Environmental and Social Benefits of Time of Use Electricity Pricing

Matt Harding, Kyle Kettler, Carlos Lamarche

January 15, 2019

Abstract

Previous research on dynamic pricing finds evidence of social benefits through peak-shaving channels; high peak periods reduce the maximum grid load and allow utilities to defer construction of costly peaker plants. However, the literature ignores a second channel through which dynamic pricing influences behavior: reallocation of consumption within a single day from high price periods to low price periods called "load shifting". This paper estimates the hourly marginal emission- and cost-changes that result from changes in the daily profile of household demand using a randomized control trial of household consumption under dynamic pricing as well as system wide observations of power generating units and grid level demand. We find that load shifting reduces hourly emissions of pollutants SO_2, CO_2, and NO_x, as well as reducing production costs for utilities. Averted production costs per household are very small, but tightly estimated.

1 Introduction

The market for electricity has the interesting feature that supply and demand must be balanced in real time to avoid damage to the distribution grid. Substantial quantities of electricity are expensive to store over time. Moreover, while consumers generally pay fixed prices, production costs can vary dramatically over the course of a day. The gap between market price and varying marginal cost leads consumers to over–or under–consume electricity at different parts of the day. This consumption pattern carries significant social costs, as electricity generation is costly in terms of both pollution and resources.

One policy suggestion intended to equalize production costs and market prices is the introduction of dynamic prices for electricity consumers. Efficiency can be improved by discouraging power consumption during periods of peak demand and encouraging greater usage during off-peak periods. Dynamic pricing therefore incentivizes households to reduce the variance of electricity demand from hour to hour. With the adoption of smart meters that allow two-way communication between utility and customers, dynamic pricing for electricity is now technically feasible.

While dynamic pricing is increasingly prevalent among utilities, it is still rare for consumers to face pricing that matches the wholesale market in real time. The volatility of the wholesale market, where prices are computed multiple times in a five minute interval, would transmit a significant amount of risk to consumers in the form of unpredictable electricity bills (Borenstein 2012). Instead, many utilities have used randomized control trials and pilot programs to investigate consumer responsiveness to "time of use" programs. Under time of use (TOU) pricing, consumers face elevated prices during peak hours and reduced prices during off-peak periods.

The desired behavioral response prompted by dynamic pricing may lead to unintended environmental consequences. Peak demand for the summer usually occurs during the afternoon, when solar power production is an important contributor to power generation. Shifting consumption to the evening or early morning hours may entail an increase in electricity generation via dirtier nonrenewable power sources. Reallocating consumption to off-peak hours may therefore offset the environmental benefits of conservation during the peak hours.

This paper uses high frequency household level observations from a randomized control trial for TOU pricing in a south central US state to estimate how households change their consumption under dynamic pricing. Marginal costs of production are collected from the utility that ran the experiment. Additionally, emissions data from the North American Electric Reliability Corporation (NERC) is used to compute the amount of pollution associated with regional electricity production. The production costs for the reallocated power, including both marginal cost and total emissions, is computed for each hour of the day and for three different peak price levels.

We find that the size of the price jump during peak hours is an important determinant for the degree to which households engage in load shifting. Households reduce their consumption during peak hours most on days with the highest price increase. Importantly, peak hours coincide with the lowest marginal rate of emissions but also the highest marginal production costs. This leads to a trade off between reducing operating expenses for the Utility and efforts to minimize pollution.

Section 2 of this paper discusses existing work on pollution externalities and on experimental trials of dynamic pricing. Section 3 discusses the scope and origins of available data on electricity consumption, demand, cost, and related pollution. Section 4 describes the strategy for estimating pollution and cost effects of dynamic pricing at the household level. Section 5 reports the results of analyzing hourly and net daily treatment effects, while section 6 concludes with a discussion of how this paper suggests additional research topics.

2 Literature Review

2.1 Emissions Analysis

A number of modeling strategies have been used to estimate the environmental effects of supplying energy to the grid. One project by Burtraw and E. Mansur (1999) looks at the health effects of trading emissions permits for sulfur dioxide. Another paper by S. Holland and E. Mansur (2008) estimates the effect of demand variability on net emission of pollutants due to energy production for various NERC regions. The inter-day variation in marginal emissions is the focus of Siler-Evans, Azevedo, and Morgan (2012). Each of these papers chooses an approach to address variation in emissions due to grid activity along geographic, temporal, and economic axes. The environmental cost of powering the grid varies due to the types of generating plants used in different parts of the country and different parts of the day.

A paper by J. Graff-Zivin, Kotchen, and Erin Mansur (2014) takes advantage of this variation in generation portfolios to compare the carbon emissions of gas powered cars and electric cars charged from the grid. Graff-Zivin et al. find that the emissions associated with recharging an electric vehicle at one part of the day can be up to double the emissions at another part of the day. Electric cars are an interesting case, economically, because the timing of electricity consumption is largely irrelevant so long as the car is charged before the next desired use. Moreover, electric cars require a large amount of power relative to the amount drawn from the grid by an average household.

More recently, Callaway, Fowlie, and McCormick (2018) estimate the emissions abatement associated with expansion of renewable energy and energy efficiency technologies across the US. Callaway, Fowlie, and McCormick examine CO_2, operating costs avoided, and capacity value generated. Critically, however, this approach estimates the marginal rate of emissions only for fossil fuel generating plants. The paper's approach to identifying the effects of renewable capacity on emissions is a useful technique in identifying how changes in the dispatch order impact marginal emissions. For the purposes of this analysis, however, the relevant parameter is how a marginal change in demand corresponds to emissions for the grid as a whole. Thus, we follow the methodology in J. Graff-Zivin, Kotchen, and Erin Mansur (2014).

Erol-Kantarci and Hussein (2010) propose an intelligent charging algorithm to minimize the strain on the grid from charging plug-in hybrid electric vehicles. They show that using demand side control measures like a time of use rate for electricity can lower carbon dioxide emissions. A look by S. P. Holland et al. (2016) at the externalities associated with electric vehicles notes that the social benefits of electric vehicles are determined by the local geography and ecological situation. The social cost of an additional mile in a gasoline powered car is higher in California than in North Dakota because local pollutants grow more damaging as concentrations rise above a baseline level.

Of course, the economic interest in pollutants centers on understanding their social costs. Cullen (2013) compares the environmental benefits of pollution abatement with the cost of subsidies for wind power. Cullen finds that subsidies for wind power overstate the benefit of offsets in fossil fuel production sought by policymakers. A key insight of the paper is that renewable power often displaces production using dispatchable fuel sources like natural gas instead of base generation like coal. The ongoing transition from coal to natural gas as base load generation is examined by Holladay and LaRiviere (2017). They find that the generation portfolio established before the advent of fracking is a main determinant in the degree to which renewable energy sources provide opportunities for pollution abatement.

2.2 Randomized Control Trials

This paper contributes to a literature about electricity demand estimation based on randomized control trials. The most important precursor to my analysis is the work published by Harding and Lamarche (2016) involving the same experimental program. Harding and Lamarche compute the treatment effects of a dynamic pricing structure with a single peak price called every afternoon. Treated households are given access to the same technologies described in this paper. Harding and Lamarche find that automation (using programmable thermostats) is associated with a large reduction in the quantity of electricity demanded during peak pricing periods. Moreover, automated households display evidence of load shifting, smoothing electricity consumption over the course of the day. The experimental group analyzed by Harding and Lamarche is a different arm of the experiment described in this paper. Load shifting is an important feature of this dataset.

