

ICC Docket No. 09-0263

Commonwealth Edison Company's Response to
MC Squared Energy Services and Constellation New Energy ("MC-CNE") Data Request
MC-CNE 1.01 - 1.13
Dated: July 2, 2009

Admi Heel
8/19/09
ICCA Cross
Ex 1-13 Cks

REQUEST NO. MC-CNE 1.04:

Reference: Customer Applications Plan, footnote 5.

The plan states that the "N" or number of customers in each research type "represents the number of participants required to yield a statistically valid estimate of behavioral response" and the footnote cites the confidence interval and confidence level.

- a. Please provide any workpapers that document the calculation of the "N"s
- b. Did ComEd consider confidence levels less than 99% (e.g. 95%) in order to minimize the research cost? If yes, why was 99% selected? If no, why did ComEd not consider levels less than 99%?

RESPONSE:

- a. The N=200 number referenced in the customer applications plan is the estimate for the number of subjects plus oversampling (N=50) to account for estimated subject mortality. This is based upon an article in the July 2009 edition of the *Electricity Journal* 2009 by Ahmad Faruqui, Ryan Hledik and Sanem Sergici entitled, "Piloting the Smart Grid" which is labeled as MC-CNE 1.04_Attach 1. In this paper, the authors make the case that a reasonable sampling strategy shall include, at minimum, an N of 150 subjects per cell. Any N=400 number was due to bifurcated applications (see Table 2, ComEd Ex. 4.0 at 9).
- b. No, given the workpaper referenced in (a), this question about confidence levels is no longer applicable. Due to an inadvertent error, the confidence levels shown in the customer applications plan reflect classical sampling design, which is NOT being used for the M&V evaluation. Instead, the M&V evaluation will be using the Bayesian method described in the above-referenced workpaper, from which the sample described in (a) is estimated.

OFFICIAL FILE

I.C.C. DOCKET NO. 09-0263
~~ICCA~~ Exhibit No. 1
 Witness _____
 Date 8/19/09 Reporter CC

Piloting the Smart Grid

To address the likely impact of the smart grid on customers, utilities, and society as a whole, it may be necessary to conduct a pilot. When should a pilot be conducted and how should it be conducted? What validity criteria should the pilot satisfy? Here are issues to consider.

Ahmad Faruqui, Ryan Hledik and Sanem Sergici

*Ahmad Faruqui is a Principal with
The Brattle Group, Inc.*

*Ryan Hledik is Associate with The
Brattle Group.*

*Sanem Sergici is Associate with
The Brattle Group.*

*The views expressed in this article are
strictly those of the authors and do
not necessarily state or reflect the
views of The Brattle Group, Inc., or
its clients.*

The transformative power of the smart grid is enormous. It is receiving much consideration from utilities and commissions across North America. Several members of the European Union, China, Japan, and other nations are also engaged in the same endeavor.

The smart grid has the potential for revolutionizing the way we produce and consume electricity but because it contains so many new elements; its core value proposition remains untested. The unanswered questions include:

- What new services will the smart grid provide customers?
- Do customers want these new services?

- Will they respond by changing their energy use patterns?

The answers to these questions will help policymakers in federal and state government determine whether the benefits of the smart grid will cover its costs.

It is widely understood that the new services enabled by the smart grid will include different rate designs that encourage curtailment of peak loads and make more efficient use of energy. Examples include dynamic pricing and inclining block rates.¹ These innovative rate designs will be enhanced by various automating technologies such as Energy Orbs, programmable communicating thermostats (PCTs), whole-building energy

management systems (Auto DR), and in-home displays (IHDs).

The smart grid will of course go beyond smart meters and rate design and enable renewable energy resources to be connected to the grid. This will allow optimal use of intermittent resources, such as wind, which often reach their peak generating capacity during off-peak hours. New off-peak loads, such as plug-in hybrid electric vehicles, which reduce overall energy consumption and improve the carbon footprint, will be energized by the smart grid.

To address the likely impact of the smart grid on customers, utilities, and society as a whole, it may be necessary to conduct a pilot. When should a pilot be conducted and how should it be conducted? To be useful, a pilot must yield credible results. This requires that the pilot satisfy various validity criteria. These issues form the focus of this article. We provide examples from several recent pilots that involved dynamic pricing, a key element of the smart grid. The concluding section discusses how a hypothetical company, Smart Power, should go about designing its own pilot.

I. Should a Pilot Be Conducted?²

Policymakers should consider implementing a pilot if there is much uncertainty in the cost-benefit analysis of proceeding with full-scale deployment.

A powerful method for resolving uncertainty is to assess the value of information that would be generated from a pilot. This point is best illustrated with a case study.

California suffered the worst energy crisis in its history in 2001. Most analysts attributed the crisis in part to the lack of demand response in the market design. When prices rose in wholesale markets, there was no incentive

for retail customers to lower demand. In the summer of 2002, the California Public Utilities Commission initiated proceedings on demand response, advanced metering, and dynamic pricing. Early in the proceedings, it became clear that the decision to deploy advanced metering was fraught with risk. The deployment would be costly and the benefits uncertain, as they depended on the customers' price elasticity of demand.

A preliminary cost-benefit analysis using price elasticities from the literature on time-of-use pricing (which ranged from -0.10 to -0.30) carried out for an

investor-owned utility showed that such deployment would provide gross benefits, ranging from \$561 million to \$2,637 million. The cost of advanced metering infrastructure (AMI), net of operational benefits, was estimated to be \$1,080 million. This suggested that the net benefits would range from a loss of \$519 million to a gain of \$1,557 million. In other words, the range of benefits, at some \$2 billion, is wide.

Classical statistics would yield sample sizes (using formulas such as those shown in Figure 1) but would not provide insights about the value of information that would be gleaned from the pilot. Thus, a decision was made to pursue a Bayesian approach to determining the optimal sample size of the pilot. *If the size was zero, it would mean the pilot would provide information of no value and should not be pursued.*³

Given the wide variation in climatic conditions in the state, a decision was made to divide the state into four climate zones and estimate the potential diversity in customer response patterns. In one of the climate zones, a critical-peak pricing rate was found to have an expected net benefit of \$331 million using prior information on price elasticities. However, there was a 12 percent probability that the program would generate negative net benefits, as shown in Figure 2.

