Does Conservation Targeting Work? Evidence from a Statewide Electricity Rebate Program in California

Koichiro Ito*
Stanford University

December, 2012

Abstract

Regulators often use conservation targeting rebate programs to promote energy conservation. Consumers participating in such programs receive a financial reward if they achieve a targeted level of conservation. The effectiveness of the programs, however, remains highly controversial because non-experimental data rarely provide reliable estimates of the treatment effect. This study examines the cost-effectiveness of conservation rebate programs by using a regression discontinuity design in the California 20/20 rebate program. In 2005, California residents received a 20% discount on their summer electricity bills if they saved electricity by 20% relative to their consumption in 2004. The program's enrollment scheme prevented self-selection and created a discontinuity of treatment status. Using household-level monthly billing records from the three largest California electric utilities, I find heterogeneous responses to the rebate incentive. In the areas where summer temperature is persistently high and income-level is relatively low, the incentive reduced electricity consumption by 5% to 10%. In the other areas, however, the treatment effect is essentially zero. I show that the cost-effectiveness is very poor in these areas because many consumers still received a rebate without conservation efforts on their part. To save 1 kWh of electricity, the program cost 2 cents in inland areas, 91 cents in coastal areas, and 14.8 cents on average for all service areas in California.

*SIEPR Postdoctoral Fellow, Stanford Institute for Economic Policy Research, Stanford University. Email: koichiro.ito@stanford.edu. I am grateful to Severin Borenstein and Michael Hanemann for their support and advice, and to Michael Anderson, Maximilian Auffhammer, Peter Berck, James Bushnell, Howard Chong, Lucas Davis, Meredith Fowlie, Catie Hausman, Larry Karp, Erica Myers, Karen Notsund, Hideyuki Nakagawa, Carla Peterman, Catherine Wolfram and seminar participants at UC Berkeley for their helpful comments. I also thank the California Public Utility Commission, Pacific Gas & Electric, Southern California Edison, and San Diego Gas & Electric for providing residential electricity data for this study. Financial support from the Joseph L. Fisher Doctoral Dissertation Fellowship by Resources for the Future and from the California Energy Commission is gratefully acknowledged.

1 Introduction

Engineering studies often suggest that improving residential energy efficiency is the least expensive way to abate global greenhouse gas emissions.[1] However, policy makers generally believe that it is difficult to find a policy that can effectively change households' electricity consumption behavior. The price elasticity of residential electricity demand is relatively inelastic, and it is often politically infeasible to introduce an electricity price that is high enough to achieve a substantial consumption reduction.

A potential solution to this problem is to provide explicit financial incentives to save electricity. As an explicit financial incentive, many electric utilities offer rebate programs. Consumers receive a rebate for purchasing efficient appliances, or weatherizing homes. Also, electric utilities often offer "conservation rebate programs" that provide a rebate for rewarding a reduction in consumption achieved during a certain time period. For example, the California state government offered "California 20/20 electricity rebate program", which provided a 20% discount on electricity bills as a financial reward for reducing electricity consumption by 20% relative to the previous year. Such conservation rebate programs allow consumers to choose how they will reduce consumption and is therefore more flexible than other rebate programs that require a purchase of specific appliances. The cost-effectiveness of these programs, however, remains controversial. There is little evidence that consumers save electricity in response to the economic incentives created by these rebate programs. Another concern is that households may receive rebates simply because of the natural year-to-year fluctuation in their electricity consumption rather than concerted efforts to conserve.

[1]For example the abatement cost curve of greenhouse gas emissions by Naucler and Enkvist (2009) indicates that the abatement cost in the residential electricity sector can be negative in the sense that improving energy efficiency at home would reduce greenhouse gas emissions and household expenditure for electricity.

This study aims to measure the treatment effect and cost-effectiveness of such conservation price-rebate programs by applying a regression discontinuity design to the California 20/20 rebate program in 2005. In the summer of 2005, most California households could receive a 20% discount on their electricity bills if they reduced their electricity consumption by 20% relative to their consumption in the summer of 2004. Nearly all California residents were enrolled in the program. However, those who started their electricity service after a certain cutoff date in 2004 were ineligible to participate. The electric utilities that offered this program strictly enforced this eligibility rule, and therefore, excluded non-eligible households from the program. Importantly, it was impossible for households to anticipate the program in advance and thus they could not strategically choose their account open date for the rebate program. Consequently, the eligibility rule excluded self-selection, and generated essentially random assignment of the program among households who opened their account near the cutoff date. I apply a regression discontinuity design to this discontinuous eligibility cutoff date to estimate the treatment effect and cost-effectiveness of the rebate program on electricity conservation.

My empirical analysis relies on a panel data set of household-level monthly electricity billing records for nearly all households in the three largest investor-owned electric utilities in California. This confidential data set was directly provided by the three electric utilities, Pacific Gas & Electric, Southern California Edison (SCE), and San Diego Gas & Electric (SDG&E). The data set includes detailed information about each customer's monthly bills in 2004 and 2005. Importantly, the data set allows me to identify the exact start and close dates of each household's electricity service so that I can conduct regression discontinuity estimation based on these dates. I combine this billing data set with weather information from the National Oceanic and Atmospheric Administration's (NOAA's) National Climate Data Center (NCDC) and demographic information from the US Census 2000 to further explore how weather conditions and income levels affect the treatment effect of the program.

Using these three sets of data, I find the following results for the effect of the 2005 California 20/20 rebate program on residential electricity consumption. First, estimates from the regression discontinuity estimation provide evidence that the rebate incentive reduced electricity consumption by 5% to 10% in the areas where the summer temperature is persistently high and the income level is relatively low. In contrast, the treatment effects are nearly zero in the areas where the summer temperature is moderate and the income level is relatively high. Second, to explore the cause of this heterogeneous treatment effect, I estimate interaction effects between the treatment variable and climate conditions, and between the treatment variable and income levels. Results from these regressions suggest that the treatment effect increases by .15 percent as average temperatures increase by 1 °F and decreases by .027 percent as income levels increase by 1%. Finally, using the estimates of the program's treatment effects, I calculate the cost and benefit of the program. Results from this exercise suggest that the program cost 91 cents to save 1 kWh of electricity in the coastal areas and 2 cents to save 1 kWh of electricity in the inland areas. In the state level, the cost per kWh reduction was 14.8 cents.

The estimated cost of reducing consumption, 14.8 cents per kWh, is larger than previous estimates reported by the electric utilities. This is partly because estimates from the electric utilities usually attribute all of the reduction in consumption by rebated customers to the presence of the rebate program. In previous studies, Reiss and White (2003) estimate the cost and benefit of the 2001 rebate program find that the average cost from June to September for San Diego Gas & Electric was 18 cents per kWh. Inc (2006) uses survey results to estimate the cost and benefit of the 2005 rebate program. Their estimate ranges from 29 cents per kWh to $1 per kWh. An important finding in my study is that I find that the cost-effectiveness is substantially different between the coastal and inland areas in California.

The results from this study provide several policy implications for the California 20/20 electricity rebate program. First, under the current rebating scheme, the expense of natural year-to-year fluctuations in electricity consumption is substantial. As a result, providing a rebate for reductions that would have happened in the absence of the program can be very costly unless the treatment effect is sufficiently large. Second, the estimation results suggest that it is important to account for heterogeneous treatment effects particularly based on different weather conditions and income levels among households. For example, my cost-effectiveness estimates for the coastal areas are by far larger than previous estimates while my estimates for the inland areas are far lower than previous estimates. Finally, the heterogeneous treatment effect suggests that the program's performance could be improved if the program focused on certain types of households to minimize rebate expenses for reductions that would have occurred in the absence of the program.

The paper proceeds as follows. Section 2 presents the background and research design. Section 3 describes the data. Section 4 presents the empirical framework. Section 5 presents the results, and Section 6 provides conclusions and future research avenues.

2 Background and Research Design

This section provides the institutional background and the research design of this study. First, I describe a brief history of the California 20/20 electricity rebate program. Second, I discuss the evidence and challenges of existing studies. Finally, I present this study's regression discontinuity design.

2.1 California 20/20 Electricity Rebate Programs

The California 20/20 electricity rebate program originates from the initial rebate program ordered by California Governor Gray Davis in 2001 during the California electricity crisis.[2] The California Public Utility Commission (CPUC) expected that a continuous electricity shortage was likely to cause rolling blackouts. To prevent rolling blackouts in the summer of 2001, the CPUC ordered the three largest California investor-owned electric utilities, Pacific Gas and Electric (PG&E), Southern California Edison (SCE) and San Diego Gas & Electric, to provide their customers financial incentives to reduce electricity consumption. In the summer of 2001 and 2002, customers of the three California investor-owned electric utilities received a 20% discount for their June, July, August, and September bill if their monthly consumption was at least 20% lower than the same billing month in 2000. The CPUC ordered the same program in 2005 with a slight change in the scheme. In 2005, the original monthly-based rule was replaced by the whole summer-based rule in 2005 where customers received a 20% discount for their bills over the entire four-month period if they reduced their entire summer consumption by at least 20% relative to 2004.

