ENERGY INSTITUTE AT HAAS

EI @ Haas WP 238R

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

Katrina Jessoe and David Rapson
February 2013

Revised version published in the *Journal of Industrial Economics* 63(3), 397–421, September 2015.

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

Katrina Jessoe[2] and David Rapson[3]

February 19, 2013

Abstract

This paper is the first to evaluate the impact of a large-scale field deployment of mandatory time-of-use (TOU) pricing on the energy use of commercial and industrial firms. The regulation imposes higher user prices during hours when electricity is generally more expensive to produce, and is by far the most common way for time-varying incentives to be transmitted to retail electricity customers. We exploit a natural experiment that arises from the rules governing the program to present evidence that TOU pricing induced no (or very little) change in customer usage or load. As such, economic efficiency was not increased by this regulation. Bill levels and volatility exhibit only minor shifts, suggesting that concerns from advocacy groups about increased expenditure and customer exposure to risk have been overstated.

JEL: D22, L50, L94, Q41

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

[1]We thank, without implicating, Lucas Davis, Kenneth Gillingham, Koichiro Ito, Chris Knittel, Aaron Smith and Jeffrey Williams. Thanks also to United Illuminating personnel for assistance in obtaining the data and shedding light on institutional details. Comments from seminar participants at AERE, RFF, Stanford, the UC Energy Institute, UCE3 and U. Melbourne helped to improve this draft. Tom Blake and Brock Smith provided excellent research assistance. We also thank the University of California Center for Environmental and Energy Economics (UCE3) for financial support. Rapson thanks the Energy Institute at Haas for support under a research contract from the California Energy Commission. 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, and an increase in the opportunity for fringe 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] The clear first-best policy would be a retail price which varies to reflect real-time fluctuations in the wholesale market. However, technological and political obstacles have forced regulators to proceed with caution, and, if deviating from flat-rates at all, implement a coarse variant of time-varying incentives called "time-of-use" (TOU) pricing.[3] Its effectiveness lies in its hypothesized ability to induce consumers to change their behavior by reducing demand during peak hours. This study documents evidence from one such deployment in the field for commercial and industrial (C&I) customers.

TOU pricing is the most common incentive-based tariff that is currently implemented (or being considered) by utilities and regulators to address peak load challenges. It may be viewed as a either a regulatory "baby step" towards dynamic pricing or a pragmatic compromise, depending on one's perspective. A perceived advantage of this rate structure over more granular versions is that customers can readily understand it and, in theory, respond to it. Its fundamental drawback is coarseness; many hours during "peak" are characterized by low wholesale price, despite also subsuming the most expensive hours. TOU pricing as currently devised can capture no more than

[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]Under a TOU tariff, two (sometimes 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.

6 percent of wholesale market price variation in Connecticut.[4] Facing such a strong constraint, it is important to understand whether TOU pricing is effective at achieving its the goal of reducing peak demand. If it is not, then either it should be replaced by a more effective policy (potentially a more granular and timely policy, like real-time pricing), or regulators should clearly state that TOU is intended as a transition path to such a policy. In this paper we present quantitative evidence that the response to a large-scale field deployment of mandatory TOU pricing for C&I customers in Connecticut was minimal.

The ideal dataset for evaluating TOU pricing would have certain key features, some of which (but not all) are present in our setting. First, TOU pricing would be randomly assigned to some firms. In our setting, the mandatory TOU assignment rule states that C&I firms whose peak load breaches a pre-determined threshold are placed on a TOU schedule and cannot, regardless of future behavior, return to a flat-rate tariff. In the neighborhood of the threshold, status as a "crosser" is plausibly random. While we do not observe sufficient density near the threshold to implement the standard regression discontinuity toolkit, the treatment assignment mechanism is plausibly exogenous.

A nuance of our setting is that in the years preceding the mandatory TOU policy, firms had the option to select onto the TOU rate. This does not impact the validity of our results, and in fact provides regulators and policy makers with results from the most relevant setting.[5] If other utilities introduce mandatory TOU pricing, as is being considered universally in the US today, it is certain to be accompanied by a voluntary adoption mechanism.[6] Thus it is not a coincidence that the setting observed in the first field deployment shares the most important features of potential future deployments being considered by policymakers in other jurisdictions.

[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.

[5]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. That the timing of voluntary adoption in our setting occurs years before the mandatory policy is activated is only problematic with respect to the effect of mandatory switchers if one 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.

[6]In fact, mandatory TOU is widely viewed as politically challenging. Borenstein (forthcoming), for example, states that "while mandatory or default dynamic pricing tari may be good policy, it is clear that it is unlikely to be politically acceptable for residential customers in the near future."

A second feature of the ideal dataset would be exogenous variation in the peak/off-peak price gradient, and any other relevant prices such as a per-kW peak demand charge. In our setting, we observe (more-or-less) two draws from the menu of potential prices.[7] Since the price schedule was produced by an actual rate determination process, we view it as the most relevant of the potential rates in our setting. However, it is reasonable to expect that rates conveying different incentives would produce different results. As such, our estimates of the customer response should be viewed as a single realization from this distribution. The near impossibility of observing a broad menu of price alternatives in a single field setting highlights the need for multiple peer-reviewed studies on this topic. To our knowledge there has not been one in the US in nearly three decades (Aigner & Hirschberg 1985).

Finally, the ideal dataset would offer cross-sectional and time-series variation that allows the researcher to account for important household characteristics and aggregate demand shifters. We use a monthly panel dataset from the universe of customers in the utility's "General Services" rate class, allowing us to exploit variation both over time and across space to estimate the impact of this policy on usage, peak load and expenditure. We take care to navigate several potential confounding factors when estimating the treatment effect. 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 regulation's effectiveness. We also address potential concerns about endogenous selection, and show that there is no evidence of treatment avoidance behavior.

