You're Getting Warmer: Impacts of New Approaches to Residential Demand Reduction

Ken Agnew and Miriam Goldberg, KEMA-XENERGY
Rob Rubin, San Diego Gas and Electric

ABSTRACT

San Diego Gas and Electric's residential Smart Thermostat (ST) Program is a pilot mandated by the California Public Utilities Commission that tests a new combination of paging, Internet, and thermostat technology to raise air conditioner set points remotely. An important feature of the program is customer control of the thermostat both manually and through the Internet, even during re-set events.

The impact evaluation of this program finds that the program savings are lower than targeted by the program plans. Savings reported for the participants in the Statewide Pricing Pilot's Critical Peak Pricing (CPP) program for the same re-set periods were higher than those for the participants in the ST program. Those in the CPP program had incentives to reduce usage during all peak periods, not just declared critical peaks, and also benefited from both air-conditioning and non-air-conditioning usage reductions.

A key factor in the lower-than-targeted savings for the ST program is the timing of the re-set events to coincide with statewide emergencies. As a result, the weather is often not hot in San Diego on these days, and many air conditioners are not being used. On the other hand, on warmer days, over-ride rates rise substantially. A stronger penalty structure could reduce the over-ride rate, but would also make program recruitment more challenging; hence, the net effect on program savings and costs is difficult to predict.

While the design constraints appear to limit the savings potential from this program, the impact analysis methods described can be applied to other re-set programs. These methods take advantage of the program's re-set structure and advanced communications technology. The analysis utilizes end-use metering data for flip-flopping re-set and comparison groups, together with program operational data on signal non-receipt and over-ride rates. The methodology provides estimates that eliminate some key sources of bias while providing good statistical accuracy. The modeling also provides projected future savings as a function of weather and re-set magnitude, taking into account both increased potential savings and increased over-ride rates at higher temperatures.

Introduction

In the aftermath of the energy crisis in California, the California Public Utilities Commission mandated that San Diego Gas and Electric (SDG&E) implement a residential demand response pilot program. The result, the Smart Thermostat (ST) program, was designed to test a customer-oriented approach to demand response using a mix of pager, Internet, and thermostat technologies. The two-way paging system re-sets the thermostat rather than directly cycling the air-conditioning (AC) compressors. Complete participant control of the thermostat, both directly and remotely through the Internet, is a key feature of the program. The thermostat responds to the sender to indicate that the re-set has occurred. A second page from the thermostat will be returned if the re-set temperature is over-ridden by the participant. The Internet also provides easy access to event summaries.

This impact evaluation addressed program design features and took advantage of the program's advanced communications technology. Event data were used to identify potential contributors, focusing the impact estimate and providing greater accuracy. The impact analysis used end-use interval metering for a sample of program participants. During each re-set event, half of the metered sample was left un-controlled, as a comparison group. The savings analysis combined customer-specific load models with comparison group adjustments.

In addition to estimating program-related impacts, we projected potential savings across a wide range of temperature and setback scenarios. To support these projections, it was necessary to estimate the percent of over-riders as a function of program variables. The results are informative regarding the effects of participant control of the thermostat under re-set conditions.

Finally, we consider the results of the impact evaluation for a separate group of ST participants that joined the Statewide Pricing Pilot's (SPP) Critical Peak Pricing (CPP) program. Though it is impossible to do a rigorous comparison based on results from two different analyses, the comparative outcomes are nevertheless thought provoking.

The full report on this study (KEMA-XENERGY 2004) provides further details on the analysis.

Background

SDG&E implemented the ST program beginning in the spring of 2002. The program was designed to include approximately 5,000 residential customers representing an estimated 4 MW in peak demand reduction before 2002 year end. In fact, the program had still not reached this goal by November 2003. The program targeted premises with above-average consumption and AC. Participants receive a $100 incentive for a program year, with a reduction of $2 for each re-set over-ride.

The program is deployed when the California Independent System Operator issues a Stage 2 Emergency Notice. There were no such events in the summer of 2003. However, when the SPP was implemented, the program was invoked on a test basis. ST customers in the evaluation metering sample had thermostats re-set during each SPP critical peak event. This made it possible to evaluate the potential impact of the program, even though the whole program was not implemented.

