Exploring Energy and Demand Impacts of a Controllable Rates Program

Miriam L. Goldberg, XENERGY Inc.

This paper presents the results of the impact evaluation of a Midwestern utility's controllable rates. Under these rates, customers agree to reduce their demand to a pre-determined level (PDL) upon notification by the utility.

This evaluation assessed both dispatchable and embedded impacts, for 1991. The dispatchable impacts are the relief realized in response to control notification. The embedded impacts are the result of customer responses to the controllable rate that affect load even outside of control periods. This evaluation includes the first comprehensive assessment by a major utility of embedded impacts of load management rates.

The dispatchable impacts are estimated from load research data, by constructing a prediction model for control days from noncontrolled days. This model is also able to isolate anticipatory effects due to actions taken prior to control notification on days control periods are likely.

Embedded impacts are estimated from engineering analysis of onsite verification data. These small-sample site-specific results are expanded to the general population by leveraging a screener survey fielded to all participating customers. Both the dispatchable and embedded analysis methods appear to work well, and to have broader applicability.

Introduction

This paper presents the results of the impact evaluation of a Midwestern utility's controllable rates. Under these rates, customers agree to reduce their demand to a pre-determined level (PDL) upon telephone notification by the utility. Customers who fail to reduce their load below the PDL pay a penalty for each kW above the PDL during the control period. Outside the control periods, customers on the rate pay a reduced rate for all kW above the PDL.

This evaluation assessed both dispatchable and embedded impacts, for 1991. The dispatchable impacts are the relief realized in response to control notification. The embedded impacts are the result of customer responses to the controllable rate that affect load even outside of control periods. This evaluation includes the first comprehensive assessment by a major utility of embedded impacts of load management rates.

Origins of the Project

The utility's controllable rates programs began over ten years ago with the utility's largest customers. Since that time, the program has been gradually extended to customers with smaller and smaller "controllable load," i.e. the amount of load above the PDL. In 1991 and 1992, the program was available to customers with peak controllable loads as small as 75 kW.

This project was motivated by several concerns regarding the controllable rates programs. One was the declining apparent "coincidence factor" as the program was extended to customers with smaller peak loads. That is, the load relief realized on control days appeared to be a smaller fraction of the contracted value of the controllable load. Understanding and quantifying this pattern is important for planning and forecasting as the program is extended to more customers.

Related to the declining coincidence factors was a perceived need to assess the magnitude of "embedded impacts." One possible explanation for the declining coincidence factors is that greater fractions of the controllable load are actually taken as permanent or seasonal reductions. These reductions are embedded impacts of the program, but are not available as dispatchable relief when needed. Thus, for planning purposes, quantifying embedded impacts is important to avoid projecting more dispatchable relief than will actually be available from future participants. In addition, if embedded impacts of an appreciable magnitude exist, the program should be given credit for them.

Another motivating concern was the need for a well founded method of measuring the dispatchable impacts that have been achieved. Accurate measures of dispatchable impacts are important both for systems operations purposes, and to quantify the benefits of the program. Utility staff were generally comfortable with the existing in-house methods from an operational standpoint. However, there was some interest in developing a more rigorous technical foundation, for these methods.

Overview of Methods Used in this Study

Three sources of data were available for the impact analysis:

1. 1991 load research data for the approximately 1500 customers who were on controllable rates during that summer;

2. screener survey data collected in this study, for the approximately 70 percent of customers responding to the survey; and

3. follow-up audit data collected in this study, for the 130 sites that received follow-up verification, by onsite visit (122 sites) or by telephone (8 sites).

The load research data are the basis of the dispatchable impact estimation. The model developed for this estimate also provides an estimate of short-term embedded impacts. The load research modeling is described in the Estimation Methods section.

The embedded analysis is based on engineering and statistical analysis of the screener survey and follow-up data.

The data analysis for this study began in the spring of 1993. During the previous summer (1992) there were no control days because of unusually cool weather in the midwest. The dispatchable analysis was therefore conducted for the summer of 1991. For consistency, the embedded analysis was designed to cover the same year.