This conservation rebate program was the largest in scale compared to similar rebate programs that pay households for reducing their consumption. Table 1 shows the scale of the 2005 rebate program for PG&E, SCE, and SDG&E. In 2005, 8% to 9% of customers received a rebate and the total rebate expense for residential customers in these electric utilities was about $25 million. More customers received at least one rebate during the 2001 and 2002 programs because it was not based on consumption over the entire summer but on each billing month.

[2] By August of 2000, wholesale energy prices had more than tripled from the end of 1999, which caused price spikes in retail electricity rates, and financial losses to electric utilities in California. Many cost factors and demand shocks contributed to this rise, but several studies have also found the market power of suppliers to be significant throughout this period. See Joskow (2001), Borenstein, Bushnell, and Wolak (2002), Bushnell and Mansur (2005), Puller (2007), and Reiss and White (2008) for more details on the California electricity crisis.

Reiss and White (2003) report that about 39% of monthly residential bills in SDG&E qualified for a rebate in June, July, August, and September, 2001. For the same 2001 rebate program, Goldman, Barbose, and Eto (2002) note that in the three investor-owned electric utilities, about 33% of their residential customers received a rebate.

Although the CPUC aimed for a substantial reduction in electricity consumption,[3] the effectiveness of the program was highly controversial. The proponents of the program argued that the simplicity of the program makes it straightforward for customers to understand the incentive and will likely encourage energy conservation. The rebate program was also more politically appealing than alternative pricing policies such as an increase in electricity price or an introduction of real-time pricing. In contrast to these alternative policies, the rebate program does not make customers feel a large economic burden even though the program's expenditure will be paid by ratepayers as an increase in electricity price.

The opponents, on the other hand, often questioned the fairness and effectiveness of the program. For example, Faruqui and George (2006) note that the program is politically popular but is likely to be inefficient for energy conservation. The first concern is that the program does not account for weather differences between the base year and target year. Therefore, if the target year turns to be cooler than the base year, many households may receive a rebate simply because of the weather difference. The second concern is that even if there turns out to be no significant weather difference between the two years, many customers will receive a rebate because of random fluctuations in their electricity consumption. For example, customers that had a friend visit in the base year or customers that traveled in the target year can reduce their target year's consumption by 20% compared to their base year without conservation efforts.

[3]For example, in the executive order, CPUC (2001) estimated that the program would help reduce energy consumption by up to 3,500 gigawatt hours in total and up to 2,200 megawatt hours during critical summer peak consumption periods.

Table 2 shows some evidence for the two concerns by the opponents of the program. I use household-level electricity consumption data to calculate what fraction of households reduce their summer electricity consumption more than 20% when there was no rebate programs. I calculate each household's change in consumption from 2003 to 2004 and from 1999 to 2000 in Southern California Edison. Note that the rebate programs were not in effect in any of the four years. From 2003 to 2004, the median household reduced consumption by 1.7% because the summer of 2004 was cooler than 2003. More importantly, 14.3% of households reduced their consumption more than 20%. This statistic suggests that 14.3% of households would have received a rebate without a conservation effort if a rebate program were in effect in 2004. In contrast, the summer of 2000 was warmer than 1999. As a result, the median household increased consumption by 7.7%. However, even in this case, 6.8% of households reduced consumption by 20% or more. Thus, random fluctuations in household electricity consumption creates necessary costs for this rebate program. This issue sometimes leads to a concern for fairness because the program could induce a simple income transfer from one household to others unrelated to their conservation efforts. Moreover, if the rebate expense for these random fluctuations exceeds the program's actual benefit, the cost-effectiveness of the program can be lower than previous estimates.

2.2 Challenges to Estimating the Treatment Effect

To examine the cost-effectiveness of the program, we need a reliable estimate of the treatment effect that is produced solely by the program incentive. The estimation of this treatment effect is, however, generally challenging with non-experimental data. Obviously, it is misleading to make a conclusion simply by looking at the total consumption reduction achieved by the customers that received a rebate. Some of the rebated customers received a rebate not because of their conservation effort. On the other hand, some un-rebated customers may have responded to the program incentive but failed to reach the 20% reduction cutoff to receive a rebate. Therefore, comparing rebated and un-rebated customers does not provide much information about the program's treatment effect. The second challenge is how to control for potential differences between the base and target years that are unrelated to the program but affected electricity consumption. For instance, differences in weather and economic conditions likely affect electricity consumption in the two years. Therefore, changes in electricity consumption between the two years include the program's treatment effect and other confounding factors that are unrelated to the program, and these two effects must be disentangled by researchers to find the treatment effect.

Previous studies acknowledge that it is difficult to estimate the actual treatment effect of the program. Goldman, Barbose, and Eto (2002) is the first study that examines the impact of the original California 20/20 rebate program in 2001. Based on a survey of 400 residential customers, the study finds that 70% of surveyed customers took some active steps to save electricity in 2001, 40% of surveyed customers knew about the program, and 57% of those who took active steps to conserve electricity knew about the program. The study concludes that the cost of purchasing savings through the 20/20 program was about 9 cents per kWh given the assumption that their estimated load reductions are solely attributable to the 20/20 program.

For the same 2001 rebate program, Reiss and White (2008) estimate the treatment effect by using household-level billing data for 70,000 households in SDG&E. The study explores how household-level electricity consumption changes from the years before the California electricity crisis in 2001 to the years after the crisis. Based on the average within-household consumption changes relative to the same month during pre-crisis years, they conclude that the rebate program lowers consumption by approximately 4% to 6%. However, they also note that it is difficult to conclude that this estimate solely reflects the program's treatment effect because there were other conservation programs and public appeals in this period.

Finally, to my knowledge, Wirtshafter Associates (2006) is only the previous study that explores the effect of the 2005 California 20/20 program. The study uses some billing data from the electric utilities and also conduct a survey of 1,177 customers. The study uses the survey results to make two adjustments for estimation: subtract the reduction achieved by the rebated customers that was not due to their conservation efforts; and add the consumption reduction achieved by the non-rebated customers who tried but failed to reach the 20% cutoff. The study concludes that the cost per kWh savings range from 29 cents to $1 per kWh because a substantial level of load reductions may or may not be attributable to the program in their estimation.

A fundamental challenge in the previous studies is that researchers usually do not observe counterfactual groups. Therefore, the previous studies compare consumption between the base and target years of the program and make adjustments for weather differences between the two years. Adjusting for weather differences is difficult without knowing a correct functional form for the effect of weather on household electricity consumption and having detailed weather data. Moreover, as noted by Reiss and White (2008) and Goldman, Barbose, and Eto (2002), it is probably even more challenging to adjust for the effect of other policies such as changes in electricity price, other conservation programs, and public appeals. The next section describes how the current study overcomes these challenges by using a regression discontinuity design for the 2005 rebate program.

2.3 A Regression Discontinuity Design for the 2005 Rebate Program

This paper exploits a discontinuous eligibility rule in the 2005 California 20/20 rebate program to estimate how the rebate program changed household electricity consumption. To be eligible for the 2005 rebate program, customers had to start their electricity service by a certain cutoff date in 2004. Figure 1 illustrates how the eligibility rules were applied to customers. In SCE, for example, the cutoff date was June 5, 2004. Therefore, customers that started their electricity service on or before June 5, 2004 received a notice letter in the spring of 2005 for the 2005 rebate program, whereas customers that started their service after the cutoff date (e.g. June 6, 2004) were not eligible for the program in 2005.

The rule includes two additional key components. First, it was impossible for customers to anticipate the 2005 rebate program when they started their electricity service in 2004 because the program was announced in the spring of 2005. Therefore, it was not possible for customers to strategically choose their start date across the cutoff date of the program. Second, as long as a customer was eligible for the program, the customer automatically participated in the program without having to apply. This automatic participation rule excludes self-selection for the program. The three electric utilities strictly enforced these rules without exception.

This quasi-experimental environment provides the following advantages in estimating the program's treatment effect. The discontinuous eligibility rule generated essentially random assignment of the program among households who started their account near the cutoff date. For example, customers that started their electricity service right before the cutoff date and right after the cutoff date are likely to have similar underlying properties for their electricity consumption, but they were assigned into different groups in terms of the treatment assignment of the rebate program. Even if there is a concern that the underlying properties might be correlated with their service start date, a regression discontinuity design (RDD) can eliminate the bias as long as the correlation between unobservable factors of electricity consumption and service start dates is continuous around the cutoff date for the rebate program.

A potential concern is whether this research design can provide enough observations to have sufficient statistical power to quantify the program treatment effect. In California, about 10,000 customers open an electric account per day. Therefore, there is a large number of observations even if I limit the samples to households that opened an account close to the cutoff date. In addition, because new accounts are generally opened in a wide range of geographical areas in California, the geographical variation allows estimating potential heterogeneous treatment effects in different regions in California. The next section explains the data sets that I use for the analysis.

