The Effects of Residential Critical Peak Pricing as a Function of Weather and Customer Characteristics: Evidence from California's Statewide Pricing Pilot

Rob Letzler

November 29, 2009

Abstract

California's Statewide Pricing Pilot (SPP) explored the impact of switching residential customers from conventional, time-invariant electricity rates to Critical Peak Pricing (CPP). CPP lowers prices most of the time, raises prices modestly during weekday afternoon peak hours and raises prices dramatically during rare "critical" periods. CPP rates better reflect significant hour-to-hour variations in the cost of generating power. California SPP customers were socioeconomically diverse and lived in diverse climate zones. This paper takes a flexible, difference-in-difference approach to estimating the impacts of the statewide pricing pilot and provides evidence about who is likely to response the most to CPP when. It finds that dynamic pricing led to larger consumption reductions on hotter days and for larger customers. It estimates that the benefits of dynamic pricing range from zero in cooler climates on cooler days to .3 (.4) kW every hour for increased afternoon ("critical peak") prices on the hottest days in hot climates. A program designed to address extreme electrical demand on hot summer days worked best in regions where most customers had air conditioning on days when temperatures above about 90^o prompted them to run their air conditioners. Thus, targeting marketing efforts at high-consumption customers in hot regions has the potential to increase the program's cost-effectiveness.

The views expressed in this paper are those of the author and do not necessarily represent the views of the Federal Trade Commission or any individual Commissioner. I would like to thank my advisers Lee Friedman, Severin Borenstein, Rob MacCoun, John Morgan, and Geno Smolensky as well as Max Aufhammer, Rich Barnes, Meghan Busse, Karen Herter, John Jurewitz, Karen Notsund, Lauren Pemberton, and Xing Zhong for comments, advice, suggestions, documents, data, and support. I would especially like to thank Pat McAuliffe for facilitating the process of making the SPP data set available for academic research; and to thank the CEC, CPUC, PG&E, SCE, and SDG&E for agreeing to share the data. I thank Susan McNicholl for alerting me to the selection problem discussed below.

Contents

1 Background: Critical Peak Pricing and the California Statewide Pricing Pilot

This paper assesses the impact of a residential Critical Peak Pricing (CPP) program. It focuses on understanding when and for whom price increases reduce consumption.

The cost of generating power fluctuates enormously from hour to hour, but most customers pay time-invariant prices for power. The wholesale cost of power is well under 10 cents per kilowatt hour (kWh) during most hours but $1.00 / kWh price caps can bind during scarcity hours and the price cap is likely below the power's social cost. Electricity systems have to balance the amount supplied and consumed on a minute by minute basis to avoid blackouts. Air conditioners consume a great deal of electricity which leads electricity demand to peak on hot weekday afternoons in much of the US. If electricity customers exhibit even modest demand elasticity, then reducing the mismatch between the fluctuating cost and the fixed retail price could eliminate billions of dollars in deadweight losses [Borenstein, 2005]. Residences consume 36% of the electricity used in the US.[1] Electricity demand to air condition (often empty) homes may be more elastic than commercial and industrial demand during the workday.

[1]2004 figures from the Energy Information Administration's Electric Power Monthly: http://www.eia.doe.gov/cneaf/electricity/epm/table5_1.html

This paper assesses the impacts of a residential critical peak pricing (CPP) pilot program. CPP uses a small menu of off-peak, peak, and "critical" prices to capture some of the variation in power cost over time. Peak prices are in effect during scheduled hours, typically every non-holiday weekday afternoon. The utility can invoke a significantly higher "critical" price a limited number of times per season. CPP is the dynamic pricing approach that most practitioners consider for customers who use only a modest amount of power. California's Statewide Pricing Pilot (SPP) field experiment exposed a sample of customers of its three major investor owned utilities to CPP from July 2003 through September 2004. The SPP cohort studied here paid a peak price from 2-7 PM on non-holiday weekdays except when the experiment used automated phone calls to inform participants that critical prices would be in effect the next day. This paper analyzes the effects of CPP on SPP customers' power use on summer weekday afternoons.

Careful empirical research can provide information about the magnitude of customers' short term response to the new prices as a function of weather, historical aggregate use, climate, and other customer and customer-day characteristics. Estimates of the magnitude of response can inform discussions about whether the benefits of an improved pricing plan justify the cost of advanced meters and recruiting customers. Analysis can identify the customers who are most worth recruiting.

1.1 Prior Work

This paper extends a literature that uses the same SPP data set, namely:

- Faruqui and George [2005] and the SPP final report [Charles River Associates, c] use a continuous elasticity of substitution demand model to estimate the impacts of the SPP's price changes.[2] These papers compare customers on CPP to the SPP's control group. They make fairly strong functional form assumptions about the nature of customer demand. The present paper explicitly tests a subset of their assumptions.

- Herter et al. [2007] looks at the combined effect of climate and temperature[3] on the impacts of dynamic pricing. Specifically, they estimate the difference in CPP customers' electricity consumption between critical-priced and ordinary, peak-priced weekday afternoons.

- Herter [2006a] estimates the impacts of the critical price relative to ordinary weekdays using just data from the "treatment" group that experienced CPP. It runs one regression per customer and does not report disaggregated impacts of customer characteristics like climate zone or central air conditioning ownership on response. It estimates use hour-by-hour.

[2]Faruqui and George [2005] is adapted from the executive summary of Charles River Associates [c] which provides a detailed documentation of their econometric approach.

[3]Climate (i.e. the typical weather pattern) drives customers in some areas to invest in air conditioners and insulation. Temperature is the weather realization on a single day. Both climate and temperature affect electricity demand. Customers in climates with frequent hot days will run air conditioning on a 95^oF day. Customers in temperate climates often **have no** air conditioner to run on a 95^oF day.

Wolak [2006] does careful applied econometric work on an experiment involving residential customers in Anaheim, California. The experiment's baseline-rebate rate offered rebates to customers whose use during critical periods was below their highest historical usage.

All of these studies suggest that dynamic pricing causes significant reductions in usage during high priced periods.

This paper extends the literature with a more detailed, flexible difference-in-difference regression approach. It combines Herter et al. [2007] and Herter [2006a]'s functional form flexibility with Faruqui and George [2005]'s ability to estimate the impact of both critical and peak prices. This project explores how temperature and customer characteristics affect response. I also report selection issues that may change the interpretation of some of the existing literature.

2 Studying Opt-in CPP: The SPP's recruitment process, selection issues, and time line

The SPP's treatment group (this document uses "CPP group" and "treatment group" interchangeably) experienced CPP while its control group continued on status quo time-invariant rates. The SPP collected data on each group's usage patterns before and after the CPP group switched from California's standard, time-invariant rates to CPP. The SPP's sampling strategy recruited treatment and control groups that could be weighted to represent the state's population. The SPP was designed to facilitate the estimation of a demand system. The SPP's design also lays groundwork for a clean, powerful difference-in-difference analysis.

2.1 Design Compromises and Differences between the Control and Treatment Groups

Unfortunately design compromises and operational challenges led SPP implementation to diverge from an ideal design, which is a common problem among field experiments,.[4] [5] The SPP involved impressive inter-organizational negotiation and coordination among three utilities, state agencies, stakeholders, and evaluation contractors. Some SPP working group members objected to making CPP nearly mandatory or even putting customers on CPP unless they actively opted out [Charles River Associates, c, 30]. Thus, potential CPP customers

[4]This kind of problem is widespread enough that there is a "broken experiments" literature (e.g. Barnard et al. [2003]) that considers ways to fix this kind of flaw.

[5]The SPP's reports [Charles River Associates, c,a][6] document the Statewide Pricing Pilot, its sampling strategy and final enrollment. Their appendices [Charles River Associates, d,b] contain examples of the recruitment materials, Welcome Kits, and surveys sent to customers. Herter [2006a] also provides useful documentation.

got detailed information and had to affirmatively decide to participate.[7] [8] In practice, residential critical peak pricing programs are and will likely continue to be opt-in until they build a track record, so this design yields valuable information about CPP's effects on customers willing to volunteer for CPP. However, the SPP randomly selected its control group, making their use an imperfect counterfactual to the treatment group's use. [9] The SPP surveyed both its treatment and control groups about their energy use and recorded the same detailed usage data for both. [10]

There is reason to think there are some substantive differences between the treatment and control groups. Many customers who demand a large amount of weekday afternoon power apparently declined offers to be treated or left the experiment early. A study of why customers chose not to participate reports that, "Virtually everyone who refused to participate in a particular pilot rate believed they would have wound up spending more - and perhaps a lot more - on electricity if they switched to the new pricing plan."[11]

Table 1 compares a variety of characteristics of the treatment and control groups.[12] The

[7]An example of the SPP's detailed invitation letter that describes the new pricing, study requirements, and $175 in participation payments is in Charles River Associates [b, 18-23]. Experiment staff spent up to two weeks trying to contact each customer. They contacted the next alternate in line if efforts to reach the customer failed, the customer refused to participate, or the customer was ineligible Herter [2006a].

[8]SPP customers had to take action by either returning the enrollment card or agreeing on the telephone to participate and that only about 20% of the customers that the experiment tried to recruit did so. The SPP final report sometimes describes that as "opt-out" while I characterize it as "opt-in." Specifically, a passage in the SPP Final Report [Charles River Associates, c, 30] characterizes the design as opt-out: "The final SPP design involved mailing an enrollment package to selected customers and obtaining an affirmative response regarding the willingness of each customer to participant (sic.). As such, it is a voluntary program but one predicated on an opt-out recruitment strategy rather than an opt-in one." In the end, refusal to participate was not the most common reason that an invitation was unsuccessful: "Ultimately, about 20% of customers accepted the invitation to participate, 15% declined to participate, and the remaining 65% were unreachable or otherwise excluded. Subsequent analyses using mean comparison and Heckman correction indicated that the final sample was a representative cross-section of California residents by appliance holdings, income, education, and 16 other measured variables" (Herter [2006a] citing a draft of Charles River Associates [a]). Thus, the differences in the use of the terms opt-in and opt-out are a minor, semantic difference.

[9]An ideal study of opt-in CPP would recruit a sample of customers willing to participate in a study of new electric rates and then randomly assigned these, willing customers to control and treatment categories.

[10]Customers completed a survey "in most cases at least one month after the Welcome Package was sent. Many surveys were not completed until the fall of 2003" [Barnes, 2007]. The surveys described the customer's appliances, home, household members, and appliance usage habits. The CPP group reports being far more likely to use their dishwasher, laundry, and air conditioning only off peak in ways that the Welcome Kit suggests. The survey's timing makes it impossible to understand to what extent these differences were preexisting. There is a copy of the survey and documentation about how Faruqui and George coded its variables for use in their work in Charles River Associates [d]. In a few cases that should be clear from the tables this paper analyzes more detailed, disaggregated data from the survey than they did.

[11]It goes on to temper the notion that this was a fully informed choice, writing that, "[N]one of the respondents had actually used the graphics to calculate whether they would be better off, or worse off, under the new pricing plan. Everyone admitting to just 'eye-balling' the bar chart and new rate plan and then deciding they probably would wind up spending more."[Focus Pointe, 6,22]

[12]Table 1 reports that there are the minimum peak period load in the data is zero. In fact, a bit less than 1% of all customer-non-holiday weekdays report zero peak period load. These entries are strange because things like refrigerators and electronics tend to draw power regardless of whether customers are home. Extensive investigations reveal no clear patterns by date or by utility. Two customers, who each report more than 100 days with zero use account for about half of the zeros. Two explanations seem plausible: these zeros could

	control subjects	treatment subjects	p-value	min	max
avg. daily use, kWh, summer 2002	17.10	16.70	0.643	2.06	78.30
weekday peak use as % of total use; 6/1-15/03	0.21	0.19	0.066	0.05	0.52
avg. use, kWh, weekdays 2-7PM, June 1-15 '03	4.24	3.82	0.144	0.41	30.50
avg. daily use offpeak usage, kWh, June 1-15 '03	10.80	10.70	0.904	1.49	54.00
avg. 4PM temperature, June 1-15 '03	74.00	73.70	0.698	60.20	99.90
# children 0 to 5	0.32	0.29	0.674	0.00	4.00
# children 6-18	0.65	0.59	0.538	0.00	5.00
# people over 65	0.27	0.32	0.507	0.00	4.00
everyone in household is > 65	0.09	0.13	0.251	0.00	1.00
home built after 1979	0.39	0.38	0.932	0.00	1.00
% work from home part/full time	0.15	0.12	0.384	0.00	1.00
agrees w/ "everyone should pay a little ...[for] a cleaner environment"	0.53	0.67	0.007	0.00	1.00
agrees that "a cleaner environment will mean fewer jobs"	0.23	0.20	0.576	0.00	1.00
agree/strongly agree that 'global warming is a threat...'	0.71	0.66	0.344	0.00	1.00
1=rates utility performance good or excellent	0.78	0.79	0.881	0.00	1.00
household head is a college graduate	0.44	0.47	0.596	0.00	1.00
has central air conditioning	0.45	0.43	0.783	0.00	1.00
has 1+ room air conditioners	0.15	0.16	0.814	0.00	1.00
electric well pump	0.03	0.03	0.813	0.00	1.00
# refrigerators + freezers	1.35	1.31	0.553	0.00	5.00
electric hot water	0.14	0.10	0.263	0.00	1.00
electric range	0.38	0.30	0.096	0.00	1.00
electric oven	0.44	0.40	0.407	0.00	1.00
electric dryer	0.37	0.31	0.199	0.00	1.00
programmable thermostat for Central AC	0.23	0.22	0.826	0.00	1.00
swimming pool	0.08	0.08	0.898	0.00	1.00
electric spa	0.07	0.05	0.318	0.00	1.00
number of customers contacted before one accepted	1.22	2.74	0.000	1.00	11.00

Table 1: Mean characteristics of households in the regression sample. With a few exceptions that are explored in depth in table 2, the treatment and control groups' observable characteristics are statistically indistinguishable. This table weights the sample to achieve the same geographic distribution as the state's population. ^{m} indicates that the p-value on equality of means comes from a Mann-Whitney rank sum test conducted on an unweighted sample of categorical answers. This non parametric test is appropriate because customers who reported having 750-1000 square feet of space have larger houses than those who checked "less than 750" but we have no basis on which to develop an accurate point estimate of the difference. The average income and square footage figures are coded as documented in [Charles River Associates, d, 113-119], typically assuming that each customer is at the midpoint of the range they selected.

		whole regression sample		apts. and low use single family		high use single family	
	cust. type	all subjects	seen > 4 months	all subjects	seen > 4 months	all subjects	seen > 4 months
kWh / day, summer '02	control	17.10	17.40	12.20	12.40	33.20	33.10**
	CPP	16.70	16.70	12.30	12.60	30.90	30.10**
weekday kWh 2-7PM, June 1-17, '03	control	4.24	4.29*	2.95	3.00	8.46**	8.42***
	CPP	3.82	3.79*	2.75	2.82	7.27**	6.90***
daily offpeak kWh, June 1-17 '03	control	10.80	10.90	8.14	8.26	19.40	19.30
	CPP	10.70	10.80	8.25	8.46	18.70	18.30
4PM temperature, June 1-17 '03	control	74.00	73.90	73.20	73.20	76.40	76.30
	CPP	73.70	73.70	73.00	73.10	75.90	75.40
# people over 65	control	0.27	0.28	0.25	0.26	0.34	0.34
	CPP	0.32	0.34	0.31	0.33	0.35	0.35
everyone in household is > 65	control	0.09	0.09	0.09	0.09	0.09	0.08
	CPP	0.13	0.14	0.14	0.16	0.08	0.08
% work from home part/full time	control	0.15	0.16	0.11	0.12	0.28	0.28
	CPP	0.12	0.12	0.10	0.09	0.21	0.20
agrees "everyone should pay [for] a cleaner environment"	control	0.53 ***	0.52 ***	0.56	0.56*	0.43***	0.43***
	CPP	0.67 ***	0.69 ***	0.66	0.68*	0.70***	0.70***
agree that "global warming is a threat..."	control	0.71	0.71	0.76	0.77	0.53	0.53
	CPP	0.66	0.67	0.67	0.68	0.63	0.63
rates utility good or excellent	control	0.78	0.78	0.80	0.80	0.73*	0.73*
	CPP	0.79	0.80	0.78	0.80	0.82*	0.83*
central air conditioning	control	0.45	0.45	0.38	0.38	0.67	0.67*
	CPP	0.43	0.42	0.39	0.38	0.58	0.56*
electric range	control	0.38*	0.38*	0.37	0.37	0.41	0.41
	CPP	0.30*	0.29*	0.28	0.27	0.35	0.34
recruited before participant found	control	1.22 ***	1.22 ***	1.22 ***	1.22 ***	1.23***	1.23***
	CPP	2.74 ***	2.76 ***	2.71 ***	2.75 ***	2.85***	2.80***
total annual household income, 1000's	control	68.17^{*m}	68.68^{*m}	59.25^{*m}	59.79^{*m}	94.52^{m}	94.16^{m}
	CPP	58.87^{*m}	59.45^{*m}	49.41^{*m}	50.38^{*m}	89.51^{m}	89.50^{m}

Table 2: Differences between control and CPP groups: at the beginning of the experiment and after the first four months of attrition. The high use CPP group uses less power during peak hours than does the high use control group. This difference grows with attrition. Statistical significance of differences between the control and CPP groups: * .10, ** .05, and *** .01. Attrition causes no statistically significant changes in mean. All reported values are weighted by region to give the sample the same geographic distribution as the state's population. As described in depth in the caption to table 1, ^{m} indicates that the p-value on equality of means comes from a Mann-Whitney rank sum test conducted on unweighted, categorical answers.

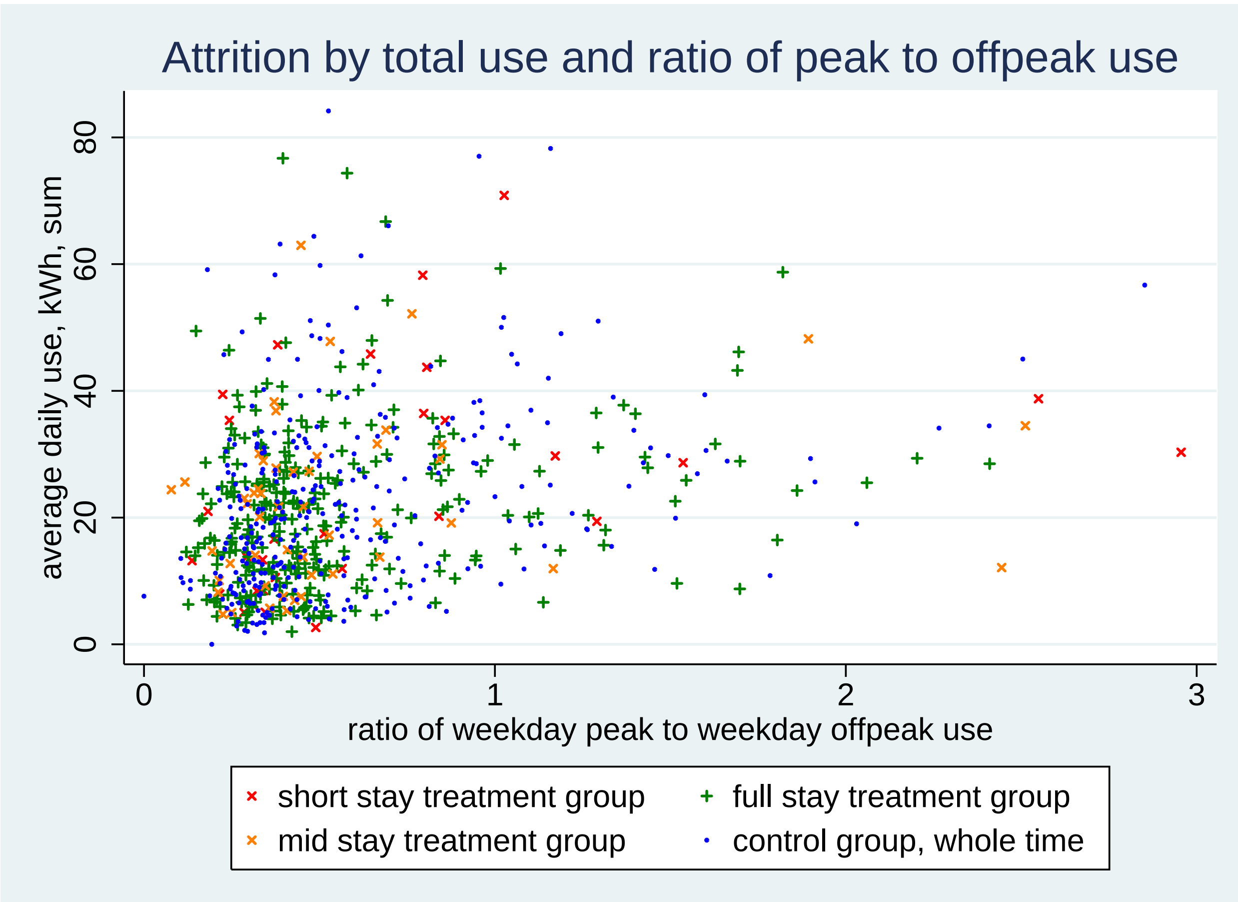

Figure 1: *The CPP group has fewer customers who both use a large total amount of power and use a large proportion of that during peak-priced periods than does the control group. There are relatively few treatment customers who used more than about 35 kWh per day and whose peak use was roughly as big or bigger than their off peak use. Thus, there are few treatment customers in the top right part of the scatter plot.*

treatment and control groups are generally quite similar, with a few notable exceptions, namely:

- The CPP group has fewer high use customers who also use a large proportion of their power during peak-priced periods. If we separate this into components, we find statistically insignificant differences in the distributions of total use or of proportion of power used on peak. The distributions start out with this difference and attrition increases the difference. Further, the SPP divided its sample into three cells: apartments, and high and low use single family homes. This difference is only significant in the high use single family home cell. People who use more power may have a better sense of when and how they use it. Holding the percentage of the power that a customer uses during each price period constant, customers who use a greater total amount of power are more likely to notice and react to changes in their bills.

- CPP and control customers express similar levels of concern about environmental problems, but treatment customers are more committed to civic action. CPP customers are more likely to agree or strongly agree with the statement: "I believe everyone should pay a little bit more to ensure a cleaner environment." The CPP and control groups are, however, indistinguishable in their propensities to agree with "The cost of a cleaner environment will mean fewer jobs and hurt the economy" (sic.) and "Global warming is a threat I am seriously concerned about." This apparent difference in civic-mindedness is unsurprising in a social experiment being sponsored by state agencies that used recruiting materials that, "[W]ere quite ineffective [marketing]. .... The materials made scant reference to any benefit - direct or indirect - that the customer might gain by participating...." [Focus Pointe, 6].