The different types of impacts for which estimates were developed in this study are described in the next section followed by a summary of the estimation approach, and finally the impact estimates.

Definitions of Impacts

Baselines for Defining Impacts

For any program, impacts are defined relative to a *baseline* of what would have occurred in the absence of the program or of particular program components. The different types of impacts correspond to different conceptual definitions of the baseline. For this study, three different baselines are of interest for any given day (Figure 1).

> The *likely-day baseline* represents what would be expected on that day if a control period was likely, but none actually occurred.
>
> The *unlikely-day baseline* represents what would be expected on that day if a control period was unlikely (and none occurred).
>
> The *no-program baseline* represents what would have been expected on that day if the program had never existed.

These definitions beg the question of how a “likely” day is defined. That question is addressed below.

Distinct Impacts Relative to Different Baselines

Three distinct impacts are defined, with respect to the three baselines. Each of these impacts is illustrated in the figure.

Dispatchable Impact: On a control day, the difference between the observed level and the likely-day baseline. That is, the difference between the actual (observed) level on a control day and what would have been observed on that day if the control period had not occurred. Demand (kW) impacts are defined for each hour of the day, including non curtailment hours. Energy (kWh) impacts are the total net difference for the day, including non curtailment hours. Dispatchable impacts exist only on control days.

Anticipatory Impacts; On a likely day, the difference between the likely-day baseline and the unlikely-day baseline. That is, the difference between the load expected (or observed) when a control period was likely, and what would have been expected on that day if a control period had not been likely. Demand (kW) impacts are defined for each hour of the day. Energy (kWh) impacts are the total net difference for the day.

Examples of anticipatory impacts include pre-chilling the building in the evening or early morning when a control

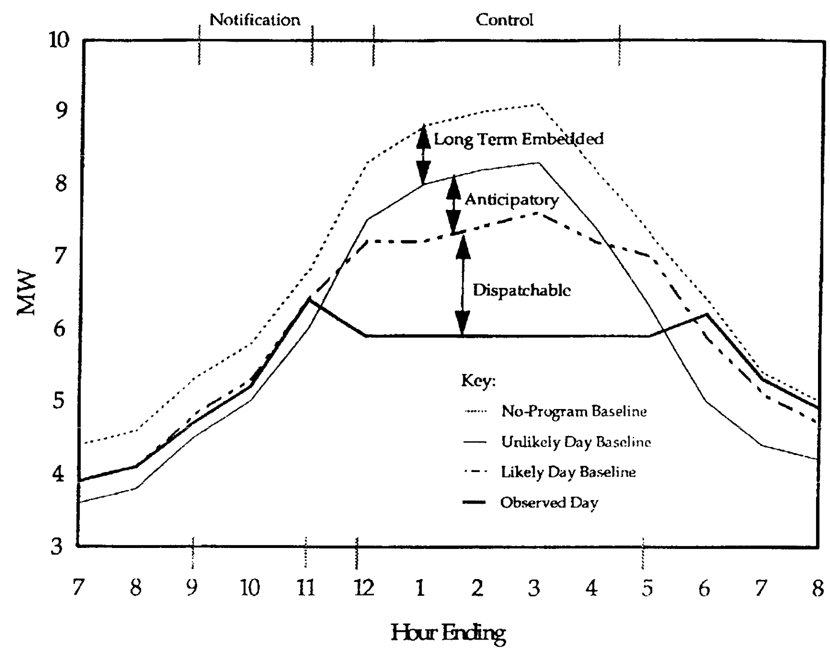

Figure 1. Distinct Impacts in Relation to Different Baselines, Hypothetical Control Day

period is anticipated for the next day; or scheduling a shift or process to begin early in the morning and finish by noon. Anticipatory impacts occur both on control days and on likely but noncontrol days.

Seasonal or Permanent, or Standard Operating Procedure (SOP) Impacts: The difference between the unlikely-day baseline and the no-program baseline. That is, the difference between what would be expected if no control occurred or was likely on a particular day, and what would have been expected that day if the program did not exist. Demand (kW) impacts are defined for each hour of the year, but are of especial interest for peak hours. Energy (kWh) impacts are the total net difference for the year.

