

Using Information to Improve the Effectiveness of Nonlinear Pricing:
Evidence from a Field Experiment
by

Matthew E. Kahn
Institute of the Environment
Department of Economics
University of California—Los Angeles
mkahn@ioe.ucla.edu

Frank A. Wolak
Program on Energy and Sustainable Development
Department of Economics
Stanford University
wolak@zia.stanford.edu

March 25, 2013

Preliminary Draft
Please Do Not Distribute

Abstract

This paper reports on the results of two field experiments that examine the impact of providing current information on a household's electricity consuming actions on how that household responds to a nonlinear retail price schedule for electricity. Across the two utilities, over 2,000 households participated in a customized on-line interactive educational program that taught them how their monthly electricity bill was determined from nonlinear retail pricing scheme they face. Each household was also told where their typical consumption monthly places it on this nonlinear pricing schedule. Households were also shown how changes in their major electricity-consuming activities would affect their monthly bill under the nonlinear pricing scheme. Using data from before and after this intervention for households that took the educational program (our treatment) and a randomly selected set of control households, we estimate the overall treatment effect associated with our educational program as well as a treatment effect for households on each specific pricing tier on the nonlinear price schedule during the pre-intervention period. For both utilities, we find that the overall impact of our treatment is a reduction in the household's daily average consumption. In addition, our price tier-specific treatment effect results are that households that learn they face a high marginal price for consuming electricity reduce their electricity consumption and households that learn they face a low marginal price increase their electricity consumption. These results emphasize that the need to provide timely and actionable information to households in order to maximize the effectiveness of nonlinear retail pricing schemes.

1. Introduction

Virtually all residential electricity customers in the United States face an increasing block tariff (IBT) pricing structure where the price paid for an additional unit of consumption, what is typically called the marginal price, varies with the amount of electricity consumed within the month. The IBT pricing schedule resembles a staircase with height of each step equal to the constant marginal price and the length equal to range of monthly consumption that the marginal price applies to.

An IBT pricing structure allow consumers to purchase their “necessary” monthly electricity consumption at a low marginal price and then charges consumers progressively higher marginal prices for more discretionary electricity uses. Under an IBT, electricity consumers that use more KWh during the month face a higher marginal price, and therefore should have stronger incentive to reduce their monthly consumption relative to consumers that use less KWh during the month and therefore face a lower marginal price.

There are a number of necessary conditions for IBT pricing to provide these incentives. First, consumers must know how their monthly electricity bills are determined from an IBT. Second, consumers must determine where on the IBT schedule they are likely to end the month in order to determine the appropriate marginal price to use for their electricity consumption decisions during the month. Third, and perhaps most important, consumers need to understand how their electricity-using actions translate into KWhs of electricity consumption, because depending on where the customer is on the nonlinear price schedule the cost of a given electricity consuming action can differ by more than a factor of two.

Electricity demand is derived from a consumer’s demand for air conditioning, space heating, water heating, lighting, television watching, computer use, and electricity-consuming

appliance services in general. Although most households are able to predict the dollar amount of their typical summer and winter monthly electricity bills, there is considerable debate among industry observers whether: (1) households understand their how electricity bills are determined from an IBT, (2) what their marginal price for the month is likely to be, and (3) how their appliance-using actions translate into KWhs of consumption and dollars on their monthly electricity bill.

For example, suppose a household did not know how much electricity it was likely to consume in the month. Even if these households understood how their monthly bill was computed under IBT pricing, they would not know the marginal price of electricity for that month and would therefore be forced to use less efficient approaches to managing their monthly electricity bill. Shin (1985) argues that lack of information about the monthly marginal price (the marginal price for the last KWh consumed in the month), may lead households to adopt rules of thumb for determining their monthly electricity consumption. Shin postulates that consumers respond to the monthly average price—the household’s monthly electricity bill divided by their monthly electricity consumption—and he presents empirical evidence consistent with this hypothesis. Borenstein (2009) argues that lack of information about the monthly marginal price, because the households cannot accurately forecast their monthly electricity consumption, causes them not to respond to the monthly marginal price. Ito (2012) proposes an empirical test of whether households respond to the monthly marginal price versus the monthly average price and finds evidence consistent with the hypothesis that households respond to the monthly average price.

Possible explanations for the empirical results of Shin (1985) and Ito (2012) are that households are unaware of their monthly marginal price or they do not understand how their

monthly electricity bill is determined from an IBT or how their electricity-consuming actions translate into dollars on their monthly electricity bill. By providing households with this information, one would expect them to make more efficient electricity consumption choices, rather than resort to ad hoc decisions rules for determining their monthly electricity consumption.

This study employs a randomized field experiment to study whether households that are provided with provided with information that educates them about the monthly marginal price they face choose to change their behavior. We provide customers in our treatment group with: (1) information about how their electricity bill is determined from an IBT pricing function, (2) the value of their typical monthly marginal price, and (3) information about how their appliance-using actions translate into dollars on their monthly electricity bill. Our hypothesis is that households can respond to IBT pricing structures and other more complex retail pricing structures such as hourly retail prices that vary with hourly wholesale prices, but they lack the information that would allow them to do so. Our on-line educational treatment provides the basic information that allows households to respond to an IBT pricing structure.

Starting in the Spring of 2011, we partnered with two California electric utilities and implemented a field experiment with households from each utility that resulted in roughly 2,000 households taking an our Internet educational course that provided with them with the minimum information necessary to respond to an IBT pricing structure. At each of the two electric utilities, we were provided with a random sample of customers who live in single-family homes. We randomly assigned the vast majority of these households to the intent-to-treat (ITT) group and the remaining households to the control group. All households in the ITT group received an encouragement letter or e-mail inviting them to complete our treatment, a 30-minute Internet educational course that provided the household with: (1) information about how its electricity

bill is determined from its IBT pricing function, (2) the value of its typical monthly marginal price of electricity, and (3) information about how the household's appliance-using actions translate into dollars on their monthly electricity bill.

After completing the Internet course, we tracked the changes in electricity consumption for the ITT group, which is composed of households that were offered the opportunity to take our educational course but refused (the offered-but-refused-treatment group) and households that were offered the opportunity to take the course and actually took it, and the control group. Besides estimating the treatment effect of our intervention across all households, we also specify a measurement framework that allows the treatment effect of our education program on electricity consumption to differ by the level of the household's typical monthly marginal price. This econometric approach is justified by an appeal to the static theory of household utility-maximizing choice subject to an IBT which implies that the level of this marginal price should impact a household's electricity consumption choices.

For both utilities, the overall the average treatment effect for our educational program is negative, implying that on-average households that took our treatment group reduced their consumption relative to households in our sample that did not take the treatment. We also estimate pricing-tier-specific treatment effects regressions for our educational program. Consistent with the view that households armed with information about their typical monthly marginal price will use this information to use electricity more efficiently, we find that households that learn they face a high marginal price for electricity *reduce* their consumption while households that learn that they face a low marginal price for electricity *increase* their consumption.

This result is consistent with the view that the households that do not complete our treatment (those in the control group or offered-but-refused group) are using rules-of-thumb based on their average price of electricity to determine their monthly electricity consumption. Those that complete our on-line educational course and are provided with information that their marginal price is less than this average price decide to consume more and those that find their marginal price is above this average price decide to consume less. If our educational treatment did not provide the treatment households with new information about how their electricity-using actions impact their monthly electricity bill, then we would expect to see no change in behavior in our treatment versus control and treatment versus offered-but-refused comparisons.

A unique feature of our study is that we ran the same experiment at two independent locations and obtained qualitatively the same results at both locations, which raises our confidence in the results that we report. There are also crucial differences across the two utilities in terms of their IBT pricing structures, climate conditions and the type of data we can access so that we learn more from combining insights from two experiments than if we had run our experiment at only one site.

This paper adds to an emerging literature on educating households about non-linear incentives such as the tax code and evaluating the impact of providing this information on household behavior. Chetty and Saez (2013) assess the impact of tax preparers giving simple, personalized information to a random sample of their clients about the Earned Income Tax Credit (EITC) on the subsequent earnings of these clients. Liebman and Luttmer (2011) use a field experiment design document how information about the Social Security system impacts older women's labor force participation.

This literature and our field experiments assess the impact of lowering the cost of information acquisition of economic behavior. In our case, the on-line educational course lowers the cost to the household of learning the information that will allow it to respond to an IBT for electricity and in the case of Chetty and Saez (2013) their treatment lowers the cost to the taxpayer of learning about the incentives to work created by the EITC. In both cases the information is randomly assigned to individuals and then a statistical model is used to test for a behavioral change in response to the provision of this information. Because both experiments assume that the only reason for the differences in the behavior of the treatment and control groups is due to the information provided, therefore it is reasonable to assume that any change in behavior between the treatment and control groups is caused by the information provided. In our case, making households aware of their monthly marginal price, how their monthly electricity bill is determined under an IBT, and how their monthly bill varies with changes in their electricity-consuming activities appears to cause changes in their monthly electricity consumption in a manner consistent with these households making beneficial use of this information.

Although our results are directly relevant to the impact of information provision on the behavior of households subject to IBT pricing structures, it does not seem to be a stretch to extend them to case of marginal prices that differ over time. Specifically, our results suggest that notifying households that they face an hourly price that is higher than their annual average retail price will cause them to reduce their consumption and notifying households that they face an hourly price that is lower than their annual average retail price will cause them to increase their consumption. This logic is also consistent with the results of Wolak (2010), which uses an experiment to study how residential consumers respond to retail electricity prices that vary with

hourly wholesale prices. Wolak (2010) presents evidence of substantial demand reductions in response to high hourly retail prices and documents that even low-income households are adept at responding to the higher hourly prices. These results highlight how appropriate price signals combined with timely information provision can play a key role in helping California achieve its ambitions AB32 goals.