3 Data

The primary data for this study consist of a panel data set of household-level monthly electricity billing records from 2004 to 2005 in the three largest investor-owned electric utilities in California. Under a confidentiality agreement, Pacific Gas & Electric (PG&E), Southern California Edison (SCE) and San Diego Gas & Electric (SDG&E) provided the complete billing history of essentially all residential customers in their service areas.[4] Each monthly record includes a customer's account ID, premise ID, billing start date and end date, monthly consumption, monthly bill, tariff type, climate zone, and nine-digit zip code. The names of customers and their exact addresses are excluded in the records made available for this study. The billing record also includes each customer's tariff information. In the following analysis, I focus on customers that

[4] A very small number of customers are not individually metered in this area. The data sets include only individually-metered customers.

are on a standard tariff schedule.[5]

Figure 2 shows the service areas of California's electric utilities. PG&E provides gas and electric service for northern California, SCE serves electric customers in most of the southern California areas, and SDG&E provides gas and electric service around the greater San Diego metropolitan area. The 2005 rebate program was applied to customers served by all three electric utilities.

The key variable for the regression discontinuity design of this study is each customer's account open date. The billing records include the exact open and close dates for each customer. Each day in California, about 10,000 customers open an electric account. For my main estimation, I use customers that open their electricity account within 90 days before and 90 days after the cutoff date. The number of households that are on the standard tariff and started their electricity service during these 180 days is 703,903 households for PG&E, 578,362 households for SCE, and 239,168 households for SDG&E.

The billing data do not include customers' exact address or demographic information. To obtain demographic information, I match each customer's nine-digit zip code to a census block group in the 2000 US Census data. The demographic information at the census block group level include median household income, the number of households, and housing characteristics.

Finally, I use daily weather data from the Cooperative Station Dataset published by the National Oceanic and Atmospheric Administration's (NOAA's) National Climate Data Center (NCDC).[6] The data set includes daily minimum and maximum temperature for 370 weather stations in California. First, I match a household's zip code with the nearest weather sta-

[5]About 80% to 85% of customers in each utility are on the utility's standard tariff schedule. The majority of the rest of the customers are on the California Alternative Rates for Electricity (CARE) program, which is a means-tested rate discount program for low income households.

[6]I thank Anin Aroonruengsawat and Maximilian Auffhammer for sharing the data.

tion by following the matching mechanism in Aroonruengsawat and Auffhammer (2009). Second, for each billing cycle, I calculate the cooling degree days (CDD), which is defined as $\sum_{t=S}^{E} Max\,\{\text{Average Temperature (t) - 65, 0}\}$ where S and E are the start and end date of the billing cycle. Figure 3 shows the CDD in one of the August billing cycles in 2005. The coastal areas have small numbers of CDD whereas the inland areas in PG&E and SCE service areas have large numbers of CDD since the summer temperatures are persistently high.

4 Identification and Estimation

This section describes the econometric models that I use to estimate the treatment effect of the California 20/20 rebate program on electricity consumption. Let y_{it} denote household i's average daily electricity consumption during billing month t, and $\triangle \ln y_{it} = \ln y_{it} - \ln y_{i,2004m9}$ the change in log of household's consumption between a billing period t and the September billing period in 2004.

Suppose that the program enrollment is randomly assigned among households. Then, the ordinary least squares (OLS) estimation of,

$$\triangle \ln y_{it} = \alpha + \beta \cdot Treat_i + \varepsilon_{it}, \tag{1}$$

produces a consistent estimate of the average treatment effect (ATE) of the rebate incentive because a random assignment assures that the error term ε_{it} is uncorrelated with the treatment dummy variable, $Treat_i$. In the California 20/20 program, however, the treatment was not randomly assigned. Instead, the treatment was determined by the following rule.

$$Treat_i = 1\left\{X_i \leq c\right\} \text{ where } \begin{cases} X_i = \text{account open date} \\ \\ c = \text{cutoff date} \end{cases} \tag{2}$$

Because the treatment assignment is a function of X_i, the OLS estimate of equation (1) is biased if $E[\varepsilon_{it}|X_i] \neq 0$. For example, for the first few months after move-in, households gradually increase their electricity consumption. This tendency is found in the billing data at any time period. As a result, $\triangle \ln y_{it}$ always has a slight positive trend in X_i, which is unrelated to the rebate program. Therefore, if this trend is ignored, the condition, $E[\varepsilon_{it}|X_i] \neq 0$ will be violated. This trend in $\triangle \ln y_{it}$ is quantitatively small and it disappears in a few month after the customer's move-in. A failure to control for this trend, however, would produce an upward bias in estimates of β, because the positive trend of $\triangle \ln y_{it}$ in X_i means that customers without the treatment are likely to have systematically higher $\triangle \ln y_{it}$ compared to customers with the treatment.

The main idea of regression discontinuity designs is that a potential bias from this trend can be eliminated as long as the relationship between the confounding trend and the error term ε_{it} is smooth and continuous in X_i. Given the condition, it is possible to consistently estimate the local average treatment effect (LATE) by including flexible parametric or nonparametric controls for X_i. Including a smooth function of X_i does not destroy the identification because the treatment variable, $Treat_i$ is a discontinuous function of X_i.

Imbens and Lemieux (2008) describe two approaches to specifying a smooth control function of X_i. The first approach is to include a flexible parametric function of X_i where the slope coefficients are allowed to be different on the left and right of the cutoff date:

$$\triangle \ln y_{it} = \alpha + \beta \cdot Treat_i + \sum_{s=1}^{S} (\gamma^s \cdot X_i^s + \theta^s \cdot Treat_i \cdot X_i^s) + \delta_{zip} + \delta_{cycle} + \varepsilon_{i,t}, \tag{3}$$

The equation includes a polynomial function of s order that is allowed to have different slopes across the cutoff point. Each of X_i^s is allowed to have different coefficients for the left and right side of the cutoff date. To control for other factors that influence electricity consumption (e.g. weather differences), I also include dummy variables at the zip code level, δ_{zip} and dummy variables at the billing cycle level, δ_{cycle}. These variables control for an unobservable shock between the two years at the zip code and billing cycle levels. Imbens and Lemieux (2008) note that this parametric approach could have a disadvantage in its parametric assumptions on the function of X_i. This concern motivates them to suggest the second approach that uses non-parametric controls for X_i by employing a local linear regression:

$$\triangle \ln y_{it} = K\left(\frac{X_i - c}{h}\right) \cdot \left(\alpha + \beta \cdot Treat_i + \gamma \cdot X_i + \theta \cdot Treat_i \cdot X_i + \delta_{zip} + \delta_{cycle} + \varepsilon_{i,t}\right). \tag{4}$$

The local linear regression is equivalent to a simple OLS regression but with higher weights on samples that are closer to the cutoff date. $K(.)$ is a kernel function for weights and h is a bandwidth. Similar to the previous equation, this regression also allows different slope coefficients on the left and right side of the cutoff date. Previous studies suggest that a triangular kernel is the most robust for discontinuous data points (Hahn, Todd, and der Klaauw 2001, Imbens and Lemieux 2008). There is, however, no rule of thumb for choosing a right bandwidth for local linear regressions. Therefore, I provide estimation results with different bandwidth choices for robustness checks.

The parametric and non-parametric approaches have advantages and disadvantages and there is still no convincing evidence for which approach works better in regression discontinuity estimation. Therefore, I estimate both equation (3) and (4) to explore how estimates will be affected by the choice between the parametric and non-parametric approaches.

5 Results

5.1 Validity Checks of the Regression Discontinuity Design

When customers opened their electricity account in 2004, nobody knew that there was going to be the 20/20 rebate program in 2005. It is therefore hard to imagine that customers strategically selected their account open date around the cutoff date of the program eligibility. Still, it can be a concern if there is a non-random discontinuous difference between customers around the cutoff date and the difference may affect the outcome variable of interest.

To assess the validity of the regression discontinuity design, I first plot the number of new accounts opened per day in Figure 4. The horizontal axis shows the account open date relative to the cutoff date of the program eligibility. For SCE, for example the horizontal axis shows a customer's account open date relative to June 5, 2004. Each dot shows the mean number of new accounts per day in SCE over the 15-day bandwidth. Everyday, about 1500 customers opened their account. The solid line shows local linear fit and the dashed lines present the 95% confidence intervals. Over the 90-day bandwidth, there is slight upward trend in the number of new accounts, although the slope is not statistically significant from zero and there is no discontinuous jump at the cutoff date.

In the second to fourth graphs in Figure 4, I plot household characteristics over the account open date relative to the cutoff date. Because the variables are from the U.S. Census 2000 at the census block group level, I cluster the standard errors of the local linear fit at the census block group level. All the three variables do not show a statistically significant discrete jump at the cutoff date.

5.2 Main Results

This section provides estimation results of the regression discontinuity estimation described in equation (3) and (4). The estimates of β can be interpreted as the program's local average treatment effect. Because the treatment effect can be different between households in different climate conditions, I present results by climate zones. The 2005 rebate program started in June and ended in September. To receive a rebate at the end of the summer, customers needed to reduce their overall electricity consumption in the four months by at least 20%. To examine whether the treatment effect is quantitatively different among the four months, I present results for each billing month separately. The results in this section suggest evidence of heterogeneous treatment effects between different climate zones. In the next section, I pool the data sets from different climate zones and estimate interaction effects of the treatment effect with other variables to examine what might explain the heterogeneous treatment effects.