Methods

Previous Approaches

A variety of methods have been used to evaluate load management programs. Most of these involve collecting interval metering data, and comparing the observed load during a control event with the predicted load based on load data from un-controlled periods. A California Energy Commission study (KEMA-XENERGY 2003) compares alternative prediction methods for demand response programs. This study found that weather-sensitive loads require methods that use information from the controlled day itself via weather models and/or load observations in the hours just prior to curtailment.

For AC direct control programs, methods that are specifically focused on that end use can be used. For AC cycling programs, evaluation methods often involve analysis of the AC duty cycle, which is directly controlled by the program (e.g., Goldberg et al. 2001). Recently, Violette and Ozog (2003) used a leveraged analysis, calibrating a large sample of (inexpensive) run-time metering with a small subsample of interval kW metering. They applied this method to both an AC cycling program and a thermostat re-set program operated similarly to San Diego's. Wright and Martinez (2003) also used a leveraged analysis for a small commercial re-set program in southern California. They concluded that run-time data together with tonnage can be used to determine impacts in the future. None of these studies utilized a comparison group.

Overview of Methods for this Study

For the ST pilot program, the design specified that kW metering would be conducted for a sample of 100 participating homes to provide high-quality information to inform the pilot assessment. Having this large, high-quality metering sample allowed use of a comparison group during each re-set event.

The metering sample was split randomly into two groups, only one of which was re-set for any SPP event. The comparison group provides protection against any systematic bias in the projected usage absent a control. Using participants as the comparison group protects against self-selection bias. The two groups were found to be very similar in terms of the distribution of AC capacity (tons). The "difference of differences" method described below combined the split metering sample with weather-based load estimates.

The evaluation also took advantage of program operational data on signal non-receipt and over-ride rates to improve the accuracy of the estimates. Because the re-set technology involves two-way communication, these data are available for all participants. For traditional AC cycling programs, by contrast, population data on signal receipt are not typically available, and over-ride is not typically an option.

Extensions of the methods and models used to analyze the actual re-set days provide projections of potential savings based on temperature and setback amount. These estimates combine the weather-based load models with an estimate of re-set over-ride rate as a function of average temperature.

Specific Methods Used

Difference of differences. The difference of differences approach used in this analysis combines the benefits of a comparison group from the split metering sample with the benefits of weather correction from a weather-based regression model.

The difference of differences equation is

$$S_{dh} = \left(\bar{\hat{L}}_{Rdh} - \bar{L}_{Rdh}\right) - \left(\bar{\hat{L}}_{Cdh} - \bar{L}_{Cdh}\right),$$

where $\bar{\hat{L}}_{Rdh}, \bar{\hat{L}}_{Cdh}$ are the mean model estimates of load for the re-set and comparison groups for hour *h* of day *d*, respectively, $\bar{L}_{Rdh}, \bar{L}_{Cdh}$ are the mean observed load for re-set and comparison groups for hour *h* of day *d*, respectively, and S_{dh} is the savings estimate for hour *h* of day *d*.

The difference $\left(\hat{\bar{L}}_{Rdh} - \bar{L}_{Rdh}\right)$ provides a raw savings estimate, the difference between (a) what the model estimates the average load would have been in the absence of a re-set $\left(\hat{\bar{L}}_{Rdh}\right)$ and (b) the average observed load $\left(\bar{L}_{Rdh}\right)$. The second difference, $\left(\hat{\bar{L}}_{Cdh} - \bar{L}_{Cdh}\right)$, is the corresponding comparison group difference between modeled and observed load. This difference adjusts for how far off the model estimates were on average in the comparison group.

To use this approach, we must believe that the comparison group average modeling error (estimate minus observed) for a given day and hour is a good estimate of the corresponding re-set group average modeling error (estimate minus theoretical hypothetical usage absent the re-set). We did, in fact, compare the modeling error for the comparison and re-set groups on non re-set days. These errors indicate that, while the groups themselves are not identical, the modeling error of one group is a good predictor of the other.

Weather-based estimates. For the estimated-load portion of the difference of differences equation, we used a weather-normalization model.