2. Impact of Marginal Price Information

To understand how notifying households of their typical monthly marginal price is likely to impact their electricity consumption consider the following two models of utility-maximizing household-level electricity consumption behavior. The first assumes linear pricing where the household faces the same marginal price for units consumed and nonlinear or IBT pricing where the household pays higher marginal prices for higher levels of monthly consumption. Specifically, we assume that the household faces an IBT for electricity with two pricing tiers. For consumption between zero and E_1 the household faces a marginal price of p_1 . For consumption greater than E_1 the household faces a marginal price of p_2 that is strictly greater than p_1 .

Suppose that household is only told its monthly electricity bill is B dollars and its consumption is E KWhs. Suppose that based on this information, the household concludes that it faces a price of $p_A \equiv B/E$ for electricity. Assume the household has a budget constraint of M dollars and that the only other good available to the household is a composite good X , that has a price of p_X . The household's preferences are assumed to be described by the utility function, $U(E,X)$, which is increasing in each argument.

Consider the following simple model of household electricity consumption subject to a linear price of electricity set equal to the average price faced by the household, p_A . In this case that household's utility maximization problem is:

$$\text{Max}_{\{E,X\}} U(E,X) \text{ subject to } p_X X + p_A E = M$$

If E^* and X^* are the solutions to this problem then they satisfy the following first-order condition: $U_X(E^*, X^*)/p_X = U_E(E^*, X^*)/p_A$, where $U_S(E^*, X^*)$ is the partial derivative of $U(E, X)$ of with respect to S for $S=E, X$ evaluated at (E^*, X^*) .

Suppose that $E^* > E_1$, which implies that $p_1 < p_A < p_2$. Now suppose that the household receives our information treatment and is notified that its monthly marginal price is p_2 and that it in fact faces an IBT pricing structure. In this case, the household utility-maximization problem is:

$$\text{Max}_{\{E, X\}} U(E, X) \text{ subject to } p_X X + p_2 \max((E - E_1), 0) + p_1 \min(E_1, E) = M.$$

If E^+ and X^+ are the solutions to this problem then they satisfy the following first-order condition: $U_X(E^+, X^+)/p_X = U_E(E^+, X^+)/p_2$ because the household has been told that p_2 is its monthly marginal price. Note that because $p_2 > p_A$, $U_E(E^*, X^*)/p_A > U_E(E^*, X^*)/p_2$, which implies that $E^+ < E^*$, if we make the usual assumption that $U_E(E, X)$ is decreasing in E and increasing in X and $U_X(E, X)$ is decreasing in X and increasing in E (the marginal utility of X is decreasing in X and increasing in E). This result implies that telling a household that formerly thought it was maximizing utility subject to a linear price equal to the average price of electricity that it is facing a IBT and has a marginal price of $p_2 > p_A$, will result in that household reducing its electricity consumption. By similar logic, telling the household it faces an IBT and a marginal price of $p_1 < p_A$ will result in that household increasing its electricity consumption.¹ We test these hypotheses about the impact of providing this information using the field experiment that we describe in the next section.

3. Field Experiment Implementation Steps

To conduct this field experiment, we partnered with two California electric utilities. For confidentiality reasons, we never accessed data that identified the household's name or street address. However, by partnering with these electric utilities we were able to access household-level data on electricity consumption before and after our intervention took place for a large

¹ Note that even with only two pricing tiers, p_A , is greater than p_1 because households must pay a number of small monthly fixed charges that make the average price of households on the first pricing tiers greater than p_1 .

random sample of households from each utility with a subsample of the households from each utility that have taken our educational program.

We now describe how we implemented this field experiment and recruited a subject pool and provide a summary of the educational course's content. Our two field experiments were accomplished in several steps. First, we chose a random subset of two electric utility's residential customers. Second, we randomly assigned these households to the intent-to-treat group and the control group. Third, we randomized a participation payment amount offer to each member of the intent-to-treat group. Fourth, we invited each member of the intent-to-treat group to take our on-line education course with the promise of the randomly assigned payment amount in an Amazon.com gift code if they completed the course. Fifth, a subset of households in the intent-to-treat group chose to complete our on-line course to become members of our treatment group. At both experimental sites the enrolled subjects took the same educational treatment tailored to the specific IBT the household faced. Sixth, we tracked the post-treatment electricity consumption for the intent-to-treat group (the treatment group and the offered-but-refused group) and the control group. These panel data allow us to estimate the econometric models we present below. We now provide details about each of these steps.

Utility A Experiment Specifics

Utility A provided a third party with the confidential data for 1,407,500 single family home owners in its service area.² This third-party had to be provided with the name and street address of each of these home owners because it sent out the invitation letters that we describe below. These households were all single-family homeowners who faced the most common residential rate structures used by Utility A.

² The sample provided by Utility A includes single family residences, domestic tariff only, continuously on a domestic rate for one year, no tariff riders such as an air conditioning cycling program and no reduced payment tariffs that were income based or voluntary.

This third party was employed to guarantee data confidentiality. They created a unique household identifier and provided data to us only using this identifier. This procedure guaranteed the privacy of all households in our treatment group and control group. From the population of Utility A single family residential customers, 25,000 households were drawn at random. The control group received no correspondence from us.

Each member of the intent-to-treat group was randomly assigned ones of the following Amazon Gift Card Amounts: \$10, \$20, \$25, \$30, \$35, and \$50. The third-party incorporated this information in the invitation letter that was sent out to each member of the treatment group. We use this household-specific randomized payment amount as an instrument for the decision of a household to take our treatment to compute a local average treatment effect for our educational program.

The third-party mailed out 7,500 invitation letters on July 29th 2011, an additional 3,000 letters on August 17th 2011 and a final batch of 2,500 letters on September 23rd 2011. Below, we will present our statistics on who chose to participate and how this propensity varies with the size of the Amazon Gift Card offered.

The subset of households that took our treatment logged into an Internet website and used a household-specific ID (rather than their name) and password to login. Their household ID and password were provided in their invitation letter. Once participating individuals logged in, they received customized information about their electricity consumption that we describe below and they answered a set of survey questions. Figure 1 gives a sample screenshot from our education program showing the IBT for an inland customer of Utility A. The grey shaded area under the curve shows how much the customer must pay for amount of energy shown on the horizontal axis under this IBT.

Upon finishing the Internet Education Program, each household was e-mailed the promised Amazon gift card and a printable file that could be posted in a prominent place in the household containing a set of tips for reducing the household's electricity bills customized based on the household's answers to a number of survey questions.³ At this point, those in the treatment group did not interact with us again.

Utility B Specifics

While the general content and intent of the Internet Education Program was identical at both sites, there are some important differences in how we implemented the experiment at Utility B relative to the Utility A. In choosing the treatment group and control group, we selected a random sample of single-family homeowners that had an electronic account with Utility B. Such households receive their communications from the utility via e-mail rather than through the United States Postal Service. According to Utility B, roughly 20% of residential customers have electronic accounts. We were told by Utility B that these households are a bit younger and more ethnically diverse than the utility's overall service population. We focused our experiment on this population of utility B households because of cost considerations. The marginal cost of sending an e-mail asking a household to take our treatment is virtually zero versus almost a dollar per letter by regular mail.

Our Utility B experiment population is composed of all electronic account customers that had an interval meter that records the households' hourly consumption installed as of the Autumn of 2011. At that time roughly 50% of the utility's meters had been installed. The interval meters were installed by geographic territories within the utility's service area and the roll out of these meters was completed by April 2012. We recognize that these selection rules

³ The third-party handled the processing of the Amazon Gift Cards so at no point did we have access to treated households' e-mail addresses. The third-party also e-mailed the customized electricity conservation tips.

mean that we do not have a random sample of single-family homeowners in the service area. Instead, our sample is randomly drawn from the population of households that had electronic accounts and whose geographic area had smart meters in place as of Autumn 2011.

We followed the same steps at Utility B that we did at Utility A to assign households randomly to be in our sample and then randomly assigned this group to the intent-to-treat group and the control group. For those assigned to the treatment group, we randomly assigned Amazon Gift Card payments of \$0, \$10, 20, \$20, \$30, \$40 and \$50 to all treatment group customers for completing the course. A \$0 Amazon amount meant that we sent a solicitation e-mail asking them to take the course without any promise of a financial payment.

On October 18th 2011, we launched the Utility B experiment by sending out the solicitation e-mails with randomly assigned Amazon amounts, including a solicitation letter with no promise of payment. On October 26th 2011 another 4,500 emails were sent out to those who did not open the first email. After a household completed our on-line education program, the third party emailed the household their Amazon Gift card and the customized electricity consumption tips. The third party worked with Utility B to provide us with anonymous household-level electricity consumption data with an experiment identifier assigned to each household in the intent-to-treat group and control group.

Table 1 gives the details of the experiment design for the Utility A and Utility B experiments. The breakdown for the intent-to-treat, treatment, and control groups reflects that fact that data errors after receiving the final data from Utility A and Utility B required deleting a number of households from each of the three groups. Table 2 gives a breakdown of the treatment acceptance rates for each Amazon Gift card amount for each experiment. For both

experiments, the null hypothesis of an acceptance probability that is monotone increasing in the Amazon Gift card amount offered cannot be rejected.

4. Valid Randomization and Potential Endogenous Selection

This section analyzes pre-intervention data on monthly consumption from Utility A and monthly consumption from Utility B aimed at demonstrating: (1) our procedure for randomizing customers into the intent-to-treat (ITT) group and the control group is statistically valid and (2) that the pre-intervention distribution of consumption for households in the ITT sample that took the treatment is no different from the distribution for those that did not take the treatment.

Our analysis relies on the two-sample Kolmogorov-Smirnov test of equality of two distributions. Specifically, let Q_{im}^k equal the average daily consumption of household i of type k during month m , where k denotes one of four groups: ITT, control, treatment, and ITT-declined (the household is in the ITT group but declined to take the treatment). We express all monthly consumption in terms of average daily consumption during that month to account for the fact that there are different numbers of days during different months of the year.

