Implementing Opt-in, Residential, Dynamic Electricity Pricing: Insights from Economics and Psychology

by

Robert James Letzler

B.A. (Brown University) 1999
M.P.P. (UC Berkeley) 2003
M.A. (UC Berkeley) 2004

A dissertation submitted in partial satisfaction of the
requirements for the degree of
Doctor of Philosophy

in

Public Policy

in the

GRADUATE DIVISION
of the
UNIVERSITY OF CALIFORNIA, BERKELEY

Committee in charge:
Professor Lee S. Friedman, Chair
Professor Severin Borenstein
Professor Robert MacCoun
Professor John Morgan

Fall 2007

Implementing Opt-in, Residential, Dynamic Electricity Pricing: Insights from Economics and Psychology

Abstract

Implementing Opt-in, Residential, Dynamic Electricity Pricing: Insights from Economics and Psychology

by

Robert James Letzler
Doctor of Philosophy in Public Policy

University of California, Berkeley

Professor Lee S. Friedman, Chair

The Impacts of Residential Critical Peak Pricing: Evidence from California's Statewide Pricing Pilot

California's Statewide Pricing Pilot explored the impact of Critical Peak Electricity Pricing (CPP) for residential customers. These 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 1) the kinds of customers and situations in which CPP is likely to generate the greatest response and 2) how CPP will affect different subgroups of the population. It finds that dynamic pricing had larger absolute kilowatt load impacts on hotter days and on 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 that were hot enough that most customers had air conditioning and in conditions that prompted them to use air conditioning. Targetting extra marketing efforts at the kind of hot-region customers who reduced electricity use the most when power prices rose is likely to incresae the program's benefits.

Applying Psychology to Economic Policy Design: Using Incentive Preserving Rebates to Increase Acceptance of Critical Peak Electricity Pricing

This project extends the idea that policy makers should address problems by improving economic incentives. This project adds that presenting incentives in a way that reflects how people make decisions can sometimes improve consumers' responses to the incentives and policy outcomes. This paper uses behavioral economics to propose ways to increase electricity policy effectiveness.

The cost of generating power fluctuates enormously from hour to hour but most customers pay time-invariant prices for power. The mismatch between the fluctuating cost and the fixed price wastes billions of dollars. Critical Peak Pricing (CPP) reduces this waste by setting offpeak, peak, and "critical" prices that better reflect the cost of power during time periods. Customers in CPP pilot programs used less power during high-priced periods than did customers on traditional, time-invariant rates. CPP customers reported high satisfaction levels and often saved 10% or more. Yet, roughly 99% of customers reject opportunities to switch to CPP. The psychology literature documents a set of decision making heuristics that people use to choose among options with uncertain payoffs. This paper describes the evidence that one or more of these heuristics explains customer reluctance to opt-in to CPP. It then suggests Incentive Preserving Rebates that change the presentation of CPP to address these heuristics. Incentive Preserving Rebates reframe scarcity "events" as opportunities to get rebates rather than as periods of extremely high prices. Incentive Preserving Rebates change the presentation, but change neither marginal incentives nor each customer's total annual payments. The paper then explores the implications of Incentive Preserving Rebates for customers who participated in a California pilot program.

Optimal Deployment of a simple menu of Incentive Preserving Rebates for CPP Rates with Heterogeneous Customers

The previous chapter proposes using Incentive Preserving (IP) Rebates to change the presentation of critical peak pricing (CPP) in a way that makes it more attractive to consumers. Adding IP Rebates to critical peak pricing maintains CPP's marginal incentives and leaves each customer's total annual bill the same. IP rebates work by selling each customer rights that they can either use to buy power at the usual price during a high-price "event" or to cash in for a rebate. Customers buy their own rights bundled with the first units of power they buy each month. An IP rebate implementation has to assign each customer a quantity of rights per event and the amount of power that the rate marks up each month to pay for these rights. Good choice of quantities make IP rebates more likely to work as promised. Simple, effective assignment rules are desirable. This project shows that it is possible to derive a small, optimized menu of IP rebate offers that assigns existing categories of customers into low, medium, and high use categories. It shows that these offers work well for the customers who participated in California's Statewide Pricing Pilot (SPP). These offers use existing customer categories. Using three optimal categories far outperforms one- and two-category offers. Four and five category offers perform modestly better than the three category offer, but perhaps not enough to justify the added complexity. The three, four, and five category offers all achieve at least 96% of the benefit level of making one offer to each of the 16 groups of customers. The optimal three category offer makes consistent offers to between 75% and 90% of the customers in most groups. An offer is consistent if the customer gets consistent rebates during each month with an event and consistently purchases the rights that the rate offers them. A customer gets consistent rebates if the offer includes enough kWh at the usual price so that the customer gets a (weakly positive) rebate during each month with an event. The customer makes consistent rights purchases if the rights come bundled with a number of units of power that is less than the customer uses each month. This offer performs far better than a single, statewide one-size-fits-all offer or making one offer for each of the four climate zones. Thus, it is possible to consderably simplify IP Rebate rates while preserving their performance. Good rates will both require differentiating among customers by energy consumption level or a good proxy for it. Even the best rates considered here require accepting that a few customers will not get offers that are ideal matches for their consumption patterns.

Professor Lee S. Friedman
Dissertation Committee Chair

In Memory of Susie Li (1976-2000): Mentor, Inspiration, Friend.

Contents

Acknowledgments

I would like to thank my advisers Lee Friedman, Severin Borenstein, Rob MacCoun, and John Morgan. Meghan Busse and Geno Smolensky also provided major, ongoing input, guidance and support.

Thanks go to Max Aufhammer, Rich Barnes, Nina Bubnova, Ahmad Faruqui, Pete Fishman, Karen Herter, Teck Ho, Lily Hseuh, John Jurewitz, Karen Notsund, Lauren Pemberton, Matt Rabin, Catherine Waddams, Brian White, Charlie Whitmore, John Winfield, 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 send thanks to Susan McNicholl for alerting me to the selection problem that I discuss at length in chapter 1.

Chapter 2 benefitted from comments from participants at the Berkeley PEPP IGERT and UCEI seminars.

Thanks go to Ramteen Shioshansi, Vishnu Narayan, Jian Yao, Rich Plevin, and Michael McElroy for help working on the optimization in chapter 3 as an integer programming problem in AMPL and in linking it to the operations research literature.

Chapter 1

Introduction

Microeconomic principles urge analysts to address market failures and policy failures that cause goods' prices to diverge from their marginal, social costs. These failures are ubiquitous, but some are far more serious than others. Interventions like trucking deregulation and emissions credit trading have been quite effective because they address these poor incentives. It often takes significant effort to devise policies that fix incentive problems and:

- work with existing institutions that may have to be changed incrementally;
- convert simple principles into approaches that can be implemented by agencies or firms with limited ability to gather and process information;
- seem attractive to people who are comfortable with the status quo and loss averse, who use simplified heuristics instead of doing all of the math to assess risky choices, and who may be concerned about the credibility of new regimes;
- and offer, if not a Pareto improvement, compelling benefits to enough stake holders to make the change politically feasible.

Public Policy scholarship has long recognized that policies are often made through incremental change. This project contributes to a literature (e.g. Robyn [1987], Hausker [1992, 1986]) that uses economic ideas and careful analysis to guide incremental change. Public policy scholarship can make useful contributions by identifying classes of problems that challenge efforts to improve incentives and developing new insights about practical ways to solve these problems.

This project draws on scholarship at the intersection of psychology and economics to explain the difficulty of implementing an economic policy. It is an early project to design a policy that creates incentives that perform well on conventional microeconomic criteria and present these incentives to be attractive to customers regardless of whether they use neo-classical expected utility maximization or heuristics identified by the Judgment and Decision Making literature. It and similar scholarship may be en route to identifying a new class of market failures where people make poor choices because the presentation or structure of a choice leads their decision-making heuristics awry in systematic, predictable ways. Such a view would urge policy designers to structure incentives, choices, and information flows in ways that consumers' decision making heuristics handle well. This approach would have clear parallels to parts of the marketing literature, but would aspire to create information flows and choice structures that lead to good customer choices that create efficient exchange rather than structures that aspire to increase firm profits, potentially at the cost of inducing customer error.

This project takes on these questions in the context of electricity pricing, and focuses on the challenges posed for implementation with residential consumers. Most customers pay the same price per unit for electric power regardless of when they use it, while the cost of generation varies enormously over time. The existing time-invariant rates squander an opportunity to use prices that reflect marginal costs to manage demand for scarce, expensive resources. Existing scholarship shows that charging the marginal cost of power each hour could achieve social savings on the order of 5-10 % of the cost of the elecriticy system [Borenstein, 2005a]. The first customers deliver the greatest benefits and marginal benefits shrink as more customers participate [Borenstein and Holland, 2005]. This new approach can increase month-to-month bill variation [Borenstein, 2007] and tends to have significant redistributive effects because it reduces cross subsidies implicit in status quo, time-invariant pricing [Borenstein, 2006], but solutions to these problems appear feasible. Finding an approach that benefits the diverse set of electricity policy stakeholders in restructured energy markets is also a significant, open challenge.[1]

This project's first essay improves estimates about how much customers respond to improved pricing. This will help equip analysts to forecast the implications of dynamic

[1] It is striking that a vertically integrated, regulated monopoly, the Southern Company, has deployed dynamic pricing programs more aggressively than utilities where regulators have proclaimed the importance of bringing market forces to energy markets and introduced formal spot markets and competition among generators.

pricing and to decide when and where to deploy it. The ubiquity of prices that diverge from social cost and the difficulty of addressing them makes evidence about potential benefits useful for identifying problems worth addressing and for making the case for the change.

This and previous work have found that residential customers who receive better pricing save money, use less power during periods when power is socially expensive, report high levels of satisfaction with the new pricing, and tend to remain on the program. It is thus surprising that most customers decline offers to sign up for improved pricing. Its second essay uses scholarship at the intersection of psychology and economics to better understand this resistance and to propose novel policy designs that address it while retaining desirable economic incentive properties. The second and third essays further explore the feasibility of implementing such an approach given limited information about customers' demand patterns and a utility's limited flexibility to implement such a program.

The first order policy suggestion is simply to get prices right – by setting them at the market clearing spot price – for everyone all the time through real time pricing. Fixed fees or Ramsey pricing might be necessary to cover the system's fixed costs. Implementing real time pricing with a possible fixed fee is straight forward.

The opt-in dynamic pricing scheme that this project studies turns out to require considerably more analysis. The analysis needs to, among other things, identify customers who provide the greatest benefits and design pricing that attracts them both when they use decision-making heuristics to evaluate whether to try a novel program and later when they experience it and are able to compare their bills and comfort levels on the new and old programs.

Economists' skepticism that many real programs get so complicated that they are inevitably seriously flawed from the start is well placed. But incremental change is often the only feasible way to address flaws in the status quo. Incremental change can deliver real benefits and can benefit from careful analysis using economic tools. A series of well chosen incremental changes can eventually achieve the economic goals that are unattainable with a single broad stroke. Rigorous analysis can identify the most important margins to work on and identify opportunities to make progress there.

1.0.1 The questions in the three essays and major conclusions

This project contains three essays designed to better understand opt-in residential dynamic electricity pricing's experience in the field and to convert insights about its performance into guidance for future policies:

1. How much did residential customers in California's Statewide Pricing Pilot (SPP) field experiment respond to dynamic pricing? What characteristics of days, climate, and households affected the size of this response? How did response vary between days with a routine, weekday afternoon price increase relative to a much higher critical afternoon price? Are benefits from customer response correlated with costs from foregone cross subsidies? What does this imply about designing an opt-in program that will deliver significant benefits while it paves the way for expanded programs? This chapter's econometric analysis finds that some customers responded far more than others. Customers in regions where air conditioning is ubiquitous responded the most to the program and the program had its greatest effect on hot days. This is propitious because dynamic pricing is largely a policy response to costly, air-conditioning-driven electricity scarcities on the hottest summer weekday afternoons. Some obvious implementation strategies would, however, ask many of the most responsive customers to give up significant cross subsidies to participate. Programs can be designed to attract and retain the most desirable customers.

2. Can the judgment and decision making literature explain why most consumers reject dynamic electricity pricing offers that improves incentives and could save them money? This uses insights from this literature to propose a new rate feature. It makes practical, revenue-neutral changes that make the offer more attractive to consumers by reframing periods of critical electricity scarcity as opportunities for customers to earn rebates without changing its marginal incentives. I simulate the implications of these changes for customers in a California pilot program.

3. Essay 2 identified three customer-level constraints that characterize desirable incentive-preserving rebate rates. This, third, essay continues the exploration of whether utilities have the data and flexibility to make customers offers that meet these constraints. Incentive preserving rebate implementations will use limited data to make offers given a limited amount of flexibility to choose offers. We explore the trade offs between sim-

plicity and program performance, and show that simple menus of offers can perform reasonably well.

Essay two presents a strong case that presentation and incentives both matter and that incentive preserving rebates can retain good incentives while improving presentation. Future research should assess their effectiveness with particular attention to whether added complexity negates the benefits of the improved presentation. Essays two and three present evidence that utilities know enough to develop rates that work well for most customers, but may need to differentiate offers by customer consumption level and geography.

Chapter 2

The Impacts of Residential Critical Peak Pricing: Evidence from California's Statewide Pricing Pilot

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

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 can hit price caps of $1.00 / kWh during a few scarcity hours and likely has an even higher social cost during these hours. If customers exhibit even a small amount of demand elasticity for power, then reducing the mismatch between the fluctuating cost and the fixed price could eliminate billions of dollars in deadweight losses [Borenstein, 2005a]. The structure of many wholesale electricity markets makes real time pricing (RTP) a nearly optimal approach to dynamic pricing. RTP sets a price typically for every hour, often as a function of the day ahead spot market cost of power.[1] This paper assesses the impact of a Critical Peak Pricing (CPP) program. CPP

[1] RTP is not quite optimal because it sets prices for hour-long periods a day ahead when electricity scarcities can happen on a second-by-second basis with little notice. In practice, most electricity wholesale markets manage brief or unanticipated scarcities through "ancillary services" markets for reserve, standby generation capacity that are separate from the wholesale power market which it uses to manage foreseeable, extended scarcities. Borenstein [2005b] discusses the selection of the price period granularity and the lag between price setting and the actual system operation in more detail.

is a simpler approach to dynamic pricing that approximates RTP using a small menu of off peak, peak, and "critical", prices. 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.[2] California ran a major field experiment, the Statewide Pricing Pilot (SPP), that exposed customers of its three major investor owned utilities to CPP from July 2003 through September 2004. The SPP generated data about how a change in pricing affects customers' power use on summer weekday afternoons.

In practice, the first generation of residential critical peak pricing programs are and will likely continue to be opt in, so it is particularly important to understand CPP's effects on customers willing to volunteer for CPP.[3] The SPP required CPP customers to agree to participate and allowed them to leave at any time. This paper explores the impacts of critical peak pricing on electricity use in this field experiment to better understand how the new prices change the quantity of power customers use during weekday afternoons. It further tries to understand what kinds of customers and situations lead to the greatest changes in electricity consumption.

It is worth considering dynamic pricing for residential customers because residences consume 36% of the electricity used in the US.[4] Residential demand may be elastic enough to create significant deadweight losses in the absence of dynamic pricing. Residential demand might be more elastic than other customer classes' demand. Reducing power consumption in many commercial and industrial uses involves advance planning, requires temporarily shutting down production, or would annoy customers if it involved reductions in air conditioning or service availability. Those uses are more expensive to change than is

[2]RTP might be a stronger option if it were bundled with technology that automates response to quirky, hour-to-hour and day-to-day price variations. Making this technology easy for users to program to get response that significantly outperforms response to CPP would be challenging. It might also be important to find ways to keep the technology from crowding out manual response during extreme price spikes.

[3]The one kind of universal dynamic pricing program that seems to be politically feasible is a baseline-rebate approach. These calculate customer-specific baseline usage levels. They create incentives for customers to reduce use during critical periods by offering rebates to customers who use less than their baseline amount of power. Universal participation is quite attractive, but baseline-rebate programs have significant drawbacks relative to CPP discussed in Chapter 3. The SPP is the kind of CPP that will generally be an opt in program and is difficult to interpret as evidence about a baseline rebate program. Wolak [2006] is a careful study of a baseline-rebate program in Anaheim California.

[4]2004 figures from the Energy Information Administration's Electric Power Monthly: http://www.eia.doe.gov/cneaf/electricity/epm/table5_1.html

residential clothes drying or air conditioning, especially of empty homes.

California is beginning to roll out advanced meters to every customer and then to offer dynamic pricing programs.[5] Other states and utilities are strongly considering residential CPP.

Careful empirical research can provide information about:

1. whether there are benefits that are large enough to justify the investments in advanced meters and customer education needed for improved pricing programs and

2. about which kinds of customers are most worth recruiting.

There are four central economic questions about the design of practical CPP programs:

1. Do customers respond to the new prices in the short term? How do weather and customer characteristics affect the way they respond?

