Peak Time Rebate's Dirty Little Secret

Craig Williamson, Energy Insights, Boulder, CO
Kelly Marrin, Energy Insights, Riverside, CA

Abstract

Peak Time Rebate (PTR) rates have attracted a lot of attention as a way to implement dynamic pricing without making some customers (the "structural losers") pay more. In this paper we reveal PTR's potential "dirty little secret" – many customers will get paid for doing nothing, based on their normal random variations in electricity use. We will look what causes this phenomenon, how utilities offering PTR are addressing this issue, and what activities a utility might consider to ameliorate the situation.

We first analyze examples of individual customer load data to show how random variations in customer usage can result in PTR payments to customers who are not participating in an event. We also explore the issues that surround designing a customer-specific baseline and use examples to show how different baseline calculations affect the rebate. Finally, we look into the experience of two California utilities moving forward with universal residential PTR. We look at their estimates of revenue loss and how the revenue loss will be handled within the regulatory system.

PTR programs are known as carrot only, or win-win programs, and are easily approved by commissions, but while the program is good for customers it poses problems for a utility. We believe that there is significant potential for revenue loss associated with a PTR program and it is important for any utility considering proposing a PTR program to fully understand any negative implications.

Introduction

Traditional Time-based Pricing

Time-based pricing for residential customers is under consideration in many jurisdictions across North America. This is driven by two main motivating factors. One factor is demand response. Dynamic pricing can reduce the highest peaks on a utility system, thereby reducing capacity needs and the corresponding costs. The other factor is the wave of advanced metering infrastructure (AMI) implementations across the industry. The demand response, load shifting, and conservation value of time-based pricing can help justify the huge expense of AMI implementation for utilities.

Time-based pricing allows customers more influence on their energy bills. With higher differentials between time period prices, people can shift load and save money – and not just a few cents here and there. Two of the more common time-based pricing approaches, time-of-use (TOU) and critical peak pricing (CPP) provide these benefits but also have features some see as problematic. When these rates are designed, the timing of the on and off-peak periods and the associated prices will inherently create structural winners and losers. This means that without changing usage patterns at all, some customers will pay more and some will pay less. Those who start with a higher percentage of on-peak energy use (structural losers) will pay more, those with a lower percentage of on-peak energy (structural winners) will pay less, when compared with their standard residential rate. For example, working couples with no children in the home, who work during the weekdays, will naturally benefit on a TOU rate because the on-peak period will occur when no one is home. A rate that is designed to be revenue neutral will take care of this on a class basis, making the class revenue the same if all customers move to the rate, but it will not help individual customers. For a rate with a few critical periods each summer, such as a CPP rate, the random variation in customer loads will also give some customers an advantage or disadvantage on a particular event day. For example, if the customer happens to be on vacation during an event their usage will be much lower than usual. The abnormally low usage will reduce what would have been a high on-peak charge. On the other hand, they could be hosting a social gathering on an event day, using more energy, resulting in a higher charge. Over time and across customers, these random variations will balance out.

For either a TOU or CPP rate, energy bills will increase for those structural losers unwilling or unable to change their energy use patterns. This can be a public relations issue for the utility, especially because these rates are often perceived (whether justified or not) to hurt the elderly and lower income customers most. Implementing time-based pricing on an opt-out or mandatory basis can be legally difficult. For example, customers in California have a portion of their usage protected from rate increases under California law AB1X. This makes the rate difficult to design; especially considering some portion of the population (about 40% in SCE territory) will never experience the on-peak rate.

Some have tried to alleviate the problem with structural losers by making TOU and CPP rates optional. This can help – those who would be hit hardest can simply stay off the rate. But this also results in much lower load shifting – people will choose to go on the rate to save money, not because they can shift load. While some may shift once on the rate, the load change will be significantly less than if the rate were mandatory.

Even for those whose usage patterns mean that they pay about the same as they would on a non-time differentiated rate, these rates include both a penalty and a reward. If the customer reduces their energy use during high-priced periods, they can save money. But if they increase use during high-priced periods, they pay a penalty. Many in the industry prefer to use rewards and incentives without penalties.

