Fiddling with Thermostats: Energy Implications of Heating and Cooling Set Point Behavior

James Woods, Portland State University

ABSTRACT

Household energy use depends heavily on how people set their thermostats. In fact, evaluations of the benefits of new cooling and heating technologies often assume specific thermostat behaviors, or set points. California's Title 24 Standards, for example, assume a certain range of settings and frequency of daily changes in those settings. Until recently, data have not been available to test such assumptions. In 2001-02, the California Energy Commission conducted a demand response experiment that produced unique, high frequency observations of residential thermostat settings and internal temperature measurements, which allow testing of assumptions about thermostat behaviors.

Comparing the thermostat settings observed in the California experiment with those commonly assumed in policy modeling indicates that people change cooling and heating set points much more frequently than has been assumed. Frequent set point changes, and the extreme diversity of set point behavior across the population, have significant energy implications. This paper uses Shannon Entropy to assess consistency of thermostat settings, which can produce both higher and lower levels of energy consumption than is conventionally assumed. The findings call into question the benefits of energy efficiency programs that focus on equipment replacement and choice.

Introduction

Modern thermostats allow for sophisticated temperature control. Even basic programmable thermostats allow household members to establish different heating and cooling set points for different times of the day, whether they are home or not. Policy-related research assumes a limited level of technical control. Households are assumed to set a maximum and minimum temperature and to leave the temperature within that range. This relatively simple level of technical control is embodied in many common residential energy simulators, notably CALRES (California Energy Commission 2001) and the DOE-2 derivatives that have been benchmarked with HERS BESTEST (Judko & Neymark 1995).

Many thermostats allow for more complicated, time-based setbacks and even the ability to react to events such as extreme weather or price signals generated by the local utility[1]. In spite of the increasing sophistication of technical controls, the simulation models used in California Title 24 verification (California Energy Commission 2001) and HERS verification still assume relatively unsophisticated thermostat behaviors[2]. This paper will show that households are much more active managers of internal temperatures through sophisticated use of their thermostats than has previously been supposed in models.

The intent is to expose flaws in the assumptions used by these common simulators to estimate the expected energy use in buildings and expected energy savings from changes. Once this deviation from assumed behavior is documented, it is the responsibility of those that use,

[1] The RCS model TL36, for example.

[2] The thermostat control in HERS BESTEST is simpler than in CALRES and specifies constant set points for both heating and cooling

support and maintain those models to show that the assumptions about thermostat use produce accurate estimates of *expected* energy use even when the assumptions are far different from observed behavior.

The data used in this paper were originally generated by Energyn as part of a peak load reduction experiment sponsored by the California Energy Commission (Enerygyn, 2004). These data are high frequency, 15-minute observations of approximately 96 houses over a period of a year. The next section will compare the observed thermostat settings with those described in California Title 24. The consistency of those settings will be characterized with Shannon Entropy. It will be shown that households, in both a statistical and a practical sense, act differently with respect to both the level of thermostat settings and the consistency of those settings than is assumed in California Title 24.

Data

The Energyn project "Field Deployment of the iPower Advanced Energy Efficiency and Information System," (Enerygyn, 2004) begun in the summer of 2002, had as its core objectives:

- Demonstration of a broadband-based residential demand response system;
- Integration of a web-based energy information system with the demand response system;
- Development of an effective approach to training people to use web-based information as part of a strategy to control their rate of energy consumption;
- Development of a platform that would allow additional control and information technologies within the home to be integrated with a demand response system;
- Development of a database of information on customer thermostat behavior;
- Exploration of how such a system might provide multiple sources of value to consumers and potential business partners, assisting in the development of a fully or partially market-based solution.

Households in Pacific Gas and Electric's territory, mostly in the San Francisco Bay Area, were recruited at events, such as home improvement shows, that would attract technology-oriented individuals with an interest in environmental improvement. Additional participants were recruited via cold calls to the target population.