2. In the long term, does having a large cohort of people on dynamic pricing lead firms to offer products that reduce peak period energy use or that automate response to prices? Will customers invest in these products? Will they become standard in new buildings? Will bundling these products with dynamic pricing make it more attractive and effective? What kind of reputation will dynamic pricing get with customers, utilities, and regulators and how will this affect the programs' growth?

3. The customers who make the greatest reduction in peak period usage may be those who were using the most power during those, now costlier periods under time invariant pricing. Do the rate designs closest to the status quo, like adding peak surcharges and off peak discounts, make participation rational for the customers who would respond the most? Section 2.7.6 finds that the customers who reduced peak and critical period use most in response to dynamic pricing were using a greater percentage of their total power consumption during these high cost periods under time invariant

[5] California's roll out may create a natural experiment that will generate much larger and richer usage data than the SPP did. Some utilities have proposed rates that do not have daily summer peak prices periods. Thus, the SPP data may still be an important source of evidence about the performance of good CPP rates. There are important opportunities to use data from dynamic pricing once widespread advanced metering begins to get better estimates of the performance of opt-in CPP by comparing the usage patterns of customers before and after they joined to the usage patterns of other, similar customers who did not join a dynamic pricing program. Studies of broad-based dynamic pricing could be even more powerful if they combined the billing and usage data with survey data on customer characteristics.

pricing than were less responsive customers. Hence, dynamic pricing implementations can deny desirable customers the significant cross-subsidies that they received under time invariant pricing, leaving customers with increased bills despite their significant response. Analysis in section 2.7.6 shows that simple changes to rate designs can reduce this problem.[6]

4. What are the trade offs between designing a rate that captures much of the variations in the hour-by-hour wholesale cost of power and designing a rate that satisfies customers' preferences for simple, predictable rates that mesh well with their lifestyle?[7]

The present paper directly addresses the questions about who responds and when that are raised in question 1. Question 4 about designing rates that are attractive both to consumers and to regulators and utilities seeking to reduce waste (deadweight loss) in the electricity sector is quite important and is largely unaddressed in the literature, but is beyond the scope of this paper. Understanding firms' (customers') interest in offering (installing) new technologies that make dynamic pricing programs more effective and easier to participate in is quite important, but this and other issues from question 2 are difficult to address in the absence of with long term experience with large scale deployments of dynamic pricing.

2.1.1 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 change in prices. These papers compare customers on CPP to the SPP's

[6] Borenstein [2006] simulates the implications of dynamic pricing for industrial customers and find that it creates broad categories of winners and losers, with some groups losing even if they have fairly elastic demand. It appears that that results for residential customers are qualitatively similar.

[7] We might worry that loss averse customers might reject a rate that raised the cost of evening cooking and climate control even if it lowered their overall bill. It is also likely that the electric demand elasticity of businesses goes up as workers start going home (since it is easier for them to reduce lighting, equipment, and climate control use when offices are mostly empty) while residential demand becomes less elastic. If that were the case, then a practical, efficiency improving policy might raise residential rates in the afternoon and commercial rates in the early evening.

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.[8]

- Herter et al. [2007] looks at the combined effect of climate and temperature[9] 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.

- Chapter 3and Herter [2006b] analyze the same database to consider rate-design related questions, but neither attempts multivariate impact analysis.

This paper adds to the existing literature by using a difference-in-difference regression framework. Its approach 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 the exceptional, critical prices and the ordinary, scheduled weekday time-of-use (TOU) peak price and to explore how temperature and customer characteristics affect response.

Wolak [2006] is careful applied econometric work using a similar data set on a universal participation, baseline-rebate residential program that offers rebates to customers who "conserve" relative to their highest historical usage. Wolak analyzes a program with a different rate and a far more homogeneous sample than the SPP. It finds significant benefits, but also strategic consumer response in the form of increased use during baseline-setting periods that make the customers eligible for more rebates later.

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

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

[9] 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. Electricity demand is a function of the interaction between temperature and climate, since customers on 95^oF days in hot climates tend to run their air conditioners, while customers on 95^oF days in temperate climates tend not to have air conditioners to run.

This paper addresses issues that the existing Statewide Pricing Pilot papers did not address. In particular, the existing papers lack a simple, multivariate analysis that explores factors' relative importance in explaining dynamic pricing customers' reductions in usage. This paper also incrementally improves on the existing models' flexibility, informativeness, and econometric defensibility. I also discover selection issues that may change the interpretation of some of the existing literature.

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

The performance of opt-in critical peak pricing rests centrally on the rate at which customers opt in, the focus of Chapter 3, and on how the new pricing changes participants' consumption relative to what they would have used under the alternative rate.[10] The SPP has a treatment group (I use "CPP group" and "treatment group" interchangeably) and a control group that did and did not experience CPP, respectively. The SPP collected data on each group's usage patterns from before and after the CPP group switched from California's standard, time-invariant rates to CPP. This design measures the behavior of control and treatment (CPP) groups before and after the treatment group experienced a change. The design could lay the groundwork for a clean, powerful difference-in-difference style analysis. Unfortunately the SPP's implementation, like most field experiments, diverged from its design.[11] Getting the SPP into the field was an impressive piece of inter-organizational work that involved negotiation and coordination among three utilities, state agencies, stakeholders, and evaluation contractors. Some members of the working group that oversaw the SPP objected to making CPP nearly mandatory or even putting customers on CPP unless they actively opted out [Charles River Associates, c, 30]. The resulting design that required CPP – but not control – customers to affirm that they wanted to participate makes the treatment and control groups less comparable[12], but has the serendipitous implication of

[10] The evaluation literature (e.g. Diamond and Sekhon [2006]) calls this the treatment effect on the treated.

[11] 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.

[12] An even better, if more expensive, design 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.

making the treatment group a better study of realistic, opt-in CPP.[13] The SPP's implementation reflected a great deal of attention to sampling and to estimating a demand system, but the experiment would have been stronger had it worked to ensure the comparability of the treatment and control groups' recruitment process and a clean delineation between the before and after price change periods by doing things like putting, "New prices begin July 1" in bold in the first few sentences of its mailings. This section describes the SPP's implementation and how its implementation complicates the analysis.[14]

The SPP aspired to recruit treatment and control groups that were representative of the state's population as a whole and comparable to each other, but it recruited treatment and control customers in quite different ways. Potential CPP customers got detailed information and a chance to opt in to the experiment. Control group customers were randomly selected.

The SPP sent potential CPP customers detailed invitation letters [Charles River Associates, b, 18-23] describing the new pricing, study requirements, and $175 in participation payments. Many customers who demand a large amount of weekday afternoon power apparently declined offers to join the CPP group 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." But it goes on to temper the notion that this was a fully rational 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] The CPP group may be representative of customers who would consider opting in to a CPP program.[15]

[13]It is possible that there are minor differences between the way I use the phrase "opt-in" and the way that the SPP final report authors understand it. Regardless of any semantic differences, we share an understanding that 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. 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."

[14]There are also extensive discussions of SPP implementation in Charles River Associates [c,d,a,b,a], Herter [2006a]

[15]Perhaps, more accurately, the CPP group equips us to construct a group that that is representative of the customers who would opt in as a whole. There are good data about how many customers in each

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

Table 2.1: **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]. 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 [c], San Diego Gas & Electric, Southern California Edison, Charles River Associates [b] While the high and low ratio summer rates are reasonably similar, they differ strikingly during the winter. Notably, during the winter, the peak surcharge is 22.3 cents for the high ratio rate and 0.7 cents for the low ratio rate.

The SPP also had a randomly-selected control group that continued to get status quo, time invariant rates. The control group is a stratified sample that can be weighted to represent customers statewide. The central statistical challenge is to use the control group’s

residence category and climate zone refused to join or failed to respond before the experiment contacted a prospective treatment customer who agreed to participate. The SPP’s experience identifies some groups that may be more difficult to recruit than others. The analysis presented here, however, does not use the data about the difficulty of recruiting customers because there are a variety of challenges in interpreting the SPP’s experience as being comparable to an opt-in rate program, including:

1. A full scale roll out of opt-in CPP by a utility that believes that it can capture some of the additional surplus created by dynamic pricing is likely to employ effective, carefully tested marketing materials. By contrast, the utilities’ research on why people refused to participate in the SPP says in part that the recruiting materials were “were quite ineffective [marketing]. .... The materials made scant reference to any benefit - direct or indirect - that the customer might gain by participating....” [Focus Pointe, 6]
2. Potential SPP subjects were offered opportunities to get paid $175 and to contribute to developing better rates for their community in a way that participants in a more conventional rate would not.
3. The SPP’s designers did not have enough field evidence about customer responsiveness to tune rate offers and recruiting materials to attract the kinds of customers who would deliver the greatest reductions in the use of socially expensive power.

Rather, the present analysis imagines the deployment of a statewide marketing effort that recruits an equal proportion of each climate zone and customer class and respond like SPP participants did. Accurate information about how difficult it is to recruit customers in each class is certainly necessary for effective deployment.

behavior to construct a valid counterfactual about what the CPP group would have done in absence of the new prices. The central requirement for the control group to provide this valid counterfactual is that, controlling for observable differences, the expected behavior of the control group is identical to what the expected behavior of the CPP group would have been during the treatment period had they remained on status quo, time invariant pricing.[16]

This section describes the process that created the data and the resulting statistical challenge. The SPP's reports [Charles River Associates, c,a][17] 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. This section briefly describes facts documented elsewhere[18] and describes in depth some issues that have received little previous attention.

2.2.1 SPP CPP Group Time line

This section describes when the SPP collected data about the CPP group relative to its provision of information and incentives that could have affected the CPP group's behavior. It also describes how the SPP's structure shapes some aspects of the analysis.

- Billing data record each customer's Summer 2002 average daily use. It is impossible to know how much of this consumption took place during weekday afternoons. These data from the year before the experiment began are, however, uncontaminated by the experiment.

- The experiment identified potential CPP customers using the sampling strategy described in Charles River Associates [c,a]. The experiment sent letters[19] to potential customers starting in April 2003. These letters described the structure of the new

[16]See Diamond and Sekhon [2006] for an extended discussion of this and an argument that genetic matching using characteristics including propensity scores is the best way to construct a group that provides a valid counterfactual. Although the current chapter does not use a matching approach, doing so is a very logical extension. It is likely that the CPP group is more civic minded, more adventurous, and more price responsive than the control group, but none of these factors threatens the validity of using the control group as a counterfactual if the condition above holds.

[17]The 2003 report [Charles River Associates, a] has more detail on some sample selection issues than does the final report [Charles River Associates, c].

[18]In particular, I do not do justice to the carefully thought out sampling strategies described in Charles River Associates [c,a].

[19]A copy of a typical letter is in the final report appendices [Charles River Associates, b, 18-23].

rates in detail on their second page but only mention the July start date for the new rates on their third page. 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].[20] Efforts to meet recruiting goals and to replace customers who exited continued throughout the experiment.

- The SPP installed "interval" electric meters on participants' residences. The interval meters recorded power consumption every 15 minutes, both during a pretreatment period before new prices began and during the experiment itself. The first 31% of all customers' meters were activated in March through May of 2003.[21]Another 23.5% started reporting data exactly on June 1, 2003 for a total of 54%. The analysis reported in this paper uses pretreatment data collected starting June 1, 2003 because only a select group of customers have data before then. A total of 75% of interval meters were on by July 1, 2003. The analysis reported here considers only the initial cohort of participants whose meters were on by June 15th. This allowing their usage to be measured on at least one pretreatment weekend day and one weekday. CPP customers are included in the analysis only if they experienced the new rate beginning July 1, 2003.

- The SPP mailed the initial cohort of CPP customers "Welcome Kits"[22] starting on June 17-18, 2003. It took between 1 and 2 weeks to complete mailing this batch [Barnes, 2007]. The Welcome Kits included detailed instructions about how to reduce electricity prices during weekend afternoons by changing the use of air conditioners and other major appliances but only mentioned that the new rates go into effect on July 1 on their 18th page. I deal with this by discarding data from between June 18

[20]The experiment listed 10 possible accounts for each slot in the experiment and ranked them from a first choice to a tenth choice and sent out invitations sequentially. "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]). However, section 2.2.4 below documents a couple of possibly important differences between the control and CPP groups.

[21]The first CPP customer meters report data starting on April 23, 2003. The first control group meters came online March 31.

[22]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.

and July 3rd inclusive and further discuss how this affects my analysis in section 2.2.5 below.

- Critical Peak Pricing began July 1, 2003 for the initial cohort of customers. This paper analyzes the initial cohort's response.[23]
- Critical Peak Pricing remained in effect until the subjects opted out, moved, or the available data set ended at the end of September 2004.
- The Summer rates that this paper analyzes were in effect from May 1 through October 31 each year.
- Customers completed a survey[24] "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 whether there were preexisting differences in appliance use habits between the CPP and control groups.[25]

2.2.2 CPP winter and summer rates

The SPP assigned each CPP customer to either a "high ratio" or "low ratio" rate. The winter season high and low ratio rates are quite different, while the summer season high and low ratio rates are qualitatively quite similar.[26] Further, California electricity

[23]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.

[24]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 I analyze more detailed, disaggregated data from the survey than they did.

[25]For example, it would be useful to know whether the CPP group over represents people who tend to work late or people who tend to do laundry on weekends.

[26]The SPP chose this rate design to meet several implementation constraints. The rate had to be revenue neutral for the class average customer over the course of the year while allowing identification with a demand model that considered total daily consumption as a function of the average cost of power over the course of the day. Requiring that the new rate be revenue neutral (i.e. maintain the average price) over the course of the year while changing the average price within each day forces shifting revenue across seasons. Charles River Associates [c, 18-20] discusses this in detail.

demand is summer-peaking, so the scarcity periods, when reducing demand has the highest value, are typically summer occurrences. Thus, the current analysis considers only summer rate data. Both high and low ratio summer rates create moderate incentives to shift use away from ordinary weekday afternoons and strong incentives to shift use away from critical afternoons. Thus, there is good reason to pool the customers into a single CPP "treatment" group and to test whether this pooling is appropriate.

By contrast, the winter low-ratio rate provides almost no incentive to shift usage away from ordinary weekday afternoon peak periods, while the high ratio rate invokes higher time-of-use (TOU) peak prices during the winter season than during the summer. This difference should be central to any analysis of winter behavior. A natural extension of this project would analyze the winter results after interacting each customer's rate with every analysis variable. This interaction is roughly what Charles River Associates [c] and Faruqui and George [2005] do.[27]

2.2.3 Control Customer Time Line

Control customers got a note from the utility indicating that their meter had been replaced and received the survey and follow up mailings, phone calls, and visits until they completed the survey. Other than that, they received no incentives or information to shift their power use. This method allowed the experiment to enroll the first control customer that it considered most of the time, while it had to contact, on average, between two and three potential CPP customers in order to enroll one for the first cohort. Table 2.2 shows that this difference is highly statistically significant.

2.2.4 Systematic Differences between the Control and Treatment Groups

Table 2.2 compares the treatment and control groups on a variety of characteristics.[28] The treatment and control groups are generally quite similar, with a few notable

[27] The approach taken below already estimates separate dummies for the impact of being on high ratio rates during TOU peak and critical hours, which is the minimal acceptable specification. But when the rates suggest that low ratio rates will use their electric heaters more like control customers than like high ratio customers, it seems compelling to use a specification that allows low and high ratio customers to have different sensitivities to cold weather.

[28] Table 2.2 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:

	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 2.2: Mean household characteristics of customers in the regression sample. With a few exceptions that are explored in depth in table 2.3, the treatment and control groups are statistically indistinguishable on observable characteristics. All reported values are weighted by region to give the sample 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. The population reported here is from the "basic survey variables" regression 2 below. This requires that they have meters turned on by June 15, 2003 and have valid answers to the people per household, number of bedrooms, and air conditioning questions.

		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 2.3: 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. I have not run a third interesting hypothesis test about whether any of these characteristics are statistically significantly correlated with a customer's likelihood of leaving the study. 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 2.2, ^{m} indicates that the p-value on equality of means comes from a Mann-Whitney rank sum test conducted on the unweighted, categorical answers.

