Can Dispatchable Pricing Options Be Used to Delay Distribution Investments? Some Empirical Evidence

Robert L. Cruz, Pacific Gas and Electric Company
Dennis M. Keane and Michael J. Sullivan, Freeman, Sullivan & Co.

In recent years, there has been growing support for the idea of geographically targeting demand-side management (DSM) programs. In particular, pricing and service reliability options designed to shave local area peaks offer great promise in reducing utility costs by deferring investments in distribution capacity.

For this promise to become a reality, however, the inherent skepticism of utility distribution planners must be overcome. Currently, many planners entertain serious doubts about the ability of DSM programs to reliably deliver load relief at the local level on critical days—especially programs which rely solely on pricing mechanisms to induce customers to curtail loads. Planners are understandably reluctant to risk overloads by trusting in demand-side programs which may not perform as promised.

This paper attempts to shed some light on the ability of a dispatchable time-of-use pricing option to effectively achieve peak load reductions, by presenting findings from a 1992 experiment conducted by Pacific Gas and Electric Company as part of its Model Energy Communities Program. The experiment involved installing price-sensitive thermostats on over 90 residential customers' HVAC systems, and monitoring load responses to utility-dispatched price signals.

Results are presented from experimental operations conducted during the summer of 1992. In addition to presenting estimates of the load impacts achieved by the program, the paper also describes efforts to evaluate customer satisfaction. The paper concludes with a discussion of the steps DSM advocates must take to allay planner concerns, if they are to accelerate the implementation of dispatchable pricing options at the local level.

Introduction

Pacific Gas and Electric (PG&E) Company's Model Energy Communities (MEC) Project is an effort to determine whether intensive marketing of demand-side management programs in a concentrated geographic location (PG&E's Delta Division) can effectively control load growth, and allow utility planners to defer investments in distribution capacity. One component of the MEC project, tested in 1991-1993, involved using price-sensitive thermostats (PSTs) to manage residential customer demand. These devices offer the potential to achieve significant reductions in the local area summer peak load, by combining a dispatchable time-of-use (TOU) rate with a "smart" thermostat which can respond to the varying prices by altering the set point on the customer's air conditioning unit. Despite this promise, though, utility distribution planners remain skeptical of the ability of PSTs to reliably reduce local peak loads. Until they are thoroughly tested and proven, planners cannot be expected employ PSTS solely because of their "potential."

The uncertainty about the reliability of PSTS stems partly from concerns about hardware performance. More importantly, though, the achievable load impacts depend crucially on customer behavior during critical peak periods—since customers have the ability to change thermostat set points and override the PSTs. Because of these concerns, PG&E decided to evaluate the effectiveness of PSTs by conducting a limited field experiment. In late 1991 PSTs were installed on a sample of 96 households, and data were collected on customer loads, override actions, and weather. The experiment continued through the end of 1993.

This paper reports the results from the first summer of operations, specifically covering the June-August 1992 study period. It focuses on two of the most important objectives of the experiment:

- Estimating Load Impacts - Evaluating the ability of PSTs to reduce individual customer and aggregate substation load, particularly on very hot days when distribution facilities are constrained and approach overloaded conditions; and

- Evaluating Customer Satisfaction - Determining whether the PSTs satisfied customers' needs for better-managed electric bills and comfortable living environments.

The paper concludes with a discussion of the steps DSM advocates must take to allay planner concerns about the reliability of PSTs as an effective means to control peak loads in local distribution areas.

Program Design

Direct vs. Indirect Control

There are two basic approaches that utilities can take when designing dispatchable load management program offerings. One philosophy is to design programs which directly control customer loads. For example, a radio-control switch can be installed on a particular residential appliance, and the utility can then remotely turn it off when desired. The second philosophy is to rely on price signals to manage loads—a type of indirect control. With this approach, the utility dispatches very high prices during critical periods hoping to induce customers to take actions to reduce their loads and avoid the high bills which would otherwise occur.

Each approach has its advantages. Direct control is simple to explain to customers and easy for them to understand. It is also generally perceived by local distribution planners to be more reliable than a pricing signal in achieving load relief.[1] The indirect or price-based approach, on the other hand, has the advantage of being much more flexible. On each curtailment day customers get to choose which appliances are most valuable (and, therefore, remain in operation) and which are low priority (and thus get curtailed). A household's preferences for using different appliances may change from day to day (and, in some cases, even from hour to hour) depending upon its lifestyle, the weather and other factors. With an indirect control scheme, the household can decide which appliances to curtail based upon its priority rankings on that particular day.[2] The PST experiment represents a test of the effectiveness of using an indirect control approach.

