



FREEMAN, SULLIVAN & CO.

A MEMBER OF THE FSC GROUP

2010 Energy Conservation Evaluation of Pacific Gas & Electric Company's Energy Alerts and Customer Web Presentment Programs

April 29, 2011

Freeman, Sullivan & Co.
101 Montgomery St., 15th Floor
San Francisco, CA 94104



Prepared for:
Pacific Gas & Electric Co.

Prepared by:
Stephen George, Ph.D.
Mike Perry, Ph.D.
Sarah Woehleke
Freeman, Sullivan & Co.

Table of Contents

1 Executive Summary	1
2 CWP and Energy Alert Marketing	3
3 CWP Population, Enrollment and Usage	4
4 Energy Alert Enrollment and the Frequency of Alerts	8
5 Analysis Methodology and Results	12
5.1 An Ideal Experimental Design	12
5.2 Limitations of Matching Methods	14
5.3 Matching Methods Used.....	16
5.3.1 Stratified Matching for Energy Alerts.....	17
5.3.2 Propensity Score Matching for CWP.....	17
5.4 Impact Analysis	18
5.4.1 Graphical Results	18
5.4.2 Regression Results.....	25
6 Recommendations	28
Appendix A. Matching Results	29

1 Executive Summary

Pacific Gas & Electric Company's (PG&E) Energy Alerts and Customer Web Presentment (CWP) are two informational energy conservation programs for residential customers that were first marketed in 2010.¹ The Energy Alerts program became available in June 2010 as an option for PG&E customers with an installed SmartMeter™ meter that is being read remotely (typically one to two billing cycles after installation). The program allows customers to have advance warning via email, phone or text message if their electricity usage is projected to push them into higher pricing tiers by the end of the current billing cycle. Projected usage is calculated on the 8th day of the customer's billing cycle, and Energy Alerts are subsequently sent out to those customers whose total usage for the billing cycle is likely to enter the higher (3rd, 4th, or 5th) pricing tiers. Energy Alerts are also sent out when the customer's usage has actually entered any of the higher pricing tiers, but total Energy Alerts per billing cycle are capped at 4 per service agreement.

PG&E's tiered pricing is structured to encourage energy conservation in line with the aims of The State of California's energy policy. There are five tiers of usage which correspond to increasing prices per kWh. The first tier is a baseline which is determined by season, climate zone, use of electric heat vs. gas heat, and typical electricity use by customers in a specific region. The baseline is meant to cover a substantial portion of the energy needs for an average customer. The tiers above the baseline are categorized as percentages of baseline usage: Tier 2 is 101-130% of the baseline, Tier 3 is 131-200%, Tier 4 is 201-300%, and Tier 5 is any use in excess of 300% of the baseline. Energy alerts are only sent out for customers projected to enter, or entering, the 3rd, 4th, and/or 5th tiers. As of the end of December 2010, there were 30,155 customers enrolled in the Energy Alerts program.

CWP is available through PG&E's *My Account* program—its online portal. Once an installed SmartMeter™ meter is being read remotely (again, typically one to two billing cycles after installation), the customer may log onto their *My Account* page to check their energy usage through the previous day. The widgets on the site allow customers to view their energy usage by hour, day, or week. The "Energy Highlights" option gives customers a quick overview of their usage characteristics, like how much their current monthly bill is projected to be and what their average daily cost of electricity is. Additionally, customers can see how much they are paying per hour of electric use during the month. As of the end of December 2010, there were 1.4 million customers with current account numbers in the *My Account* program, of which 1.1 million had SmartMeter™ meters. Of these, about 128,000 logged in to CWP at least once in 2010.

It is important to note that this is the first evaluation of these programs, and therefore there are research design issues that affect results here that will be handled better in future evaluations. Neither of these programs was undertaken with an experimental design in order to allow for evaluation of any conservation effects that they might generate. Statistical matching strategies were employed to select control groups for use in evaluating the energy conservation impacts of these programs. For CWP, propensity score matching was used and for Energy Alerts, stratified matching was used. The two strategies are very similar and either could be used for both programs. In this case, propensity score

¹ CWP is also offered to commercial customers with less than 200 kW peak demand. There were about 2,300 such customers with SmartMeters who viewed CWP in 2010. They have not been analyzed in this report.

matching was used for CWP in order to allow both for a greater number of variables to be used for matching and for a greater number of customers to be retained in the matching process. Stratified matching was chosen for Energy Alerts customers because the procedure is more transparent. As discussed in the Analysis section of the report, these choices had no effect on the primary conclusions of the evaluation. The results of the matching processes (as opposed to the estimated conservation impacts) are shown in the Appendix.

These matching strategies have weaknesses that are discussed in the main analysis section. The estimated conservation effects for both programs for 2010 are zero. However, there is significant uncertainty in these estimates, so it is possible that the programs could affect usage by 1-2% in either direction.² Future evaluations will be conducted with greater attention paid to the experimental design. Greater attention will also be paid to directing these programs towards customers who are likely to respond and to designing the most effective messaging for each program. This should provide larger and more precise estimates of conservation effects in the future.

² If the programs have the unintended effect of informing consumers that electricity is cheaper than they thought or that conservation is more effort than it is worth to them, then usage could increase due to these programs.

2 CWP and Energy Alert Marketing

Prior to 2010, neither CWP nor Energy Alerts were marketed to customers. Beginning in early 2010, PG&E began marketing both CWP and Energy Alerts using both mail and email and using several different marketing pieces. These included:

- Pre-installation bill inserts sent to customers who were to have SmartMeter™ meter installed in the near future. These inserts introduced customers to the SmartMeter™ meter and how it could help them manage their electricity use. Both CWP and Energy Alerts were highlighted in these inserts. Roughly 800,000 of these were sent out from January through April 2010.
- A SmartMeter™ Welcome Kit sent after the meter was installed and once it was being read remotely. This kit included information on how to read a SmartMeter™ meter and it introduced several programs that PG&E has to help customers reduce their energy bills, including CWP. Energy Alerts were not mentioned in the SmartMeter™ Welcome Kits. Roughly 1.7 million SmartMeter™ Welcome Kits were sent out to residential customers from April through August 2010.
- A Transition Booklet, which replaced the SmartMeter™ Welcome Kits beginning in September 2010 and which contained similar information to the Welcome Kits. The Transition Booklet advertised both CWP and Energy Alerts. About 900,000 Transition Booklets were sent to residential customers from September to December 2010.
- An email sent to about 14,000 customers in June 2010 announcing that Energy Alerts were available. These customers had previously indicated their interest in Energy Alerts.
- The Anatomy of a Rate mailing, which was sent in August 2010 to customers who had had bills in tier 3. This included an explanation of the tiered rate structure, along with discussions of both CWP and Energy Alerts and how they can be used to manage electricity use. About 560,000 of these mailings were sent out.