Examples of SOP impacts include rescheduling a shift for the entire summer; turning off certain lines or equipment for the summer; or installing more efficient equipment that is used year-round. SOP impacts occur on all days, control, likely, and unlikely.

In this study, the three distinct impacts, dispatchable, anticipatory, and SOP, were estimated separately. Estimates were developed of demand impacts by hour, and were aggregated to seasonal and annual energy impacts.

Coincidence Factors

For purposes of this study, the coincidence factor is defined as the ratio of the estimated demand impact to the maximum controllable load. The maximum controllable load is the difference between the observed summer peak for each customer and the customer's pre-determined level or PDL. A separate coincidence factor is defined and computed for each of the three distinct impacts: dispatchable, anticipatory, and SOP. Since the total program impact is the sum of these three impacts, the overall program coincidence factor is the sum of the three corresponding coincidence factors.

Estimation Methods

Dispatchable Estimation Methods

The investigation of dispatchable impacts focused on the load research data. These impacts are discernible from the recorded loads on control days, in comparison with the noncontrol-day loads. The approach of the analysis was to use load research data from noncontrol days to construct a model to predict what the load shape would have been on the control day if the control period had not been called.

Several different modeling approaches were explored. These methods started with the in-house methods the utility has used in the past, and built on these methods in developing the final models. The success of the different approaches was assessed through cross-validation, measuring the predictive capability of each method in terms of the accuracy with which noncontrol days were predicted.

The model selected on the basis of the cross-validation included an explicit additive term for likely control days.

This term represents anticipatory impacts. Thus, the load research analysis provided estimates of both dispatchable and anticipatory impacts.

Analytic Data Set

Dispatchable and anticipatory impacts were estimated from the 1991 load research data. The utility provided interval kW data for participating customers from June through August of 1991. From this data set, an analytic data set was developed, containing hourly load data (average of the 15-minute loads) for each account.

The hourly load data used in the dispatchable impact model was the average load per account, by rate class. This approach normalized for changing numbers of accounts over the summer, and also for occasional missing values due to data problems.

Weather data, including hourly wetbulb and drybulb temperature at three hour intervals, were provided by the utility for weather stations in each major region. The three-hour data were linearly interpolated to provide hourly values.

Each rate class, defined by rate type, state, and voltage, was modeled using data from the weather station in the same region. Likely days were defined for all customers based on temperatures from the largest region. Since control decisions are based on system load, not local loads, the weather affecting the largest portion of the system load was taken to be the determining factor.

The Dispatchable Impact Model

Several different models were explored for dispatchable impact estimation. All of these involved constructing an estimate of the hypothetical load curve that would have been seen had the control period not been called on a particular day. Thus, each estimation method is based on a prediction model.

Specification of the Load Model. The final model developed for the dispatchable impact analysis represented the load L_{dh} at a given day d and hour h as the sum of the following terms:

- a basic hourly load shape, corresponding to a different intercept term for each hour of the day;
- an immediate cooling load, proportional to the current cooling degree-hours;
- an additional lagged cooling load due to stored heat in the building, proportional to a weighted average of cooling degree-hours for the last 48 hours;
- a latent heat load for removal of humidity, proportional to the wetbulb temperature for the hour;
- an hourly shift in level up or down according to how different the 9 am load was from the average 9 am load on unlikely days;
- an hourly anticipatory effect, if the day was a likely day.

The dispatchable model is given algebraically as

$$L_{dh} = \lambda_h + \beta_0 C^0 dh + \beta_c C_{dh} + \beta_w W_{dh} + \gamma_h (L_{d9} - \bar{\bar{L}}_{U9}) + \alpha_h A_d \quad (1)$$

where

C^0_{dh}	is the cooling degree-hours at hour h on day d
C_{dh}	is the weighted average of cooling degree-hours from hour h-48 to hour h
W_{dh}	is the wetbulb temperature at hour h of day d
$\bar{L}_{U9}$	is the average 9 am load for unlikely days
A_d	is a dummy variable equal to 1 if day d is a likely day, zero otherwise

λ_h, β_0, β_c, β_w, Γ_h, and α_h are coefficients estimated by the regression.