Dispatchable TOU Rates

In May 1991 PG&E received permission from the California Public Utilities Commission to offer a special rate, Schedule E-B7 ("Experimental Residential Dispatchable Peak Time-of-Use Service"), to customers in the MEC planning area. During the summer months, Schedule E-B7 has three pricing periods defined as follows:

Pricing Period	Definition
High	Weekdays, 4:00 p.m. -8:00 p.m.
Low	All Other Hours
Critical	Dispatched at PG&E's discretion (but no more than four hours per day, and no more than 100 hours per year)

These prices have been designed to send customers signals which better reflect the time-varying nature of PG&E's costs. The first two pricing periods are similar to PG&E's standard time-of-use (TOU) rate, Schedule E-7, which is available to residential customers system-wide. TOU periods are pre-set according to day of the week and hour of the day, with a higher price charged during the peak period and a lower price levied during the off-peak period.[3] The third pricing period, though, is quite different in that it can be dispatched at the utility's discretion. By dispatching very high prices during critical periods, the utility has a means to induce customers to control loads for a limited number of hours per year.

During the June-August 1992 study period, the following prices from Schedule E-B7 were sent to customers:

Pricing Period	Price ($/kWh)
Low	$0.07278
High	$0.32186
Critical	$0.55856

There was a very large price differential between electricity consumed during the peak and off-peak periods, with peak power costing households over four times as much as off-peak power. During the limited number of critical hours, the price rose even more, to a level almost 75 percent higher than the on-peak price. These price differentials offered customers a distinct incentive to manage their loads to limit household usage during high and critical periods.

How the PST Works

In order to implement a dispatchable TOU rate, a means must exist for the utility to communicate price signals to customers. The PST does this, utilizing a two-way telecommunication system where the house receives the price signals sent by PG&E via phone line and modem. In addition, the PST can help the customer respond to the prices by automatically adjusting the air conditioner thermostat to reduce energy usage when the price is high.

This automatic control takes place via a set of programmed instructions in the device's memory. For each pricing period, the customer can program thermostat set points which appropriately balance its competing desires for low electric bills and environmental comfort. For example, during the summer a household might program the following set points:

Period	Thermostat Setting
Low	65 degrees
High	75 degrees
Critical	90 degrees

In this example, the household has decided that it wants to remain comfortable during the low-cost hours, since electricity is relatively inexpensive. Consequently, it has set the air conditioner thermostat to a pleasant 65 degrees. The household is concerned, though, about the effect that a 65 degree setting during the high-price period will have on its monthly bill. So during the high-price period it sets the thermostat ten degrees higher, at 75 degrees, to prevent the air conditioner from running as long when electricity is more expensive. For the critical periods, when the price is very high, the household has set its thermostat even higher, at 90 degrees, to further restrict the air conditioner's operating time.

At the time of each installation, the electrical contractor doing the work initially programmed the thermostat settings based upon the household's stated preferences. At any time afterward, though, the customer has the ability to re-program the device to alter thermostat settings to better meet its comfort/cost needs.[4] The PST also has a simple "safety valve" override capability which allows the household to temporarily assert control over the thermostat without having to go through the steps of reprogramming the device. Instead, the customer merely pushes the override button that sets the thermostat manually. The override feature is designed to make the devices easier to use and more acceptable to customers. However, the override capability has a potential downside. If, on extremely hot days the override buttons are used simultaneously by many customers, they have the potential to negate the load relief which would otherwise be obtained, and seriously reduce the value of the program to the utility.

Load Impacts

Analysis Approach

The 96 PST participants were recruited from the households living in the MEC planning area using a combination of newspaper advertising and personal recruiting by PG&E. In addition, some customers contacted PG&E to participate after hearing about the program from friends and neighbors. To qualify, customers had to own a relatively new central air conditioner. Participation was voluntary and no direct incentive was paid by PG&E.

One of the primary objectives of the PST experiment was to determine the effects of implementing a TOU rate and dispatchable critical price signal on participants' load shapes. The high peak and critical prices were expected to reduce customer loads in two ways. First, because the prices being charged (especially during the critical period) were quite high, customers were likely program their devices for fairly high indoor temperatures, resulting in lower than usual air conditioner usage. Second, customers were also expected to be less likely to utilize the override capability to increase their comfort level during peak and critical periods, again because of the financial cost of doing so.

The analysis approach generally employed by researchers to evaluate the impacts of dispatchable load management or TOU programs is to compare the "disturbed" loads of the participating customers (i. e., the loads during the period when the program was dispatched or prices were higher than normal) to estimates of what their loads would have been in the absence of the program (i.e., their "undisturbed" loads). The disturbed loads are metered directly during critical periods using load profile monitoring equipment installed on the houses of the participating customers. The undisturbed loads, however, must be estimated, since they are hypothetical: they represent what the participating customers' loads would have been during the critical period if they had not participated in the program.

To obtain estimates of undisturbed loads, researchers generally employ some sort of control group, The control group is often a collection of similar customers who are not participating in the program. The goal in selecting the control group is to choose customers who are similar to the participants in as many respects as possible, so that the only thing different between the two groups is whether or not they participate in the program.[5] In this way, any differences in the two groups' load shapes can be attributable solely to the effect of the program.