In each case, CWP was always marketed as a feature of *My Account*. Customers were directed to PG&E's website where they could sign up for or log in to *My Account*, at which point they could use the menus on the site to get information about their energy use through CWP.

3 CWP Population, Enrollment and Usage

The potential population for the CWP program consists of those customers who have signed up for *My Account* and who have an installed SmartMeter™ meter being read remotely. For background, Figure 2-1 shows the number of residential SmartMeter™ meter installations by month in PG&E's territory.

Figure 2-1
Monthly Residential SmartMeter Installations




Figure 2-2 shows the number of *My Account* sign-ups by month. It excludes the period prior to January, 2006, during which about 640,000 customers signed up for *My Account*, but for which sign-up dates are not available.

Figure 2-2
Monthly My Account Sign-Ups





Figure 2-3 shows the number of log-ins per month to CWP and the monthly number of unique visitors. There is a steady upward trend as more customers get SmartMeter™ meters and more customers sign-up for *My Account*. The number of unique visitors is consistently 20-30% of the number of log-ins each month.

Figure 2-3
Monthly CWP Log-ins and Unique Visitors




A total of 2.3 million customers (including customers with and without SmartMeter™ meters) have enrolled in *My Account* since it became available, and 8% have logged in to CWP at least once. As of the end of December 2010, there were 1.4 million customers with current account numbers in the *My Account* program, of which 1.1 million had SmartMeter™ meters. Of these, 127,643 (12%) logged in to CWP at least once in 2010.

For purposes of the energy conservation analysis, the treatment group is limited to people who logged in to CWP at least once in 2010, and for the first time in 2010. In 2010, 600,791 customers with SmartMeter™ meters enrolled in *My Account*. Figure 2-4 shows the number of new *My Account* customers who used CWP a given number of times. While 90% of customers enrolled in *My Account* 2010 but never used CWP, 10% used it at least once, and of those, 3% used it once, 5% used it 2-5 times, 1% used it 6-10 times, and 1% used it 11 or more times.

All of these figures point to a similar conclusion: regardless of how the *My Account* population is dissected, about 10% of *My Account* customers tend to ever view CWP pages.

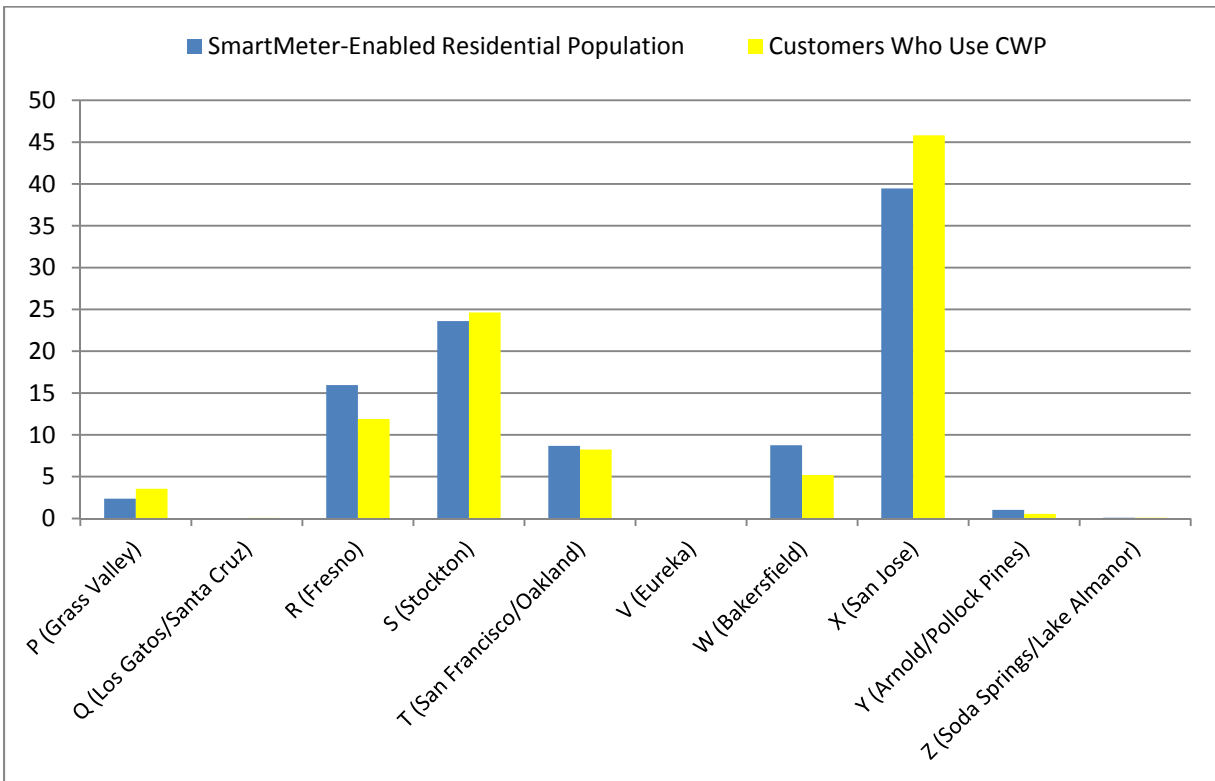
Figure 2-4
Distribution of 2010 Log-Ins by *My Account* Customers who First Signed Up in 2010



Geographically, customers who participate in CWP (who have logged on to view their usage at least once) are distributed fairly similarly to the SmartMeter™-enabled residential population. The differences in climate zone distribution between the SmartMeter™-enabled residential population as of June 2010 and customers who log-in to CWP are shown in Figure 2-5. The figure also includes the name of the city or cities in each zone with the largest PG&E residential population for context. Differences between the geographic distributions of the two populations are small.

There is a larger proportion of CARE customers in the general population than among customers who use CWP (27% in the population versus 14% of users). Nearly all of the customers in both the population and among those who use CWP are on the E1 residential rate schedule (about 95%).

Figure 2-5
Percentage of CWP Customers and of SmartMeter™-Enabled Residential Customers in each Climate Zone



Having logged in to CWP, customers have the option of looking at three charts: daily total energy; daily total energy with weather; and hourly usage. It appears that in some cases there are some other chart options available, but they account for a very small fraction of page views. The daily total energy page accounts for 73% of total (rather than unique customer) page views, followed by hourly usage with 19% and daily total energy with weather at 8%. These frequencies are quite stable over different times of year and over different years.

4 Energy Alert Enrollment and the Frequency of Alerts

Enrollment in Energy Alerts began in June 2010. As of the end of December 2010, there were 30,155 customers enrolled in the program. As shown in Figure 3-1, the rate of enrollments per month peaked in the middle of summer 2010 with about 8,000 new participants in August, and after a steep drop from August to September, remained fairly stable at a rate of less than 3,000 per month through December 2010. The peak in enrollment during August seems likely to be due to the Anatomy of a Rate mailing sent out in early August specifically to customers in tier 3 and above.

