Commercial and Industrial Demand Response Under Mandatory Time-of-Use Electricity Pricing [1]

Katrina Jessoe[2] and David Rapson[3]

October 10, 2011

Abstract

This paper evaluates the impact of a large-scale field deployment of mandatory time-of-use (TOU) pricing on the energy use of commercial and industrial firms. We find (1) little evidence of change in usage or load (2) no change in bills after adjusting for a rate-class discount, and (3) small shifts in bill levels and a slight increase in bill volatility. Economic efficiency was not increased by this policy, as the TOU regime did not generate a reduction in peak load, the intended response. Our results also suggest that concerns about increased electricity expenditure and bill volatility have been overstated.

JEL: D22, L50, L94, Q41

Keywords: Regulation; Mean Reversion; Electricity; Time-of-Use Pricing

[1]We thank, without implicating, Lucas Davis, Kenneth Gillingham, Koichiro Ito, Chris Knittel, Aaron Smith and Jeffrey Williams. Thanks to United Illuminating personnel Megan Pomeroy, Donna Burrows, and Mark Colca for assistance in obtaining the data for this project. We also thank the University of California Center for Environmental and Energy Economics (UCE3) for financial support. Comments from seminar participants at the UC Energy Institute and UCE3 helped to improve this draft. Tom Blake and Brock Smith provided excellent research assistance. All errors are our own.

[2]Department of Agricultural Economics, University of California, Davis, One Shields Ave, Davis, CA 95616; Phone: (530) 752-6977; Email: `kkjessoe@ucdavis.edu`

[3]Department of Economics, University of California, Davis, One Shields Ave, Davis, CA 95616; Phone: (530) 752-5368; Email: `dsrapson@ucdavis.edu`

1 Introduction

In the electricity market, marginal costs vary by the minute but retail prices are (for the most) part time-invariant. Because of this, the market does not function properly and substantial economic inefficiency may result, both in the short-run and the long-run.[1] The disparity between wholesale and retail prices leads to chronic over- and under-consumption at different times of the day, excess capital investment to prevent blackouts (in the absence of rationing) and an increase in the opportunity for producers to exploit market power. Estimates place the magnitude of the deadweight loss from time-invariant retail electricity prices in the tens of billions of dollars annually in the U.S.[2] Real-time pricing (RTP), in which the retail price of electricity varies to reflect the wholesale price, would fully eliminate these inefficiencies by transmitting changes in the marginal cost to retail consumers.

The need for time-varying prices is not unique to electricity. In many markets, flat retail prices cause periods of congestion in which peak demand may approach or exceed the capacity of fixed infrastructure. This congestion induces an inefficient outcome, since each user fails to internalize the costs imposed on other market participants. Countless economic settings exhibit congestion: electricity markets; automobile congestion on bridges, highways and in urban centers; airport gates and runways; Internet users and bandwidth. The phenomenon was formalized in the economics literature decades ago and in many settings.[3] The recommendation among economists follows in the spirit of Pigou (1912). Charge prices such that the user internalizes the externality. Such a price, in this setting, will vary at high frequency (i.e. RTP).

[1]During the California energy crisis in 2001, wholesale electricity prices exceeded $1,400 per mWh, or $1.40 per kWh, more than twenty times the then retail price of $0.067 per kWh.

[2]Borenstein & Holland (2005) estimate that 5-10 percent of the market is deadweight loss. In 2009, $350 billion worth of electricity was consumed in the United States, implying an annual inefficiency on the order of $17 - $35 billion, roughly 2-3 times the entire budget of the Environmental Protection Agency (just over $10 billion in 2010).

[3]By Houthakker (1951), Steiner (1957), and Williamson (1966) for electricity pricing, Walters (1961) for traffic congestion, and Carlin & Park (1970) and Levine (1969) for airport congestion.

Previous work explores the potential of RTP to reduce allocative inefficiencies and capital over-investment in electricity markets. Of particular relevance is a study by Borenstein & Holland (2005), which may be seen as a close theoretical counterpart to this paper. They build a theoretical model to analyze the effect of placing retail customers on RTP. Simulations show that time-invariant pricing presents a barrier to both allocative efficiency and efficient capital investment, though increasing the share of customers on RTP will not necessarily lead to an increase in market efficiency. Simulation results demonstrate large potential gains from switching customers onto RTP; placing one-third of customers on RTP will lead to a 3 to 11 percent increase in surplus and will reduce peak capacity requirements by 44 percent.

Despite the promise of RTP to increase economic efficiency, technological and political obstacles make it costly to implement. In practice the most common form of corrective pricing in electricity markets is a weaker version, called "time-of-use" (TOU) pricing. Under this approach, two or three periods each day are designated as high, medium or low demand (according to historic patterns) and higher prices are assigned to higher-demand hours in a coarse attempt to address the inefficiency caused by the discrepancy between retail and wholesale prices. Its elegance is in the incentives facing a prospective user: consume when your marginal benefit is highest, but pay a high price if that time corresponds to a system peak; or shift usage to off-peak hours and pay a lower price. Its fundamental drawback is coarseness; movements in the (social) shadow cost are not perfectly transmitted to the user, and thus not fully internalized by consumers. Our estimates show that TOU pricing can capture no more than 6 percent of wholesale market price variation in Connecticut leading us to question the ability of TOU pricing to reduce demand during periods of high wholesale prices.[4] Without some demand response the deadweight loss will remain (unless demand is perfectly inelastic). In this study, we analyze empirically if TOU pricing produces the intended response of a shift in when electricity is consumed, or if users continue to act as they did on a flat rate.

To date, the empirical evidence from electricity markets suggests that TOU pricing is at best moderately

[4] Borenstein (2005) performs similar calculations for California and New Jersey, and estimates that 6 to 13 percent of wholesale market variation is captured by TOU pricing.

effective among commercial and industrial (C&I) users (Aigner & Hirschberg 1985, Aigner et al. 1994).[5] Using quasi-experimental data, Aigner & Hirschberg (1985) find strong complementarity between peak and off-peak usage, and evidence of small but statistically significant substitution from peak to off peak hours in response to TOU pricing. Their results also suggest that the magnitude of the differential between peak and off-peak prices influences response. Their setting is less than ideal due to the fact that participants, though randomly assigned initially to control and treatment, are allowed to opt out if adversely affected by the treatment. In a randomized controlled trial in Israel, Aigner et al. (1994) also detect small but significant shifts in usage by firms. In both experiments, the price change is explicitly temporary in nature.

In this paper we evaluate a large-scale field deployment of mandatory TOU pricing for C&I customers in Connecticut. The setting is unique since most TOU programs in the field are introduced voluntarily, and experiments are generally temporary in nature. In contrast, we study the impact of an involuntarily-assigned and permanent rate change. In our setting, firms are switched from a flat-rate price per kWh onto a TOU schedule with peak prices that are approximately 60 percent higher than off-peak. Under the mandatory TOU assignment rule, C&I firms whose peak load exceeds a pre-determined threshold are placed on a TOU schedule and can not, regardless of future behavior, return to a flat rate tariff. Using monthly panel data from the universe of customers in United Illuminating's "General Services" rate class, we exploit variation both over time and across space to estimate the impact of this policy on usage, peak load and expenditure.[6] Despite the presence of plausibly exogenous variation in treatment assignment, estimation of the treatment effect requires navigating several potential confounding factors. Failure to account for firm-specific trends, firm life cycle considerations, and within-firm autocorrelation in unobservables leads to qualitatively different (and we argue, incorrect) conclusions about the policy's effectiveness.

The mechanism by which firms are assigned to "treatment" makes mean reversion an additional empirical

[5] A long empirical literature has also studied residential responsiveness to time-variant electricity pricing, finding evidence that some policies induce a shift in peak to off-peak usage (e.g. Wolak 2007) while others induce conservation (e.g. Allcott 2010).

[6] The General Services rate class includes flat-rate demand and non-demand schedules, as well as TOU pricing.

challenge. Under the treatment assignment rule, any flat-rate firm with peak load exceeding 100kW after June 2010 is switched onto TOU pricing, with no possibility of reversal. If the increase in peak load that causes the firm to exceed the assignment threshold is due to a transitory shock, then peak load (and positively correlated variables like usage and billed amount) will decrease in the next period, regardless of the new pricing structure. Failure to account for transitory shocks will cause the econometrician to incorrectly attribute a decline in load (usage or expenditure) to treatment, when in fact it is due to mean reversion. The problem of mean reversion arises in many contexts, and plays a central role in recent studies evaluating school funding programs (Chay et al. 2003), estimating the elasticity of taxable income with respect to marginal tax rates (Saez et al. 2009) and measuring customer response to block rate electricity pricing (Ito 2011). These studies all suggest corrections to control for mean reversion and highlight that the institutional details characterizing an empirical setting play a crucial role in determining the appropriate solution. Guided by the unique empirical features of our setting, we implement three controls for mean-reversion.

Our analysis yields three primary results. First, there is little evidence of a change in usage or load from TOU pricing in our setting, and we cannot reject the hypothesis that there is no aggregate change in usage or peak load. In our preferred specification, we can rule out aggregate decreases of greater than 1 percent in monthly consumption or peak load in response to treatment. Second, we find that on average firm electricity bills decrease. This is due to a rate-class discount implicit in the price schedule, rather than a behavioral response. After adjusting for the rate class discount, TOU pricing does not impact monthly electricity expenditure. Finally, increases in bill levels and bill volatility are minimal, with only a small number of firms being adversely affected.

These results contribute a new data point to the sparse existing literature on C&I TOU pricing. We provide an internally consistent estimate of the impact of TOU pricing, finding that this policy induces a statistically- and economically-insignificant response. The detected non-response among C&I customers may be attributable to one (or some combination) of several reasons: (1) the peak to off-peak differential is not high enough to induce a response[7]; (2) TOU prices are too coarse to be effective in some settings[8]; (3) in the short-run, firms cannot adjust their electricity load profile[9]; or, (4) an assortment of other explanations along the lines of rational inattention (if electricity was a small line-item expenditure) or principal agent contracting imperfections (when the bill-payer does not make electricity usage decisions).