Suppose there are M months during the pre-intervention period, which ends during the month that first letter was sent out in the case of Utility A and the month the first e-mail was sent out in the case of Utility B. Suppose there are N_k households in group k . Define the empirical distribution of average daily consumption in month m for group k as:

$$F_m^k(t) = \frac{1}{N_k} \sum_{i=1}^{N_k} I(Q_{im}^k \leq t),$$

where $I(X \leq t)$ is an indicator variable that takes on the value 1 if X is less than or equal to t and zero otherwise. Under the assumption that the Q_{im}^k , $i=1,2,\dots,N_k$ are independent and

identically distributed within month m for each group k with population distribution equal to $G_m^k(t)$, we can perform the hypothesis test:

$$H: G_m^k(t) = G_m^h(t) \text{ versus } K: G_m^k(t) \neq G_m^h(t),$$

using the two-sample Kolmogorov-Smirnov statistic

$$KS = \sup_t |F_m^k(t) - F_m^h(t)|.$$

Table 3A reports the KS statistic and the associated probability value (p-value) for month m for the test that the population distribution of average daily consumption for month m for the control group is equal to the population distribution of average daily consumption for month m for the ITT group for Utility A. The p-value gives the probability of obtaining a draw from the null asymptotic distribution of the test statistic greater than the realized value of the KS statistic. The table also gives the number of observations in both the ITT and control groups. For the months January to July of 2011, the probability value is never less than 0.05, indicating that a size $\alpha = 0.05$ test of the null hypothesis would not be rejected for any of these months during the pre-sample period.

Table 3B reports the KS statistic and the associated probability value (p-value) for month m for the test that the population distribution of average daily consumption for month m for the treatment group is equal to the population distribution of average daily consumption for month m for the subset of the ITT group that declined treatment for Utility A. The table also gives the number of observations in both the treatment and ITT-but-declined-treatment groups. For the months January to July of 2011, the probability value is never less than 0.05, indicating that the null hypothesis would not be rejected for any of these months during the pre-sample period.

These results provide no evidence against the null hypothesis that the average daily consumption distributions for each month of the pre-intervention period for the ITT sample and

control sample are equal and no evidence against the null hypothesis that the average daily consumption distributions of the treatment and ITT-but-declined-treatment groups are equal during any of the months of the pre-intervention period. This result is evidence consistent with the hypothesis that households did not self-select into the treatment group based on their pre-intervention consumption levels.

Table 4A repeats the analysis in Table 3A for Utility B for the months of May to September of 2011. Because of metering errors and other data recording problems throughout our sample period for Utility B, there were different numbers of households in the control and ITT groups for each month. Because the data we obtained from Utility B was during its transition from monthly metering to interval metering for the households in our sample, there were many measurement error and data recording problems that were never resolved. These problems did not arise for Utility A, because we had access to billing quality data for the entire sample period. Nevertheless, the results in Table 4A provide no evidence against the null hypothesis that the pre-intervention distributions of average daily consumption for the control group and the ITT group are equal for all of the months from May to September 2011.

Table 4B repeats the analysis in Table 3B for Utility B for the months of May to September 2011. For the reasons described above, there are different numbers of households in the treatment group and the ITT-but-declined-treatment groups each month. Different from the case of Utility A, there is substantial evidence against the null hypothesis of equality of the two distributions for the majority of months of the pre-intervention period. However, we are unable to determine if data errors is the reason for the p-values below 0.05 for a number of the months, or if the two monthly distributions are in fact different.

Taken as whole, the results of this section suggest that our randomization of households into the ITT and control groups is statistically valid for both Utility A and Utility B, the pre-invention distributions of average daily consumption for the treatment and ITT-but-declined-treatment groups were the same for all months during this time period for Utility A, and that the KS statistics provide evidence against this null hypothesis for Utility B, but it is unclear if this result is being driven by data issues or failure of null hypothesis to hold.

The next section summarizes the basic features of our educational program and the three key pieces of information that it conveys to households that should cause them to become more efficient consumers of electricity.

5. Internet Educational Treatment Specifics

The basic goal of the Internet education course was to familiarize the household with the three pieces of information described earlier: (1) the household's typical monthly marginal price, (2) how the household's bill was determined from its IBT, and (3) how the household's monthly electricity bill changes in response to changes in how the household uses its electricity-consuming appliances.

The survey was organized in three sections. Section 1 demonstrates how the IBT works and where the customer's consumption for the month is typically located on the schedule during the three months before each survey was administered. This section also shows how the marginal price of electricity and monthly bill changes depending on how much electricity the household consumes in the month. Section 2 surveys the household about the characteristics of its home and the appliances in it. Section 3 uses the answers provided in Section 2 to determine how changes in the utilization of these appliances would impact the household's monthly electricity bill and provides recommendations tailored the household's appliance stock to assist them in becoming more sophisticated electricity consumers.

Differences in IBT Pricing Between the Two Utilities

Figure 2 displays the IBT for Utility A for the inland and coastal portion of its service area. This is the dominant rate structure for single family home owners in its service territory, the lowest marginal price is 12 cents per kWh and the highest is 31 cents/KWh. The only difference between the IBT for the coastal and inland areas for Utility A is the length of each price step. Because of the more extreme temperatures in the inland the region of Utility A's service territory, the length of each pricing step in the IBT is longer for the inland versus coastal region.

Households in the intent-to-treat group that took our Internet course were presented with their IBT schedule and shown the locational of their typical month's usage in KWh and the dollar cost of this typical monthly consumption. The graphic in the on-line education program would demonstrate that the monthly bill was the area under the IBT up to the household's typical monthly consumption. The program also gave participants access to a slider that allowed them to conduct their own thought experiments to see how their monthly bill would change if they increased or decreased their monthly consumption.

Figure 2 demonstrates that the two electricity retailers in our study differ sharply with respect to the steepness of their pricing tiers. Utility A has five pricing tiers with marginal prices that differ by almost a factor of three from the lowest to highest-priced tier. As shown in Figure 2, the other retailer, Utility B, only has two pricing tiers, and the top tier is slightly more than 50 percent higher than the price on the first tier. These differences in the IBT between the two utilities allows us to implement a more robust test of whether information on a household's typical monthly marginal prices yields the anticipated behavioral response described in Section 2.

Data

This section discusses the data the utilities have provided to us. For both utilities, our unit of analysis is a household/month. For Utility A, we have billing cycle level-data that we converted to average daily electricity consumption data for each month of our sample starting in July 2010 and running through June 2012. For those assigned to the treatment group, we create a dummy variable that equals one if the month takes place after treatment and is the fraction of the month after treatment for the month in which the treatment took place. For both utilities assumed and all households in the treatment group, we assumed that treatment took place day the first letter was send for Utility A or the first e-mail was sent for utility B. For Utility A, we use data for 6 months before the treatment and 6 months after the treatment.

Although we had hourly billing cycle-level data available for Utility B, for comparability with the results for Utility A, we converted this data to average daily values for each month of our sample period, similar to how the data used for Utility A. Different from the Utility A setup, we use this average daily consumption monthly data for 5 months before and after the first e-mail was sent.

For both Utility A and Utility B, we also re-ran our analysis with the treatment dates for households in the treatment group set equal to the actual date they took the treatment rather that the first day a letter was mailed for Utility A or e-mail was sent for Utility B and obtained very similar results to the ones reported below.

6. Econometric Modeling Framework and Empirical Results

This section presents our treatment effect and tier-specific treatment effects econometric modeling framework that we use estimate the impact of our informational intervention on the

household's monthly electricity consumption for each utility. We then present results of our model estimation for both Utility A and Utility B for the control and ITT (treatment plus ITT-but-declined-treatment) sample. We obtain very similar results to those reported below when we use the treatment plus the control sample and exclude the ITT-but-declined-treatment households from the sample. However, we report treatment versus control and ITT-but-declined-treatment ordinary least squares results below so that they can be compared to the local average treatment effect (LATE) estimates that we present using the ITT indicator variable as instrument for the Treatment indicator variable. The LATE estimator can only be computed for the combined ITT and control sample.

Let Q_{im} equal the average daily electricity consumption in month m for household i . This data is constructed from the customer-level billing cycle-level data by computing an average daily consumption for each billing cycle and then taking a weighted average of the average daily consumption of each billing cycle in month i , where the weights are the shares of the days of the month i associated with each billing cycle. For example, if a month has 30 days and 20 of them are in one billing cycle and the remaining 10 are in another, then the weights are $2/3$ and $1/3$. Let $Treat(i)$ equal 1 if household i is in the treatment group and 0 for all other households.

Rather than attempt to model how a household's consumption of electricity varies across months in our sample and our estimated treatment effects rely on this assumption, we instead decided to estimate all of our treatment effects off of the cross-section of differences between average daily monthly consumption after the intervention versus before the intervention. Define Q_i^{pre} as the mean average daily consumption for all months during the pre-intervention time period for household i . Define Q_i^{post} as the mean average daily consumption for all months during the post-intervention time period for household i . Define $Y_i = Q_i^{post} - Q_i^{pre}$, which is the

difference between average daily consumption during the post-intervention period minus the average daily consumption during the pre-intervention period for household i .

In terms of this notation our overall treatment effect model is:

$$Y_i = \alpha + \beta * Treat(i) + \varepsilon_i,$$

where β is the overall average treatment effect. Our tier-specific treatment effects model for tier j assumes:

$$Y_{ij} = \alpha_j + \beta_j * Treat(i) + \varepsilon_{ij},$$

where β_j is the treatment effect for customers typically on tier j . We estimate β_j by restricting our sample to customers whose typical pre-intervention consumption was on tier j . The pre-intervention typical consumption price tier is computed for all household in both the control and ITT samples, but only customers in the treatment group are told this information through our educational program. We estimate the same two models for Utility B. The only difference is that Utility has only two pricing tiers, so we estimate two tier-specific treatment effects.