The Allure of Peak Time Rebate

Peak Time Rebate (PTR) rates have attracted a lot of attention as a way to implement dynamic pricing and induce load response while bypassing these two issues: structural winners/losers and having a penalty as well as a reward. PTR is a dynamic pricing program with event days, but no high critical price. It uses a carrot-only rebate approach, based on changes from each customer's past usage. PTR is gaining in popularity, because it is seen as a win-win for customers and is easily approved by utility commissions. It also improves participation because there are no issues with a mandatory vs. optional rate. Since there is no penalty, all customers can be defaulted onto the rate. Both Southern California Edison (SCE) and San Diego Gas and Electric (SDG&E) are instituting universal PTR with their AMI implementations.

Customers receive a rebate for reducing their usage below a customer-specific baseline. The baseline is an estimate of what the customer's energy use would have been in absence of an event. There are many ways to calculate a customer-specific baseline. The most common method is using the average of the highest three of the past five eligible event days. Both SDG&E and SCE have employed this method. However, there are some other methods to consider, including lengthening the period to draw days from, and making adjustments to the baseline to reflect weather on the event day.

So what's not to love about PTR? Customers who reduce peak usage save money, those who don't reduce do not suffer, commissions love it, and it is a PR dream. Unfortunately, PTR has a "dirty little secret." Many customers receive rebates for normal random variations in electricity use. Because PTR is one-sided, there is no offsetting of this effect across customers or across events. If, without taking any actions to reduce load, based on normal random variations in usage, a customer uses less on one event day and more on another, they receive a rebate for the first, but do not pay anything for the second. Therefore, the very nature of a PTR rate prohibits it from being revenue neutral across a class. For any season a portion of the rebates a utility pays will be for actual load reduction and another portion will be for random variations in customer usage. Preliminary utility analyses presented in cases before the California Public Utility Commission (CPUC) show that the proportion of rebates going to non-participating customers is quite large, upwards of a third of total PTR payments, depending on the level of demand response assumed. We feel it is important for the industry to fully understand this issue as it represents a significant loss to utilities offering PTR rates

The remainder of this paper discusses three issues surrounding this problem: investigating and explaining the causes of this phenomenon, learning how utilities that plan to offer PTR are addressing this issue, and examining methods that might ameliorate the situation.

Revenue Loss and Other PTR Specific Issues

Natural Variation in Customer Load Shapes

A PTR rebate is calculated based on the difference between a customer's peak period usage on an event day and their baseline peak period usage. The baseline is an estimate of a customer's usage on the event day in the absence of an event. Anyone who has worked with residential interval load data understands that this class of customers has a lot of daily variation in their individual load shapes. They go on vacation, work odd hours, purchase new appliances, and have relatives come to visit. Even if occupancy and appliance holdings are static, use of those appliances varies from day to day –use of clothes washers and dryers, ovens, and even home entertainment and computers can vary widely from one day to the next. This variability causes deviations from even the most carefully crafted baseline that are purely random and not due to demand response actions taken by the customer. Sometimes this random variation results in higher on peak energy use, and sometimes it results in lower on peak use.

To illustrate this phenomenon we use examples from real load data for two individual customers. This residential load data from the Pacific Southwest is from the Center for Electric End-Use Data (CEED). We chose customers who had higher than average usage to capture central AC load, usually a target of DR programs. These customers were not participating in any sort of demand response program. We also chose two hypothetical event days, by selecting days where the temperature reached more than 80 degrees Fahrenheit. We also calculated two baselines: the average of the highest 3 of the past 5 eligible days (eligible days are non-holiday weekdays), and the average of the highest 3 of the past 10 eligible days. We assumed a peak period of 2:00-6:00 pm. Figures 1 and 2 below show the load shapes on the event and the two baselines for customer #1 on each of the two hypothetical event days.

In Figure 1, on event day A, customer #1 uses more on the event day than the baseline averages. These are the types of load shapes we would expect to see from a customer who is not participating in an event. The event day usage is close to the baseline usage, but slightly higher perhaps due to the hot weather. In Figure 2, on event day B, customer #1 actually uses quite a bit less than either of the two baseline averages. In fact, the event day load shape looks like a demand response shape with a drop in usage during the event. This customer must have done something out of the ordinary on that day. Perhaps they went to the movies instead of turning up the AC, we will never know. What we do know is that the load drop on the hypothetical event day is purely a result of this customer's random variation in usage, not in response to an event. In this example, even though the customer did not actively participate in either event, customer #1 would receive a rebate for event B and a no rebate for event A.