We also add context to the ongoing debate about whether (conditional on using TOU pricing) TOU rates should be mandated or voluntarily implemented. Driven by concerns about excess harm from high bill levels or dramatic increases in bill volatility, most TOU tariffs in the U.S. have been introduced voluntarily, with customers having the option to select into a TOU rate.[10] Our results suggest these concerns have been overstated. Of treated firms, 95 percent experience bill increases of less than 8.5 percent[11] and bill volatility increases on average by less than 5 percent.

Finally, we believe that our setting is precisely that which regulators and policy-makers will find relevant. At first glance, one may be concerned about external validity given that, in the 30 years preceding the mandatory TOU policy, firms in UI's territory had the option to select onto the TOU rate. By the time that mandatory TOU pricing took effect, over 1,400 firms had opted into the time-varying tariff. Despite the clear preemptive self-selection of some firms into this pricing program, our empirical setting remains appropriate to measure the marginal impact of a mandatory assignment rule. If other utilities introduce mandatory TOU pricing, it is highly unlikely that firms would not be given the option to opt into TOU pricing. That the timing of voluntary adoption occurs before the mandatory policy is activated is only problematic if one

[7]Aigner & Hirschberg (1985).

[8]Borenstein (2005).

[9]As the program in CT matures, we will likely have the opportunity to examine a longer time-series of post-treatment outcomes.

[10]In February 2010 the President of the California Small Business Association warned that "...with dynamic pricing, small businesses will send workers home, tell workers not to come into work or pay large electric bills for using power on peak days." In response, PG&E has requested to delay its TOU deployment.

[11]This is an aggressive measure, having been calculated using bill levels already adjusted for the TOU rate-class discount implicitly offered by the CT policy.

believes that the composition of the voluntary group would be time-dependent. It is difficult to conceive of a rational model under which this would be true.

If mandatory programs are introduced in isolation (without giving firms the option to voluntarily enroll in TOU pricing) then we provide an upper bound estimate of financial harm since firms mandated onto the tariff are most likely to be adversely affected by the policy. Even for these firms, we find little evidence of bill increases and volatility in response to TOU pricing. In our setting, it appears that TOU pricing does not impose significant harm on C&I customers.[12]

2 Regulatory Setting

In 2006, the Connecticut Department of Public Utilities and Control (DPUC) issued an order requiring United Illuminating (UI), an electric utility serving over 324,000 residential and C&I customers in Connecticut, to phase in mandatory TOU pricing for commercial users. This policy was approved and implemented in an effort to reduce growing demand for electricity during peak periods.

In coordination with the DPUC, UI established peak load thresholds that, if exceeded, would cause a firm to be placed on mandatory time-of-use pricing (TOU). Once transferred onto the TOU tariff, a firm could not return to the flat rate schedule, regardless of future consumption.[13] The first demand threshold took effect on June 1, 2008 at 300kW. In each subsequent year the thresholds qualifying a firm for TOU pricing lowered; on June 1, 2009 the threshold was set at 200kW and on June 1, 2010 the threshold was set at 100kW. The

[12] In this paper we choose to focus on firms that are mandated onto the TOU rate. The question of what drives self-selection onto the tariff is also interesting and potentially important. We plan to investigate this question in a separate study.

[13] Firms are free to choose an alternate supplier for generation and these suppliers currently offer time-invariant rates. However, the majority of the TOU price differential is transmitted through distribution charges, which are billed through UI for all customers, regardless of the generation supplier.

majority of small commercial users did not approach even the lowest threshold for mandatory TOU pricing. Customers that crossed the threshold were charged a peak rate for electricity consumed between 10am and 6pm Monday-Friday, and an off-peak rate during all other hours.

Historically, small and medium-sized commercial users in UI's territory paid a single rate per kilowatt-hour (kWh) regardless of when electricity was consumed during the day, unless they opted into TOU pricing.[14] Since 1978, any general services customer could select into TOU pricing in lieu of the flat rate. Similar to mandatory TOU pricing, once customers volunteered for this tariff they could not return to the flat rate schedule. Also, since the early 1980s, customers reaching peak load in excess of 500kW were mandated onto TOU pricing.

Prices in UI territory are high by national standards. Table 1 reports the rate schedules in 2010 for commercial customers on a flat rate and TOU rate. For flat-rate customers the price per kWh of electricity is \$0.1659 in the winter and \$0.1864 in the summer. By comparison, TOU customers pay a higher price for electricity during peak hours, \$0.2281 in the winter and \$0.2284 in the summer, and a lower price during off peak hours, \$0.14391 in the both the summer and the winter. This amounts to a roughly 60 percent increase in peak prices relative to off-peak, and is in the low range of price differentials when compared to the C&I TOU studies referenced earlier.[15] For customers that either consume electricity primarily during off-peak hours or can readily reallocate consumption from peak to off-peak hours, TOU pricing has the potential to lower monthly electricity bills. In practice, the TOU prices are set in such a way that, on average, firms see their bills decrease upon switching (a point to which we will return later).

[14] Congestion charges vary by season but remain constant within a day.

[15] The peak to off-peak price ratios from Aigner & Hirschberg (1984) and Aigner et a. (1995) ranged from 1.2 to 2.5 and 1.9 to 8.3, respectively.

3 Data and Summary Statistics

The data are comprised of two separate databases maintained by UI. The primary data used to estimate the empirical models consist of monthly electricity usage, demand and expenditure from 1,803 commercial users serviced between January 1, 2009 and May 2011. We refer to these data as the customer billing data. A second data set supplements the billing data by providing information on peak and off-peak usage for a random sample of 1,168 commercial users between January 2009 and December 2010. We refer to these as the "load research" data.

While we rely only on the billing data to estimate our empirical model, the load research data serve two purposes. First, they provide us with an additional variable, the ratio of peak to off-peak usage, that we use to evaluate the comparability of our control and treatment groups. In contrast, the billing data do not contain information on peak and off-peak usage for (i) firms in our control group or (ii) mandatory TOU firms in pre-treatment months. Second, we use these data to calculate the TOU rate class discount: holding usage fixed, the decrease in customer bills that occurs simply from switching from a flat to TOU rate.

3.1 Customer Billing Data

We have a balanced panel of data for 1,212 customers and partial panel data for the remaining customers, creating a dataset of 42,639 firm-months. Monthly data obtained from the utility include: peak kilowatts, kilowatt hours consumed, the electricity bill and rate class. Table 2 provides descriptive results; means are reported by firm type. Customers are grouped into one of three firm types: "mandatory switchers" are customers that were mandated onto TOU pricing during the period of study; "always-TOU firms" identify customers that opt into the TOU rate before August 2007 and are on TOU pricing for the duration of our sample; lastly, "non-TOU firms" are customers that pay a flat rate throughout the period of study.

Mandatory TOU firms comprise our treatment group and in our preferred specification the non-TOU firms comprise our control group. As a robustness check, in some specifications, we expand our control group to also include always-TOU firms. Though UI provided data on the entire population of customers, we restrict our sample to mandatory switchers and the customers that most closely resemble the mandatory switchers. First, we limit the sample of non-TOU firms to those customers reaching at least 75 kilowatts of peak load in any month, since this subset of flat-rate firms most resembles mandatory switchers in terms of size and peak load. Later, we expand the control group to include firms that are always on TOU pricing, since these firms are more similar to treatment firms in observables. We discuss our choice of control group in the empirical approach, and later test the robustness of our results to this choice.

During the period of study, 102 firms are mandated onto the TOU rate. Figure 1 plots a histogram of the calendar month in which a firm first faces mandatory TOU pricing. Over 77 percent of mandatory firms face their first month of mandatory TOU pricing in the summer of 2010. The modal month in which firms cross the mandatory TOU threshold is June 2010, the first month in which the mandatory kW threshold was reduced from 200 to 100. On average the lag between when a firm exceeds the TOU threshold and first faces TOU pricing is 2 months.

As shown in Table 2, a comparison of unconditional means suggests that mandatory TOU firms are the largest firms. The 1,444 firms always on a TOU rate are the second largest group of customers in size; the difference in unconditional means between TOU firms and always-TOU firms is statistically significant. On average, firms always subject to a flat rate are the smallest firms, when size is measured using electricity usage, peak load and monthly electricity expenditure. Lastly, for firms that are mandated onto TOU pricing, a comparison of raw means suggests that electricity usage, peak load and expenditure are higher once firms are mandated onto TOU pricing.

Figures 2 and 3 describe the time path of peak demand (Figure 2) and electricity usage (Figure 3) for flat-rate customers, mandatory TOU customers and always-TOU customers. Peak demand and electricity usage are averaged across firm type and calendar month. The vertical line corresponds to August 2010, the modal month in which customers first face mandatory TOU pricing. On average, peak demand for mandatory TOU firms ranges between 112 and 164 kw, clearly exceeding the 100 kW threshold implemented in June 2010. In the raw data, we do not observe a decrease in peak demand in anticipation of the impending threshold. Additionally, this figure illustrates that on average peak demand would have needed to be reduced by 10 to 40 percent for mandatory firms to avoid crossing the TOU threshold.

3.2 Load Research Data

The load research data report monthly peak demand, total monthly usage, peak usage and off-peak usage for a sample of commercial and industrial customers. These data are comprised of customers who always face a time invariant rate structure or always face a TOU rate structure; we do not observe mandatory switchers in these data. Firms in the load research data are a random sample of the commercial user population.

