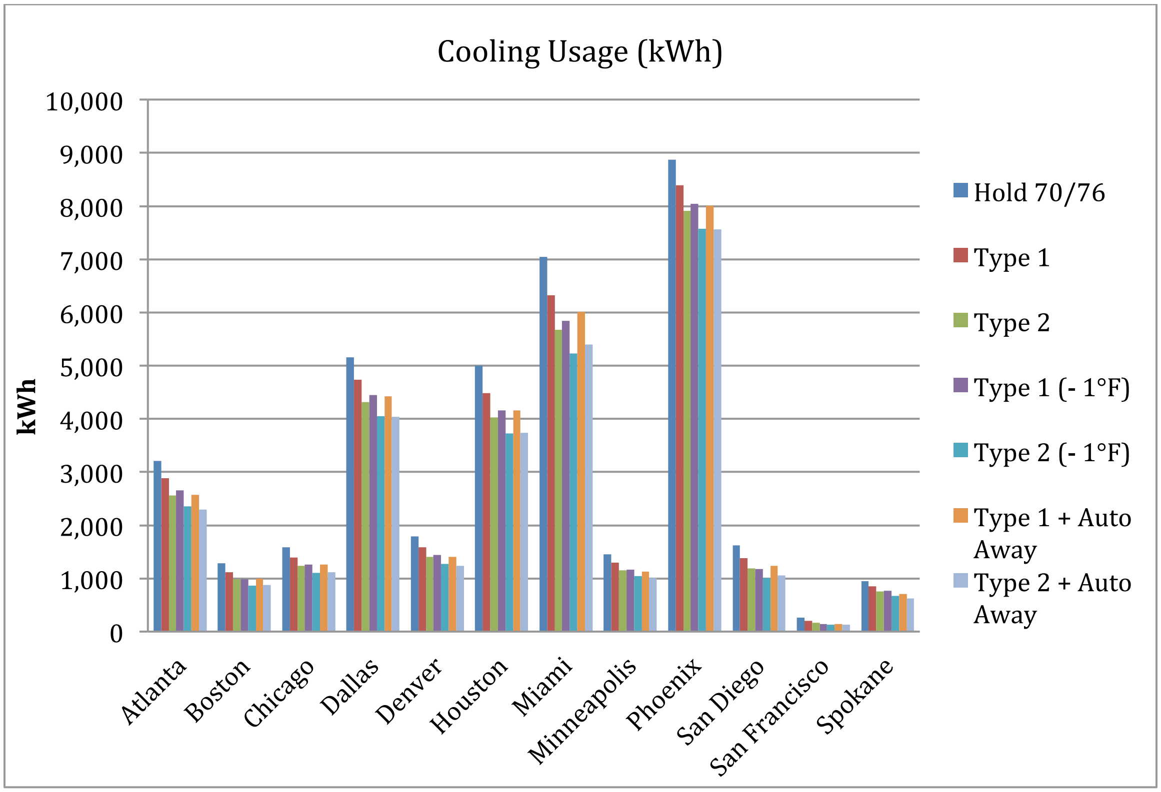

Figure 6: Annual Cooling Usage in kWh

Figure 6 shows the simulated annual energy usage for cooling in kWh for Type 1 and Type 2 occupants having a basic schedule, 1°F carving, and with auto-away absences. Also shown for comparison are the "Hold" baseline households that have no program. Less energy is used by the Type 2 households in general, and for households with 1°F less as well as those with an auto-away detected absence simulated. It is also notable that San Francisco has very low annual energy usage for heating in general.