In regression discontinuity estimation, graphical analyses are an important part of quantifying the magnitudes of treatment effects as well as checking the validity of identification strategy. The nature of regression discontinuity designs suggests that the effect of the treatment of interest can be measured by the value of the discontinuity in the expected value of the outcome at a particular point (Imbens and Lemieux 2008). Therefore, inspecting the estimated version of this conditional expectation is a simple yet powerful way to visualize the identification strategy.

Figure 5 shows regression discontinuity estimates for the September billing month in SCE by its climates zones. The horizontal axis is a household's account open date relative to the cutoff date for the program's eligibility, $X_i^c = X_i - c$.[7] Households on the left side of the cutoff date

[7]For example, if customer i in SCE started electricity service on June 25 in 2004, then $X_i^c = 20$, because

are the treatment group and households on the right side of the cutoff are the control group. I include only households that started their electricity service between 90 days before or 90 days after the cutoff date in 2004. In other words, the bandwidth is 90 days for each side of the cutoff date. The vertical axis shows $\widetilde{\triangle \ln y_{it}} = \triangle \ln y_{it} - \hat{\delta}_{zip} - \hat{\delta}_{cycle}$, the log change in average daily electricity consumption from the September billing month in 2004 to the September billing month in 2005. To control for weather and economic shocks, I subtract the zip code level mean and billing cycle level mean from the log change.

If the rebate program has a significant treatment effect on electricity consumption, the change in consumption from 2004 to 2005 should be lower for the treatment group relative to the control group. In that case, the conditional expectation of the outcome variable $\widetilde{\triangle \ln y_{it}}$ conditional on the running variable X_i should have a discontinuous jump across the cutoff date c. To see whether the expected value of the outcome variable has the discontinuous jump, I plot the local average value of $\widetilde{\triangle \ln y_{it}}$ over X_i^c. Each dot in Figure 5 shows the local average value of $\widetilde{\triangle \ln y_{it}}$ using fifteen days bandwidth of X_i^c. For the eight bins in each side of the cutoff, I take a simple local average for each bin, and plot them on the diagrams. The local averages are the estimated counterparts to the conditional mean of the outcome $E[\triangle \ln y_{it} | X_i^c, \delta_{zip}, \delta_{cycle}]$.

The top two figures show results for climate zone 10 and 17 in SCE. These climate zones include coastal areas, which have a relatively moderate summer climate condition relative to inland areas. For example, the cities of Santa Barbara, Long Beach, and Irvine are included in these climate zones. The figures suggest evidence that the program did not significantly change electricity consumption for the treatment group in these climate zones. The change in electricity consumption has a moderate positive trend in the account open date as discussed in the previous section, but it does not have a discontinuous jump at the cutoff date.

SCE's cutoff date was June 5, 2004.

In contrast, the bottom two figures indicate evidence that the rebate program had a significant effect on electricity consumption in climate zones 15 and 16. These climate zones are located in inland areas of southern California, where the summer temperature is persistently high and households typically use an air conditioner throughout the summer. The next section explores in more detail how higher temperature affects the treatment effect.

To statistically estimate the magnitude of the treatment effects, I estimate the parametric regression in equation (3) using quadratic functions and the nonparametric local linear regression in equation (4) using a triangular kernel function for the sample weight. The dashed lines show fitted value for the parametric regression and the solid lines show the fitted lines for the nonparametric local linear regression. Essentially, both econometric equations estimate β as the magnitude of the jump in the outcome variable at the cutoff date by fitting parametric or nonparametric functions of X_i^c. As it is visually clear in Figure 5, the estimates are not statistically different between the two estimation methods. Following Imbens and Lemieux (2008), I focus on the estimates from the nonparametric local linear regression in the following, but none of the estimates are sensitive to the selection of the two estimation methods.

Each diagram in Figure 5 includes the point estimate from the local linear regression and the robust standard errors in the parentheses. In climate zone 10 and 17, the point estimates are close to zero and they are statistically insignificant from zero. In climate zone 15 and 16, the point estimates are -.093 and -.101 with standard errors .04 and .032, respectively. That is, households with the rebate incentive reduced their consumption by about 9% in climate zone 15 and 16 relative to households without the rebate incentive.

I find similar results in other utility territories. For example, Figure 6 shows the regression discontinuity estimates for the September billing cycle in the two most populated coastal areas in SDG&E. The graphical evidence suggest that, similar to coastal climate zones in SCE, the rebate program did not significantly change electricity consumption for the two areas in SDG&E.

Table 3 summarizes the results for all billing months and all of the three utilities. Following Lee and Card (2008), I cluster standard errors by the discrete assignment variable of treatment status X_i, a customer's account start date, to account for specification errors in the continuous function $f(X_i)$. The table shows that there are two findings that are consistent among the three electric utilities. First, in the coastal areas, the treatment effects are not statistically different from zero in all of the four summer months. However, in the inland areas, the rebate incentive lowered household consumption by 5 to 9%. Second, in the inland areas, the treatment effect is largest in the last month and is monotonically increasing from the first month. There are two possible reasons for this differences in treatment effects between the months. The first potential reason is that some households may have gradually become aware of the program when they looked at the information about the rebate program on their monthly bills. The second possible reason is that the program's design may have created a larger incentive for customers to reduce consumption at the end of the four-month period than at the beginning. Once customers have already achieved a certain amount of reduction in the first couple of months, the possibility of receiving a rebate in return for their efforts is more certain in the later months.

Lee and Lemieux (2010) note that if a regression discontinuity design produces an essentially random assignment of treatment, including covariates in the estimation should not change the consistency of the estimate for the same reason that adding covariates in randomized controlled experiments does not change the consistency. Table 4 shows results with alternative specification for the inland climate zone in SCE in September. The main specification includes zip code and billing cycle dummy variables. Column 1, 2, and 3 show that excluding these dummy variables do not change the point estimates. There is a slight chance that including micro climate variables might reduce standard errors because the billing cycle dummy captures only the overall weather condition for each billing cycle[8]. Using daily temperature data for 4 km by 4 km cells, I calculate cooling degree days and heating degree days for each bill. Column 5 shows that including these micro climate controls still do not change the estimate but slightly reduce standard errors.

5.3 Interaction with Weather and Income

The previous section finds that the estimated treatment effects are larger in inland areas compared to coastal areas. This section explores what drives the heterogeneous treatment effect of the 2005 rebate program. In particular, I examine whether climate conditions or income differences can explain the heterogeneous treatment effects.

One of the significant differences between inland and coastal California is the summer climate conditions. For example, Figure 3 illustrates cooling degree days (CDD) in California in August 2005 by five-digit zip code areas. Generally, summer temperature is persistently high in the inland areas but quite moderate in the coastal areas. As a result, inland households typically use air conditioners throughout the summer while coastal households either use air conditioners very little or do not own them at all. For households that do not use an air conditioner, a 20% reduction in summer electricity consumption is challenging for typical residential electricity consumers. In contrast, for households that constantly use an air conditioner during the summer season, a 20% consumption reduction can be achieved by slightly changing the temperature settings or the length of usage.

Another significant difference between inland and coastal California is their demographic

[8] For example, even in the same climate zone and within the same billing cycle, micro climate conditions might be slightly different especially when the climate zone includes considerably different micro climate areas.

characteristics. For instance, income levels tend to be higher in coastal areas than inland areas. In previous studies on residential electricity demand, many studies find slightly larger price elasticity estimates for low income households (e.g. Reiss and White 2005 and Ito 2010). Because the 20/20 rebate program is essentially a price-discount rebate program, households with lower income may be more likely to respond to the incentive if their price elasticity is larger than households with higher income.

To examine how climate conditions and income levels affect the program's treatment effects, I focus on the September billing month from all climate zones and conduct two statistical tests. First, I focus on households in SCE's climate zone 10 and explore whether within-climate-zone variation can provide any evidence of heterogeneous treatment effects. I split households in this climate zone by the quartiles of cooling degree days and the quartiles of household income. Then, for each of the quartile, I estimate the local linear regression in equation (5). Table 5 shows the evidence that even in this coastal climate zone I find evidence that households with a large numbers of high temperature days show an economically small but statistically significant treatment effect. Similarly, the treatment effect is statistically significant for households with lower income.

As a second approach, I pool data from all climate zones and include interaction terms between the treatment variable and average temperature in each customer's billing cycle, and the treatment variable and household income to equation (4). This model estimates the differences in the treatment effect for different temperature values and income levels assuming that the interaction terms linearly affect the treatment effect. Column 1 of Table 6 shows estimation results with the interaction term between the treatment and average temperature. It indicates that the treatment effect increases by .15 percentage point when the average temperature increases 1 degree Fahrenheit. Column 2 includes the treatment effect and the interaction variable between the treatment effect and log of income. The estimates suggest that the treatment effect increases by .027 percent with an 1% increase in household income. These two interaction effects remain the same when both terms are included in the regression as Column 3 of the table shows. Therefore, results from the two estimation methods indicate that both climate conditions and income levels have a statistically significant effect on the magnitude of the program's treatment effect.