Thus, there was a reasonable probability that the state would make the wrong decision in the absence of better information. It

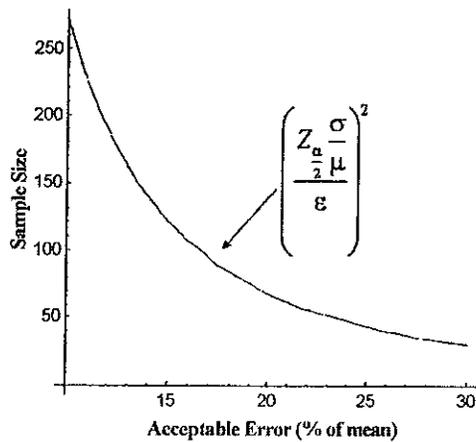


Figure 1: Classical Sampling Design

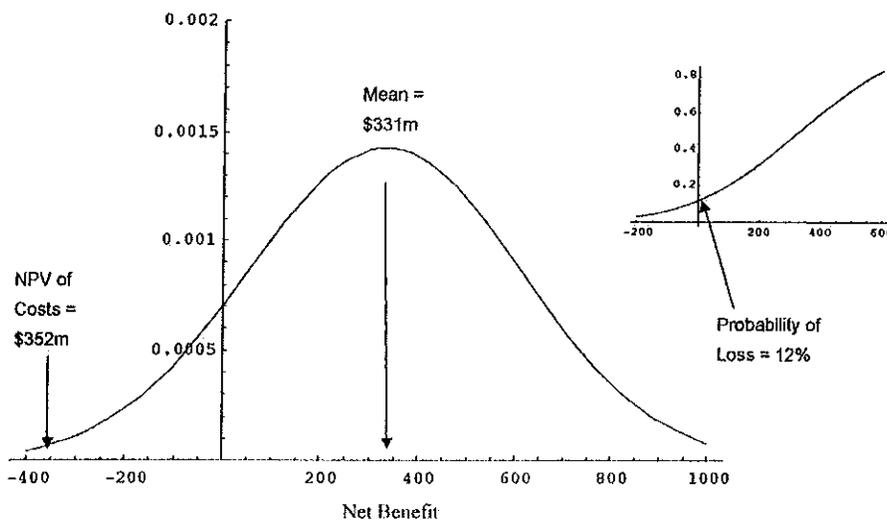


Figure 2: Pre-Pilot Distribution of Net Benefits

could proceed to implement dynamic pricing when it was not warranted, thereby burdening ratepayers with a loss of around a billion dollars. Or it could chose to stay with the status quo, thereby denying Californians the benefits of dynamic pricing and burdening them with higher power costs.

A properly drawn sample should improve the probability of making a correct decision on full-scale implementation of dynamic pricing. This involves three major steps:

- Estimating the net benefits of implementing dynamic pricing for each of the treatments that look promising based on *a priori* information about price elasticities and other aspects of customer behavior.
- Estimating the costs of implementing each treatment during the sampling phase of the study.
- Drawing the sample to maximize the probability of making the right decision, taking into account the tradeoff between value of information and cost of sampling.

The estimation of net benefits involves a computation of benefits and costs, usually as discounted present values over a planning horizon of 15 years.⁴ Benefits are estimated using the following equation:

Benefits

$$= (\text{Existing usage per customer} \times \text{Percent change in price} \times \text{price elasticity}) \times \text{Number of participants}$$

Costs are estimated using the following equation:

$$\text{Costs} = \text{Unit cost per participant} \times \text{Number of participants}$$

Both calculations involve several variables that cannot be predicted with certainty, and are best modeled in probabilistic terms. Monte Carlo simulation was used to develop the appropriate probability distributions.

With the Bayesian approach, the following sampling outcomes are possible. If Treatment A is likely to generate greater net benefits compared to Treatment B, but there is significant uncertainty in that result, the Bayesian approach would recommend drawing a larger sample than if there is no uncertainty in the result. For instance, if A will always be better than B, then sampling does not impact the final policy decision and contains very little useful information. The Bayesian approach explicitly factors the value of information into the determination of the optimal

sample size for each treatment cell. It differs from the classical statistics approach where value of information does not play any explicit role in determining the sample size. The two approaches would give similar results if the prior information on net benefits were diffuse or uncertain. The more sharply focused the prior information, the more the two sampling approaches will differ.

The Bayesian process was implemented using information from the preliminary analysis of net benefits for a single utility, scaled up to reflect statewide conditions, and information on the cost of sampling various treatments as well as the cost of full-scale implementation. The net gain from sampling as a function of sample size is shown in Figure 3 for Zone 1 with the critical peak pricing (CPP) rate. The curve rises steeply until a sample size of 50 is reached, and then increases at a decreasing rate. The maximum is reached at a sample size of 179, which would yield a net gain of \$15.2 million. In other words, a sample size of 179 would maximize the net benefits of information being generated by the sample. Once the optimal sample size is reached, the net gain curve flattens out, with a small negative slope. The flat shape of the optimal sampling curve means that one can factor in non-economic objectives such as equity and equal coverage without sacrificing economic value in the process. For logistical and budgetary reasons, the pilot

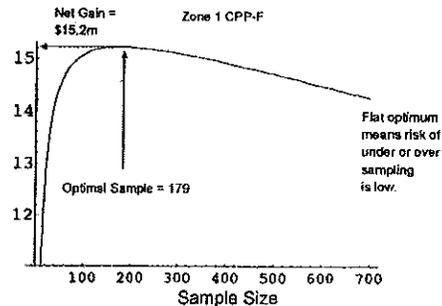


Figure 3: Optimal Bayesian Sampling

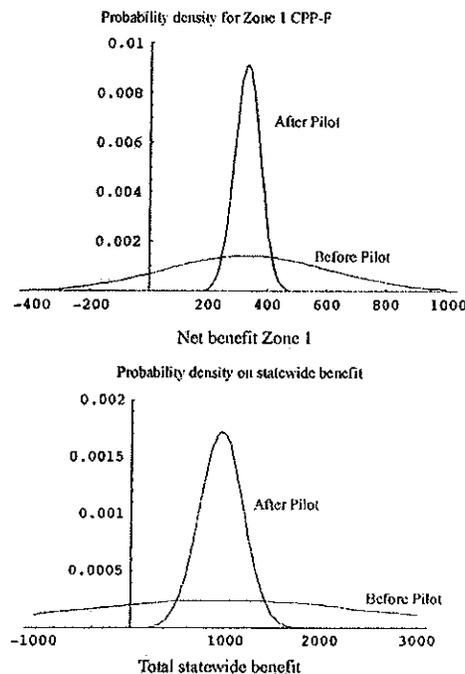


Figure 4: Narrowing of Uncertainty Due to Sampling

proposed using a smaller sample size of 119, which would sacrifice a net gain of only \$0.1 million.

In aggregate terms, the proposed sample design for the pilot was estimated to have a net gain of \$225 million. The primary benefit of sampling is that it will narrow the prior probability distribution of net benefits. This effect is shown in Figure 4, where the top panel shows the effects for Zone 1 and the bottom panel shows the effects for the state as a whole. A decision

was made to pursue what came to be known later as the Statewide Pricing Pilot (SPP). The design is described later in the article.

The SPP sample was ultimately based on a combination of factors, including the results of the Bayesian approach, the interests and issues raised by the stakeholders who voiced their views through a working group process, and practical considerations about timing and budget.