Two types of control groups were used to estimate the load impacts for the PST program. To estimate the impact of the TOU rate structure on customer loads, metered load data from the 96 participants were compared to load data from a large sample of about 600 customers in the MEC planning area who did not participate in the program. Since this sample was not scientifically selected to match the characteristics of the PST participants, it is not an ideal control group. Nevertheless, the customers in this sample are sufficiently similar to the PST program participants to make meaningful comparisons and draw conclusions. [6]

For purposes of evaluating how the critical price signal impacts customer loads, however, an alternative approach was followed. Since critical price signals occurred only on a small subset of summer days, participant data from the numerous non-critical days were employed as the control (as described in endnote 5). Using participants as their own control group is a preferable approach for estimating the effects of critical price signals, since undisturbed usage patterns are likely to be estimated more accurately from "same-customer data" on similar days than by similar-customer data on me same day.

Load Impacts of TOU Rate Structure

Figures 1 and 2 present a comparison of the average weekday load shapes for PST program participants and control group households, respectively. Each figure shows multiple load shapes (one for each weekday of the study period) arranged in order of increasing maximum daily temperature. In both cases the magnitudes of the hourly loads (i.e., the vertical heights of the load curves) increase with temperature, as expected. The shapes of the load curves differ dramatically, though, between the two figures. The load curves for the control group in Figure 2

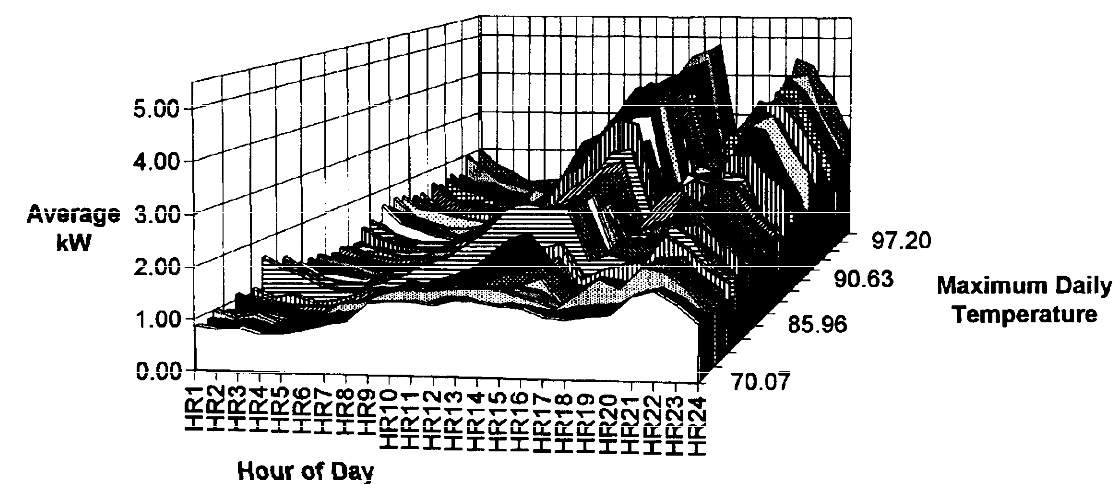

Figure 1. Average Weekday Load Curves of Program Participants, Sorted by Maximum Daily Temperature

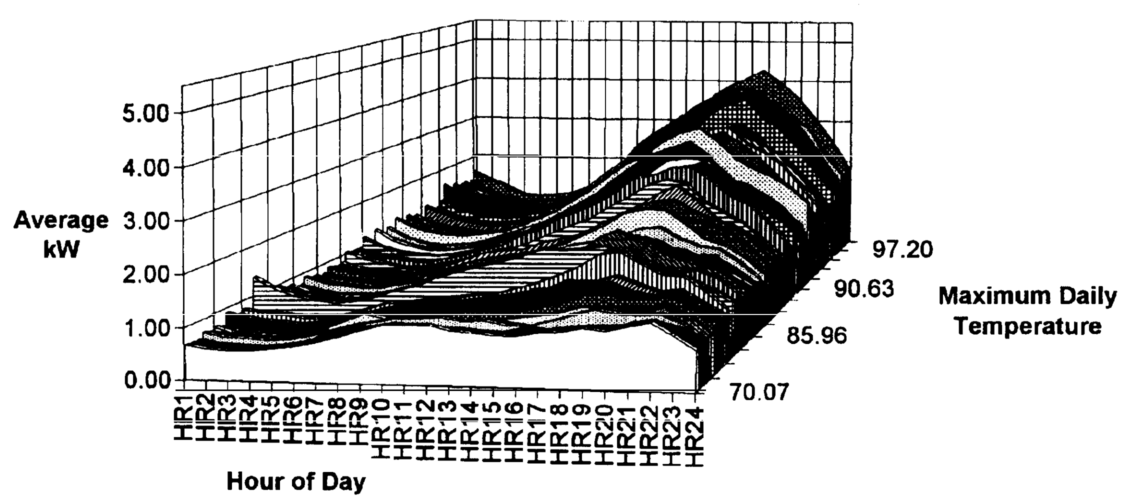

Figure 2. Average Weekday Load Curves of Control Group Customers, Sorted by Maximum Daily Temperature

have the conventional shape, rising throughout the afternoon to peak during the early evening hours, then declining. The shapes of the PST participants, on the other hand, show a pronounced "precipice" in the late afternoon-with loads dropping sharply at around 4:00 p.m., then rising slowly to a second, later peak around 9:00 p.m.

