Asymmetric Incentives in Subsidies: Evidence from a Large-Scale Electricity Rebate Program[†]

By Koichiro Ito*

Many countries use substantial public funds to subsidize reductions in negative externalities. Such policy designs create asymmetric incentives because increases in externalities remain unpriced. I investigate the implications of such policies by using a regression discontinuity design in California's electricity rebate program. Using household-level panel data, I find that the incentive produced precisely estimated zero treatment effects on energy conservation in coastal areas. In contrast, the rebate induced short-run and long-run consumption reductions in inland areas. Income, climate, and air conditioner saturation significantly drive the heterogeneity. Finally, I provide a cost-effectiveness analysis and investigate how to improve the policy design. (JEL D12, D62, H76, L94, L98, Q48)

In economic theory, negative externalities can be corrected by Pigouvian taxes that internalize external costs (Pigou 1924). However, taxpayer opposition usually prevents the introduction of such taxes in practice. Alternatively, regulators use substantial public funds to subsidize economic activities that presumably induce reductions in negative externalities. For example, many countries have failed to introduce a carbon tax on greenhouse gas emissions and decided to provide large subsidies for energy conservation and pollution abatement.[1] Likewise, regulators usually provide subsidies for smoking cessation and public transportation in lieu of high taxes on smoking and traffic congestion.[2]

[1] The American Recovery and Reinvestment Act of 2009 provided $17 billion for energy conservation programs. US electric utilities spent $26 billion on energy efficiency programs in 1994–2011, and annual spending has been continuously increasing since 2003 (US Energy Information Administration 2013).

[2] Many countries provide subsidies for energy-efficient appliances (Davis, Fuchs, and Gertler 2014; Boomhower and Davis 2014), energy-efficient vehicles (Gallagher and Muehlegger 2011, Sallee 2011, Mian and Sufi 2012, and Sallee and Slemrod 2012), and reductions in energy consumption (Reiss and White 2008, Wolak 2010,

Such subsidies, however, often create asymmetric incentives because *increases* in negative externalities remain unpriced. With this asymmetry, these policies may not be able to correct all negative externalities if the marginal decisions of many individuals are unaffected by the subsidy incentive. This adverse effect contrasts with the theory underlying Pigouvian taxation, which aims to equalize the private and social marginal costs for all individuals. Despite the importance of this problem, there is limited empirical evidence on this question largely because empirical analysis is hampered by the "additionality" problem (Joskow and Marron 1992)—some of the observed behavior are not "additional" if they would occur in the absence of subsidy incentives. It is, therefore, misleading to evaluate a subsidy program's causal effect simply by analyzing those who received a subsidy, although many previous studies take this approach.[3]

In this paper, I investigate these problems by applying a regression discontinuity (RD) design to a large-scale electricity rebate program in California. For the summer of 2005, California residents received a 20 percent discount on their monthly electricity bills if they reduced their electricity usage by 20 percent compared to the summer of 2004. The program's eligibility rule provides two advantages for my empirical strategy. First, to be eligible for the program in summer 2005, customers had to open their electricity account before a cutoff date in 2004. A strategic manipulation of the account opening dates was impossible because until the spring of 2005 the program was not announced. This rule created a sharp discontinuity in the treatment assignment between customers who opened their accounts before and after the cutoff date. Second, all eligible customers were automatically enrolled in the program, preventing the self-selection problem, which presents a major challenge in previous studies.[4]

I obtained the administrative data on customer-level monthly electricity billing records from the electric utilities that administered the rebate program. Compared to survey data, the full administrative billing records have the advantages of measurement accuracy and the comprehensive coverage of customers. The data also include each customer's nine-digit zip code, which I match with demographic and weather data to investigate potential heterogeneity in response to the subsidy incentive.

Using the RD design, I first estimate the rebate program's local average treatment effect (LATE). I find that the rebate incentive reduced electricity consumption by 4 percent in inland areas in California, where the summer temperatures are persistently high and the income levels are relatively low. Moreover, this conservation effect continued during the summers of 2006, 2007, and 2008. In contrast, I

---

and Borenstein 2013). Carbon offset programs, such as the United Nations' clean development mechanism, give firms credits if they reduce their pollution relative to a business-as-usual baseline level (Sutter and Parreño 2007, Schneider 2007, and Duflo et al. 2013). Financial incentives for smoking cessation are becoming a key policy instrument (Volpp et al. 2009). Congestion pricing is still rarely implemented in the US transportation system, and the federal and state governments use substantial public funds to subsidize public transit to address congestion (Anderson 2014).

[3] Joskow and Marron (1992) argue that many policy evaluations of utility conservation programs fail to take into account the additionality problem. Boomhower and Davis (2014) studies how this problem affects the cost-effectiveness of an energy-efficient appliance subsidy program in Mexico.

[4] In most utility conservation programs, consumers opt-in to the programs (Joskow and Marron 1992). This opt-in participation creates a self-selection bias because participants in the program are likely to be different from nonparticipants.

find precisely estimated zero treatment effects in coastal areas, where the summer temperatures are moderate and the income levels are relatively high. To explore what drives the heterogeneity in the treatment effects, I estimate the interaction effects between the treatment variable and climate conditions and those between the treatment variable and income levels. The results from the regressions suggest that the treatment effect increases by 0.15 percentage points as the average temperature increases by 1 degree Fahrenheit and decreases by 0.029 percentage points as income levels increase by 1 percent. I also use air conditioner saturation data to show that higher air conditioner saturation rates result in larger treatment effects.

The asymmetric subsidy structure introduces the possibility that the response to the subsidy differs between households whose consumption is close to the target level of consumption and households whose consumption is far from the target. To test whether the asymmetric incentive creates a "giving-up" effect for consumers far from the target level, I estimate the quantile treatment effects on the changes in consumption. I find that most of the treatment effects come from households who are closer to the target level of consumption and that the treatment effect is not significantly different from zero for consumers who are far from the target level. This finding provides evidence that the asymmetry in the subsidy schedule weakens the incentive for conservation when compared to a simple Pigouvian tax.

An advantage of RD designs is that these require relatively weak identification assumptions to estimate LATE. However, RD designs generally do not provide average treatment effects (ATE) (Angrist and Rokkanen 2012). In my research design, the RD estimates come from customers who opened an electricity account about a year before the treatment period began. An important question is whether the treatment effect differs between my RD sample and the customers who opened accounts earlier. To address this point, I use a method that combines an RD design with three-way fixed effects. This method estimates the ATE with one additional identification assumption. In general, residential electricity consumers have a small positive trend in their electricity consumption after they open accounts. This trend is translated into a small trend component in my running variable for the RD design. To estimate the ATE, I assume that the positive trend in consumption is the same for customers who opened accounts on a certain date and those who opened accounts on the same day during the previous year. With this assumption, I can isolate the small positive trend from my RD design and estimate the ATE. I use this method to estimate the ATE for consumers who opened accounts 90 days, 180 days, 1 year, 2 years, 3 years, and 4 years before the eligibility cutoff date. I find that the difference between the ATE and the LATE is small and statistically insignificant. This finding suggests that the RD estimates are not significantly different from the treatment effects for consumers who opened accounts earlier than those included in my RD sample.

This paper's findings provide several important policy implications. First, asymmetric incentives created by subsidy programs are likely to weaken the incentives to reduce negative externalities. The evidence of zero treatment effects in coastal areas is consistent with the theoretical prediction that consumers do not respond to the asymmetric incentive *at all* if the price elasticity is below a cutoff level. Second, the difference between my RD estimates and naïve estimates of the treatment effect shows that the additionality problem is a central concern in evaluations of subsidy

programs (Joskow and Marron 1992, and Boomhower and Davis 2014). While my RD estimates show precisely estimated zero causal effects in coastal areas, the naïve estimates that ignore the additionality problem indicate that a significant number of consumers responded to the incentive. This result provides evidence that careful empirical analysis is critical for evaluating energy conservation programs (Allcott and Greenstone 2012). This is particularly important for recent US energy policy because public spending for energy conservation programs has been growing rapidly.[5] Finally, my analysis of the program's cost-effectiveness suggests that the heterogeneous treatment effects result in quite different levels of costs among coastal areas (94.5 cents per kWh reduction) and inland areas (2.5 cents per kWh reduction). However, because substantial rebates were paid to customers in the areas in which I find nearly zero treatment effects, the overall program cost is 17.5 cents per kWh reduction and $381 per ton of carbon dioxide reduction, which is unlikely to be sufficiently cost-effective to reduce negative externalities for a reasonable range of the social marginal cost of electricity.

I. Conceptual Framework

A. *The Asymmetric Incentive Structures of Conservation Subsidies*

In this section, I use a simple framework to characterize the theoretical predictions of consumer behavior in the presence of subsidies for energy conservation. Suppose that consumers have quasilinear utility functions $u(y_i, n_i) = v(y_i) + n_i$ for electricity consumption y_i and a numeraire consumption good n_i. Consumers with income I_i and electricity price p maximize $v(y_i) + I_i - py_i$ and consume y_0, where $v'(y_0) = p$.

Suppose that regulators consider that electricity price p does not properly reflect the social marginal cost of electricity. For example, p may not reflect the negative environmental externalities from generating electricity or may not reflect the higher marginal cost of supplying electricity when the system faces a supply shortage. The first-best solution is to increase the price by the cost of the externalities τ. That is, increasing the electricity price by τ lets consumers choose y^*, where $v'(y^*) = p + \tau$.[6]

However, regulators often prefer to implement conservation subsidies instead of introducing a price increase. In conservation subsidy programs, regulators first determine the rebate baseline consumption b_i, which is usually a function of consumer i's past consumption level. Then, they offer a subsidy based on b_i and y_i. For example,

    [6]The constant Pigouvian tax τ is the first-best solution given the assumption that electricity users are homogeneous in the externalities they generate. For example, the marginal cost of supplying electricity is generally higher in peak hours than during off-peak hours. Suppose that the marginal price p does not reflect this time-varying marginal cost. Then, if some customers tend to use electricity more during peak hours, the externalities are higher than those of others, and therefore, their τ has to be set higher than that of others.
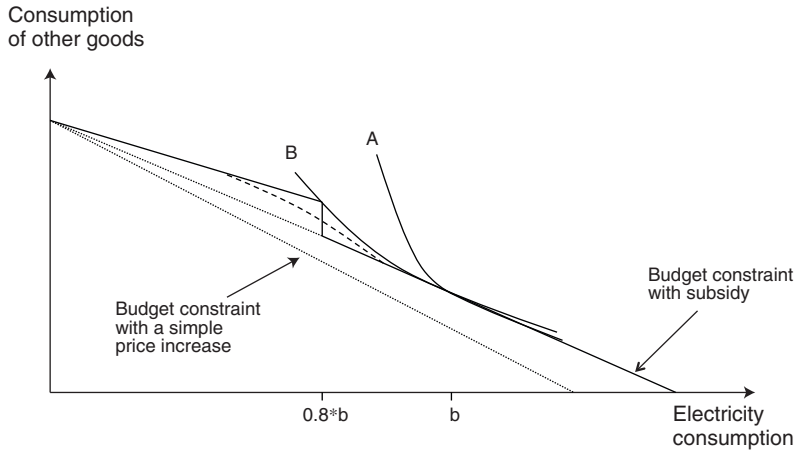
FIGURE 1. THEORETICAL PREDICTIONS

Notes: This figure illustrates the theoretical predictions of how consumers behave when they are faced with the conservation subsidy schedule of the California 20/20 rebate program. The subsidy makes a notch in the budget constraint. If consumers respond to the expected price because of the uncertainty in consumption, the budget constraint based on the expected price becomes the dotted smooth line.