The bottom half of Table 3 provides descriptive results; these are reported by (i) decile of usage for the largest two deciles of users and (ii) vigintile for the largest two vigintiles of flat-rate firms. We choose to restrict our discussion of these data to the two largest deciles (and vigintiles of flat-rate customers) because we omit from the customer data all users whose demand never exceeds 75 kW. We observe differences in total usage, peak demand and the break down in peak and off peak usage across the largest two deciles of firms and the largest two vigintiles of flat-rate firms.

In addition to providing us with information on peak usage for all customers, these data allow us to calculate the TOU rate class discount. This discount is defined as the percentage decrease in a customer's bill simply from switching from a flat to TOU rate, holding usage and the load profile constant. In our analysis, we are interested in isolating the change in electricity expenditure attributable to a behavioral response; this requires us to net out the change in expenditure due to the rate class discount. To calculate this discount, we select the largest 5 percent of flat-rate firms in the load profile data, since these firms are closest to the TOU threshold, and calculate the bill counterfactual under TOU prices. The discount is calculated by firm-month, but we average up to an annual measure.[16] Table 4 shows the TOU discount for the top two usage vigintiles among flat-rate firms. On average, the TOU discount reduces kWh expenditures by 3.5 percent and kW expenditures by 40.7 percent.

3.3 Peak Usage

Ideally, our empirical analysis would include the ratio of peak to off-peak usage as a dependent variable. However, in the billing data we do not observe this ratio for (i) flat rate firms or (ii) mandatory switchers when they face a flat tariff. In these data, we can infer this ratio for always-TOU customers and mandatory switchers, once the the latter customers are on a TOU rate. The first two rows of Table 3 provide the fraction of usage occurring during peak hours for these two firm types. On average, we find that 35 percent of monthly usage for both mandatory TOU customers (once they are on a TOU rate structure) and always-TOU customers occurs during peak hours.[17] Along this observable, mandatory switchers and always-TOU firms, one firm type (at times) in our control group are similar.

To compare non-TOU firms, our control group, to mandatory TOU firms (once they are on TOU pricing) in terms of peak usage, we rely on the load profile data. Along two observables, monthly usage and peak load, flat-rate customers in the billing data are similar to flat-rate "load research" customers in 20th and 19th vigintiles. Relying on the similarities in these observables and the random nature of the load research data, we extrapolate the load profile for the "load research" customers to the flat-rate customers in the billing data. After extrapolation, we observe that the load profile of the mandatory TOU firms differs from that of the flat-rate firms. Between 37 to 39 percent of total usage occurs during peak hours for flat-rate firms as compared to 35 percent for mandatory switchers. As expected firms on a flat rate consume a larger fraction of electricity during peak hours since they have no incentive to shift from peak to off-peak usage. In terms of the fraction of electricity consumed during peak hours, monthly usage and peak load mandatory TOU

[16] Our load profile data extend through December 2010, so we assign the 2010 discount to 2011 months.

[17] As context, 24 percent or 40 of the 168 hours in the week are priced at peak rates.

firms differ from flat-rate firms. Differences in these observables between our treatment and control groups suggest that pre-existing trends in firm usage may be an important element in the analysis.

4 Identification and Empirical Approach

In this section we describe the empirical approach used to evaluate firm response to TOU pricing. We begin by estimating a simple difference-in-differences model on monthly electricity usage (kWh), peak load (kW), and electricity expenditure. In this specification we are able to control flexibly for a wide range of potential confounders. Then, in an attempt to diagnose the extent of the mean-reversion problem, we estimate the same specification again, but this time on a placebo treatment that mimics the TOU assignment in 2009, the year before the 100kW threshold actually takes effect. These results confirm the presence of mean reversion, which we then control for explicitly using a variety of approaches.

The simple difference-in-differences model exploits variation in the rate structure over time and across firms. Our panel dataset offers billing and usage outcomes before and after the introduction of TOU pricing, and the control group allows us to exploit the fact that TOU pricing is only mandated for firms whose peak demand exceeds the pre-determined threshold. The baseline specification is as follows:

$$y_{it} = \beta I_{it}^{TOU} + \alpha_t + \eta_i t + \gamma_i + \epsilon_{it} \tag{1}$$

In this specification, y_{it} is the natural log of our dependent variable of interest: either peak load, total usage or expenditure by customer i in month t. TOU pricing is denoted by I^{TOU}, an indicator set equal to one if firm i is on the mandatory TOU schedule in month t. The dependent variable depends on such factors as weather shocks or economy wide shocks, which we capture flexibly with the inclusion of month-by-year dummies, α_t. Firm fixed effects, γ_i, and firm-specific trends, $\eta_i t$, are also included to control for time-invariant unobservable characteristics and pre-existing trends at the firm level. The idiosyncratic error term,

ϵ_{it}, is given its conventional interpretation. We cluster standard errors at the firm to allow for correlation across all observations within a firm.[18]

The parameter of interest is β. The identifying assumption is that, conditional on fixed firm characteristics, aggregate period-level effects and firm trends, unobservables are not correlated with mandatory TOU pricing; that is, $E[I^{TOU}\epsilon] = 0$. Given the rich set of flexible control variables included, potential confounders such as weather, economic activity and fixed firm characteristics are eliminated. There are few remaining plausible confounders (which we subsequently discuss) that would impede a causal interpretation of β, the parameter of interest.

One potential issue when thinking about mandatory TOU pricing as an exogenous regulatory treatment is anticipation. Estimates of the treatment effect will be downward-biased if commercial customers invested preemptively in conservation or load-shifting before being switched onto the tariff. Prior to the implementation of mandatory TOU pricing, the DPUC required that UI implement a firm-specific educational campaign to inform treated commercial users about this program as well strategies they could take to control peak demand and electricity usage. Part of this educational effort also included bill comparisons that highlighted the potential impact of TOU rates on monthly electricity bills. Clearly, commercial users anticipated this program. If they changed demand upon notification of mandatory TOU pricing but before implementation of the this program, we will understate the program's effect. We don't believe this to be the case. As highlighted in Figure 2, if anything firms increased demand in the months preceding the introduction of mandatory TOU pricing. We also investigate whether firms, in anticipation of the mandatory TOU threshold, exhibited bunching around the TOU threshold. As shown in Figure 4, a k-density plot of peak load by firm type in June 2010 (the first month of the 100 kW threshold), the data do not reveal any clustering below the threshold as one would expect if firms in anticipation of the policy change reduced peak demand.

Our empirical strategy also requires that no other policies issued by the DPUC or the utility coincided with

[18] As discussed in Bertrand et al. (2004), this is similar to using Newey-West standard errors to control for within-firm autocorrelation, and allowing all lags to be potentially important.

the introduction of mandatory TOU pricing, while also causing changes in our outcome variables of interest. If other policies were introduced in the summer of 2010, when 75 percent of customers in our sample were mandated onto TOU pricing, and differentially affected control and treatment firms, then our treatment effect may be potentially contaminated. While the DPUC and utility introduced various programs targeting energy efficiency, these programs were implemented well before 2010 and on a voluntary basis.

Consistent estimation of β hinges on the "parallel trends" assumption which assumes that in the absence of treatment changes in control and treatment group outcomes between the pre- and post-treatment periods would have been the same. We control for any differences in trends across the two groups through the inclusion of firm-specific linear trends in the difference-in-differences specification.

4.1 Selection

In our empirical setting there are two firm types - flat-rate firms and always-TOU firms - from which to create a control group. The latter firm type consists of users who voluntarily opted into TOU pricing prior to the introduction of mandatory TOU pricing.[19] Firms that selected into the dynamic pricing regime are likely to be those that stand to benefit the most from this rate structure, perhaps because they have a low peak to off-peak load profile or can readily shift load from peak to off peak hours.

Restricting our control group to include only flat-rate firms allows us to generate internally consistent estimates of the treatment effect. Framed differently, a hypothetical re-randomization of assignment to control and treatment in this sample would produce consistent estimates as long as systematic differences between flat-rate and mandatory TOU firms do not exist. While the flat-rate firms are systematically smaller (by construction of the assignment rule) and consume a larger share of electricity during peak hours, we assume that conditional on firm-specific trends and firm fixed effects, there are no systematic differences in the percentage changes in the key dependent variables between control and treatment firms. In robustness checks,

[19] Some firms also volunteered onto TOU during the period of our sample, but those are dropped from our dataset.

we expand the set of control firms to include firms that are always on TOU pricing.[20]

4.2 Mean Reversion

A final potential confounder to our empirical approach is the possibility of positive and transitory unobserved (to the econometrician) shocks in demand for those firms assigned to treatment. In our setting, assignment to "treatment" is not random; it occurs if peak load exceeds a threshold. If a large transitory shock pushes firms into the treatment group, mean reversion may bias the treatment effect in equation 1 downwards. To see this, suppose that firm i receives a large unobserved shock (ϵ_{it}) to peak load in period t, causing firm i to exceed the mandatory TOU threshold. Conditional on the distribution, probability theory suggests that this firm is less likely to experience another high draw in $t+1$. Thus, even in the absence of treatment, we would observe a decrease in this firm's peak load simply because of mean reversion. Variables that are correlated with peak load (such as kWh and monthly bill) will suffer from the same effect.

In the next few paragraphs, we describe approaches to correct for the mean reversion problem in our setting.[21] The first two are very conservative, but eliminate a large number of observations. The third approach is almost identical to a regression discontinuity design. It relies on the full sample to pin down trends and month-by-year effects and then estimates the treatment effect from changes in differences (relative to the treatment assignment period) in the dependent variable across control and treatment firms.

Correction by Selection (CS): Mean reversion is a concern only when a transitory unobserved shock induces assignment to treatment. Firms with peak load that (i) always exceeds the 100kW threshold or (ii) is above the threshold due to movements captured by observables (seasonality or fixed firm characteristics) will be assigned to treatment regardless of unobserved shocks. In a first approach to control for mean

[20] We considered using propensity score matching on the entire population of customers but it was infeasible. There are no observables aside from the outcome variables on which to match firms.