5.4 Cost-Effectiveness of the Program

The cost-effectiveness of the program is a central policy question for conservation rebate programs such as the California 20/20 electricity rebate program. Policy makers often argue that the simplicity of the programs makes it easy for consumers to understand the incentive and therefore encourage their conservation. On the other hand, many people doubt its cost-effectiveness because a considerable number of customers may receive rebates in the absence of extra efforts for conservation. In the past, the utility companies reported the cost-effectiveness by calculating the total rebate paid to customers who achieved a 20% consumption reduction. However, it is misleading to use the measurement as the cost-effectiveness because 1) some of these rebated customers reduced their consumption for reasons unrelated to the rebate incentive and 2) some customers reduced their consumption but cannot reached the 20% target level.

In this section, I define the cost-effectiveness of the program as how much money is spent to save 1 kWh of electricity consumption.[9] To obtain an estimate of the consumption reduc-

[9]Note that this cost-effectiveness does not provide a welfare measure such as the efficiency cost of the program. Because the rebate expense is simply a transfer between customers, one can see that the rebate program produces an implicit increase in electricity price. If the ex-ante price is set at the efficient level, the rebate program creates efficiency cost. If the ex-ante price is too low because it does not include environmental externalities for example, the rebate program may increase welfare. However, the existence of five-tier increasing block pricing makes it

tion (kWh) produced by the program incentive, I use the treatment effect estimated from the regression discontinuity design in the previous section. Then, I divide the estimate by the total rebate expense that was paid to customers. This exercise implicitly assumes the external validity of the RD estimate. That is, this calculation assumes that the treatment effect is the same between customers in the RD samples, who opened their account around the cutoff date in 2004, and other customers, which include customers opened their account before 2004. If the true treatment effects are different between the two types of customers, the cost-effectiveness estimates are biased. The direction of the bias is ambiguous. The treatment effect may be larger for customers who have been in the same premise long time because they are likely to have a larger incentive to invest some money on energy efficient devices. However, if these customers have already invested on such devices, their consumption in the base year is already low, which makes it harder for these customers to achieve an additional reduction in consumption in the target year. Given this assumption, I calculate the cost-effectiveness of the 20/20 rebate program in 2005. I focus on Southern California Edison's territory in this section. The result for other electric utilities are similar to the results presented in this section.

In Table 7, I calculate total electricity consumption in the summer of 2004 and 2005, the amount of rebate expenses in 2005, and the estimated consumption reduction that is produced by the program's incentive. Column 1 shows the number of households who maintained their electricity account in the summer of 2004 and 2005. Column 2 presents the total electricity consumption of these households during the four summer months in 2005 and column 3 shows how much money was spent to pay rebates to households that reduced their consumption at

harder to conclude whether the rebate program improves welfare. Even with the environmental externalities, the higher tiers of the increasing block pricing (20 cents to 35 cents per kWh) are well above the social marginal cost. However, the lower tiers (10 cents per kWh) can be lower than the social marginal cost. Therefore, the total welfare effect depends on 1) the distribution of customers in the increasing block pricing and 2) the price elasticity of electricity for these customers.

least by 20% relative to 2004.

To obtain an estimate of how much of electricity consumption was saved by the rebate program, I calculate an estimated kWh reduction K_j for each climate zone j based on the estimated treatment effect $\hat{\beta^j_m}$ in Table 3,

$$K_j = \sum_{m=6}^{9} \left(\frac{-\hat{\beta^j_m}}{1+\hat{\beta^j_m}} \cdot C^j_m \right), \tag{5}$$

where m is a billing month and C^j_m is aggregate consumption.[10] Then, I calculate the cost-benefit ratio of the program as the total rebates divided by K_j. An important assumption here is that the estimated treatment effect $\hat{\beta^j_m}$, which is estimated from the regression discontinuity design, can be applied to all households in the climate zone regardless of their service start date.

Column 4 shows the estimated reduction K_j and column 5 presents the aggregate rebate expense divided by K_j. In the coastal areas, the program cost 90.7 cents to save 1 kWh electricity. On the other hand, in the inland areas, the program spent 2 cents to save 1 kWh consumption. The average cost per kWh reduction is 14.8 cents.

Note that this average estimate of the cost and benefit does not necessarily give a fully accurate evaluation of the program's cost and benefit. Ideally, we want to know how much of the reduction happened in the on-peak and off-peak periods of the electricity load. For example, if most of the reduction occurred in the on-peak load, in which the marginal cost of electricity is relatively high, the benefit of the consumption reduction is large. However, because the monthly consumption data do not give the exact timing of consumption reductions, I discuss the costs and benefits of the program based on this average number.

[10] Let A^j_m denote the aggregate consumption in the absence of the treatment effect. Then, $(1+\hat{\beta^j_m})A^j_m = C^j_m$. Hence, $A^j_m = \frac{C^j_m}{1+\hat{\beta^j_m}}$ and $K^j_m = A^j_m - C^j_m = \frac{-\hat{\beta^j_m}}{1+\hat{\beta^j_m}} \cdot C^j_m$.

The estimated cost, 14.8 cents per kWh consumption reduction, is larger than the numbers that were given by the electric utilities because the utilities usually attribute all of the consumption reductions by the rebated customers to the presence of the rebate program. In previous studies, Reiss and White (2003) estimate the 2001 rebate program's costs and benefits and find that the overall cost (June to September in 2001) for SDG&E was 18 cents per kWh.

A potential reference point for discussing the program's cost is the average cost of electricity production, which was 13.37 cents per kWh in 2005 in SCE. Compared to this number, the average cost estimate, 14.8 cents per kWh, is still higher than the average cost of supplying electricity.

An important finding in the current study is that the cost-effectiveness is substantially different between the coastal and inland areas in California. Previous studies provide only aggregate cost-benefit estimates for all of California. Results in Table 3 suggest that the 2005 20/20 rebate program was fairly cost-effective in the inland areas but quite cost-ineffective in the coastal areas. This is because in the coastal areas, a large number of households received a rebate, but the treatment effect in these areas were nearly zero. Finally, note that these cost estimates do not include the administrative costs and advertisement fees that were associated with the rebate program. Therefore, the actual cost is likely to be higher than the cost estimates in this section when we account for these additional program expenses.

6 Conclusion and Future Work

This paper examines the treatment effect and cost-effectiveness of conservation rebate programs that are often used by electric, natural gas, and water utilities. To deal with identification problems, I apply a regression discontinuity design to the 2005 California 20/20 electricity rebate program. The discontinuous eligibility rule of the program enables me to estimate the treatment effect by controlling for confounding factors such as weather and economic conditions.

This study provides several empirical findings based on a panel data set of household-level monthly electricity billing records from the three largest electric utilities in California. First, the regression discontinuity estimates provide evidence that the rebate incentive made consumers reduce their electricity consumption by 5% to 10% in the areas where the summer temperature is persistently high and the income level is relatively low. In contrast, the treatment effects are nearly zero in the areas where the summer temperature is moderate and the income level is relatively high. Second, to explore which variables explain this heterogeneous treatment effect, I estimate interaction effects between the treatment variable and climate conditions, and between the treatment variable and income levels. Results from these regressions suggest that the treatment effect increases by .15 percent as average temperatures increase by 1 °F and decreases by .027 percent as income levels increase by 1%. Finally, using the estimates of the treatment effect, I calculate the cost-effectiveness of the program. The results from this exercise show that the program cost 90.7 cents in the coastal areas and 2 cents in the inland areas to save 1 kWh of electricity consumption. The overall cost per kWh reduction was 14.8 cents per kWh.

The results from this study provide several policy implications for the California 20/20 electricity rebate program. First, under the current rebating scheme, the expense of natural year-to-year fluctuations in electricity consumption is substantial. As a result, providing a rebate for reductions that would have happened in the absence of the program can be very costly unless the treatment effect is sufficiently large. Second, the estimation results suggest that it is important to account for heterogeneous treatment effects particularly based on different weather conditions and income levels among households. For example, my cost-effectiveness estimates for the coastal areas are by far larger than previous estimates while my estimates for the inland areas are far lower than previous estimates. Finally, the heterogeneous treatment effect suggests that the program's performance could be improved if the program focused on certain types of households to minimize rebate expenses for reductions that would have occurred in the absence of the program.