This dramatic drop-off in the loads of program participants is strong evidence that many customers have set their thermostats during the high-price period at substantially higher levels than during the low price period. The higher thermostat settings cause air conditioner load (and thus total household load, too) to decrease rapidly at 4:00 p.m. when the high-price period begins on summer weekdays. Figures 3 and 4 show additional evidence of this effect. These figures show—for the participant and control groups, respectively–how average load shapes on the ten hottest summer weekdays vary with the monthly kWh usage of the customer.[7] Examining the control group customers in Figure 4 first, the conventional load pattern is again seen for all three usage groups, with the high usage customers showing higher loads during all hours than the lower usage categories. In Figure 3, though, the

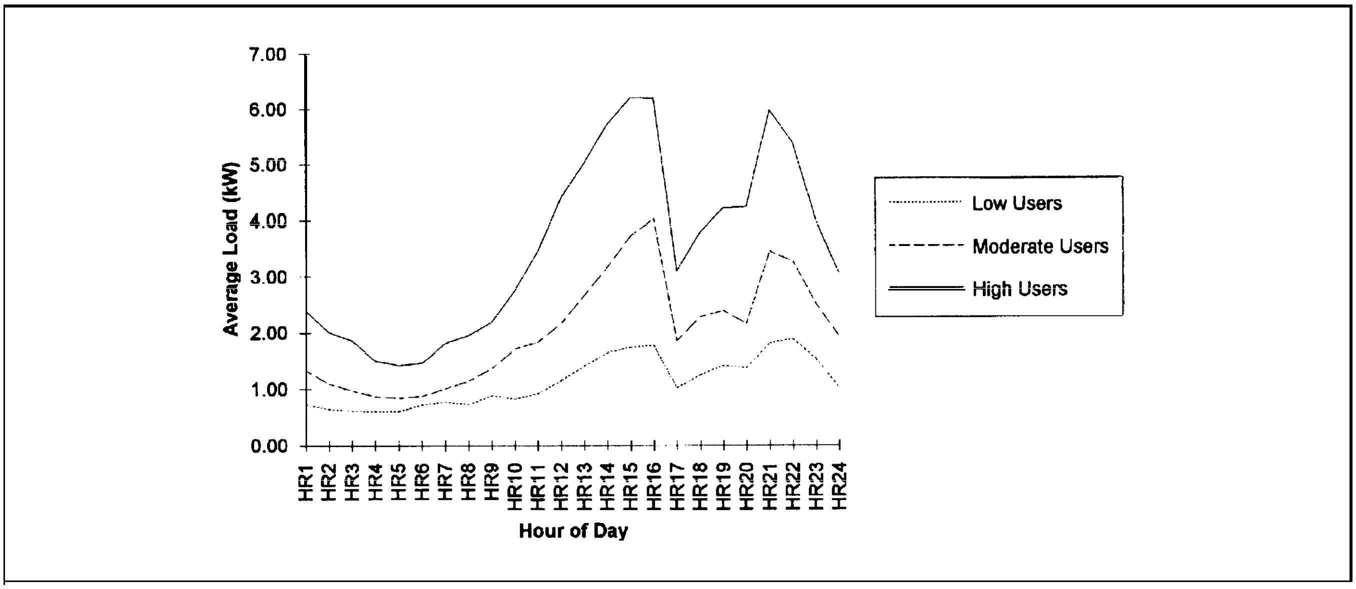

Figure 3. Average Load of Program Participants, Ten Hottest Summer Weekdays, By Energy Usage Group