We also estimate these equations by instrumental variables using the instrument ITT_i which equals 1 if household i is in the ITT sample—it receive a letter asking it to take the treatment in the case of Utility A or an e-mail asking it to take the treatment in the case of Utility B. ITT_i equals zero if household i is in the control group and therefore did not receive a letter in the case of Utility A or an e-mail in the case of Utility B. We estimate both the overall treatment effect equation and the tier-specific equations using ITT_i as an instrument for $Treat(i)$. Using the full set of indicator variables for the Amazon gift card amounts—6 indicator variables for Utility A and 7 indicator variables for Utility B--yields parameters estimates for the overall and tier-specific local average treatment effects that are very similar to the ones presented below for the single indicator variable ITT_i for both Utilities.

Because our treatment effects model is simply estimating the difference in means of Y_i , the difference in the post-intervention versus pre-intervention average consumption more informative graphical presentation the results is possible. Figure 3 plots the histogram on Y_i for the treatment sample and the histogram of Y_i for the control and ITT-but-declined-treatment sample. The distribution of Y_i for the treatment sample is almost a uniformly negative shift across all percentiles of the distribution of Y_i for the control and ITT-but-declined-treatment sample. The regression result in Table 5-1 in demonstrates that the average treatment effect associated with our educational program is precisely estimated to be -0.723 KWh per day.

Figures 4-1 to 4.5 plot the histograms on Y_i for the treatment sample and the histogram of Y_i for the control and ITT-but-declined-treatment sample for each of the five pricing tiers. Consistent the logic of Section 2, for the first pricing tier, distribution of pre- versus post-intervention daily-average consumption difference is a slight positive shift of relative to distribution for the control and ITT-but-declined-treatment sample. Table 5-1 confirms that the tier-specific treatment effect of our educational program for the first tier is 0.36 KWh per day. For the remaining 4 tiers, the treatment distribution of Y_i is an increasingly negative shift relative to the relative to the distribution of Y_i for the control and ITT-but-declined-treatment sample. Table 5-1 shows that these tiers specific treatment effects are all precisely estimated to be negative and range from -0.461 KWh per day for Tier 3 to -1.201 KWh per day for Tier 5.

Table 5-2 reports the local average treatment effect estimate and tier-specific local average treatment effect estimates using the indicator variable ITT_i as an instrument for $Treat(i)$. The instrumental variable estimates are significantly larger in absolute value that the corresponding estimates reported in Table 5-2. It is important to bear in mind that that sample the overall and tier-specific local average treatment effect can be written as:

$$\text{LATE} = [E(Y_i | \text{ITT}_i=1) - (E(Y_i | \text{ITT}_i=0))]/E(\text{Treat}(i)=1 | \text{ITT}_i = 1),$$

for the sample analogues to these population conditional expectations. This expression points to two possible explanations for the significantly larger in absolute value estimate of the local average treatment effects in Table 5-2 versus Table 5-1. First, although the difference between the mean of Y_i for the IIT sample versus the control sample is negative and roughly the same order of magnitude as the difference between the mean of Y_i for the treatment sample versus the control plus ITT-but-declined-treatment sample, from Table 2 $E(\text{Treat}(i)=1 | \text{ITT}_i = 1) = 0.12$ for the entire sample, so that this difference is divided by 0.12 to produce the LATE. Second, the fact that a household received our letter asking them to take our education program may have caused them to reduce their consumption of electricity, because the letter (reproduced in Appendix A) emphasizes that the educational program will help them reduce their electricity consumption. This difference in parameter estimates emphasizes the well-known result that the instrumental variables estimate and ordinary least squares estimates of a slope coefficient rarely are consistent for the same population magnitude.

Figure 5-plots the density of Y_i for the treatment group and the control plus ITT-but-declined-treatment samples for Utility B. Again, the treatment density appears to be a negative shift of the control plus ITT-but-declined-treatment density. The coefficient estimate in Table 6-1 confirms that the treatment effect of our educational program for Utility A is -0.328 KWh per day, although this point estimate is not precisely estimated. Figures 6-1 and 6- 2 plot the tier-specific density comparisons for Utility B. Consistent with the logic in Section 2, the tier 1 treatment density for Y_i appears to be a positive shift of the control plus ITT-but-declined-treatment density and the tier 2 treatment density for Y_i is a larger negative shift. The parameter estimates in Table 6-1 find somewhat precisely estimated tier 1 treatment effect of our

educational program of 0.670 KWh per day and more precisely estimated tier 2 treatment effect of -1.922 KWh per day.

Similar to the case of Utility A, the local average treatment effects shown in Table 6-2 are larger in absolute value than the corresponding coefficient estimates reported in Table 6-1. However, of the LATE coefficient estimates in Table 6-2 are not estimated with sufficient precise to draw any firm conclusions about their magnitude.

Although the coefficients for Utility B are not as precisely estimated as those for Utility A, the signs are consistent with the view that our educational program provides useful marginal price information that allows the household to become a more sophisticated consumer. Although in the aggregate, households reduced their electricity consumption in response to this marginal price information, our results suggest that households on the lowest marginal price tier responded to this information by increasing their consumption. As discussed in Section 2, if households are using rules of thumb to determine their monthly electricity consumption based on an average price computed from their total monthly bill and total electricity consumed, providing marginal price information is likely to have this response.

7. Conclusion

In co-operation with two California electric utilities, we have designed and implemented an educational field experiment to quantify how increased knowledge about nonlinear pricing and how the household's appliance use translates into electricity use impacts the household's electricity consumption.

Our experiment demonstrated that providing information to household about the IBT they face and where they are on that schedule can help electric utilities reduce residential electricity

consumption and to reduce greenhouse gas emissions. The average daily consumption of households for both utilities during our sample periods for both utilities is approximately 25 KWh. Our overall treatment effect of -0.724 KWh per day for Utility A is approximately 3 percent of average daily consumption, whereas the overall treatment effect for Utility B is -0.328 KWh day which is approximately 1.5 percent of average daily consumption.

Our findings emphasize that the timely provision of actionable information to final electricity consumers is a crucial to smart grid technology delivering economic benefits to electricity consumers. The “smart grid” has been touted as mechanism for reducing electricity bills, increasing energy efficiency, and reducing greenhouse gas (GHG) emissions from the electricity sector. Residential electricity consumption accounts for roughly 33% of total electricity consumption in the United States (US) and the electricity sector produces approximately 40% of US GHG emissions. Consequently, if smart grid technologies can produce modest reductions in household-level electricity consumption, particularly during periods of peak electricity demand when GHG-emissions-intensive generation units must be relied on to produce electricity, this can yield in tangible reductions in US GHG emissions.

That consumers are responsive to marginal prices when informed through our treatment supports the claim that by accompanying the universal deployment of interval metering at California’s three large investor-owned utilities with dynamic retail pricing and the timely provision of actionable information to households on these pricing plans will induce significant behavioral changes that reduce both annual and peak-period electricity consumption, which would also reduce California’s annual GHG emissions. Households with interval meters that are also provided with information on the hourly retail price electricity and their real-time

electricity consumption and understand how electricity-using actions translate into dollars on their monthly electricity bill and can make more cost-effective appliance utilization decisions.

References

- Ayers, Ian, Sophie Raseman, and Alice Shih (2009). "Evidence from Two Large Field Experiments that Peer Comparison Feedback Can Reduce Residential Energy Usage." National Bureau of Economic Research Working Paper 15386.
- Borenstein, Severin. 2009. "To What Electricity Price Do Consumers Respond? Residential Demand Elasticity Under Increasing-Block Pricing."
- Carter, D. W and J. W Milon. 2005. "Price knowledge in household demand for utility services." *Land Economics* 81 (2):265.
- Chetty, R., A. Looney, and K. Kroft. 2009. "Salience and taxation: Theory and evidence." *The American Economic Review* 99 (4):1145–1177.
- Chetty, Raj and Emmanuel Saez. 2009. "Teaching the Tax Code: Earnings Responses to an Experiment with EITC Recipients." National Bureau of Economic Research Working Paper Series No. 14836.
- Cragg, Michael, YuYu Zhou and Kevin Gurney and Matthew E. Kahn "Carbon Geography: The Political Economy of Congressional Support for Legislation Intended to Mitigate Greenhouse Gas Production." *Economic Inquiry*, forthcoming.
- de Bartolome, Charles A. M. 1995. "Which tax rate do people use: Average or marginal?" *Journal of Public Economics* 56 (1):79–96
- Fujii, Edwin T. and Clifford B. Hawley. 1988. "On the Accuracy of Tax Perceptions." *The Review of Economics and Statistics* 70 (2):344–347.
- Heckman JJ, Urzua S, Vytlacil E. Understanding instrumental variables in models with essential heterogeneity. *Review of Economics and Statistics* 2006; 88(3): 389-432.
- Ito, Koichiro. "Do Consumers Respond to Marginal or Average Price? Evidence from Nonlinear Electricity Pricing." Energy Institute at Haas Working Paper 210 (2012).
- Liebman, Jeffrey B., and Erzo FP Luttmer. *Would people behave differently if they better understood social security? Evidence from a field experiment*. No. w17287. National Bureau of Economic Research, 2011.
- Reiss, Peter and Matthew White (2008). "What Changes Energy Consumption? Prices and Public Pressure." *Rand Journal of Economics*, 39, 636-663.
- Shin, Jeong-Shik. (1985) "Perception of Price When Price Information Is Costly: Evidence from Residential Electricity Demand." *The Review of Economics and Statistics* 67 (4):591–98.

Thaler, Richard H. and Cass R. Sunstein (2008). *Nudge: Improving Decisions about Health, Wealth, and Happiness*. New Haven, CT: Yale University Press.