The collection of household temperature and thermostat settings was not the focus of the study design, but part of a field trial of an energy information system. The study monitored thermostat behavior to see whether households would react to a simple experimental treatment that did not change the energy price structure in any meaningful way. The only financial incentive that may have altered behavior was entry in a lottery to win an energy efficient clothes dryer. Participants would be entered in the lottery if they changed their consumption behavior in one of *three* curtailment events, which took place on *two* days, July 28 and September 22, 2003. This represents a relatively small financial incentive. At the same time, the study increased the quality of information households had about their energy use, which might have affected their consumption behavior.

The sub-sample used in this paper consists of 96 houses with a total of 783,459 useful observations of thermostat settings and internal household temperatures. A small subset of households with fewer than 1,000 total observations is not included in this analysis. Observations were usually taken every 15 minutes. Periodic gaps in the data did not affect this analysis.

Figure 1 shows the distribution of observations per household from February 1, 2003, through December 23, 2003. Because of a tailing off of the field experiment, data from December were not used in the analysis.

Figure 1. Distribution of Observations Per Household

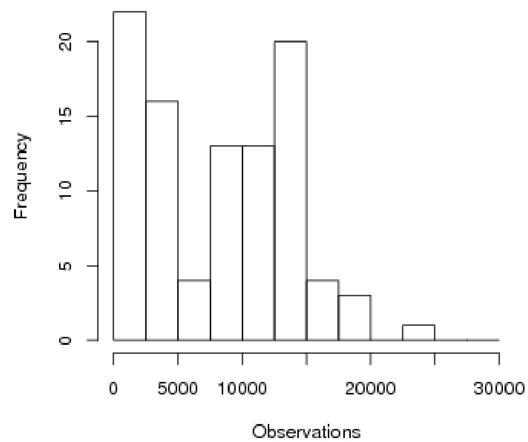

Comparison to Title 24 Assumptions

The modeling assumptions used in California's Title 24 verification specify the set point schedule shown in Tables 1 and 2. These are the thermostat settings used in residential building simulations to determine whether the building satisfies state energy efficiency standards.

The tables show Title 24 set point assumptions and the average set points across households in our dataset conditional on month.

Table 1. Cooling Set Point (F) by Month

Hour	Title 24	Feb	Mar	Apr	May	Jun	Jul	Aug	Sep	Oct	Nov
1 – 7	78	79.4	81.1	81.6	79.1	78.1	78.4	77.4	78.4	78.8	79.2
8 - 13	83	81.5	82.6	82.7	80.1	78.3	78.6	77.5	78.6	79.3	79.4
14	82	81.5	82.8	82.8	80.1	77.9	78.4	77.2	78.5	79.2	79.4
15	81	81.5	82.5	82.9	80.2	77.9	78.4	77.2	78.5	79.1	77.6
16	80	81.5	82.7	82.6	80.1	77.7	78.1	76.8	78.2	78.8	77.6
17	79	81.5	82.6	82.5	80	77.5	78.3	76.9	78.2	78.7	79.1
18-24	78	79.8	81.4	82.1	79.2	77.6	78.3	77.2	78.3	78.8	79.0

Table 2. Heating Set Point (F) by Month

Hour	Title 24	Feb	Mar	Apr	May	Jun	Jul	Aug	Sep	Oct	Nov
1 - 7	65	66.3	67.1	67.6	66.4	65.6	66.5	65	65.2	65.3	65.3
8 - 13	68	64.6	66.4	67.8	66.9	66.3	66.8	65.5	65.7	65.8	67.0
14	68	64.5	66.6	67.8	66.7	66.2	66.8	65.9	65.6	65.8	66.7
15	68	64.5	66.4	67.9	66.8	66.2	66.9	65.5	65.7	66.0	67.3
16	68	64.5	66.5	68.0	66.8	66.6	67.4	66.1	66.3	66.5	69.9
17	68	64.5	66.4	67.9	66.9	66.6	67.5	66.4	66.5	66.7	68.3
18- 24	68	66.0	67.2	68.8	67.7	66.8	67.5	65.7	66.1	66.2	66.6

Unsurprisingly, every one of the deviations from the Title 24 assumptions shown in Tables 1 and 2 is statistically significant. With the large number of observations, it is virtually impossible to have statistically insignificant differences.