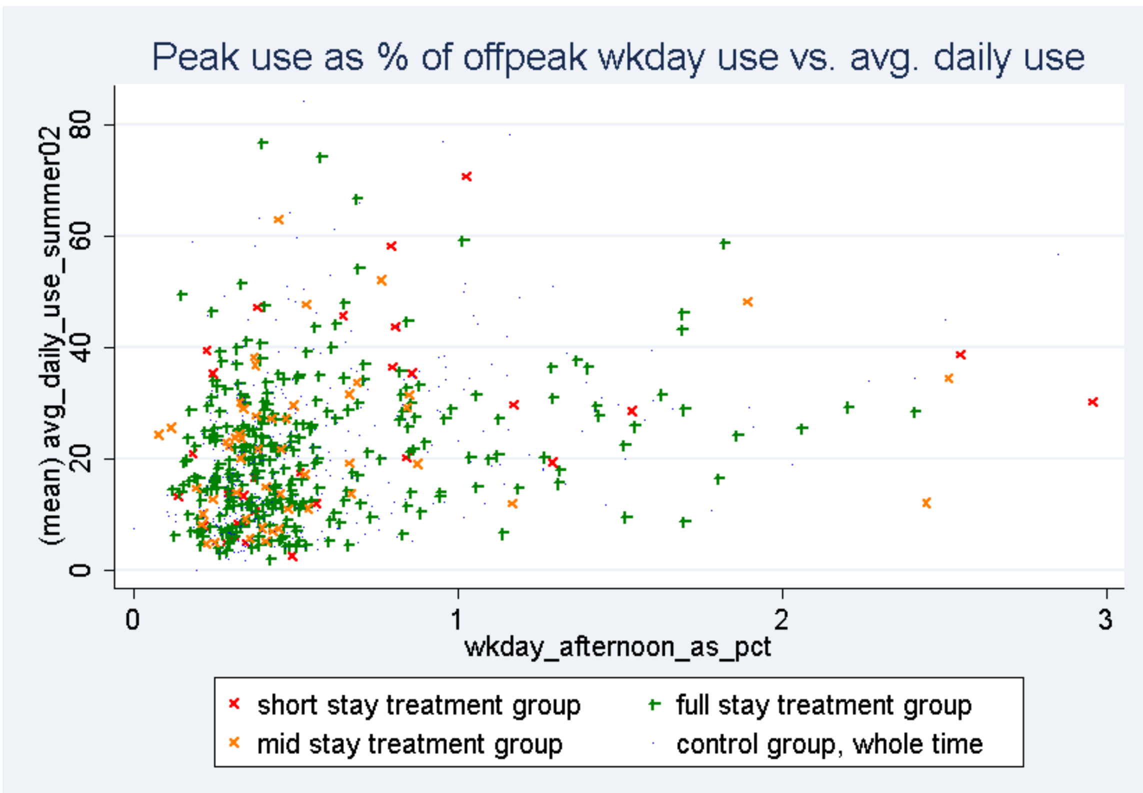

Figure 2.1: *The CPP group has fewer customers who both use a large total amount of power and use a large proportion of that during peak-priced periods than does the control group. There are relatively few treatment customers who used more than about 35 kWh per day and whose peak use was roughly as big or bigger than their off peak use. Thus, there are few treatment customers in the top right part of the scatter plot.*

exceptions, namely:

- The CPP group has fewer high use customers who also use a high 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. Controlling for the percentage of power that customers use during weekday afternoons, larger users are more likely to notice changes in their bills and to believe that their bill changes outweigh the other personal and social benefits of participation in the experiment.

- 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 and advertised through materials that "were 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 having 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.[29]

- The treatment and control groups start with statistically indistinguishable numbers of

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

[29]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.

home businesses but many treatment customers with home businesses exit the sample, creating a statistically significant difference.

- One in 20 differences between samples drawn from an identical population will be statistically significant at the 5% level purely by chance. Thus, this section may offer causal stories about random differences.

Appendix C lists the survey answers for the complete, available sample. The appendix both has more statistical power than the tables of means from the regression sample presented in the main text and may be useful for assessing whether selection problems bias findings in other papers in the literature.

2.2.5 The potential for premature response

The SPP sent CPP customers two detailed descriptions of the new pricing system that did not state the date the new prices took effect boldly, up front. This could have spawned premature reduction in afternoon loads and contaminated the pretreatment data.

- **Invitation letters:** The letters inviting customers to join the CPP group only mentioned that the new pricing went into effect in July, 2003 on their third page. The letter's second page describes the CPP rate with text, graphs, and tables. An SPP report that explored the decision making of people who refused to participate reports that "Respondents were commonly reluctant to take the time to read the pieces word for word, so they resorted to scanning them for information and to gain understanding."[Focus Pointe, 21] Premature customer response would bias impact estimates to understate the changes caused by the new rate. Unfortunately, contamination from premature response to CPP and a selection problem that causes over representation of customers with preexisting, flat load shapes will often look quite similar in the data. The appropriate statistical strategies to deal with these problems are different. 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 2.4 fails to reject the null hypothesis that there is no treatment-control difference in these relationships at the p=.2 level.[30]

- **Welcome Kits:** The Welcome Kits first mention that the rate change takes place on July 1 on page 18. The start date comes long after they explain how to reduce peak period electricity use. The usage data indicates that the CPP group reduced its peak period power use relative to the CPP group's pattern in the last week of June. The first Welcome Kits were mailed June 17, 2003 and mailing continued for a couple weeks.

2.2.6 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:

- I address late Welcome Kit arrival and apparent premature response to the Welcome Kit by dropping all data from June 18 through Thursday, July 3, 2003 inclusive.
- To deal with the gradual deployment of meters, I focus on the initial cohort of customers which started CPP on July 1, 2003 and had meters on by June 15, 2003. Further, I exclude pretreatment data from before June 1, 2003.
- Weather data for the PG&E region are missing for August, 2003. Hence, I drop those data.
- I only consider data from the Summer rate season.

Some problems with this data set are impossible to fix. The unfixable problems include the fact that customers knew too much about the coming prices when they decided whether to sign up and that appliance use habits were surveyed after the price change.

[30] Another analysis that could look at this issue would look at whether use drops and rebounds right at the edges of the 2:00 to 7:00 PM peak period in an effort to distinguish customers who were responding early from customers who had a preexisting flat load shape, but this might only spot customers who had programmable thermostats or who were at home to make manual changes to power use on weekday afternoons.

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 2.4: 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 group of authors involved in running the SPP report that they decided not to use pretreatment period data as a control (and hence to only estimate the impacts of the critical price beyond the impact of the daily TOU 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]. This paper takes an intermediate approach.

2.3 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.[31]

- 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.

- Adds a more detailed set of controls for day and weather that are not interacted with treatment status.

I estimate:[32]

$$avgLoad_{it} \quad = \quad \alpha^T \mathbf{X}^* + \delta^T \, TrtCustomer_{it}\mathbf{X}^* + \gamma^T \mathbf{T}_{it} + \kappa^T \, TrtPeriod_t \mathbf{X}^* +$$

[31] 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.

[32] Bold characters and Greek characters are vectors.

$$\nu^T \mathit{CriticalPeriod}_t \mathbf{X}^* + \beta^T \mathit{PeakPrice}_{it} \mathbf{X}^* + \psi^T \mathit{CriticalPrice}_{it} \mathbf{X}^* + \epsilon_{it}$$

Where:

- $avgLoad_{it}$ is customer i's average kW (i.e. kWh/hour) consumption from 2-7 PM on weekday t.
- $\mathbf{X}^* = \begin{bmatrix} 1 \\ \mathbf{X}_{it} \end{bmatrix}$ 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.
- $\mathbf{X}_{it}$ is a vector of customer specific, and sometimes customer-day specific, controls drawn from 1) data about the customer from the SPP classification and billing data system, 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 question, and 4) data about the customer's hour-by-hour usage during the pretreatment period.
- 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 during the period in which the first cohort of treatment customers got CPP, namely all days after July 1, 2003. It is zero before July 1, 2003.[33]
- $PeakPrice_{it}$ is 1 if customer i received a TOU peak, non-critical price on day t and is zero otherwise.[34]
- $CriticalPeriod_{it}$ is 1 if a critical event was declared for CPP customers in i's climate zone on day t and is zero otherwise.

[33] This analysis drops the cohorts of "late starting" treatment customers whose experience with CPP began after July 1, 2003 or control customers whose meters were activated after June 15, 2003.

[34] Making the $PeakPrice_{it}$ variable zero on critical days makes the standard errors on critical impacts easy to interpret, which is useful in section 2.4 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 TOU Peak and critical difference was statistically different from zero than to have the opposite situation.

- $CriticalPrice_{it}$ is 1 if the utility successfully notified customer i that t was a critical day and is zero otherwise.

2.3.1 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:

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

Where the variables are as above except that:

- $\mathbf{X^w}$ is the weather condition subset of $\mathbf{X}^*$ for customer i on day t.
- $Customer_i$ is an array of one variable per customer where variable $j \in 1..N$ is 1 for customer i if $i = j$ and is zero otherwise.

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] If we get evidence that these estimates are stable, then we can explore the impacts of different subsets and weightings of the customer population.

[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.

- **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.

2.3.2 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. Further, I downweight 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.[37],[38] This addresses the fact that, as Charles River Associates [c, 22-32] describes, 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 July, August and September in both 2003 and 2004, but just May 2004, June 2004, and October 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.[39]

5. Has a sample that changes over time. Subjects come and leave.

[37] I do not down weight the first few days of July because we drop it for 2003. And I do not down weight PG&E data from the month of August 2003 because missing weather data necessitated dropping it.

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

[39] There appears to be exactly one calendar day, August 27, that was a critical day in both 2003 and 2004.

2.3.3 Possible implications of the selection problems

This project's difference-in-difference approach compares the electric-use trajectory of control customers to that of CPP customers. Selection problems can mean that we are constructing the wrong counterfactual by using a treatment group that would have followed a different trajectory than the control group even in the absence of dynamic pricing. It is conceivable that selection bias could point in either direction:

- The people who opt in to CPP could be simply less sensitive to weather during weekday afternoons. We should be concerned that this is the case because customers who are weather-insensitive enough tend to save money on dynamic pricing. This implies that the analysis is using behavior during the relatively mild month of June as a baseline and then measuring impacts from treatment months that were, on average, hotter. If the group that switched to CPP were already less sensitive to weather changes and the 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.

- If the people who opt in to CPP tended to air condition only during very hot weather,[40] then the customers' might become much more weather sensitive during the hottest months. This might give CPP customers observables, like average daily use values or average consumption fixed effects, that would make them look like untreated customers with smaller residences. Smaller control customers' electric use would rise less than theirs later in the summer. Therefore comparing CPP customers with smaller control customers would understate CPP's impact.

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

Tables 2.6 and 2.5 report results from four specifications based on the econometric approach described in section 2.3. The regressions in tables 2.6 and 2.5 measure temperature in cooling degree hours and cooling degree hours squared and force customers from every region of the state to have the same relationship between temperature and electricity use.

[40] This kind of behavior will reduce weekday afternoon electricity use relative to use during other periods. It makes them more likely to come out ahead on the new rates. Being extremely frugal with weekday afternoon air conditioning is, however, not sufficient to ensure that a customer will pay less on CPP.

Appendix K reports the results of an improved set of regressions that fit a more flexible piecewise-linear spline to the relationship between temperature and energy use rather than the rigid quadratic form.[41] Further, they interact these splines with a variable that is 1 if a customer is in climate zone 1 or 2 where air conditioning is rare and zero if the customer is in a hotter climate.[42] Future revisions of this document will likely further improve the regressions in Appendix K and use those regressions instead of the quadratic regressions throughout the document. The regressions are as follows:

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

2. Specification 2 expands specification 1 by controlling for factors like whether the customer has 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 "everything but the kitchen sink", namely a large number of survey variables, duration of participation category dummies, 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

[41] The piecewise-linear spline allows the relationship between temperature and energy use to change at each member of a set of temperatures K. It estimates the additional impact of each degree above each "knot" K by creating a variable of the form $SplineCDH_k = \max\{0, CDH - K\}$

[42] The quadratic regressions (Specifications 1-4 below) 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 doing so by adding more than 100 regressors would use too many degrees of freedom.

dummy variables reporting that either the TOU peak price or critical prices was in effect for that customer on that day. Hence tables 2.6 and 2.5 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. Appendix D reports the complete results from the regressions. Appendix E repeats the complete results running the regressions on data from the two hottest climate zones, where summer days above 90^o are common and where more than 70% of customers have air conditioners. Appendix K presents a single regression that estimates the interaction of temperature and dynamic pricingly for climate zones 1 and 2 separately from those in zones 3 and 4. Hot summer climates are typical of many parts of the US.[43]

2.4.1 The Benefits of Dynamic Pricing Grow as the Temperature Increases

Multiple sources of evidence find that dynamic pricing leads to bigger reductions in electricity use on hot days. This is good, if perhaps unsurprising, news because dynamic pricing aspires to dampen energy use, especially air conditioning use, during hours when air conditioning demand makes electricity is scarce and expensive. The additional benefits become evident when the temperature reaches something between 85 and 100 degrees Fahrenheit. Air conditioning uses a great deal of electricity and causes demand to peak on the hottest summer weekday afternoons. Thus an increase in dynamic pricing's impacts during hot weather is propitious. It creates a positive correlation between the program's reduction in electric use and the power's wholesale price and marginal social cost that define the value of those reductions. So climate control, largely air conditioning, probably drives the increase in benefits from dynamic pricing during hot weather.

Graphs 2.2 and 2.3 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

[43] 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. (Sacramento's municipal utility district did not participate in the SPP, but PG&E assigned some of its Sacramento-area customers to a "Sacramento" weather station.)

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 FE's; 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 2-7pm	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), 2-7pm	-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 2-7pm * central AC	. .	. .	0.000017 (0.003)	. .
Critical Price in Effect * 2-7pm squared * central AC	. .	. .	0.0000065 (0.00093)	. .
Critical Price in Effect * swimming pool	. .	. .	. .	-0.289 (0.196)
Critical Price in Effect * cooling degree hours 2-7pm * 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 FE's: fixed effects				
Cooling degree hours (CDH) are base 78^o F. Heating degree hours are base 65^o F.				

Table 2.5: The impact of just 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 FE's; 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 2-7pm	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), 2-7pm	-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 2-7pm * 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 2-7pm * 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 FE's: fixed effects				
Cooling degree hours (CDH) are base 78^o F. Heating degree hours are base 65^o F.				

Table 2.6: The impact of just TOU peak prices

cooling degree hours (CDH)[44],[45]. The graphs suggest that dynamic pricing customers use less power at temperatures beyond that threshold and that these savings grow as the temperature increases. Graph 2.4 shows the pretreatment relationship between temperature and use. It and graph 2.3 illustrate the regressions' identification strategy which compares differences-by-temperature between the pretreatment and the period when CPP was in effect.

Tables 2.6 and 2.5 report qualitatively similar results from regressions. The regressions and graphs in Appendix K report a more flexible and accurate estimate of this relationship. The regressions in the main text estimate a quadratic relationship between impact and temperature, while Appendix K estimates a piecewise linear relationship and presents pictures of non-parametric estimates as well. The most authoritative evidence available comes from figures K.1 through K.5 in Appendix K. Both the regressions reported here and the work in appendix K suggests that the benefits of dynamic pricing are small and insensitive to heat at low temperatures. The quadratic regressions report that temperature-driven benefits were near zero below an afternoons with an average temperature with a point estimate that ranges from 88^o to 100^o F (i.e. 50 to 110 base-78 cooling degree hours [CDH]). Benefits appear to grow rapidly as temperatures rise beyond that point. Appendix K suggests 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 them attempting to precool their homes. The quadratic findings suggest that an increase in the average temperature from 98^o to 99.9^o (an increase of 1000 CDH^2) reduces the CPP group's load by .13 kilowatt (kW) relative to the control group's load[46], saving .67 kWh per customer per afternoon.

[44] 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 probably reflects an afternoon that had an average temperature of 88^o (90^o) F.

[45] I, in contrast to Faruqui and George [Charles River Associates, c, 41], find that using base 78^oF CDH gets better fitting results than does using base 72^oF 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 zeros in temperate climate zones where air conditioning is rare than would working with base 72^o or 75^o CDH. Further, a great deal of air conditioning is likely to be idle at temperatures between 72^o and 78^o, but active at higher temperatures.