- The CPP group begins with fewer kids under 5, but attrition makes this difference statistically insignificant. While young children might seem to make adjusting electricity use more difficult, the evidence suggests that the kids under 5 in our data do not cause much of an increase in afternoon consumption.[13]

- The treatment and control groups start with statistically indistinguishable numbers of home businesses but many treatment customers with home businesses exit the sample, creating a statistically significant difference.

- Even if the groups were in fact drawn from identical distributions, we would expect that 1 in 20 differences to be statistically significant at the 5% level purely by chance. So it is possible that this discussion overinterprets a result.

be a product of known meter problems or they reflect periods in which customers shut off their electricity e.g. for repair work.

[13]There is no data available about whether the kids under 5 are in day care. We cannot rule out the possibility that the CPP group has fewer kids actually present weekday afternoons than the control group.

SPP Summer Rates in Surcharges and Credits in Cents/kWh		
	high ratio rate	low ratio rate
critical	+60.9	+41.8
peak	+11.6	+9.8
off peak	-5.1	-1.2

Table 3: **The SPP Summer Rates.** The SPP defined its peak, off peak, and critical rates in terms of surcharges and credits relative to the standard, underlying utility rates. The underlying rates have a complicated increasing block structure. The rates changed modestly during the course of the experiment. "The average prices, expressed in cents/kWh, during the summer of 2003 were 12.7 for PG&E and, rounded, 14.1 for both SDG&E and SCE" [Charles River Associates, a, 21]. A kWh is enough energy to run a 100 watt light bulb for 10 hours or a central air conditioner for about 15 minutes. The experiment assigned each CPP customer to either a high or a low ratio rate. The high ratio rates had a bigger difference between the cost of afternoon and off peak power than did low ratio rates. This table presents PG&E and SDG&E's Summer CPP Surcharges and Credits for the SPP in cents per kWh. The SCE Rate appears to deviate from the PG&E and SDG&E rates reported here by up to two cents. Sources: author's calculations based on Pacific Gas & Electric, San Diego Gas & Electric, Southern California Edison, Charles River Associates [b]

2.2 Choosing a Reliable Subset of Data

I work around some of the deviations from an ideal experimental design by analyzing only about 60% of the data:

- Experiment materials may have caused premature response, so the analysis dropped data around the transition date: The SPP mailed instructions to the (initial) cohort of participants studied here [14] over a period lasting between 1 and 2 weeks starting on June 17, 2003 [Barnes, 2007]. These mailings included detailed instructions about how to reduce peak electricity use but only mentioned that the new rates went into effect on July 1 on their 18th page.[15] The gradual mailing and lack of a prominent discussion of the start date means it is not clear when each customer started responding. Thus, this analysis drops data from June 18 through Thursday, July 3, 2003 inclusive. [16]

[14] A typical Welcome Kit is in the Report Appendices [Charles River Associates, b, 18-23]. It appears to be for customers starting in a later cohort because its time line differs from the Welcome Kits that Karen Herter provided the author. The Welcome Kits were nearly identical across utilities.

[15] Further, the letters inviting customers to join the treatment group described the new rate on their second page, but only mentioned its effective date on their third page. If, however, premature response in early June were a serious problem, we would expect to see a statistically significantly different relationship between the CPP and time-invariant subjects' uncontaminated, pre-experiment average daily use from summer 2002 and average afternoon use during June 2003. We would also expect to see a different relationship between treatment and control customers' weekend and weekday use during June because treatment customers would be shifting use of equipment like washers and dryers from weekdays to weekends and because they would be reducing their air conditioning use only on weekdays. The regression presented in table 4 fails to reject the null hypothesis that there is no treatment-control difference in these relationships at the p=.2 level.

[16] This paper's approach to pretreatment period data falls between prior authors' approaches. One group of authors decided not to use pretreatment period data as a control (and hence to only estimate the im-

- To deal with the gradual deployment of "interval" electric meters that recorded participants' power consumption every 15 minutes, I focus on the initial cohort of customers which started CPP on July 1, 2003 and had meters on by June 15, 2003. [17] Further, I exclude pretreatment data from before June 1, 2003 because data before that date comes from a small, unrepresentative group of customers.[18]

- Weather data for the PG&E region are missing for August 2003, so the analysis drops the incomplete observations.[19]

- The SPP assigned each treatment consumer to either a high- or low-ratio rate. The two rates provided qualitatively similar incentives to change behavior during the May 1-October 31 summer season and quite different incentives during the winter season because the low-ratio rates' had almost no incentive to conserve during ordinary winter peak periods while the high ratio rate strengthened its winter conservation incentive. The summer is far more substantively important. Thus, this paper considers only the summer rate season when the rates in Table 3 were in effect so it can pool the customers into a single CPP "treatment" group and test whether this pooling is appropriate. .[20]

2.3 Summer Season Weather and Load Patterns in 2003-04

The relationship between California's 2003-04 weather, population, and electrical demand affect the results and their interpretation. The SPP used four climate zones, where zone 1 (the coastal fog belt) was the coolest and got progressively hotter through zone 2 ("the foothills", including much of greater Los Angeles and the inland suburbs of San Francisco), zone 3 (the Central Valley), and zone 4 (the desert). The SPP assigned each customer to

pacts of the critical price beyond the impact of the daily peak price) because "There is some debate about how accurately the SPP pretreatment load data reflects uninfluenced pre-experiment load, since customers received information and instructions about how to reduce peak loads prior to the pretreatment period" [Herter et al., 2005, 7]. A second group of authors appears to have used all of the pretreatment data [Charles River Associates, c,a, Faruqui and George, 2005].

[17] There were "late starting" cohorts of CPP Customers who were recruited too late to have their meter on for roughly a full month before July 1, 2003. This cohort experienced CPP about a month after their meter was turned on. Welcome kits were sent to late starting customers on an ongoing basis. I drop them from this analysis because 1) there is no data available about when late starting customers received their welcome kits and thus what part of the month of pretreatment data collected is meaningful and 2) comparing CPP pretreatment data from many time periods to control pretreatment data exclusively from the month of July may confound seasonal shocks to weather demand with preexisting differences.

[18] The first CPP customer meters report data starting on April 23, 2003. The first control group meters came on line March 31. The first 31% of all customers' meters were activated in March through May of 2003.Another 23.5% started reporting data exactly on June 1, 2003 for a total of 54%. A total of 75% of interval meters were on by July 1, 2003. Considering only participants whose meters were on by June 15th ensures that we have usable measures of pretreatment consumption from at least one weekend day and one weekday.

[19] PG&E used data from different weather stations than those in major public data sets like NOAA's Local Climactic Data, leaving no trivial way to replace the missing data.

[20] A natural extension of this analysis would include winter results and interact each customer's rate with every analysis variable. Charles River Associates [c] and Faruqui and George [2005] roughly take that approach.

Dependent variable: consumption on non holiday weekdays in kWh/h				
	Controlling for Summer 2002 Use		Also controlling for Weekend Use	
	Pretreatment	During Treatment	Pretreatment	During Treatment
Treatment Customer	-0.188 (0.155)	-0.109 (0.132)	-0.098 (0.111)	-0.034 (0.100)
electric use, kWh / day, summer 2002	0.057*** (0.005)	0.066*** (0.004)	0.010 (0.006)	0.033*** (0.005)
treatment customer * use, kWh / day, Summer 2002	0.004 (0.007)	-0.006 (0.006)	0.015* (0.008)	0.003 (0.008)
constant	0.058 (0.122)	0.123 (0.104)	0.049 (0.084)	0.123* (0.070)
avg. weekend 2-7PM use, kWh, 5/31-6/15 2003	. .	. .	0.149*** (0.013)	0.102*** (0.012)
trt. cust. * avg. use, kWh / 2-7PM wknds June 1-17 2003	. .	. .	-0.041** (0.020)	-0.032 (0.022)
R^2	0.394	0.367	0.525	0.415
N	4327.	67211.	4318.	65881.
P-value, all treatment customer coefficients=0	0.200	0.0006	0.209	0.002
Robust standard errors, clustered by customer in parentheses. Significance: *=10% ** =5% ***=1%				

Table 4: Regressions comparing the relationship between historical use and weekday peak use in the control and treatment groups. If there is no premature (i.e. before June 18, 2003) response to the price signals, there should be no difference between these relationships in the pretreatment data. If the treatment succeeds, there should be a difference in the treatment period data. We find exactly that pattern.

one of 58 weather stations around the state.[21] Hot summer climates, similar to zones 3 and 4, are typical of much of the United States.[22]

1. Statewide population-weighted[23] base-78 cooling degree hours (CDH) have a very strong positive correlation with electricity use. Figure 3 illustrates this correlation. Cooling degree hours measure the need for air conditioning. The number of cooling degrees hours is the number of degrees that the temperature is above a base level during each hour. This paper works with a base of 78^o Fahrenheit and takes a sum over the 5 hours between 2 PM and 6:59PM. In other words, it is $\sum_{T=2PM}^{6PM} max\{0, temperature_t - 78\}$ So 50 (60) CDH reflects an afternoon that had an average temperature of 88^o (90^o) F.

2. About 6.5 million of the three utilities' 8.3 million accounts are in climate zones 2 ("foothills") and 3 (Central Valley). The weather in these zones changes more over the course of a summer than does the weather in the generally cool fog-belt and generally hot desert. The variable weather and large population in Zones 2 and 3 gives them a disproportionate impact on electricity consumption in California. In particular, roughly 1 in 4 weekdays in 2003-04 had demand of 40 GW or more. All but one of these days had population-weighted 2-7 PM afternoon temperatures averaging more than 85^o (35 CDH) in zone 3, which was hotter than average since only 43% of all afternoons in May through October averaged 85^o in zone 3.

3. Hot weather in foothill zone 2 has a bigger impact on state-wide, population-weighted CDH than it does on electrical load because zone 2 has a large population, but only about 30% of these accounts have weather-sensitive central air conditioning, while more than 70% of customers in zones 3 (Central Valley) and 4 (desert) do. Figure 2 illustrates the relationship between zone and statewide temperature.[24]

4. Climate zone 4 (desert) is quite hot for very extended periods of time but has the smallest population. More than 50% of days during the experiment's summer-rate months of June - October 2003 and May-September 2004 had more than 60 CDH there, indicating that the average temperature was at least 90^o between 2 and 7PM.

5. Quite high demand days tend to have unusually hot weather in zones 3 and 4. The two highest demand days were also unusually hot in zone 2.

[21] A zone map is at Charles River Associates [c, 22] and the weather stations are listed in Charles River Associates [d, 18-19].

[22] Descriptive statistics about summer heat measured in Cooling Degree hours for cities around the US are available at: http://www.ncdc.noaa.gov/oa/climate/online/ccd/nrmcdd.html That site suggests identifies areas of the US that have total cooling degree days comparable to places in climate zone 4 such as Bakersfield and Fresno and zone 3 such as Stockton and the Sacramento suburbs. (The Sacramento Municipal Utility District did not participate in the SPP, but PG&E assigned some customers to a "Sacramento" weather station.)

[23] Appendix C describes this paper's population-weighted temperature calculation methodology.

[24] Appendix D further details these patterns by providing tables of the distribution of cooling degree hours by climate zone.

6. Twelve of the SPP's summer-season critical days were called during days with peak electricity demand in the top 10% of all 2-7 PM CAISO-region summer peaks. These critical days are probably fairly representative of the critical days that would have been called during 2003-04 had the SPP been run to minimize the costs of the energy system. The SPP was also very consciously an experiment. It thus called the other 12 critical days during periods when energy was not particularly scarce to explore how customers would react to price signals during cooler conditions and whether price signals could be used to manage scarcity created by generation or transmission problems that might coincide with a mix of temperate weather in some regions and extreme weather elsewhere. Thus, they called two critical events in October on days with between 30 and 32 GW of peak load, putting these days' peaks between the 10th and 20th percentiles of all summer season peaks. The remaining 10 critical days ranged from the 40th to 90th percentile of demand, with 6 called in July, August, and September on days between the 70th and 80th percentiles of the demand distribution. They left enough of the hottest, highest demand days non critical to allow us to estimate the impacts of both peak and critical prices under the hottest conditions seen in 2003-04.

2.4 The experiment period lacked extremely hot days

The SPP ran during two years that lacked the kind of extremely hot days that often create scarcity because demand gets so high it might outstrip supply.[25] This limits our ability to explore the impacts of dynamic pricing during the truly unusual demand events when reducing the amount of power consumed would have the greatest benefits.

The California Independent System Operator (CAISO) declares a Stage 1 emergency when it does not have enough capacity to meet standards about keeping the system robust to equipment problems, a Stage 2 emergency when shortages force it to ask customers (typically large industrial facilities) on interruptible contracts to stop using power, and a Stage 3 emergency when shortages lead to rotating blackouts. The CAISO declared no summer season Stage 1, 2, or 3 electrical emergencies during the CPP treatment period. By contrast, CAISO declared emergencies on six summer-season days in 1998 and three in 2006.[26] Four of the 1998 emergencies reached Stage 2, as did one in 2006. The Stage 2 emergency came when record setting heat hit California, especially the heavily populated climate zones 2 and 3 on July 24, 2006. [CAISO, a,b] [27]

[25]Indeed, by far the highest statewide average CDH in 2003-04 was on Sunday, September 5, 2004 when zones 1 and 2 got unusually hot. That day had an unspectacular peak demand of 38.3GW because so much commercial and industrial demand was off line for the weekend. Offpeak prices are in effect all day Sunday, so it is not in the impact analysis.

[26]A combination of extreme weather and an institutional meltdown led to power emergencies on 125 days during the 2000-2001 crisis. This number seriously overstates the number of emergencies that would have taken place had the institutions been functional.

[27]Although the author does not have directly-comparable, population-weighted 2006 weather data, it is illustrative to note that the temperatures at airports in the three zones on July 24, 2006 were least 30 CDH higher than the hottest population-weighted weekday peak period temperatures observed in zones 2-4 in 2003-04. Specifically, airports in San Jose (zone 2) recorded 76 base-78 CDH, Sacramento (zone 3) had 149, and Fresno (zone 4) reported 161.

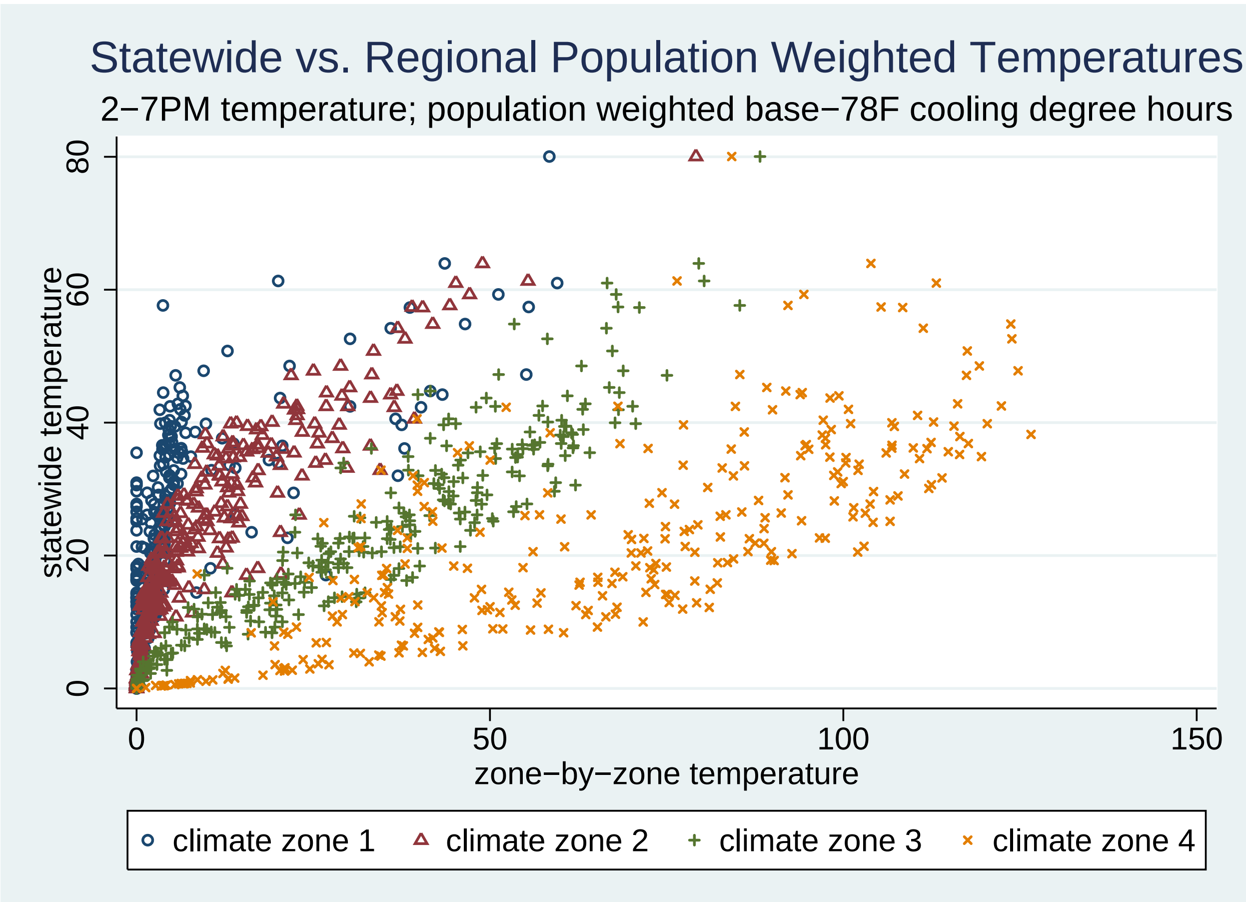

Figure 2: *The relationship between population-weighted average temperatures in each climate zone and the statewide population-weighted average temperature. Notice that each y-coordinate (a day with a single population-weighted average CDH) has an entry for each of four zones with x-coordinates indicating the population-weighted average temperature. The very top entry makes this clear. This graph includes weekdays, holidays, and weekends.*

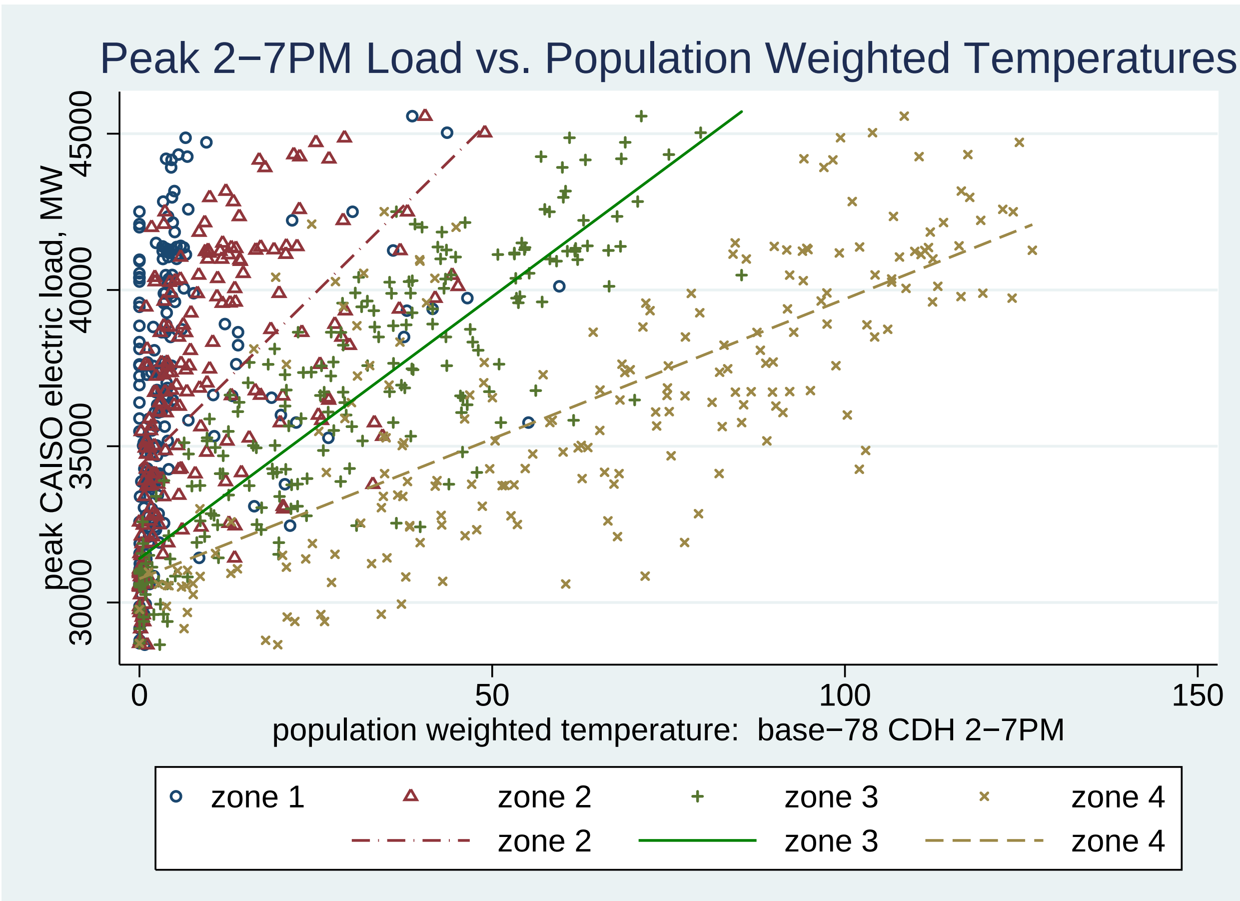

Figure 3: *The relationship between weekday zone-average afternoon CDH and maximum electric load in the CAISO system between 2 and 7PM. I have plotted the best linear fit of the 1 variable regression of each zone's CDH on load. Zones 3 and 4 fit quite well; zone 2 shows significant signs of omitted variables bias stemming from the correlation between temperature there and temperature in hotter neighboring zones. The omitted variables bias in zone 1 was so severe that the line is uninteresting. The zone 1 line is thus not displayed here.*

Another way to see the striking absence of extreme conditions during summer 2003-04 is to examine peak demand. The 99th percentile demand in a pooled sample of 2003 and 2004 was almost indistinguishable from the 99th percentile of demand in 2006, but CAISO peak demand was 50.2 GW at 4PM on July 24, 2006, which eclipsed the 45.6GW 2-7 PM peak in 2003-04.