Figure 1 illustrates the subsidy schedule of the California 20/20 rebate program.[7] Consumers receive a 20 percent discount on their summer electricity bills if they consume 20 percent less than their baseline, which is their consumption in the summer month during the previous year. This subsidy creates a notch in the household's budget constraint because it changes both the marginal and infra-marginal prices if consumers reach 80 percent of their baseline. In another type of conservation subsidy program, consumers receive a marginal subsidy for each unit of their conservation relative to a certain baseline level. Examples of this type of subsidy schedule include peak-time rebate programs in dynamic pricing (Wolak 2006, 2011, Faruqui and Sergici 2010, and Borenstein 2013). In these cases, the subsidy schedule creates a kink at the baseline rather than a notch. In both cases, consumers are subsidized for reducing consumption, but are not penalized for increasing consumption. This asymmetry creates important differences between such conservation subsidies and the first-best solution.

An inherent feature of such a subsidy schedule is that it creates asymmetry in the incentive to change consumption. In the case with the first-best solution, consumers have a simple price increase of τ, which gives all consumers the same change in the marginal incentive irrespective of where their consumption falls in the budget constraint shown in Figure 1. In contrast, the introduction of conservation subsidies creates different incentives for consumers depending on where they fall in the budget

---

[7]The figure shows the case that features a linear electricity price. In practice, residential electricity customers in California have increasing block pricing (Ito 2014). However, the insights from this section do not change with increasing block pricing because the rebate incentive changes both the infra-marginal and marginal prices by 20 percent. This implies that both the marginal and average prices change by 20 percent (note that residential electricity customers in California have zero or negligible fixed charges).

constraint, how price-elastic they are, and how much uncertainty in consumption they have.[8]

B. *Theoretical Predictions of Consumer Behavior*

I begin with a simple case that makes two assumptions which may not be realistic in practice: (a) consumers have no uncertainty about their consumption, and (b) their baseline b_i is set reasonably close to y_0, which is their optimal consumption in the absence of the subsidy incentive.

First, consumers do not respond to the subsidy incentive *at all* if the price elasticity in absolute value is smaller than a certain cutoff level. To illustrate this point, suppose that consumers have a quasilinear and iso-elastic utility function, $u(y_i, n_i) = \alpha_i \cdot \frac{y_i^{1+1/e}}{1 + 1/e} + n_i$, where α_i is a heterogeneous taste parameter and $e \leq 0$ is a constant price elasticity. In Figure 1, I illustrate two indifference curves, A and B, with $|e_B| > |e_A|$. With inelastic price elasticity e_A, households make no change in their consumption levels because the indifference curve does not reach the notch point. Price elasticity e_B is the minimum elasticity required for consumers to change their consumption. This prediction implies that the subsidy incentive induces no change in consumption when the price elasticity is smaller than $|e_B|$ in absolute value. This result contrasts with the result in the first-best solution. When consumers have a simple price increase of τ, the new budget constraint in the figure would have a steeper slope. Accordingly, all consumers would reduce consumption based on the new slope.

Second, given assumptions (a) and (b), there should be a bunching of consumers if the price elasticity is larger than $|e_B|$ in absolute value. In Figure 1, all indifference curves that have larger price elasticities than $|e_B|$ would have the optimal consumption at the notch point in the presence of the subsidy incentive.

However, in reality, assumptions (a) and (b) are unlikely to hold in the case of residential electricity demand. I begin with assumption (a). Electricity consumers have significant uncertainty about their monthly electricity consumption. When faced with this uncertainty, rational consumers do not respond to the exact non-linear budget constraint (Saez 1999, Borenstein 2009, and Ito 2014). Instead, they incorporate the uncertainty and respond to the expected price schedule, which is presented as the smoothed dotted line in Figure 1. The response to the smoothed schedule changes the first and second predictions above. First, the price elasticity's cutoff point has to be even larger than the standard case with no uncertainty. In the previous example, consumers who have price elasticity e_B no longer respond to the subsidy incentive. Second, because the smoothed schedule no longer has a notch, there can be no bunching of consumers even if the price elasticity is nonzero.

Finally, the subsidy's incentive can be further weakened if the rebate baseline b_i is set far below y_i^0, which is consumer i's optimal consumption in the absence of the subsidy. Conservation subsidy programs usually do not adjust b_i for changes

---

[8]Borenstein (2013) provides a detailed description of similar problems for peak-time rebate programs in dynamic electricity pricing.

in weather or idiosyncratic shocks to each consumer. As a result, if the base year's weather is more moderate than the target year's weather, consumers are more likely to have harder baselines to reach. Similarly, if consumers experience idiosyncratic negative shocks in consumption in the base year, they have harder baselines to reach. This endogenous baseline can introduce a "giving-up" effect because consumers whose electricity usage is far above the baseline consider the subsidy unachievable (Wolak 2010, and Borenstein 2013). Conversely, when consumers have consumption shocks in the opposite direction, b_i can be closer to y_i^0. In this case, they reduce consumption by less than 20 percent because they are now closer to the cutoff point. That is, the endogenous baseline makes the marginal incentive depend on where the rebate baseline falls in the household's budget constraint. These predictions contrast with the prediction for the first-best solution, which produces the same marginal incentive for all consumers.

In the following sections, I empirically test these theoretical predictions by applying a regression discontinuity design to the California 20/20 rebate program. In the next section, I describe the research design and data used in my empirical analysis.

II. Research Design and Data

This section provides the institutional background and research design. First, I provide a brief history of the California 20/20 electricity rebate program. Second, I discuss evidence from existing studies and their empirical challenges. Finally, I describe how I address these challenges using a regression discontinuity design to analyze the effects of the California 20/20 rebate program in 2005.

A. *The California* 20/20 *Electricity Rebate Program*

The California 20/20 electricity rebate program was originally implemented by Governor Gray Davis during the 2001 California electricity crisis.[9] To prevent rolling blackouts, the California Public Utility Commission (CPUC) ordered the state's three largest investor-owned electric utilities (IOUs)—Pacific Gas and Electric (PG&E), Southern California Edison (SCE), San Diego Gas and Electric (SDG&E)—to offer customers financial incentives to reduce electricity consumption. Every month during June, July, August, and September in 2001 and 2002, customers received a 20 percent discount on their monthly electricity bill if their consumption was 20 percent lower than their consumption during the same month in 2000. With a slight change in the scheme, the CPUC ordered the same program in 2005. The original month-based rule was replaced by a summer-based rule. Customers received a 20 percent discount on their entire summer bills if their consumption over the summer months was 20 percent lower than their consumption over the summer months in 2004.[10] This rebate program was among the largest electricity conservation rebate programs in the

TABLE 1—CHANGES IN CUSTOMER-LEVEL CONSUMPTION WHEN THERE WAS NO REBATE PROGRAM

Year	Weather conditions	Median of the percent change in customer-level consumption	Percent of customers who had 20 percent or more reductions
2003 to 2004	Cooler in 2004	−1.7 percent	14.3 percent
1999 to 2000	Warmer in 2000	7.7 percent	6.8 percent

Notes: This table shows the distribution of the percent change in customer-level electricity consumption between two summers in which no rebate program was in effect. I use customer-level monthly consumption data for summer billing months (June, July, August, and September billing months) in Southern California Edison (SCE). Note that although there was the California electricity crisis in 2000, SCE customers did not experience a price spike because their retail rates were capped (Ito 2014).

United States in terms of its expenditure and the number of customers who received rebates. In the 2005 program, about 8 percent of residential customers of the three IOUs received a rebate. The total rebate expenditure was about $25 million, excluding marketing and administrative costs.[11]

Despite the substantial expenditure, the program's effectiveness was highly controversial. Proponents have claimed that its simplicity makes it straightforward for customers to undertake energy conservation.[12] It is politically more favorable to offer a rebate program rather than raise electricity prices because the economic burden is much less salient to customers.[13] However, the 20/20 program scheme created two key concerns.[14] First, the program did not incorporate differences in weather between the base and target years. If the target year happened to be cooler than the base year, many customers received a rebate simply because of the change in weather. Second, even if there was no significant difference in weather between the two years, many customers received a rebate because of random fluctuations in their electricity consumption. For example, customers who had a friend visit during the base year or customers who traveled during the target year could reduce their electricity consumption in the target year by 20 percent for reasons unrelated to their conservation efforts.

Table 1 shows data related to the two concerns. I use household-level consumption data to calculate the fraction of customers who reduced their summer electricity usage by more than 20 percent between summers when there was no rebate program. From 2003 to 2004, the median customer reduced consumption by 1.7 percent because in 2004 the summer was cooler than in 2003. More importantly, 14.3 percent of customers reduced their consumption by more than 20 percent. This statistic suggests that 14.3 percent of customers would have received a rebate for reasons unrelated to their conservation efforts if a rebate program had been offered

---

[11] Table A.1 in the online Appendix shows more details about the scale of the 2005 rebate program. More customers received at least one rebate in 2001 and 2002 because the program was month-based. Reiss and White (2003) report that about 39 percent of SDG&E customers' monthly bills qualified for a rebate in June, July, August, and September 2001. For the same 2001 rebate program, Goldman, Barbose, and Eto (2002) find that about 33 percent of consumers received a rebate.

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

[13] Although the rebate expenditure is eventually paid by customers through future price increases, this burden is usually much less salient than raising the electricity price.