The model is

$$L_{jdh} = \alpha_{jh} + \beta_{Hjh} H_d\left(\tau_{Hj}\right) + \beta_{Cjh}\, C_d\left(\tau_{Cj}\right) + \varepsilon_{jdh}$$

where

L_{jdh}	=	AC consumption (kWh) at hour *h* of day *d* for premise *j*;
$H_d(\tau_{Hj})$	=	heating degree-days at the heating base temperature τ_{Hj} for premise *j*, on day *d*, based on daily average temperature;
$C_d(\tau_{Cj})$	=	cooling degree-days at the cooling base temperature τ_{Cj} for premise *j*, on day *d*, based on daily average temperature;
ε_{jdh}	=	regression residual;
α_{jh}, β_{Hjh}, β_{Cjh}	=	coefficients determined by the regression; and
τ_{Hj} τ_{cj}	=	base temperatures determined by choice of the optimal regression.

This model is similar to the PRISM heating-cooling model (Fels 1986) except that we allow different coefficients for base-level consumption α and heating and cooling slopes β by hour *h* of the day. The authors have used this type of model in a variety of other studies (e.g., KEMA-XENERGY 2003b).

Using regression coefficients from the fitted equation, as indicated by the overscript '^', and cooling and heating degree-days $H_d(\tau_{Hj})$ and $C_d(\tau_{Cj})$ for day *d* of the re-set event, the estimated load (without re-set) L_{jdh} , was calculated for each premise, day, and hour using

$$\hat{L}_{jdh} = \hat{\alpha}_{jh} + \hat{\beta}_{Hjh}\, H_d\left(\hat{\tau}_{Hj}\right) + \hat{\beta}_{Cjh}\, C_d\left(\hat{\tau}_{Cj}\right)$$

Fraction non-contributors. In a program like the ST program, there will always be some premises that provide little or no savings. Clearly, for a program designed to lower demand by lowering AC usage, a premise with no AC usage across the whole summer will not provide savings. For the ST program there are two additional ways participants do not contribute to savings. First, for a variety of reasons, pager signals do not always successfully reach the thermostat. Second, because participants retain control over their thermostats, they can choose to over-ride the re-set. Depending on when the over-ride takes place, there may be little or no savings. In combination, these three kinds of premises can represent a substantial percent of participants.

For non-contributors, the savings is effectively zero. Thus, the total savings S_{Tdh} per participating unit for hour h of day d is

$$S_{Tdh} = p_{\mathrm{pc}} S_{dh} ,$$

where S_{dh} is the average impact per potential contributor for a given re-set day and p_{pc} is the percent of units that are potential contributors. Separately estimating the savings per potential contributor S_{dh} and the fraction of units p_{pc} that contribute gives almost the same estimate as directly including all units in the load data analysis. However, by eliminating the effective zeroes from the load analysis, and utilizing additional data to estimate the percent non-contributors, we're able to improve the accuracy of the final estimate.

In practice, the percent of potential contributors is derived as the remainder after calculating the combined percent of non-contributors. The percent of non-contributors for each re-set day is calculated using the following equation (subscript d removed):

$$p_{NC} = p_F + (1 - p_F)(p_{OR} + p_z) = 1 - p_{PC} ,$$

where

p_{NC} = fraction of units that were non-contributors on re-set day;
p_F = fraction of units not responding to signal on re-set day;
p_{OR} = fraction of units that over-rode, of those receiving signal on re-set day; and
p_z = fraction of units with zero weekday AC usage all summer on re-set day.

That is, all units with signal failure (p_F) are non-contributors. Of the remaining units $(1 - p_F)$, those that cannot contribute to savings are those that over-ride (p_{OR}) and those that were never used (p_z). These proportions are additive because they are essentially mutually exclusive. Whether a unit has zero use is assumed to be independent of whether or not the signal was received.

Under full-scale program operation, the program-wide non-response and over-ride fractions are known with certainty from the event data. Thus, they contribute no variance to the estimate of p_{pc}. The percent of AC non-users is estimated from the whole metering sample, and does contribute variance.

For the test conditions of 2003, which did not involve full-scale operation, an alternate procedure was used, but still provided accuracy gains compared to relying on the metering sample only.

Projected savings — savings with full compliance. The weather based load model used in the difference of differences mode allowed us to make projections of impacts for future dates. With that load model, an increase in setpoint by an amount of δ is represented by an increase in the cooling reference temperature τ_C to $\tau_C + \delta$. Thus, the projected savings at ambient temperature τ_C and re-set δ is calculated as $\hat{\beta}_{Cjh}\left[C_d\left(\hat{\tau}_{Cj}\right)-C_d\left(\hat{\tau}_{Cj}+\delta\right)\right]$ using the same model parameter estimates as for the difference of differences calculation. This equation was applied to all non-zero AC users and provides projected average savings for potential contributors. With this approach, these unadjusted savings were projected for the full range of temperatures likely in the San Diego area and a range of setback amounts.