II. Ensuring Validity

To be credible and useful to policymakers, pilots need to have both internal and external validity. Internal validity means that a cause and effect relationship can be established between the various treatments being tested and the variables of interest such as peak demand and overall energy consumption. The effect of all variables needs to be purged. External validity means that the pilot results can be extrapolated to the population of interest. Both require careful design and it is generally easier to ensure internal validity than to ensure external validity.

To aid in ensuring internal validity, we can draw upon three decades of experimentation with new rate designs and technologies. The "gold standard" of pilot design stipulates that every treatment (e.g., rate design or technology associated with the smart grid) that is being tested should also have a control associated with it so that a scientifically valid "but for" world can be constructed from which deviations can be successfully measured. Without this, pilots do not have internal validity. In other words, cause-effect relationships cannot be inferred with any precision and any conclusions derived from the pilot may be subject to the charge that they simply measure spurious correlation. It is also likely that genuine cause-effect relationships (e.g., higher prices lead to lower usage by X percent)

may not be measured accurately because other factors such as a changing economy or weather may obscure the true relationship.

The best way to create a control environment is to select a matching group of customers who can serve as a proxy for the behavior of the treatment group customers. In addition, to further anchor the measurements, it is best to observe the treatment group before they are placed on the

treatments. In other words, have pre-treatment data on both the control and treatment groups as well as treatment-period data on both groups of customers. This leads to a balanced sample design. The same logic applies to test distribution automation; the test will be most definitive if we collect data from both the automated feeders, and from other, similar feeders that are not automated and will serve as "control" measures.

In the past, pilots have been carried out without matching control groups and sometimes with no control groups at all. Others have been conducted with control groups but no

pre-treatment measurements. All such inadequacies impair the internal validity of the pilot in varying degrees. Without a control group in the design, it is impossible to control for non-treatment variables that change between the pre-treatment and treatment periods (such as the weather, the economy, or general changes in attitudes toward energy use brought about by other exogenous factors). Without pre-treatment data, it is difficult to know if the treatment and control groups were balanced or unbalanced before the treatment was introduced. If systematic pre-treatment differences exist, they suggest that there is a self-selection bias in the sample that needs to be dealt with.

These are the general principles of pilot design to ensure internal validity of results. As with most things in the real world, they serve as guideposts and not mandates. Utilities will need to temper these principles in actual execution given their time and resource constraints.

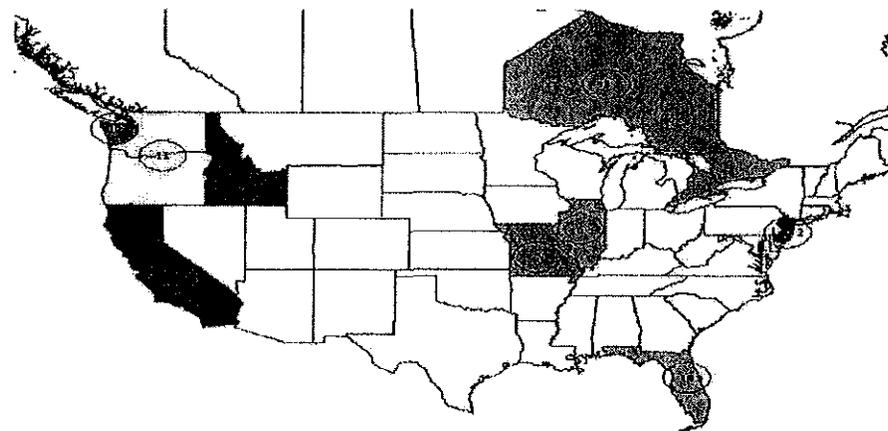
Pilots must also have external validity so that their conclusions are transferable to a real world setting. In the case of the time-differentiated rates, it will be useful to know if such rates will ultimately be offered on a universal basis, a default rate basis with opt-out provisions, or an opt-in basis. The sampling strategy for the pilot will vary across these three scenarios. For example, if universal deployment is contemplated, then both the control and treatment groups

should be chosen randomly. On the other hand, if an opt-in deployment is envisioned, then opt-in sampling would be appropriate for both groups. In addition, a random control group may also be used to contrast the results with the general population.

III. Summary of Existing Dynamic Pricing Pilots

Recent smart grid pilots have examined the impacts of the various forms of dynamic pricing rate designs and have spanned a number of customer classes and utilities. Some of the pilots have also measured the impacts of enabling technologies (such as PCTs) on customer response. To develop a better understanding of the findings of these pilots and to guide a utility's pilot design, it is useful to briefly survey the designs and some of the results from the major dynamic pricing pilots.⁵ Several of these dynamic pricing pilots are shown in Figure 5.

The following are particularly valuable in illustrating the potential impacts of the broad



- 1- PSE&G Pilot Program
- 2- CPU Pilot Program
- 3- Ontario Energy Board Smart Price Pilot
- 4- Anaheim Critical Peak Pricing Experiment
- 5- Idaho Residential Pilot Program
- 6- Energy-Smart Pricing Plan
- 7- Anterion UE Residential TOU Pilot
- 8- ADRS Pilot
- 9- Statewide Pricing Pilot
- 10- The Gulf Power Select Program
- 11- Olympic Peninsula Project
- 12- PSE TOU Program

Figure 5: Recent North American Dynamic Pricing Pilots

range of pricing and technology options that could be offered:

- California Statewide Pricing Pilot (SPP)
 - California Automated Demand Response System Pilot (ADRS)
 - California (Anaheim) Peak Time Rebate Experiment
 - Illinois's (Chicago) Community Energy Cooperative's (CEC's) Energy-Smart Pricing Plan
 - Ontario Energy Board (OEB) Smart Price Pilot in Ottawa
- Each of these pilots tested different variations of dynamic

rate designs and technologies, providing insights into what characteristics are more and less likely to be effective in achieving significant demand response. The design and conclusions of each of these four pilots are described below, followed by a summary of the findings (Table 1).

A. The California Statewide Pricing Pilot

California's three investor-owned utilities, together with the state's two regulatory commissions, conducted the

Table 1: Summary of Recent Dynamic Pricing Pilots

Pilot name	Location	Utility	Timeframe	Customer Classes	Rate Types
California SPP	California (Statewide)	PG&E, SCE, SDG&E	2003-2004	Residential, C&I	CPP, TOU, enabling technology
California ADRS	California (Statewide)	PG&E, SCE, SDG&E	2004-2005	Residential, C&I	CPP, enabling technology
Anaheim CPP Experiment	Anaheim, California	City of Anaheim Public Utilities	Summer 2005	Residential	PTR
CEC's ESPP	Chicago, Illinois	Commonwealth Edison	2003-Present	Residential	RTP
OEB's Smart Price Pilot	Ontario, Canada	Hydro Ottawa	2006-2007	Residential	CPP, TOU, PTR, enabling technology

multi-million-dollar Statewide Pricing Pilot (SPP) to test the impacts of time-of-use (TOU) and CPP pricing. The experiment ran from July 2003 to December 2004 and included about 2,500 participants consisting of residential and small-to-medium commercial and industrial (C&I) customers. The pilot was conducted across a wide range of climate zones, from the mild San Francisco Bay area to the very hot and dry Central Valley. The SPP was, and still is, the largest pilot of its kind.