This paper deals with the lack of very hot weather during the sample period by creating "synthetic" hottest peak- and critical-priced days that combines data from the day with the highest population-weighted average temperature in each climate zone. This yields nearly the hottest in-sample weather for each price level.[28]

3 Econometric Approach

This paper takes a difference-in-difference approach to estimating the impacts of dynamic pricing. Difference-in-difference estimates assume that the control and treatment groups would have maintained any preexisting differences and have experienced, on average, the same changes in consumption over time. It attributes any differences in their trajectories after the price change to dynamic pricing. The approach taken here:

- Starts with a standard four-cell difference-in-difference (before/during)*(control/treatment) setup. It generalizes this to a six-cell case with two "during" periods, representing peak and critical priced afternoons respectively.[29]
- Interacts customer and customer-day characteristics with the indicator variables for the six cells. These characteristics include the customer's summer 2002 average daily electricity use, the customer's climate zone, whether the customer has central air conditioning, and the number of cooling degree hours on each afternoon. Cooling degree hours measure the duration and intensity of heat and thus the need for air conditioning. [30]
- Adds a more detailed set of controls for day and weather but does not interact these with treatment status.
- It clusters standard errors by customer to deal with the fact that many variables are measured at the customer level and to deal with serial autocorrelation in weather and consumption.

[28] We could get marginally hotter in sample data by independently selecting the hottest day at each of 58 weather stations rather than independently selecting the hottest day in each of 4 zones.

[29] The coefficients and p-values would not change if we were to run the estimates as two conventional difference-in-difference regressions: one comparing the pretreatment period to ordinary days and one comparing the pretreatment period to critical days.

[30] This analysis, in contrast to Faruqui and George [Charles River Associates, c, 41], finds that using base $78^o F$ CDH gets better fitting results than does using base $72^o F$ CDH presumably because it avoids trying to fit the same quadratic relationship to the change from 75^o to 76^o in the Bay Area as it does to the change from 95^o to 96^o in the desert. Working with base-78^o CDH creates more zero CDH days in temperate climate zones than would base 72^o or 75^o CDH. Air conditioning is rare in climates where summer highs in the mid-seventies are common. Further, many air conditioners are likely to be idle at temperatures between 72^o and 78^o, but active at higher temperatures.

Figures 4, 5, and 6 illustrate a similar identification strategy on raw mean use data aggregated to the weather station-day level.

I estimate:[31]

$$\begin{aligned} avgLoad_{it} &= \alpha^T \mathbf{X}^* + \delta^T \, TrtCustomer_{it} \mathbf{X}^* + \gamma^T \mathbf{T}_{it} + \kappa^T \, TrtPeriod_t \mathbf{X}^* + \\ & \quad \nu^T \, CriticalPeriod_t \mathbf{X}^* + \beta^T \, PeakPrice_{it} \mathbf{X}^* + \psi^T \, CriticalPrice_{it} \mathbf{X}^* + \epsilon_{it} \end{aligned}$$

Where:

- $avgLoad_{it}$ is customer i's average kW (i.e. kWh/hour) consumption from 2-7 PM on weekday t. Electric load, measured in kilowatts (kW) is a measure of the rate of electrical use. A flow of one kW sustained for an hour is a kilowatt hour (kWh).[32] This paper often multiplies the flow of average impacts per hour that are measured in kW by five to get the average impact per afternoon peak period in kWh.

- $\mathbf{X}_{it}$ is a vector of customer- or customer-day specific controls drawn from 1) SPP administrative data, like the customer's climate zone and whether the customer lives in a single family house, 2) data from the weather station closest to the customer, 3) the customer's answers to survey questions, and 4) data about the customer's hour-by-hour usage during the pretreatment period.

- $\mathbf{X}^* = \begin{bmatrix} 1 \\ \mathbf{X}_{it} \end{bmatrix}$ Thus interacting a variable k with $\mathbf{X}^*$ yields both the base effects of k and interaction terms involving products of k and $\mathbf{X}_{it}$. The product $\alpha \mathbf{X}^*$ thus contains a constant.

- T_{it} is a vector of controls for day of week, calendar month, and year and the interactions of these variables with quadratics of cooling and heating degree hours.

- $TrtCustomer_i$ is 1 if the customer opted into the CPP group and zero if the customer is in the control group.

- $TrtPeriod$ is 1 when the CPP rate was in effect for treatment customers, i.e. the period starting July 1, 2003. It is zero before July 1, 2003.

- $PeakPrice_{it}$ is 1 if customer i received a peak, non-critical price on day t and is zero otherwise.[33]

[31] Bold characters and Greek characters are vectors.

[32] The rate of savings in kilowatt hours per hour is expressed here as kW, but elsewhere in the literature [Herter et al., 2007, Charles River Associates, c] as kWh/h.

[33] Making the $PeakPrice_{it}$ variable zero on critical days makes the standard errors on critical impacts easy to interpret, which is useful in section 4 below. This unconventional choice reflects a judgment that it was more convenient to have coefficients that tell us whether the complete impact of temperature or air conditioning ownership during a critical price event was statistically different from zero and to have to run an explicit hypothesis test to know whether the difference between the peak and critical difference was statistically different from zero than to have the opposite situation.

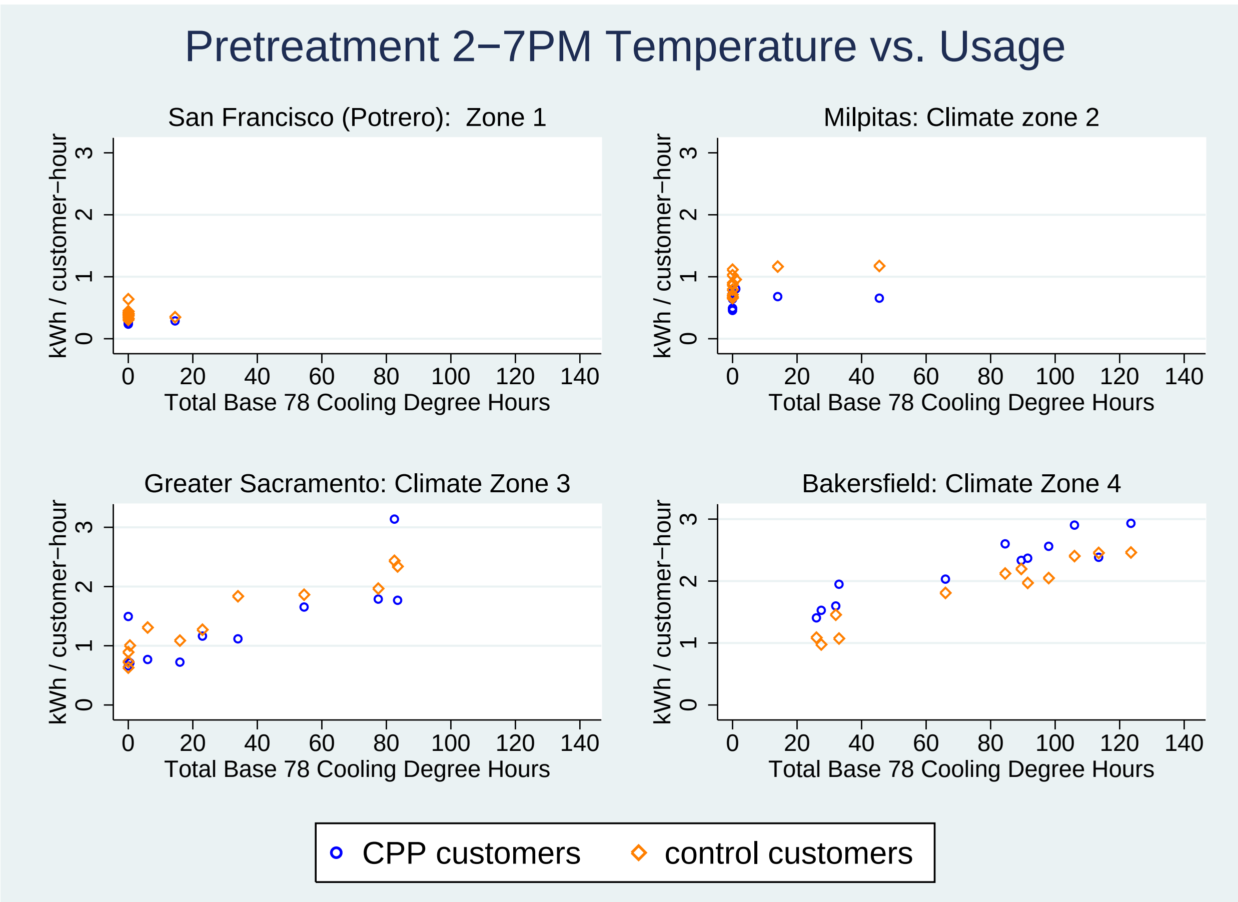

Figure 4: Identifying any pretreatment (i.e. weekdays June 1- June 17, 2003) differences. These graphs show the relationship between heat and average daily use for customers of one weather station in each of the four climate zones. The limitations of the identification strategy for this data set become clear here. There are only 12 pretreatment days, while there are roughly 200 treatment days. June is cooler than later months so the example weather station in foothill zone 2 (Central Valley zone 3) tops out at about 50 (100) base 78^oF cooling degree hours, while the same weather station has days with up to 80 (120) cooling degree hours later in the summer. Each point reflects the average use of the few dozen customers who are closest to the weather station in each graph's title. The graphs measure temperature as the sum of the base 78^oF cooling degree hours between 2 and 7PM.

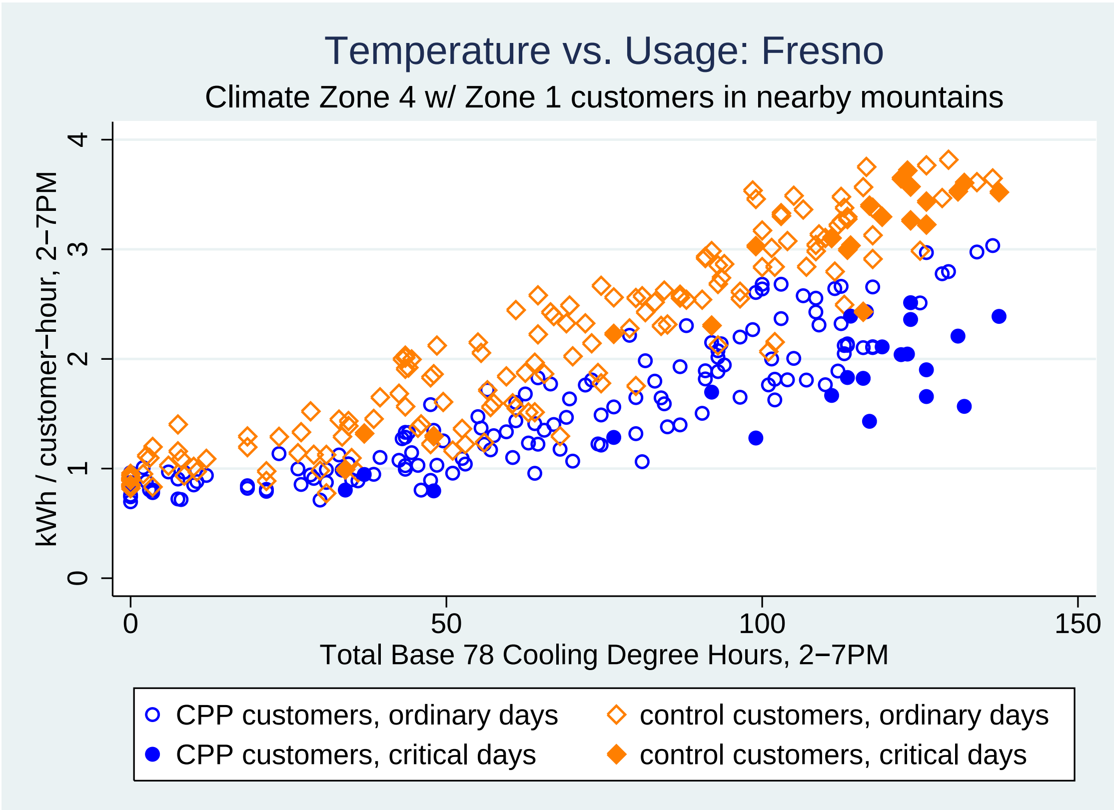

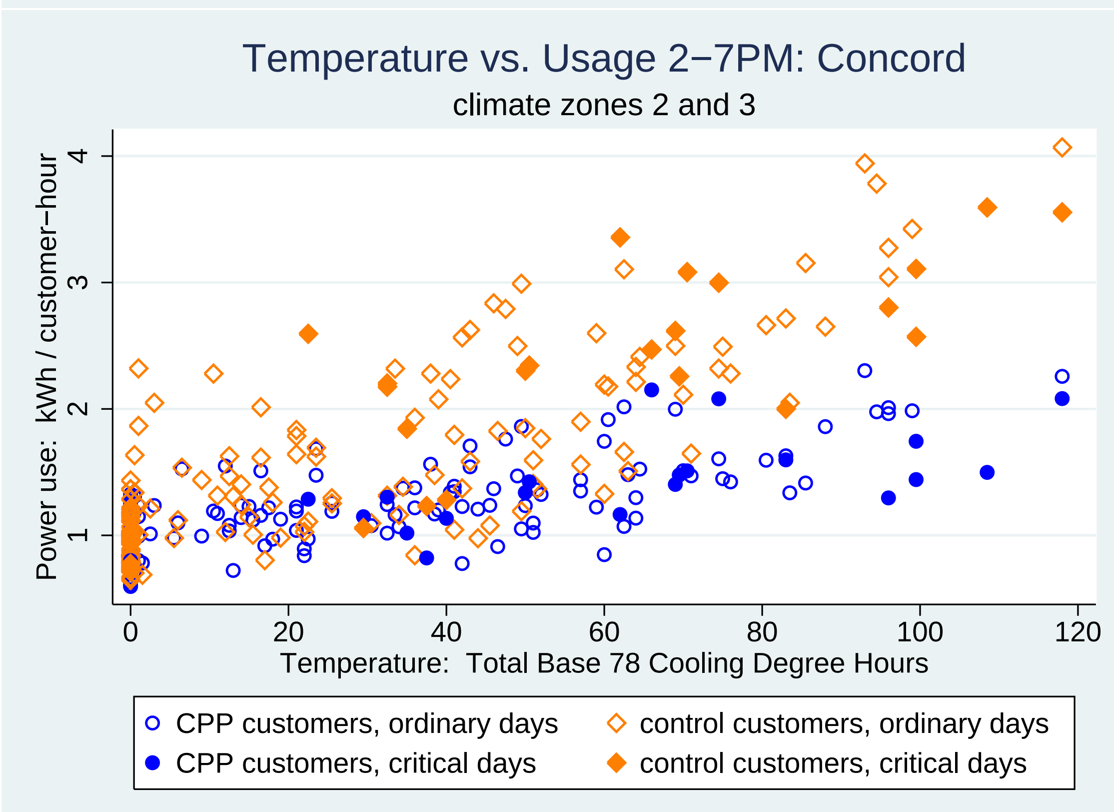

Figure 5: These figures plot the daily average hourly total electricity usage per customer between 2 and 7PM as a function of how hot the day was. They show that the difference between the average use of treatment and control customers is bigger on higher temperature days during the treatment period.

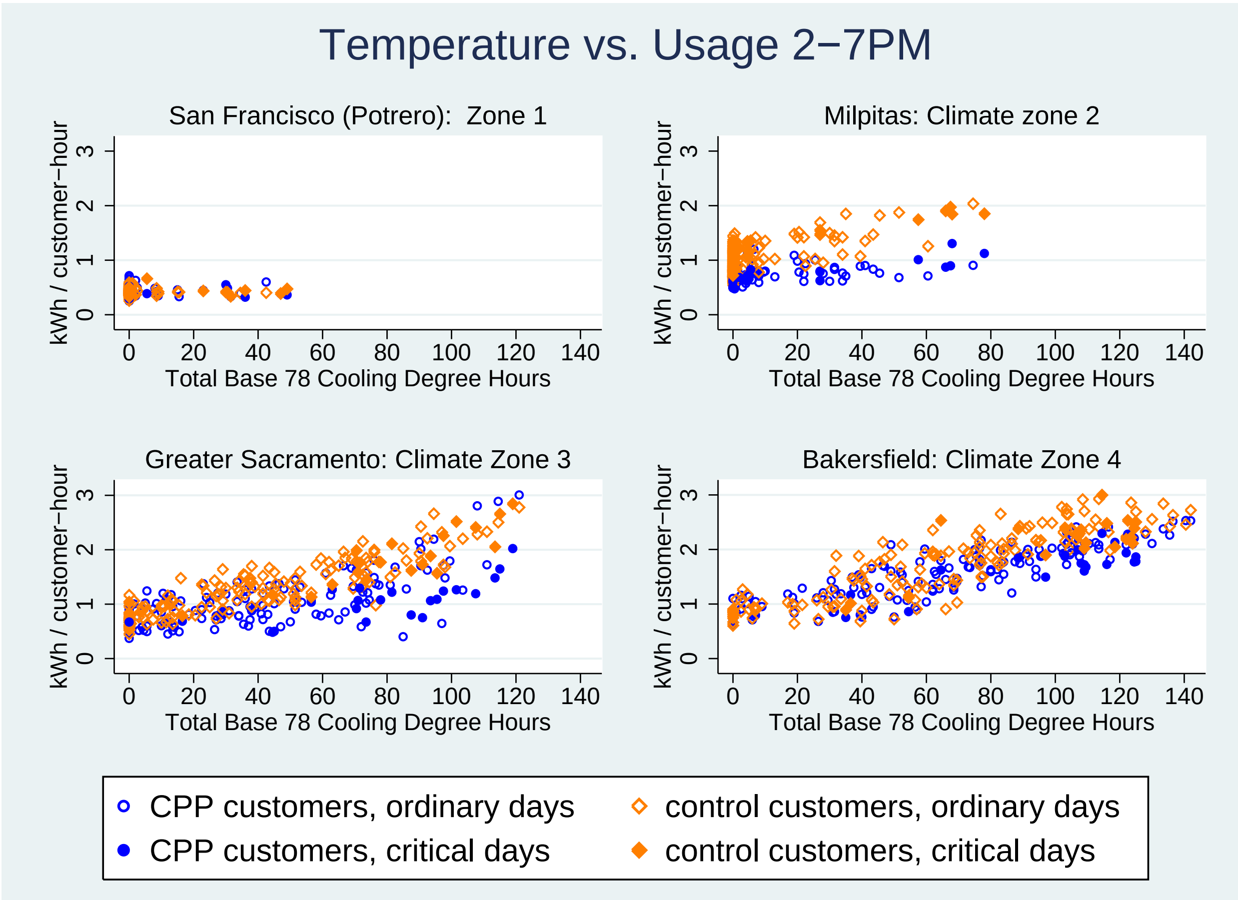

Figure 6: These graphs show the relationship between heat and average daily use for customers of one weather station in each of the four climate zones. Notice several important aspects of the identification: The vast majority of days with more than 80 base-78^oF cooling degree hours (i.e. days that have average 2-7 PM temperatures of more than 94^oF) in the sample come from climate zones 3 and 4. Indeed, many days in the temperate climate zones were below 78^oF all day and thus have zero cooling degree hours. Temperate zones 1 and 2 are less sensitive to temperature because fewer residences in those zones have air conditioners. This graph and figure 5 add to important findings from Herter et al. [2007]. First, nearly all of the data on critical events on days that are hotter than 95^oF (85 base-78 CDH) appear to come from climate zones 3 and 4, which puts the impact estimates in Herter et al. [2007] in context. Second these figures, especially figure 5, suggests that the substantive importance of response to the peak price signal on hot days is on par with the substantive importance of the additional response to the critical price signal. These graphs are not sufficient to draw conclusions, but show the likely origin of the findings below.

- $CriticalPeriod_{it}$ is 1 if a critical event was declared for CPP customers in i's climate zone on day t and is zero otherwise.
- $CriticalPrice_{it}$ is 1 if the utility successfully notified customer i that t was a critical day and is zero otherwise.[34]

3.1 Adding fixed effects to this framework

Some of the analysis reported below adds customer fixed effects to this econometric framework.[35] A fixed effects approach allows the regression to estimate customer-specific average usage level on an average day. This controls for important customer characteristics, like refrigerator efficiency, home insulation, and meal schedules, that the customer characteristic data do not measure. Controlling for customer fixed effects captures the impacts of all customer-specific characteristics that are unchanged throughout the experiment, so the estimation becomes:

$$\begin{aligned} avgLoad_{it} \;=\; & \eta^T Customer_i + \alpha^T \mathbf{X}^{\mathbf{w}} + \delta^T TrtCustomer_{it}\mathbf{X}^{\mathbf{w}} + \gamma^T \mathbf{T}_{it} + \kappa^T TrtPeriod_t\mathbf{X}^* + \\ & \nu^T CriticalPeriod_t\mathbf{X}^* + \beta^T PeakPrice_{it}\mathbf{X}^* + \psi^T CriticalPrice_{it}\mathbf{X}^* + \epsilon_{it} \end{aligned}$$

Where the variables are as above except that:

- $Customer_i$ is an array variables, one per customer, that are 1 for customer j if $i = j$ and zero otherwise.
- $\mathbf{X}^{\mathbf{w}}$ is the subset of $\mathbf{X}^*$ that remains identified for customer i on day t after we introduce fixed effects. The identified subset tracks weather conditions that vary over time.

These estimation strategies create two objects of interest:

- **The impact of the peak price** is $I_{peak} = \sum_{j\in\{1,\mathbf{X}^*\}} \beta_j \overline{x}_j$ where β_j is the coefficient on the interaction of $PeakPrice_{it}$ with the jth customer characteristic and $\overline{x}_j$ is the average value of the jth customer characteristic conditional on $PeakPrice_{it}$ being 1.[36]
- **The impact of calling a critical price** is quite similar, namely:

$$I_{critical} = \sum_{j\in\{1,\mathbf{X}^*\}} (\beta_j + \psi_j)\overline{x}'_j$$

  The differences are the addition of ψ_j the coefficient of *CriticalPrice*, and that we now calculate $\overline{x}'_j$ as the average characteristics on critical days.

[34] The experiment tracked whether its autodialer delivered its message to each participant on each critical day.

[35] Faruqui and George use customer fixed effects in their papers [Charles River Associates, c, Faruqui and George, 2005]. We can think of Herter [2006a]'s approach as being nearly equivalent to defining a fixed effect dummy variable for each customer and then interacting it with a set of control variables that describe characteristics of each day.