[21] We explored other approaches to control for mean reversion. However, these approaches were inappropriate in our empirical setting.

reversion, we segment firms into groups depending on how often (in number of months) their peak load exceeds the threshold. Specifically, we run equation (1) on three subsets of the treatment firms: those above the threshold in more than 80 percent of months, those above in more than 60 percent of months, and those above in more than 40 percent of months.

Correction by Assignment-Period Omission (CAPO): Identification of the treatment effect comes from differences in post- and pre-treatment outcomes between treatment and control firms. Mean reversion will inflate the outcome variable in the treatment assignment period. If unobserved shocks are *iid*, dropping the assignment-period observations will eliminate the variation that equation (1) incorrectly attributes to treatment. If the treatment assignment occurs in the same period for all firms, then this month could be dropped for both treatment and control firms, and equation (1) could be estimated using the remaining observations. It is important to drop control group observations as well since the distribution of shocks to these firms is skewed by the exclusion of treatment firms. In our setting, treatment assignment occurs in different months for different firms, making it difficult to determine which control group months to drop. We circumvent this problem by limiting our sample of treated firms to those that cross the threshold in June 2010 since this is the modal assignment month, and dropping all observations in that month.

Regression Discontinuity (RD): In our setting, assignment to treatment is a discrete function of peak load. We exploit the discrete nature of the assignment rule and the continuous nature of peak load to control for mean reversion. In a third and preferred approach that parallels the method used by Chay et al. (2003), we include smooth functions of peak load and still retain sufficient variation to identify the treatment effect.[22] We operationalize this by implementing a two-step approach.

To control for seasonality and firm-specific trends (that differ systematically between control and treatment

[22]The intuition and method that we implement draws from the discussion in Chay et al. (2003). In their setting, assignment into a school infrastructure investment program is a discrete function of the previous year's test scores. Failure to account for the mean reversion induced by the assignment rule leads to biased estimates that alter qualitative conclusions about the program's effectiveness.

firms), in a first step the dependent variable is de-seasonalized and de-trended (by firm) to generate a new variable denoted by $\tilde{y}$. We then restrict our sample to June 2010 (the modal assignment month) onwards and calculate the difference in the dependent variable in time t relative to June 2010; $d_t ln\tilde{y} = ln\tilde{y}_t - ln\tilde{y}_0$, where $t = 0$ corresponds to June 2010.

In the second stage, we estimate a difference-in-differences that includes a function of the treatment period level as a control.

$$d_t ln\tilde{y}_{it} = \beta I_{it} + f(\tilde{y}_{i0}) + \lambda_{it}$$

The function $f(\tilde{y}_{i0})$ may enter linearly or as a higher-order polynomial.

This approach is nearly identical to a regression discontinuity design, except that the sample is not restricted to firms "close" to the cutoff.[23] In the presence of the control, $f(\tilde{y}_{i0})$, it is reasonable to believe that $E[Cov(I_{it}, \lambda_{it}) = 0]$, which is sufficient for consistent estimation of the treatment effect, β.

5 Results

The estimates of equation 1 on usage, peak load and electricity expenditure are reported in columns 1-3, respectively, of Table 5. In column 4, the outcome variable is the adjusted bill which describes monthly expenditure after netting out the TOU discount. A log transformation of the dependent variable allows β to be interpreted as a percentage change.

While we will soon show that these results are partly driven by mean reversion, if we estimate equation 1 not accounting for mean reversion, we find an economically and weakly statistically significant 4.6 percent reduction in peak load in response to TOU pricing. There also appears to be a conservation effect, with point estimates implying a 3.8 percent reduction in monthly usage though this result is not statistically significant. Mandatory TOU pricing has a negative and significant impact on electricity expenditure, producing a 12.5

[23] The data are too sparsely distributed around the threshold to restrict the sample to just these firms.

percent reduction in monthly expenditure. Part of this reduction in expenditure is attributable to the TOU discount, which is on average 8 percent. However, some of the bill reduction can be explained by a behavioral response. After controlling for the change in expenditure attributable to the TOU discount, monthly electricity expenditure reduces by 4.5 percent. Given the nature of the treatment assignment rule, before attributing these differences to TOU pricing, we explore the presence and extent of mean reversion.

5.1 Placebo Treatments

To detect whether mean reversion poses a problem, we construct a placebo treatment. This placebo also allows us to test if the treatment effects reported in Table 5 provide evidence of a response to TOU pricing. The placebo directly mimics the treatment assignment rule, except that we now introduce the 100kW threshold in June 2009, a full year before it actually went into effect. We impose a 2 month lag between the month in which the threshold is first crossed and the first month of TOU pricing (which is the modal delay in the data).

To measure the effect of this placebo treatment on the outcome variables, we again estimate equation 1 except now an indicator variable is set equal to one if firm i is on a placebo treatment in month t. Results are reported in Table 6. Since the threshold was inactive in 2009, in the absence of mean reversion we should estimate no effect of the placebo on the outcome variables. Indeed, we find that the placebo treatment does not induce a significant change in peak load or monthly usage. Further, compared to the results reported in Table 5 the coefficient estimates are closer to zero and noisier. However, a two sample t-test comparing the coefficient estimates in Table 5 to those reported for the placebo fails to reject a difference in the coefficient estimates. This inability to distinguish between the estimated treatment effects suggests that the treatment effects reported in Table 5 may reflect some combination of a true program effect and mean reversion. We now control for mean reversion to isolate the true treatment effect.

5.2 Mean Reversion Controls

Results from the placebo treatment experiment suggest that mean reversion may bias the estimated treatment effect reported in Table 5, causing us to overstate the response to mandatory TOU pricing. To account for this, we control for mean reversion using three approaches, the results of which are reported in Table 7.

Results from our preferred approach, an adjusted differences model, are presented in columns 1 and 2 of the upper panel of Table 7. In the adjusted differences model, we control for mean reversion by including a function of the assignment period level as a control, as shown in equation 2. In this model, the dependent variable is the difference of the logs between the dependent variable in time t and time $t = 0$. Once we control for mean reversion, we find that mandatory TOU pricing does not affect patterns of energy usage in an economically significant way. The coefficient estimates reflect a mean shift in monthly usage and peak load of -0.4 percent and -0.3 percent, respectively, from mandatory TOU pricing. We can also rule out effects greater than 1 percent in magnitude. Under the preferred approach, the standard errors shrink by an order of magnitude. This occurs because our dependent variable is now the difference of the logs between time t and time $t = 0$ rather than the log of usage (peak load or expenditure) in time t. Compared to the estimates reported in Table 5, treatment induces a smaller reduction in conservation and peak load, leading us to believe that mean reversion may be partly driving the earlier results.

In a second approach, we control for mean reversion by restricting the treatment group to firms that crossed the peak load threshold in June 2010 and excluding June 2010 observations from our sample. Results for approach 2 (CAPO) are presented in columns 3 and 4 of the upper panel of Table 7. Both in terms of magnitude and significance, TOU pricing has no impact on usage or peak load. Compared to the results reported in Table 5, the estimated treatment effects are now closer to zero. However, we cannot reject a response of 5 to 7 percent.

Results for approach 3 (CS) are presented in the bottom panel. Here the treatment group is restricted to firms whose peak demand exceeds 100 kw in more than 80 percent of the months (cols. 1-2), 60 percent of the months (cols. 3-4) and 40 percent of the months (cols. 5-6). For firms in these restricted samples, TOU pricing has little or no impact on firm behavior. As the treatment rule increases in stringency from 40 to 80 percent, TOU pricing moves from having a slightly, though statistically non-significant impact on behavior to a positive but still insignificant increase in peak load and electricity usage. Given the large standard errors, we again cannot reject a response of 5 to 7 percent from mandatory TOU pricing.

Overall, these three approaches to control for mean reversion generate qualitatively similar estimates in terms of magnitude, though the standard errors vary across approaches. In our preferred specification, we can reject a response of greater than 1 percent from mandatory TOU pricing. However, in the other mean reversion approaches, we cannot reject a response of 5 to 7 percent. Using the results from the preferred specification, we demonstrate that failure to account for mean reversion will lead us to mistakenly detect a behavioral response.

5.3 Heterogeneity and Bill Volatility

Despite the fact that we do not find an aggregate effect of mandatory TOU pricing on electricity usage, peak load or expenditure, it is plausible that certain firms or industries may still respond to this policy. Further, much of the resistance to TOU pricing (and dynamic pricing in general) focuses on unexpected bill increases and volatility, both of which may have differential effects across industry type. In this section, we first attempt to explain the mechanism that may drive a heterogeneous response to TOU pricing by investigating the role of electricity intensity in the response gradient. We then create a simple counterfactual to analyze the changes in monthly expenditure from TOU pricing. Finally, we calculate the anticipated changes in bill volatility that arise from TOU pricing.

We anticipate that treatment effects may be larger in some firms than in others.[24] One hypothesis to

[24] We attempted to estimate heterogeneous treatment effects at the coarsest of industry levels (1-digit NAICS code), but the data do not provide sufficient support across industry classifications to produce reliable estimates.

explain the differential response is rational inattention: it may be rational for firms to be inattentive to the electricity bill simply because of electricity's small contribution to costs. To test this hypothesis, we gather data on three measures of electricity intensity by industry, aggregate to the three-digit NAICS code and align these data by industry code with the billing data.[25] The three measures are electricity cost as a share of inputs, as a fraction of value added, and as a fraction of total industry output. We estimate equation 2, the regression discontinuity approach, in which we interact these measures of electricity intensity with the treatment indicators. The results are displayed in Table 8.