The mechanism by which firms are assigned to treatment makes mean reversion an empirical challenge, which we address in multiple ways. Under the treatment assignment rule, any flat-rate firm with peak load exceeding 100kW after June 2010 is treated. If the peak load increase that caused the firm to exceed the assignment threshold is transitory, then peak load (and positively correlated variables like usage and billed amount) will likely decrease in the next period, regardless of the new pricing structure. Failure to account for this feature will cause the econometrician to

[7] "Control" households remain on a time-invariant rate, while the "treatment" (TOU) rate schedule in our setting imposes peak prices that are approximately 60 percent higher than off-peak. The demand charge also changes when firms are placed on TOU pricing, but this is less important due to the small bill share of the demand charge (roughly 10%).

incorrectly attribute a decline in load (usage or expenditure) to treatment.[8]

Our analysis yields three primary results. First, in the first 12 months following mandatory TOU pricing, 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 the analysis to account for the rate class discount, TOU pricing does not impact monthly electricity expenditure (lending support to our lack of evidence of measurable behavioral change). Finally, increases in bill levels and bill volatility are minimal, with only a small number of firms being adversely affected.

Our work contributes a new data point to the sparse empirical literature on C&I TOU pricing. Earlier evidence from electricity markets suggests that TOU pricing is at best moderately effective among C&I users (Aigner & Hirschberg 1985, Aigner et al. 1994).[9] 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 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. The permanence of TOU pricing in our setting is another feature that contributes to

[8] 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.

[9] 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 (Wolak 2007) while others induce conservation (Allcott 2010, Faruqui & Sergici 2011, and Jessoe & Rapson 2012). A review of residential dynamic pricing experiments can be found in Faruqui & Sergici (2010).

policy relevance.

Finally, our results shed light on the ongoing policy debate about whether TOU rates should be mandated or voluntarily implemented (if at all). 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 in our sample, 95 percent experience bill increases of less than 8.5 percent[11] and bill volatility increases on average by less than 5 percent.[12]

2 Theoretical Basis and Regulatory Setting

A vast economic literature dating back a century has described the theoretical rationale for implementing time-variant retail electricity prices. Williamson (1966) constructs a welfare framework that reveals what is now common knowledge: that the socially optimal amount of generation capacity will require some form of rationing during periods of peak demand, and that it can be achieved by equating the retail price to short-run marginal cost. This observation has been repeated in subsequent years, generally during periods when turbulent electricity markets are disruptive enough to earn popular attention. One of the clearest expositions, in our view, is from Borenstein (2005).

The argument proceeds as follows. Demand is highly variable, supply faces strict constraints in the short run, and economically-feasible storage at an industrial scale does not exist. In the event that demand exceeds supply, grid failure results in widespread "blackouts". On the other hand, excess supply damages expensive equipment. An engineering challenge in this industry is thus the

[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 successfully petitioned to delay its TOU deployment, though it nonetheless launched in November 2012.

[11] Even this is conservative, having been calculated using bill levels already adjusted for the TOU rate-class discount implicitly offered by the CT policy.

[12] High volatility could also be mitigated by the development and availability of simple hedging instruments, as suggested in Borenstein (forthcoming) and explored in Chao (2012).

necessity of equating supply and demand in real time. When demand is low the grid operator simply deploys less generation capacity; but when demand rises, generation facilities approach and then hit their capacity. At this point, either more marginal facilities must be brought online, or demand must be reduced in some way (often by implementing scheduled regional blackouts). The revealed preference in developed nations is to build sufficient excess capacity to cover demand during the peak hours. By implication, the marginal firms that are engaged in these few hours are unnecessary at all other times.

As it turns out, this situation results in large welfare losses due in large part to the time-invariant rates being charged to retail electricity customers. Marginal benefit equals marginal cost only by chance and for fleeting moments during the day or year, implying chronic over- or under-consumption with respect to the social optimum. Capacity is overbuilt as insurance against blackouts. And marginal firms (even those with low market share) face a sharply inelastic demand curve, giving them the capability to exercise market power during hours of system peak.[13] The cumulative welfare loss due to these features of the market have been estimated at approximately 5-10 percent of value in the wholesale electricity market (Borenstein & Holland 2005).

These facts are not lost on regulators, and most economists acknowledge the importance of better understanding market outcomes in this setting. In his work on peak load pricing, Steiner (1957) laments "an almost total absence of empirical evidence as to the importance of the potential shifting peak...". Recently with the proliferation of smart metering technology that allows electricity use to be measured at high frequency, there has been a renewed empirical focus on evaluating the potential of time-varying pricing. We add to this discussion by analyzing the first large-scale C&I field deployment of mandatory TOU.

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. Once transferred onto the TOU

[13] This was a major cause of the California electricity crisis in 2000, as described in Borenstein (2002).

tariff, a firm could not return to the flat-rate schedule, regardless of future consumption.[14] The first two demand thresholds took effect on June 1, 2008 at 300kW and June 1, 2009 at 200kW. The majority of small commercial users did not approach either of these. However, on June 1, 2010 the threshold declined to 100kW, and a substantial number of users were switched to TOU rates as a result of having crossed it.