[46] 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). 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. In this paper, I repeatedly multiply the flow of average impacts per hour that are measured in kW by five to

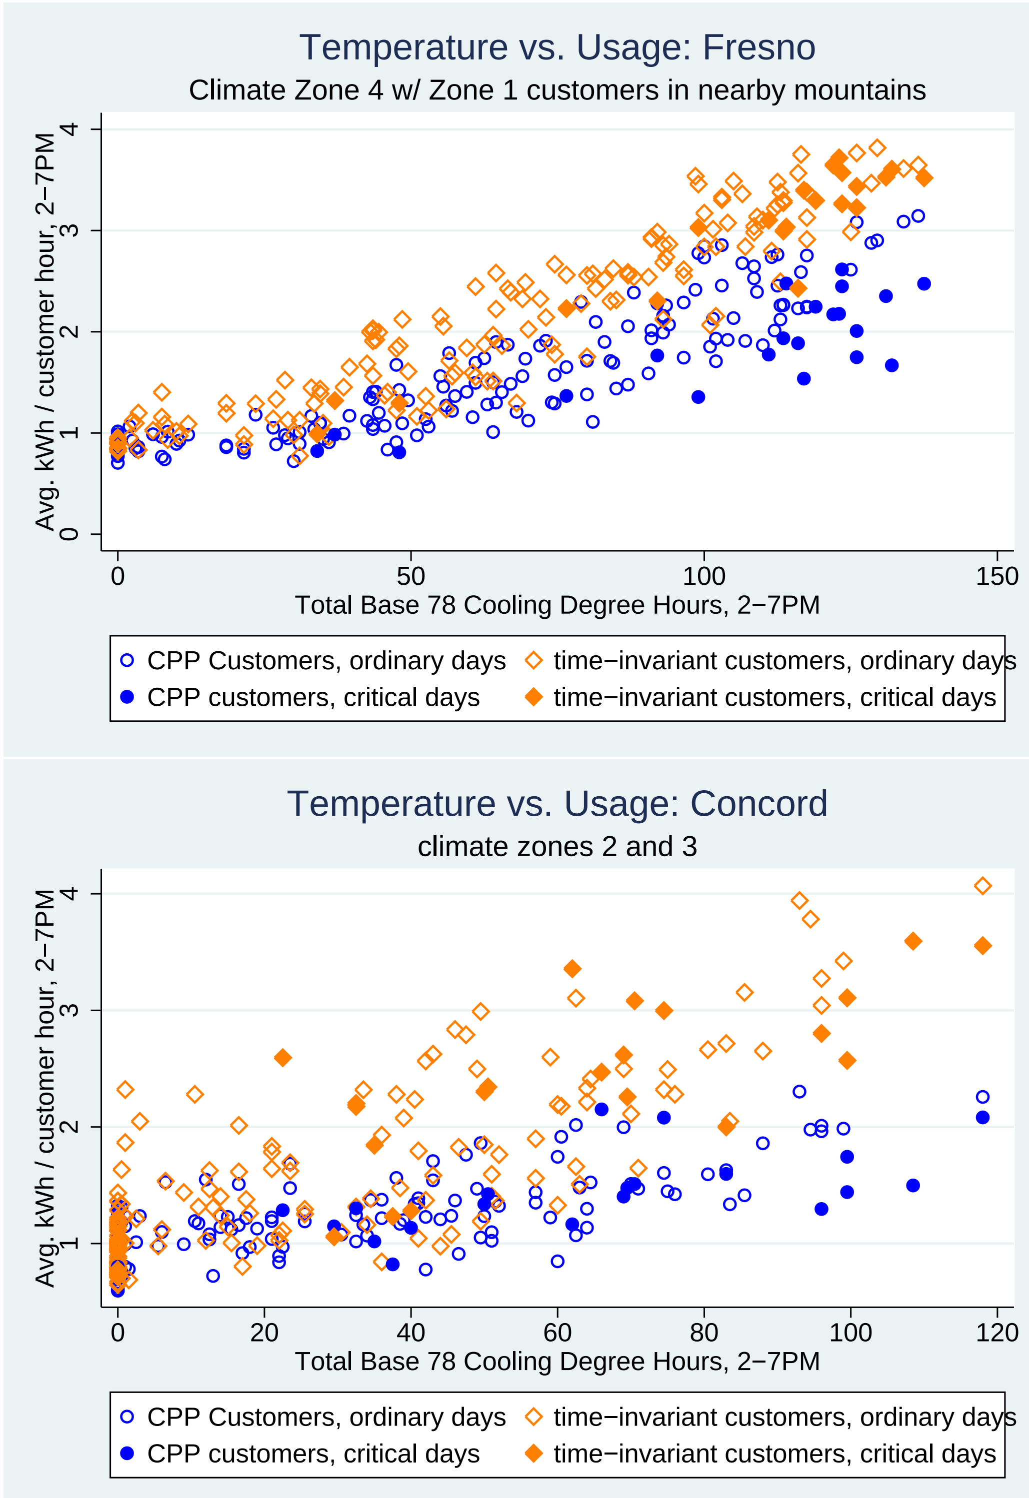

Figure 2.2: These figures plot the daily average hourly total electricity usage per customer between 2 and 7PM as a function of how hot the day was. They show that the difference between the average use of treatment and control customers is bigger on higher temperature days during the treatment period. 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.

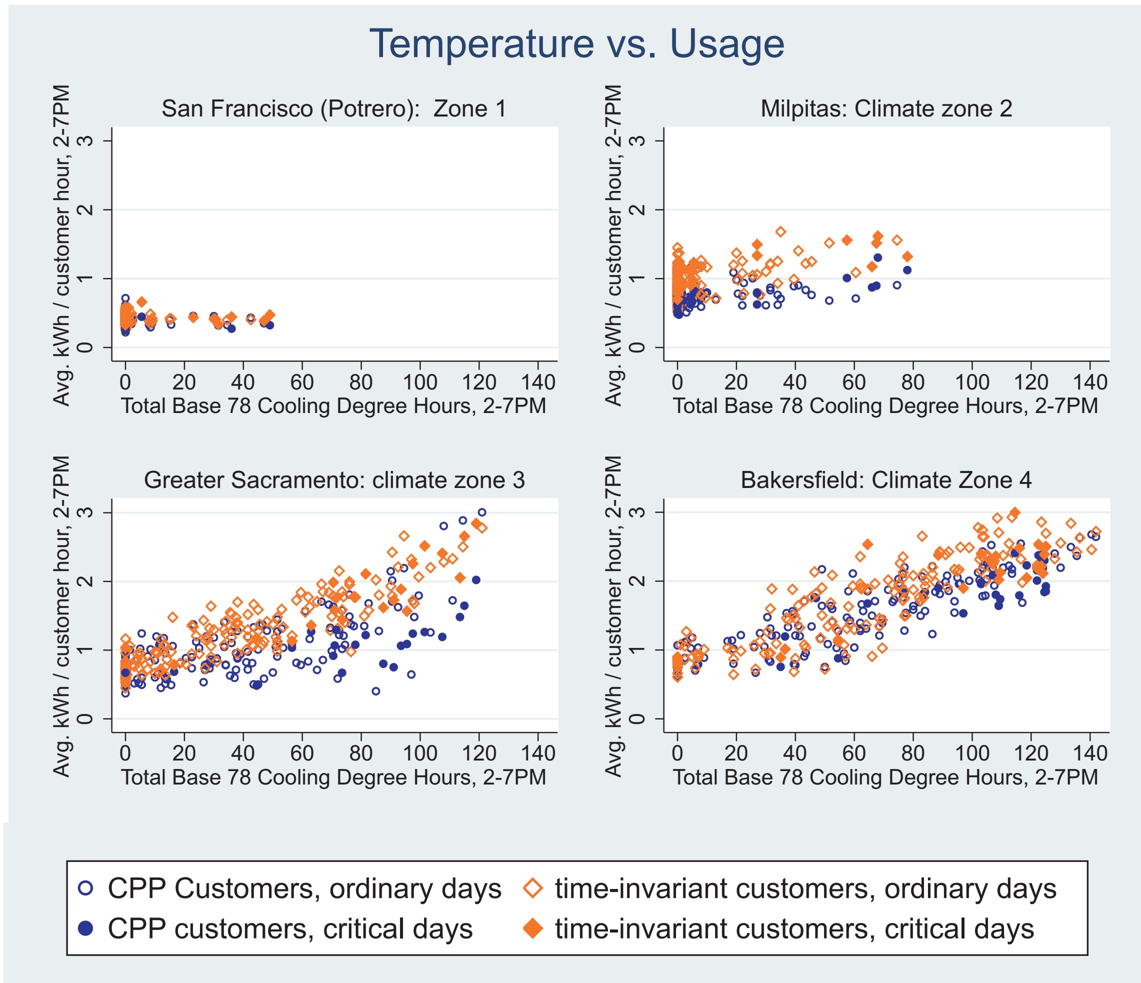

Figure 2.3: These graphs show the relationship between heat and average daily use for customers of one weather station in each of the four climate zones. Notice several important aspects of the identification: The vast majority of days with more than 80 base-78^oF cooling degree hours (i.e. days that 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 2.2 add to important findings from Herter et al. [2007]. First nearly all of the impacts of critical events on days that are hotter than 95 degrees appear to come from climate zones 3 and 4. Second these figures, especially figure 2.2, suggests that the substantive importance of response to the TOU 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.

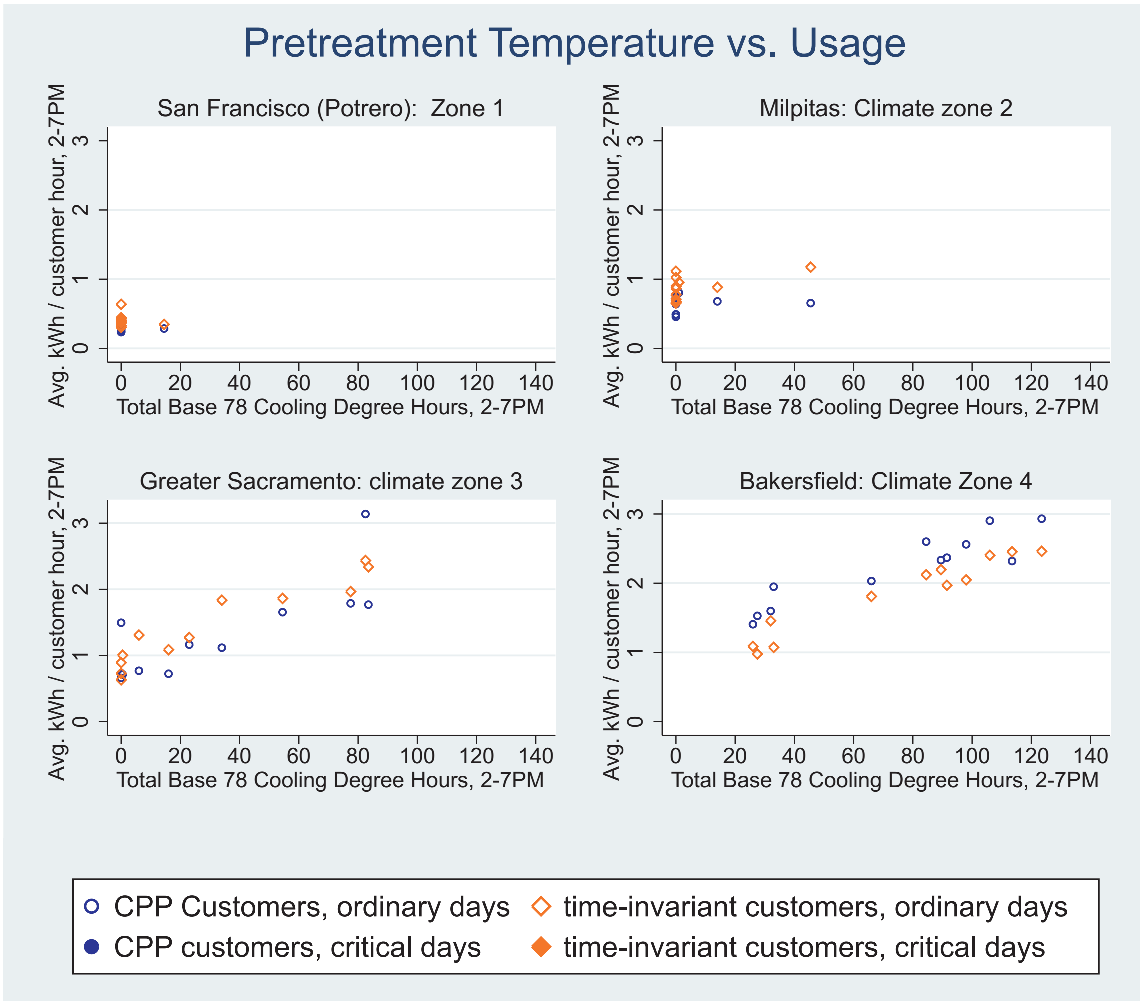

Figure 2.4: Identifying any pretreatment (i.e. weekdays June 1- June 17, 2003) differences. These graphs show the relationship between heat and average daily use for customers of one weather station in each of the four climate zones. The limitations of the identification strategy for this data set become 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.

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 during for customers experiencing TOU 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 difference-in-difference results from this work and from Appendix K seem most consistent with the interpretation that estimates found that treatment customers used, on average, far less power during moderately high temperature pretreatment afternoons than did the control group. The small size and strange structure of the pretreatment sample seem to be causing real problems here. The moderately high temperature pretreatment data come from a few weather stations on a few days when some customers' meters were being activated. These factors are likely to have created idiosyncratic correlations between treatment status, temperature, and use in the pretreatment data. Then, this difference all but disappeared during the treatment period, yielding difference-in-difference estimates that imply that the treatment caused an increase in power use during these moderately hot conditions. Regressions that consider only the treatment period difference between the control and treatment group estimates do not corroborate this finding and yield impacts that are substantively small and temperature insensitive below 60 cooling degree hours per afternoon. The treatment-period-only difference finding is a much more intuitively appealing pattern. [47]

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 TOU peak days.

get the average impact per afternoon peak period in kWh.

[47] The positive sign on the linear CDH term might achieve the best quadratic fit to a relationship where temperature has almost no effect at low temperatures and a very significant impact at high temperatures. The positive linear term thus keeps estimated impacts close to zero at low temperatures while allowing the quadratic term to fit the significant relative reductions in use at high temperatures. It could, less plausibly, correctly indicate that dynamic pricing is counterproductive on cool days.

% customers who have central air conditioning					
	climatezone				
housetype	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	
% w/ any cooling technology: room, central or building-wide AC, evaporative cooling					
all	12.8	47.5	88.4	94.7	

Table 2.7: 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 these zones as well.

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 TOU peak price and critical peak price both increase by a factor of four. This renders the point estimates of their coefficients statistically insignificant.[48] 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 TOU peak price with cooling degree hours squared falls significantly toward zero.

Table 2.7 suggests that some of 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 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

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

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.

The survey collects no data about the size, age, or efficiency of air conditioners or the quality of insulation on a house. This technical information would be valuable to researchers, but few residents know it. So it may be that the survey variables fail to capture the information about air conditioning equipment we need to understand how differences in equipment and controls affect the price sensitivity of energy use.

Appendix Tables 2.8 and 2.9 show that when we split the customers into low-use single family / apartment and high-use single family halves, the point estimate of a strong negative relationship between cooling degree hours and dynamic pricing benefits comes back among the high use single family customers. This negative impact is, however, quite imprecisely estimated. Appendix E gets the same qualitative result when we restrict the regression universe to climate zones 3 and 4. This set of findings has three possible interpretations:

- High use (hot climate) customers are generally heavy air conditioning users, so the high use (hot climate) and low use (temperate climate) customers could be generating substantively different results.
- The new point estimates with even larger standard errors could be statistical flukes.
- The high use subset of the customers as a whole, including the high use/hot climate customers bear the brunt of the selection problems discussed above. Thus, the selection problems that are more pronounced in these two subsets might be driving the finding that the customers in them responded more to dynamic pricing during high temperature periods.

Accumulated Heat

Specification 4 controls for the number of cooling degree hours on each of the three previous days and 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. This translates to an average savings of .2kW each hour.

2.4.2 The Effect of Customer Size on Response to Dynamic Pricing

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.[49] We get smaller, statistically insignificant point estimates from these three specifications when the TOU peak price is in effect.

This finding is important because utilities have billing data, but often do not know the demographics or appliance holdings of each household. This finding means that utilities can target high-use customers knowing that they will respond more on a per-account basis to critical prices than lower use customers. The results also suggest that the overall electricity use variable serves as a proxy for some of the ability to respond that more detailed data would uncover.

2.4.3 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.[50] 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, TOU peak prices after a critical event which is consistent with delaying optional activities like drying clothes and

[49]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.

[50]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 effect. The sign reversal suggests that we may be picking up impacts of things like building age, building size, or number of occupants.

running the dishwasher.[51],[52],[53]

- 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.[54]

- 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 TOU-peak (critical) price was in effect.

2.4.4 Comparing the Effects of Different Dynamic Rates

The high ratio rate charged higher prices, but the pooled sample yields no statistically significant evidence that they led to greater conservation during peak and critical periods. Most regressions find that the TOU peak period and critical price period use reductions by customers on high ratio rates differ from the reductions by low-ratio customers by a statistically and substantively insignificant amount.[55] The regression results in tables

[51] 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

[52] 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 Appendix K, which suggests that it came from an insufficiently flexible functional form for the temperature relationship.

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

[54] 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 the true spending power of retirees who may be more likely to have savings that far outweigh their incomes.

[55] 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

2.8 and 2.9 and appendix F find modest evidence that high-ratio, high-use single family customers responded more to the TOU 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.

These findings are not easy to reconcile with phone calls being irrelevant and there being a single, gently sloping (e.g. continuous elasticity of substitution) demand curve in customers' heads. These findings are consistent with customers having demand curves with regions that are near vertical. These curves would mean that a change from 20 cents to 70 cents makes a big difference in quantity demanded, but that a change from 20 to 25 or 50 to 70 cents does not make much of a difference. These patterns are also consistent with customers thinking about CPP as designating times to "use power normally," "conserve a little," and "conserve a lot."

2.4.5 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. Specification 4 uses customer fixed effects instead of identifying coefficients for preexisting differences between the control and treatment groups. 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.

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

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.

and the more pristine low-use single family and apartment customers. Tables 2.8 and 2.9 [56] and Appendix F report the 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, the point estimates of statistically significant findings from the section above retain their signs and magnitude. 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%) halves.[57]

2.5 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[58]. Electricity demand is closely correlated with temperature and most scarcity conditions take place on days when high temperatures create extreme demand.[59] Further, the results above find that the impact of dynamic pricing is quite sensitive to the temperature of the day. This section assigns days to bins by their peak 2-7 PM California Independent System Operator (CAISO) control area electric load[60], 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

[56]The results in Tables 2.8 and 2.9 are chosen to support the discussion in section 2.4.1. Appendix F reports complete results from running all four specifications.