If the rational inattention hypothesis explains firm response, then we should find negative coefficient estimates on the variable that interacts electricity intensity with the TOU indicator. Framed differently, as the fraction of electricity intensity increases, firms should become more responsive to TOU pricing. Our results provide a noisy signal of rational inattention. While the standard errors are large, usage, peak load and expenditure all decrease with an increase in electricity intensity, regardless of the measure of electricity intensity used.

To examine the impact of TOU pricing on expenditure, we implement a simple counterfactual. During the first month a mandatory firm faces TOU pricing, we calculate the monthly bill if the firm was on a flat rate (after adjusting for the TOU discount). We then compare this flat rate bill to actual expenditure. Table 9 reports statistics on the distribution of bill changes in the first month of TOU pricing by quartile of usage and industry type. In general the observed bill changes are small. On average, bills increase by 0.06 percent, and 95 percent of firms experience a bill change of less than 8.5 percent, with many experiencing savings. Larger users in terms of kWh experience savings from TOU pricing while the bottom two quartiles of customers experience a 1.8 percent bill increase. The anticipated bill changes also vary by industry. At the extremes, in the absence of a response, industrial users on average incur a 1.8 percent increase in expenditure from this policy and entertainment, food and beverage firms experience a 3.4 percent decrease in expenditure.

Figures 5 and 6 show k-density plots of the distribution of bill level changes by kWh quartile and NAICS code,

[25] The Bureau of Economic Analysis provides input/output tables by industry and includes a line-item for electric power.

respectively. The distribution of effects on the smaller TOU firms has broad support, ranging from nearly a 10 percent savings to a nearly 15 percent bill increase.[26] In this cohort, 95 percent of firms experience bill increases of less than 12.7 percent. The largest users are much more likely to benefit. Firms in the top quartile save, on average, 2.5 percent and only 5 percent of these firms experience bill increases exceeding 2.8 percent. Of the NAICS segments, firms in manufacturing, retail, services/financial and non-profit/religion experience a broad distribution of bill changes. By comparison, industrial, educational, entertainment/food/beverage and government sectors are tightly clustered around a zero change in expenditure.

While changes in bill levels appear to be small, another criticism of TOU pricing is that it will lead to substantial increases in bill volatility. Our setting is well-suited to examine the potential of this policy to increase bill volatility. We estimate the coefficient of (unadjusted) bill variation for treated and control firms before and after June 2010. Bill volatility is influenced by seasonality, so we limit our sample of treated firms to those crossing the TOU threshold in the modal month (June 2010). This allows us to compare volatility to an analogous cohort (control firms), before and after June 2010. Table 11 displays the means of the coefficients of variation for control and treatment firms before and after June 2010. Bills of treated firms exhibit an increase in volatility of 7.4 percent (on average). However, we also calculate a 3.7 percent increase in volatility of control firm bills before and after June 2010, suggesting that much of the increase in volatility for treated firms is attributable to factors aside from TOU pricing. Thus, while TOU may lead to higher bill volatility, it is on average a small change.

5.4 Robustness Checks

To check the robustness of the quantitative results to our choice of control group, we re-estimate the specifications that comprise Table 7, but this time include the always-TOU customers in the control group. These firms are the most numerous of the large firms in our dataset (they increase the control group size

[26] Recall that while these are the smallest quartile of TOU firms, this group is comprised of rather large electricity users.

by 600 percent) and could have been grouped with the flat-rate firms in our primary control group. The always-TOU firms are desirable as controls since they are closer to the treated firms in size (kWh and kW) and the peak load ratio (see Tables 2 and 3).

The results are presented in Table 10, and are both quantitatively and qualitatively similar to the results from our primary specification with mean reversion controls with one exception. Now, in our preferred approach to control for mean reversion, we find a statistically but not economically significant response to TOU pricing. Again, we can rule out a response of greater than 1 percent. The corresponding 95 percent confidence intervals using the other two approaches include moderate behavioral response, but these results are nonetheless consistent with the qualitative conclusion that firms are not responding significantly to TOU pricing.

In each of our empirical approaches, we cannot rule out the possibility that a transitory treatment effect is offset by later changes in behavior. We investigate this possibility by taking a more flexible approach to treatment timing. We estimate an "event study" model that allows for separate effects in event-time space, as defined by proximity to the month a firm is assigned to the TOU tariff. Specifically, we estimate the following equation:

$$y_{it} = \sum_{k=-\underline{m}}^{\overline{m}} D_{it}^{k}\delta_k + \alpha_t + \eta_i t + \gamma_i + \epsilon_{it} \tag{2}$$

where D_{it}^{k} are a set of dummy variables set equal to one if, in calendar month t, firm i is k months away from its first treatment assignment month. We restrict the event study window such that $k \in [\underline{m}, \overline{m}]$, where $\underline{m} = -6$ and $\overline{m} = 6$, and normalize the coefficient of event time zero to zero.[27]

Estimates of the coefficients δ_k from equation 2 are plotted in Figures 7-9. These corroborate our earlier findings that assignment into TOU pricing induces little to no response, including no transitory response. The plots of post-TOU monthly usage, peak load, and adjusted bill event-time coefficients hover around zero. There appears to be a slight downward pre-treatment trend (similar to the so-called "Ashenfelter Dip"), but

[27] Additional indicators corresponding to "outside the event window" allow us to fully capture the dynamic effects of treatment.

recall that this is conditional on firm-specific trends. It is likely that there is some residual firm-specific cyclicality in usage that is not fully captured by the calendar month-by-year effects.

6 Conclusion

In this study we measure the response of commercial and industrial customers to mandatory TOU electricity pricing. Despite a significant shift in marginal prices, customers in our setting do not adjust their overall monthly consumption or the timing of their usage during the day. The apparent lack of response implies either that these consumers are perfectly price inelastic (in which case we should not be concerned about efficiency loss in the first place), or that the pricing regime that we study is not effective at transmitting meaningful economic incentives to customers.

Two unique features of our empirical setting allow us to contribute meaningfully to the ongoing debate on how and whether to implement TOU pricing. First, earlier experimental studies that find little response to TOU pricing argue that their results may be due to the temporary nature of the rate change (Aigner & Hirschberg 1985). In our empirical setting the rate change is permanent and as a consequence more likely to induce a response where capital investment is required. Yet we continue to find little change in usage or peak load. Second, our study describes the first C&I setting in the U.S. that does not give customers the opportunity to withdraw from TOU pricing. The opt-out feature that is characteristic of other programs will bias the estimated treatment effect towards a response, since firms capable of substituting within-day usage will remain in the study and those with a low substitution elasticity will exit the program.[28] As such, we provide the first internally-consistent measure of the impact of TOU pricing on C&I firms in the U.S..

A significant number of firms in our study area opted onto the TOU tariff before the mandatory policy went into effect, and as a result are eliminated from our control and treatment groups. This would diminish the external validity of our point estimates if, in practice, mandatory TOU pricing programs were not introduced

[28] Aigner & Hirschberg (1985) are forthcoming about this drawback.

in tandem with an opt-in feature. However, it is difficult to imagine this scenario given the regulatory climate that characterizes the debate. In practice, if mandatory TOU programs are implemented for a subset of customers, these will almost certainly be bundled with a voluntary counterpart that gives the rest of the customer base the option to enroll. Under this design, the fact that some firms volunteered for TOU pricing before a mandatory TOU program was introduced does not diminish the relevance of our results to other settings.

If one were to seek to use our results to inform a setting in which mandatory programs were implemented without giving firms the option to opt in voluntarily, then the interpretation should change slightly. A natural interpretation of our estimates is that we provide an internally-consistent estimate of the upper bound on the financial harm, and a lower bound on the response relative to its implementation in the population. Firms that volunteered for TOU pricing are likely to have either a favorable load profiles (low on-peak and high off-peak usage) or the ability to shift from peak to off-peak usage at a low cost. Relative to these voluntary switchers, the remaining firms that comprise our sample are less likely to change usage patterns in response to the TOU incentive, and also more likely to incur higher electricity bills from the high on-peak rate.

There are several viable hypotheses that may explain the apparent non-response to TOU pricing. One possibility is that it is the ratio of peak to off-peak prices that influences firm behavior. Alternatively, the absolute differential between peak and off-peak prices may alter behavior. Either of these explanations may hold for firms with the ability to substitute from peak to off-peak usage. On the other hand, firms that achieve on-peak usage reductions via conservation are more likely to respond to price levels rather than price ratios or differentials. Given that we observe neither conservation or load shifting, it appears that at the current prices these hypotheses are not driving firm behavior. Another possible explanation for the absence of a response is that prices or bills are not high enough to attract the attention of most bill managers. In fact, we find that as the share of electricity increases, firms are more responsive to TOU pricing, suggesting that rational inattention may partly explain the non-response. If true, this has implications for all flavors of dynamic pricing. Peak prices may need to be set high enough for firms to take notice, or technological automation may help to guide firm response.

Finally, our results indicate that the concern over mandatory (as compared to voluntary) TOU pricing in the C&I setting has been overstated. It has been argued that firms involuntarily switched into time-varying pricing would have to either engage in investments to shift their load profile or face an increase in electricity expenditure. In our setting, this is not the case. After adjusting for TOU discount, most firms experience a small (if any) increase in bill volatility and no bill change, with less than 5 percent of mandatory TOU customers experiencing a bill increase greater than 8.5 percent in the first month of the rate change.

References

[1] Aigner, D. and Hirschberg, J. (1985). "Commercial/Industrial Customer Response to Time-of-use Electricity Prices: Some Experimental Results," *RAND Journal of Economics*, 16(3): 341-355.

[2] Aigner, D. and Newman, J. and Tishler, A. (1994). "The Response of Small and Medium-size Business Customers to Time-of-Use (TOU) Electricity Rates in Israel," *Journal of Applied Econometrics*, 9(3): 283-304.

[3] Allcott, H. (2011). "Rethinking Real-time Electricity Pricing," *Resource and Energy Economics*, 33(4): 820-842.