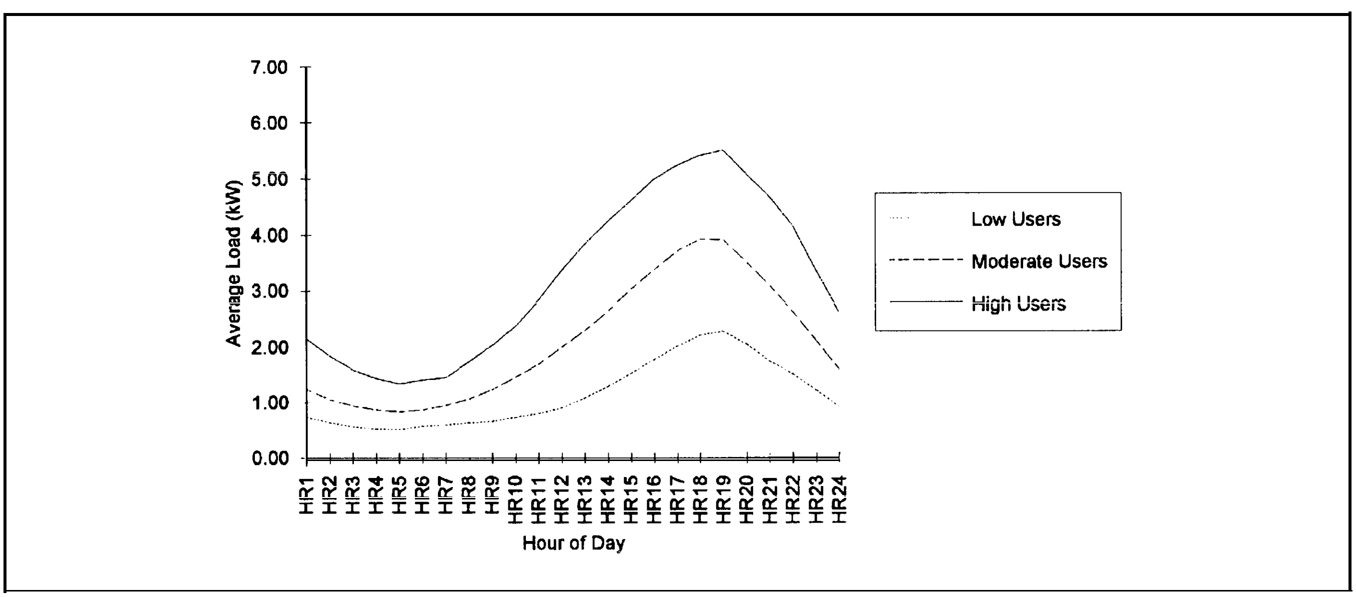

Figure 4. Average Load of Control Group Customers, Ten Hottest Summer Weekdays, By Energy Usage Group

sharp drop-off in participant loads at 4:00 p.m. is readily apparent, and it is particularly pronounced for the high usage customers.

Load Impacts of Critical Price Signals

The comparisons between participant and control group customer load curves provide clear evidence that thermostat pre-programming during the high-price period (to minimize the adverse bill impacts of the high rates) can be an effective means to reduce customer load. This effect is apparent during high-price periods on virtually all summer weekdays when hot weather conditions occur, and can be viewed as an everyday outcome of the time-of-use structure of the E-B7 rate. But also of interest is whether any *additional* load reductions were achieved as a result of dispatching the even-higher critical price. To examine this, the analysis was confined to program participants, comparing their average load shapes on the hottest critical days with their shapes on the hottest non-critical days.

Figures 5, 6, and 7 show this comparison for the high, moderate, and low energy usage groups, respectively. [8] In all three instances, the average loads on the critical days just prior to the dispatch of the critical price (i.e., at 4:00 p.m.) exceed the average loads on the non-critical days. In each case, too, average load drops by a slightly greater amount than on non-critical days upon dispatch of the critical price. [9] So the presence of the critical price signal does appear to increase the magnitude of the initial load drop by a small amount.

In addition. the price signal appears to be effective in slowing the rise in load which occurs after the initial drop-off. This is particularly apparent for the high and moderate usage groups, where the loads remain lower until after 8:00 p.m. when the critical period ends. [10] So the critical price appears to be quite effective in increasing the load impacts achieved during the later hours (from 5:00 to 8:00 p.m.) of the peak period.

As expected, participants' loads are quite high at the conclusion of the peak period, as they initiate electricity-using activities (e. g., clothes and dish washing) that they have put off during the high-price hours. This is not a problem, however, since the loads of other customers in the MEC planning area have dropped significantly by 8:00 p.m. So despite the fact that participants' peak loads have not been reduced much (if at all), the load shifting caused by the PST program is effective in reducing the local area peak.

Although these results are very promising, one cautionary note should be sounded. The summer of 1992 was quite mild compared to the norm in the MEC planning area. The temperature never exceeded 100 degrees, and there was a notable absence of heat storm activity (i.e., three or more consecutive days of very high temperatures). Because of the mild weather, the PST program was not fully tested under the type of conditions likely to occur. Consequently, its ability to produce consistent and reliable load impacts under heat storm conditions (and thus permit a planner to confidently defer a planned investment in distribution capacity) could not be determined from the 1992 data. Fortunately, weather conditions were more extreme in 1993, so analysis of these operations should produce a more conclusive determination of the PST system's peak load reduction capability.

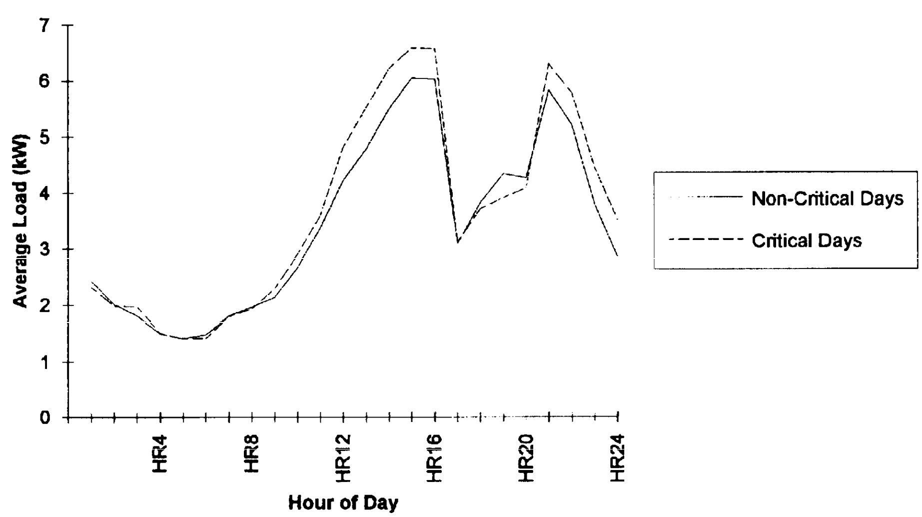

Figure 5. Average Load of High Usage Program Participants, Ten Hottest Summer Weekdays, Critical vs. Non-Critical Days