Figure 1

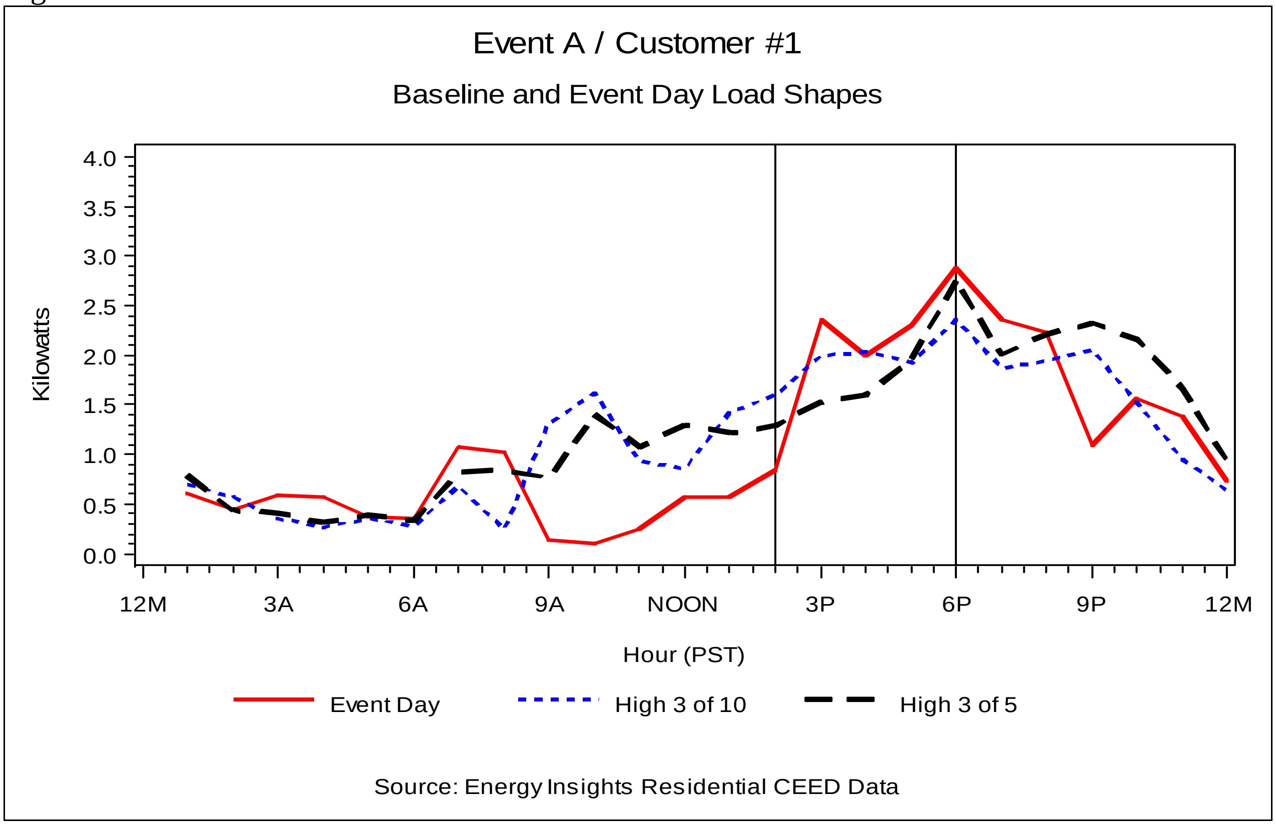

Figure 2

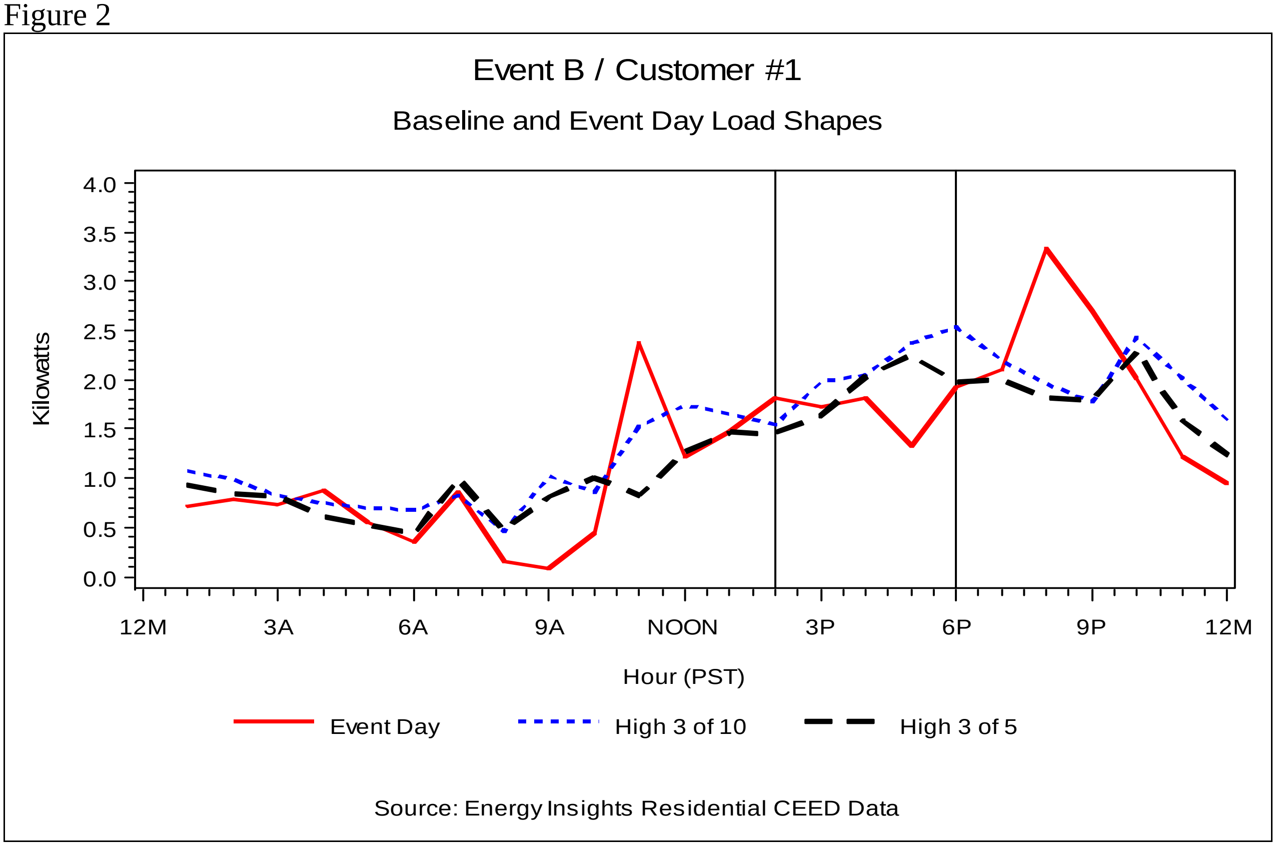

Template used with permission by IEPEC.

Figures 3 and 4 show load shapes for customer #2. On event day A, customer #2 had usage that was below the baseline. On this day the lower peak period usage should imply that the customer is responding to an event, but we know that they were not. On event day B customer #2 has usage that is very similar to the baseline average. Interestingly here, under the high 3 of 5 baseline the customer would not receive a rebate, but under the high 3 of 10 baseline the customer would receive a rebate, albeit a very small one.

The important thing to take away from these examples is that customers' activities on event days, whether the action is intentional peak-period conservation or simple coincidence, will earn PTR rebates. Because there is no increase in price to offset that rebate, and because corresponding random increases in load do not offset these rebates, the utility must absorb some level of unavoidable revenue loss. Another important point is that while a utility can do exactly what we have done here, and estimate the level of PTR payouts to non-participants, or analyze historical data from before a program goes into effect, there is no way to tell active participation from random variation once the PTR program begins.