In our setting, 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, unless they opted into TOU pricing.[15] 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 \$0.1791 in the winter and \$0.1842 in the summer. By comparison, TOU customers pay a higher price for electricity during peak hours, \$0.2237 in the winter and \$0.2364 in the summer, and a lower price during off-peak hours, \$0.14391 in the both the summer and the winter. This amounts to between a 55 to 64 percent increase in peak relative to off-peak prices, and is in the low range of price differentials when compared to the C&I TOU studies referenced earlier.[16] 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. Another feature of the regulatory setting is the discrepancy in the per kW demand charge between flat rate and TOU firms. Flat rate firms incur a demand charge of \$6.12 per kW during

[14]Firms could choose to purchase generation from an alternate supplier. However, as shown in Table 1 the majority of the TOU price differential is transmitted through distribution charges, which are charged to all customers, regardless of the generation supplier.

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

[16]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.

their hour of peak demand in a billing cycle, as compared to a \$3.63 per kW charge for TOU firms. While TOU firms face a much lower peak kW charge, this amounts to a small portion of a firm's monthly bill, comprising 8.2% of monthly expenditure for mandatory TOU firms and 12.7% of monthly expenditure for flat rate firms. 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).

3 Data

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 1,856 commercial users serviced between January 1, 2009 and August 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 dataset of 51,356 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 closely 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, 97 firms are mandated onto the TOU rate. Figure 1 plots a histogram of the calendar month in which a firm first crossed the mandatory TOU threshold. 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. Over 75 percent of mandatory firms faced their first month of mandatory TOU pricing in the summer of 2010.

As shown in Table 2, a comparison of unconditional means suggests that mandatory TOU firms are the largest firms. The 1,484 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 and peak load are higher and expenditure is lower once firms are mandated onto TOU pricing.

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.[17] 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.

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

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, and are only able to infer this ratio for always-TOU customers, and mandatory switchers once they are on a TOU rate. The first two rows of Table 3 provide the fraction of usage occurring during peak hours for TOU customers in the billing data (while they are on the TOU rate). On average, we find that 35 percent of monthly usage for mandatory TOU customers and 34 percent for always-TOU customers occurs during peak hours.[18] 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 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.

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

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.[19]

The parameter of interest is β. The identifying assumption is that, conditional on fixed firm

[19] 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.

characteristics, aggregate month 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 attenuated if commercial customers invested preemptively in conservation or load-shifting before being switched onto the tariff. This is particularly applicable given that the DPUC required that UI implement a firm-specific educational campaign prior to the introduction of the mandatory TOU policy. The campaign informed commercial users in our treatment group about the program, and strategies that they could take to control peak demand and electricity usage. The educational effort also included bill comparisons highlighting the potential impact of TOU rates on monthly electricity bills. Clearly, commercial users had sufficient information to anticipate the program. However, we now present evidence suggesting that no preemptive action was taken, which mitigates concerns about endogenous selection.

Figures 2 and 3 highlight k-density plots of firms' peak load around the 100kW cutoff in June 2010 (the modal assignment month) and various comparison months. "Bunching" just below the cutoff would be evident if marginal firms had both anticipated crossing the threshold and taken measures to avoid doing so. Similar to the approach used in Saez (2010) to detect bunching, we generate plots of the distribution of kW and check whether abnormal increases in the mass appear just below the threshold. Figure 2 compares the month of June across four years of our sample. Aside from what appears to have been a cool month in June 2009, there are no significant differences in the distribution of peak demand across years, and thus no evidence of avoidance behavior. Figure 3 compares the density of peak load in June 2010 with the adjacent months (May and July, 2010). Again, we find no evidence of bunching.

Our empirical strategy also requires that no other policies issued by the DPUC or the utility coincided with 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.[20] 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, we expand the set of control firms to include firms that are always on TOU pricing.[21]

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

[21]We considered using propensity score matching on the entire population of customers but it was infea-

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 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.

It would be natural to view the regression discontinuity (RD) design as a candidate solution to the mean reversion problem. The RD identifying assumption is that assignment to treatment is random in the neighborhood of the threshold. As the bandwidth around the discontinuity expands, the sample (and precision) increase, while the claim to random assignment weakens. Our setting does not provide support for the standard approach; the density of firms around the cutoff is too sparse. We develop three alternatives. The first two are very conservative, but eliminate a large number of observations. The third approach is close to a regression discontinuity design, but using a large bandwidth.

Correction by Selection (CS): Mean reversion is a concern only when a transitory unobserved shock induces assignment to treatment. Firms (i) with peak load that always exceeds the 100kW threshold or (ii) who are 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 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.

---

sible. There are no observables aside from the outcome variables on which to match firms.

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.

Pseudo-Regression Discontinuity (RD): This method relies on the full sample to pin down trends and month-by-year effects in a first stage, 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. We include smooth functions of the "forcing variable" (peak load) as controls.[22] We operationalize this as a two-stage approach as follows.

To control for seasonality and firm-specific trends (that differ systematically between control and treatment firms), we first de-seasonalize and de-trend (by firm) the dependent variable to generate a transformed variable denoted by $\tilde{y}$. We restrict the sample to June 2010 (the modal assignment month) onwards and calculate the difference in time t relative to June 2010; $d_t\tilde{y} = \tilde{y}_t - \tilde{y}_0$, where $t = 0$ corresponds to June 2010.

In the second stage, we estimate OLS on a many-differences model that includes a function of the treatment period level as a control.

$$d_t\tilde{y}_{it} = \beta I_{it} + f(\tilde{y}_{i0}) + \lambda_{it} \tag{2}$$

[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.

The function $f(\tilde{y}_{i0})$ may enter linearly or as a higher-order polynomial. The identifying assumption is that $E[Cov(I_{it}, \lambda_{it}) = 0]$.