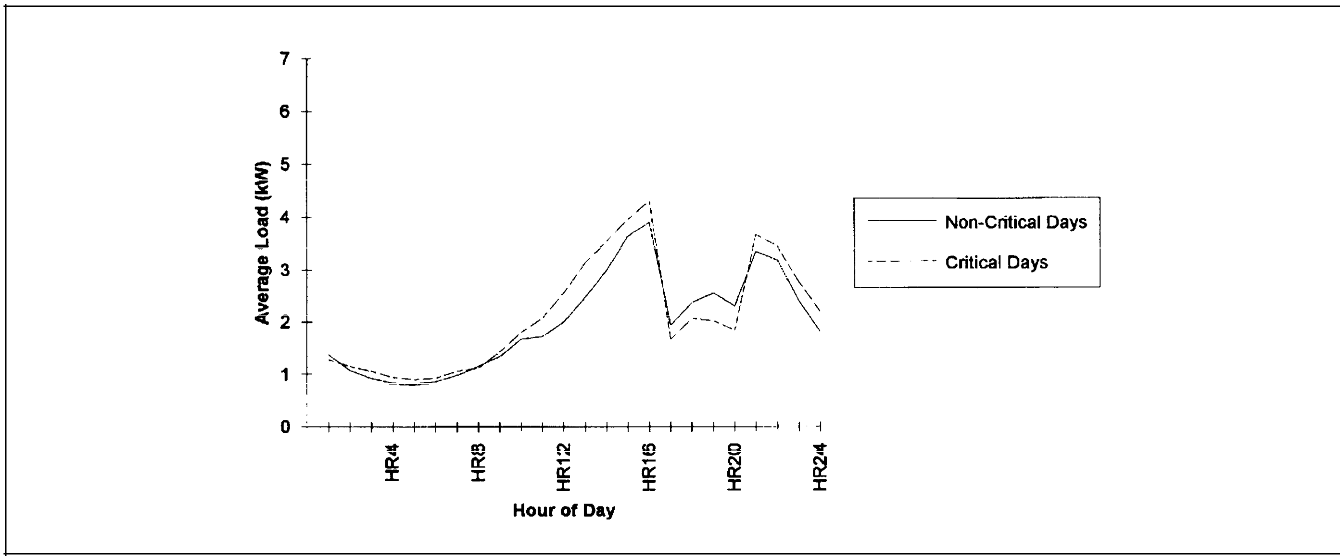

Figure 6. Average Load of Moderate Usage Program Participants, Ten Hottest Summer Weekdays, Critical vs. Non-Critical Days

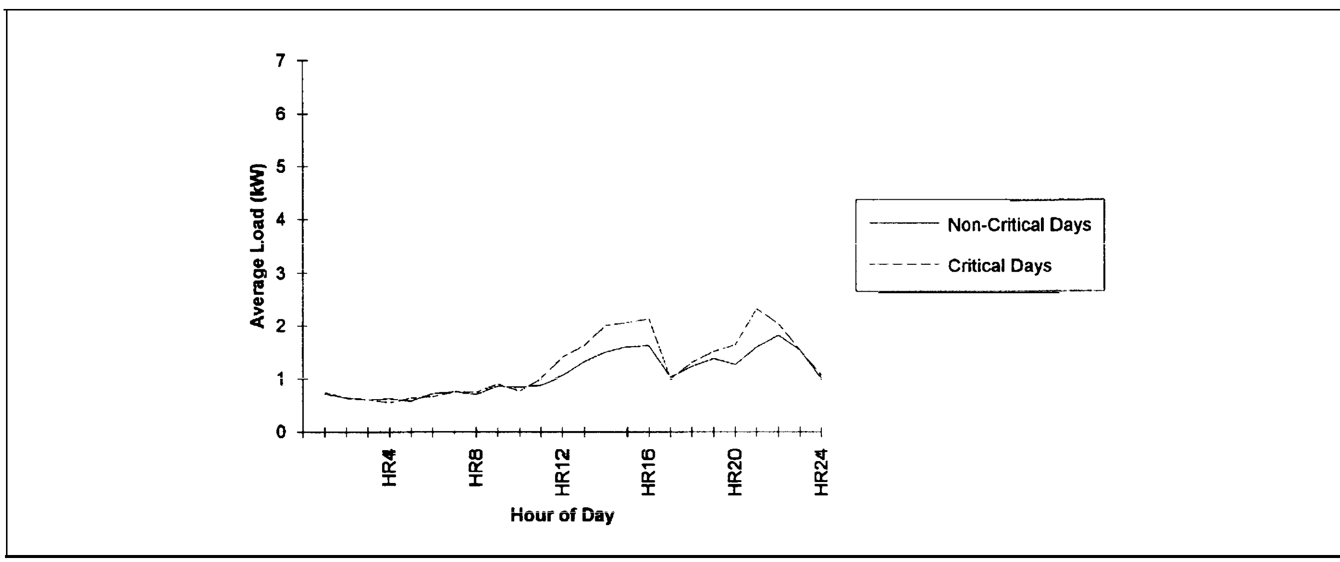

Figure 7. Average Load of Low Usage Program Participants, Ten Hottest Summer Weekdays, Critical vs. Non-Critical Days