Table 1 shows that during the primary cooling months, June through September, the average cooling set point is virtually the same throughout the day. In each of those months, the average set point does not vary more than one degree during the day. In July, for example, the highest average set point is 78.6F, between the hours of 8am and 1pm. The lowest setting is 78.1F, occurring just a few hours later, at 4pm.

The table also shows that, on average, households are using a lower set point than is assumed in Title 24 assessments. The difference between Title 24 assumptions and the average observed setting is greatest during peak hours, when electricity system costs are at a premium.

Table 2 shows a similar pattern of stable heating set points throughout the day for months when heating is used significantly. In contrast with the cooling set points, these thermostat settings are lower than the Title 24 assumptions.

These differences between what is assumed and what is observed could be pushed off as the result of differences between observed and average external temperatures – an adaptive thermostat setting model. This point of view produces several contradictions with current Title 24 practice.

First, adaptive thermostat settings cannot explain the near constant thermostat setting throughout the day in each month. The observed settings do not track the diurnal swing.

Second, a single daily program of set points would only be appropriate if the thermostat setting was a linear function of the external temperature. In that case, observed weather would be irrelevant to the extent that the difference between external and internal temperatures is a reasonable approximation of heating and cooling load. If the adaptive function is nonlinear, then there must be a different set of thermostat settings for each climate zone, which is not part of current practice.

Comparison in Entropy

Heating and cooling levels alone do not give a complete picture of thermostat behaviors in households. Changes in thermostat settings over the course of a 24-hour period, in particular the consistency of a household's thermostat settings, are also important.

Different households can have the same average temperature settings and wildly different energy use. Figure 2 shows an example with two houses. The broadly dashed line shows the cooling set point of a house with consistent thermostat settings: they are the same for all hours of the day. The solid line shows the thermostat settings for a house with wildly varying settings, but the same average temperature setting as the first house. The finely dashed line shows the external temperature.

The house with the constant thermostat setting will not require any cooling, because the thermostat setting is always higher than the external temperature. The house with the varying settings will require some cooling because the thermostat setting is sometimes above and sometimes below the external temperature. The houses have the same average temperature setting but very different energy use.

Figure 2. Consistency Does Matter

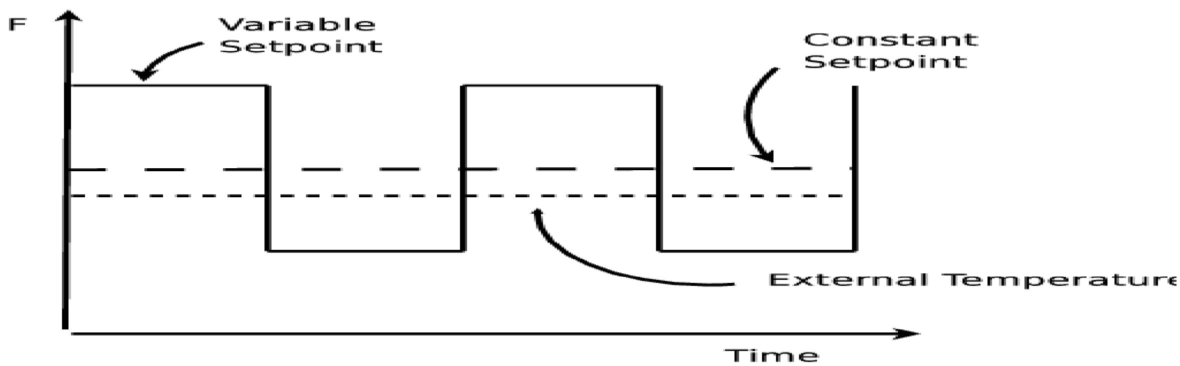

Shannon Entropy is an appropriate statistic for examining the consistency of thermostat settings within households. It may be thought of as a way of describing the consistency of a random variable. A data series with no variation, for example, would have an entropy of zero, while increasingly more variable series would have larger entropy values.