[14] See Faruqui and George (2006) for details.

in 2004. The second row provides the same statistic for the warmer summer in 2000 relative to 1999. As a result, the median customer increased consumption by 7.7 percent. However, even in this case, 6.8 percent of customers reduced consumption by 20 percent. This evidence implies that random fluctuations in electricity consumption necessarily create substantial rebate expenditures in this program design.[15]

B. *Using a RD Design to Address Empirical Challenges*

In general, there are two fundamental challenges in evaluating the causal effect of rebate programs. First, many rebate programs are offered to self-selected customers. Evaluating this type of program is difficult because households that self-select into the program are likely to be different from other households in terms of observable and unobservable factors. Second, when a rebate program is offered to all customers, it eliminates self-selection bias but creates another challenge—there is no clean control group because all customers are affected. This lack of a clean control group makes it difficult to distinguish a program's causal effect from other factors unrelated to the program.[16]

To address these challenges, I exploit a discontinuity in the eligibility rule for the California 20/20 rebate program in 2005. To be eligible for the program, households had to open their electricity account by the program's eligibility cutoff date in 2004. For example, the eligibility cutoff date for Southern California Edison customers was June 5, 2004. Customers who opened electricity accounts on or before June 5, 2004 received a notice in the spring 2005 and were automatically enrolled in the program. Customers who began their service after the cutoff date (e.g., June 6, 2004) were not eligible for the program in 2005 and did not receive the notice.[17]

The eligibility rule includes two additional key features. First, it was impossible for customers to anticipate the 2005 rebate program when they started their electricity service in 2004 since the program was only announced in the spring of 2005. It was, therefore, impossible for customers to strategically choose their start date in consideration of the rebate program. Second, all eligible customers automatically participated in the program, which eliminated any self-selection bias. Finally, the electric utilities that administered the program strictly enforced the rules without exception.

The discontinuous eligibility rule generated an essentially random assignment of the program for customers who opened their accounts near the cutoff date. The program rules allow me to use a regression discontinuity design to estimate the program's causal effect given the assumption that the conditional expectation of the

---

[15] This evidence implies that it is misleading to make conclusions about the program's effectiveness simply by calculating the number of customers receiving a rebate or the total reduction in consumption achieved by these customers. Yet, such statistics are often used in utility company reports and newspaper articles.

[16] For example, researchers need to control for changes in the weather, changes in the electricity price, other conservation programs, and macroeconomic shocks. Previous studies acknowledge this difficulty in evaluating the original 20/20 rebate program in 2001 and the later program in 2005. For example, Reiss and White (2008) and Goldman, Barbose, and Eto (2002) note that it is particularly challenging to control for the effects of other conservation programs that were active during their study periods. For evaluating the 20/20 rebate program in 2005, Wirtshafter Associates (2006) uses survey data to adjust for factors unrelated to the program. The adjustment results in a wide range of the estimated effects: the cost per kWh savings range from 29 cents to $1 per kWh.

[17] Figure A.1 in the online Appendix graphically explains how this eligibility rule was applied. The cutoff date was June 1, 2004 for PG&E customers and June 30, 2014 for SDG&E customers.

outcome variable is smooth at the cutoff date. In the empirical analysis section, I provide more details about the empirical strategy.

C. *Data*

The primary data for this study come from the panel data of customer-level monthly electricity billing records for the three largest investor-owned electric utilities in California. Under a confidentiality agreement, Pacific Gas and Electric (PG&E), Southern California Edison (SCE), and San Diego Gas and Electric (SDG&E) provided the complete billing history for essentially all residential customers in their service areas.[18] I focus on SCE in this paper and present the results for the other two utilities in the online Appendix. The conclusions are consistent across all three utilities.[19]

The monthly records include each customer's account number, premise ID, billing start and end dates, monthly consumption, monthly bill, tariff type, climate zone, and nine-digit zip code.[20] The data also include each customer's account opening and closing dates, which are key variables for my RD design. Each day in California, about 10,000 residential customers open electricity accounts—thus there are a substantial number of observations for fairly narrow bandwidths. I use the customers who opened their electricity accounts within 90 days before and 90 days after the cutoff date for my main estimation and examine the robustness with different bandwidth choices.

The billing data do not include each customer's address and demographic information. To obtain demographic information, I match the nine-digit zip codes to census block groups from the 2000 US census data. I also use daily weather data from the Cooperative Station Dataset published by the National Oceanic and Atmospheric Administration's National Climate Data Center.[21] The dataset includes the daily minimum and maximum temperatures recorded at 370 weather stations in California. I match each household's zip code with the nearest weather station by following the matching mechanism in Aroonruengsawat and Auffhammer (2011) and Chong (2012). Finally, I collect air conditioner saturation data from the 2003 Residential Appliance Saturation Study (RASS) to examine whether the program's treatment effects vary by air conditioner saturation.[22]

---

[18] A very small number of customers are not individually metered in this area. The billing datasets include only individually metered customers.

[19] The three utilities provided a similar rebate program but the programs differed substantial differences in two elements. First, the eligibility cutoff dates were different. Second, the way they calculated the outcome variable was different. Because of these differences, I conduct my analysis separately for each utility. Details are provided in the Appendix.

[20] During my sample period, residential customers in California did not have smart meters. Therefore, it was not possible for customers to see their hourly consumption data.

[21] I thank Anin Aroonruengsawat, Maximilian Auffhammer, and Howard Chong for sharing the data.

[22] RASS was funded and administered by the California Energy Commission and is based on 21,920 individually metered California customers in 2003. The variable of air conditioner saturation provides the ratio of customers who own air conditioners at the five-digit zip code level.

III. Empirical Analysis and Results

In this section, I first use a regression discontinuity design to estimate the program's local average treatment effect. Second, I examine heterogeneity in the treatment effects by investigating how income and weather affect the treatment. Third, I estimate whether the nonlinearity in the subsidy schedule induces a "giving-up" effect for consumers who are far from the 20 percent target level. An important question raised from RD estimation in general is whether the treatment effects are different in my RD sample and the overall population. In the final part of this section, I use a method that combines an RD design with three-way fixed effects to estimate the average treatment effect.

A. *A Regression Discontinuity Design to Estimate LATE*

I define customer i's natural log of electricity consumption by y_{it} for billing month t before and during the California 20/20 program. Let $D_i = 1\{i \in$ treatment group$\}$, $D_t = 1\{t \in$ treatment period$\}$, and $D_{it} = D_i \cdot D_t$. If the treatment is randomly assigned, the program's average treatment effect can be obtained by a fixed effect estimation,

$$(1) \qquad\qquad y_{it} = \alpha \cdot D_{it} + \theta_i + \lambda_t + u_{it},$$

by the ordinary least squares (OLS), where θ_i is the customer fixed effects, λ_t is the time fixed effects, and u_{it} is an error term. However, the treatment assignment was not random in the California 20/20 program. Instead, it was assigned by an eligibility rule, $D_i = 1\{x_i \leq 0\}$, where x_i is customer i's account opening date relative to the eligibility cutoff date. Because the treatment variable is a function of x_i, the estimation in equation (1) is biased if u_{it} is correlated with x_i. With the RD design, I can explicitly control for the smooth relationship between the running variable x_i and the dependent variable and can estimate the program's local average treatment effects by using the discontinuity of D_{it} in x_i:

$$(2) \qquad\qquad y_{it} = \alpha \cdot D_{it} + f_t(x_i) + \theta_i + \lambda_t + \eta_{it}.$$

The identification assumption is that the error term η_{it} has to be uncorrelated with treatment D_{it} conditional on a smooth control function $f_t(x_i)$ and other covariates.

The customer fixed effects absorb the time-invariant effects of x_i. Therefore, the potential confounding factors are the time-*varying* effects of x_i. Consider $\tilde{y}_{it}$, which is consumption demeaned by customer fixed effects. In consumption data for residential electricity, customers have a general tendency to gradually increase their electricity consumption after opening electricity accounts. This tendency creates a very small and smooth relationship between $\tilde{y}_{it}$ and x_i. Hence, I use a smooth control function $f_t(x_i)$ to control for the relationship. Imbens and Lemieux (2008) describe two approaches to specifying $f_t(x_i)$. The first approach is to include a flexible parametric function. The second approach uses a local linear regression with a triangular kernel to put more weight on data closer to the cutoff point. I use the first approach

for my main result, then use the second approach to show that my estimation is robust for both approaches. To avoid misspecifing $f_t(x_i)$ as much as possible, I focus on the data close to the cutoff date. For my main result, I use customers who opened their accounts within 90 days before or after the cutoff date. I also use 60-day and 120-day bandwidths to show the robustness.

B. *Testing the Validity of the Regression Discontinuity Design*

A threat to the validity of RD designs is that the identification assumption is violated if there is self-selection at the cutoff, although this is unlikely to be the case for my research design. In summer 2004, no households knew that the California 20/20 program would be implemented in the following summer. Therefore, there was no way for customers to self-select by strategically choosing the opening date for their electricity account.

Still, it is important to examine if there is a discontinuous difference between customers around the cutoff date. To assess the validity of the RD design, I first plot the number of new accounts opened per day in Figure 2. The horizontal axis is the account opening date relative to the eligibility cutoff date, which is July 5, 2004. Every dot shows the mean number of new accounts per day over the 15-day bandwidth. Every day about 1,500 customers opened accounts with SCE. The solid line shows the local linear fit and the dashed lines are the 95 percent confidence intervals. Over the 90-day period, there is a slight upward trend in the number of new accounts, although the slope is not statistically different from zero. The figure shows that there is no discontinuous jump at the cutoff date.

Figure 2 plots the customer characteristics against the account opening date relative to the eligibility cutoff date. I match the nine-digit zip codes in the billing data with the census block group to obtain the demographic and housing characteristics. The figures include the mean over the 15-day bandwidth, the local linear fit, and its 95 percent confidence intervals.[23] None of the three variables show a statistically significant discrete jump at the cutoff date.

C. *RD Estimates of the Program's Treatment Effects*

In RD estimation, graphical analyses play an important role in quantifying the magnitudes of the treatment effects as well as testing the validity of the identification strategy. I begin by presenting graphical analysis and then show the regression results. In California, summer electricity consumption differs between coastal and inland climate areas. In coastal areas, consumers do not have or rarely use air conditioners because the summer temperatures are moderate. In inland areas, however, consumers regularly use air conditioners because the summer temperatures are persistently high.[24] To analyze the rebate incentive's heterogeneous treatment effects,

---

[23] Because the variables are from the 2000 US census at the census block group level, I cluster the standard errors at the census block group level.

[24] As a reference, I present cooling degree days (CDD) at the five-digit zip code level for August 2005 in Figure A.3 in the online Appendix.