Customer Satisfaction

Measurement Approach

A combination of focus group techniques and telephone interviewing were used to assess the satisfaction of participating customers in the PST experiment. In-depth information about customers' opinions and attitudes concerning the PST system were obtained in a focus group which was conducted on December 3, 1992. Representatives of nine of the 96 households participated in the focus group. The opinions and attitudes expressed in the focus group were then analyzed and used to formulate a telephone survey which was administered in January 1993 to all of the households in the PST study population who agreed to participate. A total of 67 of the 87 remaining households (77 percent) were interviewed during the telephone survey. The following section summarizes the focus group and telephone survey results.

Customer Satisfaction Results

Customers expressed a high degree of satisfaction with the PST program. All nine of the focus group participants said they were generally satisfied, and 63 of the 67 respondents (94 percent) to the telephone survey stated that they were either satisfied or very satisfied with the program. Table 1 summarizes the key drivers of customer satisfaction as indicated by the aspects of the program customers mentioned when asked to explain their satisfaction ratings.

Customers' satisfaction with the PST program is strongly influenced by two things:

- their perception of its ability to achieve cost savings; and
- their perception that it enhances their ability to monitor and control energy use.

About 92 percent of the respondents reported that their energy bill was lower after joining the program, and about 64 percent said their bill was significantly lower. In explaining why they were satisfied with the program, about half of the respondents mentioned these bill savings. The respondents also liked the fact that the program made them feel more in control of their economic relationship with PG&E. The PST provided a means to determine how much electricity was costing at any point in time, and also enabled participants to monitor their energy use in real-time. Both of these features contribute to customers' perceptions of having greater control over their energy bills.

Everyone in the focus group, and all but one of the telephone survey respondents, indicated a willingness to continue in the PST program if PG&E offered it. The overwhelming majority of respondents (94 percent) also said that they would recommend the PST program to their friends—an important indicator of customer satisfaction. The majority of the respondents even said that they would be willing to pay some amount to defray the cost of the PST and associated communications equipment—although it is doubtful that they would pay the full cost of the system.

Summary and Conclusions

The first-year results from the PST experiment are quite promising, demonstrating that substantial load reductions can be obtained by providing customers with a combination of time-of-use rates and a thermostat control technology which schedules air conditioning according to those rates. The system's effect on customer load results primarily from the operation of the thermostat which allows the temperature in the house to rise during the high-priced period (i.e., from 4:00 to 8:00 p.m.). Additional load reductions can be obtained by dispatching the critical price signal. These very high prices result in a small increase in the magnitude of the initial load reductions obtained. They are also effective in slowing the rise of the air conditioner load after the initial load drop, by prolonging the period of time during which the air conditioner remains off as the indoor temperature rises to the level specified for the critical period.

The experiment also demonstrated that customers can be highly satisfied with electric service provided under the above conditions, despite the fact that they suffer obvious inconvenience and discomfort in return for the savings they receive. Customers were not just satisfied with the service they received during the experiment—they wanted to continue participating, they would recommend it to their friends, and they were even willing to pay more for it. Thus, there appears to be significant opportunity for

Table 1. Key Drivers of Customer Satisfaction

Positive Factors Mentioned	Negative Factors Mentioned
Cost Savings (23 Customers)	Transtext Equipment Failure (3 Customers)
Ability to Control Energy Use (10 Customers)	Unexpected Telephone Charges (2 Customers)
Availability of Usage Information (8 Customers)	Complicated User Interface (1 Customer)
Ease of Use (5 Customers)	Discomfort on Hot Days (1 Customer)
Service Responsiveness (1 Customer)	No Savings (1 Customer)

both the customer and PG&E to benefit from the implementation of a program which combines a dispatchable time-of-use rate and a thermostat control technology.

On the negative side, though, the customer premise equipment used in the PST experiment was very expensive, costing about $1,500 per installed point. Equipment reliability was also a problem, requiring PG&E to devote substantial management resources to identifying and correcting technical problems discovered during the experiment. So despite the promising load impacts achieved by a dispatchable rate, the high cost and uncertain reliability of the equipment make PST in its experimental configuration a cost-ineffective option.

However, the results of the PST experiment suggest that it may be possible to achieve similar load reductions and improvements in customer satisfaction using a significantly less expensive and more reliable "low-tech" approach which does not rely on a sophisticated communications system: a standard time-of-use meter in combination with a timer-controlled thermostat. The PST experiment revealed that most of the load impacts and customer cost savings obtained by the system were achieved because the customer's thermostat allowed the temperature in the home to rise significantly between the hours of 4:00 and 8:00 p.m. A communications system and sophisticated control technology are not required to accomplish this result. The only devices required to produce this effect are a residential time-of-use meter and a timer or setback-type thermostat. In theory, these two devices (costing less than $300 per installation, and virtually nothing in incremental operating cost for the utility) could achieve load reductions and cost savings similar to those obtained using the PST system—which costs $1,500 per installation and results in significantly higher incremental operating cost for PG&E.