Shannon Entropy (Khinchin 1957) is the sum of the product of the probability of an event and the log of that probability for all discrete events[3]. Algebraically this is:

$$H(x) = -\sum_{\forall x \in p(x) > 0} p(x) \ln\left(p(x)\right)$$

In the context of set points, a household that uses the most basic level of technical control – keeping a fixed heating set point of 68F, for example – would have a Shannon Entropy of $\frac{1}{\log(1)} = 0$. Households that keep two set points – perhaps 68F for a third of the day and 60F otherwise – would have a Shannon Entropy of 0.64, i.e., $\frac{1}{3}\log\left(\frac{1}{3}\right) + \frac{2}{3}\log\left(\frac{2}{3}\right)$. The Shannon Entropy would be the same if the two set points were 68F and 60F or 70F and 55F, because there is no difference in the consistency of those thermostat settings. Because of the behavior of the statistic, Shannon Entropy should be thought of as a measure of consistency rather than variance.

Figure 3 shows a household from our dataset with extremely consistent cooling set point behavior. With few exceptions, this household consistently had a cooling set point of 78F, 80F, or 88F. This household has a Shannon Entropy of 1.09.

[3] Shannon Entropy is also defined for continuous distributions.

Figure 3. Consistent Set Point Behavior (Entropy=1.09)

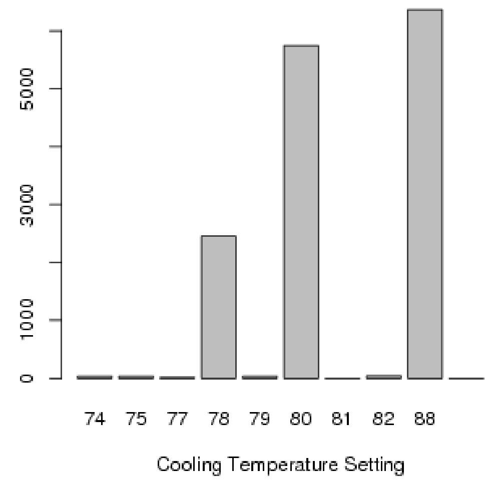

Figure 4 shows a household with much different set point behavior. Rather than establishing just a few set points, this household seems to fine-tune its thermostat, varying the setting between 70 and 90 degrees on a regular basis. Because there are many more thermostat settings in use, the entropy for this household is much higher, 2.18.

Figure 4. More Variable Set Point Behavior (Entropy=2.18)

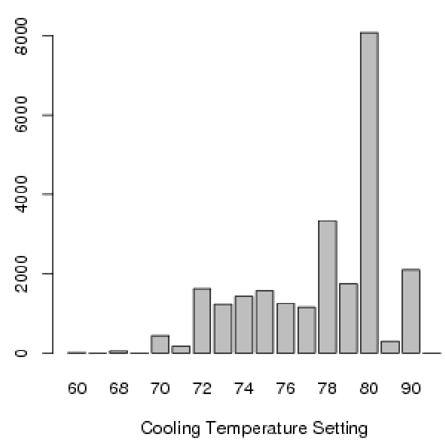

Shannon Entropy proves to be the proper metric here because it captures the underlying differences in these two examples of set point behavior. Calculation of the variances of the temperature settings would miss these differences. For example, the consistent behavior shown in Figure 3 has a standard deviation of 4.35, while the much less consistent set point behavior shown in Figure 4 is only a little more variable, with a standard deviation of 4.59.

Figures 5 and 6 show the distributions of Shannon Entropy across households in our dataset for both the cooling and heating set points[4]. In both figures there are no modes near zero entropy, what we would see if households set a single set point for the day. Few, if any, households are using the single set point method that is assumed by models based on average observed set points. Clearly, if we wish to characterize commonly observed thermostat settings in a specific hour, we cannot use the average observed settings in any single hour.

There are also very few households with a Shannon Entropy of about .64, which would be expected if a household had two set points. The most common Shannon Entropy, the mode, is just a little larger than 1.0 in both Figures 5 and 6.

These figures reveal that many households in our dataset approximate a three-set-point method for both heating and cooling: one for when someone is at home, one for nighttime, and one for when the house is vacant. This pattern is similar to the example shown in Figure 3.

Title 24 assumptions, shown as a dashed vertical line, describe heating and cooling set points with entropies of 0.1512 and 1.1709, respectively. As shown in Figures 5 and 6, in our dataset about half the probability mass in both the heating and cooling entropy distributions is to the left of 1.17. This means that, on average, households are managing the thermostat for cooling in much the same way as modeled for Title 24 compliance, but they are much more actively managing set points for heating.