The results in Harding and Lamarche (2016) concur with previous randomized control trials in that consumers are seen to respond to price incentives. Wolak (2006) examines a small sample of electricity consumers in Anaheim, CA during the summer of 2005. Customers in both the treatment and control groups face the normal block-rate price structure. Treated customers have the opportunity during "critical peak pricing" events to earn a rebate of $0.35/KWh for reducing their electricity use compared to non critical peak pricing days. Critical peak pricing for these consumers is associated with a reduction in consumption of about 12 percent. There is not a change in consumption among the treatment group on non critical peak days.

Other work on randomized control trials often features experiments with both a price component and an information component. Cosmo, Lyons, and Nolan (2012) uses data from a 2007 trial of time of use pricing accompanied by various forms of information feedback to consumers. Time of use pricing is found to be associated with significant reductions in peak usage, but there is a limited response to increasing the size of the price differential during peak hours. Different information feedback systems are associated with differential responses to the TOU pricing period. The importance of information provision is highlighted in a study examining how households respond to unpredictable price events Jessoe and Rapson (2014). Treated households are assigned to either of two groups: 'price' and 'price with in-home display'. Price events occur with varying amount of advance notice: one day ahead or thirty minutes ahead. Households with the in-home display are found to be much more responsive to price events.

3 Data

3.1 Grid-Wide Observations

The Environmental Protection Agency's Continuous Emission Monitoring System (CEMS) provides grid-wide emissions data. Emissions from all fossil-fuel generating plants with at least 25 MW of generating capacity are available. The pollutants carbon dioxide, sulfur dioxide, and nitrous oxide (CO_2, SO_2, and NO_x, respectively) are examined in this paper. Data about how much of each pollutant is emitted must be reported to the EPA with granularity at the generating unit level. Observations are reported hourly. In this analysis we use the average reported hourly emissions from all power generating units in the Southwest Power Pool NERC region between 12:01 a.m. July 1, 2011 and 11:59 p.m. September 30, 2011. This provides a grid-wide measure of how much pollution is generated by power generation activities within the pertinent NERC region. Units for pollutants are all reported in pounds.

The Federal Energy Regulatory Commission mandates that balancing authorities and planning regions must report the grid demand for each hour of every day via FERC form 714. These reports make up the second dataset used in this analysis. The sum of hourly demand for the NERC region containing the experiment (described below) is computed for each hour. To simplify the relationship between consumption and emissions, the amount of energy traded between regions is assumed to be zero. This assumption is required because of the variation in generating portfolios across the United States. Identification of grid-wide emissions per marginal MW of electricity demand follows from this simplification.

The utility that implemented the experiment in dynamic pricing has also provided data on the marginal cost of electricity production. This "locational marginal price" (LMP) is a measure of the cost to provide an additional megawatt of energy to the grid. The cost varies between nodes of the grid according to the time of day, distance to the nearest generating plant (e.g. coal vs. wind) and the amount of loss or congestion along the physical transmission infrastructure. Balancing authorities' attempts to minimize this cost results in the dispatch order, under which cheaper generating units are turned on before more expensive units. Locational marginal prices are computed in five-minute increments. This analysis uses an hourly average cost across all nodes within the coverage region operated by the utility running the experiment described below.

In total, there are 1584 observations of weekday hours in 2011. Grid level variables include electricity demand, total emissions, and marginal production cost from July to September in 2011 for the SPP NERC region.

3.2 Household-Level Observations

Household consumption data are taken from a large-scale randomized controlled trial (RCT) for a time of use pricing scheme in the residential market for electricity. Due to a confidentiality agreement, the organization running the trial will be left unnamed and referred to as 'the Utility'. The Utility tested how a new price regime impacted the daily profile of residential electricity consumption. Additionally, consumers facing the new price regime were provided with up to three technologies designed to facilitate acquiring information about and responding to price changes. Household energy use for the months of June through September during the summer of 2011 was monitored with smart meters which report consumption in fifteen minute intervals.

The price regime tested in the RCT was a variable peak price (VPP) cost structure. Households faced a low price of \$0.045/kWh during off peak hours. Peak prices were declared by 5 p.m.

the previous day. The Utility could declare a low ($0.045/kWh), medium ($0.113/kWh), high ($0.23/kWh), or critical ($0.46/kWh) price day. Peak prices occurred between 2 p.m. and 7 p.m. on weekdays excluding Independence Day and Labor Day.

To better inform households about the peak price period, members of the treatment group were provided with automation technology. All treated households had access to a web portal that reported how much electricity was consumed during the previous 15-minute usage interval, as well as the current price of electricity. Households were also provided with an "In-Home Display" (IHD), a tablet device that reported the same information in real time and fit on a counter top. Additionally, treated households were provided with an programmable controlled thermostat (PCT) which adjusted the interior temperature in response to price changes and time of day.

Control households did not receive any type of technology intervention. Control households remained on the standard block tariff, which charges $0.084/kWh for the first 1400 kW used in a month and $0.0968/kWh for usage above that level.

Assignment to treatment groups was based on a Latin Square design. Before sampling households to participate in the RCT, all customers of the Utility were assigned to a treatment group. Households were invited to enroll in the experiment. Based on a pilot study the previous year, demographics predicted to have low participation rates received additional advertising. Assignment to treatment groups was enforced successfully in most cases, although a few households changed treatment status because of technical reasons.

The data include demographic details purchased from Nielsen along two axes: the level of income and the age of the family in each household. Household income is binned into 'low,' 'middle,' and 'high' categories. Low income households have yearly earnings averaging around $30,000. The middle-income group is centered around $50,000. The high-income group has income above $75,000. Family age demographics are broken into 'younger,' 'family,' and 'mature' households. 'Younger' households are those under the age of 45 with no children. The 'Family life' segment of households consists of middle-aged families with children. The 'mature' households consist of older empty nest households, with ages of 65. Nationally, forty percent of households are under 45 and twenty percent are over the age of 65. The sample used in this study skews older than the national average, with

Control Group Households				
	Young	Family	Mature	Total
High Income	42	90	64	196
Middle Income	33	52	33	126
Low Income	43	60	27	122
Total	118	202	124	**444**

Treated Group Households				
	Young	Family	Mature	Total
High Income	21	32	29	82
Middle Income	23	19	24	66
Low Income	25	16	13	54
Total	69	67	66	**202**

Figure 1: Assignment of households to each demographic cell for both treatment and control group.

middle aged and mature households relatively over-represented. The demographic breakdowns for the control group and the treated households are given in figure 1.

For all households, smart meters report 15-minute interval measurements of electricity use to the Utility during the months of June-September. To match the frequency of grid level emissions data and locational marginal pricing, we aggregate each household's total energy consumption over each hour. The analysis restricts the sample to the months of July, August, and September to mitigate confounding learning effects at the start of the program.