The SPP tested two of the most common types of time-varying rate structures:

- TOU-only rate, where the peak price was twice the value of the off-peak price.
- CPP rate, where the peak price during the critical days was roughly five times greater than the off-peak price. Variations of the CPP were tested, in which the critical peak period was allowed to have a variable length.

The effects of day-ahead and day-of notification were also tested.

A high-level summary of some of the relevant findings of the pilot are as follows (Figure 6):

- Peak reductions were significant under both the CPP rate and the TOU rate.
- Customers produced greater peak reductions under the CPP than the TOU.
- Residential customers produced the greatest peak reductions, followed by medium-sized C&I customers.

B. California Automated Demand Response System Pilot

Related to the SPP was the California Automated Demand Response System (ADRS) pilot. It was initiated in 2004 and extended through 2005. ADRS operated under a critical peak pricing tariff that was supported with a residential-scale,

automated demand response technology. Participants in the pilot installed the GoodWatts system, an advanced home climate control system that allowed users to Web-program their preferences for the control of home appliances. Under the CPP tariff, prices were higher during the peak period (2 p.m. to 7 p.m. on weekdays). All other hours, weekends, and holidays were subject to the base rate. When the "super peak events" were called, peak price was three times higher than the regular peak price.

Results from the pilot showed that participants achieved substantial load reductions, through the use of the GoodWatts system. These reductions greatly exceeded reductions from customers who were exposed to the dynamic rates, but were not equipped with enabling technology. Ultimately, the pilot confirmed that enabling technologies are a significant driver of peak reductions.

Figure 7 summarizes the impacts of enabling technology from the SPP and the ADRS.

C. The Community Energy Cooperative's Energy-Smart Pricing Plan

The Community Energy Cooperative's (CEC) Energy-Smart Pricing Plan (ESPP), a residential real-time pricing (RTP) program, started in Illinois in 2003. The ESPP initially included 750 participants and expanded to nearly 1,500 customers by 2005. It is the only residential RTP

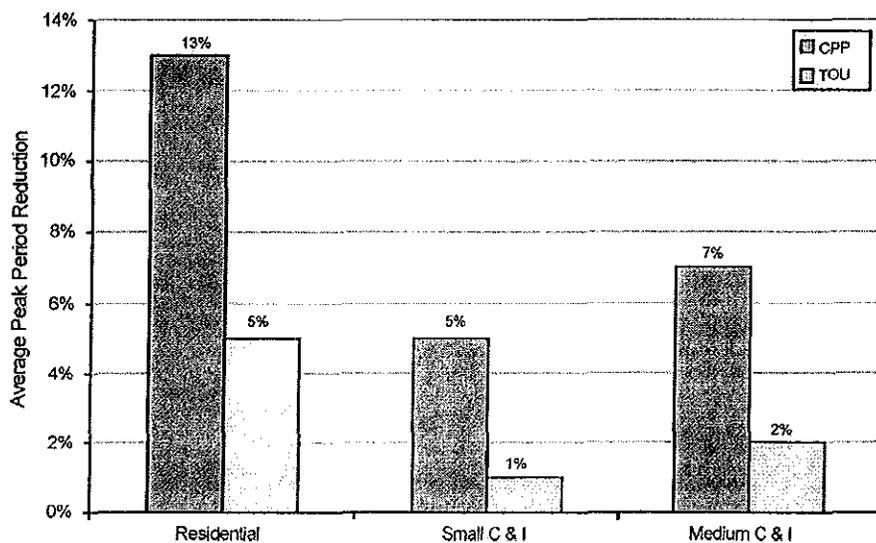


Figure 6: Summary of Impacts from California SPP

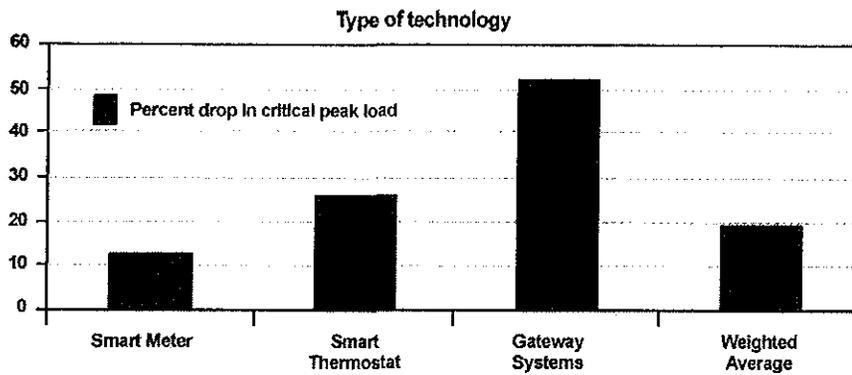


Figure 7: Impacts of Enabling Technologies from the California Pilots

program that has been tested at any scale. The ESPP focused on low technology and tested the hypothesis that major benefits may result from RTP without expensive technology adoption. The ESPP design included:

- Day-ahead announcement of the hourly electricity prices for the next day (*i.e.*, customers were charged the day-ahead hourly prices).
- Notification of “high-price days” via phone or email when the price of electricity was going to be over \$0.10 per kWh.
- A price limit hedge of \$0.50 per kWh for participants, meaning that the maximum hourly price was set at \$0.50 per kWh during their participation in the program.
- Energy usage education for participants.

Overall, the pilot found that customers were more price-responsive during the evening hours (4 p.m. to midnight) and on high-price days when they received additional notification. In response to time periods with the highest prices, customers reduced consumption by as much as 15 percent. During the summer

months, there was a slight energy conservation effect of 4 percent.

An important finding of the study was that hour-to-hour variations in the electricity price were not driving changes in customers’ consumption patterns. Rather, customers responded to blocks of time that were, on average, higher- or lower-priced times of day. Similarly, customers responded to notification of high-priced days, rather than responding to a spike in a single hourly price. This finding supports a conclusion that the complications associated with offering an hourly price may not be necessary, and a similar result could be achieved through a simpler time-period-based rate, such as a CPP or TOU rate.

D. Anaheim Critical Peak Pricing Experiment

The City of Anaheim Public Utilities (APU) conducted a residential dynamic pricing pilot between June 2005 and October 2005. A total of 123 customers participated in the experiment: 52 in the control group and 71 in the

treatment group. While the pilot’s title is somewhat of a misnomer, it actually tested the impacts of a relatively new rate design known as the peak time rebate (PTR). The PTR rewarded participants with a rebate of \$0.35 for each kWh reduction below their estimated baseline consumption (*i.e.*, their predicted consumption level in the absence of the rebate). During all other hours, their rate remained unchanged from the existing tariff.