Figure 5. Distribution of Cooling Entropy

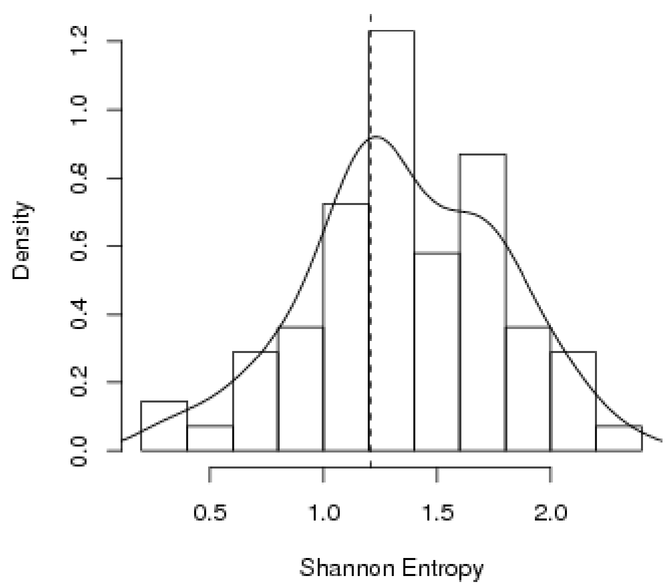

[4] The vertical is scaled so that the empirical density function integrates to one.

Figure 6. Distribution of Heating Entropy

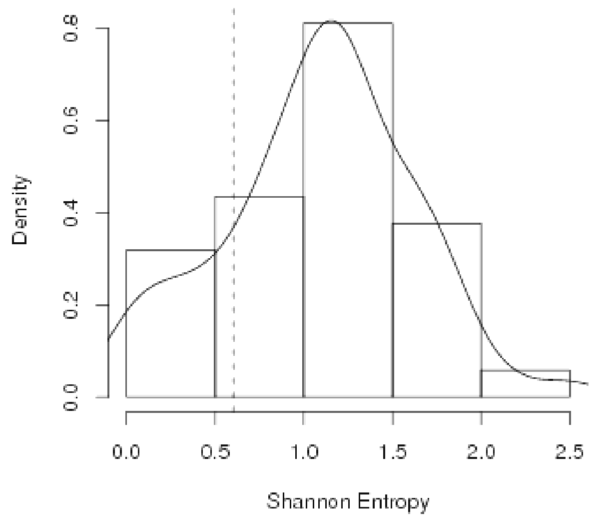

For both heating and cooling set points it is possible to test these hypotheses statistically by calculating the conditional Shannon Entropy of the set point, which is the average of the household-level Shannon Entropy weighted by the number of observations in each household. The variance of this estimate is bounded above by $\frac{\log(N^2)}{N}$ (Antos & Kontoyiannis 2001). There is a statistical difference from the Title 24 assumptions in both cases. Once again this should not come as a surprise given the large number of observations and the artificial pattern of thermostat settings.

There is also an important story to be understood about seasonal changes in thermostat behavior. You don't notice your heater during the summer and you don't notice your air conditioner during the winter. You are unlikely to make changes in the thermostat settings when you don't use them

In addition, setting a thermostat to actively manage indoor temperature takes effort, and effort usually requires some kind of payoff if it is to continue. The payoffs for actively managing cooling and heating set points are not evenly distributed throughout the year. Using additional heating set points pays off during the heating season, when the heater may turn on or off as set points are added or changed. Actively managing the cooling set point pays off during the cooling season for similar reasons.

Because of the personal effort required to actively manage temperature settings and the seasonal variation in the efficacy of that effort, it is reasonable to expect seasonal variations in the entropy of thermostat set points. Table 3 bears out that assumption, showing the monthly conditional entropy of both cooling and heating set points among households in our dataset[5].

[5] This table eliminates households with fewer than 100 observations in a month.