The terms in the dispatchable model are described further below.

Hourly Intercept l_h: The first term, as noted, is an hourly intercept term. For an unlikely day ($A_d = 0$) with no cooling degree-days, current or lagged ($C^0 dh = 0$ and $C_{dh} = 0$), with 9 am load equal to the average unlikely-day 9 am load ($L_{d9} - L_{U9} = 0$), and with constant wetbulb temperature ($W_{dh} = W$), the hourly intercepts Ah would define the load shape. Only the level would be shifted according to the humidity terms $ß_w W$.

Immediate Cooling Load $ß^o C^o_{dh}$: The immediate cooling load due to sensible heat (air temperature) is assumed proportional to cooling degree-hours, base 65°F. The cooling degree-hours at hour h is simply the difference between the outside temperature at that hour and 65°F, or zero if that difference is negative. The use of degree-days or degree-hours base 65°F is a standard, if imperfect, assumption in cooling load calculations.

Lagged Cooling Load $ß_C C_{dh}$: The lagged cooling load term accounts for the fact that heat is stored in a building, so that the cooling load depends on how hot the last few days have been, not just on the current temperature. The lagged cooling degree-hour term C_{dh} is a weighted average of cooling degree-hours for the past two days. The weights decline exponentially, so that the more distant times contribute less. The exponentially declining weights correspond to an approximation to a simple heat transfer model accounting for thermal mass (building heat storage capacity).

The lagged cooling degree-hour term C_{dh} is defined as

$$C_{dh} + \sum_{k=0}^{48} C_{d,h-k} e^{-k/48} \; / \; \sum_{k=0}^{48} e^{-k/48} \qquad (2)$$

(For k > h, hour h-k of day d is equivalent to hour h-k+24 of day d-l.)

Humidity $\beta_w W_{dh}$: The air conditioning load depends on both the sensible heat that must be removed by reducing the air temperature, and the latent heat that must be removed in reducing the moisture content of the air. The sensible cooling load is modeled by the cooling degree-hour terms. To model the humidity load, American Society of Heating, Refrigeration, and Air Conditioning Engineers (ASHRAE) enthalpy charts were reviewed. It was determined that over the range of interest the latent cooling load was roughly a linear function of wetbulb temperature. Rather than introduce a more complicated model for this component, the wetbulb temperature was included directly in the model.

Unlike the sensible heat load modeled by the drybulb temperature, the latent heat load does not include a lag term for previous hours' humidity. Standard cooling load calculations, such as those given in the ASHRAE handbooks, include lag effects for drybulb temperature, but not for humidity. These calculation procedures reflect the fact that heat is stored in buildings to a greater extent than humidity. Heat is retained in the building structure itself. Humidity is contained primarily in the air, which is exchanged several times per day, rather than being stored over several days.

Particular-Day Effect $\gamma_h(L_{d9} - L_{U9})$: Even after weather effects are accounted for, the load may be high or low on a particular day because of the activity level for the account that day. The shift up or down to account for how a particular day is different from a likely-day norm is proportional to the 9 am difference in load between the particular day and the unlikely-day average. However, a different effect of the particular day is allowed for each hour. The shape of the particular-day shift is the same for all days, and is given by the coefficients v_h.

The magnitude of the shift is determined by the 9 am load difference. If this difference is positive, the estimated load at all hours of day d is increased. If the difference is negative, the estimated load at all hours is decreased. If the difference is small in absolute terms, the amount of the hourly load shift is small.

Anticipatory Effect $a_h A_d$: The final term in the equation is the anticipatory effect. On likely days ($A_d = 1$) the estimated load at hour h is increased by the amount a_h. The estimated anticipatory effect at a given hour h is the same for all days d. The estimated model coefficients a_h directly estimate the anticipatory impact.