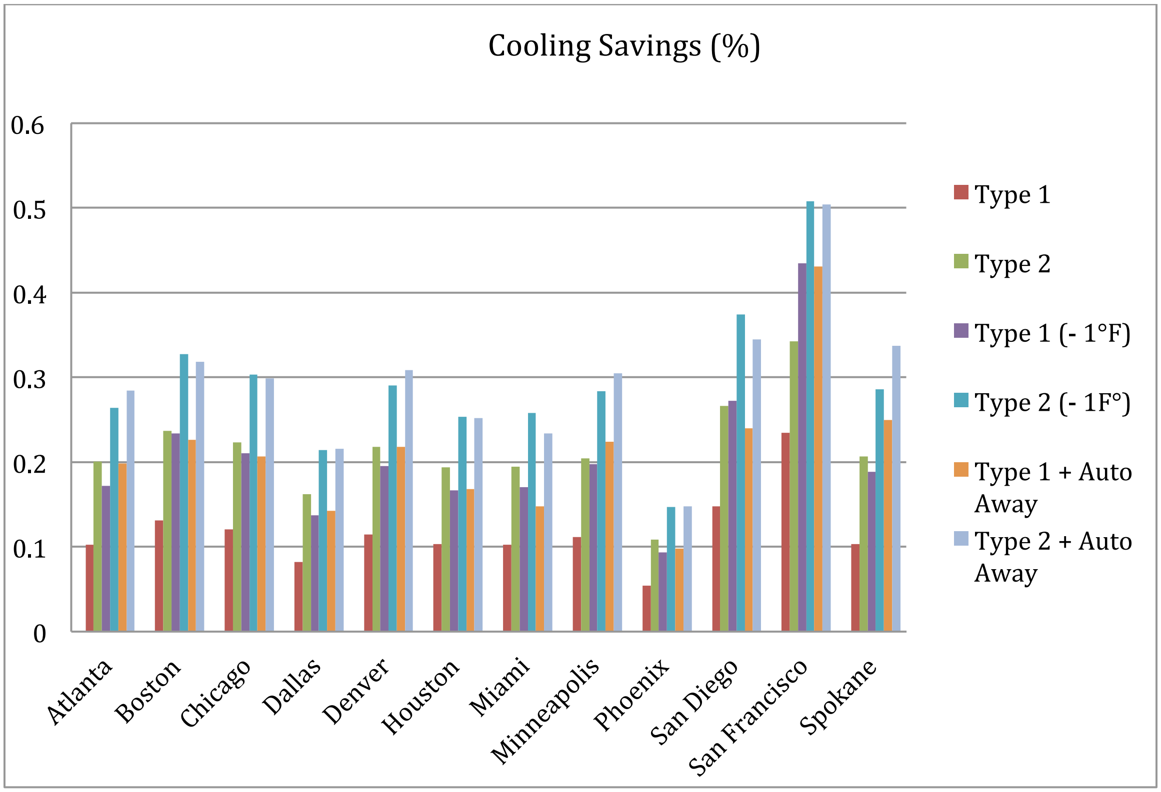

Figure 7: Annual Cooling Savings in Percent

Figure 7 shows the simulated annual cooling energy savings in percentage when compared to the "hold" baseline for Type 1 and Type 2 occupants having a basic schedule, 1°F carving, and with auto-away absences. As expected greater energy savings are realized by the Type 2 households in general, and for households with 1°F less as well as those with an auto-away detected absence simulated. Note that the values for San Francisco are not very significant since that city has very little cooling usage in any case.

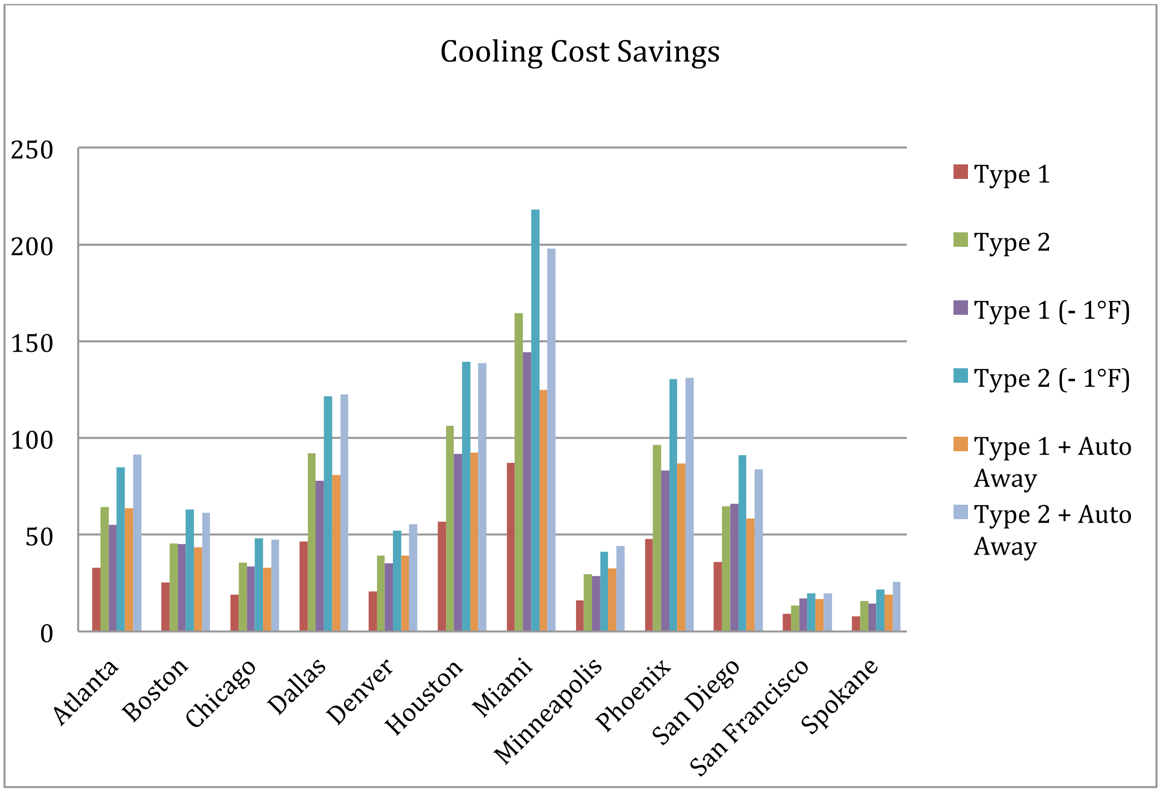

Figure 8: Annual Cooling Cost Savings

Figure 8 shows the simulated annual cooling cost savings when compared to the "Hold" baseline households, for Type 1 and Type 2 occupants having a basic schedule, 1°F carving, and with auto-away absences. The costs for energy vary by city and range from $0.08 per kWh in Spokane, to $0.15 per kWh in Boston, San Diego and San Francisco. As expected, greater savings are the Type 2 households in general, and for households with 1°F less as well as those with an auto-away detected absence simulated.