Figure 3-1
Monthly New Enrollments in Energy Alerts June-December 2010






Figure 3-2 shows the daily number of energy alerts sent out for each date in 2010. It also shows the daily number of first energy alerts sent out. In the latter case, only the first time a customer ever gets an Energy Alert is counted.

Figure 3-2
Energy Alerts by Date



Almost a quarter of the customers who enrolled in the Energy alert program in 2010 never received an Energy alert in 2010, as shown in Figure 3-3. The figure shows the number of Energy Alerts received per customer. It is interesting to note that over 25% of all enrolled customers (over 35% of customers actually receiving Energy alerts in 2010), received more than 5 Energy alerts in 2010.

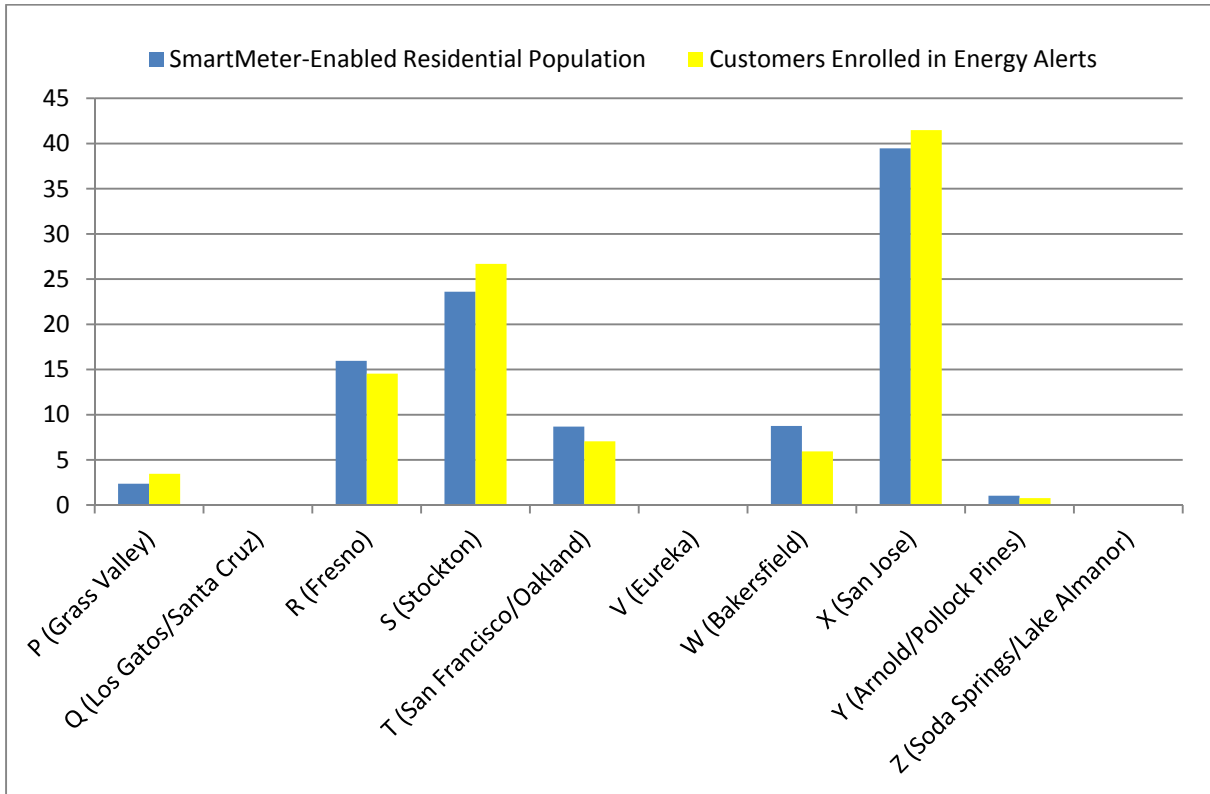
Figure 3-3
Number of Customers Receiving a Given Number of Energy Alerts



As was true for CWP customers, geographically, customers enrolled in the Energy Alert program are distributed fairly similarly to the SmartMeter™-enabled residential population. This is shown in Figure 3-4. As in Figure 2-5, Figure 3-4 shows the name of the city or cities in each zone with the largest PG&E residential population for context. As was the case for CWP customers, differences between the geographic distributions of the two populations are small.

The rate schedules of Energy Alert subscribers are almost identical to the full residential population with over 95% of customers on the E1 residential rate. Energy Alert subscribers are less likely to be in the CARE program with only 13% vs. 27% in the population.

Figure 3-4
Percentage of Energy Alert Customers and of SmartMeter™-Enabled Residential Customers in each Climate Zone



5 Analysis Methodology and Results

This section discusses the methods used to estimate the change in energy use associated with Energy Alerts and CWP and presents the estimated conservation impacts. It begins with a discussion of what an ideal research design would be to evaluate these effects. This thought exercise makes clearer what the matching methods can and cannot control for, and provides context for judging how reliable the results are. It also shows the difficulties such experiments would entail, and may bolster the case that in some cases, experiments are not worth doing.

Following the discussion of the ideal research design is a description of the matching strategies used to partially overcome the lack of an experimental design. The section concludes with graphical and regression results showing the estimated effects of the programs.

In the discussion below, a treatment refers to a program that a customer might join that might cause the customer to reduce energy use, such as Energy Alerts or CWP. A treatment group refers to a group of customers receiving a treatment by joining such a program, and a control group refers to a group of customers not receiving a treatment whose energy usage is supposed to accurately represent what the usage of the treatment group would have been if they had not been on the treatment. A treatment effect is the estimated difference in usage between the treatment group and the control group.

5.1 An Ideal Experimental Design

Before discussing the methods used to evaluate the two programs, it is worth considering what an ideal research design would look like to test the effects of these programs. An ideal design would involve a controlled experiment in which customers are randomly assigned to treatment and control conditions.

The key to an experiment is to remove from the customer's control any choice about whether they receive the treatment. Instead, whether a customer receives a treatment must be random. When one large, randomly-chosen group of customers can be selected as a treatment group and another can be selected as a control group, then it is a safe assumption that the average usage among customers in the control group is close to what average usage would have been among the treatment group if they had not been in the program. In that case, the difference between the average usage in the treatment group and the average usage in the control group is an accurate measure of the impact of the program.

In the absence of such a design, when customers themselves decide whether or not to receive a treatment, there will always be at least one important difference between the population receiving the treatment and the population not receiving the treatment. That difference is that one group elected to receive the treatment and the other group either elected not to receive the treatment or was not offered the treatment, but contains customers who would choose not to take the treatment. The fact that this difference exists makes it likely that other important differences exist that explain the difference in choices between the two groups. These other differences can cause the usage among customers on the treatment to be different from the usage of customers not on the treatment, even if the treatment had not been offered.