Figure 3

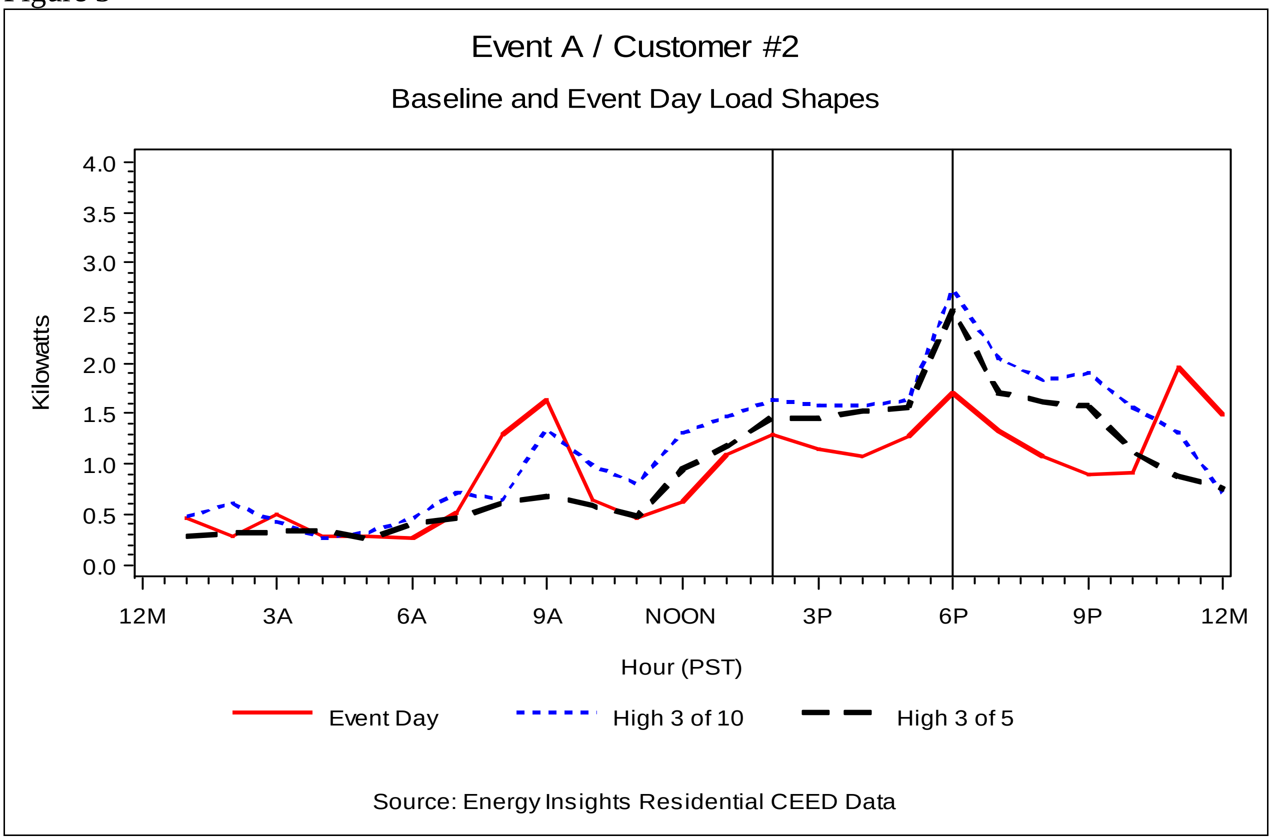

Figure 4

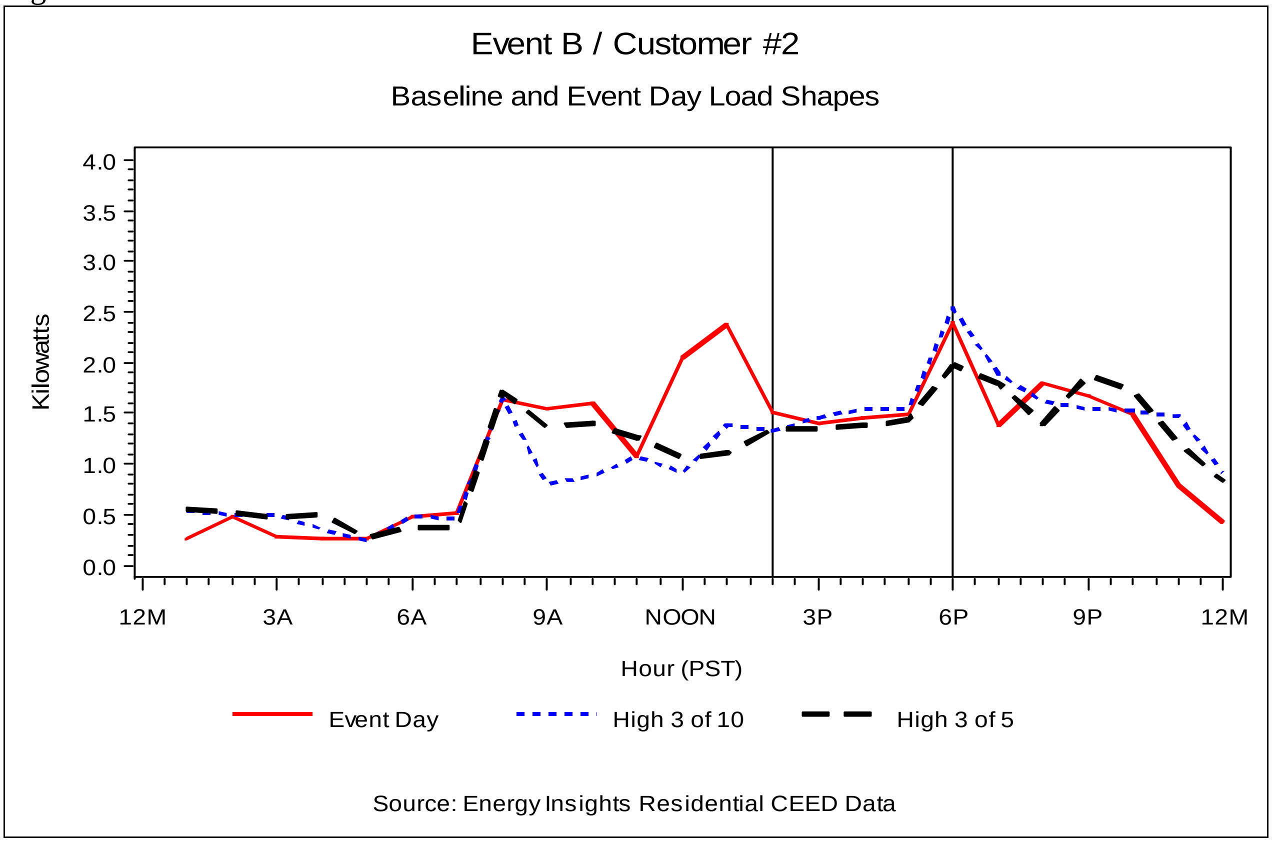

Choosing a Customer-Specific Baseline Method

Because the difference between the baseline usage and the event day usage determines the PTR rebate, the method a utility chooses for calculating the baseline must be carefully considered. Some of the key objectives to consider when developing a customer-specific baseline include:

- Accuracy: the baseline should estimate what the customer would have used on the event day as accurately as possible.
- Simplicity: the calculation should be as simple as is feasible while maintaining the desired level of accuracy. The utility will want the baseline to be easy to understand and have a low cost to calculate.
- Minimize payouts to non-participants

Several methods have been considered by utilities that plan to offer PTR rates. The method chosen by both SCE and SDG&E is the average peak period usage of the highest 3 of the previous 5 eligible days. Some others include highest 3 of the past 10 eligible days, highest 5 of the past 10 eligible days, and those methods with adjustments for event day temperature or load. Using an adjustment to the baseline to account for higher temperatures on an event day improves the accuracy of the baseline. Unfortunately, it sacrifices simplicity because it requires a customer-specific regression analysis, which has not been very popular among utilities or commissions. An adjustment based on the customer's load before the event can cause poor results if customers either pre-cool or reduce load before the start of the event.

The averaging methods are simple and easy to understand, but if events are called well (on the hottest days of the summer) they are all inherently biased. If an event is called on a very hot day, then most of the days preceding the event will be cooler than the event day. This implies that for most customers, their usage will be lower on any preceding day than it would have been on an event day under normal conditions. The more days included in the average (3 vs. 5) the lower the baseline will be. Because of this, a 3 of 10 baseline will tend to pay out a higher rebate than a 5 of 10 baseline. Conversely, the more days included in the pool of days from which the high load days are pulled, the higher the baseline will be. For example, a 3 of 10 baseline will always be the same or higher than a 3 of 5 baseline. Therefore, while any simple average will tend to underestimate usage to some extent, the type of average used will increase or decrease the rebate for the same customer on the same event. Choosing a higher baseline will make it easier to show reduction and will pay out a higher rebate than a lower baseline. Table 1 shows the rebates paid to customers #1 and #2 on event day A under both the highest 3 of 5 and the highest 3 of 10 baseline calculations, assuming a per kWh rebate of $0.60. Note that a negative demand reduction corresponds to an increase in usage.