Projected savings — non-contributors. To make these unadjusted projected average savings estimates useful from a program implementation perspective, we needed to adjust them to reflect the whole population just as we adjusted the impact estimates. Thus, we needed to predict the fraction of contributors. This, in turn, necessitates estimating the three constituent fractions.

Non-response was always a relatively small fraction of units and appeared to vary randomly. Non-response could, thus, be estimated by the mean non-response rate over the 12 re-set dates.

Zero AC use is determined for the season as a whole from the metering data. We used the 2003 observed fraction with zero use as an estimate of the fraction for future events.

The over-ride rate, on the other hand, was variable and appeared to be driven by temperature.

Using the over-ride rates and the daily average temperatures for the 12 re-set days, we estimated the relationship between temperature and percent over-ride. A log-odds specification kept the estimated fractions between zero and one. Other variables—duration of the re-set event, degrees setback, and time of the event—were tested but were found not to be significant.

These three fractions were combined using the equation for percent non-contributors. The resulting temperature-based potential contributor percents made it possible to adjust the projected savings to reflect program usage patterns across the full range of possible temperatures.

Findings

There were 12 days in the summer of 2003 when the SPP's CPP program was invoked and the ST program was re-set. The first CPP day was in July and the last in late October. Events were between 2 PM and 7 PM and ranged from 2 to 5 hours long. Average temperatures on the re-set days ranged from 68°F to 82°F and setbacks ranged between 3 and 5 degrees.

Percent Non-Contributors

The estimated fraction of potential contributors for each CPP day is shown in Table 1. The component fractions of the overarching fraction of non-contributors are also shown.

Table 1. Fraction Non-Contributing Summers 2002 and 2003

Re-set Date	Non-response: No Response Count	Non-response: ST/SPP Participants	Non-response: Fraction (P_F)	AC Non-use Fraction (P_z)	Re-set Over-ride: Over-rider Count	Re-set Over-ride: ST Sample Participants	Re-set Over-ride: Fraction (P_{or})	Fraction Not Contributing $P_{NC} = P_F + (1-P_F)(P_z+P_{or})$	Fraction Potential Contributor $P_{PC}=1-P_{NC}$
07/17/03	8	156	5%	18%	4	31	13%	35%	65%
07/28/03	11	157	7%	18%	7	41	17%	40%	60%
08/08/03	12	156	8%	18%	9	40	23%	45%	55%
08/15/03	14	165	8%	18%	14	30	47%	68%	32%
08/27/03	6	164	4%	18%	9	30	30%	50%	50%
09/03/03	19	164	12%	18%	7	29	24%	49%	51%
09/12/03	11	166	7%	18%	8	41	20%	42%	58%
09/22/03	6	152	4%	18%	6	30	20%	41%	59%
09/29/03	6	165	4%	18%	2	39	5%	26%	74%
10/09/03	15	162	9%	18%	2	30	7%	32%	68%
10/14/03	14	163	9%	18%	5	38	13%	37%	63%
10/20/03	3	151	2%	18%	4	30	13%	33%	67%

2003 Impacts

Figures 1 to 3 illustrate the difference of differences approach for the August 8 re-set event. Figure 1 illustrates the raw estimate of savings derived from the estimated and actual AC load of the re-set group. The plot of the re-set group's observed AC load diverges dramatically from the estimated (no re-set) AC load at the hour 15 start time. As the re-set period progresses, the difference between the observed and estimated load decreases as units come on to maintain even the higher re-set temperature. After the re-set period's end at hour 17, the re-set group's load jumps above the expected load as those AC units come on full time to compensate for lost cooling. The difference between observed and estimated load for the re-set group is the raw estimate of the impact for this re-set period.