To highlight the effects of time of use pricing regimes, consider the difference in daily usage profiles between the control group and the treated group. The technology allows households to automate their response to changes in price. Because cooling is such a large portion of daily consumption, automation highlights the degree to which consumers are willing to respond to fluctuations in electricity prices. Figure 2 graphs the average hourly electricity consumption of control and treated households. The panels correspond to the particular price called each day. For example, VPP2 indicates that a medium price day was called, where the peak price was \$0.113/kWh. VPP3 indicates a high price day, and VPP4 indicates a critical price day. The difference in consumption between treated households and control households is broken down by age and income in Section 6.

Comparing average consumption across treated households, there is visual evidence of load shifting. Households facing time of use pricing consume less electricity during peak periods, but consume relatively more during the following hours. On the other hand, there is less evidence of strategic pre-cooling by the treated households. This pattern is consistent with analysis of high frequency consumption data in Harding and Lamarche (2016), which examines consumption of

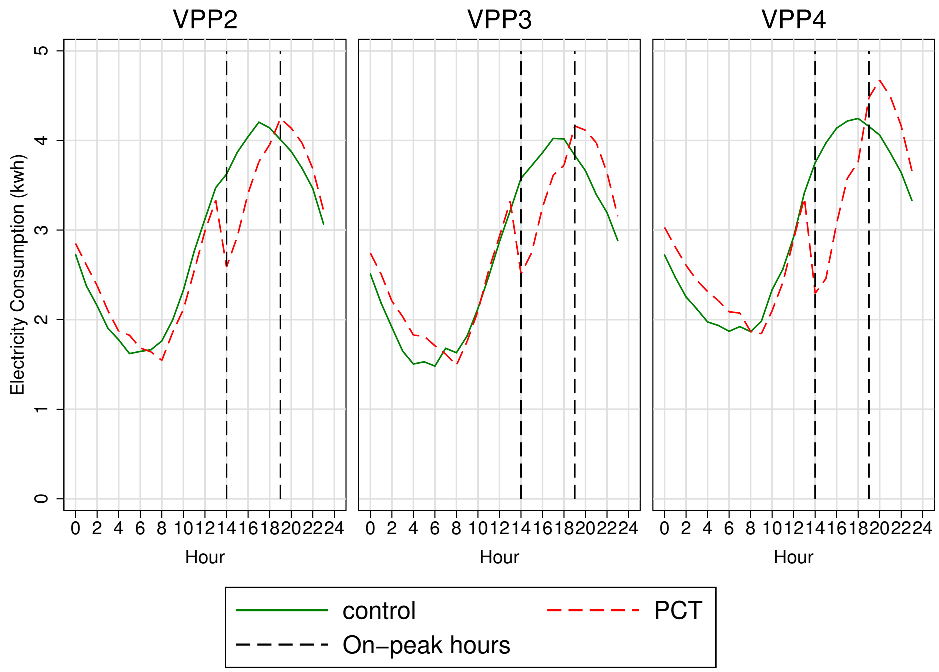

Figure 2: Average Hourly Consumption: Control and PCT Households

customers of the Utility in a parallel experiment. Households do not appear to take advantage of low prices during the morning to proactively reduce temperatures before the high peak price in the afternoon.

4 Methods

4.1 System-Wide Marginal Emissions

As discussed, time varying electricity pricing may smooth demand by shifting consumption from peak to off-peak hours. However, the environmental effects of this shift are ambiguous because different generating plants are activated over the course of the day. This analysis follows the methodology for estimating marginal emissions of electricity demand outlined in Graff-Zivin, Kotchen, and Mansur (2012). Hourly emissions data at generating units within the NERC area are aggregated each hour to provide a measure of pollution in the region. Similarly, demand is aggregated from balancing authorities within the NERC area and used as a measure of grid-load during each hour. Regional hourly emissions are regressed on regional hourly demand, with fixed effects allowing the base load to vary according to hour of the day and month of the year:

$$C_{h,m} = \alpha_{m,h} + \gamma_h * q_h + \varepsilon_h \tag{1}$$

The regression is conducted for each of the pollutants CO_2, NO_x, and SO_2 used as the dependent variable C_h. The quantity of electricity demanded in the NERC region during hour h is given by q_h. The term $\alpha_{m,h}$ is a month-by-hour fixed effect that captures changing electricity use cycles during different parts of the summer. The parameters of interest for each pollutant are given by the sequence $\gamma_0, \ldots, \gamma_{23}$ and represent the marginal emissions (in pounds) attributable to an increase of 1 kWh in grid demand at each hour of the day.

The quantity demanded and the resulting emissions are both likely to depend on price, but the Utility does not charge wholesale electricity prices to consumers. While the consumers in the smart meter study are subject to time of use pricing, the vast majority of residential customers in the NERC region are subject to a flat or block rate. Following the treatment in Graff-Zivin, Joshua S., and E. T. Mansur (2014), we assume that the grid level q_h are exogenous both to the amount of pollution emitted and to the quantity of electricity consumed by any given household in the experiment.

4.2 System-Wide Marginal Costs

One way that regulators justify dynamic prices to skeptics is the possibility of mitigating production costs due to variance in demand throughout the day. we compute both the average and standard deviation of 'locational marginal price' reported at all nodes controlled by the Utility for each hour of the day. At hour h this average is denoted as δ_h. we use only prices reported at nodes controlled by the Utility to allow for cost-benefit analysis across years.

4.3 Estimation of Mean and Quantile Household Treatment Effects

We identify the treatment effect of a programmable control thermostat paired with variable peak pricing by relying on the random assignment of households to the treatment and control groups. We pursue a modeling approach first suggested by Ramanathan et al. (1997). More recently, Harding and Lamarche (2016) compute treatment effects for another dynamic pricing model that the Utility experimented with, where the peak price is constant across days.

We model the load function as piecewise constant over each hour-long interval of consumption. Consumption for household i on day t with peak price k during hour h is used as the dependent variable to estimate the load function faced by the Utility. As is common in the electricity literature, the log of household consumption $\log(Y_{i,t,h}^k)$ is used to mitigate the effects of large positive outliers in household consumption.

Two independent variables are included in the model. An indicator variable $D_i(j)$ denotes the treatment status of each household i. The variable takes a value of 1 for treated households and 0 for control households. Additionally, a Generalized Additive Model as described by Hasie and Tibshirani (1990) is implemented to account for a nonlinear relationship between weather and power consumption. The function $f(W_{i,t,h})$ is a linear combination of the temperature and dew point approximated by 5th degree B-splines. A constant for every hour, α_h, captures the average usage among control households independent of the variable effect of weather. This model controls for the cycle of electricity demand over the course of the day.

$$\log(Y_{i,t,h}^k) = \alpha_h + \beta_h^k D_i(j) + f(W_{i,t,h}) + \varepsilon_{i,t,h},\ h \in \{0, \ldots, 23\} \tag{2}$$

The parameters of interest are the β terms for each hour of the day for each of the price types. The sequence of $\beta_0^k, \ldots, \beta_{23}^k$ provides an estimate of the treatment effect over the course of the day for each price level k. Identification of the treatment effect rests on the random assignment of households to the treatment group. The peak period between $h = 14$ and $h = 19$ is when prices are higher for the treatment group. This treatment effect is the percentage change in electricity consumption among treated households relative to the average usage among control households.