Panel A. Number of new accounts opened per day

Panel B. Median year structure built

Panel C. Median customer income ($)
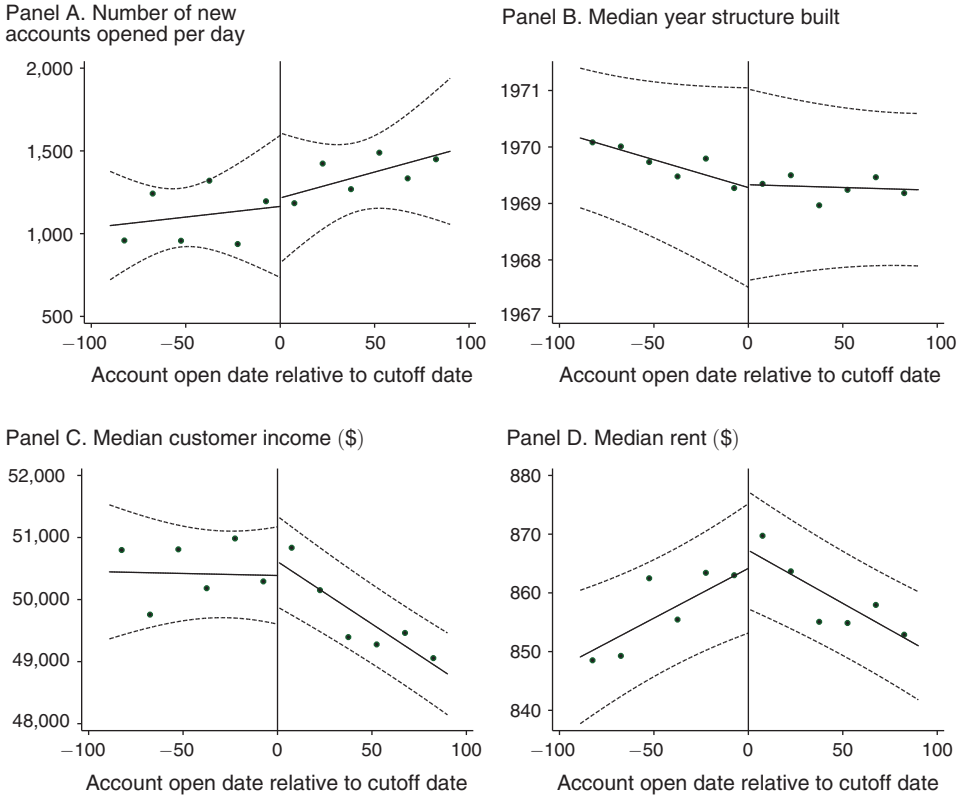
Panel D. Median rent ($)

FIGURE 2. TESTING THE VALIDITY OF THE REGRESSION DISCONTINUITY DESIGN

Notes: The horizontal axis shows the account open date relative to the cutoff date of the program eligibility, which was June 5, 2004. Each dot shows the local mean with a 15-day bandwidth. The solid line shows the local linear fit and the dashed lines present the 95 percent confidence intervals. The confidence intervals for the fitted lines for variables from census data are adjusted for clustering at the census block group level.

I begin by examining the RD estimates for coastal and inland climate zones separately, according to the climate zones defined by SCE. Then I use pooled data from all the climate zones to investigate whether heterogeneous treatment effects can be explained by differences in observable variables such as climate conditions or household income levels.

Figure 3 presents a graphical analysis of the RD estimation based on the electricity usage recorded on account statements for September, the last month of the treatment period. Using the data for before and during the treatment period, I first estimate the demeaned consumption by $\tilde{y}_{it} = y_{it} - \bar{\theta}_i$. I then calculate the local mean of $\tilde{y}_{it}$ for each 15-day bandwidth over the running variable x_i. The local means are presented as dots in the figure. Finally, I fit a local linear regression and a quadratic function to estimate $f_t(x_i)$ for each side of the cutoff date. The dashed line is the local linear fit and the solid line is the quadratic fit. On the horizontal axis, the treatment group is on the left-hand side of the cutoff date because customers who opened accounts before the cutoff date participated in the rebate program. Therefore, if the rebate incentive had an effect, there should be a discontinuous jump in the outcome variable at the cutoff point.
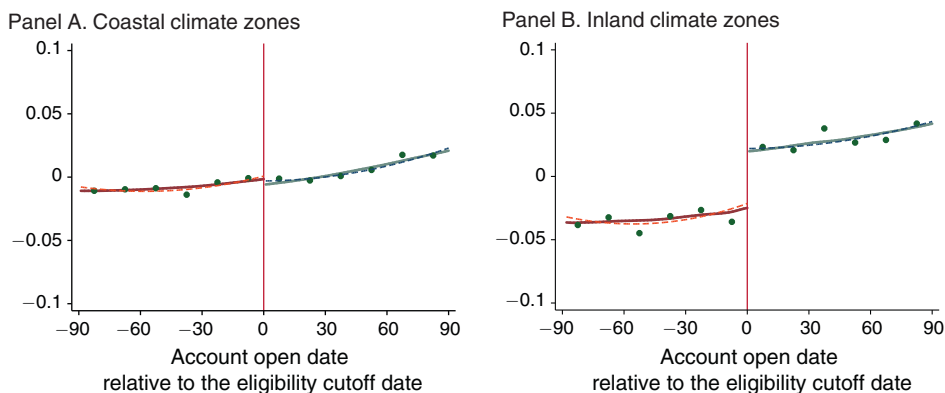
FIGURE 3. THE EFFECTS OF SUBSIDY INCENTIVES ON ENERGY CONSERVATION IN SEPTEMBER 2005

Notes: This figure presents the RD estimates of the effect of subsidy incentives on energy conservation for the September 2005 billing period. The horizontal axis shows customers' account opening date relative to the cutoff date for program eligibility, which was June 5, 2005. The vertical axis presents the log electricity consumption in September 2005, in which I subtract customer fixed effects by using the consumption data from billing months before the treatment period. Each dot shows the local mean with a 15-day bandwidth. The solid line shows the local linear fit and the dashed line shows the quadratic fit.

Figure 3 provides important insights. First, there is a slight upward trend in $\tilde{y}_{it}$ over x_i, which is the account opening date relative to the cutoff date. This upward trend comes from the general tendency in residential electricity consumption data already noted—customers tend to increase their usage gradually after they open their accounts. Because ignoring this relationship creates an estimation bias for the treatment effect, it is important to control for the trend. The fitted lines of the local linear regression and the quadratic regression over x_i indicate that the RD estimates are likely to be robust between the local linear regression and the quadratic regression.

Second, panel B shows evidence that the rebate incentive had a significant effect on lowering electricity consumption in the inland climate zones. There is a clear discontinuous change in consumption between the treatment and control groups at the cutoff point. Visually, the treatment group's usage is about 5 percent less than the control group's consumption. In contrast, panel A suggests that the rebate incentive did not significantly alter consumption in the coastal climate zones. There is no discontinuous change in consumption between the treatment and control groups at the cutoff point.

Table 2 presents the RD estimates for the effect of subsidy incentives on energy conservation. In columns 1 and 3, I estimate the program's overall treatment effect during the entire treatment period.[25] In columns 2 and 4, I allow the treatment effects to differ for each billing month in the treatment period. I report RD estimates

---

[25] The actual treatment billing months for the 20/20 rebate program were the June, July, August, and September billing months. However, because of the billing cycle systems, many May billing days fall in the calendar month of June. For example, if a customer's May billing cycle starts on May 31, most of the billing days fall in June. If customers focus on the calendar months instead of their billing cycles, a treatment effect can appear in their consumption for the May billing month. Therefore, I include the May billing month as a treatment period. To strictly focus on the actual treatment billing months, one can see the treatment effects for June through September.

TABLE 2—RD ESTIMATES OF THE EFFECT OF REBATE INCENTIVES ON ENERGY CONSERVATION

	Coastal climate zones		Inland climate zones	
	(1)	(2)	(3)	(4)
Treatment effect	−0.001		−0.042	
	(0.002)		(0.013)	
Treatment effect in May		0.003		−0.034
		(0.003)		(0.015)
Treatment effect in June		−0.001		−0.055
		(0.003)		(0.017)
Treatment effect in July		0.004		−0.041
		(0.004)		(0.019)
Treatment effect in August		−0.003		−0.041
		(0.004)		(0.018)
Treatment effect in September		−0.004		−0.056
		(0.003)		(0.016)
Observations	2,540,472	2,540,472	208,537	208,537

Notes: This table shows the RD estimates of the effect of rebate incentives on energy conservation. The dependent variable is the log of electricity consumption. I estimate equation (2) with a 90-day bandwidth and quadratic functions to control for the running variable. The standard errors are clustered at the customer level to adjust for serial correlation.

TABLE 3—ROBUSTNESS CHECKS: ALTERNATIVE BANDWIDTHS AND SPECIFICATIONS

	Coastal climate zones			Inland climate zones		
	(1)	(2)	(3)	(4)	(5)	(6)
Treatment effect in May	0.004	0.003	0.005	−0.034	−0.039	−0.029
	(0.004)	(0.003)	(0.004)	(0.015)	(0.014)	(0.017)
Treatment effect in June	−0.002	−0.001	−0.003	−0.055	−0.059	−0.05
	(0.004)	(0.004)	(0.004)	(0.017)	(0.016)	(0.019)
Treatment Effect in July	0.004	0.005	0.005	−0.041	−0.039	−0.042
	(0.004)	(0.004)	(0.005)	(0.019)	(0.017)	(0.022)
Treatment effect in August	−0.004	−0.005	−0.003	−0.036	−0.034	−0.035
	(0.004)	(0.004)	(0.004)	(0.018)	(0.016)	(0.020)
Treatment effect in September	−0.005	−0.003	−0.004	−0.056	−0.053	−0.052
	(0.003)	(0.004)	(0.004)	(0.016)	(0.015)	(0.018)
Controls for $f(x)$	Local linear	Quadratic	Quadratic	Local linear	Quadratic	Quadratic
Bandwidth	90 days	120 days	60 days	90 days	120 days	60 days
Observations	2,540,472	3,325,388	1,707,589	208,537	237,264	162,067

Notes: This table shows RD estimates with different bandwidth choices and alternative controls for the running variable. The dependent variable is the log of electricity consumption. The standard errors are clustered at the customer level to adjust for serial correlation.

with 90-day bandwidths and quadratic controls for $f_t(x_i)$. Using different bandwidths and the local linear regression do not change my results, as I show in the Table 3. To adjust for serial correlation in the electricity consumption data, I cluster standard errors at the customer level.[26]

---

[26] In fact, ignoring the serial correlation produces very small standard errors.
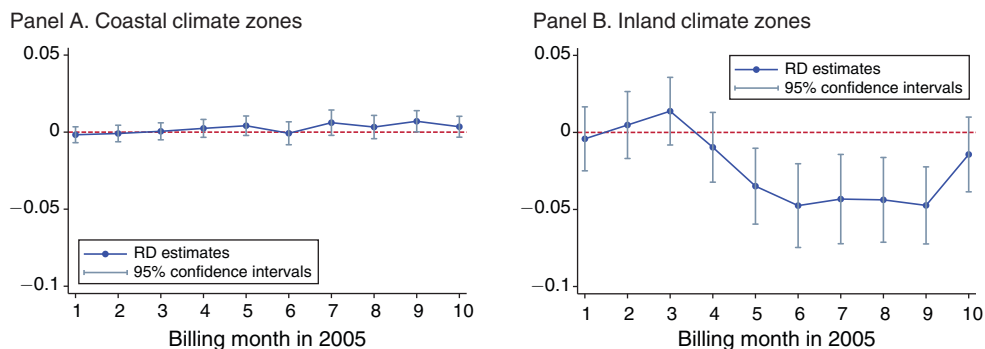
FIGURE 4. THE DIFFERENCE IN LOG CONSUMPTION BETWEEN TREATMENT AND CONTROL GROUPS

Notes: This figure presents the RD estimates of the difference in log consumption between the treatment and control groups. Customer fixed effects are subtracted by using consumption data before January 2005. I use a 90-day bandwidth and quadratic controls for the trend of the running variable, which is the same specification used to obtain my main estimation results shown in Table 2.