[36] For simplicity of discussion, I am treating the 1 as the first customer characteristic. The coefficient on 1 interacted with *PeakPrice* is the average impact of the peak price on consumption after controlling for all of the observed customer characteristics.

3.2 Weighting the Data

This analysis weights the data so that the control and treatment groups have the same geographic distribution as the state's population of electric accounts for each day of the sample.[37] Further, I down weight observations from periods – roughly July, August, and September – that we observed after the price change in both 2003 and 2004 so that the sample represents a single six month summer season.[38]

3.3 Possible implications of the selection problems

This project's difference-in-difference approach compares the electric-use trajectory of control and treatment customers. Selection problems can mean that we are constructing the wrong counterfactual by using a treatment group that would have followed a different consumption trajectory as the heat of the summer arrived than the control group even in the absence of dynamic pricing. Consumers who air condition less than average are likely to use less than the average amount of power during weekday afternoons and thus to save money on dynamic pricing, giving them reason to opt in. If customers who use less air conditioning at all temperatures were more likely to opt in to the treatment group, then selection bias could misattribute their lower use at high temperatures to CPP and overstate its benefits. If they achieve their savings at low temperatures and consume in a pretty typical way at high temperatures, then selection bias might lead us to understate the benefits of CPP. Specifically:

- The people who opt in to CPP could be simply less sensitive to weather during weekday afternoons. The analysis uses June behavior as a baseline and measures impacts from hotter-on-average treatment months. If June data did not contain enough variation to capture this difference in sensitivities, then the model would overstate the impact of the change in prices.

- Someone with a larger house who kept their thermostat set to 85^oF on summer weekday afternoons even before dynamic pricing may have summer 2002 average-daily-use observables comparable to someone with a smaller house who keeps their thermostat at 75^oF, all else equal. The customer with a higher thermostat setting and larger

[37]Weighting is necessary because the sample:

1. Undersamples low use single family customers, while oversampling high use single family customers and sampling apartments at roughly their population proportion.
2. Oversamples hot climate zones.
3. Includes August, September, and July 4-31 in both 2003 and 2004, but just May 2004, June 2004, early July 1-3 2004, and October 2003. PG&E sent usable August data for 2004 but not 2003.
4. Includes all 24 summer critical days that can be called during two summers on this 12 event per summer rate. I interpret this as being two complete years worth of data. There are no May or June critical days in either year. There are, however, three October critical days in the first year.
5. Has a sample that changes over time. Subjects come and leave.

See Charles River Associates [c, 22-32] for details.

[38]Robustness checks show that weighting does not substantially affect the regression coefficients.

house is likely to become more temperature sensitive when the outdoor temperature goes above $85^o F$, which may bias CPP impact estimates downward. In other words, if someone opted in to CPP because they use little power on relatively cool summer afternoons and our pretreatment measures capture this, then our pretreatment controls will incorrectly expect them to be misers on hot summer afternoons even in the absence of CPP and thus understate CPP impacts on hot days.

4 Results: Factors that Determine the size of Dynamic Pricing's impact

Tables 5, 6, and 7 report results from specifications based on the econometric approach described in section 3. The results in tables 6 and 5 measure temperature in cooling degree hours and cooling degree hours squared and force the relationship between impact and temperature to be the same for customers from every region. Section 5 reports the results of an improved set of regressions that fit a piecewise-linear spline to the relationship between temperature and energy use and use an interaction term to generate impact estimates for zones 1 and 2 separately from zones 3 and 4.

The regressions are as follows:

1. Specification 1 uses just billing, geographic, and weather data to predict response. Specification 5 is nearly identical, except that it uses splines as described above.

2. Specification 2 expands specification 1 by controlling for factors like the presence of central or room air conditioning, the number of people in the household, and the number of bedrooms (a proxy for house size). The survey that provides these variable is unavailable or incomplete for some customers, which reduces the sample size. Specification 6 is nearly identical, except that it uses splines as described above.

3. Specification 3 adds interactions between temperature and whether the customer has central air conditioning to specification 2.

4. Specification 4 adds many survey variables interacted with the price period indicator variables, duration of participation category indicators, and person fixed effects to specification 2. It does not, however, include the air conditioning-cooling degree hour interaction terms from specification 3. This specification should should raise a red flag if omitted variable bias drove the results above or if the results from the first three cross-section specifications were not true in a fixed effects panel model. Specification 8 is nearly identical, except that it uses splines as described above.

The difference-in-difference approach estimates impacts of dynamic pricing as the coefficients on the interactions between the characteristics of the customer-day and the dummy variables reporting that either the peak price or critical prices was in effect for that customer on that day. Hence tables 6 and 5 report just these interaction terms. Every coefficient described in this section – unless explicitly noted otherwise – is the interaction of a price dummy with a characteristic of the day or customer.

Dependent variable: consumption on non holiday weekdays in kW (kWh/h). Negative values indicate that dynamic pricing customers used less power than comparable control customers.

	Specification 1: Simplest Diff in Diff	Specification 2: Adding Survey Variables	Specification 3: Adding CAC*CDH interactions	Specification 4: Adds person FEs; controls
Critical Price in Effect	0.141 (0.097)	0.024 (0.182)	-0.024 (0.176)	0.497** (0.251)
Critical Price in Effect * day after critical price	0.052** (0.025)	0.055** (0.028)	0.037 (0.028)	0.009 (0.032)
Critical Price in Effect * electric use, kWh / day, summer 2002	-0.018*** (0.005)	-0.020*** (0.006)	-0.019*** (0.006)	-0.010 (0.007)
Critical Price in Effect * high ratio rate customer.	0.217 (0.138)	0.256* (0.153)	0.236 (0.146)	0.143 (0.108)
Critical Price in Effect * cooling degree hours	0.010*** (0.004)	0.007* (0.004)	0.007 (0.007)	0.009** (0.005)
Critical Price in Effect * cooling degree hours squared (1000's),	-0.110*** (0.038)	-0.065 (0.041)	-0.074 (0.161)	-0.107** (0.044)
Critical Price in Effect * central AC	. .	-0.218* (0.114)	-0.143 (0.123)	-0.219* (0.129)
Critical Price in Effect * room AC	. .	0.296** (0.124)	0.287** (0.132)	-0.114 (0.162)
Critical Price in Effect * cooling degree hours * central AC	. .	. .	0.000017 (0.003)	. .
Critical Price in Effect * CDH squared * central AC	. .	. .	0.0000065 (0.00093)	. .
Critical Price in Effect * swimming pool	. .	. .	. .	-0.289 (0.196)
Critical Price in Effect * cooling degree hours * room AC	. .	. .	. .	0.010*** (0.003)
Critical Price in Effect * # kids under 5 in household	. .	. .	. .	-0.221** (0.093)
Critical Price in Effect * # people over 65 in household	. .	. .	. .	-0.223*** (0.085)
Critical Price in Effect * customer stayed in expt. throughout expt.	. .	. .	. .	-0.352** (0.138)
N	121408	101981	101981	77660
R-squared	0.4915	0.5020	0.5196	0.6380
Robust standard errors, clustered by customer in parentheses. Significance: *=10% ** =5% ***=1% Abbreviations: AC: air conditioning. CAC: central air conditioning. FEs: fixed effects.				
Cooling degree hours (CDH) are 2-7PM, base 78^o F. Heating degree hours are base 65^o F.				

Table 5: The impact of critical prices

Dependent variable: consumption on non holiday weekdays in kW (kWh/h). Negative values indicate that dynamic pricing customers used less power than comparable control customers.

	Specification 1: Simplest Diff in Diff	Specification 2: Adding Survey Variables	Specification 3: Adding CAC*CDH interactions	Specification 4: Adds person FEs; controls
TOU Peak Price in Effect	0.067 (0.077)	-0.107 (0.139)	-0.100 (0.136)	0.268 (0.197)
TOU peak price in effect * day after critical price	0.024* (0.013)	0.030** (0.014)	0.030** (0.014)	0.017 (0.014)
TOU Peak Price in Effect * electric use, kWh / day, summer 2002	-0.004 (0.004)	-0.006 (0.005)	-0.005 (0.005)	0.005 (0.005)
TOU Peak Price in Effect * high ratio rate customer.	-0.011 (0.039)	-0.015 (0.043)	-0.010 (0.044)	0.018 (0.054)
TOU Peak Price in Effect * cooling degree hours	0.010*** (0.003)	0.009** (0.004)	0.006 (0.006)	0.009** (0.004)
TOU Peak Price in Effect * cooling degree hours squared (1000's),	-0.102*** (0.033)	-0.078** (0.035)	-0.013 (0.135)	-0.106*** (0.038)
TOU Peak Price in Effect * central AC	. .	-0.033 (0.079)	-0.014 (0.081)	-0.031 (0.086)
TOU Peak Price in Effect * room AC	. .	0.110 (0.084)	0.118 (0.085)	-0.086 (0.107)
TOU Peak Price in Effect * cooling degree hours * central AC	. .	. .	-0.000022 (0.002)	. .
TOU Peak Price in Effect * cooling degree hours squared * central AC	. .	. .	-0.00044 (0.00084)	. .
TOU Peak Price in Effect * swimming pool	. .	. .	. .	-0.279* (0.148)
TOU Peak Price in Effect * cooling degree hours * room AC	. .	. .	. .	0.010*** (0.003)
TOU Peak Price in Effect * # kids under 5 in household	. .	. .	. .	-0.106 (0.072)
TOU Peak Price in Effect * # people over 65 in household	. .	. .	. .	-0.117** (0.057)
TOU Peak Price in Effect * customer stayed in expt. throughout expt.	. .	. .	. .	-0.193** (0.089)
N	121408	101981	101981	77660
R-squared	0.4915	0.5020	0.5196	0.6380
Robust standard errors, clustered by customer in parentheses. Significance: *=10% ** =5% ***=1% Abbreviations: AC: air conditioning. CAC: central air conditioning. FEs: fixed effects.				
Cooling degree hours (CDH) are 2-7PM, base 78^o F. Heating degree hours are base 65^o F.				

Table 6: The impact of peak prices

4.1 The Benefits of Dynamic Pricing Grow as the Temperature Increases

The analysis shows that dynamic pricing begins to deliver significant energy use reduction benefits somewhere between 85 and 100 degrees Fahrenheit. This is fortuitous, if perhaps unsurprising, because dynamic pricing aspires to dampen energy use, especially air conditioning use, during hours when air conditioning demand makes electricity scarce and expensive. Energy-intensive air conditioning causes demand peaks on the hottest summer weekday afternoons; and allows dynamic pricing to have a larger effect on the quantity of electricity consumed on the hottest days. Thus, CPP saves the greatest quantity of power when those reductions are particularly valuable.

Graphs 5 and 6 suggest that the average use by control and treatment customers are difficult to distinguish on days with fewer than 50 to 60 base-78^o Fahrenheit cooling degree hours (CDH), but diverge at higher temperatures. The kernel and spline regression results in figures 7 through 11 corroborate this and report that benefits grow rapidly in zones 3 and 4 as the temperature rises from roughly 90^oF to about 98^oF, before leveling off or even beginning to shrink slightly. This pattern is consistent with customers increasing their thermostat settings significantly and perhaps with CPP customers precooling their homes before the peak prices go into effect.

The quadratic estimate yields two coefficients, reporting the impact of changes to the number of CDH and of the number of CDH^2. The linear term dominates on cooler days when there are small numbers of CDH, but the number of CDH^2 explodes as the temperature rises. Specifications 1, 2, and 4 show that CDH squared has a negative and statistically significant impact for customers experiencing peak prices. The linear CDH term is small, positive in every specification, and often statistically significant. For example, specification 1 reports that, relative to the time-invariant control group, dynamic pricing customers increase use by .01 kW (SE: .003) for every increase of 1 CDH and decrease use by .102kW (SE: .033) for every increase of 1000 CDH^2.

The point estimates of the relationship between temperature and the impact of critical prices are qualitatively quite similar, but the critical price impacts are only statistically significant in specifications 1 and 4. It is unsurprising that the critical price's effect is less precisely measured because the summer data set contains 24 critical days, but more than 170 peak days.

5 More Flexible Estimates of the Relationship Between Temperature, Climate Zone, and Response

Regression specifications 1 through 4 reported above model response to dynamic pricing as a quadratic function of temperature. There are troubling signs that 1) rigidities in this functional form drives some results and 2) that the relationship differs between cool and hot regions. This section uses more flexible estimation techniques to provide direct evidence about that hypothesis. This section uses two techniques. All of its estimates begin by using splines to estimate energy use as a piecewise linear function of temperature. The main estimates shown in Section 7 and figures 10 and 11 simply substitute the piecewise linear function of temperature for the quadratic function of temperature used above. Section 5.1 describes the simpler approach taken for the other figures in this section that make non-parametric kernel estimates of the difference between the control and treatment group's temperature to energy use relationships.

The regressions here need work before they are complete. Most notably, figure 10 and 11 need to display the estimates' standard errors.

5.1 The Kernel Estimates

The kernel estimate in figures 7, 8, and 9 make estimates as follows:

1. They analyze just the treatment period. This approach compares the treatment and control groups without attempting to control for any preexisting differences that remain after controlling for observable characteristics. The model here estimates a difference, not a difference-in-difference. This section takes a conventional regression approach except that it makes nonparametric estimates of the impact of the interaction between temperature and treatment status.

2. It runs a very simple model of the relationship between customer-day characteristics and electricity use in the control group. Roughly, I take the variables from specification 2[39] and simplify the section 3 regression[40] to be: $avgLoad_{it} = \alpha^T \mathbf{X}^* + \gamma^T + \epsilon_{it}$. I replace the quadratic function of temperature with a piecewise linear spline. The spline creates variables of the form: $SplineCDH_{k,t} = \max\{0, CDH_t - K\}$ where knot location $K \in \{0, 20, 40, 60, 70, 80, 90, 100, 110, 120\}$ for hot climate zones 3 and 4 and $K \in \{0, 20, 40, 60, 70, 80\}$ for temperate climate zones 1 and 2. I use no interactions between calendar days and temperature.

3. It confirms that these models are flexible enough to capture the shape of the temperature-driven changes in energy consumption by fitting a lowess, non-parametric, kernel estimator to the relationship between the temperature and the residuals from those models. Lowess estimators are, in essence, a sophisticated way of calculating a moving local average and a local slope. Here, the local average considers the closest 15% of the data. If the spline model fits well, then the lowess local average of the control group residuals will stay close to zero everywhere. The set of splines above is chosen to correct some disturbing regional deviations from zero that appear with smaller sets of splines. Specifically, this set of splines has 20 CDH between knots below 60 CDH, and 10 CDH between knots in the more interesting region above 60 CDH. Sixty base

[39] Like specification 2, the dependent variable is average 2-7PM weekday load in kW. The control variables are: average daily electricity use, Summer 2002, kWh; number of people in the household; and afternoon heating degree hours. The regressions also control for dummies for day before a critical day; day after a critical day; apartment; being the hotter climate zone of the zone 1-2 or 3-4 pair; having central air conditioning; having room air conditioning; day of week; month; and year. These estimates are run separately for the hot and cool climate zones, which is equivalent to interacting a hot climate zone dummy with every control variable.

[40] Stata's lowess kernel estimator does not support the kind of weights used in the difference-in-difference estimators, so the regressions underlying figures figures 7, 8, and 9 are unweighted.

78^oF CDH between 2 and 7PM roughly corresponds to a 2-7PM average temperature of 90^oF.

4. The analysis then uses the control group's coefficients to make out-of-sample predictions of treatment group consumption for each customer-afternoon observation. It then makes kernel estimates of the mean residual at each temperature level. They show that the treatment group used less on average than the control group, creating negative residuals. If, conditional on the characteristics controlled for above, the control and treatment groups have no preexisting differences, then the resulting graphs will show the average impact of dynamic pricing at each temperature.

5.2 The Difference-in-Difference Spline Estimates

The difference-in-difference spline estimates are an extension of the general econometric model described in Section 3.[41]

The kernel estimates include knots at 110 and 120 CDH, but the difference-in-difference framework leaves them out. Including them would be dicey because it would require estimating the interaction between being a treatment customer and there being more than 120 CDH in an afternoon off from only about 30 pretreatment customer-days with temperatures above 120 which have a maximum of 123.5 CDH. The available June data between 110 and 120 is similarly thin.

All of the high temperature difference-in-difference estimates yield suspect results. They yield large magnitude but statistically insignificant point estimates of the treatment-control difference in pretreatment sensitivity to weather at many of these line segments. The estimation further reports that the "treatment effect" almost exactly negates these coefficients.

Figures 10 and 11 visualize this issue by reporting both difference-in-difference lines and "difference" lines that add the treatment customers' pretreatment coefficient on temperature to the coefficient on the new price being in effect. In the notation of Section 3, the difference lines report $(\hat{\delta}+\hat{\beta})\mathbf{X}^*$ for peak impacts and $(\hat{\delta}+\hat{\psi})\mathbf{X}^*$ for the impact of the critical price where $\hat{\delta}$ is the preexisting control-treatment difference. By contrast, the difference-in-difference estimates are $\hat{\beta}\mathbf{X}^*$ and $\hat{\psi}\mathbf{X}^*$ respectively. The "difference" estimators yield more intuitively appealing results, especially in conditions that were fairly hot for June.

5.3 Results

This analysis makes several tentative findings:

- It appears that the benefits of dynamic pricing increase in temperature in the 90's, but stop growing in temperature at higher temperatures and may even shrink when temperatures become extreme. This result is consistent with plausible stories about

[41] The quadratic regressions (Specifications 1-4) include controls that interact dummies for day of week, calendar month, and year with quadratics of cooling and heating degree hours. The piecewise regressions do not interact these dummies with the temperature splines out of concern that adding more than 100 temperature-interaction regressors would introduce too many collinear variables and introduce too many opportunities to pick up spurious results.

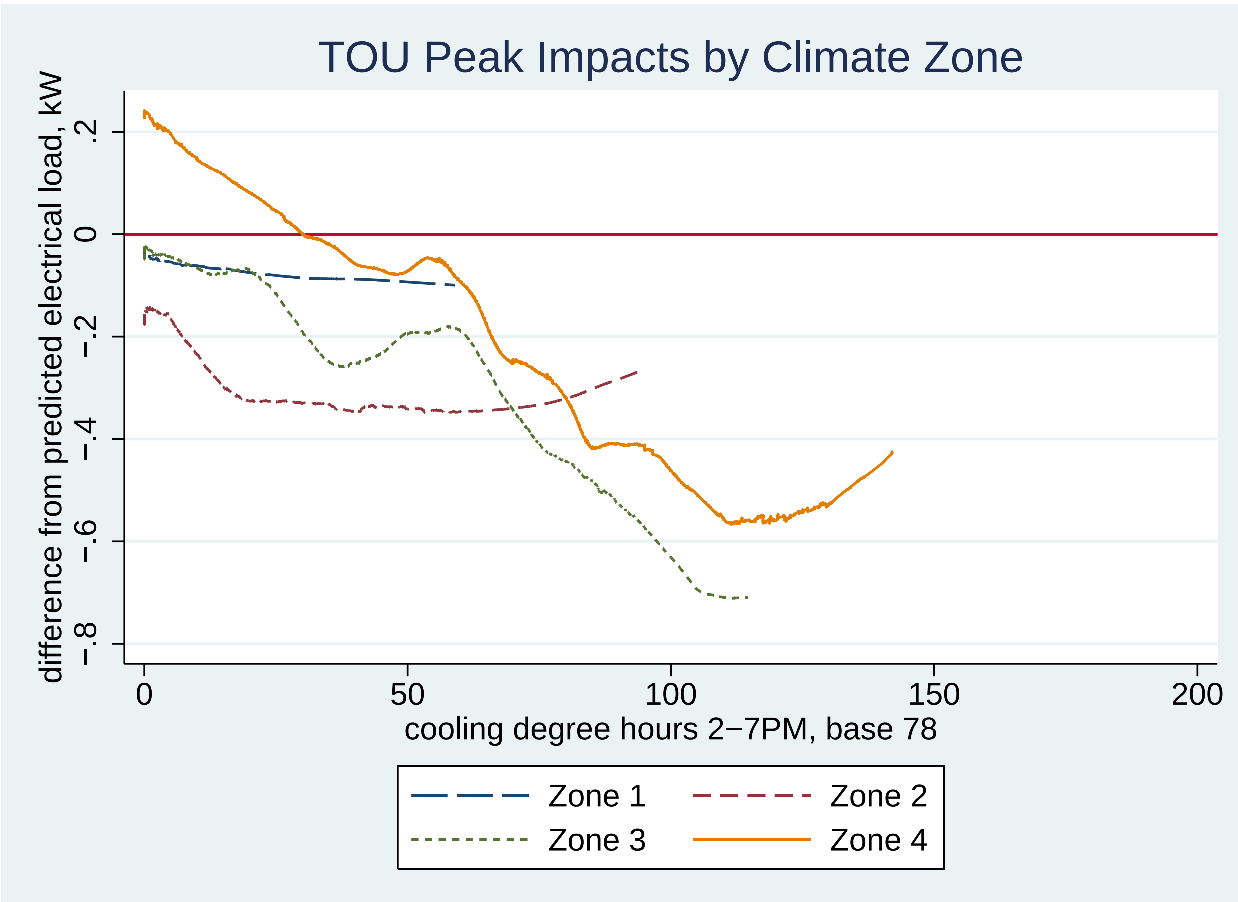

Figure 7: *Kernel estimates of the relationship between the impact of peak prices and temperature by climate zone. Plotted for temperatures up to the 99th percentile of the customer-day temperatures distribution for each climate zone. There are many observations with the lowest temperature, 0 cooling-degree hours, but a long thin tail of high temperature observations estimated off of unusually hot days in unusually hot places. (This project has a strong substantive interest in understanding program impacts on days when broad areas are quite hot without claiming to have made meaningful estimates for a wide temperature range based on a tiny handful of data points from unusual climates. A future revision might plot impacts based on the average temperature for each climate zone – which should yield a better compromise and escape the need to drop observations. Full temperature range graphs and disaggregated graphs available upon request.)*