Errors are clustered by household to account for serial correlation.

Price responsiveness is an important dimension of heterogeneity in the size of the treatment effect. Consumers may only change their behavior on days when the higher price is called. To address this concern, the sequence of hourly treatment effects is computed for each elevated price level called by the Utility. Consumers face different marginal prices during peak hours on each day, which affects their incentive to shift or reduce electricity consumption at every hour of that day. There is a stronger incentive to engage in load-shifting or pre-cooling on a critical price day than on a medium price day. Thus, we calculate a sequence of treatment effects for each price level. This estimation is accomplished by running the regression in 2 on the subsample of calendar days with price level $k \in \{\text{medium, high, critical}\}$. Households have multiple observations of consumption at hour h for each price level k because there are multiple days of each price type observed throughout the summer.

In addition to the pooled OLS regression for each price level, we compute a quantile regression following Koenker (2005). we model the τ-th quantile of the conditional distribution of $\log(Y_{i,t,h})$ during each hour h on days with price level k.

$$Q_{log(Y^k_{i,t,h})}(\tau|D_i(j), W_{i,t,h}) = \alpha_h I(h) + \beta^k_h(\tau) D_i(j) + f(W_{i,t,h}) + \varepsilon_{i,t,h},\ h \in \{0, \ldots, 23\} \qquad (3)$$

The parameters of interest are the sequences $\beta^k_h(\tau)$ for days at each peak price level k (medium, high, or critical price days). As before, these are estimates of the change in behavior at hour h among treated households relative to control households. The objective function is weighted to estimate the treatment effect at the τ-th percentile of the consumption distribution. For example, $\widehat{\beta^k_{12}(10)}$ is the estimate of how much treatment changes the electricity use of households at the 10th percentile of consumption in hour 12 on a day with price level k.

4.4 Projecting Hourly Changes Under Dynamic Prices

we estimate the change in ecological and economic costs associated with the effect of dynamic pricing on household electricity demand. First we compute how many kilowatts are conserved by implementing the dynamic price tariff. Then we use the marginal emissions rate from 1 to compute how many pounds of pollutants would be associated with the change in energy consumption. Adding up the hourly point estimates provides the net change in household emissions for a single day's electricity consumption. This process is repeated for each peak price level that can be called.

The change in the level of power consumption is computed relative to the average consumption in hour h among control households on days with price level k. Denote this as $\bar{y_h}^k$. The treatment effect β_h^k is the percent change in consumption due to a change in peak price of level k. The level change in power consumption is therefore given by:

$$\Delta Y_h^k = \beta_h^k * \bar{y_h}^k, \tag{4}$$

$$\Delta Y_h^k(\tau) = \beta_h^k(\tau) * \bar{y_h}^k \tag{5}$$

This level change gives the amount of power in hour h affected by implementing a dynamic price of k during the peak hours. As before, the value of ΔY_h is allowed to vary across each of the peak price levels. The marginal rate of emissions for CO_2, SO_2, and NO_x are given by γ_h. Marginal emissions are not computed separately for each price level because grid-wide dispatch decisions are not determined at the same level as the Utility's peak price decision. Thus the change in household emissions for price day k is given by:

$$\Delta C_h^k = \Delta Y_h^k * \gamma_h \tag{6}$$

$$\Delta C_h^k(\tau) = \Delta Y_h^k(\tau) * \gamma_h \tag{7}$$

We apply a similar approach to estimate the change in hourly production costs. We multiply the average 'locational marginal price' in hour h, δ_h, by the hourly level change in consumption, ΔY_h. This product, ΔB_h is the change in the Utility's production costs needed to supply the household's demanded power to the grid in hour h.

$$\Delta B_h = \Delta Y_h^k * \delta_h \tag{8}$$

$$\Delta B_h(\tau) = \Delta Y_h^k(\tau) * \delta_h \tag{9}$$

Finally, the policy analysis is chiefly interested in how a dynamic pricing program influences the emissions and costs associated with meeting household electricity demand. For each price level k, We compute the net effect on emissions and costs from a day of behavior under the new price regime:

$$\sum_{h=0}^{23} \Delta C_h^k = \sum_{h=0}^{23} \Delta Y_h^k * \gamma_h \tag{10}$$

$$\sum_{h=0}^{23} \Delta C_h^k(\tau) = \sum_{h=0}^{23} \Delta Y_h^k(\tau) * \gamma_h \tag{11}$$

$$\sum_{h=0}^{23} \Delta B_h^k = \sum_{h=0}^{23} \Delta Y_h^k * \delta_h \tag{12}$$

$$\sum_{h=0}^{23} \Delta B_h^k(\tau) = \sum_{h=0}^{23} \Delta Y_h^k(\tau) * \delta_h \tag{13}$$

This summation offers a point estimate of the daily average household effect of dynamic pricing. The sign of $\sum \Delta C$ indicates how the household's behavior influenced costs overall. To provide context for these estimates, we use the EPA greenhouse gas equivalency calculator[1] to find the effect of the policy in a scenario where ten million residential households adopted the dynamic price from the Utility's experimental program. (There are roughly ten million households in Texas, which has a similar climate to the metropolitan area served by the Utility. The EPA equivalency calculator provides a back of the envelope maximum for the effect of dynamic pricing.)

[1]https://www.epa.gov/energy/greenhouse-gas-equivalencies-calculator

Confidence intervals for marginal emissions effects γ_h are correctly estimated using OLS. However, the standard errors for ΔY_h^k cannot be taken directly from a regression procedure because of a correlation between the treatment effect β and the average level of consumption among control houses $\bar{y_h}$. We implement a block bootstrap procedure to address this problem. The block bootstrap samples with replacement from households in both the control group and the treated group. (Draws equal the number of households in each group in the observed experiment: 202 treated households and 444 control households.) It is important that each draw preserves selected households' entire consumption path over the summer in order to account for learning effects, varying consumer engagement with the experiment, and idiosyncratic inter-day correlations from each household's daily schedule.

In each iteration ΔY_h^k is computed and stored to construct the sampling distribution. The block bootstrap provides an estimate of the variance of ΔY_h^k. We compute the hourly level change in kWh among treated households for 500 iterations of the block bootstrap as the sampling distribution. The standard deviation of the sampling distribution is the standard error for the ΔY_h^k. Confidence bands for this parameter are constructed as ± 1.96 standard errors around the point estimate constructed using the original (observed) data.