[57]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.

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

[59]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.

[60]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.

	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), 2-7pm	-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 2-7pm * central AC	. .	. .	-0.00013 (0.003)	-0.003 (0.005)
TOU Pk Price in Effect * cooling degree hrs 2-7pm 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 2.8: The impact of TOU 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 base 78^o F. Heating degree hours are base 65^o F.				

Table 2.9: The impact of critical prices when we separate apartment / low use single family customers from high use single family customers

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.

2.5.1 Summer Season Weather and Load Patterns in 2003-04

Facts about California's 2003-04 weather, population, and electrical demand are useful in understanding the results below. The SPP assigned each customer to the four climate zones presented on a map at Charles River Associates [c, 22]. This section takes the SPP data that assigns each customer to one of the 58 weather stations around the state listed in Charles River Associates [d, 18-19] and uses them to calculates population-weighted temperatures. Appendix G describes this paper's population-weighted temperature calculation methodology.

1. Statewide population-weighted base-78 cooling degree hours (CDH) have a very strong positive correlation with electricity use. Figure 2.6 illustrates this.

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 other two zones (which tend to be fairly cool and quite hot most of the time, respectively). This combination of variable weather and large population mean that zones 2 and 3 appear to drive the variation in 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. By contrast, roughly 57% of all days May through October had fewer than 35 CDH 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 2.5 illustrates this relationship.

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 afternoon temperature was at least 90^o between 2 and 7PM.

5. Climate zone 1 (the coastal fog belt) is generally quite cool, with a few warm days.

6. Quite high demand days tend to reflect a confluence of unusually hot weather in the hot inland zones 3 and 4. The two highest demand days were also unusually hot in zone 2.

7. Appendix H further details these patterns by providing tables of the distribution of cooling degree hours by climate zone.

8. 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 TOU peak and critical prices under the hottest conditions seen in 2003-04.

2.5.2 The experiment period lacked extremely hot days

The SPP ran during two years that lacked the kind of extremely hot days that tend to create extreme demand peaks. This limits our ability to explore the impacts of dynamic pricing during the truly unusual demand events when reductions in power demand have

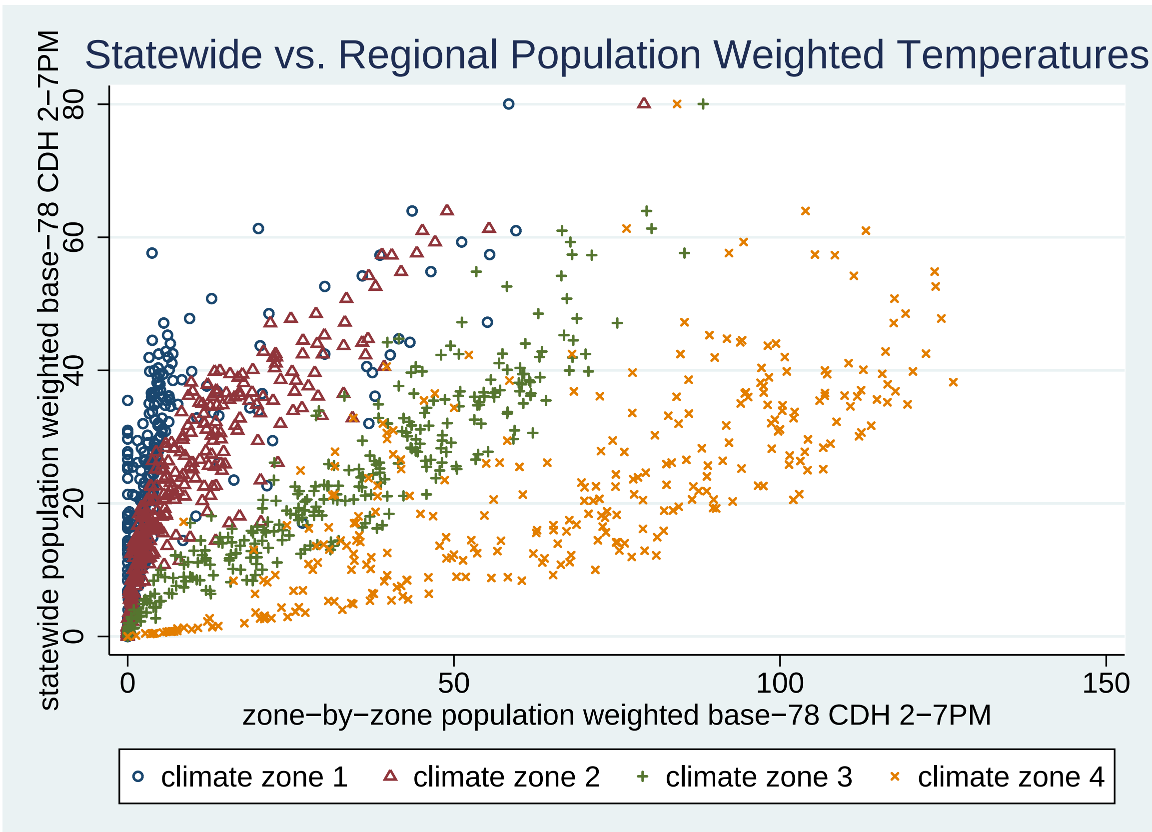

Figure 2.5: *The relationship between population-weighted average temperatures in each climate zone and the statewide population-weighted average temperature. Notice that each y-coordinate (a day with a single population-weighted average CDH) has an entry for each of four zones with x-coordinates indicating the population-weighted average temperature. The very top entry makes this clear. This graph includes both weekdays, holidays, and weekends.*

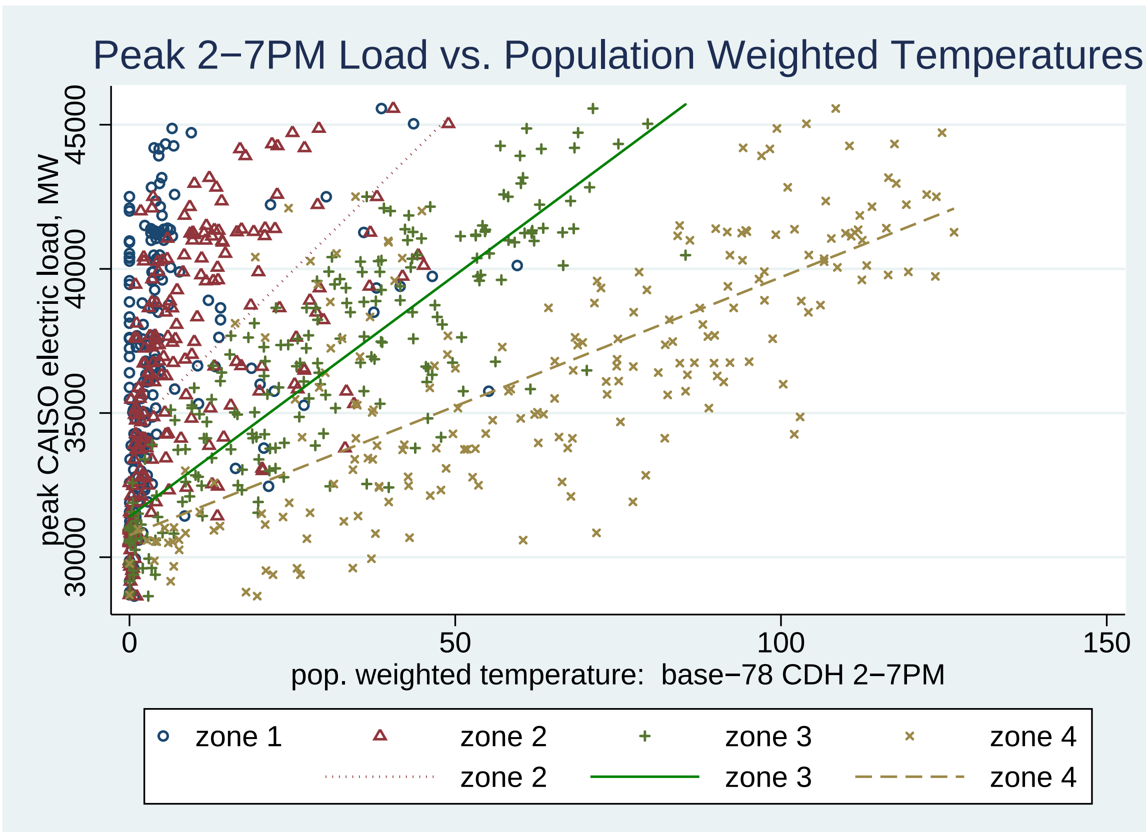

Figure 2.6: *The relationship between weekday zone-average afternoon CDH and maximum electric load in the CAISO system between 2 and 7PM. I have plotted the best linear fit of the 1 variable regression of each zone's CDH on load. Zones 3 and 4 fit quite well; zone 2 shows significant signs of omitted 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.*

the greatest benefits.[61] 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.[62] By contrast, CAISO declared emergencies on six summer-season days in 1998 and three in 2006. 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] 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 limits our ability to explore the impacts of dynamic pricing during the truly unusual demand events when reductions in power demand have the greatest benefits.[63] By contrast, record setting heat hit California, especially the heavily populated climate zones 2 and 3, in July, 2006. Record demand caused a Stage II Power Emergency on July 24, 2006. The 99% percentile demand in a pooled sample of 2003 and 2004 was almost indistinguishable from the 99% percentile of demand in 2006, but CAISO peak demand was 50.2 GW at 4PM on July 24, 2006, which eclipsed the 45.6GW peak in 2003-04.

There is good reason to think that July 2006 got far hotter than any day in the SPP sample, although the author does not have directly-comparable, population-weighted weather data for 2006. It, however, is illustrative to note that record setting day had 76, 149, and 161 base-78 CDH at airports in San Jose (zone 2), Sacramento (zone 3), and Fresno (zone 4) respectively. These are at least 30 CDH higher than the hottest population-weighted weekday temperatures observed in zones 2-4 in 2003-04.

[61] 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. It had an unspectacular peak demand of 38.3GW because so much commercial and industrial demand was offline for the weekend. As a Sunday, it is not in the impact analysis.

[62] 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.

[63] 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. It had an unspectacular peak demand of 38.3GW because so much commercial and industrial demand was offline for the weekend. As a Sunday, it is not in the impact analysis.

This paper pushes its exploration of the impacts of dynamic pricing during very hot weather by creating a "synthetic" hottest TOU peak priced day that combines data from the day with the highest population-weighted average temperature in each climate zone. It repeats the process to create a synthetic hottest day from the sample of critical-priced days. This yields nearly the hottest in-sample weather for each price level.[64]

2.5.3 Calculating Total Impacts from the Econometric Results Above

The estimation strategy described in Section 2.3 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.[65]

- **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.

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 TOU peak

[64] We could get marginally hotter in sample data by independently selecting the hottest day at each of 58 weather stations rather than the independently selecting the hottest day in each of 4 zones. Creating that day will yield no great surprises. That day would would marginally increase estimated benefits by allowing us to walk slightly further on a quadratic relationship that reports that benefits increase in temperature.

[65] 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.

afternoon prices.[66] Appendix J 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:

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

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

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

Table 2.10 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.

[66] 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 2.10: Point estimates of the total impacts of dynamic pricing by climate zone and load scenario. Significance: *=10% ** =5%

- The point estimates of the TOU 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 TOU peak customer-day. These estimates are imprecisely estimated during both critical and TOU 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 2.4.5 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.[67] If this preexisting difference is, in part, a premature response, then 2.10 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 Appendix K 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. Appendix E drops zones 1 and 2 and finds that the weather sensitivity grows slightly, but that the results are qualitatively nearly identical. Further, converting statewide data to base-78 cooling degree hours means that we are largely fitting the temperature curve to data

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

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% more benefits than they delivered during days that were just a few percentage points lower in the load distribution. The figures in Appendix K strongly suggest that the parabola we fit overstates 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 K.1, K.2, and K.4 in Appendix K 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 Appendix K shown in figure K.4 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 Appendix K piecewise linear estimates.

2.6 Comparing this approach to existing papers

The analysis that is closest to this paper's work is the project 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. Section 2.4.4 above and Herter [2006a] find evidence that the demand curve, rather than being smooth, has some significant bends.

- 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.

2.7 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 explores the implications of cross subsidies for the design of rates that can attract responsive customers. Customers' apparent insensitivity to small changes of peak and critical prices has implications for rate design.

2.7.1 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.[68] We can maximize the reductions in the cost of the electricity system by signing up the customers who offer the most beneficial change in demand multiplied by the locational cost of power per unit of recruiting effort.[69] The analysis here reports that this implies targeting large customers in hot climates if 1) locational prices are the same everywhere and 2) the cost of recruiting customers of every type is the same. Future work could consider whether and how differences in recruiting costs and locational energy prices change the optimal targeting.

Targeting high value customers is particularly valuable if enrolling each customer requires an expensive meter installation. Places like California are, however, deploying advanced meters for customers on all rates.

One propitious implication of this analysis is that the most important variables in determining which customers to target are regional temperature and a customer's historical use. These are both readily available to utilities.[70]

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 demonstrate that the plan can maintain long term customer satisfaction and effectiveness for a

[68]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.

[69]These 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.

[70]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%.

diverse group of customers in all the climate zones that might be part of future programs. Thus, a good midterm approach is a compromise between targeting highly responsive customers to increase programs benefits while recruiting a customer pool that is representative enough of the likely participants in future programs.

2.7.2 Benefits: primarily, not exclusively, air conditioning

The SPP's Summer-Fall 2003 survey found that treatment customer's report being far less likely to use dryers and dishwashers during peak hours than do control customers. Similarly, treatment customers report being much more likely to actively manage their use of air conditioners. The electric use data find evidence of differences in air conditioning management, but cannot detect differences in major appliance use. This probably reflects the fact that major appliances consume a smaller share of energy than does air conditioning. For example, shifting a once a week, 5 kWh dryer load from peak to off peak hours would yield an average benefit of .2 kWh per hour over the 25 weekly peak hours. This magnitude of impact is generally too small to be statistically significant in this analysis. Further, only about 31% of the treatment group even owns an electric dryer, so we would expect that if those 31% shifted an average of one whole load per week, that we would see about .06 kW per customer decrease in peak load. But, if the SPP sample generalizes statewide, then there are well over 2.5 million electric dryers in California, and that shifting a fraction of a load per dryer per week from peak to off peak periods could lead to quite substantial social savings. Post-SPP customer surveys indicated that most customers found it easy to shift major appliance use, but that many found shifting air conditioning use difficult and unpleasant [Lineweber, 2005, 5,14-15]. This suggests that rescheduling major appliance use offers a greater deadweight loss reduction per kWh shifted than does changing air conditioning habits.

2.7.3 Targeting seasons with the greatest potential social value

The largest potential social benefits from dynamic pricing appear to come from providing good incentives during hot periods that make electricity scarce and customers price sensitive. If there is a simple, feasible way to target weather conditions that seems fair to customers, that approach may be compelling because some of the periods that create the highest costs and greatest reliability problems are unseasonably hot days in the spring and fall when generators that are only online during the summer peak season are not available.[71] Advocates may suggest simpler rates with more muted peak-offpeak differences that largely preserve existing cross subsidies and reduce the risk that customers will feel inconvenienced or experience bill spikes. A potentially elegant compromise could impose the largest difference between daily TOU peak and offpeak rates during the summer, while allowing critical days to be called for unseasonably hot spring or fall conditions. A customer-friendly presentation of year round incentives to shift dryer, dish washer, and pool pump use into low cost hours is also a compelling program component.

2.7.4 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[72] 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 customers in zone 3 are more likely to avoid the need a peaker that would run a few dozen hours a year.[73]

2.7.5 The Implications of Having Peak Prices Extend into the Evening

The SPP's 2-7 PM peak period includes afternoon hours when many home are empty and most people are at work. It also includes evening hours during which offices empty as workers go home. Electricity prices are similarly high during both periods because these are the hottest hours in the day and because people are cooling both the empty and the full buildings. Good policy analysis should notice the differences in demand during these periods. Empty buildings have more elastic demand and thus suffer greater deadweight

[71] This kind of scenario led to a 2 hour, 1000MW rolling blackout when temperatures went over 100^o F in Texas on April 18, 2006 [Quinn, 2006]. Interestingly, the critical days with the two highest statewide, population weighted temperature took place on September 8 and 22, 2003.

[72] 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^oF for those hours.

[73] 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.

losses from mispriced electricity.[74]

Other studies (e.g. Momentum Market Intelligence) gather more direct evidence about customers opinions about early evening electricity use. The slice of the SPP data discussed here offers a modest amount of evidence that extending peak prices until 7PM led to enrollment resistance and attrition among people with electric stoves. It is also possible that it drove resistance and attrition among customers who tend to get home early in the evening. We would need to disaggregate afternoon from evening peak period SPP data to get a clear sense of this. Doing so would give us a sense of how much demand elasticity differs during these periods.