Varian, Hal R. (1992) *Microeconomic Analysis (Third Edition)*, W.W. Norton & Company.

Wolak, Frank A. (2010) "An Experimental Comparison of Critical Peak and Hourly Pricing: The PowerCentsDC Program,, available at <http://www.stanford.edu/~wolak> .

Table 1: Experiment Design for Utility A and Utility B

Experimental Design Utility A	
Intent-to-Treat Group Size (Invited)	12,273 customers
Treatment Group Size (Invited and Accepted)	1,227 customers
Control Group Size	10,964 customers
1 st Letter Sent	August 1, 2011
2 nd Letter Sent	August 18, 2011
3 rd Letter Sent	September 26, 2011
Period of Analysis	2/1/2011 to 1/31/2012 (monthly data)

Experimental Design Utility B	
Intent-to-Treat Group Size (Invited)	5,715
Treatment Group Size (Invited and Accepted)	785
Control Group Size	1,000
1 st E-mail Sent	October 18, 2011
2 nd Letter Sent	October 26, 2011
Period of Analysis	5/1/2011 to 2/29/2012 (monthly data)

Table 2: Treatment Response Rates for Utility A and B

Utility A Gift Card Amounts and Response Rates				
Amazon Amount (\$)	Number Offered	Cards Sent	Sent/Offered %	Standard Error %
0	1800	108	6.00%	0.56%
10	1500	181	12.10%	0.84%
20	1100	142	12.90%	1.01%
30	900	150	16.70%	1.24%
40	700	121	17.30%	1.43%
50	500	83	16.60%	1.66%
Overall	6500	785	12.10%	0.40%

Utility B Gift Card Amounts and Response Rates				
Amazon Amount (\$)	Number Offered	Cards Sent	Sent/Offered %	Standard Error %
10	700	34	4.90%	0.82%
20	3437	283	8.20%	0.47%
25	3161	291	9.20%	0.51%
30	3053	298	9.70%	0.54%
35	2949	299	10.10%	0.55%
50	200	22	11.00%	2.21%
Overall	13500	1227	9.10%	0.25%

Table 3: Kolomogorov-Smirnov Test of Equality of Pre-Invention Distributions of Average Daily Consumption by Month for Utility A

Table 3A: Test of Equality of Monthly ITT and Control Distributions During Pre-Invention Period for Utility A		
Number in Control Group = 11,500 Number in ITT Group = 13,500		
Month of 2011	KS Statistic	P-Value
January	0.0137	0.193
February	0.0172	0.051
March	0.0131	0.236
April	0.0170	0.052
May	0.0152	0.115
June	0.0149	0.126
July	0.0167	0.061
Table 3B: Test of Equality of Monthly Treatment and ITT-But-Declined Treatment Distributions During Pre-Intervention Period for Utility A		
Number in Treatment Group = 1,227 Number in ITT-But-Declined Treatment Group = 12,273		
Month of 2011	KS Statistic	P-Value
January	0.0238	0.551
February	0.0303	0.256
March	0.0296	0.283
April	0.0265	0.415
May	0.0318	0.208
June	0.0195	0.789
July	0.0317	0.211

Table 4: Kolomogorov-Smirnov Test of Equality of Pre-Invention Distributions of Average Daily Consumption by Month for Utility B

Table 4A: Test of Equality of Monthly ITT and Control Distributions During Pre-Invention Period for Utility B				
Month of 2011	KS Statistic	P-Value	Number in Control	Number in ITT
May	0.0232	0.742	993	6,461
June	0.0234	0.732	996	6,474
July	0.0175	0.968	914	5,986
August	0.0223	0.823	916	5,997
September	0.0284	0.542	917	5,992
Table 4B: Test of Equality of Monthly Treatment and ITT-But-Declined Treatment Distributions During Pre-Intervention Period for Utility B				
Month of 2011	KS Statistic	P-Value	Number in Treatment	Number in ITT-But-Declined
May	0.0531	0.038	783	5,678
June	0.0455	0.110	782	5,692
July	0.0551	0.033	747	5,239
August	0.0581	0.005	747	5,250
September	0.0541	0.042	747	5,245

**Table 5-1: Overall and Tier-Specific Treatment Effects
Estimates for Utility A**

Overall Treatment Effect		
Variable	Estimate	Standard Error
Treat	-0.723	0.082
Constant	-0.668	0.019
Tier 1 Treatment Effect		
Treat	0.136	0.060
Constant	-0.184	0.040
Tier 2 Treatment Effect		
Treat	-0.483	0.175
Constant	-0.291	0.046
Tier 3 Treatment Effect		
Treat	-0.461	0.134
Constant	-0.528	0.032
Tier 4 Treatment Effect		
Treat	-0.912	0.155
Constant	-0.786	0.378
Tier 5 Treatment Effect		
Treat	-1.201	0.220
Constant	-1.200	0.052

Table 5-2: Overall and Tier-Specific Local Average Treatment Effects Estimates for Utility A Using ITT Indicator as Instrument

Overall Treatment Effect		
Variable	Estimate	Standard Error
Treat	-12.77	0.516
Constant	-0.621	0.025
Tier 1 Treatment Effect		
Treat	1.136	0.051
Constant	-0.184	0.056
Tier 2 Treatment Effect		
Treat	-5.32	0.974
Constant	-0.043	0.060
Tier 3 Treatment Effect		
Treat	-9.84	0.842
Constant	-0.072	0.042
Tier 4 Treatment Effect		
Treat	-14.39	1.024
Constant	-0.070	0.051
Tier 5 Treatment Effect		
Treat	-22.31	1.635
Constant	-0.080	0.071

Table 6-1: Overall and Tier-Specific Average Treatment Effects Estimates for Utility B

Overall Treatment Effect		
Variable	Estimate	Standard Error
Treat	-0.328	0.282
Constant	-4.830	0.084
Tier 1 Treatment Effect		
Treat	0.670	0.271
Constant	-3.077	0.086
Tier 2 Treatment Effect		
Treat	-1.922	0.419
Constant	-6.276	0.128

Table 6-2: Overall and Tier-Specific Local Average Treatment Effects Estimates for Utility B Using ITT Indicator as Instrument

Overall Treatment Effect		
Variable	Estimate	Standard Error
Treat	-1.712	1.834
Constant	-4.676	0.218
Tier 1 Treatment Effect		
Treat	1.509	1.710
Constant	-3.183	0.234
Tier 2 Treatment Effect		
Treat	-5.122	3.200
Constant	-5.963	0.336