In coastal climate zones, the treatment effects are essentially zero with tight standard errors. Because of the tightly estimated point estimates, the 95 percent confidence intervals do not include 1 percent treatment effects, suggesting that the 20/20 program did not have a significant effect on customers in coastal climate zones. In contrast, the subsidy incentive had a significant effect on electricity consumption in the inland climate zones. The overall treatment effect is about 4 percent and the treatment effect of each month ranges between 4 percent and 5 percent.[27]

Because I have data for several months before the treatment period began, a useful robustness check is to produce the RD estimates for the billing months before the treatment period. Figure 4 presents the RD estimates of the difference in log consumption between the treatment and control groups for the billing months of January 2005 through October 2005. In the coastal climate zones, the RD estimates are essentially zero both before and during the treatment period. In inland climate zones, the RD estimates are not statistically different from zero before the treatment period. In contrast, the estimates during the treatment period suggest that the customers in the treatment group reduced their usage by about 5 percent. This figure provides evidence that the reduction in consumption is unlikely to come from factors unrelated to the program.

Another important robustness check is to examine how the choice of bandwidths and the method to control for $f_t(x_i)$ affects the estimates.[28] In Table 3, I present RD

---

[27] I find similar results for PG&E and SDG&E customers, although the definitions of the running variable and outcome variables differ between the three companies (see the Appendix for details). In PG&E, I find nearly zero effects for coastal customers and 3 percent to 4 percent effects for inland customers. In SDG&E, the majority of customers are in coastal areas, and I find a nearly zero effects for them.

[28] Although many studies, including Lee and Lemieux (2010), recommend reporting a number of specifications to illustrate the robustness of the results, another approach is to use the Akaike information criterion (AIC), which provides guidance on the choice for the polynomial. In my RD estimation, the quadratic controls produce the lowest AIC, although it is not substantially different from linear or third-order polynomial control functions. This is because the relationship between the running variable and the outcome variable is smooth and approximately linear in the data (see Figure 3).

estimates with 60-day and 120-day bandwidths and the RD estimates with the local linear regression. Consistent with the suggestive evidence presented in Figure 3, these estimates are not sensitive to the bandwidth choice or to the method used to control for $f_t(x_i)$. While the standard errors change slightly when using different bandwidths, the RD estimates are fairly stable between different bandwidth choices.[29] In addition, using the local linear regression instead of quadratic controls does not significantly change the estimates.

D. *Heterogeneity in the Treatment Effect*

Income, *Climate Conditions*, *and Air Conditioner Saturation*.—In the previous section, I find significant treatment effects for inland customers and do not find significant effects for coastal customers. This section explores what drives the rebate program's heterogeneous treatment effects. In particular, I examine whether climate conditions, income differences, and air conditioner saturation can explain the heterogeneous treatment effects.

A significant difference between inland and coastal California is the summer climate. Summer temperatures are persistently high in inland areas but quite moderate in coastal areas.[30] Inland customers, therefore, typically use air conditioners (AC) throughout the summer, while many coastal customers rarely use AC. It is likely to be a challenge for customers who do not use AC to reduce their summer electricity consumption by 20 percent. In contrast, if customers constantly use AC, a 20 percent reduction can be achieved by changing the temperature settings or the length of their AC usage.

Demographic characteristics are another significant difference between inland and coastal California. For instance, income levels tend to be higher in the coastal areas than the inland areas. Many previous studies of residential electricity demand find slightly larger price elasticity estimates for low-income customers (Reiss and White 2005). Because the 20/20 rebate program is essentially a price-discount rebate program, lower-income customers may be more likely to respond to the incentive if their price elasticity is larger than that of higher-income customers.

To examine how climate conditions and income levels affect the program's treatment effects, I pool data from all climate zones and estimate the interaction effects. First, I calculate the average temperature at the nine-digit zip code level by calculating the mean of the daily mean temperature for the summer days in 2004 and 2005. Second, I obtain the median per capita income at the census block group level from the 2000 US census. Column 1 of Table 4 shows the RD estimate of the interaction term between the treatment variable and the average temperature in degrees Fahrenheit. The estimate implies that the treatment effect increases by

---

[29] As Figure 3 suggests, using even narrower bandwidths (30-day bandwidths, for example) also does not change the RD estimates. Narrower bandwidths result in larger standard errors compared to the estimates with the baseline bandwidths.

[30] Figure A.3 in the online Appendix shows the cooling degree days (CDD) by five-digit zip code areas. For example, the average daily maximum temperatures are 69, 71, 73, 74, and 74 degrees Fahrenheit for May, June, July, August, and September for Santa Barbara (which is in a coastal climate zone), which are quite moderate. In contrast, they are 96, 104, 108, 107, and 102 degrees Fahrenheit for May, June, July, August, and September for Palm Springs (which is in an inland zone), which are quite high.

TABLE 4—RD ESTIMATES INTERACTED WITH INCOME, CLIMATE, AND AIR CONDITIONER SATURATION

	(1)	(2)	(3)	(4)
Treatment	0.095	−0.297	−0.199	−0.478
	(0.051)	(0.055)	(0.077)	(0.056)
Treatment × avg.temp.(°F)	−0.0015		−0.0016	
	(0.0007)		(0.0008)	
Treatment × ln(income)		0.029	0.031	0.044
		(0.005)	(0.005)	(0.003)
Treatment × air conditioner				−0.014
				(0.005)
Observations	2,749,009	2,749,009	2,749,009	2,749,009

Notes: This table presents the RD estimates of the effect of rebate incentives on energy conservation interacted with income, climate conditions, and air conditioner saturation. The dependent variable is the log of electricity consumption. I use a 90-day bandwidth and quadratic controls for the trend in the running variable. Income is at the census block group level. Average temperature and air conditioner saturation (the ratio of customers who own air conditioners) are at the five-digit zip code level. The standard errors are clustered at the customer level to adjust for serial correlation.

0.15 percentage points with an increase in the average temperature of 1 degree Fahrenheit. The estimate in Column 2 implies that the treatment effect *decreases* by 0.029 percentage points with a 1 percent increase in income.[31] These two interaction effects remain the same when both terms are included in the regression in column 3. Finally, I examine the interaction effect with the air conditioner saturation. The 2003 RASS data provide the proportion of customers who own air conditioners at the five-digit zip code level. Column 4 shows evidence that higher AC saturation rates result in larger treatment effects. Overall, these results indicate that climate conditions, income levels, and air conditioner saturation have statistically significant effects on the program's treatment effect.

Nonlinearity in the Subsidy Schedule.—A theoretical prediction in Section II implies that the nonlinearity in the subsidy schedule may induce a "giving-up" effect. Even if a consumer has a large price elasticity, the consumer may not respond to the incentive at all if the consumption is far from the cutoff point required to earn the rebate. This implies that the treatment effect may not come from all consumers equally. Consider Δy_{it}, the change in log consumption from 2004 to 2005. If there is a giving-up effect, I expect the different parts of the distribution of Δy_{it} will have different treatment effects. In particular, I would expect there to be no change in consumption for higher percentiles in the consumption distribution because the treatment intervention is likely to have no effect on these percentiles if there is a giving-up effect. To test the prediction, I estimate the quantile treatment effects in my RD design (Frandsen, Frölich, and Melly, 2012):

$$(3) \qquad\qquad \Delta y_{it} \,=\, \alpha \cdot D_i + f(x_i) + \lambda_t + \epsilon_{it},$$

[31] The income variable is the median income at the census block group level. Each census block group in my sample consists of about 500 households. The income variable from the census data may have a measurement error in the sense that I observe the median household income instead of each household's income. This measurement error implies that the estimated interaction effects in the paper are possibly underestimated. If the income variable includes a classical measurement error, my estimate of the interaction effect will be attenuated toward zero. Therefore, my estimate can be considered as a lower bound.

Table 5—Quantile Regressions on the Change in log Consumption for Inland Climate Zones

	p5	p10	p25	p50	p75	p90	p95
Treatment	0.034	−0.099	−0.078	−0.007	−0.020	−0.019	−0.025
	(0.056)	(0.035)	(0.018)	(0.018)	(0.019)	(0.033)	(0.063)
Observations	37,914	37,914	37,914	37,914	37,914	37,914	37,914

Notes: This table presents the quantile RD estimates of the effect of rebate incentives on energy conservation. The dependent variable is the change in the log of electricity consumption from 2004 to 2005. I use a 90-day bandwidth and quadratic controls for the trend in the running variable. The standard errors are clustered at the customer level to adjust for serial correlation.

where $D_i = 1\{x_i \leq c\}$. Note that this is a quantile regression on the changes in consumption—it estimates how the treatment intervention changes the distribution of the changes in consumption. I estimate the equation for the changes in consumption during August and September for customers in inland climate zones.

Table 5 presents the quantile treatment effects for the 5th, 10th, 25th, 50th, 75th, 90th, and 95th percentiles.[32] The treatment effect is larger in the lower tails of the distribution for the change in consumption. In contrast, the treatment effect is not statistically significant in the median and the higher tails of the distribution. This evidence suggests that the treatment effect mainly comes from the lower tail of the distribution and that customers whose consumption is relatively far from the 20 percent reduction target are likely to give up responding to the incentive. In theory, the rebate incentive can increase consumption in the left tail of the distribution because if consumers are sure to receive a rebate, the rebate program acts as a decrease in price for these consumers. I find a positive point estimate for the fifth percentile, but the estimate is too noisy in the tails of the distribution to be statistically different from zero.