Comparison with Other Models for Estimating Impacts of Load Management Programs. Despite the increasing importance of load management programs in utilities' DSM portfolios, relatively few impact evaluations of such programs have been conducted utilizing observed load data from controlled customers. Joyce et al. (1993) also analyzed impacts from a set of load management programs using a degree-hour regression model to fit hourly loads on noncontrol days. Their model was fit to individual sites, across multiple years and seasons. It did not isolate an anticipatory effect.

EPRI (1988) has an impact estimation model that uses Fourier functions to extract cyclic patterns from load data. This model also includes temperature effects, is fit across all seasons, and does not identify anticipatory effects. A survey of 25 utilities conducted at the outset of the study described here indicated that the EPRI model has generally not been used even by utilities that own it, because of the time requirements and complexity of implementation.

The survey also indicated that a common approach to determining the load that would have occurred if an interrupt had not been called is to develop an average load shape from hot noncontrol days, then shift this load shape up or down to coincide with observed load prior to the control notification. The effectiveness of this approach as implemented by the utility sponsoring this study was the motivation for the inclusion of the particular-day effect in the load model (l).

Definition of Likely Days

A key to the estimation of the load models is the definition of reference or likely days. In previous work, the utility had classified days as high, medium, or low demand. Reference days were then taken as the high days within the same month, if there were enough, or else as the high and medium days within the month.

For the present work, all weekdays with an average outside temperature greater than 70°F were classified as "likely" control days. This definition resulted in all control days and nine noncontrol days being classified as likely days. The nine likely (noncontrol) days were used as the reference days for estimation of model (1), without attempting to match by month.

Other definitions of likely days were explored. These definitions included some based on the temperature-humidity index, and others based on the maximum outside temperature for the day. A maximum outdoor temperature above 90°F turned out to be a near perfect predictor of control days. However, this basis for classifying days was not useful for defining reference or likely days. For one thing, the likely days are intended to be days when customers would have anticipated a likely control period. The maximum temperature for the day is not determined until after the anticipatory action, if any, would have occurred. More importantly, an index that nearly perfectly predicts control days does not provide any noncontrol likely days for comparison.

The "average daily temperature" used to define likely days is actually computed, per convention, as the mean of the minimum and the maximum dry bulb temperatures for the day. This mean reflects conditions prior to any notification period. The cut-off of 70°F correctly captured all true control days, and provided a roughly equal number of likely days as comparison. Classifications based on a temperature-humidity index did no better.

Impact Findings from the Load Model

The load model was applied to each rate type. As noted, this model provided both dispatchable and anticipatory impacts.

Dispatchable Impacts. Once the model was fit to the 1991 hourly load data for a particular rate type, dispatchable impacts were estimated for each 1991 control day in two steps:

1. The fitted model was used to predict load for each hour of the control day, substituting the actual weather variables for that day and setting the likely-day dummy A_d to 1.

2. Impacts were computed for each hour as the difference between the predicted and actual loads. With this definition, positive impacts correspond to demand reductions, and negative impacts to increases.

Impacts for all customers on controllable rates were estimated by summing the estimates of total impact across the separately estimated rate types. This method was expected to be more accurate than fitting a single model for all customers combined.

Impacts were estimated for each hour of each 1991 control day, by rate type. For the controllable population as a whole, the analysis showed a mean coincidence factor at the hour of the system peak (4 p.m.) of 47.1 percent, averaged over all control days, and a maximum of 49.8 percent (Table 1).

The standard error of the program-wide coincidence factor was 3 percentage points, or about 6 percent of the estimated dispatchable impact. For most individual rate types, the relative standard error was on the order of 20 to 30 percent. For the largest single rate type (Rate Class A), the relative error was also 6 percent. For a few rate types, however, the standard errors were fairly large compared to the mean loads. For these rate types, the model fit was poor, and the resulting estimates of dispatchable impacts of low accuracy.