The paper leaves at least two important research questions for future work. First, because the rebate program required households to reduce their overall electricity consumption by 20% over the four summer billing months in 2005, there could be a dynamic response to their incentive throughout the four month period. For example, households that achieved a large reduction in the first three months do not have to reduce much of their last month's consumption to reach an overall 20% reduction. On the other hand, households that consumed too much in the first three months have no way to get a rebate regardless of their effort in the last month. Second, in this paper, I do not fully specify a consumer's electricity demand model with their rebate incentive. Instead, I estimate the average treatment effect of the program. If I model a consumer's decision more precisely in a demand model, consumers could get two different incentives from the rebate program. This is because the rebate program may affect consumption in two different ways. In one hand, the program provides an incentive to reduce consumption because households can receive a rebate only if they reduce their consumption by 20%. On the other hand, once a household receives a rebate, the rebate works as a price discount for total consumption. The household, therefore, may increase consumption in response to the price discount. For example, in an extreme case where a household is sure to use much less electricity in 2005 relative to 2004, the household has almost no incentive to reduce consumption and is likely to increase consumption in response to the expected discount for electricity price. My future work would incorporate this behavior to provide more comprehensive understanding of the program's effect on electricity demand.

References

Aroonruengsawat, Anin and Maximilian Auffhammer. 2009. "Impacts of Climate Change on Residential Electricity Consumption: Evidence From Billing Data." *California Energy Commission Report* (CEC-150-2009-001).

Borenstein, S., J. B Bushnell, and F. A Wolak. 2002. "Measuring Market Inefficiencies in Californias Restructured Wholesale Electricity Market." *The American Economic Review* 92 (5):1376–1405.

Bushnell, James B. and Erin T. Mansur. 2005. "Consumption under Noisy Price Signals: A Study of Electricity Retail Rate Deregulation in San Diego." *Journal of Industrial Economics* 53 (4):493–513.

CPUC. 2001. "Resolution E-3733." *California Public Utility Commission* .

Faruqui, Ahmad and Stephen S. George. 2006. "Pushing the Envelope on Rate Design." *The Electricity Journal* 19 (2):33–42.

Goldman, Charles A., Galen L. Barbose, and Joseph H. Eto. 2002. "California Customer Load Reductions during the Electricity Crisis: Did They Help to Keep the Lights On?" *Journal of Industry, Competition and Trade* 2 (1):113–142.

Hahn, J., P. Todd, and W. Van der Klaauw. 2001. "Identification and estimation of treatment effects with a regression-discontinuity design." *Econometrica* 69 (1):201–209.

Imbens, Guido W. and Thomas Lemieux. 2008. "Regression discontinuity designs: A guide to practice." *Journal of Econometrics* 142 (2):615–635.

Inc, Wirtshafter Associates. 2006. "Evaluation of the California Statewide 20/20 Demand Reduction Programs." Tech. rep.

Ito, Koichiro. 2010. "Do Consumers Respond to Marginal or Average Price? Evidence from Nonlinear Electricity Pricing." *Energy Institute at Haas Working Paper 210* .

Joskow, P. L. 2001. "California's electricity crisis." *Oxford Review of Economic Policy* 17 (3):365.

Lee, David S. and David Card. 2008. "Regression discontinuity inference with specification error." *Journal of Econometrics* 142 (2):655–674.

Lee, David S. and Thomas Lemieux. 2010. "Regression Discontinuity Designs in Economics." *Journal of Economic Literature* 48 (2):281–355.

Naucler, T. and P. A. Enkvist. 2009. "Pathways to a Low-Carbon Economy: Version 2 of the Global Greenhouse Gas Abatement Cost Curve." *McKinsey & Company* .

Puller, Steven L. 2007. “Pricing and Firm Conduct in California’s Deregulated Electricity Market.” *Review of Economics and Statistics* 89 (1):75–87.

Reiss, Peter C. and Matthew W. White. 2003. “Demand and Pricing in Electricity Markets: Evidence from San Diego During California’s Energy Crisis.” *National Bureau of Economic Research Working Paper Series* No. 9986.

———. 2005. “Household electricity demand, revisited.” *Review of Economic Studies* 72 (3):853–883.

———. 2008. “What changes energy consumption? Prices and public pressures.” *RAND Journal of Economics* 39 (3):636–663.

Wirtshafter Associates, Inc. 2006. “Evaluation of the California Statewide 20/20 Demand Reduction Programs.” .

Figure 1: Program Eligibility Rule for the 2005 California 20/20 Electricity Rebate Program

	Cutoff Date
PG&E	June 1, 2004
SCE	June 5, 2004
SDG&E	June 30, 2004

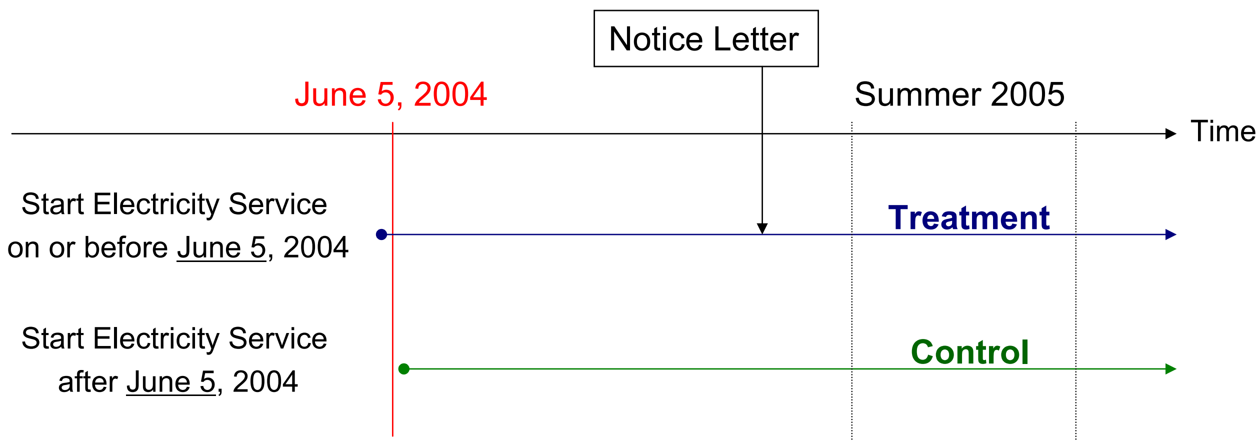

Note: Households who opened their account on or before the cutoff date in 2004 received a notice letter around April, 2005 and were automatically enrolled in the 2005 California 20/20 electricity rebate program program. These households were eligible for a 20% discount on their summer electricity bills if they reduced their electricity consumption by 20% relative to their consumption in 2004. Households who opened their account after the cutoff date were excluded from the program. The three electric utilities have slightly different cutoff dates.

Figure 2: Electric Utility Service Areas in California

Note: This figure shows the service areas of electric utilities in California. The original source file is available at the California Energy Commission's website. Three investor owned electric utilities, Pacific Gas & Electric, Southern California Edison, and San Diego Gas & Electric, participated in the 2005 California 20/20 electricity rebate program program.

Figure 3: Cooling Degree Days in August 2005 in California

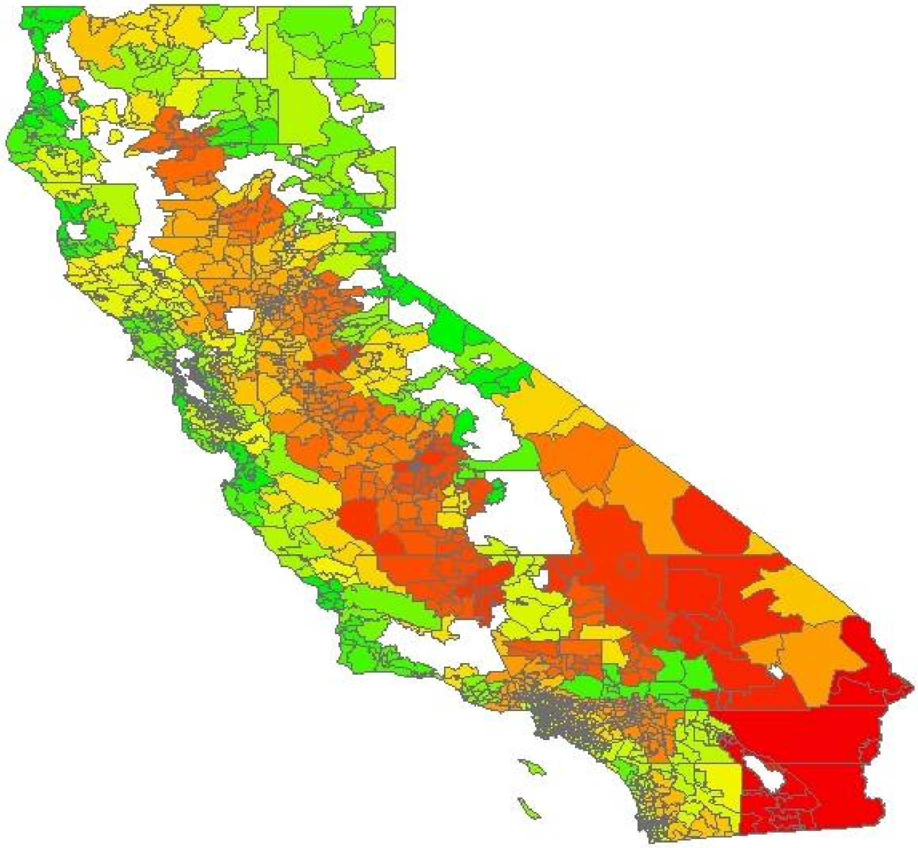

Note: This figure shows the cooling degree days (CDD) in August 2005 in California by zip code boundaries.