E. *Potential Long-Run Effects*

While the rebate program was in effect only during the summer of 2005, the program may have had a long-run impact on future summer electricity consumption.[33] Customers may have learned how to become more energy efficient and then reduced their subsequent electricity consumption during future summers. Another possibility, though probably less plausible, is that the program may have induced marginal consumers to undertake capital investments in energy-efficiency that pay off over multiple years. Because the electricity billing data do not include information on a household's durable goods purchases, I cannot distinguish between the two potential effects. However, I can test for the overall long-run effects by using consumption data for the treatment and control groups in the summers of 2006, 2007, and 2008.

I use the RD estimation in equation (2) with one modification. Instead of using summer 2005, I use subsequent summers as a treatment period. For example, to

---

[32] Note that the customers with a 20 percent reduction in consumption fall near the 25th percentile.
[33] I thank a referee for suggesting this analysis. The billing data allow me to test the potential long-run effects for 2006, 2007, and 2008, which is the final year of the available data.

TABLE 6—POTENTIAL LONG-RUN EFFECTS

	Coastal				Inland			
	2005	2006	2007	2008	2005	2006	2007	2008
Treatment effect	−0.001	0.001	−0.002	0.002	−0.042	−0.040	−0.048	−0.043
	(0.002)	(0.003)	(0.004)	(0.004)	(0.013)	(0.018)	(0.021)	(0.022)

Notes: This table shows the RD estimates of the potential long-run effect of rebate incentives on energy conservation. The dependent variable is the log of electricity consumption. The treatment variable is the interaction of the treatment group and the summer of 2006, 2007, and 2008, which are one, two, and three years after the rebate program. The standard errors are clustered at the customer level to adjust for serial correlation.

estimate the effect for the summer of 2006, I define $D_{it} = 1$ if $i \in$ treatment group and $t \in$ the summer of 2006, and $D_{it} = 0$ otherwise. From the dataset used for the main RD estimation, I exclude the data for the summer of 2005 and include the data for the summer of 2006. The rest of the procedure is the same as for the main RD estimation.

Table 6 shows the rebate program's long-run effects. Consistent with the main RD estimation, I find nearly zero treatment effects for coastal customers. In contrast, I find long-run conservation effects for inland customers, resulting in about a 4 percent reduction in consumption. Although the long-run effects are slightly different from the short-run effects, these are not statistically different at conventional significance levels. The results imply that for inland customers the rebate program induced persistent conservation effects.[34]

F. *Using an RD Design with Three-Way Fixed Effects to Estimate ATE*

An advantage of RD designs is that these require relatively weak identification assumptions to estimate local average treatment effects. However, RD designs generally do not provide average treatment effects. For example, my RD estimates are the LATE for customers who opened accounts a year before the program began. Is the LATE different from the ATE for customers who opened accounts earlier? This is an important question because if possible the policy should be evaluated based on the entire affected population, and because the difference between the LATE and ATE is not obvious without empirical investigation.[35]

In RD designs, it is challenging to estimate the treatment effects for samples that are away from the treatment cutoff, although recent studies provide several potential approaches to address this important question (Jackson 2010; Angrist and Rokkanen 2012). In my research design, ATE can be estimated by making an additional

---

[34] The persistent effects are consistent with the findings in Ito, Ida, and Tanaka (2014). In their randomized field experiment, consumers who received economic incentives continued to conserve energy after the treatment period ended. Ferraro, Miranda, and Price (2011) and Allcott and Rogers (2014) also find similar persistent effects on water and energy conservation from information about peers' consumption.

[35] The difference between the LATE and ATE is not obvious because many factors can affect them in different directions. For example, one could expect that the ATE may be larger than the LATE (in absolute value) if customers who had lived longer in their homes were less likely to move and therefore had a larger incentive to invest in energy efficiency in response to the program. However, such customers may have already made this investment prior to the program (because they had resided there longer). Then, the ATE may be smaller than the LATE because newer residents had a larger incentive to invest in response to the program.

identification assumption. I use a method that combines an RD design with three-way fixed effects. The idea behind this method is similar to the approach used in Jackson (2010).

Recall that in my RD estimation there is a slight upward trend of the outcome variable over the running variable. As mentioned above, this trend comes from the general tendency for residential electricity customers to gradually increase their consumption after opening electricity accounts. Consider that customer i's consumption can be modeled as $y_{it} = \theta_i + \lambda_t + g(t - d_i) + \epsilon_{it}$, where d_i is the account opening date. Consumption depends on customer fixed effects, time fixed effects, and the growth of consumption $g(t - d_i)$. The actual functional form of $g(t - d_i)$ is unknown. Consider that customer A opened his account on the date of d_i in a year and that customer B opened her account on the *exact same date* in the previous year. The identification assumption that I make in this section is that $g(t - d_i)$ is common to customers A and B.

To estimate the ATE, I make two datasets. The first dataset is the electricity consumption data for summer 2005. This is the same dataset used for the main RD estimation. Recall that the running variable ($x_i = d_i - c$) is the account open date (d_i) relative to the enrollment cutoff date ($c =$ June 5, 2004). I define a series of 10-day bins for the running variable. For $j = \ldots, -20, -10, 0, 10, 20, \ldots$, I construct bin$_j$, which includes customers whose running variables satisfy $j < x_i \leq j + 10$. For example, bin$_{j=0}$ includes customers who opened accounts between June 6 and June 15 in 2004.[36]

The second dataset is the electricity consumption data for the summer of 2004. For this dataset, I define the running variable by $x_i = d_i - c_{2003}$, where c_{2003} is June 5, 2003. Using this running variable, I define bin$_j$ in the same way as the first dataset. For example, bin$_{j=0}$ includes customers who opened accounts in 2003 between June 6 and June 15. Thus, the second dataset can be considered a placebo dataset as if there was a rebate program for the 2004 summer based on an enrollment cutoff date in 2003.

Let y_{isjt} be electricity consumption for customer i in dataset $s = \{1, 2\}$ in bin j for billing month t.[37] Define the treatment dummy variable by D_{isjt}, which equals one if $s = 1, j < 0$, and $t \in$ treatment period. Pooling the first and second datasets, I estimate

$$(4) \qquad\qquad y_{isjt} = \alpha \cdot D_{isjt} + \theta_i + \lambda_{st} + \mu_{jt} + \eta_{isjt}.$$

Given my identification assumptions, α provides the program's ATE. I include three-way fixed effects: customer-level fixed effects (θ_i), time fixed effects that are specific to each dataset (λ_{st}), and those that are specific to each bin (μ_{jt}). The idea is that μ_{jt} controls for potential confounding effects of the running variable in a nearly nonparametric way given that the growth pattern of consumption is shared among

---

[36] Using a narrower bin size does not significantly change the estimation results.

[37] I define t for the first and second datasets as follows. Suppose that I set $t = 0$ for the June 2005 billing month for the first dataset. Then, $t = 0$ for the June 2004 for the second dataset. By defining t in this way, I can use bin fixed effects μ_{jt} to control for the running variable in a nearly nonparametric way.

Table 7—Average Treatment Effects (ATE) for Inland Climate Zones

Estimates: Bandwidth:	LATE 90 days (1)	ATE 180 days (2)	ATE 1 year (3)	ATE 2 years (4)	ATE 3 years (5)	ATE 4 years (6)
Panel A. All months						
Treatment effect	−0.042	−0.041	−0.043	−0.042	−0.037	−0.034
	(0.013)	(0.010)	(0.009)	(0.008)	(0.008)	(0.008)
Panel B. Each month						
Treatment effect in May	−0.034	−0.036	−0.037	−0.042	−0.036	−0.037
	(0.015)	(0.011)	(0.009)	(0.009)	(0.009)	(0.008)
Treatment effect in June	−0.055	−0.046	−0.050	−0.048	−0.041	−0.038
	(0.017)	(0.013)	(0.011)	(0.010)	(0.010)	(0.010)
Treatment effect in July	−0.041	−0.037	−0.039	−0.037	−0.030	−0.027
	(0.019)	(0.013)	(0.012)	(0.011)	(0.011)	(0.010)
Treatment effect in August	−0.037	−0.040	−0.041	−0.035	−0.030	−0.024
	(0.018)	(0.013)	(0.011)	(0.010)	(0.010)	(0.010)
Treatment effect in September	−0.056	−0.048	−0.053	−0.052	−0.050	−0.047
	(0.016)	(0.012)	(0.010)	(0.009)	(0.009)	(0.009)
Observations	208,537	420,149	640,415	978,707	1,257,978	1,508,618

Notes: The dependent variable is the log of electricity consumption. Given the identification assumptions described in the main text, this estimation produces the average treatment effect (ATE) for the samples included in each estimation. Standard errors are clustered at the customer level to adjust for serial correlation.

customers who opened accounts on a specific day and those who did so on the same day during the previous year.

I begin by estimating the ATE for customers with $-180 \leq x_i \leq 90$. This estimation includes customers who opened accounts in the period that was 180 days before and 90 days after the cutoff date.[38] Similarly, I estimate the ATE for customers with $x_i \geq -365$ (one year before the cutoff date), $x_i \geq -730$ (two years), $x_i \geq -1,095$ (three years), and $x_i \geq -1,460$ (four years) to examine how the treatment effects differ for customers who opened accounts earlier than those households used in my RD sample.

Table 7 presents the LATE and the ATE for inland climate zones. Panel A reports the overall treatment effects for all summer billing months and panel B reports the separate estimates for each month. As a reference, column 1 shows the LATE from the previous section. Column 2 shows the ATE for customers who opened accounts in the period that was 180 days before and 90 days after the cutoff date. The ATE and LATE are not statistically different, although their point estimates are slightly different. The standard errors are slightly tighter for the ATE because the ATE is estimated from a broader range of samples, while the LATE is estimated from the samples close to the eligibility cutoff date.

---

[38] To understand the intuition, it is useful to consider the four groups that are included in the estimation: (1) the treatment group (whose account open dates were between January 5 and June 5, 2004), (2) the control group (between June 6 and September 5, 2004), (3) the placebo treatment group (between January 5 and June 5, 2003), and (4) the placebo control group (between June 6 and September 5, 2003). Bin fixed effects (μ_{jt}) are used to control for the running variable given that the growth pattern of consumption is common between customers who opened accounts on a day and those who did so on the same day in the previous year.