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. The coefficient of interest β can 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 statistically significant 6.9 percent reduction in peak load in response to TOU pricing. There also appears to be a meaningful conservation effect, with point estimates implying a 5.6 percent reduction in monthly usage that is statistically significant at 90 percent. Mandatory TOU pricing has a negative and significant impact on electricity expenditure, producing a 14.9 percent reduction in monthly expenditure. Part of this reduction in expenditure is attributable to the TOU discount, 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 7.0 percent. Given the nature of the treatment assignment rule, before attributing differences in usage, peak demand and expenditure 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.[23] 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, a regression discontinuity model, are presented in columns 1 and 2 of the upper panel of Table 7. In this 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 between the dependent variable (which is defined in logs) 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

[23] For the placebo exercise, we restrict the sample to pre-June 2010 so as not to include subsequent treatment months.

estimates reflect a mean shift in monthly usage and peak load of -0.5 percentage points and -0.3 percentage points, respectively, from mandatory TOU pricing. We can also rule out effects greater than 1 percentage point 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, confirming our suspicion that mean reversion was 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 and reveal that TOU pricing has a small impact on usage and peak load, the latter of which is statistically significant at 90 percent. However, compared to the results reported in Table 5, the estimated treatment effects are closer to zero. While we cannot reject a response of 6 to 9 percent under this specification, our robustness checks (implemented below with an alternate choice of control group) will generate much smaller treatment effects using the same mean reversion correction.

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. Using this approach, we can rule out an average load reduction of more than 6 percent from mandatory TOU pricing.

Overall, these three approaches to control for mean reversion generate qualitatively similar estimates, with each implying little to no demand response from mandatory TOU pricing. In our preferred specification, we reject a response of greater than 1 percentage point. Using the other mean reversion corrections, we cannot reject somewhat larger responses (5 to 9 percent); however as we will show below, these effects may be driven by selection of the control group. In any case, we demonstrate that failure to account for mean reversion will lead one to falsely detect a behavioral response.

5.3 Bill Changes and Volatility

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 create a simple counterfactual against which to analyze the changes in monthly expenditure from TOU pricing. We also calculate the anticipated changes in bill volatility that arise from TOU pricing.

The counterfactual exercise consists of calculating what the TOU customer's bill would have been if instead the firm had stayed on a flat-rate tariff. 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 8 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 4 and 5 show k-density plots of the distribution of bill level changes by kWh quartile and NAICS code, 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.[24] In this cohort, 95 percent of firms experience bill increases of less than 12.7 percent. The largest users are much

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

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 9 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 Retail Competition

In 2000 electricity generation in Connecticut was opened up to retail competition, offering customers in UI's territory the option to purchase electricity from retailers other than UI.[25] Customers opting to purchase service from an alternate retailer continue to receive their electricity bill from UI, but the generation charge is calculated based on the rate charged by the retailer. In our sample, almost 70 percent of firms purchased electricity from an alternate retail provider at some point in time.

The option to purchase generation at a flat-rate may influence our results if firms purchase generation from an alternate supplier to avoid the UI TOU rate. And in fact public listings of current alternate suppliers show that all rates offered by alternative retailers are time-invariant, confirming

[25] If a customer does not choose to purchase electricity from an alternate supplier, UI is the default provider.

an assertion made to us by a UI account manager. Regulators in Connecticut may have recognized the possibility of avoidance behavior, since the price differential between peak and off-peak hours is primarily transmitted through transmission and distribution charges. Still as shown in Table 1, approximately one-third of the peak/off-peak differential in prices is transmitted through the generation charge. When firms on TOU pricing purchase electricity from an retail supplier, the overall price differential drops from 8-9.25 cents to 5-6 cents.

For two reasons this institutional feature would, if anything, bias our estimates towards zero. First, the response to time-varying prices depends on the peak to off-peak price differential. Since this differential is muted for firms relying on alternate suppliers, we would expect them to be less responsive to mandatory TOU pricing. Second, firms with a more "peaky" load profile or less capacity to shift usage to lower price hours of the day may have selected into the alternate retailer as a cost-saving measure. In either case, one expects that firms purchasing generation from an alternate supplier would be less responsive to mandatory TOU pricing.

To test the extent to which the estimated non-response in Table 7 is driven by avoidance behavior, we estimate how the response to mandatory TOU pricing differs across firms purchasing electricity from UI and from other suppliers. To do this, we estimate a variation of equation (2) (which implements the regression discontinuity correction for mean reversion) in which we interact the TOU pricing indicator with an alternate supplier indicator set equal to 1 if a firm purchases electricity from an alternative supplier in month t. Results are reported in Table 10, where the control group is defined as flat-rate firms purchasing generation from UI. Compared to control firms, there is a small increase in kW for non-TOU firms purchasing electricity from an alternate supplier, though percentage point changes of more than 1.5 percent can be ruled out. Turning to look at the impact of mandatory TOU pricing, we find no significant impact of the policy on peak load for firms purchasing electricity from UI or from an alternate supplier. Further, we fail to reject a differential response across firm type. Results for monthly usage are economically similar, though less significant. These results highlight that avoidance behavior among firms using alternate suppliers is not driving the overall result of no response.

5.5 Robustness and Dynamic Effects

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 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 11, and for the most part are qualitatively similar to the results from our primary specification with mean reversion controls. There are two exceptions. Now, in our preferred approach (regression discontinuity) to control for mean reversion we find a statistically significant response in kWh to TOU pricing. However, the point estimate remains very small (less than 1 percent) and we can rule out usage decreases of more than 1.5 percent. The second mean reversion correction (excluding June 2010) produces point estimates for load and usage response that are much smaller now than those presented in Table 7. We can rule out load reductions of more than 4 percent. These results remain 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 first faces the TOU tariff. Specifically, we estimate the following equation:

$$y_{it} = \sum_{k=-\underline{m}}^{\overline{m}} D_{it}^{k}\delta_k + \alpha_t + \eta_i t + \gamma_i + \epsilon_{it} \tag{3}$$

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 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.[26]

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

Estimates of the coefficients δ_k from equation 3 are plotted in Figures 6-8. 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, which should be interpreted in light of the fact that these specifications control for firm-specific trends. The relatively wide 95 percent confidence bounds reflect the fact that we estimate a much larger number of treatment coefficients (one per pre- and post-event month) than in the previous specifications, which erodes precision.