[4] Bertrand, M. and Duflo, E. and Mullainathan, S. (2004). "How Much Should We Trust Differences-in-Differences Estimates?," *Quarterly Journal of Economics*, 119(1): 249-275.

[5] Borenstein, S. and Holland, S. (2005). "On the Efficiency of Competitive Electricity Markets With Time-Invariant Retail Prices," *Rand Journal of Economics*, 36(3): 469-493.

[6] Borenstein, S. (2005). "Time-Varying Retail Electricity Prices: Theory and Practice," in Griffin and Puller, eds., *Electricity Deregulation: Choices and Challenges*, Chicago: University of Chicago Press.

[7] Carlin, A. and Park, R. (1970). "Marginal Cost Pricing of Airport Runway Capacity," *American Economic Review*, 60(3): 310-319.

[8] Chay, K. and McEwan, P. and Urquiola, M. (2003). "The Central Role of Noise in Evaluating Interventions that use Test Scores to Rank Schools," *American Economic Review*, 95(3): 1237-1258.

[9] Faruqui, A. and Sergici, S. (2008). "The Power of Experimentation: New Evidence on Residential Demand Response," Working Paper, Brattle Group (May).

[10] Hurwicz, L. (1950). "Least Squares Bias in Time Series," in T. Koopmans, eds., *Statistical Inference in Dynamic Economic Models*, New York: Wiley.

[11] Houthakker, H. (1951). "Electricity Tariffs in Theory and Practice," *The Economic Journal*, 61(241): 1-25.

[12] Levine, M. (1969). "Landing Fees and the Airport Congestion Problem," *Journal of Law and Economics*, 12(1): 79-108.

[13] Nickell, S. (1981). "Biases in Dynamic Models with Fixed Effects," *Econometrica*, 49: 1417-1426.

[14] Pigou, A. (1912). *Wealth and Welfare*, London: Macmillan.

[15] Steiner, P. (1957). "Peak Loads and Efficient Pricing," *The Quarterly Journal of Economics*, 71(4): 585-610.

[16] Walters, A. (1961). "The Theory and Measurement of Private and Social Cost of Highway Congestion," *Econometrica: Journal of the Econometric Society*, 29(4): 676-699.

[17] Williamson, O. (1966). "Peak-load Pricing and Optimal Capacity Under Indivisibility Constraints," *The American Economic Review*, 56(4): 810-827.

[18] Wolak, F. (2007). "Quantifying the Supply-side Benefits from Forward Contracting in Wholesale Electricity Markets," *Journal of Applied Econometrics*, 22(7): 1179-1209.

Table 1: UI Electricity Price Schedule

	Non-TOU		TOU			
Basic Service Charge	39.19		66.82			
Per-Kw Demand Charge	6.12		3.63			
Per-Kwh costs	Summer	Winter	Summer		Winter	
			Peak	Off-peak	Peak	Off-peak
Standard Service Generation	0.1159		0.1360	0.1060	0.1360	0.1060
Delivery Charges	0.0074		0.0074		0.0074	
Distribution Charges	0.0208		0.0153		0.0153	
Competitive Transmission Assessment	0.0152		0.0152		0.0152	
Congestion Costs	-0.0011	-0.0010	-0.0024	0.0000	-0.0022	0.0000
Transmission Charges	0.0260	0.0208	0.0650	0.0000	0.0520	0.0000
Total per-kwh charge	0.1842	0.1791	0.2364	0.1439	0.2237	0.1439

Table 2: Summary Statistics

	Monthly Consumption (000s kWh)		Monthly Peak Load (kW)		Bill ($)			
	Mean	St.Dev	Mean	St.Dev	Mean	St.Dev	# firms	N
Mandatory TOU firms	39.1	29.2	134.6	78.6	7,708	5,405	102	2,209
Always TOU firms	33.6	32.5	117.9	158.3	6,331	6,278	1,444	34,642
Non-TOU firms	13.2	9.4	52.7	26.5	2,775	1,861	257	5,788
Mandatory TOU firms pre-treatment	35.4	24.5	120.5	51.7	7,330	4,831	102	1,370
Mandatory TOU firms post-treatment	45.1	34.9	157.6	105.2	8,326	6,182	102	839

Table 3: Usage and Demand for Customer and Load Research Data

	Monthly Consumption (000s kWh)		Fraction Peak Usage		Monthly Peak Load (kW)			
	Mean	St.Dev	Mean	St.Dev	Mean	St.Dev	# firms	N
Billing Data								
Mandatory TOU firms	39.1	29.2	0.35	0.11	134.6	78.6	102	2,209
Always TOU firms	33.6	32.5	0.34	1.81	117.9	158.3	1,444	34,642
Non-TOU firms	13.2	9.4			52.7	26.5	257	5,788
Load Profile Data								
10th Decile usage	94.8	70.9	0.32	0.08	285.1	350.8	100	2,238
9th Decile usage	29.1	9.8	0.36	0.08	88.2	45.2	101	2,246
20th Vigintile usage flat rate firms only	25.2	11.9	0.37	0.08	74.8	40.7	34	626
19th Vigintile usage flat rate firms only	10.9	3.2	0.39	0.11	38.2	16.7	36	682

Notes: The billing data are used to estimate the empirical models and the load profile data provide an additional observable - peak usage - along which to compare control and treatment firms. In the load profile data, means are reported for the largest two deciles of kwh users and the largest two vigintiles of flat rate users.

Table 4: TOU Rate Class Discount (2010)

	kWh	kW	Fixed Fee ($)
20th Vintile, Flat Rate Firms Only	3.5%	40.7%	-$27.63
	-3.7%	0.0%	
19th Vintile, Flat Rate Firms Only	2.7%	40.7%	-$27.63
	-5.0%	0.0%	

Note: kWh and kW figures are mean percentage bill reductions from flat rate firms switching to TOU, holding usage and load constant. The fixed fee discount is constant in dollar terms (and negative).

Table 5: Effect of TOU Pricing on kWh, kW, and Bill

	Mandatory TOU			
	ln(kWh)	ln(kW)	ln(bill)	ln(bill_adj)
TOU Indicator	-0.038	-0.046*	-0.125***	-0.045*
	(0.027)	(0.028)	(0.027)	(0.027)
Firm FEs	Y	Y	Y	Y
Firm Trends	Y	Y	Y	Y
Month-by-year FEs	Y	Y	Y	Y
R-Squared	0.295	0.245	0.355	0.353
Observations	7,971	7,971	7,997	7,997

** Significant at the 0.10 level, ** Significant at the 0.05 level, *** Significant at the 0.01 level. Standard errors clustered at the firm level. Control group: large flat rate firms.*

Table 6: Effect of Placebo Treatments on kWh, kW, and Bill

	ln(kWh)	ln(kW)	ln(bill)	ln(bill_adj
Placebo: Impose 100kW threshold beginning in June 2009, with 2 month delay before TOU activated				
Placebo	-0.027	-0.022	-0.015	-0.015
	(0.039)	(0.031)	(0.037)	(0.037)
Firm FEs	Y	Y	Y	Y
Firm Trends	Y	Y	Y	Y
Month-by-year FEs	Y	Y	Y	Y
R-Squared	0.93	0.85	0.91	0.91
Observations	7,897	7,873	7,938	7,938

** Significant at the 0.10 level, ** Significant at the 0.05 level, *** Significant at the 0.01 level.*
Standard errors clustered at the firm level. Control group: large flat rate firms.

Table 7: Controlling for Mean Reversion: Effect of TOU Pricing on kWh and kW

	Approach 1: Regression Discontinuity				**Approach 2: exclude June 2010**	
	Δln(kWh)	Δln(kW)			ln(kWh)	ln(kW)
TOU Indicator	-0.004	-0.003			-0.005	-0.033
	(0.003)	(0.002)			(0.027)	(0.024)
R-Squared	0.42	0.37			0.30	0.24
Observations	2,258	2,258			7,018	7,018
Approach 3: Treatment=1 if peak demand exceeds 100 kW in over…						
	40% of periods		60% of periods		80% of periods	
	ln(kWh)	ln(kW)	ln(kWh)	ln(kW)	ln(kWh)	ln(kW)
TOU Indicator	-0.028	-0.021	-0.006	-0.006	0.004	0.005
	(0.025)	(0.027)	(0.027)	(0.030)	(0.033)	(0.033)
R-Squared	0.926	0.848	0.923	0.842	0.922	0.837
Observations	6,868	6,868	6,558	6,558	6,349	6,349

Notes: Each specification includes firm fixed effects, period fixed effects and firm trends as additional controls.
** Significant at the 0.10 level, ** Significant at the 0.05 level, *** Significant at the 0.01 level.*
Control group: large flat rate firms.