The Rate Class A had the most accounts and largest total controllable load of any single rate type. However, the average controllable load for this group was fairly low. The mean coincidence factor for this group was 36 percent, somewhat lower than for the total population. This finding is consistent with previous estimates that indicate lower coincidence factors for customers with smaller controllable loads. As further corroboration of that tendency, all rate classes with mean controllable load below 200 kW had coincidence factors below 40 percent.

The utility has also been concerned that primary and secondary schools tend to have lower coincidence factors, because they tend not to operate during the time control periods than are likely. Applying the load model to aggregates defined by SIC, rather then rate type, yielded a coincidence factor of only 22 percent for a group of primary and secondary schools. For a group of colleges and universities, by contrast, the coincidence factor was 59 percent.

For non controlled hours, control-day impacts were generally negative, reflecting some prior and subsequent compensation for the reduced load during control periods. However, not all the load reduction was made up at other times. Summing positive and negative impacts over the control day yielded a positive balance of energy savings. Averaged over all hours, the control-day energy savings was equivalent to a 24-hour load reduction equal to about 10 percent of the total controllable load.

Anticipatory Impacts. The estimated model yields anticipatory impacts as the coefficient a_h of the likely-day dummy variable Ad for each hour h. That is, the coefficients a_h estimated by the model are themselves the estimates of anticipatory impacts for each hour h. The accuracy of the estimated anticipatory impact is indicated by its t-statistic, which is the ratio of the coefficient to its standard error. This ratio is the inverse of the relative error. As for dispatchable impacts, the anticipatory impacts for all customers combined was estimated by summing the estimates of total impact across rate types.

Table 1. Dispatchable and anticipatory Impacts at 4pm from the Load Research Model for Selected Customer Groups

Customer Group	Fraction of Total Controllable Load	Dispatchable Coincidence Factor				Anticipatory Coincidence Factor	
		Mean	Max	Std Err	Rel Err	CF	t-statistic
All customers	100%	47.1	49.8	3.0	6.3%	-0.5	-0.47
excluding extremes	84%					2.6	4.05
Rate Class A	36%	35.9	39.4	2.3	6.3%	1.9	1.95
Rate Class B	70%	51.1	56.0	4.7	9.2%	3.8	2.14
Rate Class C	50%	74.2	107.7	31.9	43.0%	-20.6	-1.73
Primary/Secondary School	13%	21.5	24.0	1.8	8.2%	3.7	5.61
College/University	80%	58.6	75.6	5.5	9.4%	5.1	2.33

For all customers combined, the anticipatory impact was negative, though by less than 1 percent. This result for the whole group was driven mainly by extremely large negative impacts for two rate types. These two rate types had 2 and 1 customer, respectively, but they were very large customers. For both these rate types, the load model had relatively poor fits, with relative standard errors of the regression of 31 percent and 43 percent (Rate Class C in Table 1). With the two most extreme rate types removed, accounting for 16 percent of the total controllable load, the overall anticipatory coincidence factor was 2.6 percent.

Adding this estimate to the dispatchable coincidence factor gives a mean actual coincidence factor of 49.7 percent, and a maximum of 52.4 percent. The latter estimate coincides with the utility's previous in-house estimate, which also combined dispatchable and anticipatory impacts.

The anticipatory impact for the total program was not significantly different from zero (t-statistic = -0.5). However with the two extreme rate types removed, the coincidence factor for the remaining customers had a t-statistic of 4.0 percent, indicating good evidence of anticipatory impacts for the bulk of the program.

In summary, while the load research data shows evidence of anticipatory impacts, these impacts are not large overall, and are difficult to identify with confidence. The anticipatory impacts is explored further below, through the analysis of the survey and onsite data.

Embedded Impacts from Survey and Onsite Data

The analysis of embedded impacts addressed both anticipatory impacts and long-term embedded impacts. Each of these impacts was assessed by using the screener survey to determine the prevalence of the type of effect, and using the analysis of the follow-up verification to quantify the magnitude of the effect for accounts that had it. Impacts were estimated separately by rate class and by SIC group. However, because of the relatively small number of follow-up verifications completed with each type of effect, some factors were estimated only for the population as a whole, then applied to individual rate types or SIC's.