Finally, we present a slightly different representation of the potential dynamic effects of treatment. The event study methodology does not correct for mean reversion, and also is generally implemented to reflect a relatively short time-period before and after treatment (in our case six months). One may be concerned that response to treatment may take time if firms are responding by investing in capital. If this hypothesis is correct, then we would expect to see larger (negative) treatment effects over time. Table 12 presents estimated treatment effects from the regression discontinuity specification (equation (2)) in three-month bins according to duration from treatment. Post-treatment data for firms treated in the modal month will appear for 13 months. We find no evidence of the treatment effect increasing as might be expected in a time-to-build capital investment scenario. In fact, in the 10+ month post-treatment bin (which corresponds to the summer months of 2011) there is a small but insignificant *increase* in usage. Of course, this analysis will not capture investments that require more than a one-year time horizon.

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 exhibit reductions in peak load or overall usage. 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.

Three unique features of our empirical setting allow us to contribute meaningfully to the ongoing debate on how and whether to implement TOU pricing. First, we examine a mandatory deployment, which is a rare contrast to the more frequent strategy of requiring customers to opt in. Second, 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 in the first full year following mandatory TOU pricing. Third, 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.[27] As such, we provide the first credible measure in several decades of the impact of a mandatory 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. Since our empirical focus is on the mandatory nature of the rollout, the voluntary adopters are eliminated from our sample (ie. from both control and treatment groups). One may wonder if this diminishes the external validity of our point estimates if, in practice, mandatory TOU pricing programs were not introduced 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-valid estimate of the upper bound on the financial harm, and a lower bound on the response relative to

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

its implementation in the population. Firms that volunteered for TOU pricing are likely to have either a favorable load profile (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.

Finally, our results indicate that concerns over mandatory (as compared to voluntary) TOU pricing in the C&I setting have 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. Even after adjusting for the 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. As such, despite the apparent lack of behavioral change induced by the TOU prices, regulators may still view TOU pricing as a way to smooth the path to more timely and granular dynamic pricing (such as RTP) without exposing customers to high bill volatility.

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. (forthcoming). "Effective and Equitable Adoption of Opt-In Residential Dynamic Electricity Pricing," *Review of Industrial Organization*.

[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] 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.

[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] Chao, H. (2012). "Competitive Electricity Markets with Consumer Subscription Service in a Smart Grid," *Journal of Regulatory Economics*, 41: 155-180.

[10] Faruqui, A. and Sergici, S. (2010). "Household Response to Dynamic Pricing of Electricity: A Survey of 15 Experiments," *Journal of Regulatory Economics*, 38: 193-225.

[11] Faruqui, A. and Sergici, S. (2011). "Dynamic Pricing of Electricity in the Mid-Atlantic Region: Econometric Results from the Baltimore Gas and Electric Company Experiment," *Journal of Regulatory Economics*, 40: 82-109.

[12] Jessoe, K. and Rapson, D. (2012). "Knowledge is (Less) Power: Experimental Evidence from Residential Energy Use," *Working Paper*.

[13] Saez, E. (2010). "Do Taxpayers Bunch at Kink Points?" *American Economic Journal: Economic Policy*, 2(3): 180-212.

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

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

[16] 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 (Fixed)	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
Total per-kWh charge, net generation	0.0682	0.0632	0.1005	0.0379	0.0877	0.0379

Table 2: Summary Statistics

	Monthly Consumption (000s kWh)		Monthly Peak Load (kW)		Bill ($)		
	Mean	St.Dev	Mean	St.Dev	Mean	St.Dev	# firms
Mandatory TOU firms	38.3	25.1	131.1	58.2	7,439	4,657	97
Always TOU firms	33.8	31.5	117.6	159.0	6,280	5,935	1,484
Non-TOU firms	13.3	9.8	52.9	26.1	2,796	1,918	275
Mandatory TOU firms pre-treatment	37.6	24.8	129.2	56.5	7,792	4,962	97
Mandatory TOU firms post-treatment	39.0	25.5	133.0	60.0	7,087	4,328	97

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
Billing Data							
Mandatory TOU firms	38.7	24.3	0.35	0.11	131.9	55.2	97
Always TOU firms	33.8	31.5	0.34	0.73	117.6	159.0	1,484
Non-TOU firms	13.3	9.8			52.9	26.1	275
Load Profile Data							
10th Decile usage	96.0	63.1	0.32	0.07	295.0	360.0	100
9th Decile usage	29.0	7.3	0.36	0.08	88.2	34.9	101
20th Vigintile usage flat rate firms only	26.5	11.3	0.37	0.08	80.3	39.9	34
19th Vigintile usage flat rate firms only	10.7	2.7	0.38	0.10	37.6	15.0	36

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 Vigintile, Flat Rate Firms Only	3.5%	40.7%	-$27.63
19th Vigintile, Flat Rate Firms Only	2.7%	40.7%	-$27.63

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	-0.056*	-0.069**	-0.149***	-0.070**
	(0.029)	(0.034)	(0.029)	(0.030)
Firm FEs	Y	Y	Y	Y
Firm Trends	Y	Y	Y	Y
Month-by-year FEs	Y	Y	Y	Y
R-Squared	0.270	0.241	0.292	0.290
Observations	9,512	9,513	9,553	9,553