Figure 1: IBT for Utility A Inland Households

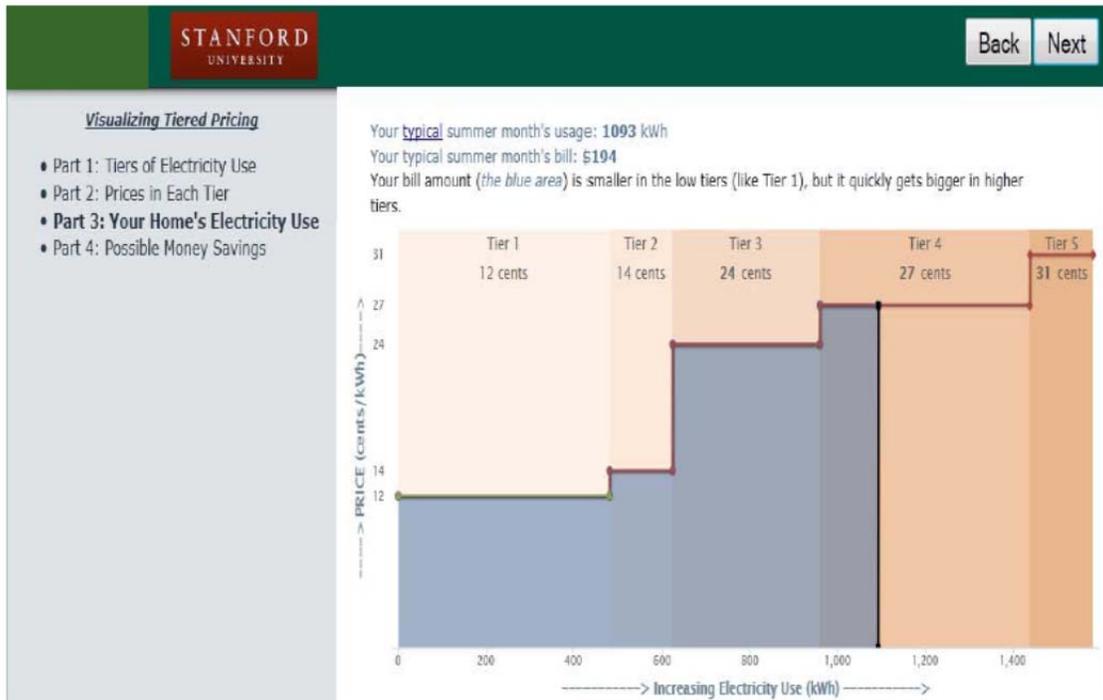


Figure 2: IBTs for Utility A and Utility B

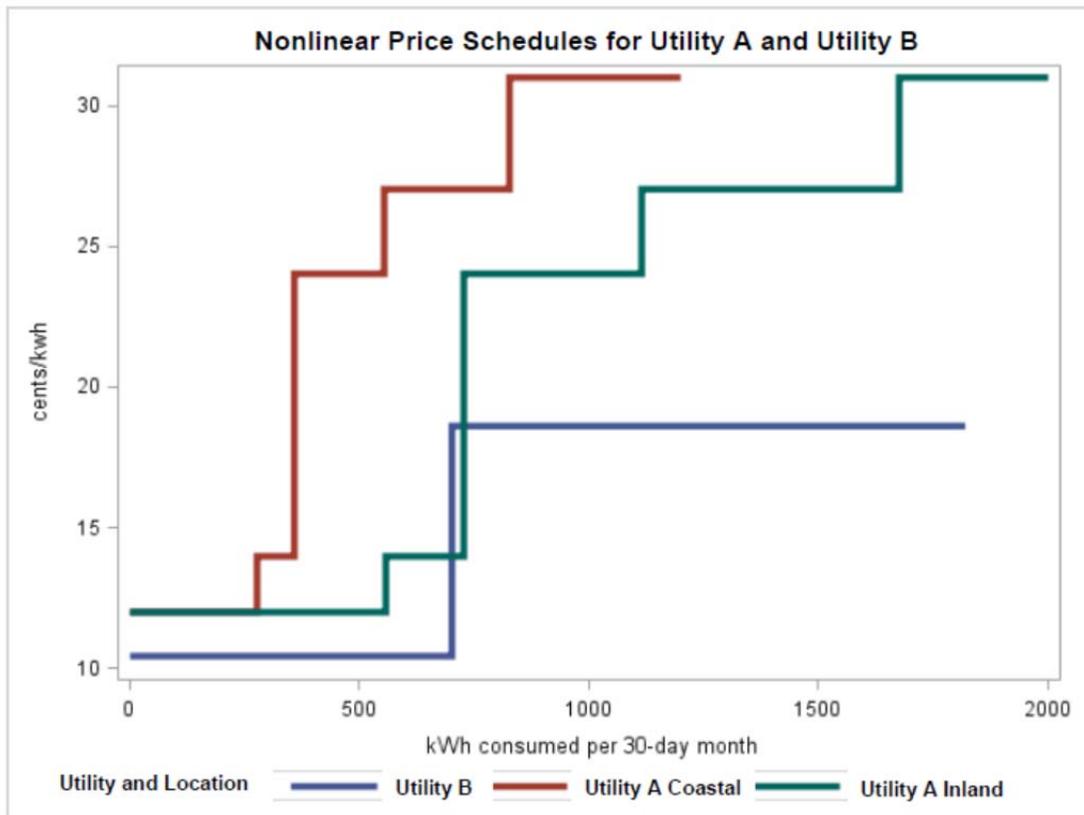


Figure 3: Densities of Post- Minus Pre-Intervention Average Consumption for Treatment and Control plus ITT-But-Declined Treatment Samples for Utility A

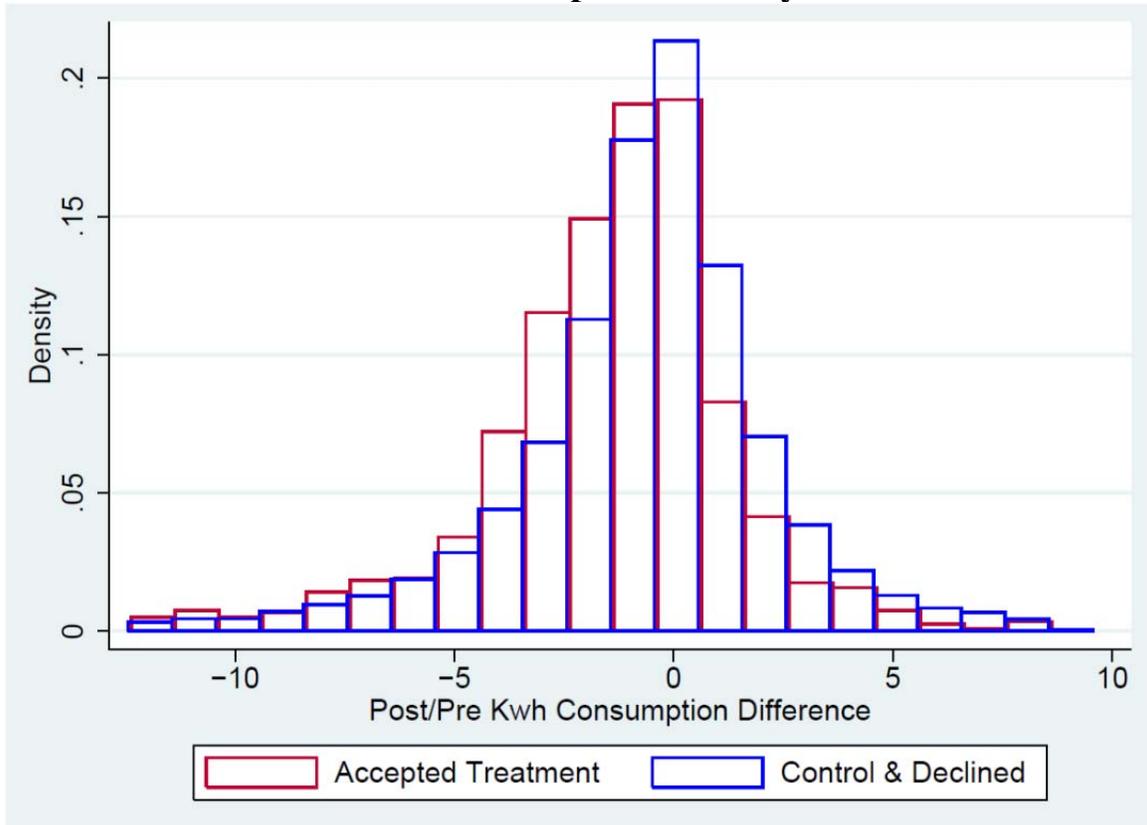


Figure 4-1: Densities of Post- Minus Pre-Intervention Average Consumption for Tier 1 Treatment and Control plus ITT-But-Declined Treatment Samples for Utility A

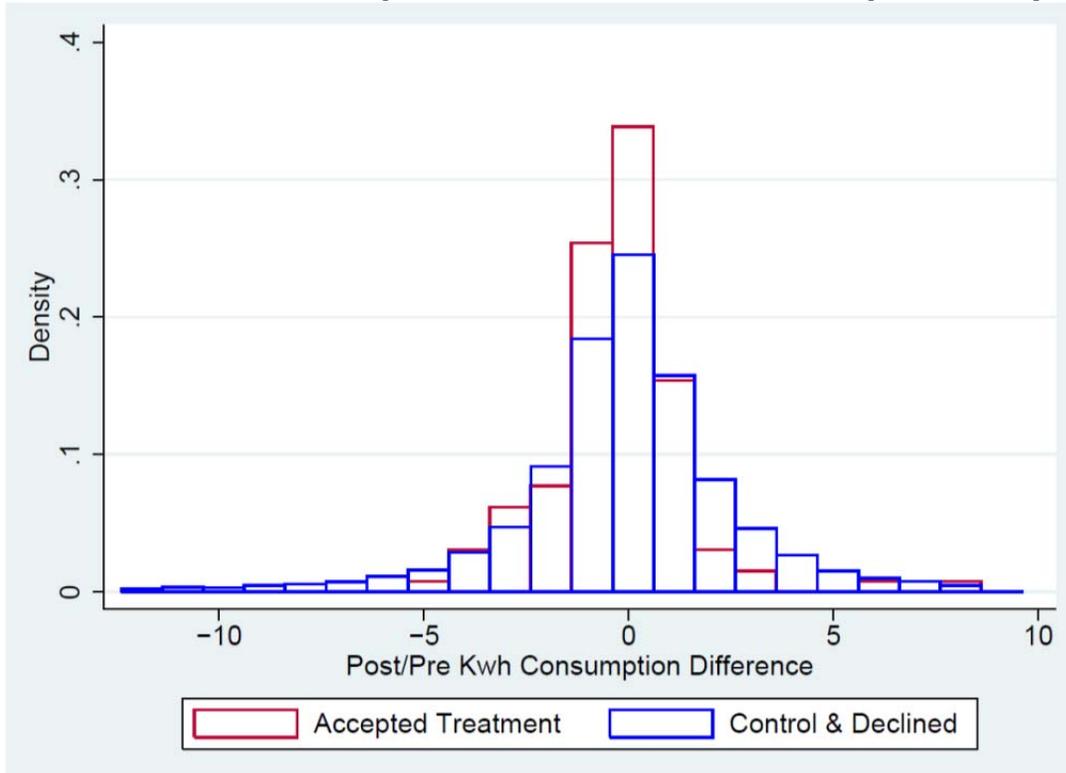


Figure 4-2: Densities of Post- Minus Pre-Intervention Average Consumption for Tier 2 Treatment and Control plus ITT-But-Declined Treatment Samples for Utility A

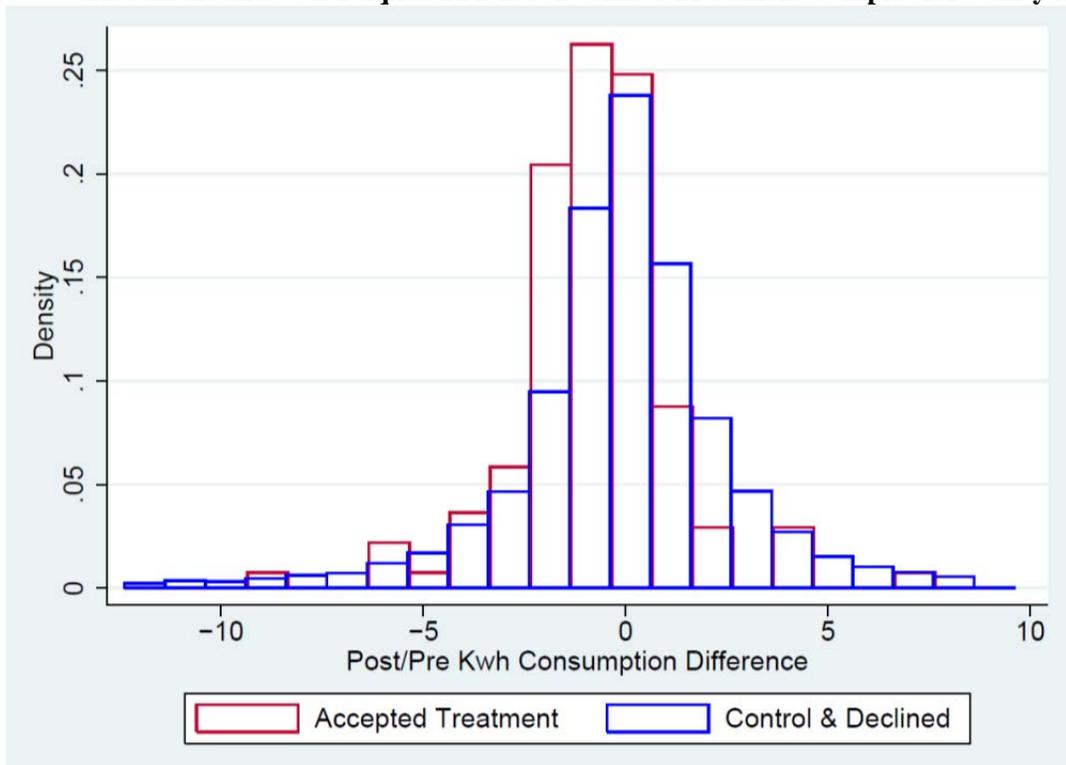


Figure 4-3: Densities of Post- Minus Pre-Intervention Average Consumption for Tier 3 Treatment and Control plus ITT-But-Declined Treatment Samples for Utility A0