The results of the experiment showed that customers on the PTR reduced their peak consumption by an average of 12 percent. On days with high temperatures, the peak reduction was even greater. However, it should be noted that the sample size for this experiment was very small. While the peak impacts appear to be comparable to those under a CPP rate, that conclusion cannot be drawn with the same level of certainty with which the CPP impacts were estimated. Several experimental pilots to more precisely compare the impacts of CPP and PTR rates are currently under development.

E. Ontario Energy Board Smart Price Pilot

The Ontario Energy Board operated the residential Ontario Energy Board Smart Price Pilot (OSPP) between August 2006 and March 2007. The OSPP used a sample of Hydro Ottawa residential customers and tested the impacts from three different price structures: a pure TOU rate, a CPP layered on a TOU, and a PTR

layered on a TOU. A total of 373 customers participated in the pilot: 124 in TOU-only, 124 in TOU-CPP, and 125 in TOU-PTR. The control group had 125 participants with installed smart meters, but who continued to pay non-TOU rates.

The pilot found large average peak reductions under all three of the rate designs, with the CPP-TOU impact being the highest, followed by the PTR-TOU, and then the pure TOU. However, standard errors surrounding these estimates were very large and it is difficult to draw meaningful conclusions from the results. What is more interesting to consider is the impact of real-time feedback monitors, which were also tested through the experiment.

Real-time feedback monitors are installed in the home and can display information about current electricity consumption, the price of electricity, and the cumulative amount that has been spent on electricity. A specific type of real-time feedback monitor is the PowerCost Monitor by Blue Line Innovations. This device also allows customers to view an estimate of the carbon-dioxide emissions that are being produced as a result of their electricity consumption. The device can be self-installed on the electric meter by the customer. Information is then wirelessly transmitted to the monitor, which can be installed anywhere in the house. The PowerCost Monitor can be purchased and installed by individual customers for under \$150.

The effectiveness of the PowerCost Monitor was tested in the OSPP. Five hundred of Hydro Ottawa's customers were equipped with the PowerCost Monitor and data on the customers' electricity usage was collected over a period of two and half years. The results of the pilot suggested that, on average, customers with the device installed reduced their electricity consumption by 6.5 percent (at a

high level of statistical significance). This reduction was sustained over the duration of the pilot.

Within the sample group, customers with non-electric space heating were found to reduce consumption at a much higher level than those with electric space heating. This might suggest that for the customers with electric space heating, the feedback from the electric heating load would need to be separated from load at other end uses in order to effectively encourage conservation for these customers. Overall, customers expressed a high level of satisfaction with the device,

with over 60 percent indicating that the device was useful in helping to conserve energy. These results were achieved in the absence of any accompanying incentives or price schemes.

IV. Pilot Design Principles

To provide a frame of reference regarding the principles of pilot design, we consider the case of a hypothetical utility called Smart Power that wishes to pursue a smart grid pilot. As it contemplates the development of a pilot, it would find it useful to begin by recalling the well-established principles of experimental design. The salient ones are summarized below:

1. In order to measure the impact of the new rate designs (called "treatments" in the literature on social experimentation), the design should: (a) control for the effect of other factors such as weather and the economy, and (b) be capable of inferring what the customers on the treatments would have done in the absence of the treatments. Otherwise a valid cause-effect cannot be established between treatment and result.

2. This is best accomplished in two ways: (a) by including a control group in the design, comprised of customers who are similar in all other respects to customers in the treatment group, and (b) by measuring the load profiles of customers in both the control and treatment groups before the new rates

(or "treatments") are initiated and during the time the treatments are initiated.

3. Sufficient numbers of customers should be recruited to fill the control and treatment groups. This often means that at least 100 customers should be in each cell. Too few customers in the cells will result in the inability to detect the effect of the treatment through statistical means (*i.e.*, the signal-to-noise ratio will be poor).

4. Customers should be randomly selected and assigned, to the extent practical, to the treatment and control groups. This will allow valid inferences to be drawn about the behavior of the target population.

5. Data should be collected not only on customer load profiles but also on their socio-demographic characteristics and their attitudes toward energy use.

6. Multiple treatments should be used to construct a model of customer price response (commonly called a "demand model") and to derive price elasticities; if only a single treatment is included, then the experiment will yield specific impact estimates for that single treatment.

7. Customers should be encouraged to stay in the pilot for as long as possible.

8. If any payments need to be made to customers to ensure that they stay through the end of the pilot, these payments should be (a) made only toward the end of the pilot, (b) unrelated to the level of their monthly usage, or (c) tied to the amount of bill savings generated by their actions.

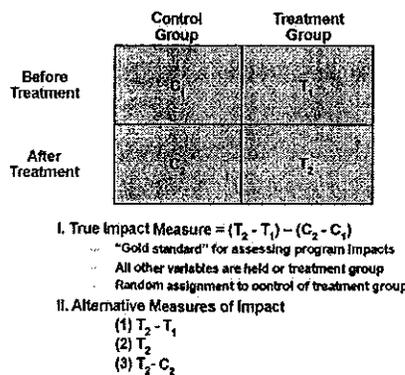


Figure 8: The Gold Standard for Experimental Design

If all these design principles are followed, then Smart Power's experiment will yield the best possible measurements. A design that conforms to all of these principles is often referred to as the "gold standard" against which other designs can be benchmarked. The "gold standard," which is recommended for Smart Power's smart grid pilot, is illustrated in Figure 8. To see why other designs may yield inferior results, imagine the following deviations from its precepts:

- A design without a control group. In this case, only before/after measurements can be carried out on the treatment group. Such a result would be subject to criticism for not having netted out the effects of other factors (such as the weather, the economy, or an energy crisis) that may have changed since the experiment began.

- There is a control group but no pre-treatment measurements are available. Such a design will not allow an assessment of whether the control and treatment groups are well-matched. The experiment may be regarded by

critics as suffering from a self-selection bias. This is a fairly common problem that can be easily rectified by building in time for pre-treatment data collection.

- A simple design that only includes a treatment group after a treatment has been initiated. This will yield the poorest results since it does not have a control group, nor has it measured the treatment group before the treatments were initiated.

- A design that includes a non-matching control group. This is called a quasi-experiment, and while it is better than having no control group at all, it is far from ideal. This problem is fairly common in dynamic pricing experiments. It can be avoided by using standard sampling techniques.

Through our extensive experience in dynamic pricing pilot design we have identified some key characteristics that are common among well-designed dynamic rates. Since well-designed rates produced significant peak reductions and high levels of customer acceptance, we propose

to incorporate these characteristics into the rates that are developed for the Smart Power pilot.

Revenue neutrality: Each dynamic pricing rate option should be revenue neutral. In other words, in the absence of any load shifting, the rate should not lead to an over- or under-collection of utility costs on a system-wide basis. At the customer level, some bills would increase and others would decrease, but the average customer's bill should not change.