Table 3. Conditional Set Point Entropy by Month

	Cooling	Heating
Feb	0.49	1.16
Mar	0.02	0.29
Apr	0.49	1.13
May	1.01	0.81
Jun	0.91	0.70
Jul	1.17	0.70
Aug	0.93	0.68
Sep	0.90	0.61
Oct	0.51	0.61
Nov	0.60	1.32

Table 3 shows a clear pattern. Cooling set point entropy is low during the winter and begins to rise in May. During the summer the average entropy indicates temperature behavior similar to the three-set-point method. Once winter begins, in about October, cooling set point entropy begins to fall to the level associated with the two-set-point method, i.e., 0.64.

Heating set point entropy shows a similar pattern. During the summer, heating set point entropy falls to a level consistent with the two-set-point method. During the winter, the heating set point entropy is consistent with the three-set-point method.

What Title 24 Modelers Need to Address

Title 24 modelers have been using a single set of thermostat settings to produce estimates of either expected energy use or expected changes in energy use. This means they are approximating the expected value of a function of a random variable, with that function evaluated at the mean of the random variable. On the face, this seems perfectly reasonable. But you can see the fallacy by applying the same reasoning to craps.

Note that the average of two dice is seven. You should then be able to walk up to a craps table, bet $10 seven, and win, “on average,” $40 per roll. Obviously this is not right. Estimates that pile averages on averages are usually only accurate when the function of the random variable is very simple – near linear.

If Title 24 modelers wish to continue the current practice of using a single set of thermostat set points, they must firmly establish that the energy use calculated with those set points is very close to the average estimated energy use using a wide variety of observed set points. In other words, they will have to simulate the same house with many many thermostat settings.

If this cannot be accomplished, the only effective way of estimating expected energy use with a simulation requires each modeled house to be simulated with a large variety of thermostat settings. The averages of the calculated energy use are proper estimates of expected energy use.

Summary

This paper used data from the Energyn project "Field Deployment of the iPower Advanced Energy Efficiency and Information System" to show that residential thermostat set point behavior is much more complicated than is generally supposed. The Energyn project collected information on household thermostat settings every fifteen minutes. Average thermostat settings across households in the dataset were constant within one degree Fahrenheit throughout the day within a given month. This does not indicate that each household kept a constant thermostat setting, however.

Shannon Entropy was introduced as a measure of the consistency of thermostat settings. Entropy is more appropriate than variance as a measure of consistency because it depends on the number of different settings rather than differences among the settings. Entropy figures demonstrate that few, if any, of the households in the dataset used the single, constant set point method, and that frequent temperature adjustments are the rule.

Clearly, Title 24 assumptions are both statistically and, with respect to heating set points, practically different from our findings. The onus is on Title 24 modelers to demonstrate that the thermostat assumptions are right "on average" or "for all practical purposes" with respect to expected energy estimates.

This should bring home the basic observation that our estimates of the benefits of equipment changes are always larger than our uncertainty about the behavior of the individuals running the equipment. Our certainty about the benefits of an air conditioner installation are at least as uncertain as the benefits of teaching people to raise their cooling set point. Conservation portfolios need to reflect this relative uncertainty in benefits.

References

Antos, A and I. Kontoyiannis. 2001. "Convergence Properties of Functional Estimates for Discrete Distributions." *Random Structures and Algorithms*. 19:163-193.

California Energy Commission. 2001. *CALRES2 User's Manual version 1.4*. Sacramento, Calif.: California Energy Commission.

Enerygyn. 2004. *Innovative Peak Load Reduction Program: Field Deployment of the iPower Advanced Energy Effciency and Information System*. Technical Report to the California Energy Commission, 16 July 2004.

Harrell, F. E. Jr. 2005. *Hmisc: Harrell Miscellaneous*. R package version 3.0-7.

James, D. A. and S. DebRoy. 2004. *RMySQL: R interface to the MySQL database*. R package version 0.5-5.

Judko, R. and J. Neymark. 1995. *Home Energy Rating System Building Energy Simulation Test*. Technical Report NREL/TP-472-7332a. Golden, Colorado: National Renewable Energy Laboratory.

Khinchin, A. I. 1957. *Mathematical Foundations of Information Theory*. New York: Dover Publications, Inc.