For example, suppose that recently unemployed customers were looking for ways to save money and were therefore likely to use less electricity in the future and more likely to sign up for Energy Alerts. In this case, customers who receive Energy Alerts would have less average usage than customers who did not. However this would be partially because they were already trying to save energy rather than purely due to receiving the alerts. In such a situation, without a control group, there is no way to separate the impact of the alerts from the fact that treatment group customers would have used less energy anyway. Conversely, suppose that many customers anticipated large increases in their electricity usage during the summer of 2010 due to purchase of new appliances. Suppose these customers were more likely to sign up for Energy Alerts in order to try to keep their new, larger electricity budgets under control. Then customers in the Energy Alerts program might have larger average usage values than customers not on the Energy Alerts program. This could be true even if the usage values were larger by a smaller amount than they would have been had the customers not had the option of being on Energy Alerts. Again, without a properly randomized control group, fully separating these two issues presents significant challenges and cannot be accomplished with certainty.

In the real world, there can be many such correlations between customer characteristics and customer willingness to sign up for a particular program. Knowing how all these potential factors balance out without a control group is frequently not possible. The same examples could be used for CWP as for Energy Alerts.

Now, consider what an experiment to evaluate the effect of CWP would entail. Suppose the question is, “Did customers who used CWP use less energy in 2010 than they would have if CWP was not offered, and if so, how much?” An experiment to answer that question would require very large sample sizes due to low utilization of CWP among eligible customers and due to low expected effect sizes among users. As was shown in Figure 2-4, the vast majority of *My Account* customers never use CWP. In order to measure the effect of CWP, one randomly-chosen group of customers would have to be offered CWP (the treatment group) and another would have CWP withheld (the control group). An important aspect of this design is that energy usage is recorded for all customers in both groups, including those in the treatment group who did not actually log in. This preserves the random aspect of treatment. This is known as a randomized-encouragement design (RED). The RED works best in situations where take-up rates (the number of customers logging in, in this case) and effect sizes are large. CWP satisfies neither of these conditions, which means that large samples would be necessary to distinguish treatment effects from random variation in usage. Based on simulations done for experimental designs to evaluate other information-based energy conservation programs, the combined size of the treatment and control groups for such an experiment might be not significantly smaller than PG&E’s entire residential population.

There are other ways that the experiment could be conducted, such as offering a treatment group a monetary incentive to log in to CWP or by focusing the experiment on customers particularly likely to use CWP. This would solve the problem of the low take-up rate and therefore reduce the necessary sample size. However, the results of such an experiment would be harder to generalize to the larger population.

An experiment to measure the effects of Energy Alerts would be simpler because PG&E has more control over whether customers receive the treatment; although PG&E cannot force customers to actually pay attention to alerts. Suppose the question is, “Did customers who signed up for Energy Alerts use less

energy than they would have had Energy Alerts not been offered, and if so, how much?” An experiment to answer that question could be done in at least two ways. First, a RED could be used. In that case, two large random groups of customers would be offered the opportunity to sign up for Energy Alerts. In practice this method would face the difficulty of extremely low sign-up rates, as evidenced by the fact that only about 30,000 customers signed up for Energy Alerts in 2010 out of more than 1 million customers marketed that the program was marketed to. Second, a group of customers could be offered Energy Alerts, but a randomly-chosen half of those customers could be told that their entry in the program would be delayed for a period long enough to measure their usage as a control group. This “recruit and delay” strategy could work well in situations where take-up rates are low. However, it also entails some risk in terms of customer relations as customers may not like having their enrollment delayed.

Finally, suppose the question is, “Is there an effect of receiving multiple Energy Alerts, over and above receiving only one, and if so, how large is that effect?” Answering this question cannot be done using any of the above experiments because as the program currently exists, whether a customer receives multiple alerts is up to the customer. The customer implicitly decides whether to receive multiple alerts by either conserving electricity after receiving the first alert or not.

Here the treatment is defined as receiving multiple alerts. Again, the key to an experiment is to take out of the customers’ hands any control over whether they receive the treatment. This can be accomplished by an experiment in which two groups of customers are both in the Energy Alerts program, but one group only ever receives their first alert and then the program ends, while the other group receives alerts according the standard program rules.

As discussed below, the matching strategies which constitute the primary analysis only partially replicate any of the experimental designs outlined above.

5.2 Limitations of Matching Methods

Matching is used to partially eliminate selection bias that results from customers themselves deciding whether to take part in the two programs of interest. The basic idea is that customers who elect to sign up for Energy Alerts, for example, made that decision for reasons that may also lead them to have different usage profiles than the population at large. Matching strategies rely on an assumption that the differences between customers who receive the treatment being evaluated can be fully described and controlled for using observable variables. This assumption can never be proven because it is always possible that some unobservable variables exist that are correlated with both the treatment and the customer’s usage (an example is given below). In some situations it is plausible that the assumption is close enough to reality that a matching strategy can be assumed to be reliable. This is primarily true in situations where the treatment itself is expected to have a large effect as compared to the effect of these unobservable variables.

To illustrate the assumption that selection bias can be controlled for using observable variables, consider the example above where customers who were recently unemployed were both more likely to sign up for Energy Alerts and more likely to reduce their electricity usage regardless of Energy Alerts. As was discussed, in this example a naïve comparison of usage among customers receiving Energy Alerts with

customers not receiving Energy Alerts would over-state the effect of Energy Alerts because customers receiving Energy Alerts were intending to reduce their usage anyway. Now suppose variables that are observable, such as climate zone, CARE status, and previous usage behavior are also correlated with customers becoming unemployed. In this case, if a comparison is made between usage among customers in the Energy Alerts program with customers not in the Energy Alerts program, but who live in the same climate zones, have the same distribution of CARE status and have similar previous usage behavior, then the amount of bias in the estimate of the Energy Alerts effect will be reduced. This will be because, by supposition, customers with similar observable characteristics also have similar recent unemployment status. If the assumption seems reasonable that after controlling for observable factors there is no remaining correlation between customers who sign up for Energy Alerts and recent unemployment, then that implies that there is no more selection bias due to unemployment. This would be the case if, for example, within a given climate zone, customers with the same CARE status and previous usage behavior were no more likely to sign up for Energy Alerts whether they were recently unemployed or not.

This assumption that selection bias can be controlled for with observable variables cannot be verified using the data because there could always be remaining unobservable variables, such as unemployment in the previous example, that are correlated with the treatment. In situations where the treatment effect is large and there is a fairly large set of observable variables to use for matching, it can be plausible that most selection bias can be removed from an analysis. However, the programs evaluated here are likely to have conservation effects of less than 5% of usage. It is very difficult to be confident that selection bias has been so completely removed by matching that an effect size less than 5% can be measured accurately. Given that effect sizes of these programs are likely small, there is almost no result that would be more likely to be due to the program than due to remaining selection bias after matching.