Estimation Methods

Each of the impacts (anticipatory and SOP) was assessed in the following major steps:

1. The proportion of accounts reported to have any anticipatory (SOP) impacts was determined from the screener survey. This proportion $P_{AC}(Ps_C)$ was measured for each rate class c in terms of the fraction of controllable load represented by these accounts. That is,

$$P_{Ac} = \sum_{j\epsilon c} C_j A_j / \sum_{j\epsilon c} C_j, \tag{3}$$

where

C_j is the controllable load for account;

$A_j (S_j)$ is a 0/1 dummy indicating that account j reported anticipatory (SOP) impacts on the screener survey.

2. The anticipatory (SOP) embedded impact per kW of controllable load for those accounts that reported this type of impact was estimated from the follow-up verification analysis as

$$V_A = \sum_{j \in F_A} D_{Aj} / \sum_{j \in F_A} C_j, \tag{4}$$

where

$F_A (F_s)$ denotes the follow-up sample for customers reporting anticipatory (SOP) impacts

$D_{Aj} (D_{Sj})$ denotes the anticipatory (SOP) impact (kW) based on the verification for customer j.

For customers who reported an embedded effect but were determined not to during the follow-up verification, the customer's controllable load C_j is included in the denominator of V, and the corresponding impact D_j in the numerator is zero. Thus, the verification ratio V reflects the proportion of survey-reported embedded effects that were verified in the field, as well as quantifying the magnitude of the effects that were found.

3. Multiplying the proportion of controllable load that reported an effect by the impact per kW for accounts that reported effects gave the anticipatory embedded coincidence factor for each rate type C. That is,

$$CF_{Ac} = P_{Ac} V_A. \tag{5}$$

In the formula for the embedded coincidence factor, only the reported proportion P_{AC} varies by rate type c. The ratio $V_A (V_s)$ was computed across all sites that had follow-ups. Separate ratios were not computed by rate type because of the limited number of sites with verified impacts within any single rate type.

This approach leverages the screener survey data collected for a majority of accounts with the more limited follow-up data. Follow-ups were conducted with only 130 of the roughly 1,000 accounts that had completed surveys. Of those that did receive follow-ups, only about half had verified anticipatory effects. Roughly one-fifth had verified SOP effects. The use of the anticipatory and SOP ratio estimators V_A and V_S allows the verified results from relatively few sites to be projected to the population as a whole.

Screener Survey Questions on Embedded Impacts. To identify accounts that have embedded impacts, the screener survey asked customers whether they have anticipatory or SOP effects, and what types of equipment is affected. For anticipatory effects, customers were asked:

> On those days when you believe a control period to be likely, but have not yet been notified of a control period, do you take any special actions as precautions in case a control period is called?

Customers who answered yes were then asked which of a list of types of anticipatory actions are taken. The list included changing operating schedules; turning on back-up generation; switching to alternate fuels; reducing use or turning off equipment; and pre-chilling the building.

To identify SOP effects, customers were asked:

> On a seasonal or permanent basis, would any of your equipment or operating practices be different if you weren't on the controllable rate? This could include different operating schedules, types of equipment, or how it's used.

Customers who answered "Yes" were then asked which of a list of operating practices would be different seasonally or permanently if they were not on the controllable rate. The list included operating schedules, fuels used, and use of different types of equipment.

Customers on peak-controlled time-of-use rates were specifically asked what they would be doing differently if they still had the time-of-use feature, but did not have the possibility of control periods.

Computation of Embedded Impacts for Verified Sites. For a substantial fraction of the accounts that received follow-up visits or calls, it turned out that there was no anticipatory or SOP impact, even though such an impact was reported on the screener. The follow-up spent more time with the customer clarifying the types of impacts that were being asked about. The impacts determined to be present during the follow-up are considered to be "verified." These impacts should be much more reliable than those reported on the screener originally.