There is inconclusive evidence that some customers react to price increases by considering whether to shut off their air conditioning altogether without considering the possibility of making marginal changes to their electricity use. One customer justified the choice to switch from CPP to a time of use (TOU) rate with no critical events:, "I live in the desert, and it was hell not to turn on the air conditioner until 7:00 at night." Another customer who left CPP for TOU observed, "We roasted, it was horrible."[75] [Lineweber, 2005, 14] We need to find dynamic pricing programs that are more attractive to customers than the time-invariant alternative regardless of whether customers are using neo-classical or behavioral decision rules. Approaches with frequent high evening prices are likely to be less attractive to customers unless customers understand that the high evening prices allow for lower prices during other periods or enable other features that make the whole package more attractive.

Electric demand between 5:30 and 7PM is a real problem, but it may be possible to deploy a portfolio of strategies that incent residential customers to reduce their usage during the workday and perhaps incent workplaces to reduce use after the end of normal business hours.[76] A residential shoulder rate or heavily discounted overnight rate might convince residences to reschedule appliance use later in the evening without precluding evening air conditioning. That period might be a useful part of this portfolio as well.

[74] The first best solution is to mandate real time pricing at marginal cost for everyone and to impose fixed fees to recover any additional costs. The discussion here focuses on finding the greatest deadweight loss reduction in a politically constrained, opt-in context with customers who may be loss averse.

[75] Neoclassical economics could explain quotes. But the fact that these customers choose to "roast" rather than run an air conditioner for something on the order of \$3.50 per hour during a critical period may suggest that they are using a rule of thumb that does not consider intermediate options between normal operation and no air conditioning.

[76] The combination of the halves of the portfolio are necessary to reduce overall peak demand and provide reliability benefits and to reduce the need for peakers.

2.7.6 Understanding and Dealing with Cross Subsidies

An optimal opt-in dynamic pricing program has to offer compelling savings to the customers it most wants to recruit, convince to change their consumption patterns, and retain.[77]

A customer with an air conditioner may use far more expensive, afternoon power than a neighboring customer without an air conditioner, but with the same total power consumption. If the two customers pay the same price per unit of power that covers the average variable costs of their joint power consumption, then the customer without air conditioning will pay more than the average variable cost of the power they consume, implicitly subsidizing the expensive afternoon use of the peaky[78] customer. To illustrate in a very simplified context, we could imagine that the peaky customer uses 100 kWh of afternoon power that costs the utility 12 cents / kWh, while the other customer uses all of his 100 kWh of power off peak when it costs the utility of 8 cents / kWh. The utility charges a single, time invariant rate of 10 cents per kWh to recover its costs. Time invariant pricing creates cross subsidies from less peaky to more peaky customers who are on an identical rate. Differences in when customers use power across climates combine with regional differences within each California utility's rate structure to creates cross subsidies among regions.

The analysis in this section finds a perverse pattern that can make naively-designed dynamic rates unattractive to very responsive customers. Specifically, big customers in hot climates use more than the statewide average proportion of their power during weekday afternoons. They tend to be quite responsive to price signals. They also give up the substantial cross subsidies from customers with flatter load shapes when they switch to a "naive" dynamic rate. "Naive" dynamic rates are dynamic rates that fail to account for regional variations in load shapes. For example, a "naive" dynamic rate for a hot region might yield revenue identical to the time invariant rate for a customer 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[79] and more likely to participate. Thus, it is crucial to

[77] The whole policy implications 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).

[78] "Peaky" ("flat") customers use a larger (smaller) than average percentage of their total power consumption during weekday afternoon hours.

[79] One compelling response to problems that arise when time invariant rates create cross subsidies may

take these cross subsidies into consideration when designing a rate structure that allows responsive customers to save.

This section presents simple example rates that achieve revenue equivalence and then discusses their cross subsidy implications. Consider a set of customers who, on average, use a proportion $\overline{\alpha}_H$ of their total power use during the set of high cost, peak hours denoted H. Consider a simple example. Its naive time-differentiated rate would markup prices by $\mathcal{M}_H$ during the set of peak hours, H, and then lower rates by $\frac{\overline{\alpha}_H}{1-\overline{\alpha}_H}\mathcal{M}_H$ at all other times. These rates would be revenue neutral for the utility if customers did not change usage in response to the price change. [80] All customers whose peak power use was larger (smaller) than $\overline{\alpha}_H$ times their total use would experience a bill increase (decrease). These "structural losses" mean that peakier customers give up receiving significant cross subsidies to go on dynamic pricing and have to respond significantly before they begin saving money. Region-specific rates that address cross subsidies are important if naive rates

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 prioritiy 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.

[80] Converting this stylized example into a more robust rate that can meet utility revenue requirements under all plausible customer responses is straightforward. The firm can set $\mathcal{M}_i$ to equal the average variable cost power during period i and set a uniform per kWh charge, U, during all periods to recover its fixed costs. The analogous time invariant rate would be to use the same uniform charge U and to set the time-specific markup equal to the round-the-clock average variable cost for the average profile among its time-invariant customers. It will then be indifferent between selling peak and offpeak power and between dynamic and time-invariant rates.(Adams and Yellen [1976] show that pricing that makes the firm indifferent between selling two different products can be part of a profit maximizing bundling strategy.) If we set $U_t = R_t/Q_t$ where R_t is the firm's revenue requirement for this class of customers during time t and Q_t is the total number of kWh the firm sells, then periodically updating U_t can ensure that it meets its revenue requirement. We can do even better by using Ramsey pricing or a practical approximation of it to meet revenue requirements while minimizing the rate's consumption distortions.

If the utility uses this approach to create, on average, revenue-equivalent dynamic and time invariant rates, then customers who respond to dynamic pricing purely by shifting their consumption will come out ahead on dynamic pricing if the proportion of their power that they use on peak is less than the population proportion: $\alpha_{H,i} < \overline{\alpha}_H$. If the customer also changes the total quantity of power that they use, then they will come out with a lower bill if their total time invariant use times the time invariant price is greater than the sum of what they would spend at the peak and offpeak dynamic rates: $Q_{i,d,H}P_H + Q_{i,d,L}P_L < (Q_{i,inv,H} + Q_{i,inv,L})P_{inv}$

The profit issue can be secondary if utility regulators commit to ensure that the utility earns its rate of return despite changes in quantity and timing. Any welfare improving change in pricing creates a potential Pareto improvement and careful rate design can capture some of its benefits to increase firm profits. In my conversations with utilities, however, concerns that the regulators will not adjust rates to restore profits has been a central reason that they have either feared dynamic pricing or chosen to recover fixed costs during off peak hours.

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 2.11: 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.

impose disproportionate structural losses on responsive customers.

Tables 2.11 and 2.12 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.[81] The control group use patterns offers insight onto the distribution of structural

[81]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 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

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 2.12: 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.[82],[83]

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.[84] Doing so would mean that between 60 and 75% and between 55% and

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

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:

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

[82]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.

[83]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.

[84]Baseline-rebate rates can be thought of as setting a customer specific $\alpha_{i,H}$. Using each customer's

Average Percentage of each customer's summer-season power used on peak for CPP customers who stayed and left			
climate zone	chose to stay	exited	stayed by default
2	18.9	19.0	16.7
3	21.3	25.3	23.9
4	17.8	23.8	24.1

Table 2.13: This is a sample of 157 customers for whom we have continuation data or who would have been in that sample had they not exited early. Results are weighted to make them representative of the state population as a whole.

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 reduced weekday afternoon use from extremely peaky starting points.

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.

Evidence from Attrition

Customer choices about whether to continue on dynamic pricing[85] at the end of the experiment confirm the importance of designing rates that allow responsive customers to save money. At the end of the SPP, customers got letters asking them to choose whether to continue on dynamic pricing or return to time invariant rates.[86] These choices offer direct evidence about whether customers thought that dynamic pricing outperformed time invariant prices. They are a meaningful measure of whether CPP worked well for a customer. These customers faced incentives that are quite close to the incentives that customers in dynamic pricing programs would face if they were prompted to choose whether to renew their participation. Table 2.7.6 shows data on a small, unrepresentative sample of about 200 end-of-experiment choices. The results discussed here are not as definitive as those elsewhere in the discussion.

- In climate zones 3 and 4, the customers who opted out by mail or phone or left the

behavior to set a customer-specific baseline level has notable drawbacks discussed at length in Chapter 3.

[85] The SPP did qualitative interview research about why customers stayed on or left dynamic pricing at the end of the experiment. Schultz and Lineweber [2006] and Lineweber [2005] report the results of this work.

[86] After the conclusion of the experiment, dynamic pricing customers 1) had to pay a new, modest, monthly meter charge and 2) got no more payments for their participation.

experiment early were, on average, peakier than the statewide average.[87]

- In all climate zones for which continuation data are available, the customers who returned a form asking to remain on dynamic pricing used, on average[88], a smaller proportion of their power on peak than the statewide average, meaning that they saved money under this set of dynamic rates.

- Customers who participated in the whole experiment and did not return the form remained on dynamic pricing by default. They were slightly peakier on average than the control group average.[89]

This discussion neglects many of the complicated facets of California's rate structure and the SPP's intentional summer-to-winter bill shifting, but its main ideas hold in the more complicated reality.

The SPP created deployed rates that were ill suited for the most responsive customers because it required that "the experimental rates . . . be revenue neutral for the class-average customer over a calendar year, in the absence of any change in the customers load shape" [Charles River Associates, c, 18]. 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.[90]

See Borenstein [2006] 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 importance of rate design for making commercial solar photovoltaic installations attractive.

[87]Some of the customers who left the experiment early moved. Others were unhappy with the rate.

[88]I get this result after dropping one extremely small (150kWh/ month), customer in zone 2 who used more than half of his power on peak. This customer is so small that a temporary spike in use during peak hours, like a construction project, could switch him from appearing extremely flat to extremely peaky.

[89]There is significant evidence in the behavioral economics literature [Choi et al., 2003] that people tend to accept the default choice even if it is not the best choice for them.

[90]If cool climate customers are barely responding to dynamic pricing, then we can offer them the structural savings in the form of bill credits without needing to enroll them in dynamic pricing. This would let us direct marketing, education, and (possibly) metering spending toward highly responsive customers in hotter climates.

2.7.7 Setting Prices

The current, imprecisely estimated results suggest that customers' response to high and low ratio CPP rates seem to be quite similar.[91] If these results hold up, it would suggest that a variety of rates that offer customers incentives to reduce demand during high cost periods and stronger incentives during scarcity periods can significantly outperform time invariant rates. This gives policy makers flexibility to choose rates to meet concerns beyond economic efficiency. However, before we take this to mean that all reasonable CPP rates are indistinguishable, it is important to note that the SPP data show that modest rate differences drive few differences in the short term, but the current analysis provides little evidence about longer term participation and investment issues:

- Customers will learn over years of experience whether responding to price signals generates savings that justifies staying on the rate and changing electricity use on a day-to-day basis.
- The relationship among the critical, peak and off peak prices and the price variation among seasons affect which customers save under the program and stay on it. Rates that cause bill shocks or change seasonal bill patterns may cause customers who analyze one bill at a time to wrongly believe that they are losing under the new rate and to exit.[92],[93]
- In the long term, rates determine whether it makes economic sense for firms to offer and customers to adopt response technology.

Finding that different choices of peak prices elicit fairly similar customer responses has significant implications for forecasting the impact of candidate rates. It may suggest

[91] One specification finds that customers who faced a higher peak and critical prices, counterintuitively, used more critical-period power than did customers who faced lower critical prices during critical events at the p=.10 level. Other critical period results were substantively similar, but statistically insignificant. The TOU peak results were imprecisely estimated, but substantively small and statistically indistinguishable from zero. All of the current results allow us to reject the hypothesis that moving customers from the low-ratio to high ratio rate would prompt them to reduce consumption at least another .1 kW at the p=.05 level. So we can state that immediate, substantively striking impacts of modest adjustments to peak and critical prices are unlikely.

[92] Bold statements on bills can address this problem. They might say: "You are on track to save $X this year. This rate tends to lower winter bills and raise summer bills."

[93] The SPP's high ratio rate tended to lower summer bills and raise winter bills. This creates evidence about bill shifts' impact on attrition. We would look for differences between high and low ratio rates in customer's propensity to exit that are unrelated to what the customers would save. The data to do this analysis exist, but UCEI is still in the process of obtaining them.

that, in the absence of more flexible demand estimation models fitted to larger datasets, predictions based on point estimates of the price-independent TOU peak and critical price impacts may be more accurate than plugging the prices and a point estimate of a CES demand elasticity into a demand model.

2.7.8 Selling Responsible Energy Citizenship

The survey results discussed in section 2.2.4 suggest that the CPP group, and especially the high use subset of it, has a stronger belief in sacrifice for the common good than did the average control customer. The survey reports that treatment customers were more likely to agree that "everyone should pay a little" to protect the environment but were no more likely to believe that the environment was at risk. This suggests that some customers are receptive to signing up for dynamic pricing because of its social benefits. Thus marketers should test materials that describe benefits like increased grid reliability, reductions in the operations of old, dirty peakers[94], and the reduced need to build peaking plants in the customer's community.

2.7.9 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.[95]

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 2.2 on page 18 and 2.3 on page 19 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 in the pilot

[94] California's peakers are dirtier than its base load capacity. This is not true in other regions with coal-fired base load capacity [Holland and Mansur, 2005].

[95] 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.

so that this difference was 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.[96]

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

There is considerable evidence that most customers who opted in responded and came out ahead under the terms of the experimental rate and participation payments. [97] Many customers who paid more under dynamic pricing opted out during the experiment when some customers left, and at the end of the experiment when about 60% of customers chose to continue on CPP and the remaining customers returned to more conventional tarriffs. SPP customer choices about entering and leaving dynamic rates deserve a careful study in their own right.

[96] The results in 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.

[97] A striking deviation from rationality is that many structural winners from the temperate climate zones refused invitations to participate. And the opt in rates from this experiment were far higher than those that similar, fully deployed programs have gotten without big incentives to participate. See Chapter 3for an extended discussion of this.

2.8 Conclusions

Air-conditioning-driven electricity demand is a large part of the 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 TOU peak (critical) prices on the hottest days in the two hot climate zones. 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. 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.

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 that offering the most responsive customers enough savings to convince them to participate requires careful attention to cross subsidies and differences in regional usage patterns. The SPP's experience suggests that a wide variety of rates can convince customers to "conserve" every weekday afternoon and to "conserve a lot" during critical periods.

Chapter 3

Applying Psychology to Economic Policy Design: Using Incentive Preserving Rebates to Increase Acceptance of Critical Peak Electricity Pricing

3.1 Introduction: The economics and psychology of dynamic electricity pricing

This project extends the idea that policies should address problems by improving economic incentives. Insights about how people make decisions suggest that careful presentation can help consumers understand incentives and make the individually rational responses that economists expect them to. This project applies insights about economics and psychology to understand and address costly consumer resistance to improved residential electricity pricing. This project proposes Incentive Preserving (IP) Rebates to sidestep heuristics that can cause mistaken resistance to critical peak pricing (CPP) of electricity.

Most customers are on time-invariant pricing that charges the same price per unit of power during high and low cost hours. "Real time" electricity pricing sets hourly prices that reflect the marginal cost for that hour. Signing up every residential customer for real time pricing could deliver an estimated $6-12 billion in annual social benefits.[1]

CPP is a simplification of real time pricing that announces a schedule containing a handful of peak and offpeak periods and sets a price for each period. CPP makes prices reflect some of the enormous variations in marginal social costs between periods when power is scarce and when it is plentiful. CPP allows the utility to address scarcity by designating roughly 1% of all hours as critical periods which invoke a significantly more expensive critical rate. Policy makers consider CPP an attractive rate for residential and small commercial customers. This project takes CPP rates' prices and schedules as given.

CPP works. Residential customers who switch to CPP reduce their usage during higher-priced, peak periods and the highest-priced, critical periods. CPP customers report high satisfaction levels. Indeed, the majority of customers who received $175 to participate in a California CPP experiment chose to stay on CPP at the conclusion of the experiment.(Faruqui and George, 2005; Charles River Associates, c)

Consumers, however, resist signing up for CPP. Mailings offering Florida customers a CPP rate that saves participants an average of $90 a year get a 1.3% opt-in rate (White, 2006). Customers are more receptive to baseline-rebate programs that create similar incentives by offering rebates to customers who use less than a baseline amount of power during critical periods, but calculating baselines from the customer's recent usage creates perverse incentives for customers to use more power during baseline-setting periods in order to increase their eligibility for rebates. Customers who resist CPP but are open to rebate programs with the same average bill appear to have preferences about elements of the presentation that affect neither incentives nor total bills.[2]

Residential CPP will generally be an opt-in program until it develops a track record that justifies making it the default. Current sign-up rates limit CPP's ability to generate a compelling track record. It is difficult to generalize from the unusual customers who opt

[1] American residences spent $116 billion on electricity in 2004. (http://www.eia.doe.gov/cneaf/electricity/epm/table5_2.html) Borenstein (2005a) reports that real time pricing could yield 5-10% annual savings on energy over the long term. There are – to the best of my knowledge – no academic papers that estimate the potential benefits of CPP and compare those to real time pricing.