To illustrate this, it is worth considering what the interpretation of some hypothetical results would be. Suppose the matching analysis found a 1% difference in usage between customers who use CWP and a matched control group. In that case, a 1% difference in usage could easily be due to remaining selection bias after matching. Therefore a 1% difference between groups would not be clearly interpretable as an effect of the program. Suppose the matching analysis found a 30% difference between groups. In that case, it would not be plausible that the CWP program caused customers to reduce usage by 30%, and the likely conclusion would be that matching had failed to eliminate selection bias.

The difference between the two cases above is that the case where the difference between the two groups is 1% of usage conveys much more information about the actual effect of the program than the case where the difference between the two groups is 30%. In the case where the difference is 1%, it is not clear whether there is any effect of the program or not, but it is very likely that if the program has an effect then the effect is small—probably within a couple of percentage points of zero in either direction. Any other possibility requires that the remaining selection biases after matching almost perfectly balance with the effect of the program. Much more likely is that both effects are small. This situation describes the results below.

5.3 Matching Methods Used

The CWP analysis was conducted for two groups of customers: those who logged in at least once and for the first time in 2010; and those who logged in more than five times and for the first time in 2010.

The Energy Alerts analysis was conducted for three groups of customers: those who signed up for Energy Alerts in 2010 between June and September; those who received at least one Energy Alert between June and September; and those who received more than two Energy Alerts between June and September. The reason that September was chosen as a cut-off date was because the effect of Energy Alerts may require some amount of time to occur as customers adjust their behavior.

The results for customers receiving three or more Energy Alerts are particularly difficult to interpret given the non-experimental nature of the analysis. Customers who receive several alerts do so because they continue to use a lot of energy despite getting a first and second alert. It is hard to imagine a matching strategy based on observable variables that could properly control for that type of self-selection.

Two different matching strategies were used for CWP and Energy Alert customers. These strategies, known as stratified matching and propensity score matching, are quite similar, require similar assumptions and in situations where either can be used, generally produce similar results. Both strategies construct control groups by finding customers not exposed to the treatment that have observable characteristics similar to customers exposed to the treatment. The difference is that propensity score matching combines many observable characteristics into one index (the propensity score) and then finds control group customers with index values similar to customers in the treatment group. Stratified matching consists of finding control group customers who have similar characteristics along each observable characteristic used for matching. Stratified matching is more transparent because it does not require estimating an index function, but it works less well when there are many observable dimensions to match on because there may not be any customers in the population to match a given treatment customer along every observable characteristic used for matching.

The use of both matching strategies here is primary for the purpose of illustration. Propensity score matching works well for both programs, but the procedure is less clear, both conceptually and mechanically. Stratified matching is used because it is more transparent exactly what was done and because it better illustrates the conceptual basis of matching. Both propensity score and stratified matching were initially used for both programs to test whether it made any difference to the final conclusions. The matching strategy had no effect on the final conclusion, so each strategy is included once here to strike a balance between brevity and exposition.

For both CWP and Energy Alerts, a similar set of matching variables was used (as detailed in the Appendix). These variables are electricity usage prior to the study period, the probability of having central air-conditioning (CAC), LCA, CARE status and rate schedule.

The probability of having CAC is a variable calculated by FSC for another evaluation.³ It is developed using data from the Residential Appliance Saturation Survey and its development is documented in that report. The probability is determined primarily by analyzing the way that the customer's electricity use changes with changes in temperature over time. This is based on the basic idea that customers with CAC will use significantly more electricity as temperatures increase than customers without air conditioning. In order to avoid matching customers based on electricity usage during the periods of interest for this evaluation, CAC probability for this evaluation was calculated using customers' monthly usage values through December 2009 and May of 2010 for CWP customers and Energy Alert customers, respectively.

5.3.1 Stratified Matching for Energy Alerts

For each group of Energy Alert customers evaluated, the matching procedure consisted of three steps:

1. Assign each customer in the treatment group to a stratification cell based on their quartile of average usage during the summer of 2009, their respective quartiles of usage in January, March and May of 2010, their quartile of CAC probability, their CARE status and their climate region.⁴ Usage up to May 2010 was used because no customer's usage could have been affected by Energy Alerts prior to June 2010.
2. Assign each customer in the general population to a cell based on the same criteria.
3. For each customer in the treatment group, randomly choose one customer in the general population from the same cell. The set of randomly chosen customers constitutes the control group for the respective treatment group.

The results of the matching process for both Energy Alerts and CWP customers are shown in the Appendix.

5.3.2 Propensity Score Matching for CWP

For each group of CWP customers evaluated, the propensity score for each customer was calculated using a probit regression on the entire residential population for whom all matching variables were available. The probit was used to estimate the probability of a customer being in the treatment group based on the observed variables. The estimated probabilities from the probit regression were the propensity scores used for matching. The observed variables used in the probit function were the customer's monthly usage for each month in 2009, quartile of CAC probability, CARE status and climate region. Usage from 2009 was used because customers who did not log in prior to 2010 could not have had their 2009 usage affected by CWP.

³ 2009 Load Impact Evaluation for Pacific Gas and Electric Company's Residential SmartRate™—Peak Day Pricing and TOU Tariffs and SmartAC Program Volume 2: Ex Ante Load Impacts. FSC Group, April 1, 2010.

⁴ Although billing periods vary by customer, monthly usage was determined by finding the mid-date of each billing period and assigning the month based on that date, thereby creating a uniform system to compare usage among the full population of customers.

Having calculated each customer’s propensity score, for each customer in each treatment group, the customer in the general population with the closest propensity score was chosen to be part of the control group. The set of customers selected in this way constituted the control group for each respective treatment group.

5.4 Impact Analysis

In this section graphical results are presented first that show that the effects of CWP and Energy Alerts are small. Following that, results from regressions are presented that more precisely measure the effect of each treatment.

5.4.1 Graphical Results

Table 4-1 shows the number of customers exposed to each treatment as defined here and the number of customers actually used in each analysis. The number of customers exposed to each treatment is significantly larger than the number of customers actually analyzed using matching. This is because customers were excluded from the analyses due to missing or outlying data. For example, many customers were excluded from the analysis of CWP due to missing values in 2009 monthly usage or due to bills not coming at monthly intervals. For example, in many instances a customer’s bills might be separated by 50 days then by only 10 days. Such cases were excluded, although could probably be accommodated in future analyses. Also, some customers in each population had such high pre-treatment usage values that it was not possible to find appropriate matches for them in the population. These customers were excluded.