Table 1

		Without Demand Response		***With Demand Response***	
Event A		**Highest 3 of 5**	**Highest 3 of 10**	**Highest 3 of 5**	**Highest 3 of 10**
Customer #1	kWh Reduction	-1.697	-1.225	-0.659	0.683
	PTR Rebate	**$0.00**	**$0.00**	**$0.00**	**$0.41**
Customer #2	kWh Reduction	1.88	2.331	2.918	3.369
	PTR Rebate	**$1.13**	**$1.40**	**$1.75**	**$2.02**

In Table 1, for the columns assuming some demand response, we assumed a reduction of 20% of on peak usage. It is important to keep in mind that an individual customer who chooses to completely curtail their AC load is likely to reduce much more than 20%, but to keep in line with the assumptions made by utilities 20% was used here.

For customer #1, neither baseline pays a rebate without demand response because that customer's usage increased when compared with either baseline. Interestingly, even after assuming a 20% reduction in usage, customer #1 still does not receive a rebate under the highest 3 of 5 baseline. Under the highest 3 of 10 baseline customer #1 only receives a rebate of $0.40. So in this case, the more lenient baseline might be preferred because it recognizes customer action even if the rebate is small. This example demonstrates an undesirable aspect of baseline calculations. Some customers who successfully reduce their usage on an event day may not receive the rebate they deserve because the baseline usage is too low relative to their actual event day use.

Customer #2 starts out with a sizable rebate of $1.13 for doing nothing under the highest 3 of 5 baseline. Using the highest 3 of 10 baseline increases that rebate by $0.27 to $1.40. Assuming a 20% reduction in usage the rebates increase to $1.75 and $2.02 respectively. Adding the demand response, which is the "real reduction", does not do much to increase the rebate. For this customer, the majority of the rebate is a result of random variation, and a smaller portion is a result of customer action. We will see this result again on a class level when examining the estimates of revenue loss made by both SCE and SDG&E. In this example it might be better to use the more stringent baseline to minimize the revenue loss.

These two customers illustrate an important point. Choosing the baseline that minimizes the revenue loss, highest 3 of 5, also minimizes the rebate for customers who choose to respond to the event. This is an important consideration for utilities that need to balance the incentive for participation with minimizing revenue loss.

Utility Experience

SCE and SDG&E plan to offer PTR as a default rate for all residential customers. They have addressed the issue of revenue loss in similar ways in their General Rate Case (GRC) Filings with the CPUC. It is important to note that California IOUs have a special mechanism called the Energy Resource Recovery Account (ERRA) to record and recover revenue deficiencies and surpluses. The account is allowed to fluctuate between plus or minus 5% of a utility's recorded generation revenues for the prior year. If either a deficiency or surplus exceeds 5% then an application is triggered to settle those costs. Otherwise an annual filing is used to deal with the balance in the ERRA.

SCE plans to default all residential customers to PTR as their AMI meters are installed. In their AMI business case filing they estimate total PTR payments made to customers without taking any action will be $27 million. The total estimated PTR payments after assuming a 20% reduction in peak-period usage across the board is estimated at $68 million. So payments to non-participants make up nearly 40% of the total estimated payout. SCE plans to take initial estimates of payments into account ($27 million) when setting residential rates to address any inter-rate group subsidization. Subsequent revenue deficits will be accounted for in the annual ERRA filing and most likely allocated back to the residential class in the form of higher rates.

SDG&E is also defaulting residential customers onto PTR as meters roll out. For SDG&E customers, rebates will be paid on a seasonal basis, with each event rebate independent from all others. SDG&E chose to pay rebates only once per year to reduce confusion that might be associated with multiple payments. They chose to keep the payment for each event independent so as not to penalize customers who do not reduce for some of the events, or choose to consume more on a single event day. In their 2008 GRC filing, SDG&E estimated the PTR rebate payments associated with no demand response to be between $10 and $12 million. After assuming 70% of customers will reduce peak period usage by 14% those payments increased by only $5 million. For SDG&E payments to non-participants are about two thirds of total PTR payments. This revenue loss will be rolled into an ERRA and eventually recouped through increased rates in the residential class.