Table 8—Program Cost per Estimated Reductions in Consumption and Carbon Dioxide

	Coastal	Inland	Total
Number of customers	3,190,027	299,178	3,489,205
Consumption in summer 2005 (kWh)	8,247,457,920	1,154,292,248	9,401,750,168
Direct program cost for rebate ($)	9,358,919	1,250,621	10,609,540
Estimated reduction (kWh)	9,908,840	50,605,714	60,514,555
Estimated reduction in carbon dioxide (ton)	4,459	22,773	27,232
Program cost per kWh ($/kWh)	0.945	0.025	0.175
Program cost per carbon dioxide ($/ton)	2,099	55	390
Program cost per carbon dioxide ($/ton) (Adjusted for noncarbon external benefits)	2,090	46	381

Notes: This table reports the cost-benefit analysis of the 20/20 program for SCE's coastal areas, inland areas, and all service areas. Row 1 shows the number of residential customers who maintained their accounts in the summer of 2004 and 2005. Row 2 presents the aggregate consumption in the summer months. Row 3 reports the aggregate amount of rebate paid to customers. Row 4 shows the estimated kWh reduction from the treatment effect of the program. Row 5 translates this reduction into the reduction in carbon emissions by using the average carbon intensity of electricity consumed in California, which is 0.9 lb per kWh according to California Air Resources Board (2011).

In columns 2 through 6, I include the customers who opened accounts earlier than 180 days before the cutoff date. For example, column 6 includes customers who opened accounts four years before the rebate program began. I find that the point estimates of the ATE are not statistically different from the LATE at conventional significance levels, although their point estimates are slightly different. The results imply that the program's overall treatment effect is about a 4 percent reduction in consumption both for the RD sample and for customers who opened accounts earlier than those included in the RD sample.[39]

IV. Policy Implications

A. *The Rebate Program's Cost-Effectiveness*

In the literature, evaluations of energy conservation programs usually report two measures of cost-effectiveness: (1) the program's cost per unit of energy saved and (2) the program's cost per ton of emissions abated (Joskow and Marron 1992, and Boomhower and Davis 2014). Although these are not direct measures of welfare, these provide a valuable starting point and are widely used in policy discussions. I provide these values below and discuss their welfare implications in the next section.

Table 8 shows the program's cost-effectiveness based on the RD estimates of the effect of rebate incentives on energy conservation.[40] The third row shows the direct cost for the rebate payment, which is the total rebate amount paid to customers. Note that this direct cost does not include indirect costs such as administrative and marketing costs. This direct cost also does not include costs undertaken by households who reduced their energy consumption. Therefore, the cost-effectiveness measure in this analysis should be considered a lower bound of the program's cost. The fourth row shows the estimated reductions in consumption based on the RD estimates.

---

[39] Similarly, I find that the ATE is not significantly different from the LATE for coastal customers.
[40] I use the LATE from my RD estimation. Using the ATE from the previous section does not significantly change the results because my ATE and LATE are not statistically different.

The fifth row translates the estimates into reductions in carbon emissions by using the average carbon intensity of electricity consumed in California, which is 0.9 lb. per kWh (California Air Resources Board 2011). The sixth row shows the program cost per kWh of electricity saved. The seventh row provides the program cost per ton of carbon dioxide abated, which may be overestimated because noncarbon externalities are also abated (e.g., particulate reductions). In the final row, based on Greenstone and Looney (2012) I calculate the noncarbon external benefits, subtract these from the program cost, and divide the adjusted program cost by the abated carbon dioxide.[41]

The results provide important implications. First, the program's cost-effectiveness differs substantially between coastal and inland areas. In coastal areas, the program is a very expensive way to reduce electricity consumption. This is because the program did not induce significant reductions in usage but paid substantial rebates to consumers who reduced their consumption for reasons unrelated to the program's incentive. The program cost, 94.5 cents per kWh, is large relative to any reasonable range for the marginal cost of electricity. In contrast, the program cost per kWh reduction is much smaller in inland areas, where it costs 2.5 cents to obtain a kWh reduction in consumption.

Second, the overall program was unlikely to be cost-effective within a reasonable range of assumptions regarding the private and social costs of electricity. The overall program cost was 17.5 cents per kWh reduction. The average cost of electricity supplied by SCE was 13.37 cents per kWh in 2005. To justify the program's cost-effectiveness by the externality from carbon emissions, the social cost of carbon has to be larger than \$92 per ton of carbon dioxide. If I subtract the noncarbon external benefits from the program cost, the social cost of carbon emissions has to be larger than \$82 per ton of carbon dioxide, which is larger than most estimates in the literature (Greenstone, Kopits, and Wolverton 2011). In addition, this calculation does not include the program's indirect costs such as administrative and marketing expenses. According to the study by Wirtshafter Associates (2006), SCE spent about \$4 million to administer and advertise the program. The overall program cost is 24.1 cents per kWh including the indirect costs.[42]

Note that I use the estimates for short-run treatment effects to calculate the cost-effectiveness shown in Table 8. For inland areas, I also find significant long-run

[41] Table 1 in Greenstone and Looney (2012) provides estimates for the noncarbon external cost from electricity generation from existing coal plants (34 cents/kWh) and natural gas plants (2 cents/kWh). In 2005, 33.6 percent of electricity consumed in California came from natural gas and 9.8 percent came from coal (California Energy Comission 2014). I assume that other types of power plants have zero external cost, although this assumption might be invalid when taking into account other external costs, such as the cost from potential nuclear accidents. With these numbers, the noncarbon external benefits from the estimated reductions in consumption are \$39,739 (coastal), \$202,950 (inland), and \$242,689 (total). I subtract them from the direct program cost to calculate the adjusted program cost. This adjustment does not change my results substantially because only a small amount of the electricity consumed in California comes from coal.

[42] An important caveat is that the California 20/20 rebate program provided a rebate for consumers based on their monthly electricity consumption. The marginal cost of electricity is generally higher in peak hours and lower in off-peak hours. If the reductions mostly come from peak hour usage, the benefit comes not only from reductions in emissions but also from savings of the relatively high marginal cost of electricity. In this case, the cost-effectiveness would be better than in my calculation. On the other hand, if the reduction mostly comes from off-peak usage, the cost-effectiveness would be worse than in my calculation. The number of reductions that come from peak and off-peak hours cannot be quantified from the monthly billing data.

treatment effects for the summers of 2006, 2007, and 2008. How does the cost-effectiveness change if I account for the long-run effects? Given a set of assumptions, I calculate the long-run cost-effectiveness for the inland areas.[43] The estimates of the long-run treatment effects imply that the program produced reductions in electricity consumption by 156,305 GWh and carbon dioxide by 70,337 tons in the 4 years from 2005 to 2008. With a discount rate of 4 percent, the long-run program cost for the inland areas is 0.9 cents per kWh and $20 per ton of carbon dioxide. This calculation implies that it is important to consider the program's potential long-run effects when conducting the cost-benefit analysis. However, even with accounting for the long-run considerations, the program's overall cost is still expensive ($160 per ton of carbon dioxide) because the long-run treatment effects were essentially zero in coastal areas.

B. *Welfare Implications*

The high program cost found in the previous section does not necessarily mean that the program is not welfare improving because the rebate expense can be considered a transfer between customers. The utility companies passed the cost to customers by increasing the electricity price afterward.[44] The rebate expense, therefore, can be considered a transfer from all customers to rebated customers through the electricity price.

Consider two simple cases. Suppose that customers pay a linear electricity price that is lower than the social marginal cost of electricity. In this case, the rebate program can improve social welfare even if there is no treatment effect. The rebate expense slightly increases the electricity price afterward and can improve welfare if the new price is closer to the social marginal cost of electricity. Conversely, if customers pay a linear electricity price that is higher than the social marginal cost of electricity, the rebate program is likely to lower welfare because there is a greater price distortion after the electricity price increases.

Increasing block pricing complicates the welfare analysis of the California 20/20 program (Borenstein 2012, and Ito 2014). Here, the marginal price of electricity is an increasing step function of a household's monthly consumption. That is, customers pay a higher marginal price when they consume more electricity during their billing month. In 2005 and 2006, the marginal prices for the first through fourth tiers were 12, 14, 17, and 20 cents per kWh. A key question then becomes what is the correct social marginal cost of electricity? To estimate this, I make three assumptions. First, suppose that the long-run private marginal cost of electricity is equal to the average cost of electricity under the existing tariff schedule. Then, the private marginal cost is 13.37 cents per kWh for SCE in 2005. Second, suppose that the externality from

[43] Every year, about 15 percent of customers terminate their electricity accounts when they move. The reductions in consumption in future years come from customers who maintained the same electricity account. I use a discount rate to compare the rebate expense in 2005 and benefits in future years. I use the Federal Reserve's discount rate, which was 4 percent in the summer of 2005.

[44] Most utility conservation programs in the United States recover the cost by increasing the electricity price. This is a notable difference from the energy-efficient appliance replacement program in Mexico (Boomhower and Davis 2014), where the cost is paid by tax revenue.

carbon emissions is 0.95 cents and that the noncarbon externality is 0.4 cents per kWh.[45] Then, the social marginal cost of electricity is 14.72 cents per kWh.

In the billing data, about half of the customers are in the first and second tiers and half are in the third and fourth tiers. Half the customers, therefore, pay marginal prices that are slightly lower than the social marginal cost, while the other half pay marginal prices that are significantly higher than the social marginal cost of electricity. In theory, the rebate program can improve welfare if the cost is recovered from the two lower tiers. However, in practice, California's electric utilities have taken the opposite approach because of a regulatory constraint. After the 2000–2001 California electricity crisis, regulators and state legislators were concerned about the impact of price increases on lower-income customers, and the first two tiers were virtually frozen. In fact, SCE increased only the third and fourth tier rates in 2006 (from 17 to 23 cents and from 20 to 32 cents), while it did not change the first and second tier rates. It is therefore difficult to argue that the program improved social welfare unless the externality from electricity is substantially larger than the estimates in the literature.

V. Conclusion

Subsidy policies that intend to correct negative externalities often create asymmetric incentives because increases in externalities remain unpriced. I study the implications of such asymmetric incentives by using a RD design for the California 20/20 electricity rebate program. Using customer-level administrative data, I find precisely estimated zero causal effects in coastal areas and a significant 4 percent consumption reduction in inland areas. In addition, I find evidence that income, climate conditions, and air conditioner saturation significantly drive the heterogeneity and that asymmetric subsidy structures weaken incentives because consumers whose usage is far from the rebate target respond very little to the program.

The heterogeneous treatment effects result in the program's cost-effectiveness being very different between coastal areas (94.5 cents per kWh reduction) and inland areas (2.5 cents per kWh reduction). However, because substantial rebates were paid to customers in the areas for which I find nearly zero treatment effects, the overall program cost was 17.5 cents per kWh reduction and $381 per ton of carbon dioxide reduction. Therefore, the cost of the program was unlikely to be effective in reducing externalities over a reasonable range of the social marginal cost of electricity.