But will planners place a sufficient level of trust in a "low-tech" version of the PST program? The answer to this question is critical, since utility savings from any targeted DSM program depend upon planners trusting in that program's ability to perform. Before deciding to defer an otherwise-required investment, prudent planners will want convincing evidence that the program will deliver the required load relief when needed.

In particular, planners are concerned about whether the "low-tech" program can produce consistent and reliable load impacts during prolonged periods of three or more very hot days (when local area loads typically peak). Because these heat storm conditions did not occur in 1992, this question remains unanswered for the PST program. A "low-tech" version of the program, without the ability to dispatch critical price signals, would be even more suspect.

These concerns on the part of planners are legitimate and must be addressed by DSM advocates hoping to implement TOU pricing as a means to defer local capacity investments. The best way to persuade planners that a "low-tech" program is workable is to actually demonstrate its technical and economic feasibility through a carefully planned and managed field experiment. The following steps are recommended to accomplish this task:

1. Design a program combining an area-specific TOU rate with a "low-tech" thermostat which permits customers to program set-points which vary by day-type (i.e., weekday vs. weekend) and time of day.

2. Identify a geographic location for the experiment which has plenty of excess distribution capacity (so that the program, if it fails, will not have any adverse consequences).

3. Recruit from 200 to 300 households to participate in the experiment, place them on the area-specific TOU rate, and install thermostats and load profile metering equipment at their premises.

4. Select a separate sample (also numbering 200 to 300) of similar customers to serve as the control group and install load profile metering equipment at their premises (but do not place them on the TOU rate and do not install thermostats).

5. Conduct the experiment during the summer period for one year (or more, if the summer turns out to be mild and without heat storms) and collect hourly load data from participants and control group customers.

6. Estimate the load impacts attributable to the experimental program and extrapolate the results to produce estimates of what a full-scale program could achieve.

By following these steps, distribution planners can be shown hard empirical evidence that the program either works or does not. If the results are favorable, then the experimental results can be much more persuasive than mere assertions by DSM advocates that "the program will work" in convincing planners to implement a full-scale program to achieve the benefits of deferring distribution capacity investments.

Endnotes

1. This perception is due to the fact that direct control does not depend upon customer price-responsiveness to be effective. There is a widespread concern among planners that on very hot days many customers will continue using electricity despite facing a high price. At some point, though, the price can be set high

enough to ensure that most customers will respond by curtailing load.

2. This is contrasted with end-use specific direct control programs (e.g., air conditioner control) where the end-use in question will be curtailed regardless of how valuable it might be to a household on a particular day.

3. The peak (or high price) period for Schedule E-B7, though, occurs later in the day and is shorter than the peak period for Schedule E-7 (which is defined as noon to 6:00 p.m. on weekdays). PG&E designed this later peak period in order to target the tariff's high prices to the hours when the Delta District local distribution peak occurs.

4. It should be noted, though, that a certain amount of training is required before a typical household understands how to re-program its device, a task which is somewhat similar to programming a video tape recorder to record a television program automatically.

5. Sometimes participants can act as their own control group, for programs where the "treatments" do not occur every day. For example, with load control programs, where operations are only dispatched on a small number of days, participant load data on non-operations days can be the control, forming the basis for the "undisturbed" load estimate to which the "disturbed" or "treatment" load is compared.

6. Load profile meters were installed on the control group households to collect load data for evaluating other MEC programs. PG&E did not actively market demand-side management programs to these customers, so that they could be used as a standard of comparison for evaluating program impacts. Since these households were located in the same MEC planning area neighborhoods as the PST participants, they are likely to be similar in terms of dwelling size and age, appliance stocks, and household size. However, no formal tests of comparability were performed. When analyzing the data, customers in both the participant and control samples were stratified by usage to further ensure comparability between the two groups.

7. Customers are grouped into three usage categories according to their average monthly usage: high (more than 1,500 kWh per month), moderate (between 1,000 and 1,500 kWh per month), and low (less than 1,000 kWh per month).

8. Of the ten hottest summer weekdays, three were critical days and seven (including the hottest day) were non-critical.

9. After dispatch, the average load drops to below its level on non-critical days for the moderate usage group. For the high and low usage groups, average load on critical days drops to about the same level as on non-critical days.

10. At first glance, this result does not seem to hold for the low usage group. However, the average load for this group on critical days is substantially higher than on non-critical days even prior to the initiation of the critical price. If the critical load curve is normalized (i.e., proportionately reduced at each hour) to match the level of the non-critical curve just prior to the start of the critical period (i.e., at 4:00 p.m.), then the resulting normalized curve would lie below the non-critical curve throughout the 4:00-8:00 p.m. period.

Reference

Freeman, Sullivan & Co. 1993. *1992 TranstexT Experiment Evaluation.* Final Report Prepared for Pacific Gas and Electric Company, December 1993.