Table 8: Treatment Gradient by 3-Digit NAICS Electricity Intensity

	Δln(kWh)			Δln(kW)			Δln(bill_adj)		
Electricity Share of…	Inputs	Value Add	Ind. Output	Inputs	Value Add	Ind. Output	Inputs	Value Add	Ind. Output
TOU Indicator	-0.003	-0.004	-0.004	0.005	0.003	0.005	0.000	0.000	0.001
	(0.009)	(0.007)	(0.009)	(0.008)	(0.007)	(0.009)	(0.009)	(0.007)	(0.009)
Electricity Share	-0.107	0.005	-0.195	-0.093	-0.011	-0.175	-0.099	0.037	-0.130
	(0.095)	(0.107)	(0.283)	(0.063)	(0.080)	(0.189)	(0.085)	(0.081)	(0.228)
TOU*Electricity Share	-0.059	-0.034	-0.101	-0.208	-0.186	-0.533	-0.091	-0.178	-0.331
	(0.255)	(0.316)	(0.729)	(0.236)	(0.299)	(0.705)	(0.254)	(0.309)	(0.713)
Assignment Period DepVar Level	Y	Y	Y	Y	Y	Y	Y	Y	Y
Firm Trends	Y	Y	Y	Y	Y	Y	Y	Y	Y
Month-by-year FEs	Y	Y	Y	Y	Y	Y	Y	Y	Y
R-Squared	0.42	0.42	0.42	0.36	0.36	0.36	0.42	0.42	0.42
Observations	1,556	1,556	1,556	1,556	1,556	1,556	1,556	1,556	1,556

** Significant at the 0.10 level, ** Significant at the 0.05 level, *** Significant at the 0.01 level. Standard errors clustered at the firm level. Control group: large flat rate firms. The specification here is directly analagous to the differences approach used to control for mean reversion. It uses treated firms from the modal treatment month (June 2010) and the dependent variable is in log-differences. Source: Bureau of Economic Analysis.*

Table 9: Summary Statistics of Bill Changes

	Mean	Std. Deviation	5th percentile	95th percentile
Overall	0.06%	4.72%	-5.75%	8.46%
By kWh Quartile				
Quartile 1 (Low)	1.78%	6.20%	-6.95%	12.73%
Quartile 2	1.79%	4.47%	-4.72%	10.01%
Quartile 3	-0.85%	3.13%	-5.15%	4.11%
Quartile 4 (High)	-2.54%	2.97%	-5.75%	2.79%
By NAICS				
Industrial	1.80%	3.51%	-1.53%	5.46%
Manufacturing	1.02%	5.19%	-7.29%	9.35%
Retail	1.20%	4.68%	-5.38%	8.46%
Services/Financial	0.37%	5.24%	-5.08%	11.03%
Educational	1.95%	3.65%	-5.23%	10.01%
Entertainment/Food/Beverage	-3.41%	1.72%	-5.75%	0.45%
Non-Profit/Religious	-0.93%	6.20%	-8.02%	12.73%
Government	1.29%	5.09%	-4.49%	5.11%

Table 10: Robustness Checks: Large Flat-Rate and Always-TOU Firms as Controls

	Approach 1: Regression Discontinuity				**Approach 2: exclude June 2010**	
	Δln(kWh)	Δln(kW)			ln(kWh)	ln(kW)
TOU Indicator	-0.005**	-0.003**			-0.034	-0.028
	(0.002)	(0.002)			(0.023)	(0.023)
R-Squared	0.43	0.38			0.90	0.82
Observations	13,822	13,816			43,771	43,751
Approach 3: Treatment=1 if peak demand exceeds 100 kW in over…						
	40% of periods		60% of periods		80% of periods	
	ln(kWh)	ln(kW)	ln(kWh)	ln(kW)	ln(kWh)	ln(kW)
TOU Indicator	-0.022	0.001	0.001	0.017	0.011	0.027
	(0.020)	(0.020)	(0.022)	(0.024)	(0.029)	(0.027)
R-Squared	0.90	0.81	0.90	0.81	0.90	0.81
Observations	41,910	41,890	41,426	41,406	41,214	41,194

Notes: Each specification includes firm fixed effects, period fixed effects and firm trends as additional controls.
** Significant at the 0.10 level, ** Significant at the 0.05 level, *** Significant at the 0.01 level.*
Control group: large flat rate and always-TOU firms.