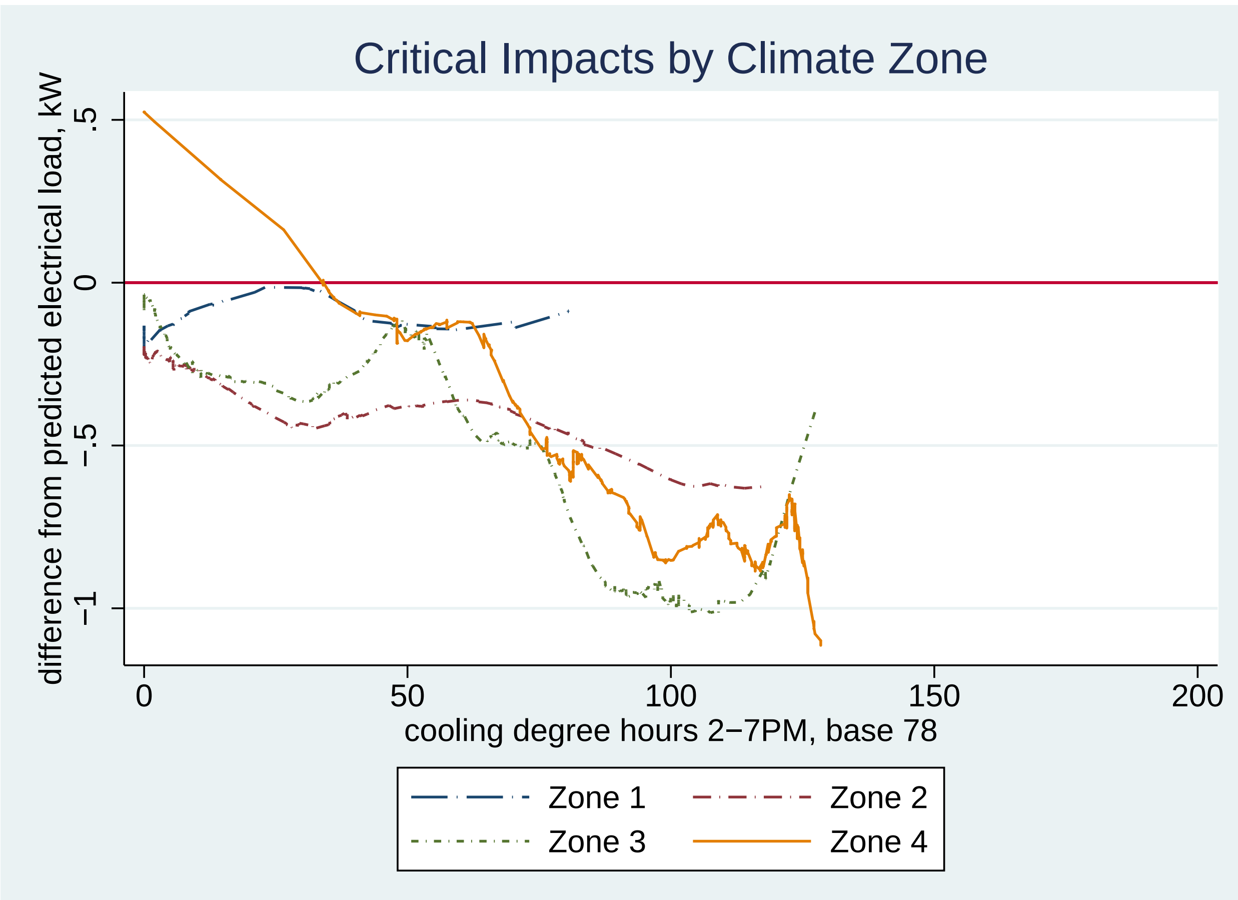

Figure 8: *Kernel estimates of the relationship between the impact of critical prices and temperature by climate zone. Plotted for temperatures below the 99th percentile of the observed customer-day temperatures distribution for each climate zone.*

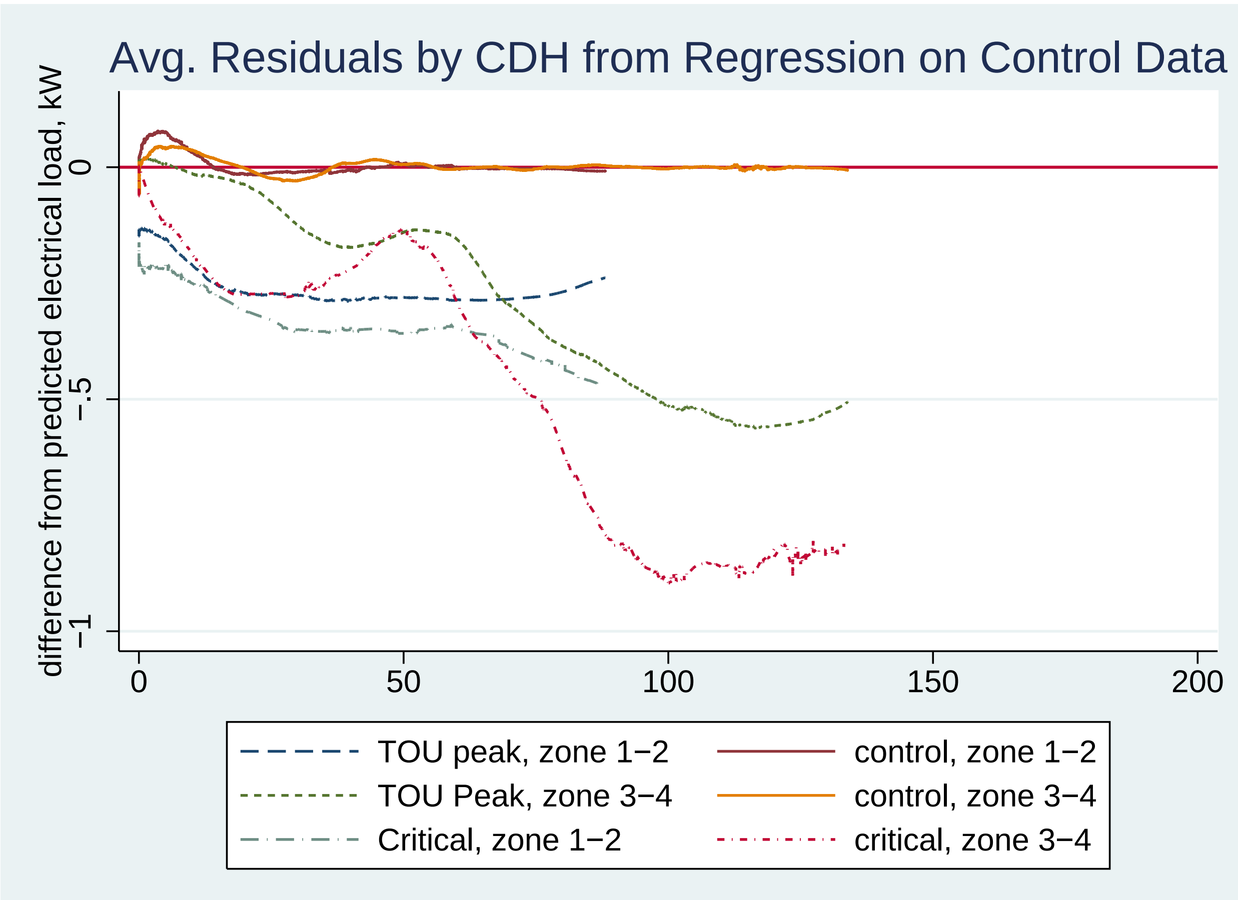

Figure 9: *Nonparametric estimates of the impacts of dynamic pricing, by climate zone and the price that is in effect. The two control group lines stay quite close to zero everywhere, suggesting that the functional form captures the average temperature-driven variation in the control group's electric use. The other four lines approximate the impacts of critical and peak prices by temperature. The results are plotted for temperatures less than the 99th percentile of customer-days from the applicable climate zones.*

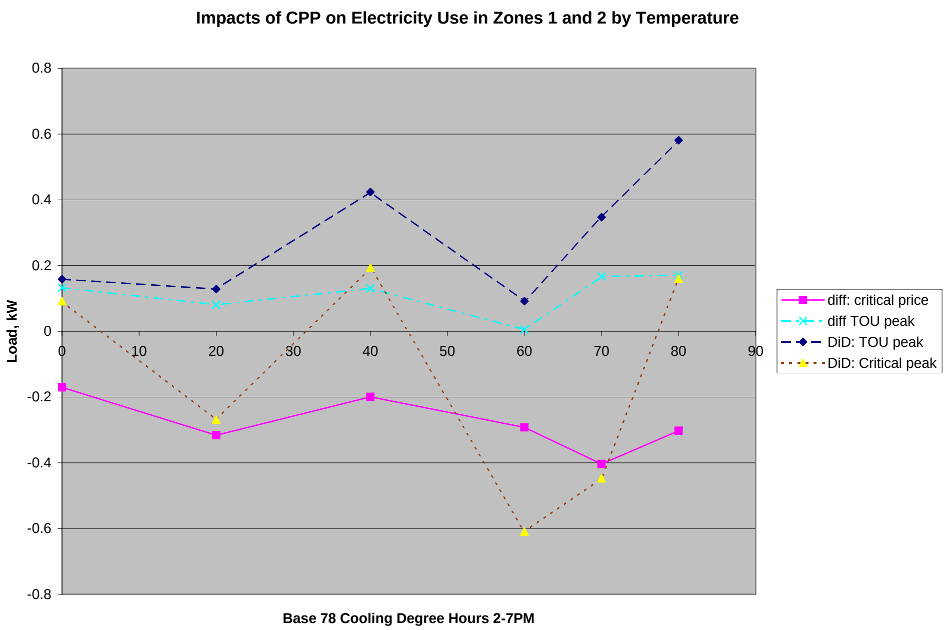

Figure 10: Splines fit to the temperature-energy use relationship for temperate climate zones 1 and 2. The Difference-in-Difference (DiD) lines measure impact relative to the "preexisting differences" measured from behavior during the quite brief, relatively cool pretreatment period. The difference lines show the sum of the impact and treatment customer interaction terms that identify control-treatment differences during the pretreatment and treatment periods respectively. The intercept in this graph comes from the average statewide customer characteristics, except that the weights on the zone 1 and 2 dummies have been scaled so that they sum to 100%, while the weight on the zone 3 dummy is set to zero. This is clearly a flawed approach. Future revisions will use customer characteristics conditional on being in zones 1 and 2 and will display the standard errors.

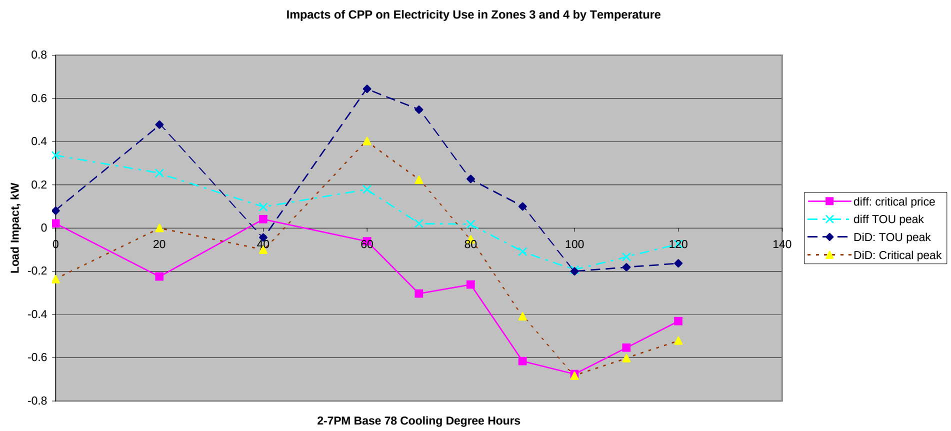

Figure 11: Splines fit to the temperature-energy use relationship for temperate climate zones 3 and 4. This uses the same techniques and conventions as Graph 10. The Difference-in-Difference regression finds that the treatment group used far less power than the control group during (rare) very hot conditions in the pretreatment period. These are large magnitude but quite imprecisely estimated effects. Then we find that during the treatment period, this difference all but disappears. This suggests that the finding that CPP is counterproductive during cool temperature conditions is, most likely, spurious.

how consumers' air conditioner operation changes as a function of temperature and price.

- CPP impacts appear to be particularly sensitive to temperatures between roughly 70 and 100 to 110 CDH (i.e. the range between an average afternoon temperature of 92 and $100^o F$).

- There are important differences in the relationship between the cooler (1-2) and hotter (3-4) zones, but more modest differences between zones 1 and 2 and between 3 and 4. Thus, zone dummies will capture the difference between zones 1 and 2 reasonably well.

- Impacts in zone 1 and, to a lesser extent, zone 2 are fairly temperature insensitive.

- The difference estimators make very few point estimates suggesting that dynamic pricing was counterproductive and raised energy use.[42] Plotting the best fit quadratic relationship on these difference estimates reveals a very small region in which the point estimates have the "wrong" sign. Much of the counter productivity finding comes from fitting the difference-in-difference estimator's controls for preexisting differences to thin and idiosyncratic high temperature, pretreatment data. The difference-in-difference estimates found that treatment customers used, on average, far less power during moderately high temperature pretreatment afternoons than did the control group. The moderately high temperature pretreatment data that drove these estimates come from a few weather stations on a few days for the subset of customers who had meters activated, which is likely to have created idiosyncratic correlations. The spline-based point estimates showing dynamic pricing to be counterproductive almost exactly negate strange, imprecisely estimated pretreatment relationships.

[42] The exception involves the impact of critical pricing on consumption in desert zone 4 on very cool days. These customer-days are in the first percentile of all zone 4 summer, critical customer-days – and should be considered imprecisely estimated.

5.4 Regressions with Splines

	Specification 5: Simplest Diff in Diff	Specification 6: Adding Survey Variables	Specification 8: Adds person FEs; controls
Critical Price in Effect	-0.117 (0.256)	-0.146 (0.321)	0.419 (0.333)
Critical Price in Effect * day before critical price	0.084*** (0.024)	0.082*** (0.027)	0.047 (0.032)
Critical Price in Effect * day after critical price	0.019 (0.025)	0.023 (0.028)	-0.007 (0.032)
Critical Price in Effect * electric use, kWh / day, summer 2002	-0.018*** (0.005)	-0.020*** (0.006)	-0.011 (0.007)
Critical Price in Effect * high ratio rate customer.	0.227* (0.136)	0.255* (0.152)	0.125 (0.109)
Critical Price in Effect * apartment	-0.016 (0.091)	0.020 (0.125)	-0.020 (0.146)

	Specification 5: Simplest Diff in Diff	Specification 6: Adding Survey Variables	Specification 8: Adds person FEs; controls
Critical Price in Effect * climate zone 1	0.274 (0.233)	0.176 (0.263)	0.072 (0.264)
Critical Price in Effect * climate zone 2	0.264 (0.230)	0.135 (0.257)	0.014 (0.244)
Critical Price in Effect * climate zone 3	0.221 (0.205)	0.173 (0.221)	0.076 (0.217)
Critical Price in Effect * cooling degree hours	0.012 (0.010)	0.012 (0.011)	0.016 (0.011)
Critical Price in Effect * impact of a CDH beyond 20	-0.017 (0.019)	-0.016 (0.021)	-0.032 (0.024)
Critical Price in Effect * impact of a CDH beyond 40	0.030 (0.024)	0.025 (0.026)	0.038 (0.028)
Critical Price in Effect * impact of a CDH beyond 60	-0.043 (0.046)	-0.027 (0.049)	-0.006 (0.049)
Critical Price in Effect * impact of a CDH beyond 70	-0.010 (0.065)	-0.012 (0.065)	-0.069 (0.074)
Critical Price in Effect * impact of a CDH beyond 80	-0.008 (0.061)	-0.026 (0.061)	0.035 (0.071)
Critical Price in Effect * impact of a CDH beyond 90	0.008 (0.059)	0.048 (0.059)	-0.020 (0.060)
Critical Price in Effect * impact of a CDH beyond 100	0.036 (0.046)	-0.002 (0.046)	0.047 (0.045)
Critical Price in Effect * CDH * zone is 1 or 2	-0.018 (0.012)	-0.024* (0.013)	-0.025* (0.014)
Critical Price in Effect * zone 1 or 2 * impact of a CDH beyond 20	0.041* (0.025)	0.052* (0.026)	0.079** (0.031)
Critical Price in Effect * zone 1 or 2 * impact of a CDH beyond 40	-0.063* (0.038)	-0.068 (0.042)	-0.130*** (0.050)
Critical Price in Effect * zone 1 or 2 * impact of a CDH beyond 60	0.056 (0.102)	0.029 (0.113)	0.126 (0.126)
Critical Price in Effect * zone 1 or 2 * impact of a CDH beyond 70	0.044 (0.109)	0.085 (0.117)	0.052 (0.123)
Critical Price in Effect * heating degree hours	0.007 (0.007)	0.003 (0.010)	-0.029* (0.017)
Critical Price in Effect * central AC	. .	-0.224** (0.114)	-0.252* (0.130)
Critical Price in Effect * room AC	.	0.250**	-0.145

	Specification 5: Simplest Diff in Diff	Specification 6: Adding Survey Variables	Specification 8: Adds person FEs; controls
	.	(0.119)	(0.170)
Critical Price in Effect * number of bedrooms	. .	0.040 (0.059)	0.036 (0.058)
Critical Price in Effect * # people in the household	. .	0.033 (0.027)	0.085* (0.047)
Critical Price in Effect * heating degree hours squared (1000's)	. .	. .	1.132 (0.692)
Critical Price in Effect * cooling degree hours, previous day	. .	. .	-0.002 (0.001)
Critical Price in Effect * cooling degree hours, two days before	. .	. .	0.001 (0.001)
Critical Price in Effect * cooling degree hours, three days before	. .	. .	-0.002 (0.001)
Critical Price in Effect * swimming pool	. .	. .	-0.259 (0.193)
Critical Price in Effect * # kids under 5 in household	. .	. .	-0.225** (0.092)
Critical Price in Effect * # people over 65 in household	. .	. .	-0.230*** (0.087)
Critical Price in Effect * electric cooktop	. .	. .	0.352* (0.191)
Critical Price in Effect * customer stayed in expt. throughout expt.	. .	. .	-0.328** (0.135)
Critical Price in Effect * cooling degree hours * room AC	. .	. .	0.009*** (0.003)
Critical Price in Effect * heating degree hours *electric heat	. .	. .	0.046** (0.021)
Critical Price in Effect * electric heat	. .	. .	-0.167 (0.138)
Critical Price in Effect * # kids over 5 in household	. .	. .	-0.055 (0.066)
Critical Price in Effect * work from home 0-10 hrs/wk	. .	. .	-0.020 (0.165)
Critical Price in Effect * electric oven	. .	. .	-0.218 (0.180)
Critical Price in Effect * number of refrigerators and freezers	. .	. .	-0.252** (0.103)

	Specification 5: Simplest Diff in Diff	Specification 6: Adding Survey Variables	Specification 8: Adds person FEs; controls
Critical Price in Effect * work from home 11-30 hrs/wk	. .	. .	-0.021 (0.161)
Critical Price in Effect * work from home >30 hrs/wk	. .	. .	-0.242 (0.280)
Critical Price in Effect * spa	. .	. .	0.104 (0.188)
Critical Price in Effect * customer stayed in expt. < 4.5 months	. .	. .	0.135 (0.161)

Robust standard errors, clustered by customer in parentheses.
Significance: *=10% ** =5% ***=1%
Abbreviations: AC: air conditioning. CAC: central air conditioning. FEs: fixed effects.
Cooling degree hours (CDH) are 2-7PM, base 78^o F. Heating degree hours are base 65^o F.
Splines for impact of a CDH beyond K are defined as: $SplineCDH_k = \max\{0, CDH - K\}$

Table 7: Spline estimates of impact of peak and critical prices. Dependent variable: 2-7 PM kW consumption on non-holiday weekdays. Negative values indicate that CPP customers used less power than comparable control customers.

6 Other results from the analysis

6.1 Controlling for the Interaction of Air Conditioning and Heat

Specification 3 adds interactions between central air conditioning ownership, CDH, and CDH^2. This specification has several unexpected results:

- The interactions between central air conditioning ownership, CDH, and dynamic pricing are fairly precisely estimated zeros. We fail to reject the joint hypothesis that the impact of central air conditioning ownership, and the interaction of central air conditioning ownership with CDH and CDH^2 are all zero (p=0.96).

- The standard error on the interactions between cooling degree hours squared with the peak and critical prices both increase by a factor of four, rendering their coefficients statistically insignificant.[43] We can, however, reject the hypothesis that the base effects of CDH and CDH^2 are both zero (p=0.04).

- The point estimate of the interaction of the peak price with cooling degree hours squared falls significantly toward zero.

Table 8 suggests that these unexpected results might come from the fact that we are identifying the effect of dynamic pricing on days when temperatures are above 90^o from

[43] Adding more control variables did not restore statistical significance, nor did adding person fixed effects.

Percentage of customers who have central air conditioning					
	Climate zone				
House type	1	2	3	4	all
Apartments	3.3	24.6	57.6	80.6	32.5
High use single family	3.8	49.2	82.2	97.2	62.7
Low use single family	5.0	24.4	71.5	63.5	41.6
All	4.1	29.7	71.4	77.0	
Percentage with any cooling technology					
All	12.8	47.5	88.4	94.7	

Table 8: Air conditioning and cooling technology by climate zone and house type. More than 70% of climate zones 3 and 4 customers have central air conditioning and more than 88% of them have some kind of cooling technology, with the vast majority of those paying for the electricity for either central or room air conditioners. Almost all of the air-conditioning-intensive, very hot weather takes place in those zones.

a population that overwhelmingly has air conditioning, typically, central air conditioning. Nearly all of the days with the highest CDH^2 values, which matched to the highest dynamic pricing impacts came from climate zones 3 and 4. More than 70% of those customers have central air conditioning and more than 88% of them have some kind of cooling technology.

These results may also suggest that responses to dynamic pricing are sensitive on the details of the air conditioner and house, but the survey collects no data about the size, age, condition, or efficiency of air conditioners and insulation with which to test that hypothesis.

6.2 Accumulated Heat

Specification 4 controls for the number of cooling degree hours on each of the three previous days. The amount of energy required to cool buildings is a function of both instant the ambient outdoor air temperature and the temperature from previous periods that has accumulated in building components like the walls and roof. Heat accumulates The analysis finds that dynamic pricing customers use less power in sustained heat than do customers on time invariant rates. These savings of .001 kW per lagged CDH can be substantively quite important, because customers experiencing a Central Valley “heat storm” may have experienced 100 or more cooling degree hours on each of the three previous days.