This paper's findings imply that one way to improve the program's cost-effectiveness is to target lower-income customers and customers in areas with high summer temperatures. Another important way to improve the program's efficiency is to target consumption during peak hours, when the marginal cost of electricity is likely to be high. For the 2005 California 20/20 rebate program, regulators could not target peak

---

[45] I use Greenstone, Kopits, and Wolverton (2011) for the social marginal cost of a ton of carbon emissions ($21 per ton), California Air Resources Board (2011) for the average carbon intensity of electricity consumed (0.9 lb. per kWh), and Greenstone and Looney (2012) for the noncarbon external cost for generating electricity from existing coal plants (34 cents/kWh) and natural gas plants (2 cents/kWh). In 2005, 33.6 percent of electricity consumed in California came from natural gas and 9.8 percent came from coal (California Energy Commission 2014).

hours because residential customers in California did not have smart meters, which record their hourly electricity consumption. Recently, a growing number of customers in many countries have gained access to smart meters, which makes it possible for regulators to target consumption in particular hours.[46] This new technology can also be used to inform consumers about their real-time consumption, which is likely to enhance their response to economic incentives. However, even with smart meters, some of the fundamental problems characterized in this paper remain if regulators continue to provide asymmetric incentives that subsidize reductions in consumption but do not address increases in consumption.[47]

REFERENCES

Allcott, Hunt, and Michael Greenstone. 2012. "Is There an Energy Efficiency Gap?" *Journal of Economic Perspectives* 26 (1): 3–28.

Allcott, Hunt, and Todd Rogers. 2014. "The Short-Run and Long-Run Effects of Behavioral Interventions: Experimental Evidence from Energy Conservation." *American Economic Review* 104 (10): 3003–37.

Anderson, Michael L. 2014. "Subways, Strikes, and Slowdowns: The Impacts of Public Transit on Traffic Congestion." *American Economic Review* 104 (9): 2763–96.

Angrist, Joshua, and Miikka Rokkanen. 2012. "Wanna Get Away? RD Identification Away from the Cutoff." National Bureau of Economic Research (NBER) Working Paper 18662.

Aroonruengsawat, Anin, and Maximilian Auffhammer. 2011. "Impacts of Climate Change on Residential Electricity Consumption: Evidence from Billing Data." In *The Economics of Climate Change: Adaptations Past and Present*, edited by Gary D. Libecap and Richard H. Steckel, 311–42. Chicago: University of Chicago Press.

Boomhower, Judson, and Lucas W. Davis. 2014. "A credible approach for measuring inframarginal participation in energy efficiency programs." *Journal of Public Economics* 113 (C): 67–79.

Borenstein, Severin. 2002. "The Trouble With Electricity Markets: Understanding California's Restructuring Disaster." *Journal of Economic Perspectives* 16 (1): 191–211.

Borenstein, Severin. 2009. "To What Electricity Price Do Consumers Respond? Residential Demand Elasticity Under Increasing-Block Pricing." http://faculty.haas.berkeley.edu/borenste/download/NBER_SI_2009.pdf.

Borenstein, Severin. 2012. "The Redistributional Impact of Nonlinear Electricity Pricing." *American Economic Journal: Economic Policy* 4 (3): 56–90.

Borenstein, Severin. 2013. "Effective and Equitable Adoption of Opt-In Residential Dynamic Electricity Pricing." *Review of Industrial Organization* 42 (2): 127–60.

Borenstein, Severin, James B. Bushnell, and Frank A. Wolak. 2002. "Measuring Market Inefficiencies in California's Restructured Wholesale Electricity Market." *American Economic Review* 92 (5): 1376–1405.

California Air Resources Board. 2011. *California Greenhouse Gas Emissions Inventory: 2000–2009*. California Environmental Protection Agency. Sacramento, December.

California Energy Commission. 2014. *California Electrical Energy Generation: Total Production, by Resource Type*. Energy Almanac. Sacramento, April.

California Public Utility Commission (**CPUC**). 2001. *Energy division resolution E-3733*. California Public Utility Commission. May. http://docs.cpuc.ca.gov/publishedDocs/PUBLISHED/FINAL_RESOLUTION/6842.htm (accessed June 9, 2015).

Chong, Howard. 2012. "Building Vintage and Electricity Use: Old Homes Use Less Electricity in Hot Weather." *European Economic Review* 56 (5): 906–30.

---

[46] See Wolak (2010); Ito, Ida, and Tanaka (2014); and Jessoe and Rapson (2014).

[47] Despite this problem, in practice, peak-time rebate (PTR) programs are often politically favored over peak-time pricing in ongoing policy discussions of dynamic electricity pricing. There is strong political opposition to peak-time pricing. In contrast, there is relatively less political opposition to PTR programs, which is equivalent in principle to the California 20/20 rebate program in the sense that the rebate subsidizes reductions in consumption but does not penalize increases in consumption (Borenstein 2013).

Davis, Lucas W., Alan Fuchs, and Paul J. Gertler. 2014. "Cash for Coolers: Evaluating a Large-Scale Appliance Replacement Program in Mexico." *American Economic Journal: Economic Policy* 6 (4): 207–38.

Duflo, Esther, Michael Greenstone, Rohini Pande, and Nicholas Ryan. 2013. "Truth-telling by Third-party Auditors and the Response of Polluting Firms: Experimental Evidence from India." *Quarterly Journal of Economics* 128 (4): 1499–1545.

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

Faruqui, Ahmad, and Sanem Sergici. 2010. "Household response to dynamic pricing of electricity: A survey of 15 experiments." *Journal of Regulatory Economics* 38 (2): 193–225.

Ferraro, Paul J., Juan Jose Miranda, and Michael K. Price. 2011. "The Persistence of Treatment Effects with Norm-Based Policy Instruments: Evidence from a Randomized Environmental Policy Experiment." *American Economic Review* 101 (3): 318–22.

Frandsen, Brigham R., Markus Frölich, and Blaise Melly. 2012. "Quantile treatment effects in the regression discontinuity design." *Journal of Econometrics* 168 (2): 382–95.

Gallagher, Kelly Sims, and Erich Muehlegger. 2011. "Giving green to get green? Incentives and consumer adoption of hybrid vehicle technology." *Journal of Environmental Economics and Management* 61 (1): 1–15.

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 (1–2): 113–42.

Greenstone, Michael, Elizabeth Kopits, and Ann Wolverton. 2011. "Estimating the Social Cost of Carbon for Use in U.S. Federal Rulemakings: A Summary and Interpretation." National Bureau of Economic Research (NBER) Working Paper 16913.

Greenstone, Michael, and Adam Looney. 2012. "Paying Too Much for Energy? The True Costs of Our Energy Choices." *Daedalus* 141 (2): 10–30.

Imbens, Guido W., and Thomas Lemieux. 2008. "Regression discontinuity designs: A guide to practice." *Journal of Econometrics* 142 (2): 615–35.

Ito, Koichiro. 2014. "Do Consumers Respond to Marginal or Average Price? Evidence from Nonlinear Electricity Pricing." *American Economic Review* 104 (2): 537–63.

Ito, Koichiro. 2015. "Asymmetric Incentives in Subsidies: Evidence from a Large-Scale Electricity Rebate Program: Dataset." *American Economic Journal: Economic Policy*. http://dx.doi.org/10.1257/pol.20130397.

Ito, Koichiro, Takanori Ida, and Makoto Tanaka. 2014. "The Persistence of Intrinsic and Extrinsic Motivation: Experimental Evidence from Energy Demand." http://www.abef.jp/event/2014/pdf_abst/PR0005.pdf.

Jackson, C. Kirabo. 2010. "Do Students Benefit from Attending Better Schools? Evidence from Rule-based Student Assignments in Trinidad and Tobago." *Economic Journal* 120 (549): 1399–1429.

Jessoe, Katrina, and David Rapson. 2014. "Knowledge is (Less) Power: Experimental Evidence from Residential Energy Use." *American Economic Review* 104 (4): 1417–38.

Joskow, Paul L. 2001. "California's electricity crisis." *Oxford Review of Economic Policy* 17 (3): 365–88.

Joskow, Paul L., and Donald B. Marron. 1992. "What Does a Negawatt Really Cost? Evidence from Utility Conservation Programs." *Energy Journal* 13 (4): 41–74.

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

Mian, Atif, and Amir Sufi. 2012. "The Effects of Fiscal Stimulus: Evidence from the 2009 'Cash for Clunkers' Program." *Quarterly Journal of Economics* 127 (3): 1107–42.

Pigou, Arthur Cecil. 1924. *The Economics of Welfare.* London: Macmillan.

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 (NBER) Working Paper 9986.

Reiss, Peter C., and Matthew W. White. 2005. "Household Electricity Demand, Revisited." *Review of Economic Studies* 72 (3): 853–83.

Reiss, Peter C., and Matthew W. White. 2008. "What changes energy consumption? Prices and public pressures." *RAND Journal of Economics* 39 (3): 636–63.

Saez, Emmanuel. 1999. "Do Taxpayers Bunch at Kink Points?" National Bureau of Economic Research (NBER) Working Paper 7366.

Sallee, James M. 2011. "The Surprising Incidence of Tax Credits for the Toyota Prius." *American Economic Journal: Economic Policy* 3 (2): 189–219.

Sallee, James M., and Joel Slemrod. 2012. "Car notches: Strategic automaker responses to fuel economy policy." *Journal of Public Economics* 96 (11): 981–99.

Schneider, Lambert. 2007. "Is the CDM fulfilling its environmental and sustainable development objectives? An evaluation of the CDM and options for improvement." Berlin, Öko-Institut. www.oeko.de/oekodoc/622/2007-162-en.pdf

Sutter, Christoph, and Juan Carlos Parreño. 2007. "Does the current Clean Development Mechanism (CDM) deliver its sustainable development claim? An analysis of officially registered CDM projects." *Climatic Change* 84 (1): 75–90.

U.S. Energy Information Administration. 2013. *Electric Power Annual 2011*. U.S. Department of Energy. Washington, DC, January.

Volpp, Kevin G., Andrea B. Troxel, Mark V. Pauly, Henry A. Glick, Andrea Puig, David A. Asch, Robert Galvin, et al. 2009. "A Randomized, Controlled Trial of Financial Incentives for Smoking Cessation." *New England Journal of Medicine* 360 (7): 699–709.

Wirtshafter Associates. 2006. *Evaluation of the California Statewide 20/20 Demand Reduction Programs*. Energy Market Innovations. Rydal, PA, June.

Wolak, Frank A. 2006. "Residential Customer Response to Real-Time Pricing: The Anaheim Critical-Peak Pricing Experiment." Unpublished.

Wolak, Frank A. 2010. "An Experimental Comparison of Critical Peak and Hourly Pricing: The PowerCentsDC Program." Unpublished.

Wolak, Frank A. 2011. "Do Residential Customers Respond to Hourly Prices? Evidence from a Dynamic Pricing Experiment." *American Economic Review* 101 (3): 83–87.