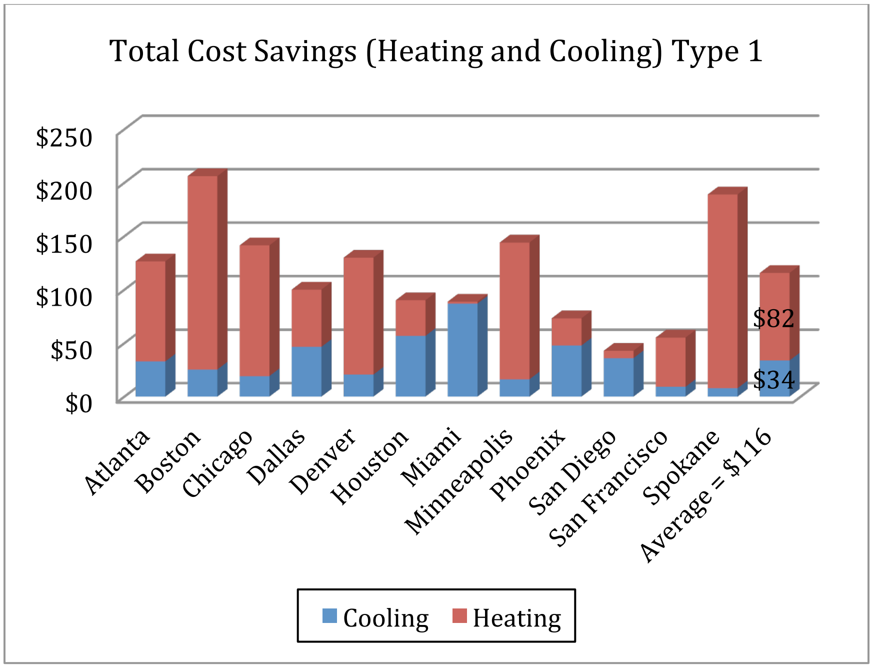

Figure 9: Annual Total Cost Savings - Type 1

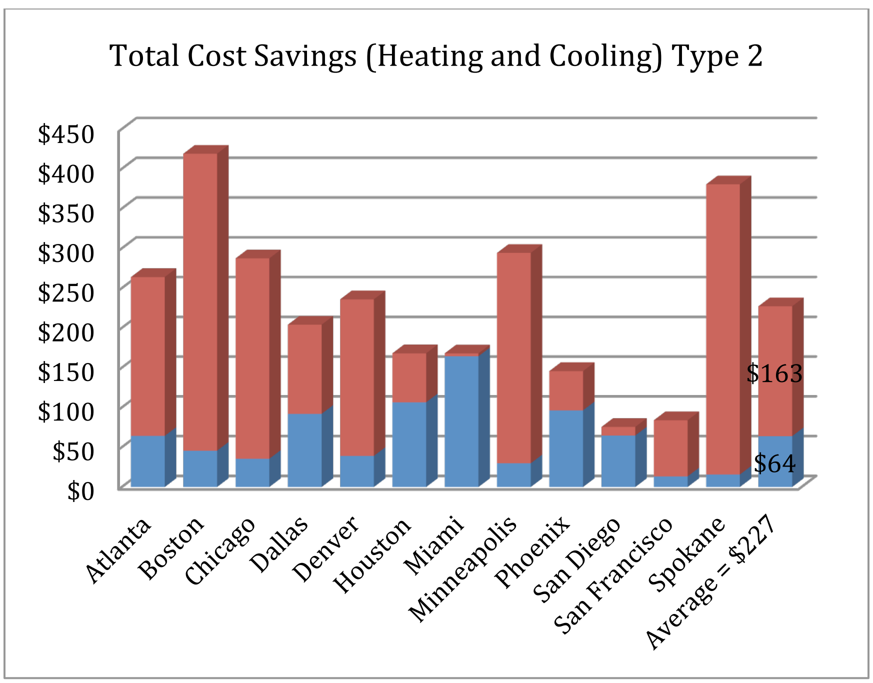

Figure 10: Annual Total Cost Savings - Type 2

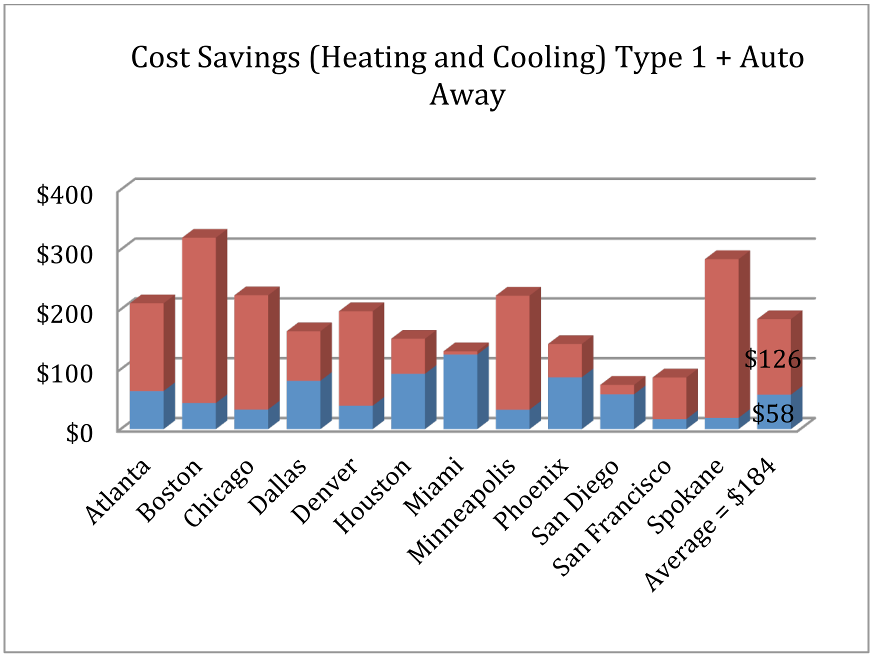

Figure 11: Annual Total Cost Savings - Type 1 with Auto-away

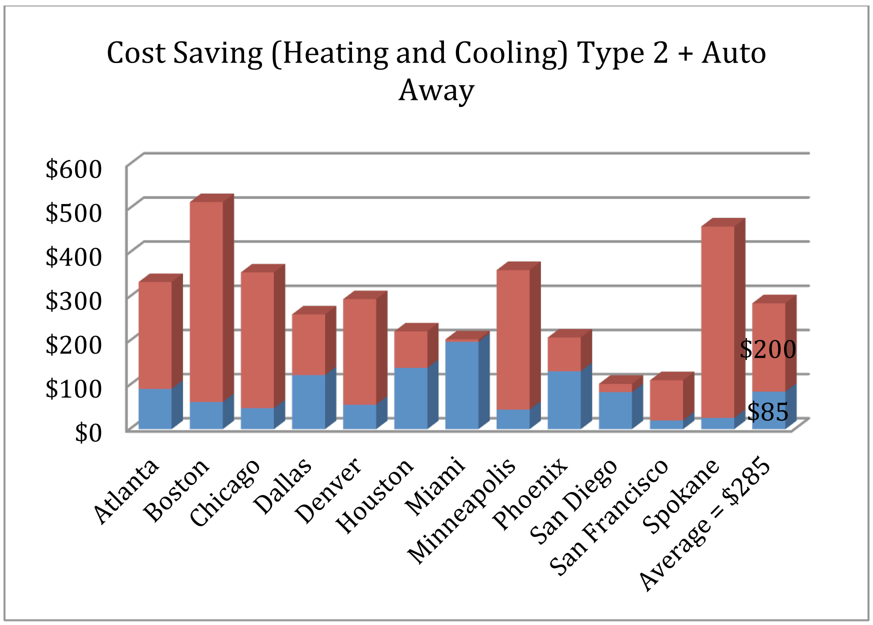

Figure 12: Annual Total Cost Savings - Type 2 with Auto-away

Figures 9 through 12 show the simulated annual costs savings for heating and cooling for Type 1 and Type 2 occupants without and with auto-away. In each of the Figures, costs for heating are shown in red and costs for cooling are shown in blue. It is notable that for most cities heating costs savings are greater than cooling costs, with the exceptions being Houston, Miami, Phoenix and San Diego. The savings for Dallas are nearly evenly split between heating and cooling. Average annual cost savings ranges from $184 for Type 1 occupants to $285 for Type 2 occupants with an auto-away absence.