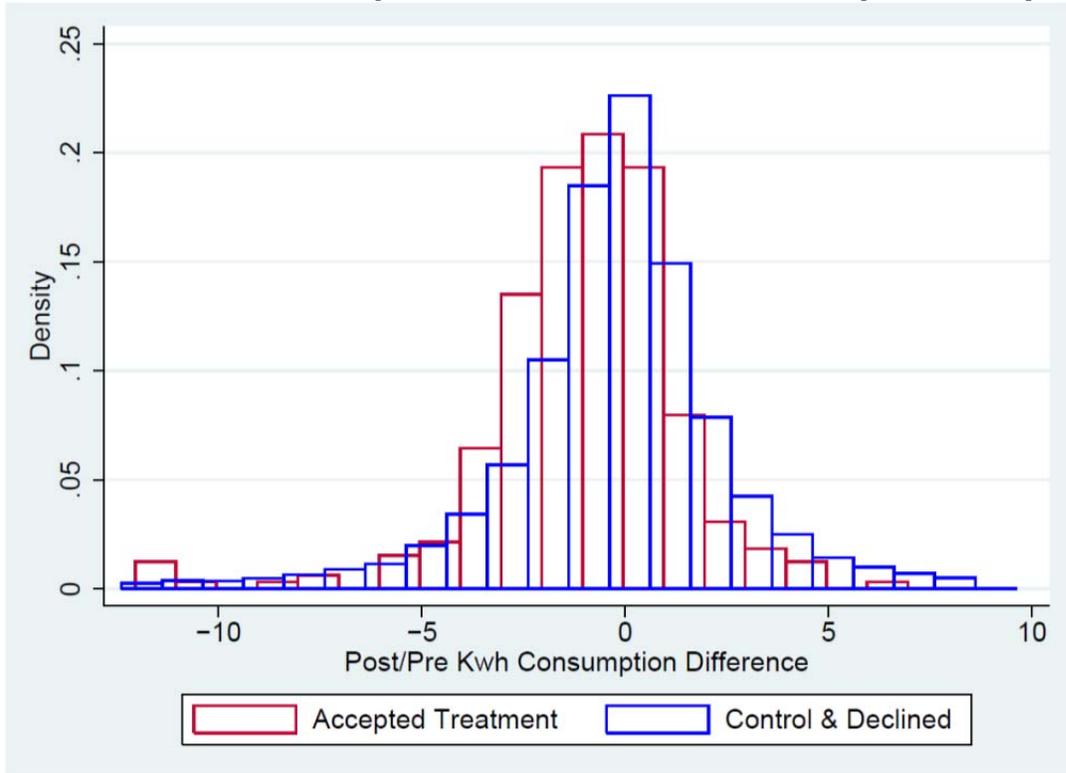


Figure 4-4: Densities of Post- Minus Pre-Intervention Average Consumption for Tier 4 Treatment and Control plus ITT-But-Declined Treatment Samples for Utility A

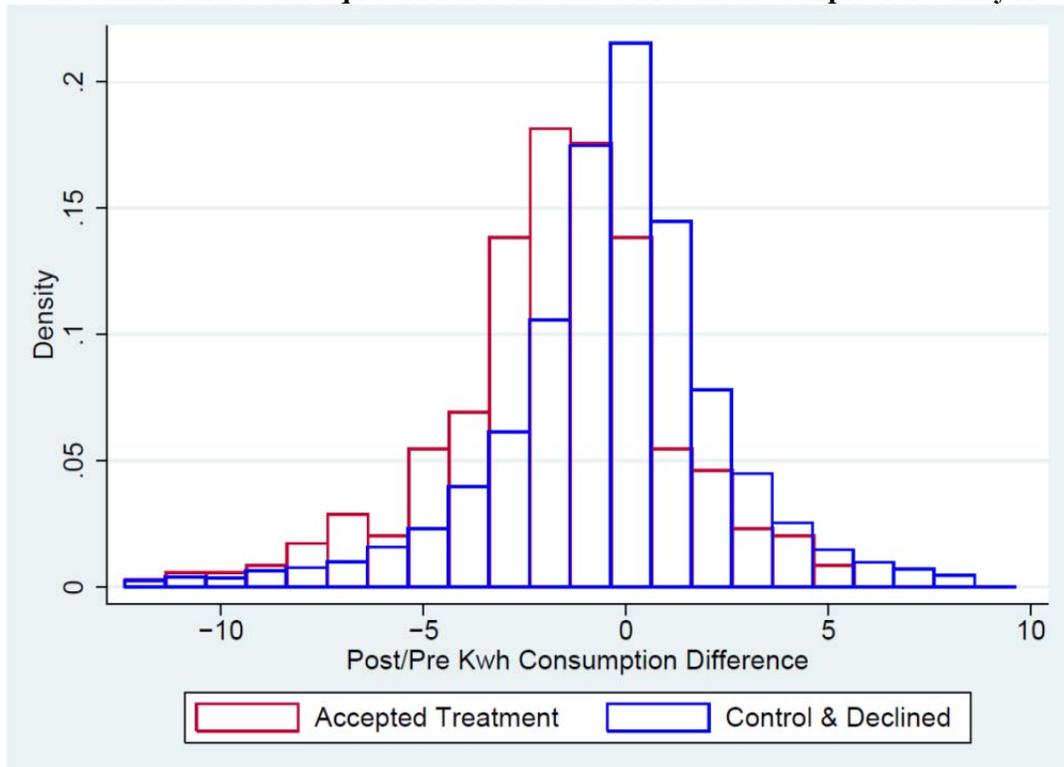


Figure 4-5: Densities of Post- Minus Pre-Intervention Average Consumption for Tier 5 Treatment and Control plus ITT-But-Declined Treatment Samples for Utility A

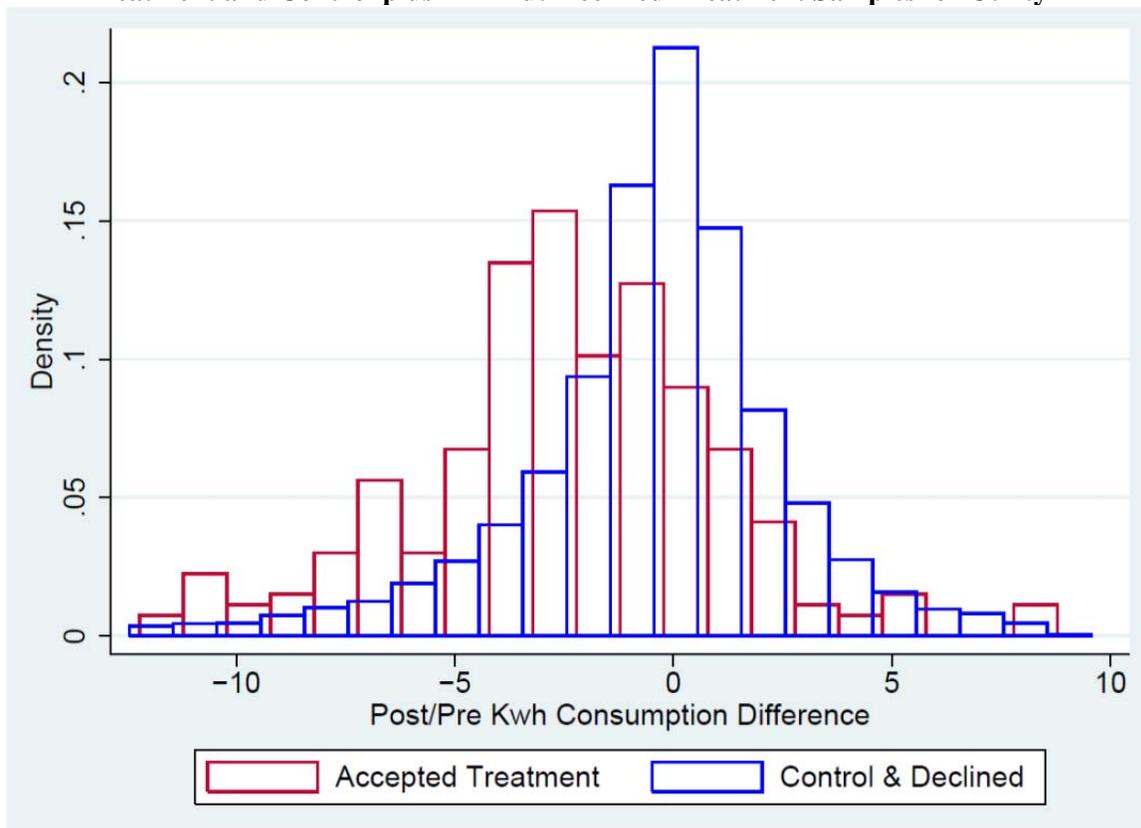


Figure 5: Densities of Post- Minus Pre-Intervention Average Consumption for Treatment and Control plus ITT-But-Declined Treatment Samples for Utility B