** 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	-0.005	-0.003			-0.023	-0.041*
	(0.003)	(0.002)			(0.023)	(0.022)
R-Squared	0.45	0.34			0.28	0.25
Observations	3,403	3,402			8,435	8,434
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	-0.013	-0.003	0.007	0.008	0.011	0.022
	(0.027)	(0.027)	(0.029)	(0.030)	(0.039)	(0.034)
R-Squared	0.922	0.845	0.921	0.843	0.918	0.833
Observations	8,245	8,244	7,899	7,898	7,529	7,528

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: 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 9: 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

Table 10: Treatment Effect by Retail Supplier Type

	Mandatory TOU	
	Δln(kWh)	Δln(kW)
MandTOU*1[No Alternate Supplier]	-0.003	-0.005
	(0.009)	(0.008)
MandTOU*1[Yes Alternate Supplier]	0.001	0.005
	(0.005)	(0.004)
1[Yes Alternate Supplier]	0.006	0.006*
	(0.004)	(0.003)
R-Squared	0.451	0.341
Observations	3,403	3,402

Results generated from a regression discontinuity specification; dependent variable is the difference in demeaned and deseasonalized ln(kwh) or ln(kw) with respect to the modal switching month, June 2011.

** 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 11: 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	-0.007*	-0.006			-0.007	-0.011
	(0.004)	(0.004)			(0.017)	(0.015)
R-Squared	0.45	0.46			0.22	0.21
Observations	21,473	21,455			50,260	50,223
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	0.003	0.025	0.026	0.038	0.032	0.051*
	(0.022)	(0.021)	(0.026)	(0.024)	(0.036)	(0.028)
R-Squared	0.90	0.81	0.90	0.81	0.90	0.81
Observations	50,478	50,446	49,941	49,910	49,521	49,490

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.

Table 12: Dynamic Treatment Effect (RD Specification)

	Mandatory TOU	
	Δln(kWh)	Δln(kW)
1-3mths post-treatment	0.015	0.005
	(0.017)	(0.018)
4-6mths post-treatment	-0.02	-0.007
	(0.017)	(0.019)
7-9mths post-treatment	-0.046	-0.026
	(0.033)	(0.031)
10+mths post-treatment	0.023	0.011
	(0.016)	(0.016)
R-Squared	0.452	0.341
Observations	3,403	3,402

*Results generated from a regression discontinuity specification; dependent variable is the difference in demeaned and deseasonalized ln(kwh) or ln(kw) with respect to the modal switching month, June 2011. * 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.*