We compute the variance of ΔC_h^k and ΔB_h^k as the product of two random variables: ΔY_h^k and γ_h. The variance of γ_h is the square of the standard error taken from the OLS regression in 1, and the variance of δ_h is the sample variance of reported node prices in hour h.

$$\operatorname{var}(\Delta C_h^k) = \operatorname{var}(\Delta Y_h^k)\operatorname{var}(\gamma_h) + \operatorname{var}(\Delta Y_h^k)E(\gamma_h)^2 + E(\Delta Y_h^k)^2\operatorname{var}(\gamma_h) \tag{14}$$

$$\operatorname{var}(\Delta B_h^k) = \operatorname{var}(\Delta Y_h^k)\operatorname{var}(\delta_h) + \operatorname{var}(\Delta Y_h^k)E(\delta_h)^2 + E(\Delta Y_h^k)^2\operatorname{var}(\delta_h) \tag{15}$$

The analytic solutions for the estimates for the net daily effects $\sum \Delta C_h^k$ and $\sum \Delta B_h^k$ are slightly more complex, since the level change in electricity consumption in hour a is unlikely to be independent of the level change in hour b (with $a, b \in \{0, \ldots, 23\}$). Thus an analytic solution for the variance of these terms involves the covariance between the change in emissions or cost in different hours.

$$
\begin{aligned}
\mathrm{var}(\sum \Delta C_h^k) &= \mathrm{var}(\Delta Y_1^k \gamma_1 + \ldots + \Delta Y_{23}\gamma_{23}) \\
&= \sum_{a=0}^{23} \mathrm{var}(\Delta Y_a^k \gamma_a) + \sum_{b=0}^{23}\sum_{c=0}^{23} cov(\Delta Y_b^k \gamma_b, \Delta Y_c^k \gamma_c)
\end{aligned}
$$

We bootstrap the variance of $\sum \Delta C_h^k$ and $\sum \Delta B_h^k$ rather than compute all 276 possible pairwise hourly covariances. The term $\sum_h \Delta C_h^k$ is computed in each bootstrap iteration as the sum of the point estimates of hourly effects in that iteration. This results in a sampling distribution of 500 draws at each price level. We compute the standard deviation of this sampling distribution and use it as the standard error for the sum of hourly emission effects estimated using the observed data set. This approach relies on a simplifying assumption that the estimated marginal household emissions and the level change in consumption among treated households are independed random variables. Such an assumption is plausible because the Utility represents a small fraction of the demand within the NERC region. Thus, treatment effects from the dynamic price are unlikely to affect the dispatch order that determines which power plants are active in any given hour.

5 Results

Treated households' electricity consumption is reported as a change (in kilowatts) relative to the average consumption among control group households in each hour of the day. The reallocation of energy consumption throughout the day is illustrated for days with a medium, high, and critical peak price level. Figure 3 reports estimates of the treatment effect for treated households at the mean of the consumption distribution as well as the 10th and 90th quantiles of electricity consumption.

A consistent pattern with respect to the peak price level appears for each regression. The first row of figure 3 indicates the average change in behavior for treated households. Days with a medium peak price ($0.113/kwh) have a modest and statistically significant reduction in consumption during peak hours. However, there is not statistically distinguishable evidence of load shifting (even to the

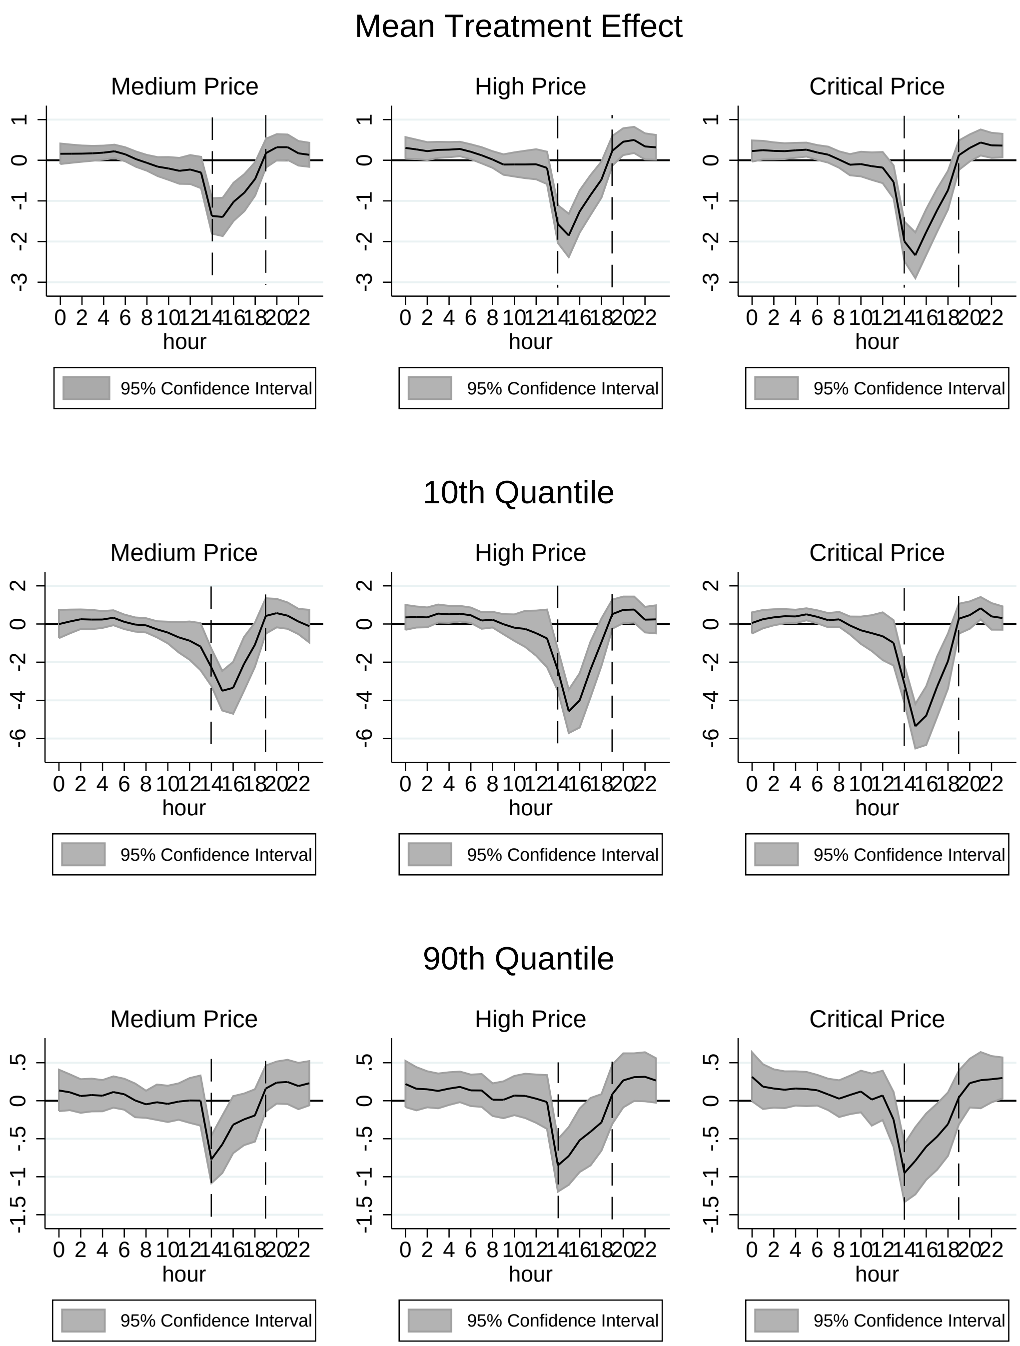

Figure 3: Level change in electricity consumption among treated households for each hour of the day, measured in kw. 95% confidence intervals use standard errors taken from block bootstrap with 500 iterations.

hour when prices fall back to off-peak levels). However, on days with high ($0.23/kwh) or critical ($0.46/kwh) peak prices there is at least one hour with statistically significant load shifting. The amount of conservation during peak hours increases as the price rises. Days with the highest prices also have the largest reduction in energy use among treated households.