Short peak period: The on-peak or critical peak periods should be kept as short as possible, while still reasonably spanning the period during which the system peak occurs. A shorter peak period makes it easier for customers to shift load to the off-peak period when demand reductions are not as critical. For example, a four-hour peak period from 2 p.m. to 6 p.m. would reasonably allow customers to shift the use of some of their appliances, such as dishwashers or clothes dryers, before or after the period's duration. A long peak period would be less likely to induce customer response, as they would need to shift usage to the early morning or late night hours, requiring more significant behavioral changes. Many voluntary TOU rates in the industry feature very long peak periods, and it is no surprise that very few customers are enrolled on such rates.

Strong price signal: The rate should convey a strong price signal to customers. In other

words, the differential between the peak and off-peak prices should be large. This large differential gives the customer a significant incentive to reduce consumption when the price is high, and produces the opportunity for greater bill savings by creating a large off-peak discount. The customer needs to notice that there is a substantial difference in prices during these two periods. A small

differential sends a weak price signal to customers and could be too insignificant for them to care about changing their consumption patterns. This problem was encountered by PSE in its pilot with TOU rates. The customers claimed to have reduced their peak usage by large percentage amounts but this often translated into trivial bill savings.

Rates should reflect system costs: While a significant price signal is important, the rate should still reflect the cost of providing power to the customer. The peak period rate typically reflects both the higher average variable cost of generation, as well as the cost of

capacity necessary to meet peak demands. The off-peak rate is a reflection of the lower average cost of meeting customer demand during hours with lower loads. This is what drives the differential between the peak and off-peak rates.

Opportunity for significant bill savings: Customers are less likely to voluntarily enroll in the dynamic rates if they do not see an opportunity for significant bill savings. Similarly, once customers are on the rates, they are more likely to produce large peak reductions if doing so allows them to save money. To create such a rate, the off-peak discount should be large (or in the case of the PTR, the peak rebate should be large).

Simplicity: Dynamic pricing rates should be easy for the customer to understand. If the customer does not understand how the pricing works, or is overburdened with information, then he or she will not be able to appropriately respond to the price signals and shift load. The residential RTP is an example of a rate design that has been shown to provide more information than customers need to provide demand response.

V. Developing an Implementation Plan

The smart pilot will allow Smart Power to quantify the impact of dynamic rates, incentive plans, and technologies on its system and its customers. To illustrate the steps involved in

designing and implementing such a pilot, we have developed a strawman pilot design. This is a high-level description. The final pilot design will need to be developed in detail through a number of working sessions with Smart Power staff, the public service commission staff, and any other stakeholders that are involved in the public process who usually attend the development and execution of such projects.

However, this strawman design is useful for understanding the pilot design process, highlighting some of the important decisions to be made in designing the pilot, and identifying some of the questions that could potentially be answered through the pilot's implementation.

Some of the key features of our strawman pilot design include:

- Testing a variety of dynamic rate designs.
- Including high and low variations of each rate type.
- Measuring the impact of enabling technologies.
- Measuring the difference in response from single-family and multi-family homes.

Our strawman pilot design is summarized in the remainder of this section.

A. The rate designs

There are a wide variety of dynamic rates that meet the rate structure design parameters that Smart Power has identified for the pilot. A description of these rate designs is provided below.

Time-of-Use (TOU): A static TOU rate divides the day into time periods and provides a schedule of rates for each period. For example, a peak period might be defined as the period from 12 p.m. to 6 p.m. on weekdays, with the remaining hours being off-peak. The rate would be higher during the peak period and lower during the off-peak, mirroring the variation in the cost of supply. There would be no

uncertainty as to what the rates would be and when they would be incurred.

Critical Peak Pricing (CPP): Under a CPP rate, participating customers pay higher peak-period prices than they would on their otherwise applicable tariff during peak hours on the few days when wholesale prices are the highest. In return, the customers pay a lower off-peak price that more accurately reflects lower off-peak energy supply costs for the duration of the season (or year). Thus, the CPP rate attempts to convey the true cost of power generation to electricity customers, and

provides them with a price signal that more accurately reflects energy costs as well as the opportunity to minimize their electricity bills. This rate form is particularly effective when elevated supply costs are limited to only a few (under 100) hours of the year, and their onset is quite predictable.

Peak Time Rebate (PTR): If a CPP tariff cannot be rolled out because of political or regulatory constraints, some parties have suggested the deployment of PTR. Instead of charging a higher rate during critical events, participants have the opportunity to buy through at the existing rate; however, they have a significant incentive for reducing critical-peak usage in the form of cash rebate that is expressed in cents per kWh of load reduction during the critical period. This, of course, requires the establishment of a baseline load from which the reductions can be computed.

CPP-Variable (CPP-V): CPP-V is similar to the CPP rate, with the exception that the duration of the peak period is not fixed. The event notification is generally provided to participants on a day-ahead basis at the same time that they are notified of the upcoming critical event. This provides utilities and ISOs with the flexibility to respond to emergencies and high priced periods of varying lengths occurring at different times of the day.

Variable CPP (VPP): It is also possible to vary the critical peak price, rather than locking it in at a pre-specified level. CPP rates with

this characteristic are called VPP rates. They provide a price signal to customers that more accurately reflects contemporaneous system conditions and marginal costs.

Real-Time Pricing (RTP):

Participants in RTP programs pay for energy at a rate that is linked to the hourly market price for electricity. Depending on their size, participants are typically made aware of the hourly prices on either a day-ahead or hour-ahead basis. Typically, only the largest customers —above 1 MW of load — face hour-ahead prices. These programs post prices that most accurately reflect the cost of producing electricity during each hour of the day, and thus provide the best price signals to customers, giving them the incentive to reduce consumption at the most expensive times. Over 70 utilities have offered RTP either in a pilot or as a permanent program.

In our strawman pilot design, we are proposing that a TOU, CPP, and PTR be tested. RTP, CPP-V, and VPP were excluded from the list of rates in order to reduce the number of customers that would need to be recruited to the pilot. Generally, these rates have been shown to provide more information than customers can use when making their consumption decision. In other words, residential customers tend to respond to prices in time blocks rather than on an hour-by-hour basis. However, it may be desirable to include one or more of these rates in the final pilot design.