Figure 1: Histogram Indicating Month Firms Exceed Mandatory TOU Threshold

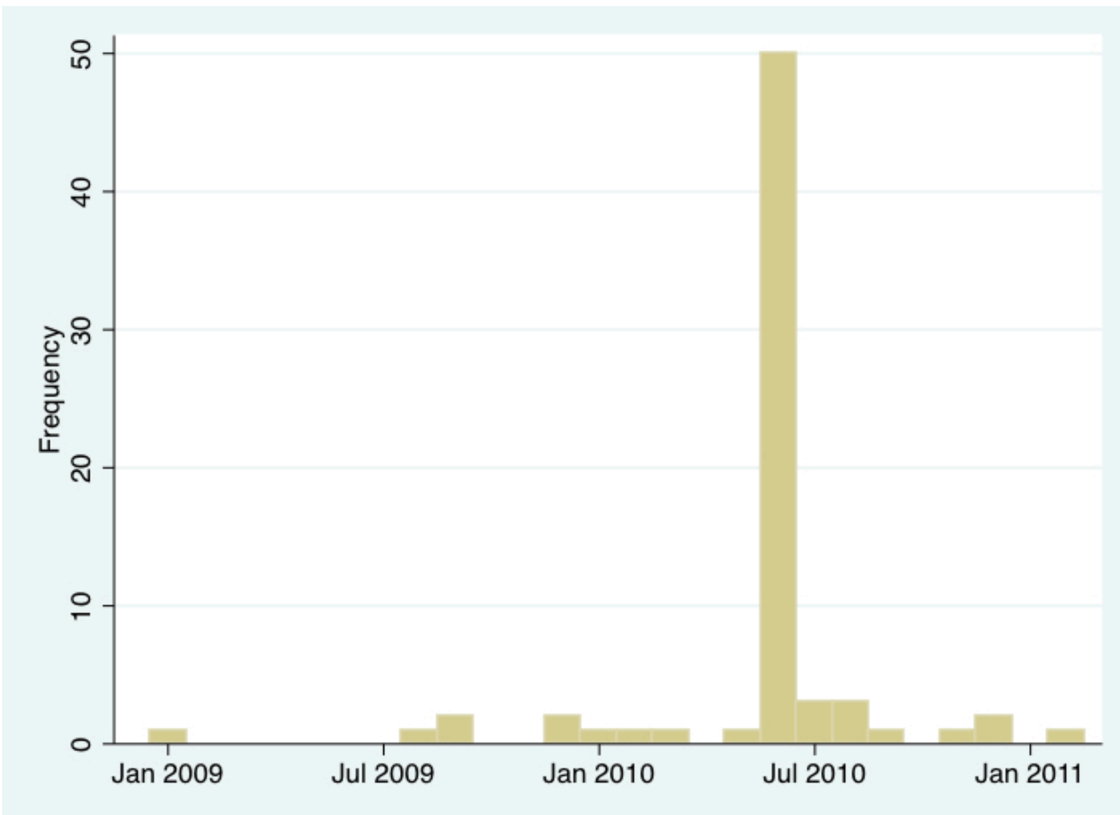

Figure 2: Kernel Density of kW, June 2008-2011

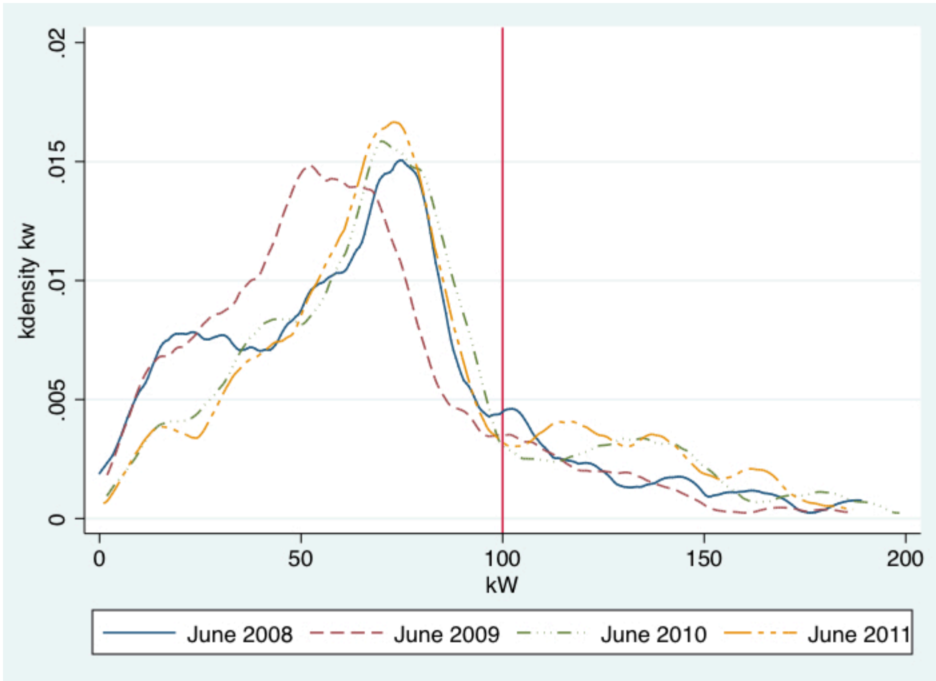

Figure 3: Kernel Density of kW, May-July 2010

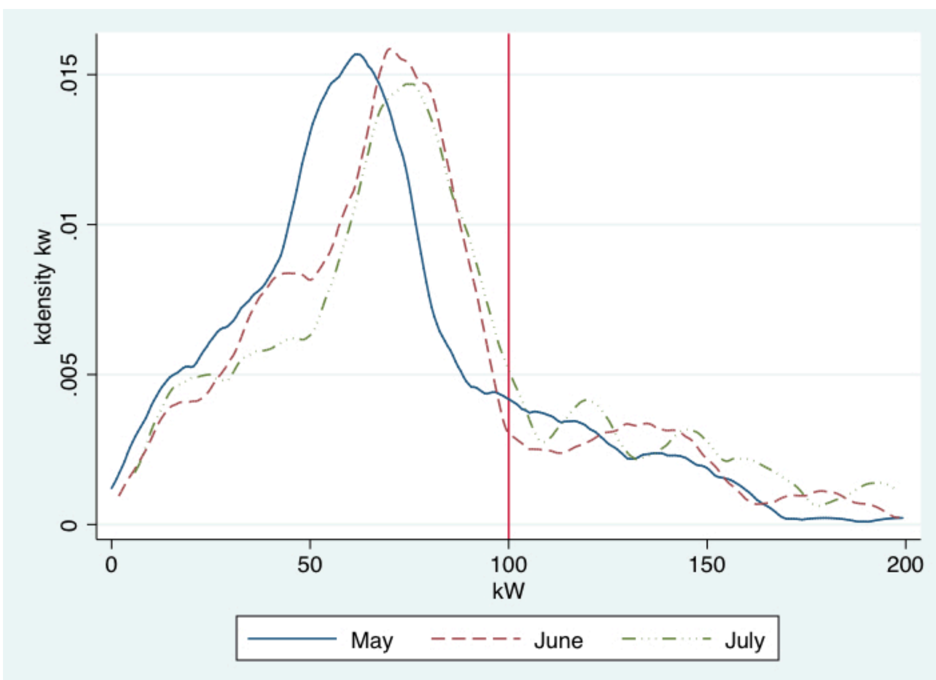

Figure 4: Kernel Density by kWh Quartile

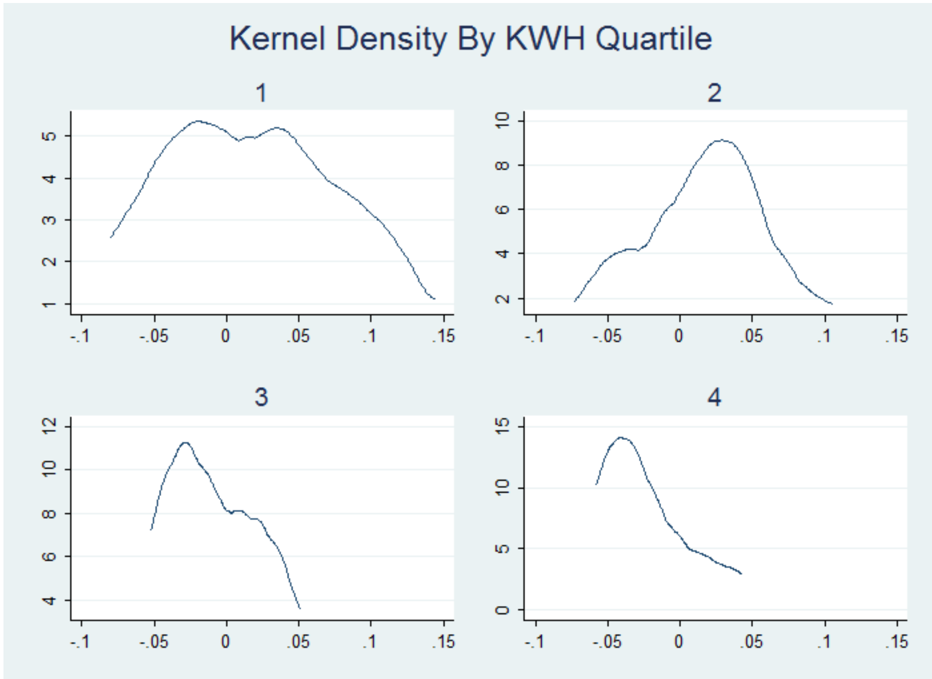