[2] It is rational for consumers to prefer time-invariant pricing to CPP if CPP's transaction costs or higher prices during peak periods reduce the consumers' overall utility. By the same token, a significant number of consumers who use a smaller-than-average proportion of their electricity during weekday afternoons could save money under CPP without changing their consumption patterns. Yet, the majority of these consumers do not sign up to claim these savings.

in to existing CPP programs to the customer base as a whole. Pilot programs attempt to recruit more representative customer pools, but they provide limited evidence because they only have tens or hundreds of customers. Policy makers are likely to want field experience with broad-enrollment programs before they consider making dynamic pricing the default rate. Presenting the rate in a way that helps customers to make better enrollment choices will both improve the performance of opt-in programs and may be a stepping stone toward making dynamic pricing the default rate. Hence, this project seeks to present CPP in a way that elicits good enrollment choices when shown side by side with the status-quo time invariant rate.

CPP presents good incentives, but does so in a way that biases several heuristics toward choosing not to enroll. CPP delivers subtle savings by lowering prices most of the time. Its critical events inflict visible losses by notifying customers that it is raising prices so much that customers spend more on power or get less power during the critical period than they would had it been an ordinary period. California customers reduced usage 12% when the CPP pilot experiment events more than doubled prices.(Faruqui and George, 2005; Charles River Associates, c; Herter et al., 2007) Thus, during the 1% of all hours that were events, many customers paid more for power despite conserving. CPP reduced most participating customers' total annual bills because the savings from small price reductions nights, mornings, and weekends more than offset any bill increases during the rare, but visible critical events. People's heuristic decision making procedures are likely to notice and overweight the high priced periods and either not notice or underweight the gains. Concentrating losses in a few high cost months and diffusing gains over the calendar repels loss-averse customers if they "narrowly bracket", meaning that they consider bills one cycle – or even one day – at a time, rather than over the long term (Thaler, 1999; Read et al., 1999). CPP also repels customers who believe it is unfair to charge very high prices for air conditioning when they need it the most (Kahneman et al., 1986). Incentive Preserving Rebates change the presentation of CPP to avoid these biases.[3],[4]

IP rebates present critical events as opportunities to earn rebates through sacrifice.

[3] IP rebates can work with a wide family of dynamic pricing programs. This project presents them in the context of CPP because CPP is a simple, illustrative, and policy relevant application. Section 3.6 describes the generalization.

[4] Adding IP rebates to CPP may slightly change the amount of power that customers buy because IP rebates charge a few extra dollars during some months and return those dollars as bill reductions in other months which creates income effects.

IP rebates change neither the total annual bill nor marginal incentives.[5] IP rebates add a fixed amount to each month's CPP bill. This monthly payment buys the customer rights to buy a fixed quantity of power at the usual price during critical events. If customers use less power than they had rights to during an event, they get a rebate for the value of the unused rights.

We can see how this works in practice by considering a rate that sets the opportunity cost of a kilowatt hour (kWh)[6] during a critical event at 60 cents and charges 24 cents per kWh during normal peak periods. A customer who has the right to buy one kWh during a critical event for the normal price of 24 cents can use the right for either 36 cents worth of power or for a 36 cent rebate. We can offer a customer the right to 8 kWh at the usual price during each of 15 events if we charge the customer $3.60 (which buys 10 kWh of rights) each month. A customer who exhausts their rights during an event has to pay the full price of 60 cents per kWh.

The IP rebate-rate design collects monthly fees to purchase rights. It does so through declining block pricing that adds a markup to the first, fixed number of kWh that a customer purchases each month. Section 3.4.2 describes this approach and literature reporting that customers make better choices under this presentation.

IP rebates maintain the right marginal incentives while using fixed transfers of cash or property rights to adjust the size of monthly bills and to use forgone credits rather than price increases to raise opportunity costs during events. Variations on this fixed transfer strategy underlie the Coase Theorem, hedging to manage risks in financial markets, and policies like cap-and-trade pollution permit systems.

Psychological and economic factors are important throughout the life cycle of a dynamic pricing program. This project suggests a way to present good economic incentives that is compatible with participants' decision-making heuristics. This project focuses on psychology at the opt-in stage and on economics of creating the right incentives at the participation stage. This prioritization reflects several insights:

- Decision making heuristics appear to cause mistakes at the opt-in stage that hurt

[5]This project, like much of the policy-oriented behavioral economics literature, assumes that consumer errors and biases are a common but not universal problem and thus seeks interventions that improve biased consumers' choices without affecting rational consumers' choices. (c.f. Camerer et al. (2003) "Regulation for Conservatives: Behavioral Economics and the Case for 'Asymmetric Paternalism'" and Sunstein and Thaler (2003) "Libertarian Paternalism is Not an Oxymoron")

[6]A kWh is enough energy to run a 100 watt light bulb for 10 hours, or a central air conditioner for about 15 minutes.

consumers and society. Customers who have experienced dynamic pricing appear to respond to its prices and consumption incentives in roughly economically rational ways. Customers who experience dynamic pricing reduce use during high-priced periods, save money, and report high satisfaction levels, which is consistent with them being economically rational.

- If there are bad incentives for participating customers, enough customers find and take advantage of them to cause problems. For example, some Anaheim customers reacted to incentives by using extra power on ordinary summer afternoons to become eligible for larger rebates (Wolak, 2006).
- Many customers suffer "projection bias" which causes them to overestimate how difficult a new situation will be to get used to (Loewenstein et al., 2003). This makes loss-averse heuristics more important when consumers decide whether to opt-in than after they have some experience with dynamic pricing.

Increasing the credit size has no effect on the utility's annual revenue or the consumers total annual payments, but ensuring that customers get rebates rather than pay high prices during events makes the offer more attractive to loss-averse customers. This project aspires to design a rate that offers most customers a rebate during any month with an event. This goal comes from considering psychological factors that are likely to affect the satisfaction of customers who have already signed up for the program.

3.1.1 Implementation

There is good reason to think that the IP rebates are feasible. IP rebates add revenue neutral charges and credits to an underlying rate that regulators can tailor to meet local needs.

This project proposes a rate that asks customers to pay for their rights through a small markup on a fixed number of units of power per month. This creates a trade off between the customer's likelihood of getting rebates and their ability to pay for their rights. For example, a rate might work well for a customer if they used fewer than 20 kilowatt hours during each five hour event and used at least 300 kilowatt hours per month.

An IP rebate policy has to decide whether to address the challenge of making an appropriate offer to each customer by offering each customer a customized offer or by

Table 3.1: Price and demand are very high during one percent of all hours. California ISO Electricity Market: October 2005 through September 2006; Spot market prices are for the NP 15 Northern California Region

	Median	99%	Max
Usage, Megawatts	27,064	43,779	50,198
Wholesale Price, $/kWh	$0.045	$0.163	$0.396

splitting customers into categories and making one offer to each category. Making offers to broad categories of customers defined by use and geography can perform adequately and has compelling advantages over individualized offers. Specifically, assigning rebate eligibility by category seems fair since neighbors who live in superficially similar houses will generally get the same offer. Offers to categories of customers facilitate analysis and discussion among utilities, regulators, and advocates who may be able to use categories already in use for other regulatory purposes. Categorical offers reduce the likelihood that customers will demand extra power in a (misinformed and fruitless) attempt to profit by becoming eligible for more rebates. An analysis of California data shows that we can make adequate offers to the vast majority of customers even if we crudely split up customers using readily observable characteristics like climate zones and monthly summer electricity usage.

This paper proceeds in two stages. The first stage describes the economic case for improved electricity pricing, the behavioral challenges to implementing it, and proposes Incentive Preserving (IP) Rebates. The second stage explores whether IP rebate deployment is feasible by simulating IP rebates' impacts on a set of California CPP customers.

3.2 Background: Improved electricity pricing can deliver significant savings

Providing better incentives for customers to shift power use away from periods of electricity scarcity has the potential to save billions of dollars, to deliver significant environmental benefits in some markets (Holland and Mansur, 2005, 2006), and to facilitate the integration of wind generation into electricity systems.

Electricity storage is generally not cost effective, but electricity supply has to meet demand minute by minute to prevent blackouts. This creates enormous variations from hour to hour in the cost of generating electricity.

Most consumers are on time invariant rates, which do not depend on when the customer uses the power. Time invariant rates offer customers no incentive to shift use away from high cost periods. The combination of time-invariant rates and the need to meet demand minute by minute creates extremely inelastic demand and can give suppliers a great deal of market power. These factors require electric system operators to maintain extra generating capacity that they only use when extreme weather or equipment problems tax the system a few hours a year. (Joskow (2000) and Borenstein (2005a) discuss this background in detail.) Table 3.1 shows that the 1% of all hours with the highest demand are very costly. California's electricity demand was more than 6,400 megawatts higher in its maximum hour than it was in the 99th percentile hour. The Energy Information Administration estimates that building a megawatt of generating capacity costs roughly $400,000 and keeping that facility maintained and ready to operate costs $11,000 a year (Conti et al., 2006). The maximum spot market price was more than twice the 99th percentile price and roughly nine times the median price.[7]

Dynamic electricity rates are a family of rates that vary prices over time to better reflect hour-to-hour differences in the marginal cost of power. Dynamic residential electricity pricing has the potential to save up to 5-10% of the $116 billion that American residences spent on electricity in 2004 – or $6-12 billion a year. Dynamic pricing was generally not deployed in the past because it was not cost effective. It requires meters that record when customers use power as well as their cumulative total use, but new computer technology makes these meters much cheaper. There is compelling evidence that dynamic pricing can save billions of dollars (Borenstein, 2005a). Residential customers consume 36% of the electricity used in the US[8] and respond well to the incentives in dynamic electricity prices (Faruqui and George, 2005; Charles River Associates, c; Herter et al., 2007).

Critical peak pricing (CPP) is the dynamic rate that gets the strongest consideration for residential and small commercial customers. CPP announces a schedule of peak and offpeak periods and prices, and allows the utility to address scarcity by designating roughly 1% of all hours as critical periods that invoke an expensive, critical rate. Table 3.2

[7]The wholesale market price of energy during extremely high demand periods is a lower bound on the true marginal cost of producing the energy. Electricity spot markets include features – like price caps – that control prices during scarcity periods and cover the true costs of maintaining capacity to meet demand during those hours through other payments. Borenstein (2005a) discusses how current electricity prices can understate the cost of scarcity and simulates more accurate prices. Wholesale energy prices are lower than retail prices because wholesale prices omit costs of infrastructure like the electrical distribution network, customer service, and utility sunk costs from things like nuclear power plant construction and the California electricity crisis.

[8]source: Electric Power Monthly 2004 figures: http://www.eia.doe.gov/cneaf/electricity/epm/table5_1.html

on page 91 presents an example of CPP.

The technology and economics are largely ready to support widespread dynamic pricing. Two recent National Town Meetings on dynamic pricing brought together more than a hundred regulators, utility staff members, academics, and suppliers. Few utilities have successfully implemented dynamic pricing – largely because they have struggled to present dynamic pricing in ways that consumers find attractive (Barbose et al., 2004). Many have not done a good job of risk management, marketing, and implementation. Most implementations have struggled to get customers.

For example, field experience in Florida and Illinois shows that residential resistance is a serious problem.[9] Gulf Power offers GoodCents Select residential CPP in Florida. Gulf Power and its parent, the Southern Company, are considered leaders among utilities in marketing and customer service. The Community Energy Cooperative's Energy-Smart Pricing Plan offers Illinois residences "real time prices" based on the hourly market rate. Both notify customers of high priced periods. The two retailers report that most customers who sign up save significant amounts of money, are satisfied, and stay enrolled. But both programs get sign up rates of only about 1%.

Some consumer resistance is rational. Customers who use a large proportion of their power during peak periods would pay more under CPP. The transaction costs of responding to price signals could deter some customers. Practitioners and scholars are working to address both issues (see e.g. Borenstein (2006); Wright et al.). Gulf Power's Good Cents Select program creates winners by saving participants an average of $90 a year (White, 2006), and reduces transaction costs by providing a computerized "set it and forget it" thermostat that automatically shifts air conditioning away from critical periods (Gulf Power). However, conventional economic reasons cannot explain the high rejection rate among customers who use a larger-than-average proportion of their power off peak and would save money on CPP even if they did not respond to price signals. Customer retention and customer recruiting are related, but fundamentally different challenges. Retention involves customers who have experienced CPP. They know the program's implications for their total bills, lifestyles, decision making heuristics, and preferences. Recruiting involves decisions by customers who know significantly less. Further, retaining customers requires

[9]Large commercial and industrial customer resistance is also a problem. Commercial and industrial customer are beyond the scope of this project because large consumers should hire analysts and otherwise analyze economic decisions in different ways than small customers do. Large customers may also suffer principal agent problems.

1) that responsive customers save money under the new program which is a function of the CPP rate; 2) that the program be explained clearly and that customers know when they are saving money; 3) that responding not be too onerous – which requires thinking carefully about whether to include evening hours in events and whether to call events on consecutive days; and 4) – potentially – that the program use something like IP rebates to present critical events in a way that meshes with the way people make decisions. Gulf Power, the Community Energy Cooperative, and California's Statewide Pricing Pilot have high customer retention rates. Field experience suggests, however, that recruiting is a harder, unsolved problem. This project concentrates on designing CPP rates that facilitate recruiting customers.

This project tries to address the recruiting challenge by designing policies that overcome psychological resistance to efficient policies among consumers who stand to gain from signing up for them. It parallels a literature that reports that economic efficiency does not sell itself in the political marketplace. That literature reports the success of clever designs and compromises that protect features that achieve efficiency while letting political concerns drive other design choices (Robyn, 1987; Hausker, 1992).

Residential CPP will generally be an opt-in program until it develops a track record that justifies making it the default.[10] The benefits of signing an additional customer up for dynamic pricing are greatest when there are few customers on the program and diminish as dynamic pricing's market share expands (Borenstein and Holland, 2005). Getting a significant fraction of customers signed up could deliver compelling benefits even if it is well under 100%. Hence, this project seeks to present CPP in a way that elicits good choices when shown side by side with the status-quo time invariant rate.

3.3 A variety of psychological theories suggest that CPP's presentation of incentives will repel customers.

CPP delivers subtle benefits by modestly lowering prices most of the time while it occasionally inflicts visible losses during critical events when the average customer uses

[10]CPP's initial opt-in status is a political reality, but making it opt-out or mandatory might be better policy. The strand of literature that suggests that IP rebates might work also reports that changing defaults or forcing people to decide can lead to considerably better choices in retirement savings. (See Choi et al. (2003) and the literature it cites.) Wood (2002a,b) argues that changing the default may be sufficient to get large scale participation in pure-pricing, time differentiated electricity rates.

less power, but pays more in total for it. A variety of psychological theories suggest that presenting subtle gains and visible losses is flawed and will repel consumers.

- "Narrow bracketing" consumers base their decision on short term outcomes like a billing period or an afternoon rather than considering the appropriate long term outcomes (Thaler and et. al., 1997; Read et al., 1999). Narrow bracketing underlies most of these psychological theories because many customers will come out ahead on CPP in the long run. (The intervention proposed here specifically addresses narrow bracketing by delivering significant monetary benefits during the critical periods that ask people to make the most salient consumption sacrifices).
- A reference dependent loss averse customer codes outcomes as gains or losses relative to an anticipated (reference) outcome. These consumers consider that losses relative to the reference point loom larger than gains. A critical event that leads the average customer to pay more to buy less than they would on their reference, non-critical day would seem quite painful (Kahneman and Tversky, 1979).
- Studies of choice under risk suggest that consumers not only exhibit something akin to loss aversion, but also often consider just the worst case rather than the whole outcome distribution. (March and Shapira, 1987; Lopes, 1987)

Field evidence is consistent with these factors playing a role: "A number of program managers suggested that the modest participation rates in their RTP [real time pricing] program were a result of the fact that .... the vast majority of eligible customers view the risks of RTP as too great and/or the potential benefits as too small." (Barbose et al., 2004)

Customers sometimes appear to choose the option with the greater number of attributes that compare favorably (Redden and Hoch, 2005). A typical CPP rate defines three-periods: offpeak, peak, and critical periods. Two of these periods – peak and critical – are more expensive than the time invariant price but account for less than 20% of all hours. Hence, three period CPP compares unfavorably to time invariant prices to customers using this heuristic. Gulf Power's decision to use four CPP periods that sets two off-peak prices – "medium" and "low" – lower than the time invariant price – might reflect this customer decision-making heuristic.

Many consumers find it unfair to raise prices to deal with a shortage stemming from a shock that has increased their demand for a product. CPP often invokes critical periods that raises prices when heat waves maximize air conditioning demand. This is nearly the exact summer analog to Kahneman et al. (1986)'s finding that most consumers found it unfair to raise the price of snow shovels during a blizzard. Many customers consider conventional, efficient pricing an unfair way to deal with shortages (Kahneman et al., 1986).[11] Gulf Power (Gulf Power) assures customers that its Critical Peak Pricing price levels "[R]eflect the actual cost of producing electricity during those periods."