6.3 The Effect of Customer Size on Response to Dynamic Pricing

Billing data measures each customer’s aggregate use the summer before the experiment, yielding a variable that does not disaggregate weekday afternoon consumption, but is uncontaminated by the experiment. The simple regression specifications in 1, 2, and 3 find that bigger customers respond more to critical price signals. The 75th-percentile customer used 12.7 kWh/day more than the 25th percentile customer in Summer 2002, and the regression coefficients imply that this increase translates to a savings of roughly 1.2kWh per critical afternoon (or .24 kW). Specification 4 adds many more controls and reduces the point estimate by about half, which renders the impact statistically insignificant.[44] We get smaller, statistically insignificant point estimates when the peak price is in effect.

Thus, utilities can use the overall consumption from readily available billing data to identify and target customers likely to respond even if they do not have detailed appliance holding data.

6.4 Other Dynamic Pricing Impacts

- On critical days people with central air conditioners conserve more than do comparison customers who do not pay for the electricity for compressor-based air conditioning.[45] The comparison groups includes people who have no air conditioning, those who have evaporative cooling, and those who have building-wide air conditioning provided as part of their rent. Further, people with room air conditioning responded less to both peak and critical price signals as the number of cooling degree hours rose.

- Customers use more power on a day with ordinary, peak prices after a critical event which is consistent with delaying optional activities like drying clothes and running the dishwasher.[46],[47],[48]

- There is evidence that people with swimming pools reduced peak electricity use more in response to dynamic pricing than other customers did. Swimming pool pumps use considerable energy and it is easy to set their timers to run them off peak. The point estimates are substantively quite big at about .28 kW or 1.5 kWh per day. This impact is imprecisely estimated, but is statistically significant with a similar coefficient in a specification that adds the variables considered for this section to specification 2. That specification also corroborates the next two findings.

- Households responded more for each member above the age of 65 than for household members between the ages of 5 and 64. If we control for income, the sign remained the same and the relationship remained statistically significant on critical days but had a p-value of 0.14 on ordinary days.[49]

---

[44] The point estimate and standard error are unaffected by the addition of removal of the person fixed effects used in the reported regression specification 4. In other words, all the action comes from adding more control variables.

[45] There is some evidence at the p=.05 level that people with room air conditioning conserve less in response to critical price signals, but the sign reverses as we add more controls and person fixed effects. The sign reversal suggests that we may be picking up impacts of things like building age, building size, or number of occupants.

[46] There is some evidence that critical days after prior critical days experience a similar rebound, but these findings get weaker as we add more controls and as we subdivide between high and low use customers as a robustness check

[47] It is worth noting, however, that most regressions find that customers use more power on a critical day before a second critical day. This goes away in the results in Section 5, which suggests that it came from an insufficiently flexible functional form for the temperature relationship.

[48] Herter [2006a] makes a similar finding.

[49] About 40% of households with at least one senior reported being in the lowest income category, while only 26% of other households did. However, the income data that we have may be a poor measure of retirees' true spending power

- Customers who stayed in the experiment to the end responded roughly 1 (1.8) kWh / day more than those who left early during hours when the peak (critical) price was in effect.

6.5 Comparing the Effects of Different Dynamic Rates

The high ratio rate charged higher prices, but most regressions find a statistically and substantively insignificant but imprecisely estimated difference between the impacts of high and low ratio rates.[50] The regression results in tables 9 and 10 find modest evidence that high-ratio, high-use single family customers responded more to the peak price than did similar low-ratio customers.

The customers in this sample got automated phone calls the day before each critical price went into effect, which may have increased customer response to the high price. This design feature means that the SPP dataset is a good source of evidence about how customers react to a CPP program with telephone notification, but a poor source of evidence about how customers would react to price changes alone.

6.6 Robustness to Selection Bias

Specifications 1, 2, and 3 find that treatment customers were using less peak power during the pretreatment period by about .07 kW which is significant at the 4-6% level in each of the specifications. These differences are not particularly disturbing if they reflect preexisting differences. If these savings come from premature response, then the results reported here understate the true value of dynamic pricing. Specification 4 uses customer fixed effects instead of identifying coefficients for preexisting differences between the control and treatment groups.

6.6.1 Evidence from Splitting the Dataset

There is evidence that the whole sample suffers from a selection bias problem among its high-use, single family customers. One way to explore whether selection bias is driving the results is to divide the sample into the suspect high use, single family customers and the more pristine low-use single family and apartment customers. Tables 9 and 10 [51] report results from taking this approach.

The results from the whole sample, just apartments / low use single family, and just high-use, single family are all qualitatively quite similar to each other. In almost all cases, statistically significant point estimates above retain their signs and magnitude in the split

[50] We make this finding during critical periods via a slightly circuitous route. The regressions find that high ratio customers use significantly less power than low ratio customers during critical periods that they did not participate in. This typically happens when they could not be successfully reached by the automated telephone notification system. The CPP group included roughly 0.05% of the state's 8.3 residential utility accounts, so customers were unlikely to learn of the event through a channel other than the direct notification. They return to the CPP group average when they are successfully called. This is probably a statistical artifact despite being nearly statistically significance at the 10% level.

[51] The results in Tables 9 and 10 are chosen to support the discussion in section 4.1. Letzler [2007] reports complete results from running all four specifications.

sample estimates. The results are less precisely estimated, which is not surprising given that we have split the sample into high use (45%) and low-use / apartment (55%).[52]

7 Estimates of the Total Impact of Dynamic Pricing

The section above shows how a variety of factors affect dynamic pricing's electricity use. This section aggregates them to calculate the average impact of dynamic pricing for customers in important scenarios.

A disproportionate part of the value of dynamic pricing comes from days when electricity is scarce, creating high energy prices[53]. Electricity demand is closely correlated with temperature and most scarcity conditions take place on days when high temperatures create extreme demand.[54] Further, the results above find that the impact of dynamic pricing is quite sensitive to temperature. This section assigns days to bins by their peak 2-7 PM California Independent System Operator (CAISO) control area electric load[55], determines the population weighted average temperature in each bin, and then calculates the average impact of dynamic pricing for the temperatures from each bin. It disaggregates the top end of the load distribution, because high demand days are likely to yield the greatest benefits and will generally be of the greatest practical importance.

7.1 Calculating Total Impacts

The estimation strategy described in Section 3 creates two total-impact objects of interest:

- **The impact of the peak price** is $I_{peak} = \sum_{j \in \{1, \mathbf{X}^*\}} \beta_j \overline{x}_j$ where β_j is the coefficient on the interaction of $PeakPrice_{it}$ with the jth customer characteristic and $\overline{x}_j$ is the average value of the jth customer characteristic conditional on $PeakPrice_{it}$ being 1.[56]

- **The impact of calling a critical price** is quite similar, namely:

[52] Premature response that inadvertently "treated" CPP customers with the perception that they could save money by reducing peak use during the pretreatment period would tend to cause these estimates to understate the impacts of dynamic pricing. A useful way to bound the magnitude of this bias would be conceive of the SPP as having treated weekday afternoon hours with higher prices. We can then repeat the regressions using weekend afternoon consumption as the "untreated" period instead of early June weekdays. Control group weekend afternoon use turns out to be a very strong predictor of control group weekday afternoon use. This approach would tend to overstate the impacts of the SPP because 1) customers will shift laundry and other major appliance use from peak periods to weekend afternoons and 2) CPP treats weekend afternoons with prices slightly lower than the time invariant price.

[53] See e.g. Borenstein [2005] for an extended discussion of this

[54] Vacation patterns are also important: people who are in town use more power, but are more able to respond to telephone-based critical peak signals. Future revisions to this work could consider vacations.

[55] The absolute daily peak took place between 2 and 7PM on 88% of weekdays during the June -October 2003 and May-September 2004 experiment period. Further, in 2003 and 2004, the absolute peak took place between 2 and 7PM on every day in July, August, and the first half of September.

[56] For simplicity of discussion, I am treating the 1 as the first customer characteristic. The coefficient on 1 interacted with *PeakPrice* is the average impact of the peak price on consumption after controlling for all of the observed-customer characteristics.

	Specification 2: Survey Variables		Specification 3: CAC*CDH interactions	
	Low Use/Apt.	High Use	Low Use/Apt.	High Use
TOU Peak Price in Effect	-0.134 (0.156)	-0.183 (0.377)	-0.148 (0.154)	-0.063 (0.363)
TOU Peak Price in Effect * day after critical price	0.042*** (0.016)	0.012 (0.030)	0.041*** (0.015)	0.013 (0.030)
TOU Peak Price in Effect * elec. use, kWh / day summer '02	-0.003 (0.008)	-0.007 (0.011)	-0.002 (0.008)	-0.007 (0.011)
TOU Peak Price in Effect * high ratio rate customer.	0.034 (0.046)	-0.170* (0.097)	0.040 (0.047)	-0.165* (0.096)
TOU Peak Price in Effect * cooling degree hours 2-7pm	0.006 (0.004)	0.013** (0.006)	0.00036 (0.006)	0.020* (0.012)
TOU Pk. Price in Effect * cooling degree hrs squared (1000's)	-0.071* (0.038)	-0.092 (0.064)	0.070 (0.133)	-0.204 (0.292)
TOU Peak Price in Effect * central AC	0.030 (0.087)	-0.017 (0.175)	0.043 (0.091)	0.012 (0.171)
TOU Peak Price in Effect * room AC	0.129 (0.094)	0.170 (0.167)	0.143 (0.095)	0.156 (0.163)
TOU Peak Price in Effect * cooling degree hours * central AC	. .	. .	-0.00013 (0.003)	-0.003 (0.005)
TOU Pk Price in Effect * cooling degree hrs squared * central AC	. .	. .	-0.00089 (0.00081)	0.00067 (0.002)
N	54446	47535	54446	47535
R-squared	0.3715	0.4331	0.3964	0.4436

Robust standard errors, clustered by customer in parentheses.
Significance: *=10% ** =5% ***=1%
Cooling degree hours are base 78^o F. Heating degree hours are base 65^o F.

Table 9: The impact of peak pricing when we separate apartment / low use single family customers from high use single family customers

	Specification 2: Survey Variables		Specification 3: CAC*CDH interactions	
	Low Use/Apt.	High Use	Low Use/Apt.	High Use
Critical Price in Effect	-0.093 (0.210)	0.389 (0.467)	-0.189 (0.206)	0.470 (0.451)
Critical Price in Effect * day after critical price	0.048 (0.031)	0.071 (0.058)	0.033 (0.031)	0.055 (0.058)
Crit. Price in Effect * elec. use, kWh / day summer 2002	-0.015 (0.011)	-0.020 (0.012)	-0.012 (0.011)	-0.020 (0.012)
Critical Price in Effect * high ratio rate customer.	0.280 (0.173)	0.190 (0.216)	0.238 (0.162)	0.275 (0.209)
Critical Price in Effect * cooling degree hours 2-7pm	0.004 (0.005)	0.012 (0.007)	0.002 (0.007)	0.020 (0.014)
crit. price in effect * cooling degree hours squared (1000's)	-0.058 (0.046)	-0.092 (0.074)	-0.010 (0.163)	-0.251 (0.330)
Critical Price in Effect * central AC	-0.102 (0.127)	-0.545** (0.227)	-0.040 (0.139)	-0.272 (0.250)
Critical Price in Effect * room AC	0.219* (0.128)	0.534** (0.248)	0.224 (0.136)	0.520** (0.248)
crit. price in effect * cooling degree hours 2-7pm * central AC	. .	. .	-0.00074 (0.003)	-0.005 (0.005)
Critical Price in Effect * CDH 2-7pm squared * central AC	. .	. .	-0.00035 (0.00092)	0.001 (0.002)
N	54446	47535	54446	47535
R-squared	0.3715	0.4331	0.3964	0.4436
Robust standard errors, clustered by customer in parentheses. Significance: *=10% ** =5% ***=1% Cooling degree hours are 2-7PM, base 78^o F. Heating degree hours are base 65^o F.				

Table 10: The impact of critical prices when we separate apartment / low use single family customers from high use single family customers

$$I_{critical} = \sum_{j \in \{1, \mathbf{X}^*\}} (\beta_j + \psi_j)\overline{x}'_j$$

The differences are the addition of ψ_j the coefficient of *CriticalPrice*, and that $\overline{x}'_j$ is an average characteristic on critical days.

Within this framework, I proceed as follows:

- I calculate the total distribution of non-holiday, weekday 2-7 PM peak loads in the CAISO control area.

- Then I use this distribution to assign each day to a load-based bin. I create two sets of bins differentiated by whether the CPP group was paying critical or peak afternoon prices.[57] Appendix F provides descriptive statistics for each bin including average temperatures and peak loads. The appendix shows that the high load bins generally have higher temperatures, meaning that the resulting set of load-based bins resembles the set of temperature based bins reported in Herter et al. [2007].

- Within each bin, I predict the impact of dynamic pricing for customers in each climate zone and statewide, conditional on the temperature conditions being the average seen within each load bin. I modify the framework above to make the average day-customer characteristics, $\overline{x}_j$, the appropriate conditional mean for each climate-load bin. The climate-zone-specific estimates use average customer-level characteristics within each climate zone. All estimates calculate the average population-weighted weather from days with loads in the current bin. The approach works as follows:

$$\begin{aligned} I_{pricetype,bin,zone} &= \\ &\sum_{j \in \{CDH, CDH^2 1, \mathbf{X}^*\}} \beta_j(\overline{x}_j | \mathit{price{=}pricetype,\ load\ in\ bin,\ account\ in\ zone}) + \\ &\sum_{j \notin \{CDH, CDH^2 1, \mathbf{X}^*\}} \beta_j(\overline{x}_j | \mathit{account\ in\ zone}) \end{aligned}$$

These temperature dependent point estimates generalize Faruqui and George's and Herter's efforts to calculate a single point estimate of the CPP effect.

7.2 Point Estimates of the Impact of Dynamic Pricing: Summer Season Weather and Load Patterns in 2003-04

[NOTE: Future work will replace this section with estimates based on the spline-based estimates.]

[57] Most critical events were called simultaneously for all customers statewide. There are a handful of exceptions documented on Charles River Associates [c, 21]. This allows us to provide direct, if imperfect, answers to crucial policy questions about the level of peak-use reduction that dynamic pricing will provide under important statewide load scenarios. I categorize days when any customers paid critical prices as being critical days. The noise in the link between load and temperature adds uncertainty beyond the standard errors on the relationship between temperature and dynamic pricing benefits.

Impacts of Critical Prices on avg. customer demand, kW					
Specification 4: More controls and customer fixed effects					
percentiles of peak load distribution	zone 1	zone 2	zone 3	zone 4	statewide
0-40	0.085	-0.072	-0.135	0.069	-0.073
40-60	0.167	0.017	-0.076	0.031	-0.008
60-80	0.128	-0.052	-0.112	-0.219	-0.082
80-90	0.117	-0.057	-0.146	-0.302*	-0.106
90-95	0.100	-0.064	-0.144	-0.419**	-0.122
95-99	0.080	-0.019	-0.176	-0.453*	-0.119
99-99.99999	0.156	0.013	-0.143	-0.335*	-0.072
max load	0.216	-0.028	-0.205	-0.361*	-0.106
maximum statewide CDH^2	0.186	-0.079	-0.254	-0.414**	-0.154
max zone-by-zone CDH^2	-0.077	-0.207	-0.295	-0.672**	-0.203

Impacts of TOU Peak Prices on avg. customer demand, kW					
Specification 4: More controls and customer fixed effects					
percentiles of peak load distribution	zone 1	zone 2	zone 3	zone 4	statewide
0-40	0.017	0.045	0.004	0.074	0.022
40-60	0.030	0.064	0.047	0.003	0.039
60-80	0.017	0.053	0.029	-0.123	0.013
80-90	0.008	0.084	0.004	-0.233	0.006
90-95	0.018	0.070	0.005	-0.245*	0.0000032
95-99	0.021	0.110	0.027	-0.245	0.026
99-99.99999	0.119	0.088	-0.078	-0.263	-0.008
max load	0.119	0.088	-0.078	-0.263	-0.008
maximum statewide CDH^2	-0.002	-0.041	-0.102	-0.411**	-0.104
max zone-by-zone CDH^2	-0.002	-0.041	-0.079	-0.603**	-0.037

Table 11: Point estimates of the total impacts of dynamic pricing by climate zone and load scenario. Significance: *=10% ** =5%

Table 11 shows just how important the interaction of temperature and dynamic pricing is:

- Customers in climate zone 4 (desert) show statistically and substantively significant benefits from critical prices during all the high demand conditions. The credible point estimates of the impacts range from 1.5/kWh to over 2 kWh per customer-day.

- The point estimates of the peak impacts in zone 4 are consistent with substantively important benefits, but the results are imprecisely estimated and often not statistically significant.

- Climate zone 3 (Central Valley) shows similarly large reductions of .75 to 1.5 kWh per critical customer-day and up to .35 kWh per peak customer-day. These estimates are imprecisely estimated during both critical and peak periods.

- Climate zone 2 shows some response to dynamic pricing, but the benefits are not as impressive as those in the hotter zones and are statistically insignificant. The estimates find zero benefits for the average customer in climate zone 1.

To put these results in perspective, it is useful to note that the statewide average consumption from 2-7PM weekdays was about .8 kW during the (moderately cool) early June pretreatment period, although many high use customers consumed more than twice that much.

These findings are qualitatively similar to those reported at Charles River Associates [c, 61]. They find what may be slightly larger impacts in zone 4 and less impact in zones 1 and 2. Section 6.6 explains that the difference-in-difference estimates find that the treatment group was using less power on peak than was the control group. It treats this as a preexisting difference between the control and treatment groups and subtracts the preexisting difference from the impact estimates, driving much of this difference in findings.[58] If this preexisting difference is, in part, a premature response, then Table 11 understates the benefits of dynamic pricing. The present paper's findings have larger standard errors.

The nature of this estimation approach suggests some caution in interpreting these results. This approach puts a single best fit quadratic function of dynamic pricing's impact as a function of temperature through data drawn from climate zones that differ strongly in the prevalence of air conditioning. Future revisions of this work will use the regressions from Section 5 which address these concerns. On the one hand, there is reason to think that these estimates – especially those that are in the interior of the sample – are quite robust. Letzler [2007] reports results from dropping zones 1 and 2 and finds that the weather sensitivity grows slightly, but that the results are qualitatively indistinguishable. Further, converting statewide data to base-78 cooling degree hours means that we are largely fitting the temperature curve to data from the hotter, high-air-conditioning climates since cool days in cool climates have zero base-78 CDH. However, a couple of results from extreme conditions should be approached with caution:

- The point estimates suggest climate zone 4 (desert) offered up to 3.5 kWh of benefits per customer-day during the most extreme weather within the sample. This is 50%

[58] Faruqui and George's use of before and after data with customer fixed effects makes a somewhat similar correction.

more benefits than they delivered during days that were just a few percentage points lower in the load distribution. The figures in Section 5 suggest that the quadratic estimates overstate benefits at the edge of the data cloud in order to get the best global fit.

- The tables report that dynamic pricing caused a counterproductive increase in zone 1 (coastal fog belt) energy consumption during the state's hottest days. These estimates reflect the rigid functional form which put these days' temperatures at the top of the inverted-U shaped statewide relationship between temperature and dynamic pricing impact. The control-treatment difference estimates shown in figures 7, 8, and 10 in Section 5 suggest that benefits in zone 1 are insensitive to temperature. The lack of correlation between dynamic pricing impacts and temperature makes sense because only 4.1% of customers in zone 1 have air conditioning. The piecewise linear difference-in-difference estimates in Section 5 shown in figure 10 reveals a slightly more complicated story: the treatment group shows large, imprecisely estimated savings during unusually hot conditions for zone 1 during the pretreatment period. These benefits largely disappear during the treatment period, yielding the strange difference-in-difference impact estimate patterns that are typical of the Section 5 piecewise linear estimates.

8 Comparing this approach to existing papers

The analysis that is closest to this paper's work is reported in Faruqui and George [2005] and the SPP final report [Charles River Associates, c]. The present analysis makes some of the same choices as the prior papers to maximize comparability. For example, I consciously emulate their use of the number of bedrooms as a proxy for house size. The nature of the experiment and the available econometric tools drives this paper to make similar choices to Faruqui and George like aggregating each weekday's 5 peak hours into a single observation and analyzing the summer and winter rate periods separately.

The present paper, however, made the opposite choice about the strength of the price sensitivity assumption. Faruqui and George's papers use the SPP data to estimate well-behaved continuous elasticity of substitution demand functions. If their assumptions about demand function are correct then the parameters they estimate predict the implications of a wide variety of rates. The present paper makes weaker assumptions about the nature of demand, meaning that it only attempts to describe the impacts of dynamic rates similar to those used in the SPP. Its approach, however, allows explicit tests of some of Faruqui and George's assumptions.

- Faruqui and George's approach fits a smooth demand curve with a single elasticity to the data, which they plot on pages 62-66 of Charles River Associates [c]. Having a single elasticity forces their estimates to find that customers on high ratio rates respond more to peak and critical events than do customers on low ratio rates.

- Faruqui and George's CES demand curves allow them to decompose response into two parts:

1. substitution between peak and off peak periods as a function of the ratio between the day's afternoon peak and off peak period prices and
2. total daily use as a function of the day's average price.

This means their estimates of the impacts of changes to peak and critical prices are potentially too sensitive to the off-peak prices. Modeling total average use as a function of a weighted average of the afternoon and off peak prices is a strong assumption. For purposes of illustration, assume that the appropriate daily average price is the simple average of the peak and offpeak price. Then their assumption implies that customers would use the same total amount of power if the price were 10 cents during both periods or 19 cents on peak and 1 cent off peak. In other words, the 19 cent price's reduction in lighting and air conditioning usage would have to be exactly offset by increases in electricity use spurred by the 1 cent offpeak price. As a practical matter, if customers reduced usage on critical days which have significantly higher average prices, then this functional form forces the prediction that high-ratio CPP customers increased their use on non critical weekdays relative to control customers because the high ratio rate lowered average prices slightly [Charles River Associates, c, 44].

- Faruqui and George deal with autocorrelation by first differencing their data. Thus, much of their identification comes from changes to and back from critical prices. This may also push them toward analyzing the impacts of customer-level characteristics one at a time rather than many at a time, although they do not explain things this way [Charles River Associates, c, 73]. The present paper deals with autocorrelation by clustering observations by customer, so we can identify coefficients from all of the data. Using the whole data set makes it easy to control for many customer characteristics at once. Neither approach is perfect.
- The approach taken in the present paper controls for more customer characteristics and controls for weather in a more flexible way.

The present paper's regression approach is somewhat similar to Herter [2006a] but extends its work by reporting the impacts of a variety of covariates and by reporting standard errors that reflect uncertainty of the estimates both within each customer as well as across customers. The present work extends the estimates of response by temperature bin in Herter et al. [2007] by decomposing temperatures by climate zone and by using load scenarios to guide the choice of temperature bins.