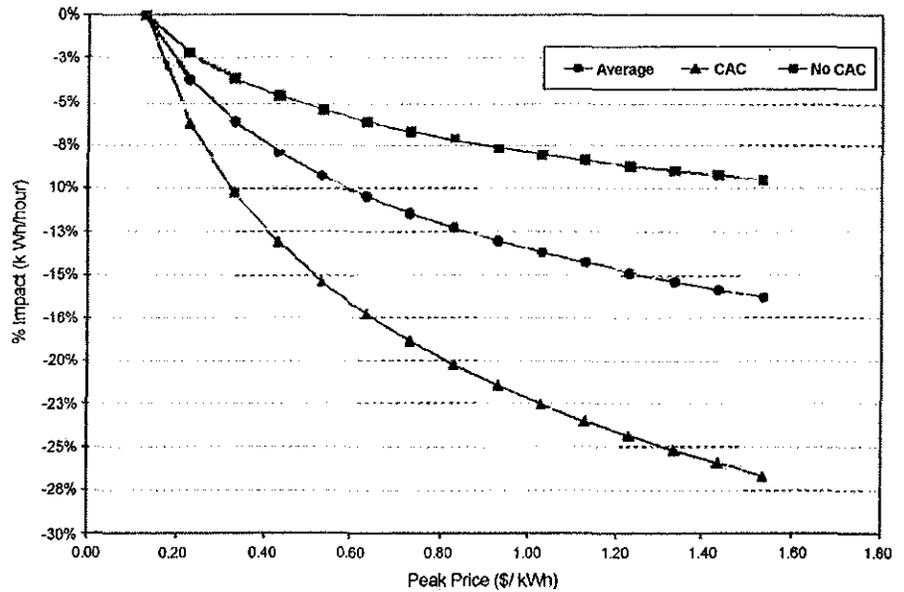


Figure 9: Residential Customer Price Response Curves on Critical Days

B. "High" and "low" rates

For econometric purposes, multiple rate configurations will be chosen within each rate design. This will allow the estimation of customer demand curves and price elasticities, as is common in such experimental pilots. Specifically, there will be a "high" and a "low" scenario for each rate design. The "high" rate will have a higher on-peak price and, as a result of the revenue neutrality constraint, a lower off-peak price. The "low" rate will have a lower on-peak price and higher off-peak price.

The purpose of these "high" and "low" rate scenarios is to produce sufficient variation in data for econometric analysis. With an estimation of customer response at more than one price level, it is possible to estimate the curvature of the customer price elasticity, rather than simply assuming that customer response increases linearly with an increase

in price at all price levels. In fact, the most thorough pricing experiments have suggested that peak reduction increases as the peak rate increases, but the reductions grow at a rate that gets incrementally smaller. This is illustrated in Figure 9. The design of the rates is an important aspect of the pilot. To ensure that the rates are designed appropriately, their impacts will be simulated using the Price Impact Simulation Model (PRISM). Simulation of the rate impacts is a critical step in the rate design process, as it ensures that the high and low rates are sufficiently different in order to produce a wide range of customer response. This is necessary from a statistical analysis perspective to estimate a robust price response curve.

C. Enabling technology

Enabling technologies such as programmable communicating thermostats (PCTs) and A/C



switches automate demand response and increase the effectiveness of dynamic prices in promoting demand reduction. Evidence from several pricing pilots reviewed in a recent survey reveals that while CPP tariffs lead to a drop in peak demand of 13 to 20 percent, pairing pricing tariffs with enabling technologies lead to a peak demand reduction in the range of 27 to 44 percent.⁶

Figure 10 demonstrates the impact of enabling technologies on critical peak demand reduction in a recent pilot program.

Given this evidence, it is important to identify the role of enabling technologies on price response by equipping certain customers with these technologies and testing the incremental impact. This is possible by having another group of customers who are also subject to the dynamic rates but do not have enabling technologies. Any pricing pilot design aimed at testing the incremental impact of enabling technologies should have at least two treatment samples: one with enabling technologies and the other without.

Single-Family vs. Multi-Family

Socio-demographic characteristics of customers may impact their price responsiveness. The customers' type of residence is one such characteristic that can potentially impact their sensitivity to dynamic rate designs. Customers who live in single-family homes can be expected to have more control over their electricity consumption patterns, and therefore have more

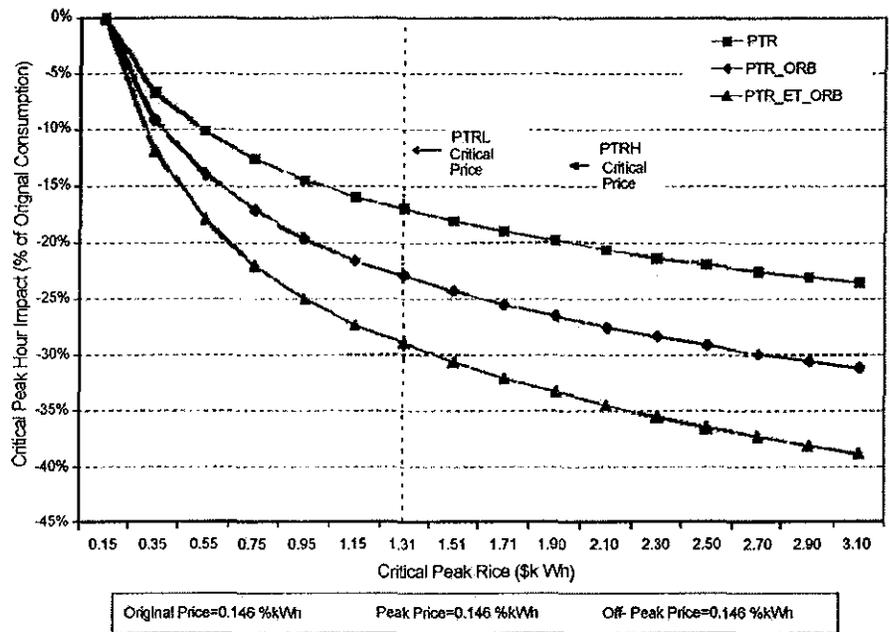


Figure 10: Impact of Enabling Technologies on Demand Response

potential to respond to the dynamic prices. Alternatively, customers who live in multi-family homes can be expected to be less price-responsive since they may not internalize all the benefits from changing their consumption patterns. Moreover, they generally have less control over the usage patterns of common electricity consuming equipment.

Low Income vs. High Income

The income of the participating customers is another socio-demographic characteristic that may impact the degree of price-responsiveness. Customers in the low-income brackets are generally more price-inelastic in terms of their necessity purchases. Electricity consumption is a necessity purchase for low-income customers who usually consume the required minimum and therefore don't have much

price-responsiveness potential. Also, these customers do not possess much of the electrical equipment whose load they could shift from peak to off-peak periods. On the other side, these customers could be more likely to shift any possible load from peak to off-peak in order to be able to realize the financial gains. Which of these effects would dominate is an empirical question. High-income customers have much more discretionary electricity consumption with which they can respond to dynamic rates. However, they may not be as highly motivated as low-income customers to realize the financial gains by changing their consumption patterns. Again, resulting impact can be determined through empirical investigation.

Description of Sample Selection

Smart Power is required by legislation to carry out a pilot

Table 2: Illustrative Experimental Design

Group	Low/High Rate	Enabling Technology	No Enabling Technology	Multi-family	Low Income	Control	TOTAL
TOU	Low Rate	150	150	–	–	–	300
	High Rate	150	150	150	150	–	600
CPP	Low Rate	150	150	–	–	–	300
	High Rate	150	150	150	150	–	600
PTR	Low Rate	150	150	–	–	–	300
	High Rate	150	150	150	150	–	600
TOTAL		900	900	450	450	300	3000

program with a minimum sample size of 2,750 customers. In light of the principles discussed above, we have crafted an illustrative experimental design for Smart Power.