Figure 1: Histogram Indicating Month Firms Exceed Mandatory TOU Threshold

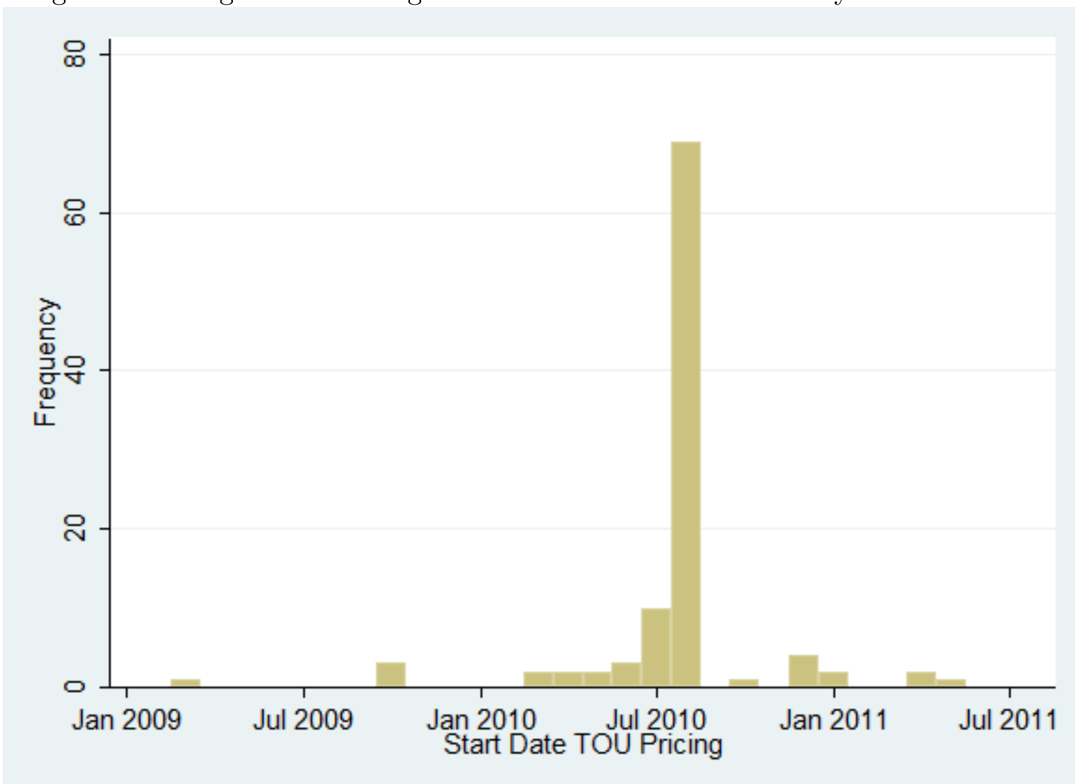

Figure 2: kW by Firm Type, June 2009 - May 2011

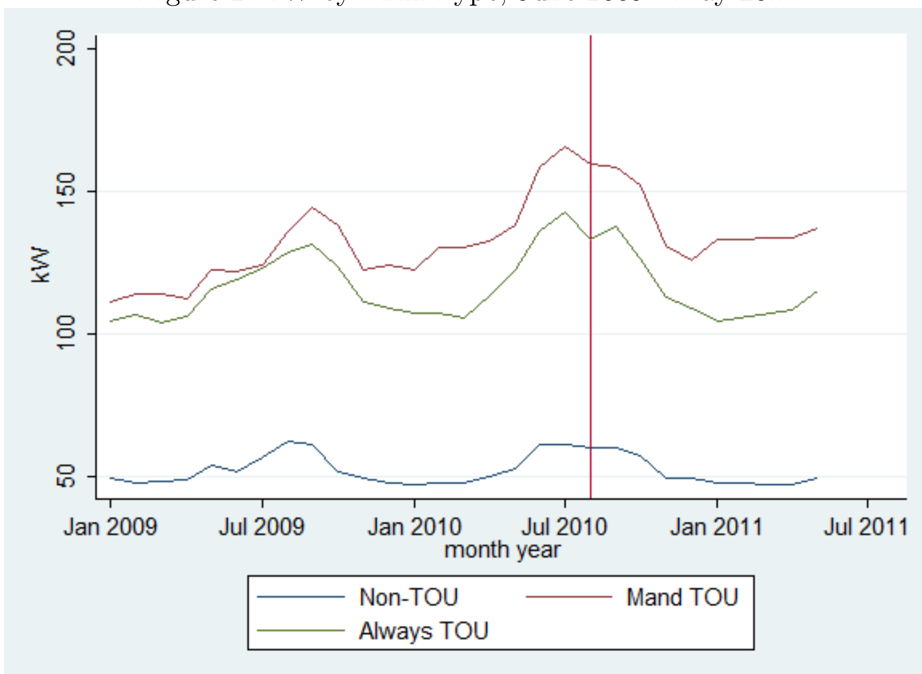

Figure 3: kWh by Firm Type, June 2009 - May 2011

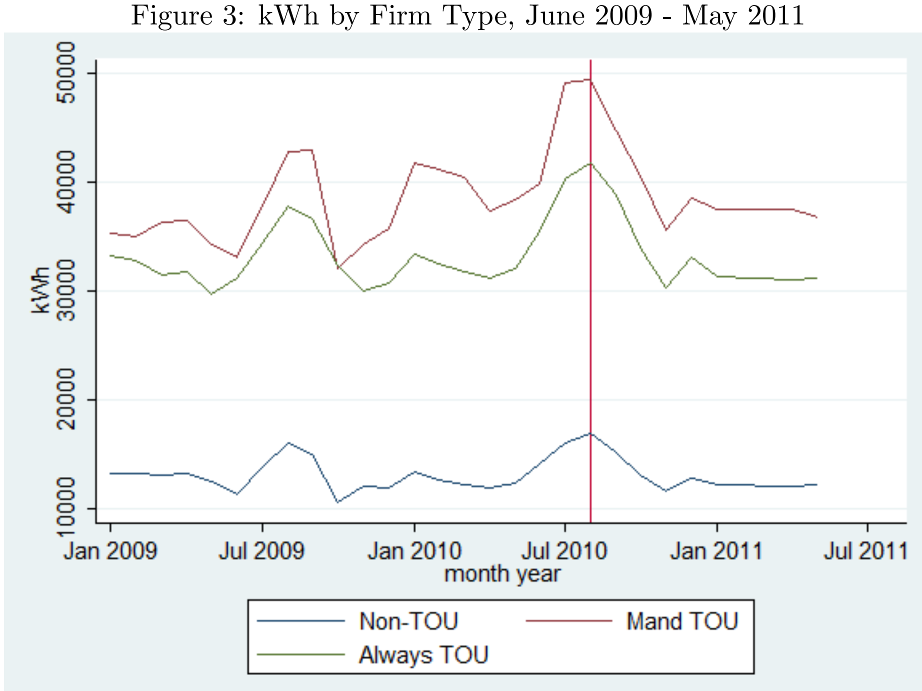

Figure 4: Kernel Density of kW in June 2010

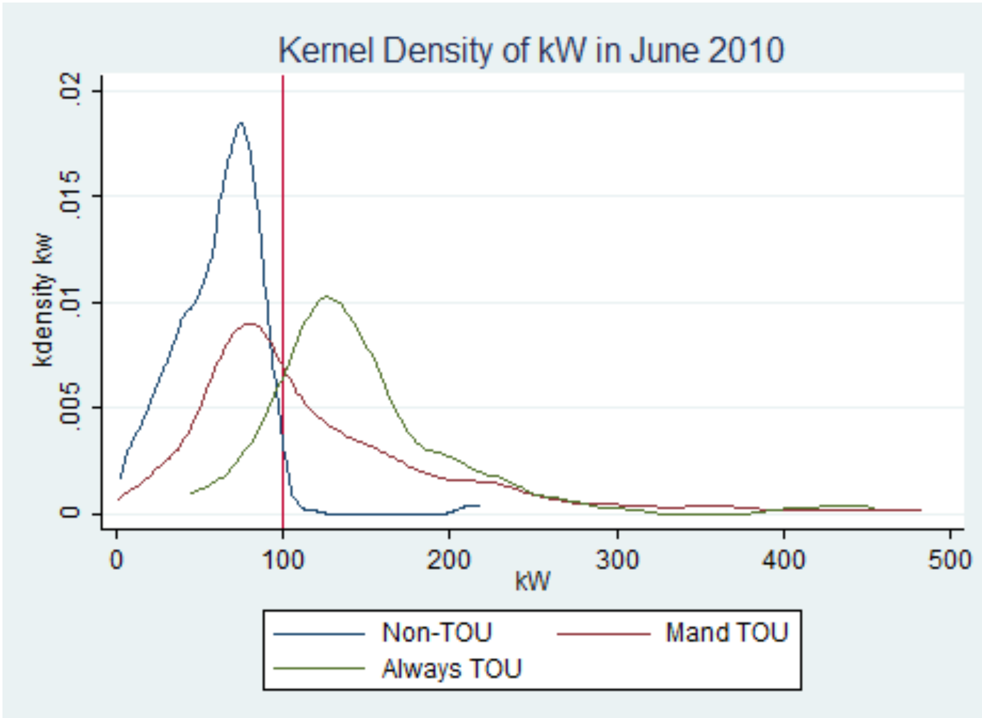

Figure 5: Kernel Density by kWh Quartile

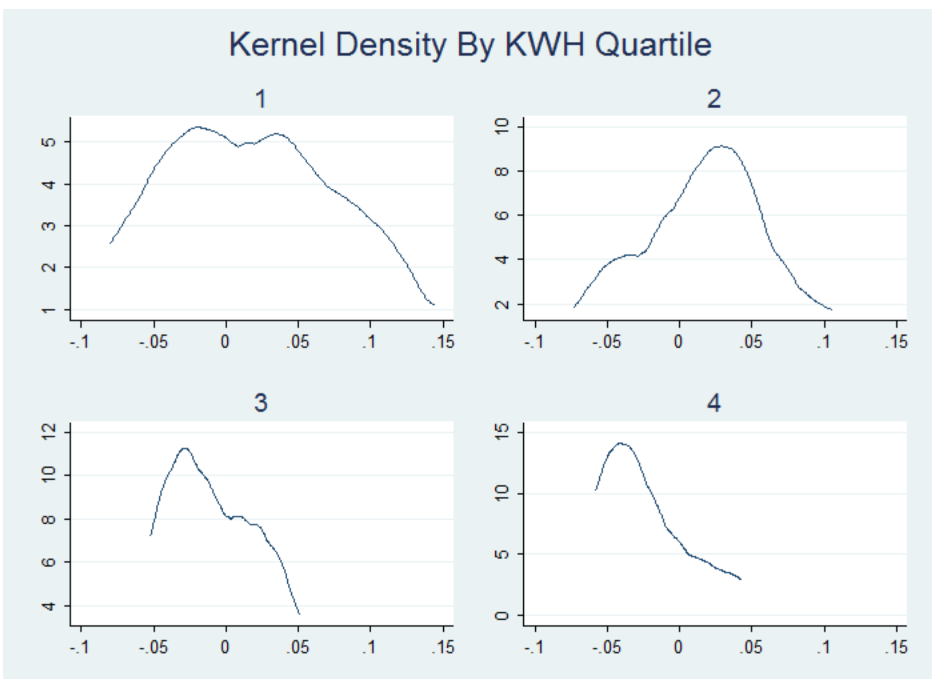

Figure 6: Kernel Density of Behavior-Invariant Bill Changes by Industry

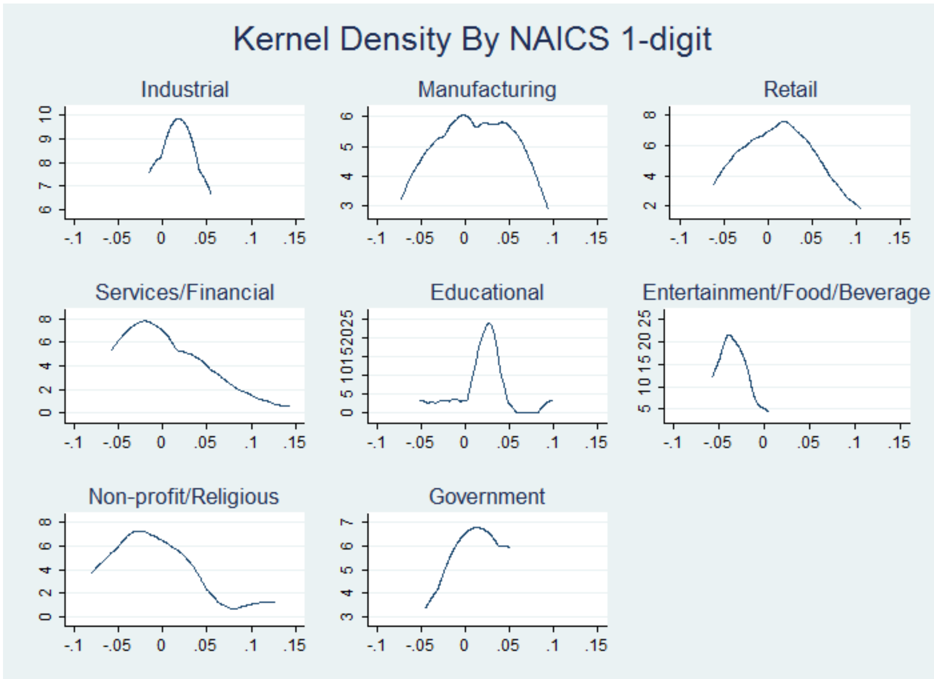

Table 11: Coefficient of Bill Variation: Treatment vs. Control Firms

	Pre-June 2010	Post-June 2010	Change
Non-TOU Firm	0.239	0.247	3.7%
TOU Firm	0.169	0.181	7.4%

Treated firms restricted to those assigned to TOU in June 2010.
Source: UI Billing Data

Figure 7: Dynamic Effects in Event Time: Monthly Consumption (kWh)

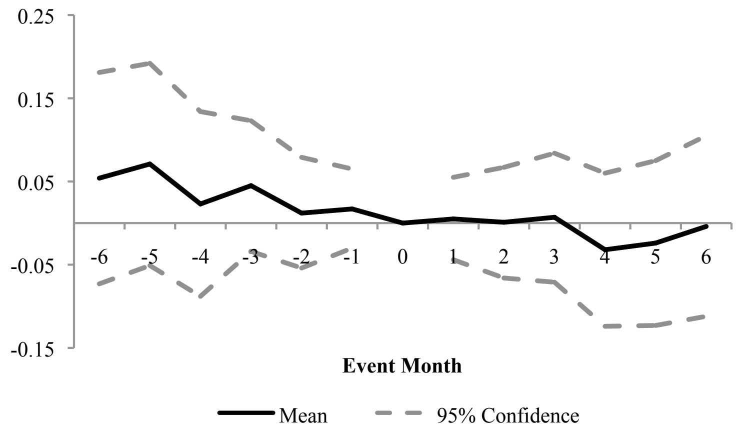

Figure 8: Dynamic Effects in Event Time: Peak Load (kW)

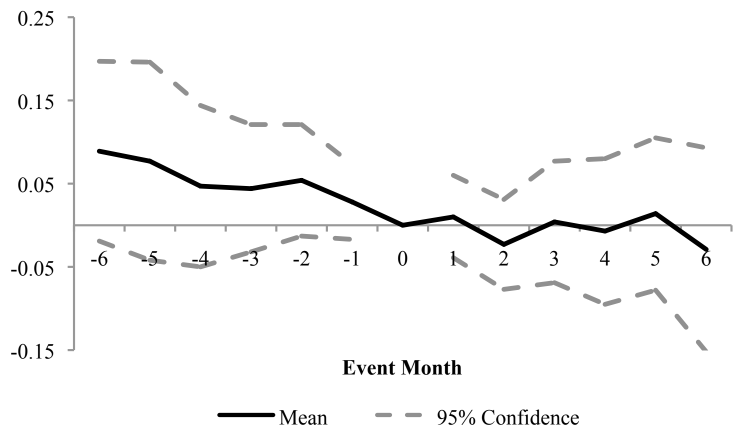

Figure 9: Dynamic Effects in Event Time: Adjusted Bill