For each follow-up site with a verified embedded impact, the magnitude of the impact was estimated using a spreadsheet designed especially for this purpose. Copies of the spreadsheet input screens were used as data collection forms.

The spreadsheet inputs included equipment capacity and loadings at each hour of the day. For anticipatory impacts, hourly loadings were determined by the auditors for both likely and unlikely days. For SOP impacts, current practice and the hypothetical practice if the program had not existed were recorded.

The spreadsheet calculation consisted of multiplying capacity by part-load to determine the load contributed by the equipment under study, for each hour the equipment was scheduled to be on. The loadings and schedules were determined for two conditions, for which the auditors entered the necessary data on the input screens. These conditions were likely versus unlikely days for anticipatory effects, and current operation versus hypothetical operation if the program did not exist for SOP effects. For the SOP effects, current and hypothetical conditions were determined for each day type. After the loads for the particular equipment under the two alternate conditions were calculated for each hour, the difference between these loads gave the impact, also by hour.

The outputs from the spreadsheet were hourly SOP or anticipatory kW impacts for each of seven day types: summer likely day, unlikely day, and weekend; winter weekday and weekend; and shoulder weekday and weekend. The impacts were computed in the spreadsheet using engineering models, with the input loading and schedule data. The kW impacts were aggregated to daily kWh impacts for each day type. Multiplying daily energy impacts by the assumed number of each day type per year gave the total annual energy impact. Anticipatory and SOP impacts were input and estimated on separate spreadsheets.

Results of the Embedded Impact Analysis

The survey and follow-up analysis yielded a 3 percent anticipatory coincidence factor at 4 p.m., similar to that found in the load research data excluding two extreme accounts. The SOP impacts estimated from this analysis were much smaller. The total program SOP impact had a coincidence factor of less than 1 percent. The utility's previous work had indicated larger (SOP) embedded effects. That preliminary finding was one of the motivations for the present study.

One difference from the previous work is that long-term embedded impacts were defined in this study according to very tight criteria. In the audits that formed the basis for quantifying embedded impacts, customers were probed to be sure the effect described could be unambiguously attributed to the controllable rate. This approach was adopted so that evidence of these impacts would not be subject to question.

Another difference is that the impacts were quantified in this study by collecting explicit equipment operating characteristics and schedules, both for current practice and for hypothetical operations if the customer were not on the controllable rates. This information was used systematically to calculate changes in demand by hour and day types. The utility's previous estimates were based on impacts reported directly by customers or their marketing representatives in an impromptu phone survey.

A limitation of the information on long-term embedded impacts collected for this study is that, for logistical reasons, the information was collected nearly two years after the last control period. As a result, customers may have had difficulty recalling how the program affected their operations.

Summed over hours of the day, and totaled across likely days, the anticipatory energy impact was negative. However, in magnitude it was only 10 percent as large as the positive dispatchable savings. the annual SOP energy impact was also negative, but was less than 2 percent as large as the dispatchable savings.

Conclusions

New methods were developed in this study for the estimation of both dispatchable and embedded impacts. These methods offer several advantages for future applications. The software developed for the dispatchable analysis may be used on a routine basis for this or similar programs. The method can be applied to any aggregate of interest, and provides statistical measures of the accuracy of the resulting estimates.

The method applied to embedded impact estimation has broad applications to customer research. This method includes a screener survey to identify customers with particular characteristics, onsite follow-up to collect detailed engineering data for custom-designed analysis, and ratio estimation techniques to expand the follow-up sample to the general population. This procedure allows the systematic extension of site- and equipment-specific engineering studies to a general population.

References

Electric Power Research Institute (EPRI) 1988. *Customer Response to Interruptible and Curtailable Rates.* EPRI EM-5630, RP2343. Electric Power Research Institute, Palo Alto, California.

Joyce, J., Dion, S., Moe, R., Johnson, T., and Spada, M. L. 1993. "Impact and Process Evaluation of Commercial and Industrial Load Management Programs." *Proceedings of the 1993 Energy Program Evaluation Conference*, pp. 664. Energy Program Evaluation Conference, Chicago, Illinois.