Figure 4: Validity Checks of the Regression Discontinuity Design

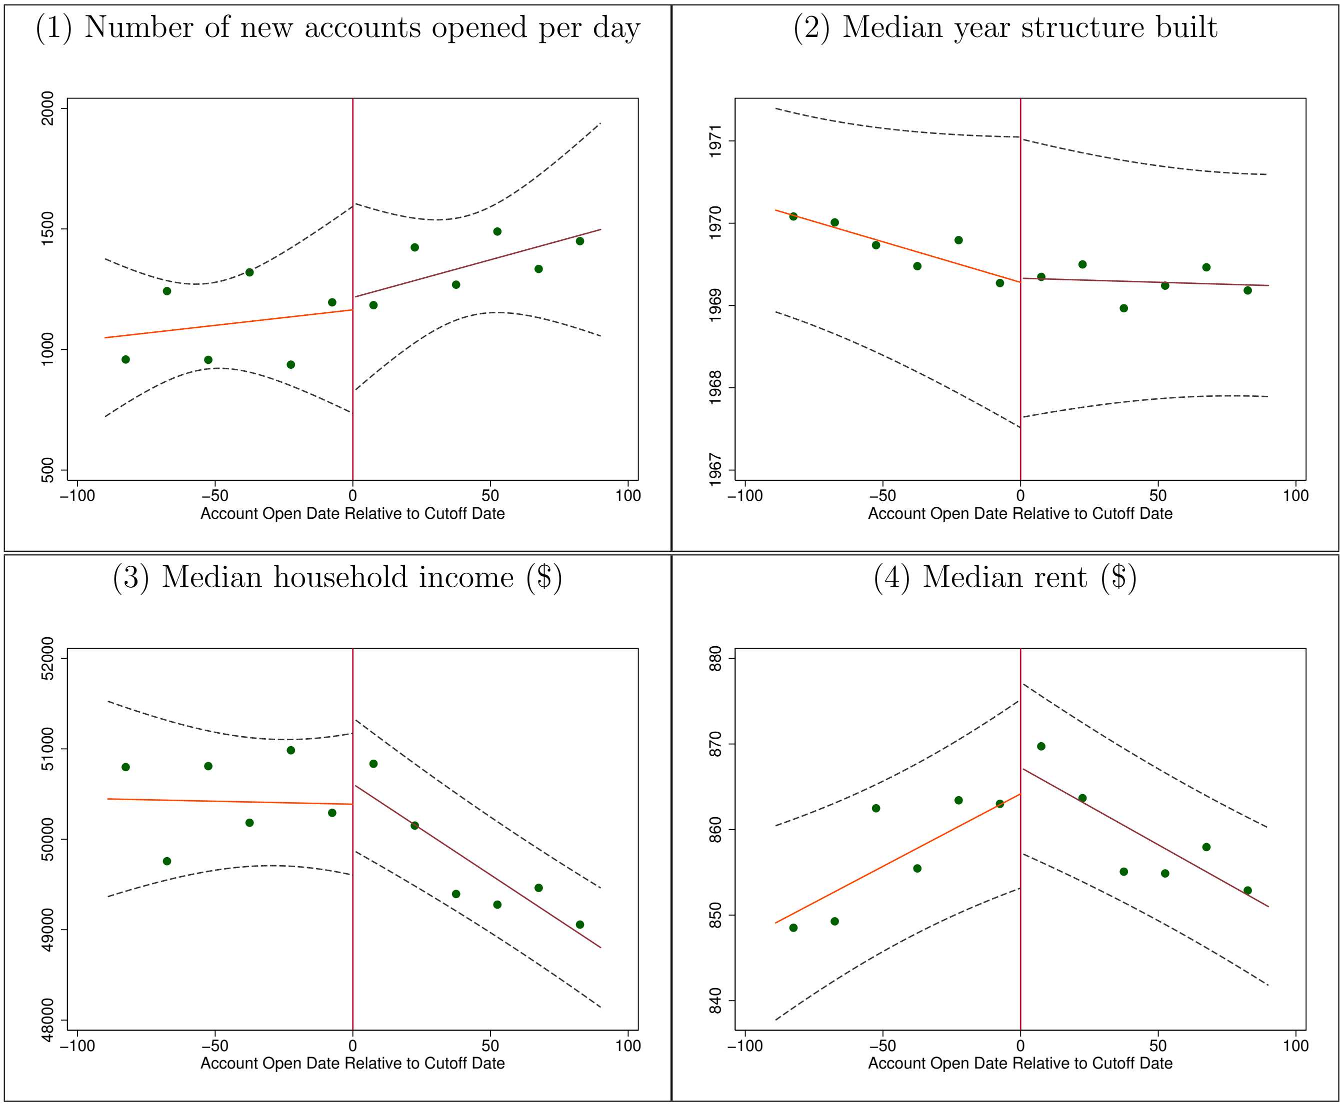

Note: The horizontal axis shows the account open date relative to the cutoff date of the program eligibility. For SCE, for example the holizontal axis shows a customer's account open date relative to June 5, 2004. Each dot shows the local mean using the 15-day bandwidth. The solid line shows local linear fit and the dashed lines present the 95% confidence intervals. I obtain each customer's account open date from the electricity billing data. Other three variables are from the U.S. Census 2000. The confidence intervals for the fitted lines for the three variables are adjusted for clustering at the census block group level.

Figure 5: Regression Discontinuity Estimates: SCE September Billing Month

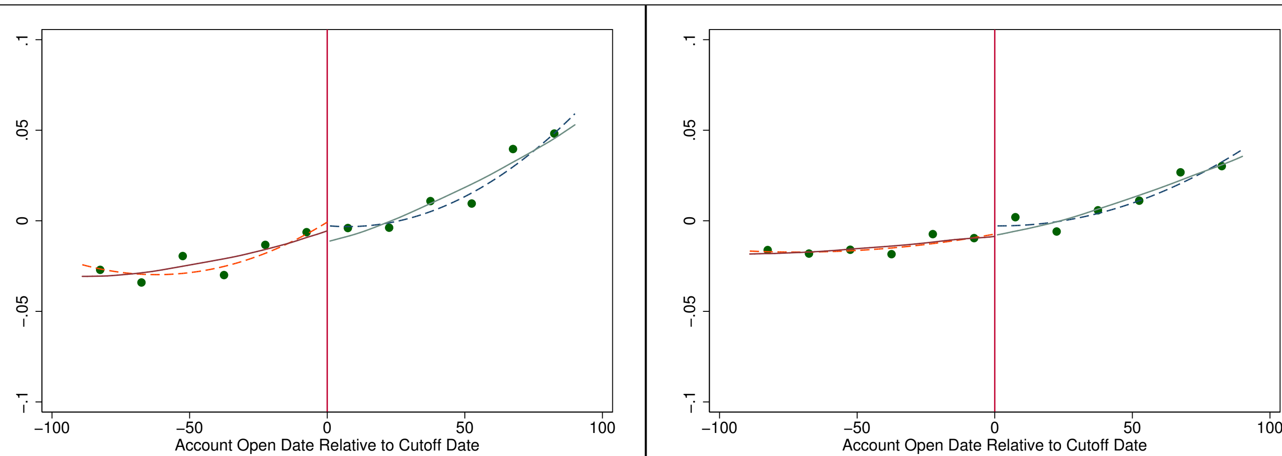

Climate Zone 10 (Santa Barbara, Irvine)

.005 (.007)

Climate Zone 17 (Riverside)

$-.002(.008)$

Climate Zone 15 (Palm Desert, Death Valley)

$-.093^{**}(.040)$

Climate Zone 16 (Bakersfield)

$-.101^{***}(.032)$

Note: This figure presents the regression discontinuity estimates for the September billing month in SCE by its climate zones. The horizontal axis shows households' account open date relative to the cutoff date for the program eligibility. The vertical axis shows the log change in September consumption from 2004 to 2005 where zip code level mean and billing cycle level mean are subtracted. Each dot presents the local mean using fifteen days window and the solid and dashed lines are the fitted lines by equation (3) and (4), respectively. I also include representative cities for each climate zone in parentheses. Finally, the figure includes the point estimate of the treatment effect with the robust standard errors in parentheses.