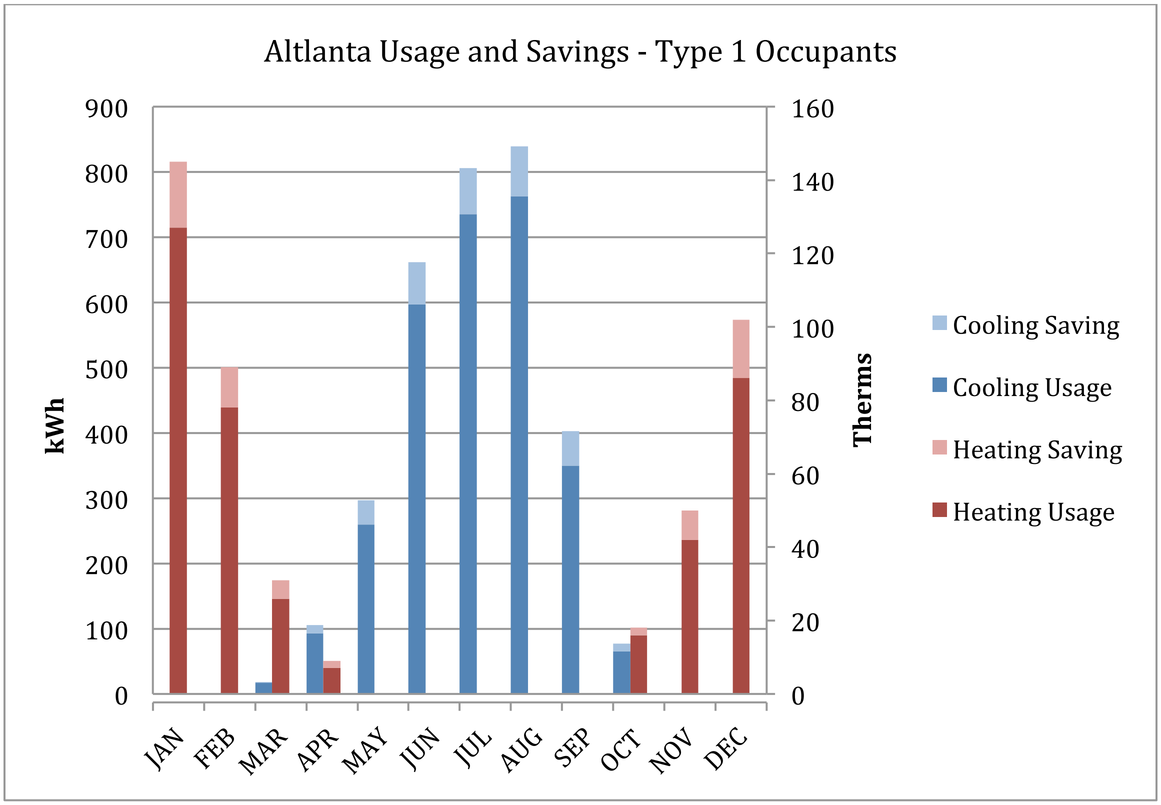

Figure 13: Monthly Energy Usage and Savings – Atlanta, Type 1

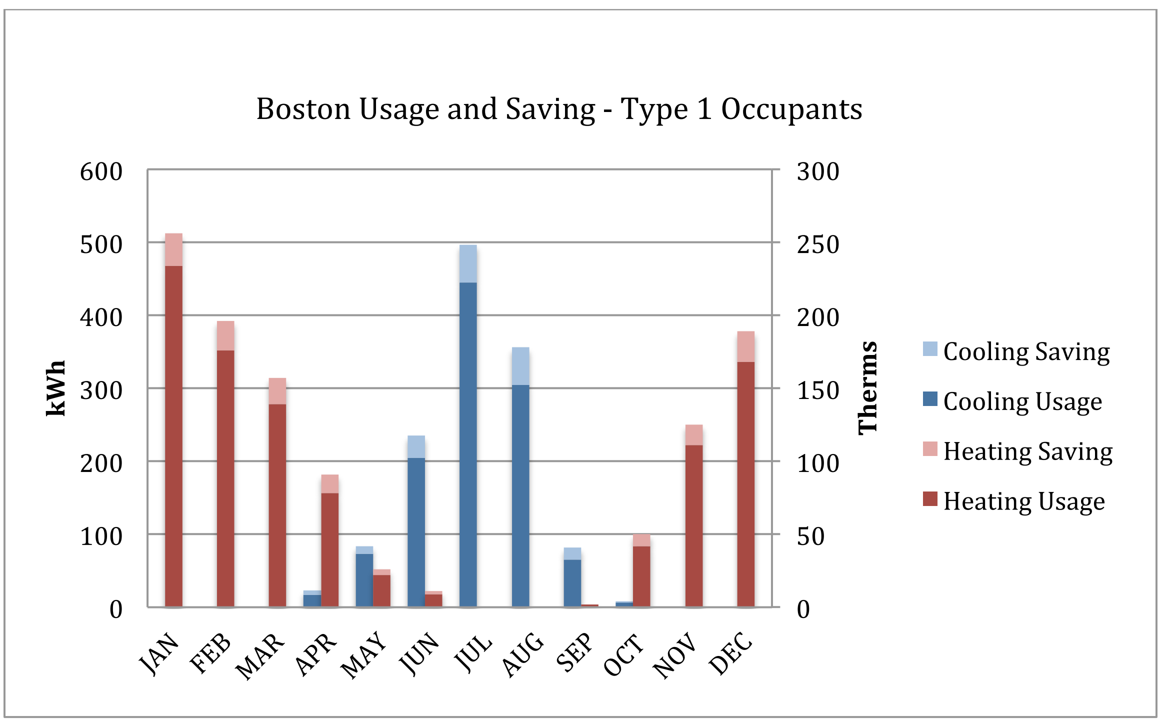

Figure 14: Monthly Energy Usage and Savings – Boston, Type 1

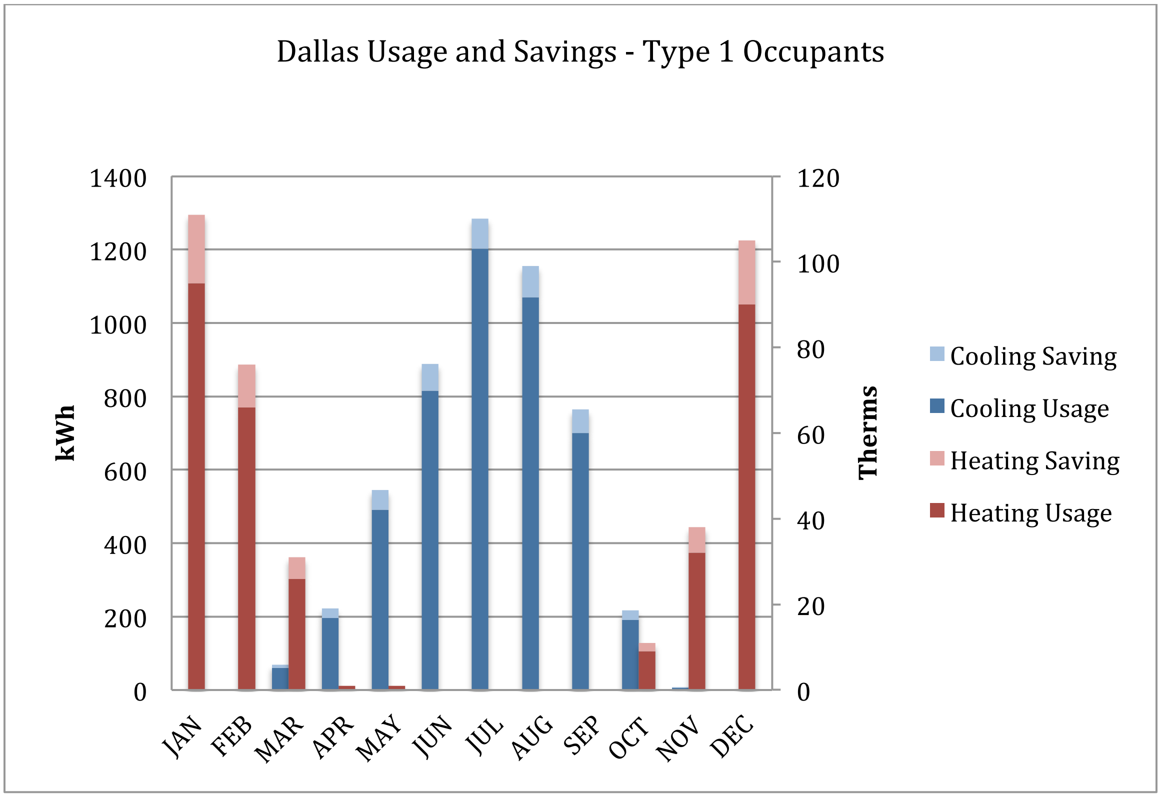

Figure 15: Monthly Energy Usage and Savings – Dallas, Type 1

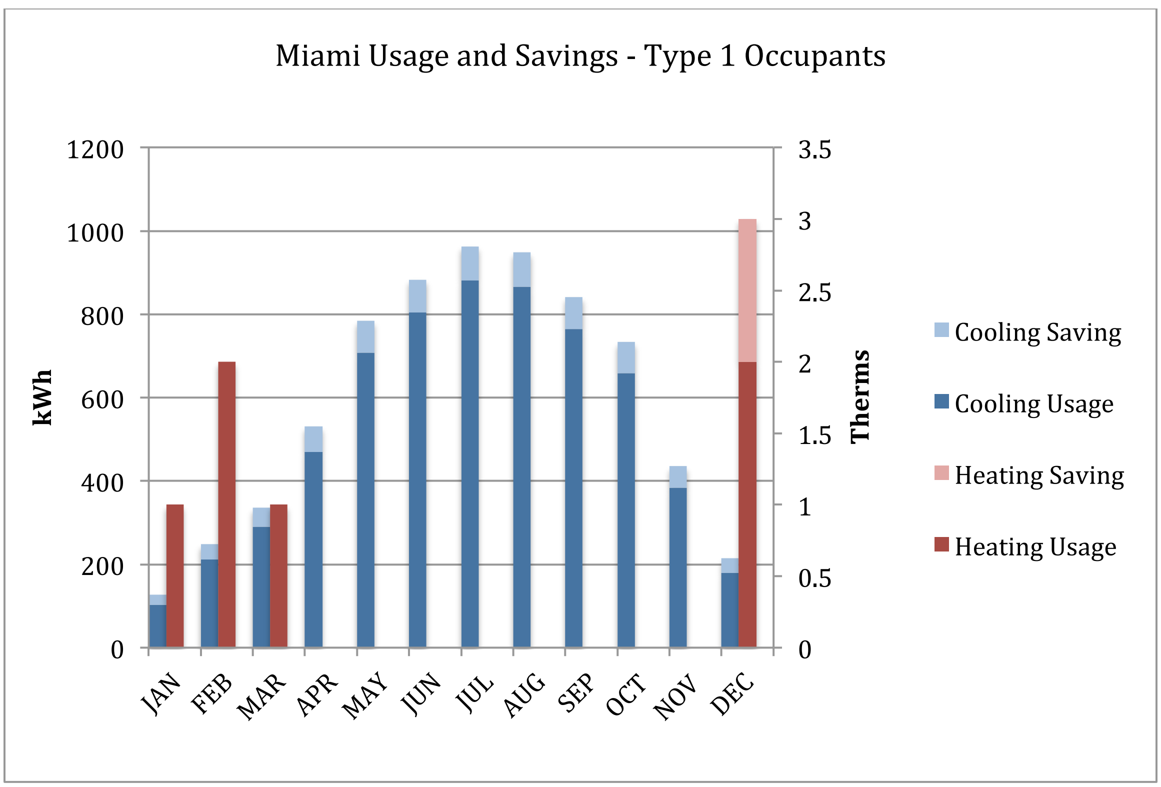

Figure 16: Monthly Energy Usage and Savings – Miami, Type 1

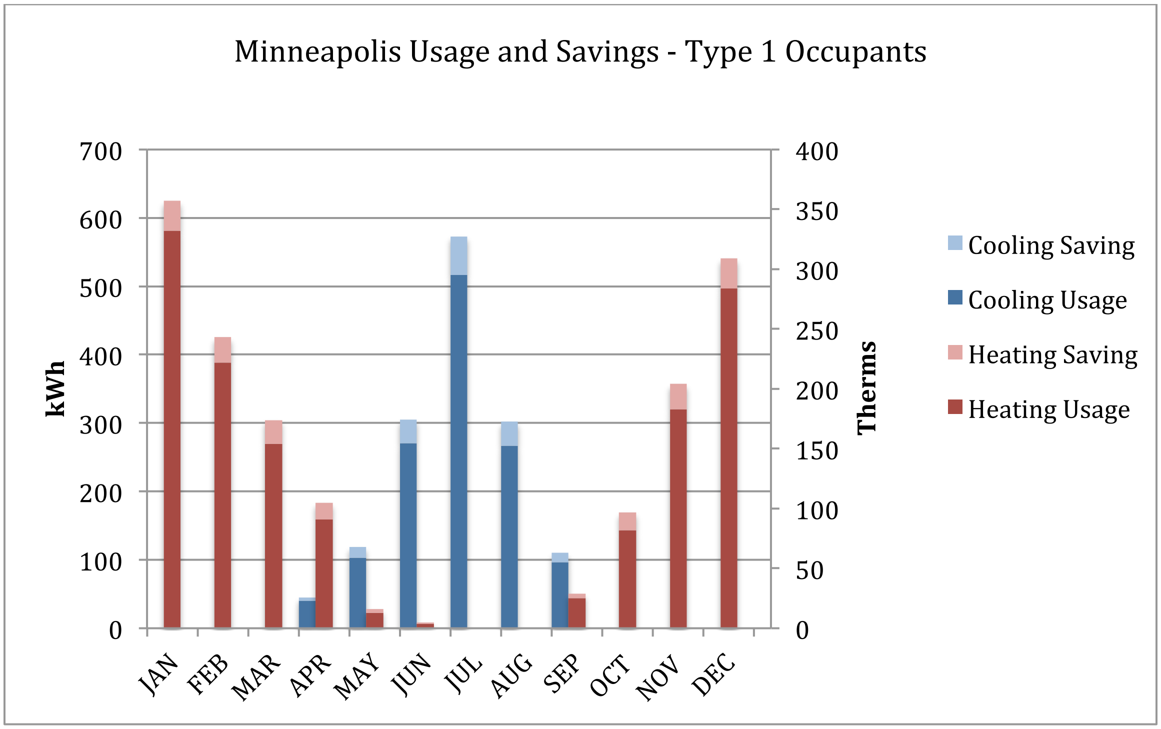