**Table 4-1
Number of Customers in Each Population and Treatment Group**

Group	CWP Users	CWP, 6 or more Log-ins	Energy Alert Enrollment	Energy Alert Recipients	Energy Alert, 3 or more Alerts
Customers subject to treatment	127,643	26,589	21,495	15,963	2,848
Customers analyzed	55,542	11,898	13,638	10,513	1,880

The treatment group sizes in Table 4-1 are such that if there was a certainty that selection bias had no effect, then effect sizes of less than 1% could be detected for CWP users, and effect sizes of roughly 2% could be detected for CWP users who logged in 6 or more times, for customers enrolled in Energy Alerts and for customers receiving Energy Alerts. Effect sizes of roughly 5% could be detected for customers receiving three or more Energy Alerts. These values show that lack of sample size is not the primary impediment to precise impact estimates. Selection bias is the primary issue. The only reason reduction in sample size would be likely to have an effect on this analysis would be if there was some strong reason to think that customers who were excluded from the analysis also happened to be customers who reduced usage in reaction to the treatments. This is unlikely.

The results of the matching strategies for each treatment group are shown in the appendix. In each case, the control group matches well with the treatment group based on observable characteristics. Also, for both CWP customers and Energy Alert customers, focusing on higher users means focusing on customers with characteristics more unlike the rest of the population. For example, customers receiving more than two Energy Alerts have significantly higher average usage values than the average in the population. They also have significantly higher usage values than customers who simply enroll in the program.

Figures 4-1 through 4-10 show graphically the primary results of each matching analysis. Figure 4-1 shows average monthly usage for 2009 and 2010 for customers who logged in to CWP for the first time in 2010. It also shows monthly usage for the matched control group (referred to as the population sample). That the matching strategy worked well is demonstrated by the fact that the CWP group's monthly usage for each month of 2009 is very close to the control group's monthly usage. The primary result of interest is that as the number of log-ins in the population increases during the year, as shown in Figure 4-2, there is no downward trend in the usage of the CWP group as compared to the control group. Figures 4-3 through 4-10 show the same basic result for CWP customers who logged in more than five times, for customers who Enrolled in Energy Alerts, for customers who received at least one Energy Alert and for customers who received more than two Energy Alerts prior to October 2010. In each case, usage between the two groups matches very well in 2009. Then as the treatment becomes more prevalent among the treatment group, there is no noticeable effect on usage as compared to the matched control groups.

Another aspect of the graphs to note, and which is discussed in more detail in the appendix, is that even during the pre-treatment period of 2009 for CWP and 2009-May 2010 for Energy Alerts, usage between the treatment and matched control groups only matches to within 0.5-2%. This is a problem if the goal is to accurately measure an effect of less than 5%. This is the limit of precision that the matching can provide, and it is not adequate for accurately measuring small effect sizes. This is a separate issue from selection bias due to unobservable variables, but just as important.

The shortcoming of these graphical results is that they do not provide a single number as an estimate for the effect on each treatment group. They cannot be used for that purpose because the level of the treatment in each treatment group changes multiple times over the time period shown. They have the virtue that they illustrate both that the matching strategies appear to have worked well in terms of finding customers with similar pre-treatment usage and that any effect estimated using regressions on these customers will be small.

Figure 4-1
Monthly Usage for CWP Customers Logging-in For the First Time During 2010




Figure 4-2
Fraction of Customers Having Logged In Among CWP Customers Logging In For the First Time During 2010




Figure 4-3
Monthly Usage for CWP Customers Logging-in at least 6 times in 2010

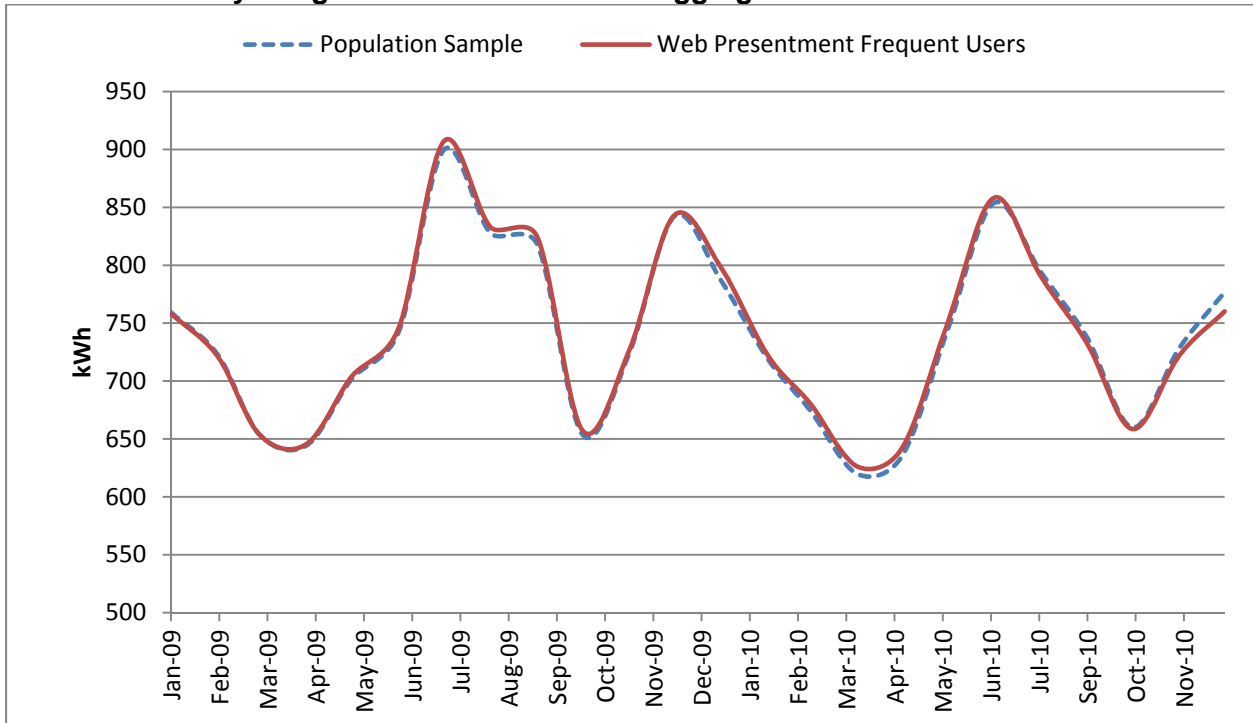


Figure 4-4
Fraction of Customers Having Logged In at least 6 times Among CWP Customers Logging-in for the first time in 2010




Figure 4-5
Monthly Usage for Energy Alert Customers Enrolling between June and September of 2010




Figure 4-6
Fraction of Customers Having Enrolled Among Energy Alert Customers Enrolling between June and September of 2010




Figure 4-7
Monthly Usage for Energy Alert Customers Receiving at Least 1 Alert between June and September of 2010




Figure 4-8
Fraction of Customers Having Received an Energy Alert Among Energy Alert Customers Receiving an Alert between June and September of 2010




Figure 4-9
Monthly Usage for Energy Alert Customers Receiving at Least 3 Alerts between June and September of 2010





Figure 4-10
Fraction of Customers Having Received at Least 3 Energy Alerts between June and September of 2010, Among Those Receiving at least 3 Between June and September



5.4.2 Regression Results

Regressions were also used to produce a single numerical estimate of the treatment effect for each treatment group. Regressions were performed on the same monthly datasets that produced the graphical results in the previous section. Usage was measured monthly at the individual customer level. Effects were estimated both in terms of kWh and in terms of percentage of usage.⁵ Each regression included indicator variables for each customer, indicator variables for each month of each year and an indicator that equaled one once the customer had received the treatment and zero prior to that.