While some utilities have a mechanism in place to recover the lost revenue attributable to PTR payments to non-participants, the fact that such a large portion of total payments is lost revenue is an important issue. We feel that utilities should consider taking steps to mitigate revenue losses on PTR rates, and perform the essential analysis to uncover what those losses might be. Utilities that have no way of recovering lost revenues through an automatic process should either try and adjust rates to cover the expected lost revenue, or consider forgoing PTR altogether.

Recommendations for Utilities Considering a PTR Rate

Know What You're Getting Into

PTR rates are quickly replacing CPP as the residential DR program of choice. The rates are carrot only, or win-win programs, and are easily approved by commissions, but while the program is good for customers, it poses risks for the utility. There is a significant potential for revenue loss associated with a PTR program and it is therefore important that any utility considering a PTR program fully understands any negative implications. The following steps can help utilities to assess the situation:

- The utility should use its residential load research sample to estimate the PTR payments being made to non-participants. To do this, choose hypothetical event days in historical data and calculate the payments to customers without making any assumptions about the level of demand response.

- Going one step further and repeating the calculation while assuming some level of DR will allow the utility to estimate the portion of total payments being made to non-participants.
- Consider performing this analysis for several baseline scenarios, keeping in mind that while setting a lower baseline will mitigate revenue loss it will also reduce the rebate for those customers who do respond on an event day.
- Choose a baseline calculation carefully. Remember that all simple averages will be biased if event days are called on the hottest days of the year.
- Consider a baseline that adjusts usage for weather on the event day. While they are more complicated to calculate and explain, they are more accurate.

Three Suggestions for Mitigating Revenue Loss

All PTR rates by design will be associated with some level of revenue loss for the utility. Whether the utility has a mechanism in place to recover that lost revenue and the magnitude of the loss will determine next steps in each case. If, like the California utilities already mentioned, the utility does automatically recover the lost revenue, then no mitigation is necessary. We have three suggestions that alter the PTR rate in such a way as to mitigate the revenue loss experienced by the utility. Two are simple, and do not change the basic nature of the rate, and one is a bit more radical.

Consider paying PTR payments on a net seasonal basis. If customers' response to PTR events were measured cumulatively across the season, then the random increases and decreases in usage over the baseline should equalize somewhat over the course of the season. This would reduce the payments for random variation. Some might argue that doing this could discourage customers from participating by allowing high usage on one event to erase savings from previous events. However, in principal, as long as a customer does not do anything out of the ordinary, and assuming that the baseline is a decent approximation of normal usage, a customer choosing not to participate in an event should not lose anything because their usage will be close to the baseline.

A second option would be to establish some minimum level of reduction below the baseline before a rebate is paid. Looking at Figure 4 above customer #2 showed usage that fell between the two baselines. These small deviations are usually a result of random variation or noise, not actions taken by the customer. Also, there are likely to be a lot of customers with these small deviations from the baseline. If a utility established some minimum reduction in usage they could eliminate many small payments to customers for what is essentially noise in the data. For example, if we conservatively assume that 20% of customers might have some small variation from baseline and that half of those deviations would be negative, for SCE that translates to roughly 400,000 customers. A $0.25 rebate to that large number of customers would be $100,000 in revenue loss per event, over a ten event summer that adds up to $1,000,000. While this does not eliminate the revenue loss, it does mitigate it somewhat.

Our final option would be to change the structure of the PTR rate from a carrot-only program to a carrot-and-stick program. A payment could be added so that any reduction below baseline would receive a credit and any increase above baseline would incur a payment. This would allow the rate to be designed in a revenue neutral fashion similar to a CPP rate. However, using a customer-specific baseline eliminates the problem of structural winners and losers. So those who could not reduce load would not pay more, as they would with a CPP rate – they would pay about the same as their standard rate. Those who use more on event days would see higher bills. This is similar in concept to a two-part real-time pricing (RTP) rate. It is important to note, however, that this approach would necessitate more careful consideration of the baseline method, as the inherent bias in the unadjusted average methods would have a more serious impact.