Figure 17: Monthly Energy Usage and Savings – Minneapolis, Type 1

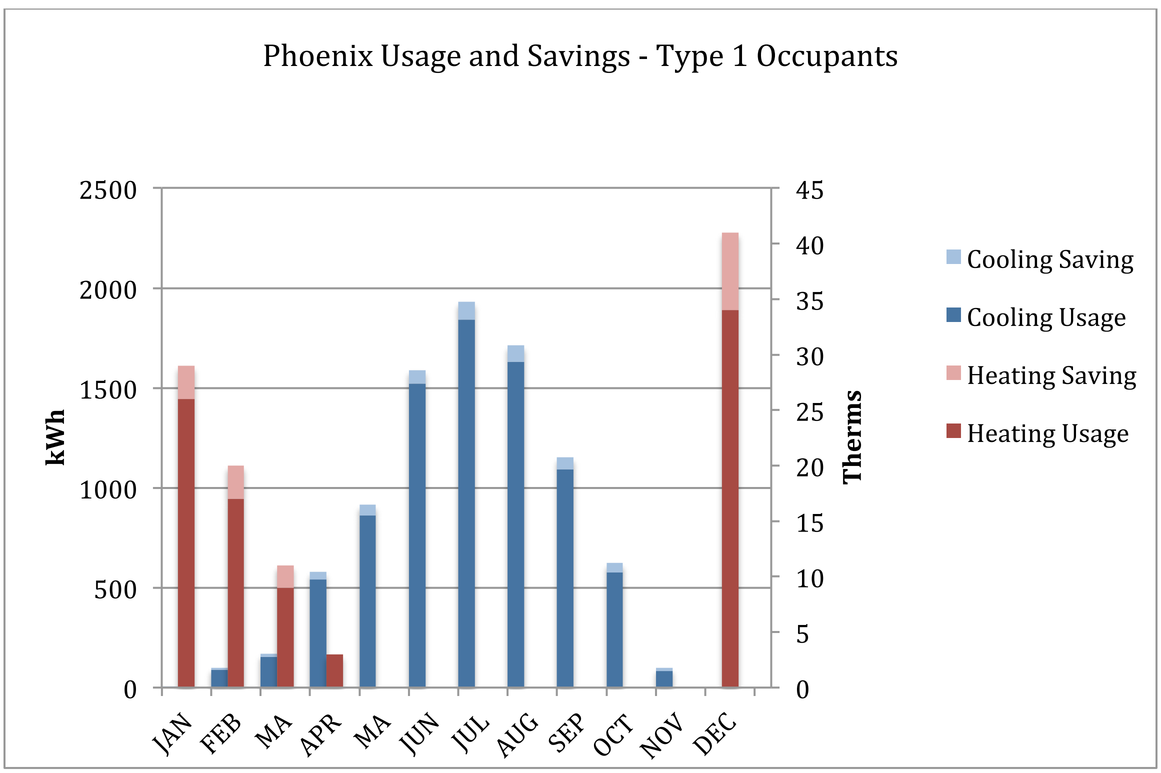

Figure 18: Monthly Energy Usage and Savings – Phoenix, Type 1

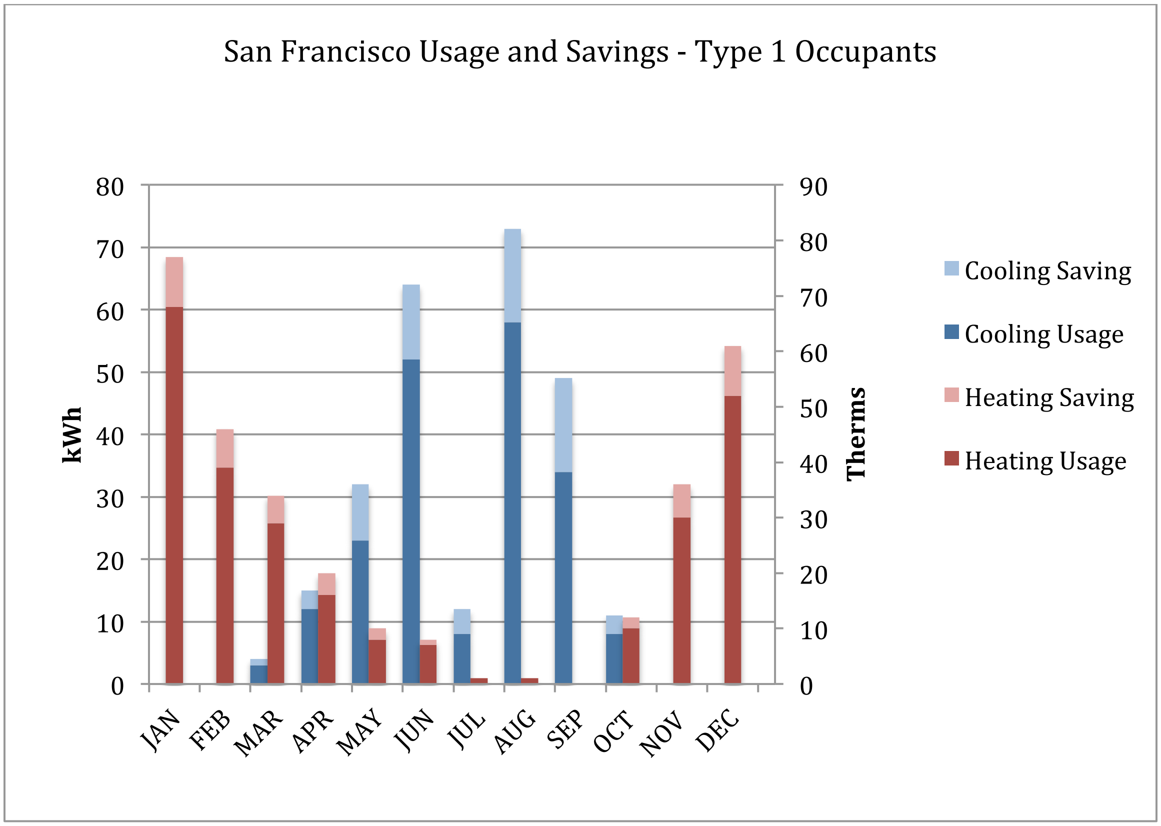

Figure 19: Monthly Energy Usage and Savings – San Francisco, Type 1

Figures 13 through 19 show the simulated monthly energy usage and savings for heating and cooling for Type 1 occupants for selected seven cities. In each of the Figures, costs for heating are shown in red and costs for cooling are shown in blue. Lighter red and blue show savings.

5. Summary of Simulation Findings

Average annual savings over the twelve modeled cities for those whose prior practice was not to program their thermostats is $116 for occupants adopting a single set back (Type 1) and $227 for occupants adopting a dual setback (Type 2) both without any vacation period accounted for. If two 2-week away periods are simulated, the average annual savings raises to $184 for Type 1 occupants and $285 for Type 2 occupants. With milder two 2-week away period temperatures the average annual savings is $133 for Type 1 occupants and $235 for Type 2 occupants. For occupants taking 1°F off (lowering set point by 1°F for heating and raising the set point by 1°F for cooling) but without any vacation period, the average savings is $176 for Type 1 occupants, and $281 for Type 2 occupants.