For CWP users, receiving the treatment was defined as logging in prior to the relevant billing date. Changing the definition by defining the treatment as having logged in 15 days prior to the billing date had no material effect on the results.

For CWP users logging in more than five times, treatment was also defined as having occurred once the sixth log-in took place. For these customers, two alternative definitions of the treatment were also used. These were: the treatment was defined to start after the customer's first log-in; and the treatment increased steadily from zero to one from the time the customer logged in first to the time the customer logged in for the sixth time. The rationale for the second definition was that perhaps customers react more to CWP as they have logged in more times and spent more time thinking about the information. Neither of these alternative treatments yielded very different coefficient estimates.

For Energy Alert customers, the treatment was defined as starting once the customer enrolled, once the customer had received the first Energy Alert or once the customer had received the third Energy Alert, for the three treatment groups, respectively. For the Energy Alert customers receiving three or more alerts, two alternative treatment definitions were also used. These were: the treatment was defined to start after the customer's first alert; and the treatment increased steadily from zero to one from the time the customer received the first alert to the time the customer received the third alert.

Table 4-2 shows the estimated coefficients for each treatment using the kWh model and the percentage model. P-values for the coefficient estimates from the regression models are displayed as well; however these must be viewed as under-estimates of the true p-values. P-values are calculated directly from the estimated standard errors of the regression models without any adjustment for the matching procedures. The authors are not aware of any currently accepted method for calculating standard errors that result from a propensity score matching procedure plus regression.⁶ There is a method for calculating standard errors that result from a stratified matching procedure plus regression, but it is fairly new, it is not included as a standard part of any statistical package that the authors are aware of, and there was not adequate time to manually implement this procedure for this evaluation. The issue for both propensity score matching and stratified matching is that the matching procedure itself produces estimates in the form of matched customers. There is associated variance with those matches that a naïve regression does not include in its standard error estimates. This issue regarding the standard errors and p-values is worth

⁵ The model is estimated in percentage terms by performing a regression on the logarithm of kWh rather than kWh itself.

⁶ Even bootstrap methods, which are typically the most reliable and robust methods for calculating standard errors, have not been shown to be valid in the context of propensity score matching.

being aware of, but it has little impact here because the regression coefficients themselves are close to zero, and to the degree that they differ from zero, selection bias is just as likely a reason as any actual effect of the programs. Statistical significance is not important when there is no practical significance.

For the kWh model, the coefficient estimates can be interpreted as the estimated monthly difference due to the treatment between the treatment group and the control group in terms of actual kWh. For example, the estimated coefficient for CWP customers who log in at least once suggests that those customers use, on average 2 kWh less per month after logging in than they would have had CWP not been available. For the percentage model, the coefficients can be interpreted directly as percentages. For example, the regression estimate for customers receiving at least one Energy Alert suggests that they use, on average 2.3% more electricity per month after receiving an alert than they would if the program had not been in place.

For both models, p-values can be interpreted as percentages. For example, the p-value on the coefficient for customers enrolled in Energy Alerts using the percentage model is 1%. This indicates that the probability of estimating a coefficient at least as large as -0.3% if the true value was zero is 1%.

Due to the possibility of unobservable selection bias, there is no way to say whether any of the coefficients in the table are due to the treatments or whether they are due to selection bias or a combination of both. The coefficients with the largest magnitude in the table are the kWh coefficient for CWP customers logging in at least 6 times and the percentage coefficients for the same customers and for customers receiving at least one Energy Alert. A naïve interpretation of the CWP result might lead someone to believe that CWP frequent users have been shown to reduce usage by 26 kWh per month or by about 2%. However, that interpretation is problematic, because the same logic then implies that customers who receive an Energy Alert increase usage by 2%. That seems unlikely. More likely is that both results have an unknown amount of selection bias in them.

As discussed above, it is very unlikely that results such as those in the table would arise due to programs with large effects that just happen to have selection biases among customers that almost perfectly balance those effects. Far more likely is that the effects of these programs are close to zero and that the selection biases are also fairly close to zero after matching. Given that the results shown are all in a fairly narrow band of 1-2% of usage on either side of zero, it is likely that the true effects of the programs are somewhere in that range.

Table 4-2
Regression Estimates of the Effects of CWP and Energy Alerts

Program	Treatment	Model			
		kWh		Percentage	
		Coeff.	P-value (%)	Coeff.	P-value (%)
CWP	Log in at least once	-2	0.0	0.5	0.0
	Log in at least 6 times	-26	0.0	-2.2	0.0
Energy Alert	Enrollment	-13	0.0	-0.4	1.0
	Receive at least one alert	13	0.0	2.3	0.0
	Receive at least three alerts	4	36	1.3	1.0

6 Recommendations

For future evaluations of the CWP program, it is unlikely that any precise measurement of the effect of simply making CWP available to residential customers can be performed. As discussed above, such a study would require very large sample sizes. It is too late for such an experiment because CWP is already available to the majority of the residential population. In any case, it seems unlikely that CWP alone has a large impact on residential usage. Future evaluations should focus on whether there are ways to drive more customers to CWP and whether there are different CWP offerings that actually lead to significant usage reduction. For example, customers could be given incentives to log on to the website a certain number of times over a set period of time. Also, certain groups of customers could be offered richer sets of online information about their usage, suggestions for ways to reduce their bill or comparisons of their usage with other customers' usages. In all of those examples, the key to usage reduction would be to find ways to get customers to actually use the site and think about the information presented. All of those suggestions have the virtue that they could be implemented as experiments, without significant possibility for self-selection bias. This would lead to more certain impact estimates in the future.

Recommendations for Energy Alerts are similar to those for CWP. Currently, few customers have signed up for Energy Alerts of those marketed to. There is no evidence of any conservation among those customers who did sign up. Future evaluations should focus on whether there are ways to interest more customers in receiving alerts and on whether there are ways to better capture the attention of those customers who receive them. There is substantial potential both to test different Energy Alert marketing strategies to new customers and different Energy Alert messages to previously signed-up customers in controlled experiments. Doing so would allow for much better estimates of program impacts in the future.