Figure 1. Re-Set Group Mean Observed and Estimated AC Loads on August 8, 2003

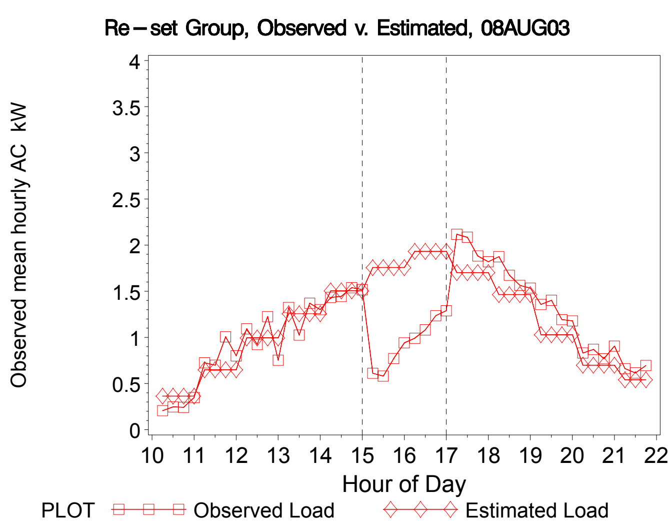

Figure 2 shows the observed and estimated AC load for the comparison group. The plot indicates that, on this particular day, the estimated load underestimated the observed load. This may have been caused by high humidity or some other systematic effect. The difference between observed and estimated load for the comparison group provides the adjustment for the raw savings indicated in Figure 1.

Figure 2. Comparison Group Mean Observed and Estimated AC Loads, August 8, 2003

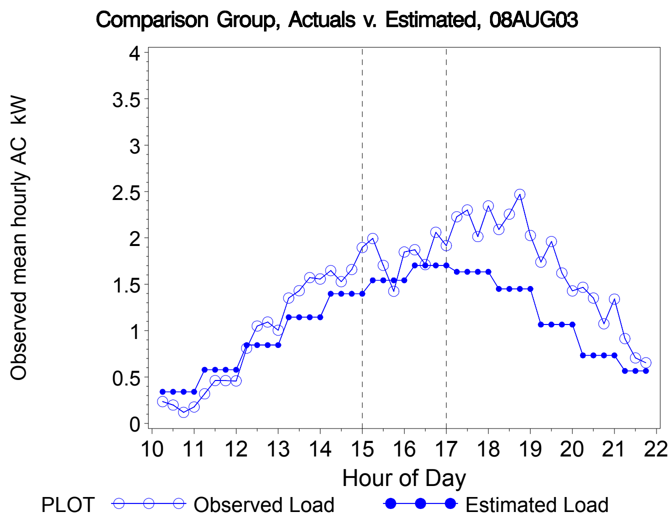

The plots above reflect the mean usage of potential contributors; that is, premises that had non-zero consumption for at least some part of the summer, received a re-set signal, and did not over-ride. It is still necessary to adjust the combined result so that it reflects a per-unit impact for all units in the program. This adjustment involves multiplying the savings per potential contributor by the fraction these potential contributors represent of the whole. The percent of potential contributors for each day was given in Table 1. The final impact estimates for the August 8 re-set event are displayed in Figure 3.

Figure 3. Mean Impacts per Participating Unit on August 8, 2003, vs. Time

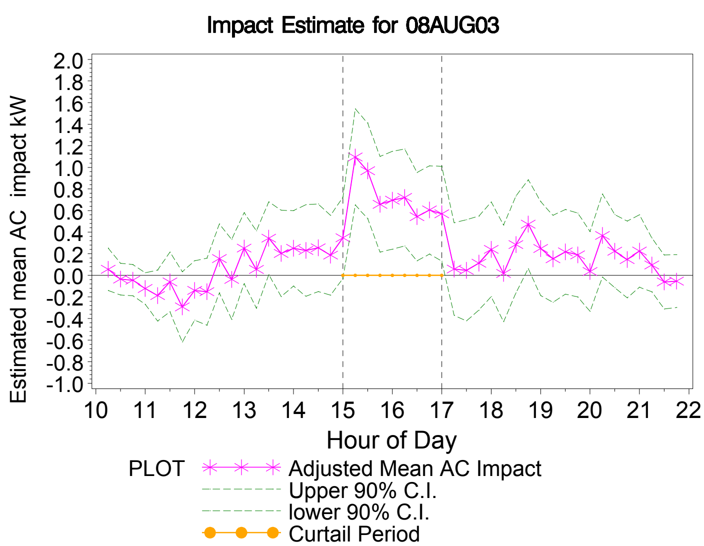

Table 2 presents the 2003 AC load impact results in tabular form. The shading indicates results that were not statistically significant from zero. Impacts ranged from 0.06 to 0.73 kW and 7 of 12 were statistically significant at the 90 percent confidence level.