Using the simulation findings and weighting them across all U.S. zip codes, it's estimated that the Nest Learning Thermostat can save an average of $173 per year, before applying any of Nest's additional energy-saving features such as Auto-Away™ and the Leaf.

Single versus Dual Setback: On average, about twice as much cost saving ($112, 97% increase in savings) is realized when moving from a single setback (Type 1) to a dual setback (Type 2) for occupants who do not have a vacation absence. The increase in savings from a single to dual setback is less both in dollars and percentage increase when the occupants have already accepted to use 1-degree less for setpoint ($105, 60% increase in savings) or have four weeks of away time accounted for ($101, 55% increase in savings).

1°F Setpoint Change: Average increased annual savings for accepting a 1°F change in setpoint values is $60 or 52% savings increase for occupants having a single setback (Type 1), and $54 or 24% savings increase for occupants having a dual setback (Type 2). The smaller amount of increased savings for dual setback households makes sense since those occupants use less energy to begin with.

Auto-away savings: Average increased annual savings for accounting for two 2-week absences from the home (one in August and one in December) is $68 or 59% savings increase for occupants having a single setback (Type 1), and $58 or 25% savings increase for occupants having dual setbacks (Type 2). With milder temperature choices, $17 or 15% savings increase for occupants having a single setback (Type 1), and $8 or 4% savings increase for occupants having dual setbacks (Type 2). Similar to the case above, the smaller amount of increased savings for dual setback households makes sense since those occupants use less energy to begin with.

People who already program their thermostats: Occupants who already have a suitable programmed thermostat (about 8% of households) will still benefit from features such as the 1-degree carving, and auto-away detection. The average savings over the basic program for the 1-degree change in set back is $60 (6%), and $54 (7%) for single and dual setback occupants respectively. The average saving over the basic program for the two 2-week absences is $68 and $58 for single and dual setback occupants, respectively. For both single and dual setback occupants, the auto-away feature leads to a 7% average cost saving.

6. Model Calibration Using Data from the First Three Months

6.1 Sampled Data

Data from January-February 2012 was collected and analyzed. For purposes of this study, installed Nest Learning Thermostats were excluded from the data if they met one or more of the following exclusion criteria:

- Device was not connected to Wi-Fi, or the Wi-Fi connection is intermittent (i.e. <90% connection time).

- Device was first installed less than 10 days ago.
- Device did not have a valid U.S. ZIP code entered.
- Device controlling heat pump, electric heat, and second stage heat.
- Device in mild climate (hdd65 < 3000) or device not in heating mode for at least 99% of the period of study.
- Local weather information was missing for more than 3 hours a day in the location where the device was installed.

After excluding these devices, a sample of more than 10,000 devices, selected from the Nest Learning Thermostat installed base, were used for the analysis below. Personally Identifiable Information, if any, was removed before the data was processed.

6.2. Definition of a Good Schedule for Purposes of Study

A schedule with significant setbacks can provide greater comfort while occupants are home and active, while saving energy while the occupants are away or asleep. The Nest Learning Thermostat was designed specifically to facilitate energy savings through setback scheduling. In fact, 98.6% of Nest Thermostats have a setback schedule (compared to the national average of approximately 10%).

To assess the quality of Nest's setback schedules, we had to define an "energy-efficient schedule." There are many ways to do this, but in this paper, we define an energy-efficient schedule in terms of significant setpoint setback(s). More specifically, the following criteria was used to define devices with an energy-efficient schedule:

- The devices have a setback of at least 8°F for at least 4 hours, for at least 94% of the days sampled.
- The devices have actual effective indoor temperature reduction. The criterion is to have the average reduction of the indoor temperature to be more than 8°F during day and night setback period. A device could satisfy this criterion with a single setback with a very deep actual temperature reduction or two setbacks with smaller reductions.
- The devices have indoor temperature reduction due to the setback. The criterion is to have the average reduction of the indoor temperature to be more than 8°F total. A device could satisfy this criterion with a single setback with a very deep actual temperature reduction or two setbacks with smaller reductions.

6.3. Model Calibration Methods

To capture the estimated savings of these energy-efficient devices, the model was calibrated with the following techniques:

1. Calculate model parameters to match each device behavior.
2. Calculate fixed setpoint (90th percentile of the device setpoints) for baseline model.
3. Run simulation using input adjustments as developed in step 1 and again using baseline setpoints set in step 2.
4. Summarize simulation results for calibrated simulation and for baseline variation.

6.4. Estimated Savings

Devices with an energy-efficient schedule, as defined in Section 6.2, had an average savings per device of 19.5%, with the highest saving of 36.1%. This is consistent with the EPA's estimate that a properly programmed thermostat can save 20% on heating and cooling costs.

Interestingly, the data shows that users who adopt sufficient setbacks are also more likely to adopt energy-saving maximum setpoint temperatures as well.

7. Future Simulation Improvements

Section 6 was conducted when the first set of field data became available. The simulation model used in Sections 2-5 was not fully updated based on data. Based on our customer-based distribution, this adjustment will be made in the future. Example of these improvements are listed below:

- Additional types of HVAC equipment can be added to increase accuracy of the results. These other types of heating systems will be weighted according to the prevalence of use in different cities. Additionally, multi-stage heating and cooling systems will be added to the model.
- Many houses in the U.S. are larger than (or smaller than) 1,800 sqft. Energy usage generally scales with the size of the home. Percentage savings remain approximately the same while the dollars savings will scale proportionally to the size of homes. With multiple thermostats installed at different locations in the house, we will simulate more accurate model representing different home sizes and number of thermostats in each home.
- When calculating average savings, the population of different cities was not accounted for in sections 2-5. As the model becomes more accurate, it will be normalized by the regional population as well.
- A range of building characteristics that affect efficiency will be varied. The current model assumes an exactly the same house in every city. The type/model of homes will be varied and matched to regional differences.