Appendix A. Matching Results

The reason for using matching strategies for impact estimation is that self-selected treatment groups have characteristics that do not match the rest of the population. After matching, the treatment group is reduced in size, due to missing data as described in Section 5.4.1. After matching, the matched group has characteristics much closer to the characteristics of the treatment group than the entire population does. This is shown in Tables A-1 through A-5. The tables show average values for four different populations for the characteristics on which each control group was matched to each treatment group. As was described in the main text, different matching strategies were used for CWP and for Energy Alerts, which is why the set of variables in Tables A-1 and A-2 differs from the set in Table A-3 through A-5.

The four groups shown in each table are:

- PG&E's residential population in the first column of each table ("Population");
- The full population of customers who received the treatment ("Pre-drop Treatment Group");
- The final group of customers who received the treatment and who had sufficient data for matching ("Treatment Group"); and
- The matched control group ("Matched Group").

The primary aspect of the tables to notice is that the two columns on the right in each case have similar values for each characteristic while the two columns on the left do not. Also, it is important to note that the characteristics of customers in the "pre-drop treatment group" are somewhat different in mean value from the characteristics of customers in the final treatment group. Again, this is because some customers were dropped, as described in Section 5.4.1.

Finally, it is important to be aware that, even after matching, the control groups have characteristics that still differ systematically from the treatment groups. In each of the tables, the treatment group average usage values are higher than the control group average usage values by around 0.5-2%. This shows that, even neglecting the issue of selection bias due to unobservable variables, the matching strategies used here are not adequate to precisely measure usage impacts of less than 5%. If the treatment group only matches the control group to within 2% during the pre-treatment period and based on the variables that were themselves used to match, then it is easy to imagine that the two groups could naturally have future usage differences of 5%, unrelated to the treatments. In such a case, the regression analyses would measure innate differences between the groups rather than treatment effects.

Table A-1
Matching Results for Customers Using CWP for the first time in 2010

Matching Characteristic	Population	Pre-drop Treatment Group	Treatment Group	Matched Group
January 2009 Usage (kWh)	647	791	741	733
February 2009 Usage (kWh)	616	752	703	695
March 2009 Usage (kWh)	551	679	634	626
April 2009 Usage (kWh)	528	662	624	616
May 2009 Usage (kWh)	568	735	679	669
June 2009 Usage (kWh)	594	778	716	705
July 2009 Usage (kWh)	711	956	871	857
August 2009 Usage (kWh)	663	886	801	788
September 2009 Usage (kWh)	639	849	796	784
October 2009 Usage (kWh)	524	668	635	626
November 2009 Usage (kWh)	593	739	699	690
December 2009 Usage (kWh)	686	850	813	802
Average AC Propensity	0.48	0.63	0.59	0.59
% in each LCA				
<i>Greater Bay Area</i>	46%	56%	62%	64%
<i>Greater Fresno</i>	10%	12%	12%	12%
<i>Humboldt</i>	1%	N/A	0%	0%
<i>Kern</i>	4%	4%	N/A	N/A
<i>Northern Coast</i>	10%	3%	4%	4%
<i>Other</i>	18%	9%	10%	9%
<i>Sierra</i>	6%	11%	12%	11%
<i>Stockton</i>	5%	6%	N/A	N/A
% in each Climate Zone				
<i>P</i>	4%	4%	4%	4%
<i>R</i>	10%	11%	11%	12%
<i>S</i>	16%	24%	20%	17%
<i>T</i>	24%	8%	10%	15%
<i>W</i>	5%	4%	1%	1%
<i>X</i>	38%	48%	54%	50%
<i>Y</i>	1%	1%	1%	1%
% CARE				
<i>CARE</i>	27%	14%	11%	11%
<i>Non-CARE</i>	73%	86%	89%	89%
% in Each Rate Schedule				
<i>E1</i>	95%	94%	96%	96%

Table A-2

Matching Results for Customers Logging In at least 6 times Among CWP Customers Logging-in for the first time in 2010

Matching Characteristic	Population Values	Pre-drop Treatment Group	Treatment Group	Matched Group
January 2009 Usage (kWh)	647	810	758	759
February 2009 Usage (kWh)	616	773	721	723
March 2009 Usage (kWh)	551	699	653	653
April 2009 Usage (kWh)	528	685	646	644
May 2009 Usage (kWh)	568	762	704	702
June 2009 Usage (kWh)	594	810	746	742
July 2009 Usage (kWh)	711	995	907	900
August 2009 Usage (kWh)	663	919	833	828
September 2009 Usage (kWh)	639	878	824	818
October 2009 Usage (kWh)	524	692	657	654
November 2009 Usage (kWh)	593	770	727	724
December 2009 Usage (kWh)	686	883	843	842
Average AC Propensity	0.48	0.65	0.62	0.61
% in each LCA				
Greater Bay Area	46%	55%	61%	62%
Greater Fresno	10%	13%	13%	12%
Kern	4%	3%	N/A	N/A
Northern Coast	10%	3%	4%	4%
Other	18%	9%	10%	10%
Sierra	6%	11%	12%	12%
Stockton	5%	7%	N/A	N/A
% in each Climate Zone				
P	4%	4%	3%	5%
Q	0%	0%	0%	0%
R	10%	12%	12%	12%
S	16%	24%	21%	18%
T	24%	8%	9%	13%
W	5%	4%	1%	1%
X	38%	47%	53%	50%
% CARE				
CARE	27%	10%	9%	8%
Non-CARE	73%	90%	91%	92%
% in Each Rate Schedule				
E1	95%	94%	96%	96%

Table A-3

Matching Results for Customers Enrolling in Energy Alerts between June and September, 2010

Matching Characteristic	Population Values	Pre-drop Treatment Group	Treatment Group	Matched Group
Summer 2009 Usage ⁷ (kWh)	656	955	920	911
January 2010 Usage (kWh)	662	807	778	775
March 2010 Usage (kWh)	563	696	669	665
May 2010 Usage (kWh)	523	679	654	648
Average AC Propensity	0.48	0.70	0.69	0.69

Table A-4

Matching Results for Customers Receiving Energy Alerts between June and September, 2010

Matching Characteristic	Population Values	Pre-drop Treatment Group	Treatment Group	Matched Group
Summer 2009 Usage ⁷ (kWh)	656	1079	1044	1027
January 2010 Usage (kWh)	662	884	853	844
March 2010 Usage (kWh)	563	766	738	728
May 2010 Usage (kWh)	523	756	731	718
Average AC Propensity	0.48	0.74	0.74	0.73

Table A-5

Matching Results for Customers Receiving 3 or More Energy Alerts between June and September, 2010

Matching Characteristic	Population Values	Pre-drop Treatment Group	Treatment Group	Matched Group
Summer 2009 Usage ⁷ (kWh)	656	1291	1245	1216
January 2010 Usage (kWh)	662	993	940	930
March 2010 Usage (kWh)	563	864	818	809
May 2010 Usage (kWh)	523	870	829	818
Average AC Propensity	0.48	0.79	0.79	0.78

⁷ Summer usage here refers to average monthly usage for the months June through August, 2009.