Table 2. Mean AC Impacts with Confidence Intervals Per Participating Unit and 5000 Units

		Mean kW Per Unit				kW for 5000 Units		
Date	Sample Group	Impact	Standard Error	90% Confidence Lower Bound	90% Confidence Upper Bound	Impact	90% Confidence Lower Bound	90% Confidence Upper Bound
7/17/03	A	0.22	0.22	-0.14	0.59	1,124	-690	2,939
7/28/03	B	0.17	0.14	-0.06	0.41	869	-318	2,057
8/8/03	B	0.73	0.27	0.27	1.19	3,662	1,362	5,962
8/15/03	A	0.30	0.27	-0.15	0.75	1,482	-761	3,726
8/27/03	A	0.68	0.21	0.33	1.04	3,414	1,632	5,195
9/3/03	A	0.26	0.14	0.02	0.50	1,302	109	2,495
9/12/03	B	0.32	0.16	0.04	0.59	1,577	202	2,951
9/22/03	A	0.30	0.16	0.03	0.56	1,475	154	2,797
9/29/03	B	0.30	0.12	0.10	0.50	1,491	497	2,485
10/9/03	A	0.06	0.11	-0.12	0.24	308	-597	1,213
10/14/03	B	0.10	0.09	-0.05	0.24	493	-227	1,213
10/20/03	A	0.53	0.25	0.11	0.94	2,637	563	4,712

In part, the generally low savings result from relatively mild temperatures on most of the re-set days. These days are typical of the days on which the program is likely to be operated. San Diego has a mild climate, and the program is invoked only for statewide emergency conditions, which don't necessarily coincide with hot weather in San Diego. Moreover, as shown below, at hotter temperatures the increased over-ride rate appears to dominate the increased potential savings.

Projected Savings

Estimated over-ride percent. To produce adjusted projected savings, we first had to estimate the over-ride rate. Figure 4 shows the results of estimating over-ride percent as a function of outdoor temperature. The observed data are also included. Two lines are actually shown in the figure. One is for a regression using data from all 12 days. The second excludes the day with the highest temperature and over-ride rate. The two curves are nearly indistinguishable. Moreover, both are very close to the observed value of 47 percent on this most extreme day. Thus, while the over-ride rate of 47 percent may seem anomalous, it is nearly exactly what the model would predict based on the other days.

This finding gives some confidence that the model is giving reasonable results for temperatures through the low 80s, despite the fact that we have only one observation above 76°F. At higher temperatures, where no observations have been made, the estimates are less certain. Between 75°F and 85°F, each 1°F increase in temperature is associated with an increase of 3.6 percentage points in the over-ride rate.

The importance of the role of over-ride rate is made clear when we consider the projected savings. Because of the nature of the linear weather regression underlying the projected savings, savings per potential contributor increase with temperature to a point where all units are on all the time. Beyond this point, savings do not change with higher temperatures. For the summer of 2003, the projected savings per potential contributor rise until the temperature hits 83°F. If we adjusted this projection by a fixed percent of potential contributors, maximum savings would still be at a constant rate at 83°F and above.

Figure 4. Predicted and Observed Over-Ride Rates vs. Outdoor Temperature

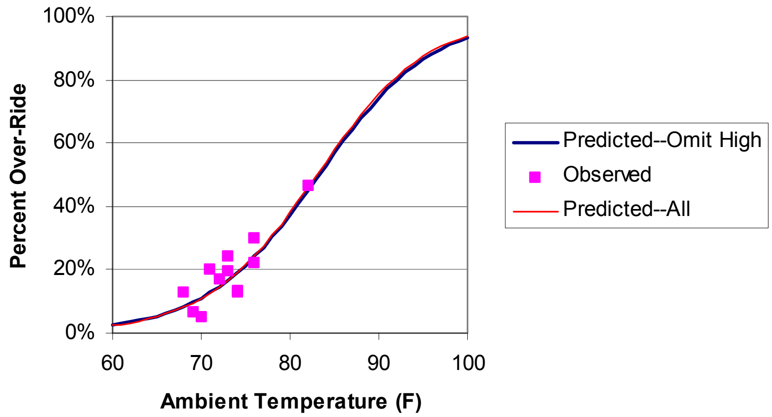

When projected savings per potential contributor are adjusted by temperature-based potential contributor percent, the result is quite different. As temperature increases, the increasing savings are countered by decreasing percentages of potential contributors. Figure 5 provides the adjusted projected impacts, per all participating units, as a function of temperature. When the effect of temperature on over-ride rate is taken into consideration, peak savings for a 3° re-set occur at 76°F rather than at 83°F. Above 76°F, despite still-increasing savings per potential contributor, adjusted savings per participant fall dramatically.