9 Policy Implications

The substantial diversity among the electric use patterns sensitivity to dynamic pricing across climate zones suggests ways to improve the design of CPP programs. This section discusses ways to focus the program on recruiting the right customers and on spurring response when it will have the greatest value. This section then explores the implications of cross subsidies for the design of rates that can attract responsive customers.

9.1 Recruiting and Targeting Customers

Some residential customers respond to dynamic prices far more than others. Utilities and regulators should use their limited marketing resources and limited program complexity to maximize social benefits of dynamic pricing.[59]

Targeting highly responsive customers is likely to increase the cost effectiveness of a CPP program.[60] Targeting the customers who respond the most increases cost effectiveness if 1) wholesale electricity prices are similar enough in each region, 2) different types of customers have similar enough recruiting costs, and 3) a targeted enrollment effort is more cost effective than a universal effort.[61]The analysis above suggests that readily available regional temperature and customer-level historical use data can play a central role in targeting.[62]

It may be useful to take a mid-term view that opt-in CPP is a stepping stone toward making CPP an opt-out or default offering. Sound mid-term policy would work to build a compelling track record. This requires cost effectiveness. It also suggests enrolling enough customers who look like the eventual participants to demonstrate that each type of customer can remain satisfied in the long term. Thus, a good midterm approach might compromise between targeting highly responsive customers and diversifying the customer pool.

9.2 Dynamic Pricing has Different Implications in Zone 3 and Zone 4

Customers who experience daily peak prices in places that are consistently hot provide different benefits from customers on the same prices in places that are selectively hot, but where heat tends to drive statewide demand peaks. Zone 3 (Central Valley) was very hot[63]

[59]This follows from taking the standard microeconomic approach of using separate policy instruments to meet efficiency and equity goals. The efficiency instruments make the net social benefits of the electricity business as big as possible. Separate instruments can then achieve equity goals through transfers chosen to minimize distortions. My choice of this approach explains many of the differences between my conclusions about targeting and the conclusions in Herter [2006b]. If we believe that it is impossible to bundle separate redistributive programs with dynamic pricing, then the approach taken in Herter [2006b] may be more appropriate than the approach taken here.

[60]Highly responsive customers also provide the greatest reduction in deadweight loss from the mispricing of electricity under moderately strong assumptions about the nature of demand. If elastic customer e changes demand more than each inelastic customer i for a change of prices from P_L to the highest price, P_H, then a sufficient assumption is that the responsive customer decreases demand weakly more than the unresponsive customer at any price in that interval. Formally $\frac{\partial Q_e}{\partial p^*} \geq \frac{\partial Q_i}{\partial p^*} \forall p* \in [P_L, P_H]$. This global condition rules out the possibility that less responsive customer actually had a big deadweight loss because they experienced a big change in quantity demanded after a small price change from the status quo, while the more responsive customer had a small deadweight loss because they experienced almost all of their change in demand in a very narrow price interval near the new price.

[61]Targeting high value customers is particularly valuable if the benefits of enrolling each customer has to justify an expensive meter installation. Places like California are, however, deploying advanced meters for customers on all rates.

[62]Targeting efforts might further explore how to use available data to identify the most valuable customers. It is not clear how fruitful these efforts can be. I conjectured that big users in hot climates would be highly responsive people with air conditioners. Operationalizing this conjecture by interacting summer 2002 kWh / day and summer 2002 kWh / day squared and with climate zone only increased R^2 by 1%.

[63]I define very hot as more than 60 CDH between 2 and 7 PM. A day with 60 or more CDH means that temperature averaged at least 90^{o}F for those hours.

an average of 3 weeks per summer while desert zone 4 was very hot an average of 10 weeks per summer. Thus, zone 3 dynamic pricing customers reduce peaks during many (but not all) of the highest demand days when there is a significant chance of scarcity. Using dynamic pricing to reduce peaks in zone 4 avoids the need to run a (not-so-inefficient) peaker every afternoon for weeks on end. Dynamic pricing in zone 3 is more likely to avoid the need for a peaker that would run a few dozen hours a year.[64]

9.3 Understanding and Dealing with Cross Subsidies

An optimal opt-in dynamic pricing program has to offer compelling savings to customers who provide the largest consumption reduction benefits. [65] Large users with air conditioners in hot climates are desirable participants because they use a great deal of power during peak periods and are quite responsive to price signals.

Big customers in hot climates use more than the statewide average proportion of their power during weekday afternoons, so they would give up the substantial cross subsidies from customers with flatter load shapes (i.e. intertemporal electricity consumption pattern) if they switched to a dynamic rate that keeps bills the same for people with the statewide average load shape. Adjusting dynamic rates for regions within each utility's service territory can make customers in peaky regions more likely to save relative to the alternative, time invariant rate[66]

Tables 12 and 13 show that larger customers and customers in hotter climates use a significantly larger percentage of their total power consumption during weekday 2-7 PM hours.[67] The control group use patterns offers insight onto the distribution of structural

[64]The present analysis has aggregated customers to the climate zone level. Practitioners might explore whether a different set of boundaries is the best, practical way to identify areas where dynamic pricing will have compelling benefits.

[65]This section could be cast in the language of mechanism design, with a benefit-maximizing rate designer attempting to design a set of incentives that cause customers to shift power away from high cost periods (incentive compatibility) while making it rational for the customers who will respond the most to these programs to accept the rate offer (individual rationality).

[66]One compelling response to problems that arise when time invariant rates create cross subsidies may be to reduce the cross subsidies in the alternative, time invariant rates. Eliminating cross subsidies so most customers face the full social cost of their decision to air condition their homes while providing carefully targeted safety nets for vulnerable customers is generally an equitable and appropriate if politically difficult policy. Deploying effective dynamic pricing should be the first priority because it will almost certainly provide far larger economic benefits than will reducing cross subsidies. Some in the industry use the term "free riders" to describe customers from cool climates who switch to dynamic pricing to avoid cross subsidizing air conditioning users. This inappropriately implies that the rival, excludable, pricable capacity to run air conditioners full blast on the hottest weekday afternoon of the summer is a public good.

[67]Customer attrition and the fact that the SPP collected nine months of data to represent a six month summer season complicate calculating a meaningful average percent of power used on peak. This is especially true for the treatment group where particularly peaky customers may have exited the experiment early. Every approach to this problem has significant flaws. One approach would be to report the average value of the ratio of peak to total consumption over all customers i for each week t, or $\frac{Q_{i,t,H}}{Q_{i,t,H}+Q_{i,t,L}}$. This would have underweighted early-exiting peaky customers because we see people who left early for fewer weeks. Another approach would be to calculate total use during peak and offpeak periods for each customer. This would give equal importance to each day that we observe a customer. Days from June would get equal importance to days from July, but the treatment period sample typically contains one observation of May, June, and

Control group: weekday peak hour use as a % of customer's total summer-season power use					
	climate zone				
house type	1	2	3	4	all
apartments	15.7	21.6	22.9	24.9	21.1
Low use single family	17.2	18.8	24.9	24.2	21.1
High use single family	19.9	25.9	27.6	29.1	26.3
all	17.0	21.2	25.2	25.8	22.3
CPP group: weekday peak hour use as a % of customer's total summer-season power use					
	climate zone				
house type	1	2	3	4	all
apartments	15.0	16.4	25.4	28.5	19.1
Low use single family	15.7	19.4	22.4	20.6	20.1
High use single family	16.8	18.0	23.3	26.2	20.8
all	15.6	18.1	23.3	23.9	20.0

Table 12: Weekday peak hour use as a proportion of each customer's total power use. These numbers are weighted to have the same house size and climate zone distribution as the customer-base of the three major utilities. The tables in this section use the same universe that I report in regression 2 and the main means table.

Percentage of each customer's summer-season power used on peak									
control group									
climate zone	min	25%	40%	45%	median	55%	60%	75%	max
1	11.5	14.7	15.8	15.8	16.1	16.2	16.9	18.4	36.7
2	11.6	16.3	18.2	18.9	19.2	20.4	20.8	23.2	52.2
3	6.8	19.5	22.2	23.4	24.7	25.2	27.3	31.5	47.3
4	15.0	19.4	22.5	23.0	25.2	26.3	27.7	31.0	39.8
CPP group									
1	11.1	14.5	15.1	15.2	15.8	15.8	16.1	17.4	21.5
2	6.7	14.6	16.0	16.2	17.1	17.3	17.5	19.8	48.4
3	7.1	16.1	19.3	20.1	21.3	22.2	23.7	28.2	55.1
4	5.5	17.9	21.0	21.6	22.4	24.4	26.3	29.8	47.0

Table 13: Power used during weekday peak hour use as a proportion of total power use. These numbers are weighted to have the same house size and climate zone distribution as the customer-base of the three major utilities.

losses. The average customer in every control group in the hotter two climate zones begins facing bigger bills because they used more than the statewide average of 22.3% of their summer-season power during peak hours. Peak hours are less than 15% of all hours. Further, even after the treatment group adjusted its load shape, more than 55% (45%) of customers in climate zone 4 (3) come out behind.[68],[69],[70]

We could preserve inter-regional cross subsidies by simply calculating a region specific proportion of power used during peak periods $\overline{\alpha}_{r,H}$ for each region r and using those to set peak and offpeak markups that preserve revenue neutrality within each region as well as statewide.[72] Doing so would mean that between 60 and 75% and between 55% and 60% of the treatment group would come out ahead in zones 3 and 4 respectively. Further analysis could examine how many customers who would still come out behind responded so little that there is little value in designing a program attractive to them and how many made significant, beneficial, but insufficient-to-come-out-ahead changes to their heavy use of afternoon power.

Making participation attractive to customers who get cross subsidies in the status quo is a common problem with a variety of solutions. Offering customers who live in hot climates fixed credits of roughly the amount of cross subsidy that they enjoy under time invariant pricing could also address this problem. See Borenstein [2007] for a careful discussion of directly analogous dynamic pricing wealth transfer issues in the context of the implementation of real time pricing for large industrial customers. Similarly Wiser et al. [2007] explores the

October (from either 2003 or 2004) and two of July, August, and September (2003 and 2004). This approach would overstate the importance of the hottest (and often peakiest) summer months which are measured in both 2003 and 2004. The ratios reported here give each customer equal weight and but downweight observations from repeated months in an attempt to give each month equal importance:

First I calculate each customer i's average use during weekday peak and offpeak periods and weekends ($j \in \{L, H, w\}$) separately for the over and undersampled months. Let $t1$ be the set of oversampled time periods (September and most of July for all utilities, August for SCE and SDG&E). Let λ_u represent the percentage of the summer season that comes from the oversampled times for customers of utility u. I then average total use during each kind of period, weighting by λ_u, the proportion of the summer season during each period

$$\overline{X_{i,j}} = \lambda_u \overline{X_{i,j,t1}} + (1 - \lambda_u)\overline{X_{i,j,t2}}$$

Then I use the average total use during each kind of period to calculate the ratio of peak to total use, letting ω be the percentage of all days in the sample that are weekends and holidays:

$$\overline{\alpha}_{i,H} = \frac{(1-\omega) * \overline{X_{i,H}}}{(1-\omega) * (\overline{X_{i,H}} + \overline{X_{i,L}}) + \omega * (\overline{X_{i,w}})}$$

[68]There is evidence that the smaller treatment customer classes in zones 3 and 4 are peakier. These small, peaky customers may have stayed in the pilot because benefits like $175 in participation payments and a potential sense of contributing to society outweighed their modest increase in bills.

[69]By contrast, more than 75% (60%) of control customers in temperate, climate zone 1 (2) would be structural winners on this simple rate. It further suggests that the treatment group may have increased its gains by further flattening its load shapes in response to price signals.

[70]This kind of cross subsidy may reconcile Herter [2006b]'s finding that customers from the cooler climates got significant bill savings with my finding that they barely changed their load in response to dynamic pricing incentives.[71]

[72]Baseline-rebate rates can be thought of as setting a customer specific $\alpha_{i,H}$. Using each customer's behavior to set a customer-specific baseline level has notable drawbacks discussed at length in Letzler [2006].

importance of rate design for making commercial solar photovoltaic installations attractive.

Ideal solutions eliminate the cross subsidies that make it artificially cheap to consume expensive peak period power in ways that increase the risk of blackouts. This section recognizes that continuing troubling cross subsidies – at least temporarily – may be an expedient way to get consumers to participate in peak load management programs.

9.4 Concerns about vulnerable populations

There is legitimate concern about whether a simplistic, mandatory CPP implementation would harm vulnerable customers who have low incomes and inflexible demand. A disproportionate number of elderly customers and families of small children might have these characteristics.[73]

It is harder to make a case that we need to protect customers from a program that customers have to opt-into and that they can leave at any time. The SPP's evidence confirms the belief that there is little reason for concern about the impacts of an opt-in program on the elderly and families with small children:

- Tables 1 on page 7 and 2 on page 8 show that children under the age of 5 were underrepresented in the CPP population relative to the control group, but the families with small children that did participate were more likely to stay rendering this difference statistically insignificant among customers who stayed in the experiment at least four months. Indeed, the results reject the intuition that families with small children are less flexible in their energy use. All specifications find that families with small children either use the same or less peak power than did households with the same number of people, but with one more member between ages 5 and 65 instead of a child. Specification 4 finds that households responded to price signals by about 1.1 kWh per person per day (p=.01) more for each child under age 5. A simpler difference-in-difference specification that adds the number of children under 5 to specification 2, gets the same sign but no statistical significance.

- Senior citizens were a larger, but statistically indistinguishable, proportion of the the CPP group than they were of the control group. Attrition did not change this pattern. Households with members over 65 responded to price signals significantly more – by about 1.1 kWh per person per day (p=.001) – than did similar households with the same number of people, but with that member being between ages 5 and 65.[74]

Gulf Power's Good Cents Select residential CPP program's experience is similarly reassuring. Thirty percent of its customers are over 65 [White, 2005, 11]. Its customers respond

[73]A simple, effective way to address concerns about a mandatory CPP program would be to improve everyone's incentives with mandatory CPP and then to add a second equity program that would identify categories of vulnerable people and give them a fixed bill credit roughly equal to the cross subsidy they gave up to participate in CPP.

[74]Results reported in Letzler [2007], appendix F.2 suggest that seniors in high use households were far more responsive than the average person in a high use household, while seniors in low use households were statistically indistinguishable from the average person in those houses. The selection problems in the high use category suggest the use of some caution in believing that the average high use senior will respond better, but do suggest that high use seniors who opt-in really can benefit from this program.

well, save money, report very high satisfaction, and rarely leave the program. [White, 2006]

10 Conclusions

The high social cost of air-conditioning-driven peaks in electricity demand is a major justification for dynamic pricing. Customers in California's Statewide Pricing propitiously responded the most to dynamic pricing during hot weather in regions where most customers have air conditioning. Customers in the desert appear to have provided sustained savings over many weeks per summer, while customers in the Central Valley appear to have provided more focused reductions in demand during peak periods. All else equal, bigger customers responded more during critical events. It estimates that the benefits of dynamic pricing range from zero in cooler climates on cooler days to .3 (.4) kW every hour for peak (critical) prices on the hottest days in the two hot climate zones. This suggests that opt-in dynamic pricing programs should be designed to recruit customers from hot climates and to provide good incentives during hot weather.

It finds a difference of .07 kW in use between the treatment and control groups beginning during the 'pretreatment period' that it treats as a preexisting difference, but may in fact be further impacts of dynamic pricing that began early when customers received documents that were clearer about the nature of the new prices than about their timing.

Dynamic rates need to make it rational for highly responsive consumers to participate in the program and give them incentives to reduce usage during high cost periods. The SPP's experience suggests offering the most responsive customers enough savings to convince them to participate requires careful attention to cross subsidies and differences in regional usage patterns.

References

John Barnard, Constantine E. Frangakis, Jennifer L. Hill, and Donald B. Rubin. Principal stratification approach to broken randomized experiments: A case study of school choice vouchers in new york city. *Journal of the American Statistical Association*, 98, 2003.

Rich Barnes. personal e-mail correspondance. Rich Barnes is Senior Vice President of Demand Side Services at Kema-Xenergy, a firm which helped design SPP mailings and conduct the SPP survey, May 2007.

Severin Borenstein. The long-run efficiency of real-time electricity pricing. *Energy Journal*, 26(3), 2005.

Severin Borenstein. Wealth transfers among large customers from implementing real-time retail electricity pricing. *The Energy Journal*, 28(2), 2007. Also: Center for the Study of Energy Markets Working Paper 156 http://www.ucei.berkeley.edu/PDF/csemwp156.pdf, University of California Energy Institute.

CAISO. Cumulative totals of restricted maintenance operations, alert, warning, emergency and power watch notices issued from 1998 to present, May 2007a. http://www.caiso.com/docs/2001/06/01/200106011228581047.html.

CAISO. Alerts, warnings, and emergencies. White Paper, April 2004b. http://www.caiso.com/awe/AlertsWarnings-WhitePaper.pdf.

Charles River Associates. *Statewide Pricing Pilot Summer 2003 Impact Analysis.* Charles River Associates, August 2004a. http://www.energy.ca.gov/demandresponse/documents/group3_final_reports/2004-10-29_SPP_REPORT.PDF.

Charles River Associates. Statewide pricing pilot summer 2003 impact analysis appendices. *California Energy Commission Demand Response Docket: Working Group 3*, August 2004b. http://www.energy.ca.gov/demandresponse/documents/group3_final_reports/2004-10-29_SPP_REPORT-APPEN.PDF.

Charles River Associates. *Final Report: Impact Evaluation of the California Statewide Pricing Pilot*, March 2005c. energyarchive.ca.gov/demandresponse/documents/group3_final_reports /2005-03-24_SPP_FINAL_REP.PDF.

Charles River Associates. *Impact Evaluation of the California Statewide Pricing Pilot Appendices.* Charles River Associates, March 2005d. http://energyarchive.ca.gov/demandresponse/documents/group3_final_reports/2005-03-24_SPP_APPENDICES.PDF.

Ahmad Faruqui and Stephen George. Quantifying customer response to dynamic pricing. *The Electricity Journal*, 2005.

Focus Pointe. Statewide pricing pilot: Enrollment refusal follow-up research. Report, November 2003. Project SEMP03-02; Conducted for Southern California Edison, San Diego Gas & Electric and Pacific Gas & Electric.

Karen Herter. *Effects of Critical Peak Pricing on Residential Electricity Use in California.* PhD thesis, Energy and Resources Group, UC Berkeley, Fall 2006a.

Karen Herter. Residential implementation of critical-peak pricing of electricity. *Energy Policy*, 2006b.

Karen Herter, Patrick McAuliffe, and Arthur Rosenfeld. Observed temperature effects on hourly residential electric load reduction in response to an experimental critical peak pricing tariff. *Lawrence Berkeley National Laboratory Working Paper*, 2005.

Karen Herter, Patrick McAuliffe, and Arthur Rosenfeld. An exploratory analysis of California residential customer response to critical peak pricing of electricity. *Energy*, 2007.

Robert Letzler. Applying psychology to economic incentive design: Using incentive preserving rebates to increase acceptance of critical peak pricing. *Center for the Study of Energy Markets Working Paper 162*, December 2006. http://www.ucei.berkeley.edu/PDF/csemwp162.pdf.

Robert Letzler. *Essays on Opt-in Dynamic Electricity Pricing.* PhD thesis, UC Berkeley, 2007. http://www.allorama.org/rl/letzler_dissertation.pdf.

Pacific Gas & Electric. Schedule e-3 experimental residential critical peak pricing service. http://www.pge.com/nots/rates/tariffs/advice/tariffsheets/20515-20591.pdf California P.U.C. Sheet No. 20520-E see also: http://www.pge.com/rates/tariffs/Res_030401-040229.xls, July 2003.

San Diego Gas & Electric. Schedule EECC-CPP-F experimental electric energy commodity cost domestic critical peak pricing service - fixed. California P.U.C. Sheet No. 17208-E http://www.sdge.com/tm2/pdf/1605-E.pdf, July 2004.

Southern California Edison. Experimental schedule tou-d-cppf-1. California P.U.C. Sheet No 34512-E through 34521-E http://www.sce.com/NR/sc3/tm2/pdf/1722-E.pdf See also: http://www.sce.com/NR/sc3/tm2/RPA/Reg_Info_Ctr/Historical/D.pdf http://www.sce.com/NR/sc3/tm2/pdf/1695-E.pdf, July 15 2003.

Brian White. Gulf power. Personal Interview, 2006.

Brian White. GoodCents select program overview. Presentations to the Oregon PUC and IEEE, January 2005.

R. Wiser, A. Mills, G. Barbose, and W. Golove. The impact of retail rate structures on the economics of commercial photovoltaic systems in california. *Lawrence Berkeley National Lab 63019*, July 2007. http://eetd.lbl.gov/ea/ems/reports/63019.pdf.

Frank Wolak. Residential customer response to real time pricing: The anaheim critical peak pricing experiment. *Working Paper, presented at the POWER conference*, March 2006. ftp://zia.stanford.edu/pub/papers/anaheim_cpp.pdf.

A Appendix: Customer Recruitment Process

Figure 12: The recruiting process that generated the control and CPP groups.