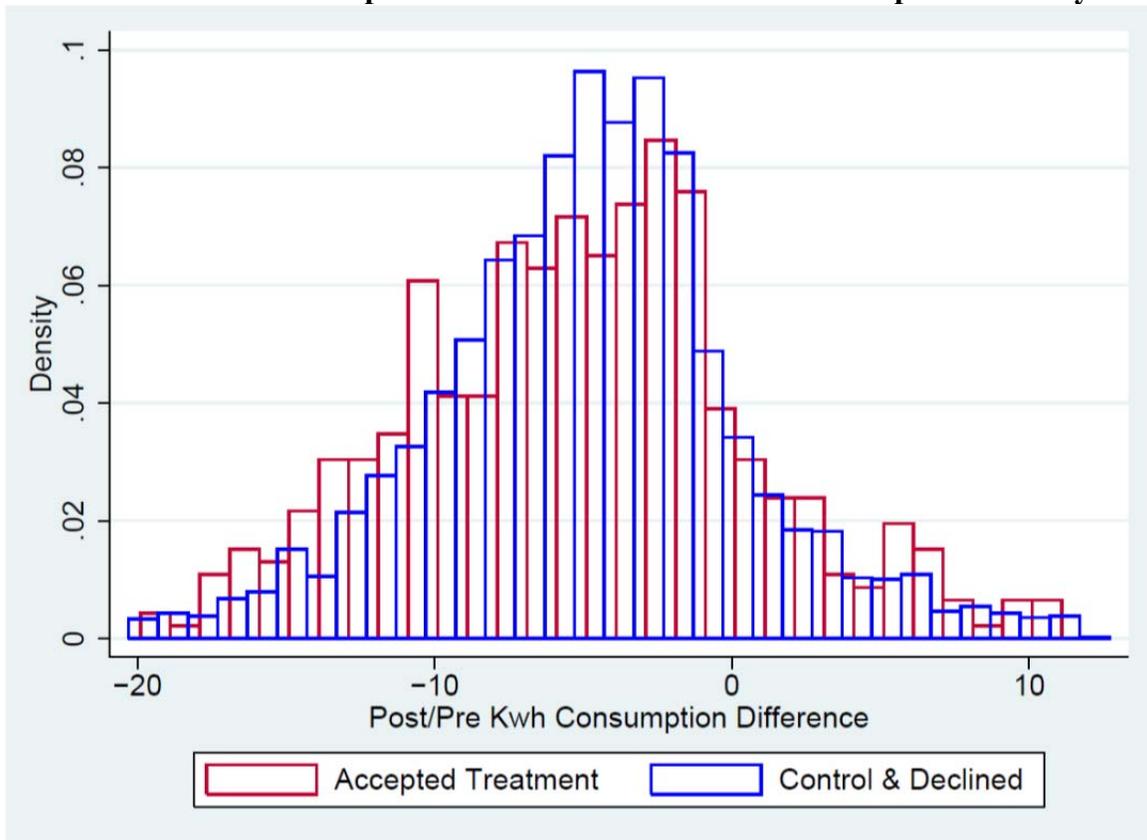


Figure 6-1: Densities of Post- Minus Pre-Intervention Average Consumption for Tier 1 Treatment and Control plus ITT-But-Declined Treatment Samples for Utility B

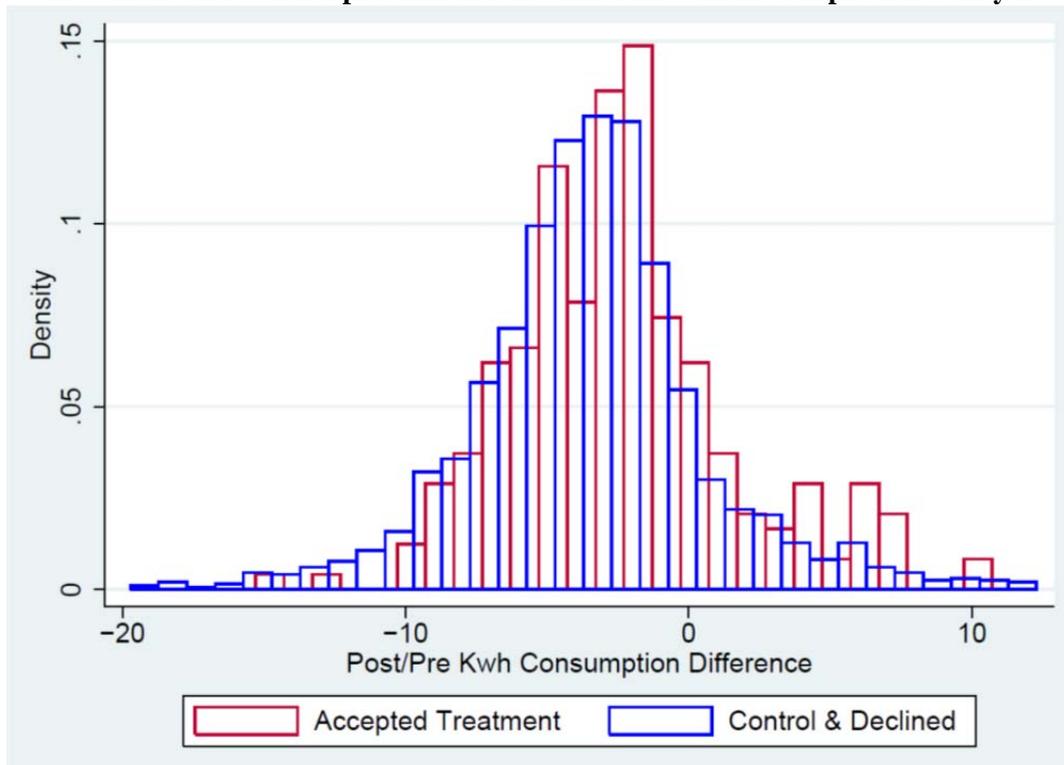
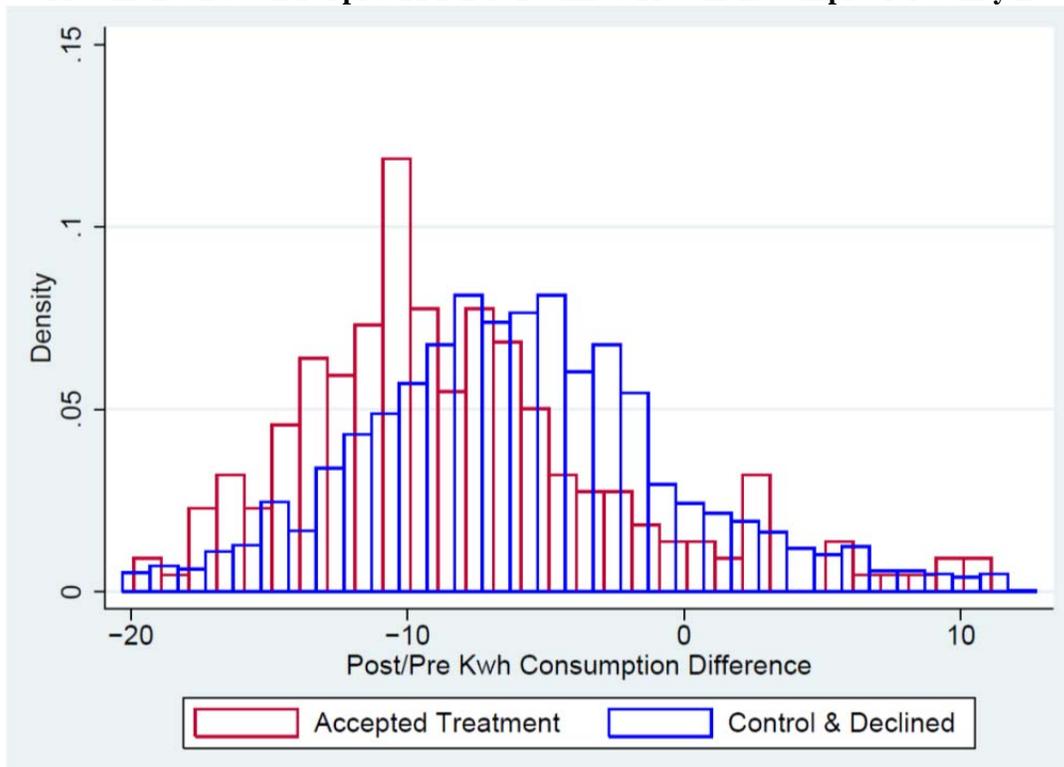


Figure 6-2: Densities of Post- Minus Pre-Intervention Average Consumption for Tier 2 Treatment and Control plus ITT-But-Declined Treatment Samples for Utility B



Appendix A

Dear [insert name],

Utility A is partnering with researchers from Stanford University and UCLA to develop a home energy savings workshop. Your valuable input will help Utility A create similar educational tools in the future. This online workshop is a 15-20 minute tutorial that could help you save money on your next electricity bill. For completing the workshop, you will receive a **\$50 gift card to Amazon.com**.

The workshop starts by showing how your electricity use affects your electricity bill. Then, using a brief survey of your household's characteristics, the workshop generates the customized suggestions you can use to reduce your household's electricity bills.

We hope you will try this innovative program today. To begin the workshop simply click the link below or paste it into your browser. If you have any questions, please contact me at the number below.

[insert link]

Sincerely,

Mr. Smith
Utility A

P.S. You will receive your Amazon Gift Card by e-mail within 10 days of completing the survey. If you do not receive the card within this time period, please check your spam filter.