Figure 5. Adjusted Projected Impacts by Temperature, Average per AC Unit (for 3°F Re-Set)

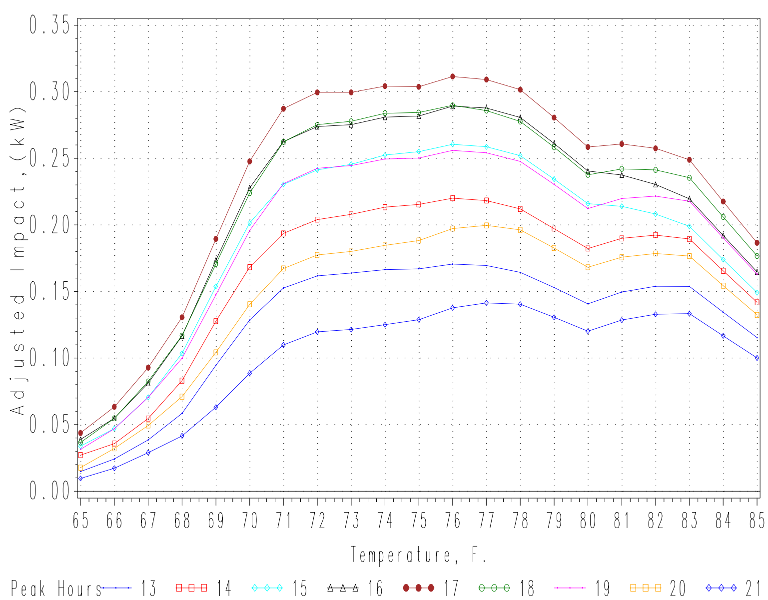

Figure 5 illustrates the challenge of participant control of the thermostat setting. Under more extreme circumstances, perhaps those most likely to motivate a re-set event, the potential savings decrease. At 83°F, savings have already been reduced from the maximum by 20 percent and will continue to drop by approximately 10 percent for each of the next five 1° degree increments. On the other hand, the relatively flat effect of temperature on savings per participant in the range of 71°F to 78°F can make program goals easier. Eight of the twelve re-set events in 2003 had a temperature in this range. Three of the remaining four re-set events were days with temperatures lower than 71°F, a reminder that San Diego weather is not a driver of the SPP critical peak events.

Comparison with SPP

The SPP established its new price-based demand response program in 2003. The CPP-V tariff of the SPP allowed San Diego area ST participants to opt into this new program. The impact analysis for summer of 2003 was recently published (Charles River Associates 2004).

CPP-V rate participants were put on time-of-use rates for non-CPP days (10.8 ¢/kWh off-peak, 27.8 ¢/kWh on-peak). On CPP days, the "super" peak period energy price increased to 76.8 ¢/kWh. The SPP impact evaluation identified per premise savings for the San Diego area premises with respect to a comparison group for two periods: peak hours on non-CPP days and super peak hours on SPP days. The evaluation identified average savings over the summer of 0.74 kW during non-CPP peak hours and savings of 1.2 kW during CPP super peak hours. Thus, the CPP-V rate shows incremental savings of 0.46 kW for the CPP day super peak period over and above the TOU rate savings of 0.74 kW during non-CPP peak periods.

By comparison, the San Diego ST program produced AC savings per unit between 0.06 kW and 0.73 kW, averaging 0.29 kW over all re-set hours of the 12 re-set days. Thus, the re-set appears to yield higher savings per unit for those on the SPP rate than for those in the ST program.

Several factors can account for this difference. One is that the SPP rate provides incentives for load reduction in other end uses, apart from AC, and for AC reduction beyond the re-set amount. Another factor may simply be self-selection. ST participants had an option whether to change to the SPP rate.