Households at the tenth quantile of the usage distribution exhibit the most extreme response to changes in the peak price. The estimated reduction among these households is roughly two times larger than the average estimate for each price level reported in the first row. The quantile estimates are much noisier than the average estimates but there is again (weak) evidence of load shifting found in the quantile regression for the tenth percentile of consumers.

In contrast to households that use less electricity, households in the ninetieth percentile of consumption are much less responsive to elevated peak prices. Moreover, the relationship between the jump in prices during peak hours and hourly conservation efforts is much weaker for these households. The estimated conservation effect for households at the ninetieth percentile does not appreciably change in response to higher peak prices. Notably, these effects (while statistically significant) are small relative to the standard error of the estimates even during peak hours.

The treatment effect (at the average as well as the tails) is characterized by a large reduction in consumption at the very beginning of the peak period that tapers off over the afternoon. Where there is evidence for load shifting it occurs at the very earliest part of the morning (5 a.m.) or in the hours immediately following the peak hours (8 - 11 p.m.). These responses are both compatible with a household that takes advantage of the programmable control thermostat to reallocate cooling hours away from the period of elevated prices. Automation technology reduces consumption at the beginning of the peak period by raising the allowable interior temperature. As homes warm up, thermostats turn on to maintain the new (higher) temperature. Thermostats then reset to the default temperature when the peak period ends, prompting an increase in energy consumption to return the household to the initial lower temperature.

This paper uses observations of power-plant emissions to estimate the environmental consequences of this type of load shifting. We regress all emissions reported by power plants that identify themselves as operating within the region managed by the SPP NERC region. The EPA's

Continuous Emissions Monitoring System measures the total emissions of sulfur dioxide, carbon dioxide, and nitrous oxide in each hour. The independent variable is the hourly demand for electricity from the grid in the SPP NERC region. Figure 4 reports the twenty-four hourly estimates of how much pollution (in pounds) is emitted in response to a 1 kw increase in electricity demand.

Additionally, Figure 1 reports the hourly average locational marginal price to provide 1 kw of electricity to customers of the Utility. The confidence bands for cost are taken from the standard deviation of observed prices at each hour. This provides a measure of the variability of cost within the Utility's service area at each hour of the day.

The daily pattern of marginal emissions for SO_2, CO_2, and NO_x is unsurprising and reflects the efficiency of base load generation relative to generating plants that can be dispatched. During the early morning and late night demand across the SPP region is low, so only coal power plants are directed to produce electricity. During peak hours, relatively cleaner gas power plants are dispatched to respond to fluctuations in grid demand so the marginal emissions are lower. Moreover, solar power is available during the day and is has zero marginal emissions.

This pattern, then, suggests that conservation behavior by residential consumers produces the smallest change in emissions during the afternoon and the largest effects during the night. On the intensive margin, where no power plant is dispatched to address shifts of the daily profile of demand, we can compute the trade off between conservation in the afternoon and extra use in other time periods as a net change in emissions of greenhouse gasses.

The most interesting feature of Figure 1 is that the peak period has the lowest marginal emissions rates, but the highest marginal costs. This presents a dilemma for the Utility, in that the treatment effect's pattern of load shifting reduces total production costs but may be accompanied by a higher amount of pollution. On the other hand, the variance of reported node prices is highest during the peak hours chosen by the Utility. In the wholesale price of electricity is observed to fall below zero when producers are willing to pay customers to accept the output of their generating plant rather than incur ramping costs when they want to restore output later.

The product of the average treatment effect and the marginal emission rate generates the hourly change in emissions associated with treatment with a dynamic price and programmable control

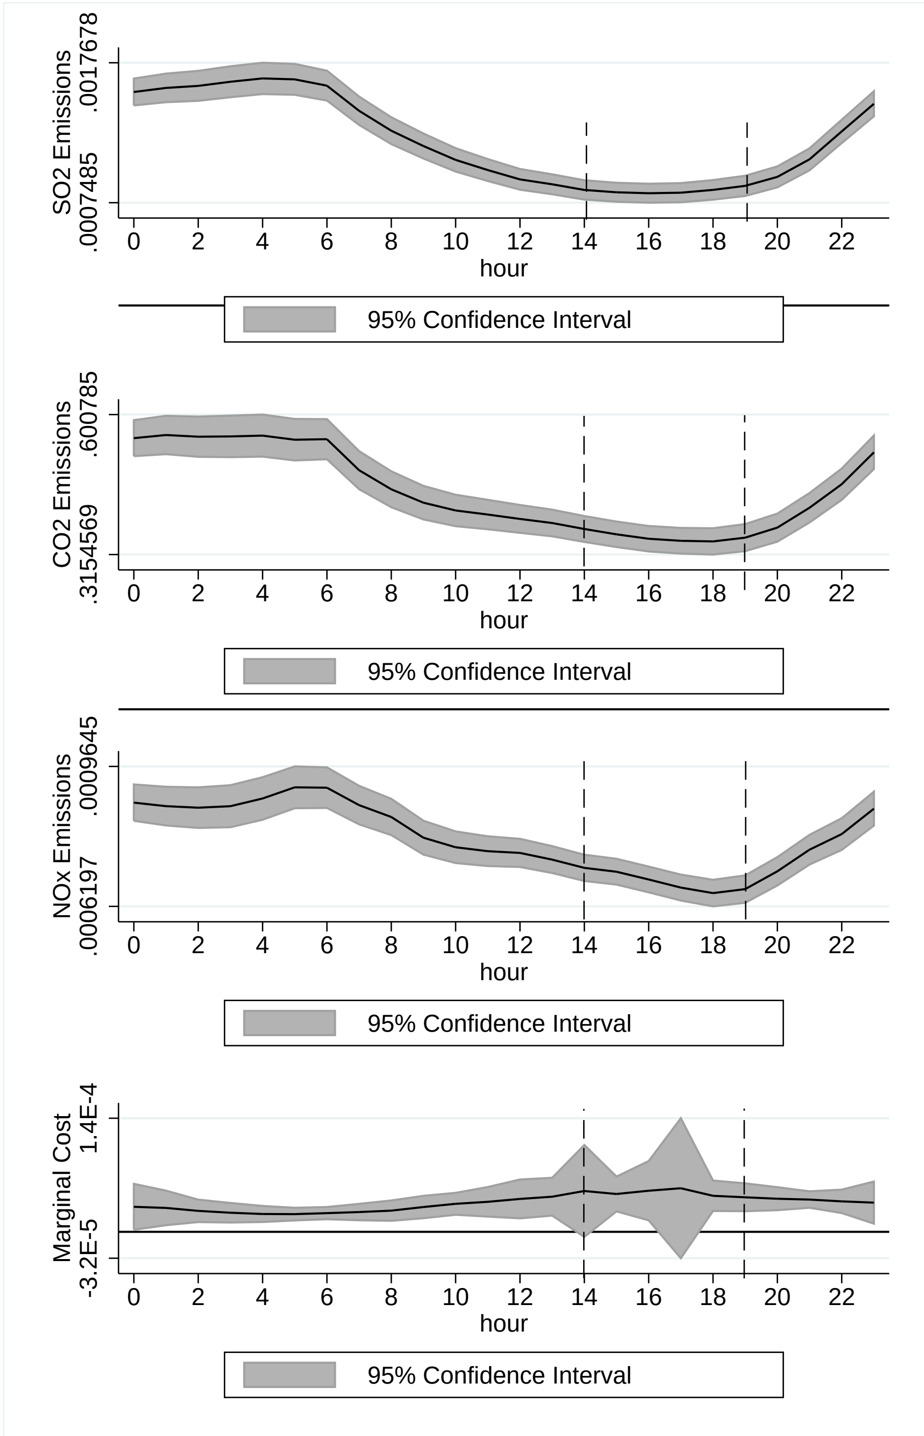

Figure 4: Marginal emissions associated with electricity consumption in the SPP NERC region during summer 2011. Marginal cost to provide 1 kw of electricity averaged across all nodes in the Utility's service area.