In this strawman design, we propose to test the impacts of enabling technologies, multi-family home ownership, and low income on low and high rates from each of the TOU, CPP, and PTR rate designs. It is necessary to have low and high rates so that price

response curves can be estimated with statistical precision and so that results can be obtained for rates other than just the ones that are included in the experiment.

In each pricing pilot experiment, a group of control customers is required to anchor the impacts of the dynamic prices. These control customers remain on the existing rates and act as a comparison base representing how the treatment customers would behave but for the

dynamic prices and other pilot features. Table 2 presents a strawman experimental design for Smart Power. There are 18 program cells, each containing 150 customers, and a control group cell containing 300 customers. In total, there are 3,000 customers in the experiment sample.

This is, of course, only an illustrative design. A complete and successful experimental design is only possible after sufficient interaction with the utility through which the priorities and intricacies of the utility can be determined. A series of focus group meetings may be required with potential participants to understand their priorities and the incentives to which they react. For reference purposes, the design parameters of a few other pilot programs are included in the appendix to this article. ■

Appendix A. Experimental Designs from Recent Pricing Pilots

Table 3 Baltimore Gas and Electric (BGE) SEP Experimental Design

		Enabling Technology		No Enabling Technology	Control	TOTAL
		Orb Only	Orb + A/C Switch			
DPP	Normal Rate	–	111	148	–	259
PTR	Low Rate	141	113	126	–	380
	High Rate	137	118	127	–	382
TOTAL		278	342	401	354	1375

Track A: Random Sampling With Opt Out Design

	Control	CPP-F	CPP-F (info)	CPP-V (SDG&E)	Info Only	TOU	Total
Residential							
Zone 1	63	52	0	0	0	50	165
Zone 2	100	188	0	0	0	50	338
Zone 3	207	188	0	125	126	50	696

Appendix A (Continued)

Track A: Random Sampling With Opt Out Design

	Control	CPP-F	CPP-F (info)	CPP-V (SDG&E)	Info Only	TOU	Total
Zone 4	100	114	0	0	0	50	264
Total	470	542	0	125	126	200	1,463
Commercial SCE			CPP-V (SCE)			TOU (SCE)	
<20 kW	88	0	0	58	0	50	196
>20 kW	88	0	0	80	0	50	218
Total	176	0	0	138	0	100	414
All Sectors Total	646	542	0	263	126	300	1,877

Track B: SF Cooperative

Residential	Control	CPP-F	CPP-F (Info)	CPP-V	Info Only	TOU	Total
PG&E	63	64	126	0	0	0	253
Total	63	64	126	0	0	0	253

Track C: AB 970 Sub-Sample

Residential	Control	CPP-F	CPP-F (info)	CPP-V (SDG&E)	Info Only	TOU	Total
SDG&E	20	0	0	125	0	0	145
Total	20	0	0	125	0	0	145
Commercial SCE		CPP-F	CPP-F (Info)	CPP-V (SCE)	Info Only	TOU	Total
<20 kW	42	0	0	56	0	0	98
>20 kW	42	0	0	76	0	0	118
Total	84	0	0	132	0	0	216
All Sectors Total	104	0	0	257	0	0	361

Summary

	Control	CPP-F	CPP-F (Info)	CPP-V	Info Only	TOU	Total
Total Sample Size	813	606	126	520	126	300	2,491

Number of Residential Customers in the Experiment and Estimating Sample

Customer Segment	Climate Zone	Track	Tariff	Load Data		Load & A/C Ownership Data			
				Summer 2003	Winter	Summer 2003	Winter	Summer 2004	
R	1	A	Standard	68	62	64	51	47	48
R	2	A	Standard	106	107	108	90	92	90
R	3	A	Standard	105	108	108	89	88	81

Appendix A (Continued)

Number of Residential Customers in the Experiment and Estimating Sample

Customer Segment	Climate Zone	Track	Tariff	Load Data			Load & A/C Ownership Data		
				Summer 2003	Winter	Summer 2004	Summer 2003	Winter	Summer 2004
R	4	A	Standard	106	109	105	87	83	81
R	1	A	CPP-F	59	59	61	54	54	56
R		A	CPP-F	212	214	217	205	206	202
R	3	A	CPP-F	214	215	219	200	201	203
R	4	A	CPP-F	129	128	136	121	120	124
R	2	A	CPP-V			58			53
R	3	A	CPP-V			41			40
R	2	A	Info Only (Standard)	70	64	68	65	60	64
R	3	A	Info Only (Standard)	68	68	69	63	62	63
R	1	A	TOU	57	57	58	55	55	56
R		A	TOU	56	56	57	54	54	55
R	3	A	TOU	58	57	63	54	53	58
R	4	A	TOU	55	55	56	53	53	53
R	2	A	Standard			26			21
R	3	A	Standard			17			16
R	1	B	Info Only (Standard)	71	53	52	48	34	33
R	1	B	CPP-F	135	133	133	104	102	102
R	1	B	CPP-F	78	78	78	71	71	71
R	2&3	C	Standard	20	21	20	18	19	19
R	2&3	C	CPP-V	131	142	135	121	127	124
R	2&3	C	Standard	94	97	87	80	80	77

Endnotes:

1. Dynamic pricing rates are discussed by Ahmad Faruqui and Ryan Hledik in *The Power of Dynamic Pricing*, ELEC. J., April 2009, and in Ahmad Faruqui and Ryan Hledik, *Transitioning to Dynamic Pricing*, PUB. UTIL. FORTNIGHTLY, March 2009.

Inclining block rates are discussed in Ahmad Faruqui's *Inclining Toward Energy Efficiency*, PUB. UTIL. FORTNIGHTLY, Aug. 2008.

2. This section heavily relies on a Working Group 3 Report preceding

the implementation of the California Statewide Pricing Pilot. See Report of Working Group 3 to Working Group 1, "Proposed Pilot Projects and Market Research to Assess the Potential for Deployment of Dynamic Tariffs for Residential and Small Commercial Customers," Dec. 10, 2002.

3. HOWARD RAIFFA AND ROBERT SCHLAIFER, *APPLIED STATISTICAL DECISION THEORY* (MIT Press, 1961).

4. It is important to note that while estimated cost-effectiveness of each

sample cell is used in determining the optimal experimental design, this cost-effectiveness estimate in no way prejudices the ultimate cost-effectiveness results following implementation of the pilot and based on those experimental results.

5. For more information, see Ahmad Faruqui and Sanem Sergici, *Household Response to Dynamic Pricing of Electricity: A Survey of the Experimental Evidence*, Jan. 2009, at <http://www.hks.harvard.edu/hepg/>

6. *Id.*