Figure 6: Regression Discontinuity Estimates: SDG&E September Billing Month

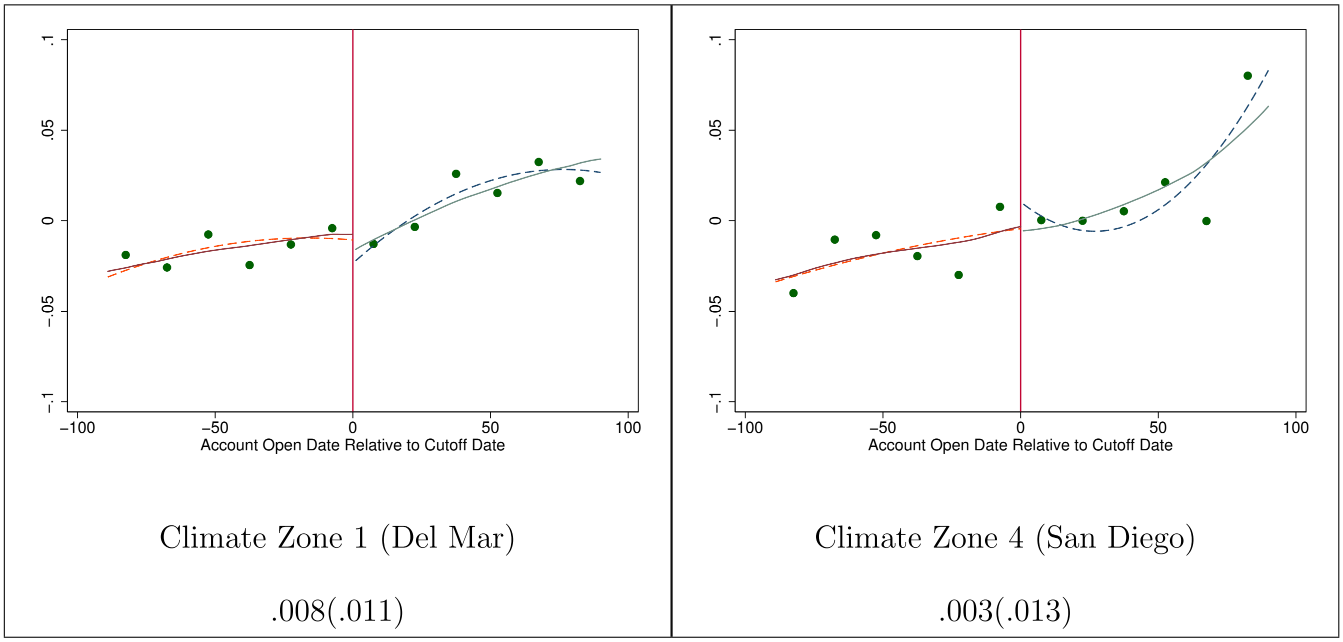

Note: This figure presents the regression discontinuity estimates for the September billing month in SDG&E by its climate zones. The horizontal axis shows households' account open date relative to the cutoff date for the program eligibility. The vertical axis shows the log change in September consumption from 2004 to 2005 where zip code level mean and billing cycle level mean are subtracted. Each dot presents the local mean using fifteen days window and the solid and dashed lines are the fitted lines by equation (3) and (4), respectively. I also include representative cities for each climate zone in parentheses. Finally, the figure includes the point estimate of the treatment effect with the robust standard errors in parentheses.

Table 1: Aggregate Consumption and Rebates in the Summer Billing Months in 2005

Utility	Consumption (kWh)	Revenue ($)	Rebated Households	Rebate ($)
PG&E	10,065,216,512	1,320,995,584	8.24%	10,786,594
SCE	9,401,883,648	1,257,056,768	7.91%	10,609,540
SDG&E	2,284,046,848	363,180,320	9.07%	4,325,000

Note: This table reports the statistics based on the actual residential billing data in the June, July, August, and September billing months in 2005. I include customers who maintained their account both in the summer of 2004 and 2005. The rebate expenditure does not include the administrative and advertising costs of the program. All expenditures are in nominal 2005 dollars.

Table 2: Changes in Summer Electricity Consumption in SCE

Year	Changes in Summer Weather	Median of % Changes in Consumption	% Households with 20% or More Reduction
From 2003 to 2004	Cooler in 2004	-1.7%	14.3%
From 1999 to 2000	Hotter in 2000	7.7%	6.8%

Note: This table reports statistics of within-household changes in summer electricity consumption in Southern California Edison (SCE). I first calculate the change in consumption for each household between the two years. I then calculate the median value of the change and the percentage of households who reduced their consumption more than 20%. Note that SCE customers did not encounter a price spike during the California electricity crisis in 2000 because their retail rates are capped.

Table 3: Treatment Effects in Each Billing Months by Climate Zones

Billing Month	6	7	8	9
PG&E				
Coastal Areas	-.002	-.001	.003	-.002
	(.004)	(.003)	(.004)	(.005)
Inland Areas	-0.01	-.016*	-.032***	-.059***
	(.013)	(.011)	(.011)	(.012)
SCE				
Coastal Areas	.001	-.001	-.001	-.002
	(.009)	(.010)	(.009)	(.008)
Inland Areas	-.019*	-.032**	-.056***	-.092***
	(.015)	(.016)	(.016)	(.015)
SDG&E				
Coastal Areas	.005	-.001	-.002	.008
	(.009)	(.010)	(.009)	(.011)
Mid-Inland Areas	-.002	-.001	.002	.003
	(.011)	(.012)	(.011)	(.013)

Note: This table presents the regression discontinuity estimates of the local linear regression of equation (5) with the triangular kernel and 90 days bandwidth. Each estimate comes from separate regressions for each billing month and climate zones. The number of observations is 535,741 (Coastal, PG&E), 168,162 (Inland, PG&E), 492,244 (Coastal, SCE), 86,118 (Inland, SCE), 138,718 (Coastal, SDG&E), and 100,450 (Mid-inland, SDG&E). Standard errors clustered by the assignment variable (account open date) are in parentheses. ***, **, and * show 1%, 5%, and 10% statistical significance respectively.

Table 4: Alternative Specification with Different Covariates

	(1) Baseline model	(2) No zip dummy	(3) No bill cycle dummy	(4) No zip and bill cycle dummy	(5) Add detail weather controls
Treat	-.092***	-.091***	-.091***	-.090***	-.093***
	(.015)	(.017)	(.017)	(.017)	(.014)
Cooling degree days					.037
					(.001)
Heating degree days					.003
					(.001)
Zip code dummy	Yes	No	Yes	No	Yes
Billing cycle dummy	Yes	Yes	No	No	Yes
Observations	86,118	86,118	86,118	86,118	86,118

Note: This table presents the regression discontinuity estimates of the local linear regression of equation (5) with the triangular kernel and 90 days bandwidth for SCE's inland climate zone in September. Standard errors clustered by the assignment variable (account open date) are in parentheses. ***, **, and * show 1%, 5%, and 10% statistical significance respectively.

Table 5: Treatment Effect by Different Weather and Income Quartiles

Quartile	1	2	3	4
Average Temperature				
Treat	.013	-.005	-.019	-.030***
	(.012)	(.009)	(.011)	(.012)
Observations	380,355	380,361	380,358	380,359
Income				
Treat	-0.017**	-.011	-.005	-.0003
	(.009)	(.008)	(.007)	(.0009)
Observations	380,335	380,363	380,397	380,338

Note: This table presents the regression discontinuity estimates of the local linear regression of equation (5) with the triangular kernel and 90 days bandwidth for September 2005. The number of observations is 231,318. The dependent variables is the log change in household daily electricity consumption from September 2004 to 2005. For the first row, I divide the samples into quartiles based on the cooling degree days (CDD) in September 2005 and run separate regressions for each quartile. Similarly for the second row, I divide the samples into income quartiles and run separate regressions for each quartile. Standard errors clustered by the assignment variable (account open date) are in parentheses. ***, **, and * show 1%, 5%, and 10% statistical significance respectively.

Table 6: Treatment Effect Interacted with Temperature and Income

	(1)	(2)	(3)
Treat	.095**	-.297***	-.199***
	(.051)	(.055)	(.077)
Treat*Ave.Temp.	-.0015**		-.0016**
	(.0007)		(.0008)
Treat*ln(Income)		.027***	.027***
		(.005)	(.005)
Observations	1,521,433	1,521,433	1,521,433

Note: This table presents the regression discontinuity estimates of the local linear regression of equation (4) with the triangular kernel and 90 days bandwidth for September 2005. The sample includes households in all climate zones in each of the electric utilities. The number of observation is 1,521,433. The dependent variables is the log change in household daily electricity consumption from September 2004 to 2005. I calculate average temperature for each billing cycle by taking mean of the average temperature during the billing days. Standard errors clustered by the assignment variable (account open date) are in parentheses. ***, **, and * show 1%, 5%, and 10% statistical significance respectively.

Table 7: Cost-Benefit Analysis for SCE

	Customers	Consumption (kWh)	Rebate ($)	Estimated Reduction (kWh)	Rebate/ Reduction ($/kWh)
Coastal	3,190,027	8,247,457,920	9,358,919	10,323,778	0.907
Inland	299,178	1,154,292,248	1,250,621	61,486,108	0.020
Total	3,489,205	9,401,750,168	10,609,540	71,809,886	0.148

Note: This table reports the cost-benefit analysis of the 20/20 program for SCE's coastal areas, inland areas, and all service areas. Column 1 shows the number of residential customers that maintained their account in the summer of 2004 and 2005. Column 2 presents the aggregate consumption in the four summer months. Column 3 reports the aggregate amount of rebates. Column 4 shows the estimated kWh reduction by the program and the last column presents the rebate expenditure per kWh reduction.