thermostat. We report the average treatment effects for each pollutant as well as the change in production costs in figure 5. The effects are computed for each outcome at the three different price levels.

The first row of Figure 5 reports the change in sulfur emissions for each hour under medium, high, and critical peak prices. Driven mostly by the level changes computed earlier, the size of the reduction is largest on critical price days. The confluence of high marginal emissions and marginally significant increases in consumption during the early morning hours results in a statistically significant increase in sulfur emitted for off-peak hours on days with high or critical prices. The same pattern is borne out in the estimates for carbon dioxide and nitrous oxides. The driving force is the size of the level change in consumption, with variation in the marginal rate of emissions largely overpowered by the shape of households' behavior under treatment.

The change in production costs for the utility, reported in row 4 of figure 5, reflect the very high savings in operating costs associated with peak shaving due to dynamic prices. The Utility reduces the cost to supply a treated household during the peak hours, and cost increases during off-peak hours are, individually, statistically indistinguishable from zero. There is some evidence of a small increase in costs during evening hours on high and critical peak price days. To determine the net daily effect, we sum the point estimate in each hour. Standard errors for these sums are taken from the daily sum in each of the 500 iterations of the block bootstrap. Because the daily estimate includes uncertainty from 24 different point estimates, the errors for these aggregates are relatively wide and deliberately conservative.

We find that the change in cost is statistically significant at the 5% level on medium and critical price days, but other estimates about net effects are too noisily estimated under the block bootstrap to be confident that the observed treatment effects here are statistically distinguishable from zero. The changes in behavior under treatment leads to a reduction in production costs of between two and three ten-thousandths of a cent per household per day. These effects are very small: There are 10,000,000 households in Texas. If the entire state was was treated with a dynamic price and programmable control thermostat, the production costs associated with supplying power for a single day would fall by only $2,100.

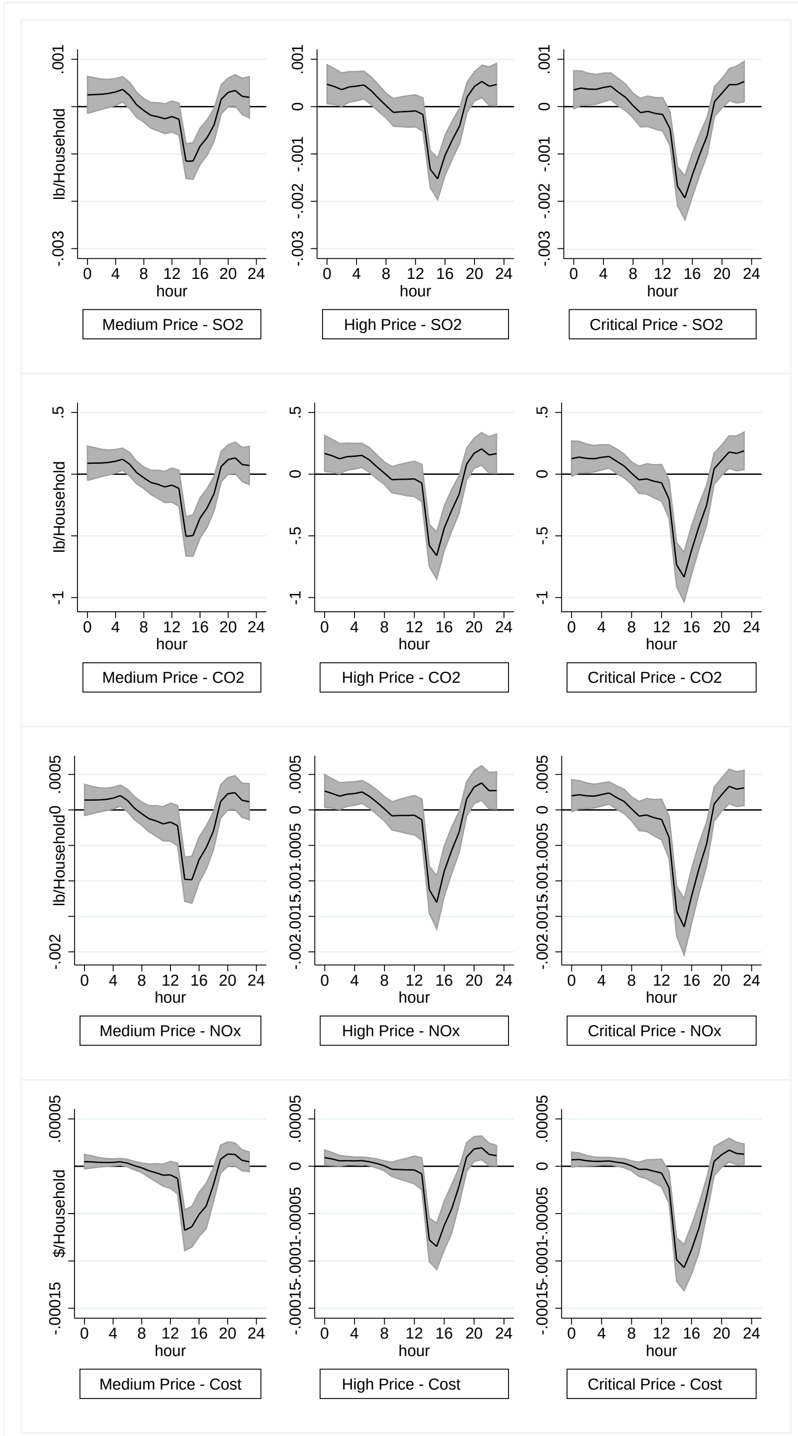

Figure 5: Level changes in pollutants under treatment with dynamic price. 95% confidence intervals use standard errors taken from block bootstrap with 500 iterations.

	Medium Price	High Price	Critical Price
SO_2/Household	-0.00219	-0.00042	-0.00300
	(.00350)	(.00369)	(.00365)
CO_2/Household	-1.12526	-0.52481	-1.60004
	(1.34081)	(1.41111)	(1.39857)
NO_x/Household	-0.00248	-0.00151	-0.00358
	(0.00243)	(0.00255)	(0.00254)
Cost/Household	-0.00021*	-0.00019	-0.000323***
	(0.00012)	(0.00013)	(0.00013)
Carbon Equivalent: # of Cars	-1,812	-948	-2,590
Days in Summer	30	22	12
Summer Total		-106,296	

Figure 6: Sum of point estimates, standard errors from block bootstrap. Cars taken off the road calculated with the EPA greenhouse gas equivalency calculator assuming treatment of 10,000,000 households.