Figure 5: Kernel Density of Behavior-Invariant Bill Changes by Industry

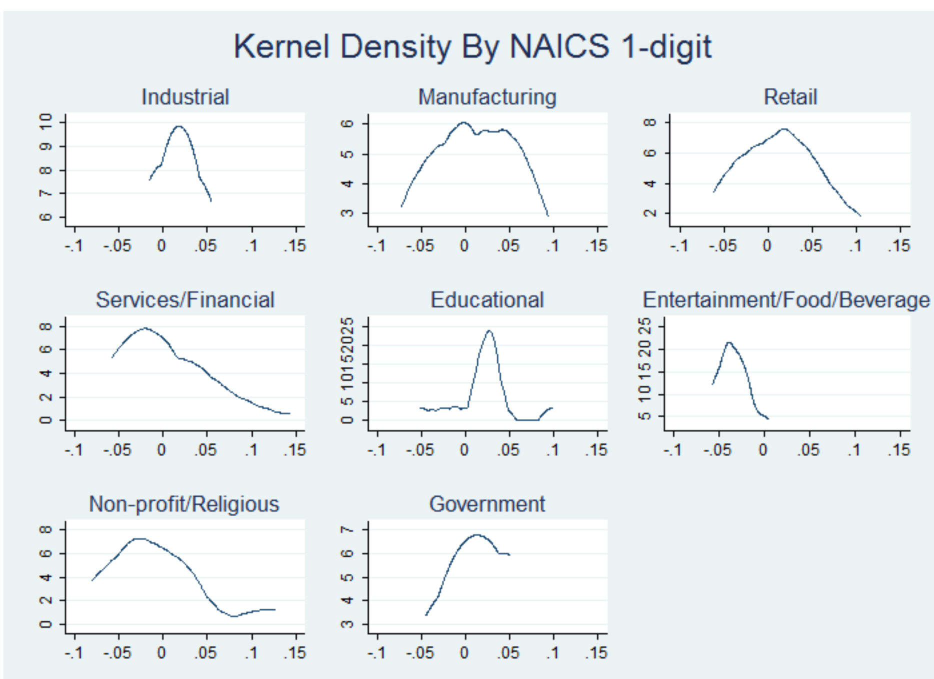

Figure 6: Dynamic Effects in Event Time: Monthly Consumption (kWh)

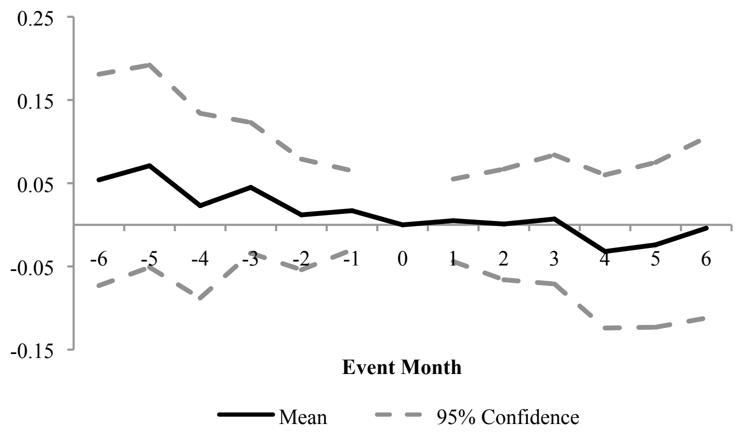

Figure 7: Dynamic Effects in Event Time: Peak Load (kW)

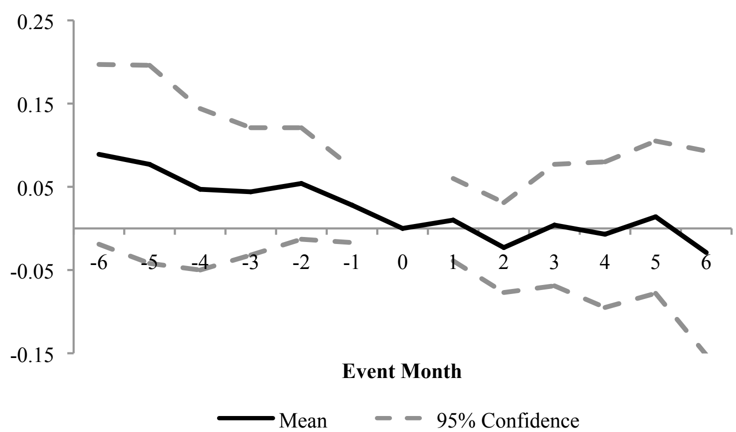

Figure 8: Dynamic Effects in Event Time: Adjusted Bill