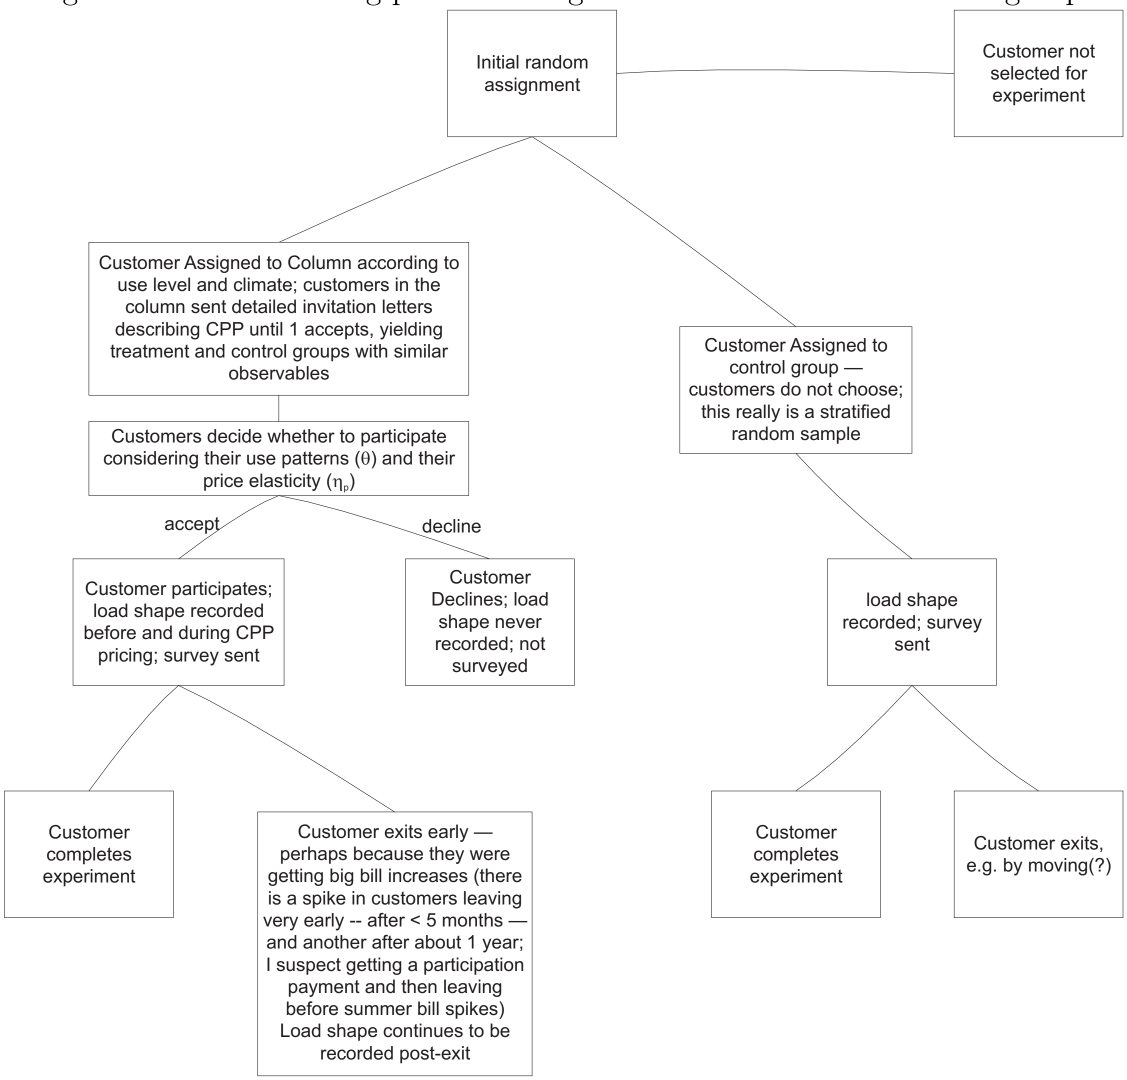

B Appendix: Graphs of the relationship between average daily electricity use the summer before the experiment and peak use during the pretreatment and treatment periods

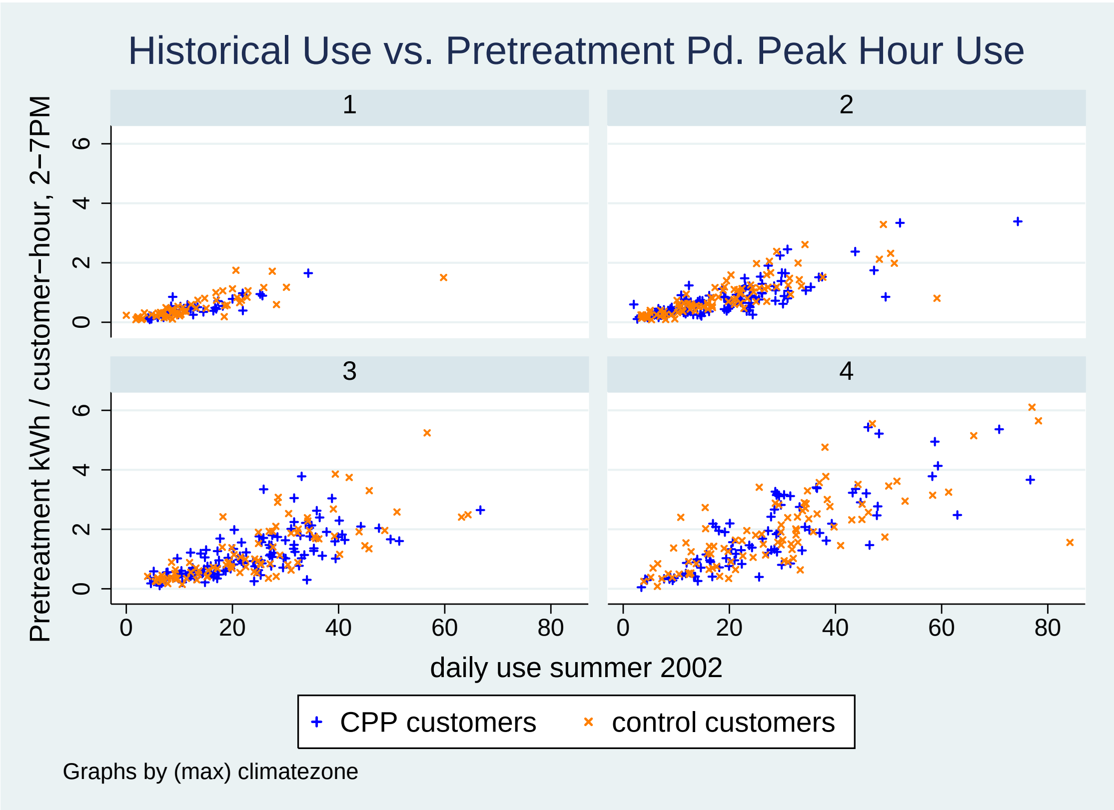

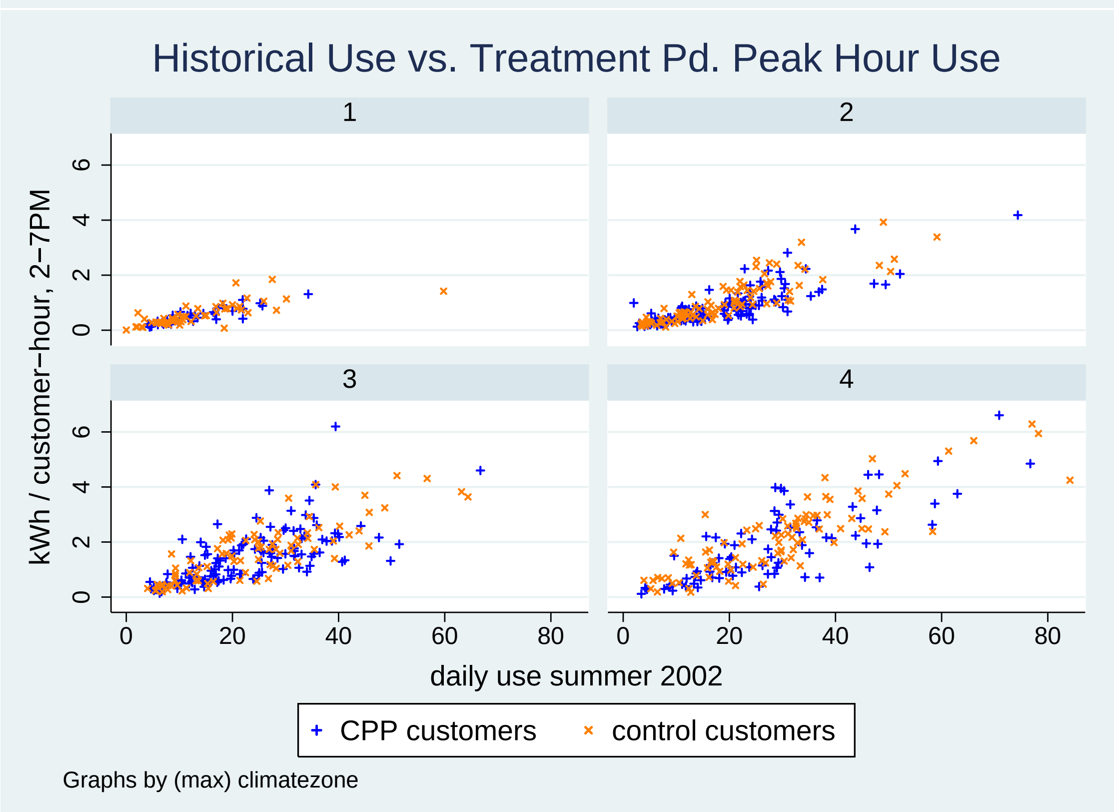

Figure 13: This graph shows raw average daily use plotted against raw average daily afternoon use before and after the experiment began. This illustrates the identification strategy for the relationship between customer historical usage level and customer use during weekday peak hours. The simple, graphical results are not nearly as striking as those that in figures 4, 6, and 5 which support the regression finding that dynamic pricing has greater impact on hotter days.

C Appendix: Methodology used to Calculate Population Weighted Temperatures

This section describes the methodology used to calculate the population weighted temperatures that I use in Section 7.

The population numbers reported here come from Charles River Associates [d, 18-19].

Despite the fact that "each ... customer in the experiment was assigned by the relevant utility to a specific weather station located in close proximity to the customer" [Charles River Associates, d, 18], California is known for its micro climates. Thus, slightly more than half of the 58 weather stations in the sample contain customers from more than one climate zone. For example, customers to the north of the Oakland weather station are in climate zone 1, while customers to the south and east of it are in climate zone 2. More disturbingly, a handful of climate zone 1 customers in the mountains above desert weather stations are assigned to those stations. I emulate the SPP's methodology that: "When a weather station was included in more than one climate zone, the distribution of control group customers in the experiment assigned to that weather station was used to allocate the station population to each climate zone" [Charles River Associates, d, 18].

Specifically, I begin with the whole sample of control group customers that the utilities considered recruiting, as documented in the SPP Database "Table 5."

Each utility appears to have set a target number of customer candidates for each climate-zone-by-customer-type "slot" in the experiment.[75] The database includes idiosyncratic numbers of additional candidates, perhaps added to deal with recruiting problems. There is reason to fear that these idiosyncrasies may be correlated with customer characteristics, like difficulty installing advanced meters in certain kinds of multifamily buildings. Thus, I drop the idiosyncratic customers and obtain 2 candidate control customers per slot in SCE, 4 in SDG&E, and 24 in PG&E. Each weather station covers just one utility, so this approach yields the most statistical power possible given the design. Estimates for PG&E may, however, be considerably more reliable than those for the other two utilities.

Charles River Associates [d, 18-19] contains the population data used here, but is missing an entry for SDG&E weather station S10. Disturbingly, attempts to back out the population of this station using statewide population numbers reveal that the population by weather station table reports a slightly larger population than does population sampling table at Charles River Associates [d, 22]. The distributions of the populations as reconstructed here are, however, qualitatively quite similar.

I calculated the percentage of customers in each zone for each weather station as follows: Let $PctPop_{a,z}$ be the percentage of the statewide population of accounts that is of account type $a \in \{$*apartments, low use single family, high use single family*$\}$ and in climate zone z. The $PctPop_{a,z}$ values came from the statewide weights spreadsheet. Let $PctSlots_{a,z}$ be the percentage of the total experimental slots assigned to account type a in zone z.

Then, I calculated a weight, $\omega_{a,z}$ representing the ratio between the number of slots who would have come from that zone and account type had the sample been representative of the population and the number of actual slots assigned:

[75] The utilities had clear recruitment targets documented at Charles River Associates [d, 22]. For example, the sample design called for PG&E to recruit 17 low-use, single family control customers for climate zone 1.

$$\omega_{a,z} = \frac{PctPop_{a,z}}{PctSlots_{a,z}}$$

Then I constructed the weighted count of people $C_{z,s}$ in each zone for weather station s, $N_{a,z,s}$ within each weather station by weighting and adding up the number of people of each type in that zone:

$$C_{z,s} = \sum_a \omega_{a,z} N_{a,z,s}$$

I could use this to assign the station's population to climate zones where $Pop_{z,s}$ is the population that lives in zone z closest to station s and $TotalPop_s$ is the total population closest to station s:

$$Pop_{z,s} = TotalPop_s \frac{C_{z,s}}{\sum_z C_{z,s}}$$

D Appendix: Distribution of Cooling Degree Hours by Climate Zone

Distribution of 2-7 PM Base-78 Cooling Degree Hours by Climate Zone

ordinary days	minimum	25%	median	75%	90%	95%	98%	max
zone 1	0.0	0.4	2.2	4.0	8.5	18.7	43.6	59.5
zone 2	0.0	1.0	3.7	12.5	20.8	27.7	41.9	49.0
zone 3	0.0	9.4	25.0	42.5	56.9	62.1	70.2	85.4
zone 4	0.0	32.8	60.2	88.9	103.9	112.1	119.5	126.6
statewide	0.0	8.8	16.9	27.6	36.7	41.9	54.8	64.0
critical days	minimum	25%	median	75%	90%	95%	98%	max
zone 1	0.0	3.8	4.8	22.2	37.9	38.7	41.6	41.6
zone 2	0.3	9.3	13.8	28.7	37.0	38.0	40.5	40.5
zone 3	2.7	41.6	50.6	61.5	68.9	71.1	75.0	75.0
zone 4	27.8	58.5	103.0	114.0	117.7	123.9	124.7	124.7
statewide	5.6	29.0	36.5	44.5	52.6	54.2	57.3	57.3

E Appendix: Main Impacts Tables for Other Regression Specifications

Impacts of Critical Prices on avg. customer demand, kW					
specification 1 Simplest Diff in Diff					
percentiles of peak load distribution	zone 1	zone 2	zone 3	zone 4	statewide
0-40	0.128	-0.012	-0.053	0.00020	-0.033
40-60	0.219*	0.090	0.026	-0.029	0.047
60-80	0.181	0.016	-0.013	-0.276*	-0.031
80-90	0.168	0.011	-0.042	-0.350**	-0.051
90-95	0.148	0.003	-0.041	-0.473**	-0.070
95-99	0.125	0.051	-0.068	-0.500**	-0.063
99-99.99999	0.212**	0.092	-0.034	-0.384**	-0.010
max load	0.282**	0.052	-0.095	-0.408**	-0.042
maximum statewide CDH^2	0.250*	-0.002	-0.148	-0.461**	-0.093
max zone-by-zone CDH^2	-0.018	-0.133	-0.185	-0.720**	-0.141

Impacts of Critical Prices on avg. customer demand, kW					
Specification 2: Adding Survey Variables					
percentiles of peak load distribution	zone 1	zone 2	zone 3	zone 4	statewide
0-40	0.186	-0.013	-0.071	-0.027	-0.036
40-60	0.251*	0.065	0.008	-0.030	0.033
60-80	0.230	0.016	-0.020	-0.157	-0.014
80-90	0.221	0.014	-0.025	-0.181	-0.020
90-95	0.205	0.008	-0.028	-0.257	-0.034
95-99	0.188	0.039	-0.032	-0.258	-0.024
99-99.99999	0.254**	0.076	-0.012	-0.199	0.014
max load	0.309**	0.058	-0.044	-0.209	0.001
maximum statewide CDH^2	0.288*	0.023	-0.079	-0.238	-0.031
max zone-by-zone CDH^2	0.135	-0.053	-0.093	-0.379	-0.055

Impacts of Critical Prices on avg. customer demand, kW					
Specification 3: Adding CAC*CDH interactions					
percentiles of peak load distribution	zone 1	zone 2	zone 3	zone 4	statewide
0-40	0.131	-0.051	-0.052	-0.013	-0.039
40-60	0.189	0.013	-0.008	-0.038	0.008
60-80	0.162	-0.036	-0.033	-0.211	-0.044
80-90	0.154	-0.040	-0.056	-0.268	-0.060
90-95	0.142	-0.045	-0.055	-0.350	-0.072
95-99	0.128	-0.013	-0.077	-0.372	-0.069
99-99.99999	0.182	0.011	-0.054	-0.291	-0.035
max load	0.225	-0.018	-0.097	-0.308	-0.059
maximum statewide CDH^2	0.204	-0.054	-0.131	-0.345	-0.092
max zone-by-zone CDH^2	0.021	-0.143	-0.159	-0.524	-0.126

Impacts of TOU Peak Prices on avg. customer demand, kW					
specification 1 Simplest Diff in Diff					
percentiles of peak load distribution	zone 1	zone 2	zone 3	zone 4	statewide
0-40	0.010	-0.041	-0.006	-0.011	-0.031
40-60	0.028	-0.014	0.061	-0.036	0.003
60-80	0.016	-0.021	0.057	-0.142	-0.014
80-90	0.004	0.013	0.057	-0.227	-0.010
90-95	0.016	-0.003	0.053	-0.247*	-0.019
95-99	0.018	0.044	0.086	-0.224	0.016
99-99.99999	0.167*	0.066	0.012	-0.238	0.018
max load	0.167*	0.066	0.012	-0.238	0.018
maximum statewide CDH^2	0.073	-0.064	-0.029	-0.367*	-0.078
max zone-by-zone CDH^2	0.073	-0.064	0.020	-0.533**	-0.003

Impacts of TOU Peak Prices on avg. customer demand, kW					
Specification 2: Adding Survey Variables					
percentiles of peak load distribution	zone 1	zone 2	zone 3	zone 4	statewide
0-40	-0.017	-0.063	-0.011	0.013	-0.037
40-60	0.00038	-0.037	0.058	0.022	-0.00045
60-80	-0.007	-0.039	0.062	-0.050	-0.008
80-90	-0.019	-0.011	0.078	-0.102	0.002
90-95	-0.009	-0.024	0.072	-0.123	-0.007
95-99	-0.007	0.017	0.105	-0.090	0.026
99-99.99999	0.142	0.061	0.066	-0.100	0.051
max load	0.142	0.061	0.066	-0.100	0.051
maximum statewide CDH^2	0.085	-0.041	0.022	-0.190	-0.026
max zone-by-zone CDH^2	0.085	-0.041	0.077	-0.305	0.039

Impacts of TOU Peak Prices on avg. customer demand, kW					
Specification 3: Adding CAC*CDH interactions					
percentiles of peak load distribution	zone 1	zone 2	zone 3	zone 4	statewide
0-40	-0.023	-0.079	-0.038	0.010	-0.060
40-60	-0.003	-0.048	0.056	0.145	0.001
60-80	0.001	-0.035	0.097	0.181	0.025
80-90	-0.014	-0.019	0.175	0.238	0.061
90-95	-0.005	-0.027	0.160	0.205	0.051
95-99	-0.006	0.007	0.202	0.285	0.088
99-99.99999	0.191	0.148	0.283	0.292	0.202
max load	0.191	0.148	0.283	0.292	0.202
maximum statewide CDH^2	0.253	0.114	0.217	0.318	0.176
max zone-by-zone CDH^2	0.253	0.114	0.311	0.359	0.220

F Appendix: Population Weighted Cooling Degree Hours and Cooling Degree Hours Squared

Peak

These tables report base-78 cooling degree hours and base-78 CDH^2 in 1000's. Both Jensen's inequality, the fact that some customers get weather data from the nearest weather station which may be in a quite different weather zone[76] and the fact that CDH cannot go negative affect these estimates. Notice, for example, that if there were no variance among the readings, that we would expect .036 thousand CDH^2 – rather than the observed .30 – in climate zone 1 for the 40-60 % of peak load scenario.

[76]For example, the city of Fresno is in climate zone 4, but the Fresno weather station includes some climate zone 1 customers who live high in the mountains above the city.

Characteristics of TOU Peak-Priced Days by Zone and Bin

	N	Peak	Zone 1		Zone 2		Zone 3		Zone 4		State	
percentiles of peak load distribution		MW	CDH	CDH^2	CDH	CDH^2	CDH	CDH^2	CDH	CDH^2	CDH	CDH^2
0 to < 40th Percentile	83	32131	1.80	0.07	3.11	0.11	10.94	0.47	36.84	2.33	9.05	0.46
40 to < 60th Percentile	41	36380	5.69	0.27	8.83	0.41	29.34	1.64	66.76	5.54	21.19	1.33
60 to < 80th Percentile	35	39000	6.87	0.51	12.01	0.79	38.25	2.56	77.40	7.64	26.76	2.05
80 to < 90th Percentile	18	41142	3.95	0.33	14.32	0.69	55.17	4.24	91.88	9.91	34.42	2.77
90 to < 95th Percentile	7	42230	5.61	0.39	13.17	0.73	51.94	3.96	85.47	9.47	32.39	2.65
95 to < 99th percentile	3	44117	5.28	0.33	19.16	0.86	60.03	4.44	101.94	10.88	39.39	3.01
99th percentile to max	1	45033	43.61	2.68	48.95	3.61	79.55	7.10	103.91	11.22	63.96	5.42
scenarios by temperature												
day with maximum statewide CDH^2	1	40117	59.52	5.17	45.17	4.50	66.58	6.22	113.16	13.40	61.00	6.08
max zone-by-zone CDH^2	-	40495	59.52	5.17	45.17	4.50	85.35	7.60	126.56	16.35	68.46	6.84

Characteristics of Critically-Priced Days by Zone and Bin

	N	Peak	Zone 1		Zone 2		Zone 3		Zone 4		State	
percentiles of peak load distribution		MW	CDH	CDH^2	CDH	CDH^2	CDH	CDH^2	CDH	CDH^2	CDH	CDH^2
0 to 40th Percentile	2	31297	0.72	0.02	2.18	0.05	11.23	0.42	41.39	2.06	9.07	0.38
40 to 60th Percentile	2	35795	14.58	0.49	22.89	1.06	48.76	3.21	58.35	3.92	33.95	1.98
60 to 80th Percentile	8	39503	16.51	1.02	17.20	1.20	43.02	3.03	79.76	8.18	32.05	2.51
80 to 90th Percentile	3	41103	14.46	0.94	18.15	1.33	56.49	4.55	101.69	10.91	38.89	3.34
90 to 95th Percentile	4	42255	9.75	0.68	15.78	1.18	53.02	4.22	97.39	11.63	35.67	3.21
95 to 99th Percentile	5	43876	5.67	0.51	19.20	1.07	66.54	5.73	114.10	13.43	42.84	3.81
99 to 100th Percentile	2	45216	22.61	1.30	34.78	2.15	66.05	5.38	103.89	11.42	50.67	4.07
max load	1	45562	38.67	2.17	40.49	3.05	71.14	6.41	108.41	12.06	57.32	4.98
scenarios by temperature												
maximum statewide CDH^2	1	41264	35.96	2.21	36.96	3.21	66.48	6.46	111.32	12.82	54.21	5.15
zone-by-zone max CDH^2	-	42183	41.55	2.73	37.98	3.45	75.04	6.96	123.85	15.70	59.44	5.79