We use the EPA greenhouse gas equivalency calculator to compute the effect of universal adoption on polluting emissions. We provide a benchmark for the effect on total emissions by assuming that 10,000,000 households have been treated, and counting the number of days that the Utility called each price level. There are 30 medium peak price days, 22 high peak price days, and 12 critical peak price days. Under universal adoption, the point estimates for daily net emissions imply savings of greenhouse gas emissions equivalent to 106,296 cars taken off the road for an entire year.

6 Conclusion

This paper contributes to the literature in two primary ways. First, we implement the approach used by J. Graff-Zivin, Kotchen, and Erin Mansur 2014 in the context of dynamic pricing, providing the first look at how dynamic pricing influences both social and private operating costs. We find a small but detectable reduction in production costs to the Utility from load-shifting patterns under a variable peak dynamic price. Though net daily emissions effects are noisily estimated, we provide a trajectory of hourly emissions effects that illustrate the trade off between cost minimization and pollution externalities.

This emphasis on operating costs is made possible by combining the results of a randomized control trial with federally reported operating data from power generating units and the SPP Independent System Operator. The paper's second contribution is an extension of the randomized control trial literature, especially with regards to electricity pricing policy. The combination of individual-specific data with institutional constraints such as the dispatch curve offers a promising approach to tabulating the magnitude of externalities using micro-level data.

Future research can build upon this analysis by estimating the extensive margin of the dispatch curve; the determinant of the marginal rate of emissions is what kind of power plant is active on the margin. If many households smooth their demand over the course of the day, different power plants may be needed in any given hour. Establishing the magnitude of the treatment effect on the extensive margin would cement the tentative conclusion of this paper: estimates of the costs and benefits (either private or social) are not meaningfully changed by including operating expenses resulting from dynamic pricing.

References

Agency, U.S. Environmental Protection (2011). Clean Air Markets - Emissions Monitoring. https://www.epa.gov/airmarkets/clean-air-markets-emissions-monitoring.

Borenstein, Severin (2012). "Effective and Equitable Adoption of Opt-In Residential Dynamic Electricity Pricing". In: DOI: 10.1007/s11151-012-9367-3.

Burtraw, D. and E. Mansur (1999). "Environmental Effects of SO2 Trading and Banking". In: Environmental Science and Technology 33 (20), "3489–3494". DOI: 10.1021/es9902726.

Callaway, Duncan, Meredith Fowlie, and Gavin McCormick (2018). "EnvironmentalLocation, Location, Location: The Variable Value of Renewable Energy and Demand-side Efficiency Resources". In: Journal of the Association of Environmental and Resource Economists 5 (1), "39–75".

Cosmo, Valeria, Sean Lyons, and Anne Nolan (2012). "Estimating the impact of time-of-use pricing on irish electricity demand". In: Energy Journal 35.2, pp. 117–136. DOI: 10.5547/01956574.35.2.6. URL: http://mpra.ub.uni-muenchen.de/39971/.

Cullen, Joseph (2013). "Measuring the Environmental Benefits of Wind-Generated Electricity". In: American Economic Journal: Economic Policy 5 (4), "107–133". ISSN: 1945-7731. DOI: 10.1257/pol.5.4.107.

Erol-Kantarci, M. and T. M. Hussein (2010). "Prediction-based charging of PHEVs from the smart grid with dynamic pricing". In: IEEE Local Computer Network Conference, pp. 1032–1039. DOI: 10.1109/LCN.2010.5735676.

FERC, Federal Energy Regulatory Commission (2011). Form No. 714 - Annual Electric Balancing Authority Area and Planning Area Report. https://www.ferc.gov/docs-filing/forms/form-714/data.asp.

Graff-Zivin, Joshua, Matthew Kotchen, and Erin Mansur (2014). "Spatial and Temporal Heterogeneity of Marginal Emissions". In: Journal of Economic Behavior and Organization 107, pp. 248–268. DOI: 10.1016/J.JEBO.2014.03.010.

Graff-Zivin, Matthew J. Joshua S. Kotchen, and Erin T. Mansur (2014). "Spatial and Temporal Heterogeneity of Marginal Emissions: Implications for Electric Cars and Other Electricity-Shifting Policies". In: Journal of Economic Behavior and Organization 107, "248–268". ISSN: 0167-2681. DOI: 10.1016/J.JEBO.2014.03.010. URL: www.sciencedirect.com/science/article/pii/S0167268114000808.

Harding, Matthew and Carlos Lamarche (2016). "Empowering Consumers Through Data and Smart Technology: Experimental Evidence on the Consequences of Time-of-Use Electricity Pricing Policies". In: Journal of Policy Analysis and Management 35.4, pp. 906–931. ISSN: 02768739. DOI: 10.1002/pam.21928. URL: http://doi.wiley.com/10.1002/pam.21928.

Holladay, J. Scott and Jacob LaRiviere (2017). "The impact of cheap natural gas on marginal emissions from electricity generation and implications for energy policy". In: Journal of Environmental Economics and Management 85, pp. 205–227. ISSN: 0095-0696. DOI: https://doi.org/10.1016/j.jeem.2017.06.004. URL: http://www.sciencedirect.com/science/article/pii/S0095069617303893.

Holland, S. and E. Mansur (2008). "Is Real-Time Pricing Green? The Environmental Impacts of Electricity Demand Variance". In: Review of Economics and Statistics 90 (3), "550–561". ISSN: 0034-6535. DOI: 10.1162/rest.90.3.550.

Holland, Stephen P. et al. (2016). "Are There Environmental Benefits from Driving Electric Vehicles? The Importance of Local Factors". In: American Economic Review 106.12, pp. 3700–3729.

DOI: 10.1257/aer.20150897. URL: http://www.aeaweb.org/articles?id=10.1257/aer.20150897.
Jessoe, Katrina and David Rapson (2014). “Knowledge is (Less) Power: Experimental Evidence from Residential Energy Use”. In: American Economic Review 104.4, pp. 1417–1438. DOI: 10.1257/aer.104.4.1417. URL: http://dx.doi.org/10.1257/aer.104.4.1417.
Ramanathan, Ramu et al. (1997). “Short-Run Forecasts of Electricity Loads and Peaks”. In: International Journal of Forecasting 13 (2), "161–174". DOI: 10.1016/S0169-2070(97)00015-0.
Siler-Evans, Kyle, Ines Lima Azevedo, and M. Granger Morgan (2012). “"Marginal Emiissions Factors for the U.S. Electricity System”. In: Environmental Science and Technology 46 (9), "4742–4748. ISSN: 0013-936X. DOI: 10.1021/es300145v.
Wolak, Frank A (2006). “Residential Customer Response to Real-Time Pricing: The Anaheim Critical-Peak Pricing Experiment*”. In: URL: http://www.stanford.edu~wolak.