Another factor may be the difference in perceived cost of an over-ride. At an average incremental load reduction of 0.46 kW for the average re-set duration of 3.65 hours, the cost to an SPP participant of over-riding the re-set and otherwise ignoring the critical peak price signal would be $1.29. However, if the participant anticipates a load reduction of 1 kW and a re-set period of 5 hours, the cost of not reducing load during the critical peak period would be $3.84. An ST participant over-riding the re-set and using 5 kWh would face a combined cost of roughly $2.75 to $3.00 depending on the rate tier. SPP over-ride rates were 30 to 50 percent lower than the ST over-ride rates for almost all the re-set periods.

Conclusions

Program Impacts

In 2003, the ST program produced savings below its design goals. This is a result of a number of factors identified in the 2003 impact analysis.

- Eighteen percent of premises never used their AC during summer weekdays, therefore provided no savings.
- The over-ride rate increased dramatically with temperature, reaching nearly 50 percent on the hottest re-set day.
- The incentive structure provides a net benefit to ST participants even if they over-ride every re-set.
- The choice of re-set days is not driven by San Diego weather, so that re-set days are not always on days with substantial potential savings.
- A small fraction of premises, 3 to 6 percent, did not receive the re-set signal.

As a result of all of these factors, potential contributors (AC users that received the re-set signal and did not over-ride) only represented between 32 and 74 percent of the population across the 12 re-set days. If all premises were potential contributors with the observed savings, the program would have surpassed its goal on two of the re-set days.

These findings do not imply that the concept of load management via thermostat re-set is flawed. San Diego's mild climate is not necessarily well suited to providing savings from AC peak-day reductions other than with strong participation and compliance incentives.

Methods

While the design constraints appear to limit the savings potential from this program, the impact analysis methods described can be applied to other re-set programs. These methods take advantage of the program's re-set structure and advanced communications technology.

Load models fit separately for each AC unit provided unit-specific estimates of what usage would have been in the absence of a re-set. These models also provided the basis for projecting savings under future conditions.

The "difference of differences" comparison group adjustment corrects for any systematic tendency for the weather model to over- or under-predict for a particular day and hour. Comparison of the modeling errors for the two groups indicates that these systematic effects do exist and need to be corrected for. The alternating re-set/comparison group assignment eliminates any bias because one or the other group is more or less inclined to reduce usage on emergency days. Use of program participants as the comparison group also eliminates self-selection bias.

The use of operational data on non-contributor rates improved the accuracy of the estimates. The model of over-ride rate as a function of weather conditions allows savings projections from the load models to be moderated by likely changes in behavioral response at higher temperatures.

While these features provide several advantages, further advantages may be available with other metering approaches. A leveraged approach using run-time data and a small interval kW subsample might allow similar methods to be used, with less costly metering.

References

Charles River Associates. 2004. *Statewide Pricing Pilot Summer 2003 Impact Analysis*. Oakland, Calif.: Charles River Associates.

Fels, Margaret F. (ed). 1986. "Special Issue Devoted to Measuring Energy Savings: The Scorekeeping Approach." *Energy and Buildings* 9 (1 and 2).

Goldberg, M. L., M. O'Kelly, V. Goepfrich, C. Neme, S. Ringhof, and F. Lynk. 2001. "Knowing Your Limits: Application of Censoring Distributions to Impact Estimation for a Residential Air Conditioner Direct Load Control Program." In *Proceedings of the 2001 International Energy Program Evaluation Conference* (poster). Salt Lake City, Utah.

KEMA-XENERGY. 2003a. *Protocol Development for Demand Response Calculation: Findings and Recommendations.* Prepared for California Energy Commission.

KEMA-XENERGY. 2003b. *2002 Smart Thermostat Program Impact Evaluation.* Prepared for San Diego Gas and Electric.

———. 2004. *Final 2003 Smart Thermostat Program Impact Evaluation*. Prepared for San Diego Gas and Electric.

Violette, D., and M. Ozog. 2003. "Mass-Market Demand Management Offerings: Evaluation Methods Assessment and Results." In *Evaluation: Meeting Diverse Needs—2003 International Energy Program Evaluation Conference*, 131–142. Seattle, Wash.

Wright, R., and M. Martinez. 2003. "Results from California's Small Commercial Demand-Responsiveness Pilot Program." In *Evaluation: Meeting Diverse Needs—2003 International Energy Program Evaluation Conference*, 119–129. Seattle, Wash.