The marketing literature reports that customers prefer declining block pricing to fixed fees that generate identical revenues and incentives for most consumers (Ho and Zhang, 2004). This is not a challenge to CPP but constrains the ways we can we modify it.

A variety of decision making heuristics in the psychology literature can explain customers rejection of a CPP rate that would save them money in the long term, but creates the possibility of larger, more salient short term losses on critical days or during months when they already use the most power; that raises prices in two out of three kinds of pricing periods and that raises prices during periods when customers need air conditioning the most. Indeed, narrow bracketing and any one of the other heuristics could explain the resistance.

The evidence about these heuristics comes from studies in other, often simpler, contexts so it is not immediately clear how those studies play out in the CPP context. A paper in development explores how people think about dynamic pricing to clarify whether and how these heuristics applies to electricity pricing choices. The study seeks to understand the thought process that drives resistance to CPP and how IP rebates change customer thinking.

3.3.1 Behavioral interventions can significantly reduce irrational choices without affecting the decisions of rational players

Scholars have proposed interventions which, like IP rebates, change the timing of decisions, costs, and benefits; presentation; and information flows. Studies show that these interventions improve consumer choices in areas like retirement savings and investment choices. For example, the "Save More Tomorrow" program increased employee savings at companies by asking them to precommit to increase their retirement savings rate the next time they got a raise. This timing sidesteps loss aversion. (Thaler and Benartzi,

[11] Rabin (1993) and Charness and Rabin (2002) formally model preferences for fairness.

2004) A series of lab experiments (Gneezy et al., 2003; Thaler and et. al., 1997; Gneezy and Potters, 1997) show that customers invest more in riskier, but higher expected return instruments when experimenters send them aggregated information which forces them to broadly bracket. This reduces the likelihood that subjects will learn of temporary losses, which reduces resistance from reference dependent loss aversion. Most of the increase in risk taking happened as soon as the subjects learned that they will be getting aggregated feedback but before they experienced the aggregated reports (Gneezy and Potters, 1997) which offers hope that merely promising to reduce the experience of losses can recruit customers.

Marketers frame costs as gains by presenting sales below a “regular” reference price or by offering rebates. They use deferred payments to deliver benefits now and costs later. Public policy designers who consider decision making heuristics in designing policies should aspire to use interventions that correct, rather than cause self-destructive mistakes. The policy oriented behavioral economics literature is yet to consider whether changing the presentation of prices in economically neutral ways could help people make better decisions about whether to sign up for CPP, whether to invest in energy efficient products, or whether to own a car or rely on pay-per-use transportation. The policy oriented literature uses many of the same insights, but aspires to help people avoid mistakes rather than just changing decisions to maximize profits.

Behavioral economics studies and commentaries on their policy implications (Sunstein and Thaler, 2003; Camerer et al., 2003) seek interventions that improve biased consumers’ choices without affecting rational consumers’ choices because consumer errors and biases are a common but not universal problem. This project shares that outlook.

Interventions like changing the presentation of choices and information flows and the combination of CPP and IP rebates (CPP-IPR) proposed here reflect the emerging idea that we should address problems by not only using good incentives – like CPP – but also by presenting incentives in ways that reflect how consumers decide.

3.3.2 Similar interventions to improve decision making by addressing biases are possible for electricity pricing

Assigning customers property rights that allow them to earn rebates and come out ahead during critical peak events can sidestep resistance from all of these heuristics:

- Rebate opportunities can transform the presentation of critical events into opportunities to earn rebates through sacrifice rather than obvious losses when customers pay more to get less.

- Moving monetary gains into the same evaluation period as consumption losses largely sidesteps narrow bracketing.

- Customers who count the number of favorable attributes will see outcomes during critical events that look better than the status quo time invariant rate instead of looking worse. That means that the majority of CPP-IPR periods look better than the status quo rate.

- A rebate program seems fair because customers have the right to use power during shortages and the program pays willing customers to conserve, rather than making customers face the likelihood of paying more to get less during each event.

- The property rights that creates rebate opportunities will typically improve outcomes on both the costliest days and costliest months[12] which will make the offer more attractive to customers focused on the worst case. As 3.10 describes, CPP-IPR is particularly effective in reducing the worst monthly outcomes in climates with the hottest where customers use the most power and concentrate that use during the season of scarcity.

Rebate programs seem popular with consumers. Members of a PEPCO focus group of residential customers from Washington DC considered conventional real time pricing and a similar rate presented with rebates. When asked which of the rate features presented that night they liked, most customers mentioned rebates. (King and Harper-Slaboszewicz, 2006)

We need to worry that using IP rebates to reframe critical events as opportunities rather than threats will make customers more willing to sign up, but less responsive to critical events. A series of experiments have documented an "endowment effect" that makes customers demand far more to give up mugs that they have just been handed than they are willing to pay for mugs that they do not have, so there is reason to be concerned that telling customers that they own the right to low priced power will make them less likely

[12] Critical events happen disproportionately when the average customer is using the most power and paying the most for it. CPP-IPR dampens seasonal variations for these customers. Some customers' usage does not follow these patterns.

to give it up. Anaheim customers shifted a significant amount of consumption away from critical periods to earn rebates in a pilot test of a a critical peak rebate program (Wolak, 2006). A CPP-IPR program that created an endowment effect might still deliver more benefits than a CPP program with no endowment effect, if the CPP-IPR program got a higher participation rate than the CPP program.

The remaining challenge is to develop an intervention along these lines that preserves CPP's incentives and is feasible to implement.

3.4 Improving the presentation while preserving incentives

Psychology suggests a need to change the presentation of dynamic pricing to make critical events into opportunities to gain and to do so without imposing fixed fees. This section seeks ways to deliver that presentation in a way that preserves marginal incentives and charges each customer a total annual bill identical to what would have been charged under CPP.

CPP rates are attractive because they are a simple rate that creates a uniform and reasonably good marginal incentive during each time period. Their incentives are independent of the customer's power usage in this time period or a previous time period. Preserving marginal incentives is important because Wolak (2006) reports that many customers exploited the incentives in Anaheim's rebate program. A senior regulator characterized these customers as "mini-Enrons" when she discussed his findings at the University of California Energy Institute POWER Conference in March 2006.

3.4.1 Using fixed credits and fixed charges to offer rebates while preserving CPP's marginal incentives and annual total bills

Incentive preserving rebates transform the presentation of CPP while preserving CPP's revenue streams and generally preserving its marginal incentives.

IP rebates make critical events into opportunities for customers to gain by selling each customer rights to a block of power at the regular price during critical events and offering rebates for the value of any unused credits. The rebate value gives customers the right incentives to choose between using their rights and cashing them in. Each customer pays a fixed monthly fee to buy these rights. For example, a customer might pay $5 a month to buy the rights to $4 worth of power during each of 15 events a year. So the customer pays a total of $12 * \$5 = \$60$ per year to get credits worth $60.[13]

The essential insight here is that we can maintain the right marginal incentives while adjusting daily and monthly bills through fixed transfers of cash or property rights. Since CPP defines the number of events per year and the prices during the critical events the value of rights to use power during each event is clear and there is no reason to charge a risk premium. Hence, there is a well defined price of the property rights to power, so transfers of cash are identical to transfers of property rights. Versions of this insight underlie the Coase Theorem, hedging in financial markets, and policies like cap-and-trade pollution permit systems.

We can make these transfers of rights revenue neutral for each customer – every penny they put in they get back either as power or as a rebate.[14] This revenue neutrality means that it is impossible for a customer to profit by strategically manipulating the number of rights that the utility assigns to them. Hence, the only new economic incentive that the program creates is the desired incentive to shift use away from peak and critical periods. Revenue neutrality allows designers to offer quantities of rights that depart from a customers' likely usage. They can use this freedom to ensure that most customers get rebates and to make the same offer to a broad class of customers. Cross subsidies, where some types of electricity customers pay more than their share of total system costs, while other types of customers pay less than their share, are a ubiquitous flaw of electricity rates.[15] Revenue neutrality means that there are no cross subsidies in the rebate program, making CPP-IPR as transparent and as equitable as the underlying CPP rate. Further, it means that the rebate program should deliver revenue that is identical to CPP revenue – so none of the design parameters that the rebate program dictates lessens the utility's revenue stream or makes revenues harder to predict.[16]

[13]The rate in Table 3.2 has customers contribute $2.5\ cents * 450kWh = \$11.25$ per month, which is $135 per year, and offers them $\$9 * 15\ events = \$135\ per\ year$

[14]Achieving true revenue neutrality requires paying interest on contributions to equate the net present value of the dollars each customer pays to the net present value of the rights that they get later in that calendar year. I omit this straightforward but small and tedious adjustment for brevity.

[15]Going to real time pricing that sets a price of power each hour is a necessary, but not sufficient condition to eliminate cross subsidies. Critical peak pricing reduces cross subsidies.

[16]CPP and CPP-IPR raise identical revenue assuming that customers behave identically under them. Section 3.3.2 discusses endowment effects that might cause CPP-IPR customers to use more critical period power. This might cause the kind of shift in revenue that would reflect the relationship between the CPP price's critical price and the cost of providing power during those times. If utilities charged a critical price equal to the marginal cost of critical-period power, then this change would also be (net) revenue neutral.

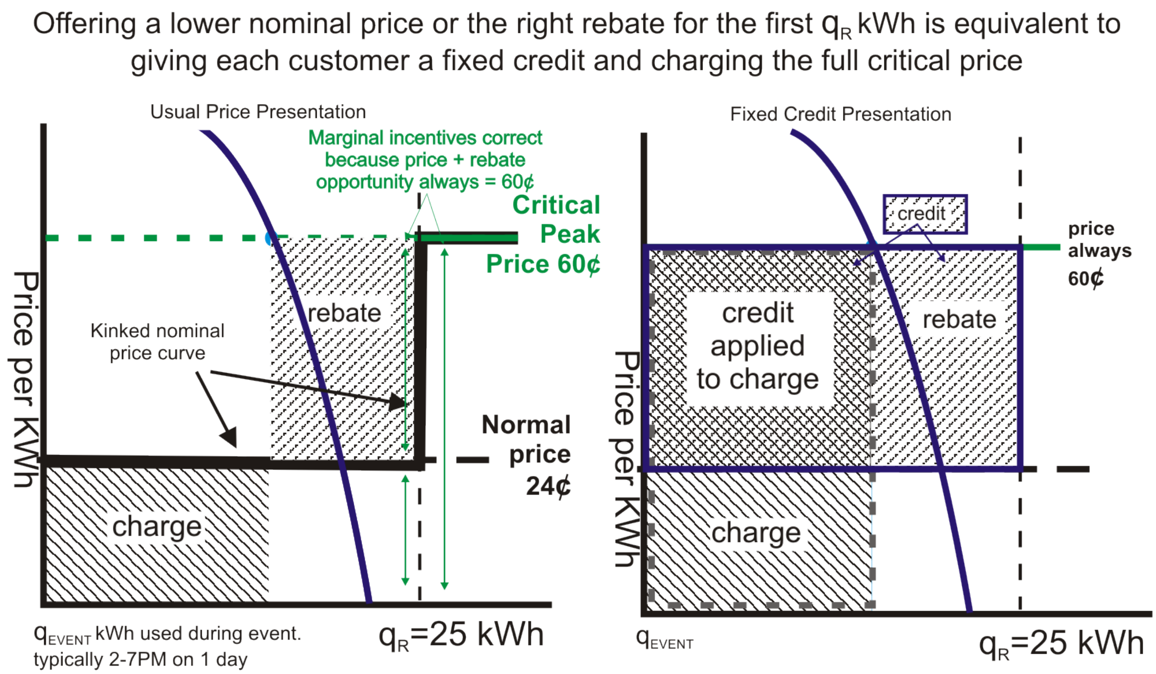

Figure 3.1: *Presenting IP rebates as the right to choose between power at the usual price or a rebate during each event is equivalent to presenting it as a fixed credit during each critical event. Both presentations keep the marginal incentives equal to the critical peak price.*

We have a great deal of flexibility in how to describe the rights that the customers own, but one attractive way to do so is to describe it not as a fixed bill credit during each event, but as a property right to access a set number of units of power at the reference price during each critical event. Customers can cash in unused rights for the value of the discount that the rights provide. Customers on the rate in table 3.2 get \$9 worth of rights during each critical event, which lets them access up to 25 kWh of power for 24 cents each instead of 60 cents and cash in the unused part of these rights for a rebate of 36 cents per kWh. Figure 3.1 shows the equivalence of the fixed credit and regular-priced-units presentations while section 3.5.2 proves their equivalence.

Loss-averse customers may be more receptive to the “usual-price” explanation (used in table 3.2) than to the “fixed-credit” presentation. The usual price presentation casts critical event incentives as rights to buy up to a fixed number of kWh at the usual

Rates are often marked up to recover fixed costs and regulators deal with shifts in demand on a regular basis.

price, and to earn rebates by automatically cashing in any unused rights. The usual price is the strongest candidate for the customer's reference price. The usual price presentation means that the customers who use less power than they had rights to buy at the usual price during an event will buy at their reference price and experience no losses. Further, this presentation allows bills to indicate that a customer never paid for power at the critical price whenever they had rights to more critical period power in that billing cycle than they used during events, even if they used more power than they had rights to during some events. By contrast, the fixed-credit presentation sells power to customers at a higher price, creating a perceived loss and then returns "lost" dollars through a credit that customers will code as a gain.[17] Since loss-averse customers put more weight on losses than gains, customers may perceive the offsetting gains and losses in the fixed-credit presentation as a net loss.[18]

This presentation of IP rebates preserves CPP's marginal incentive by keeping the sum of the price of the marginal unit and any foregone rebate equal to the CPP price for that time period. It is efficient and fair that all customers face the same CPP incentives to use power regardless of whether the customer is eligible for a rebate. The critical rate presented in table 3.2 achieves this by having customers with no rights left pay 60 cents per kWh, while those with rights face a 60 cent opportunity cost because they pay 24 cents and forgo a 36 cent rebate for each kWh.

3.4.2 Implementing the monthly fixed fee through declining block pricing

The marketing literature reports that customers are averse to paying the kind of fixed monthly fees that would be the natural way to charge customers for power rights without changing their incentives, but that customers are more receptive to almost identical incentives and charges presented through declining block prices (Ho and Zhang, 2004). A declining block rate marks up the first few units that each customer uses each month. For example, instead of using the $5 a month markup in the example in section 3.4.1, we could mark up the first 200 kWh of power the customer used by 2.5 cents each.

This markup does not change the consumer's incentive to buy the efficient amount since IP rebates' revenue neutrality means they return every cent that customers pay

[17]Prospect theory suggests that customers will have diminishing marginal sensitivity to gains which will further diminish the perceived difference in size between the gain from the full credit of presented by the fixed-credit presentation and the smaller rebate emphasized by the usual-price presentation.

[18]We could use a model like Koszegi and Rabin (forthcoming) to formalize this.

through it. Section 3.5.5 proves this. We aspire to have each customer buy the entire marked-up quantity each month in order to ensure that each customer buys the rights that the offer promised them. If the rate fails to offer a customer the rebates that it promised because the customer did not fully purchase their rights, customers may experience the kind of unexpected loss that the rate structure is designed to avoid.

It is desirable to keep marked-up offpeak power less expensive than the time invariant rate so 3-period CPP-IPR rates compare favorably to time invariant rates during both off peak and critical periods. This is attractive to customers who simply count the number of periods during which one rate outperforms another (Redden and Hoch, 2005). It also lets marketers claim that CPP-IPR offers lower prices more than 80% of the time – as Gulf Power does.

IP rebates that use declining block pricing to collect money and return it during critical events reschedule a significant part of CPP's savings. CPP delivers savings in a subtle way year round during offpeak times. CPP-IPR delivers many of those benefits in a visible way during critical events. Table 3.3 compares the proportion of customer bills that come from peak, offpeak, and critical periods under time invariant, CPP, and CPP-IPR rates. This change in the timing of charges over the year may cause small income effects, but these effects are likely to be negligible.

Table 3.2 provides an example of how CPP-IPR works in practice and how offer letters might explain it to consumers.[19] In this example rate, customers on the time invariant rate pay 14.6 cents per kWh regardless of when they use it. Customers on CPP pay more – a 24 cent per kWh "high rate"– during non-holiday weekday afternoons. Customers on CPP pay less – a 12 cent per kWh "low rate"– during off peak periods. More than 85% of all hours are offpeak. All hours except for weekday afternoons are offpeak inlcuding weekends, holidays, nights and mornings. The utility can notify customers by telephone that a period – typically a weekday afternoon – will be a critical period, invoking a 60 cent per kWh critical rate. Roughly 1% of all hours are critical. The CPP-IPR example rate is identical to CPP with three modifications:

1. The customer's first 450 kWh per month are marked up by 2.5 cents.

[19] The CPP rate in table 3.2 is based on Pacific Gas and Electric's "low ratio" experimental CPP rate (Pacific Gas & Electric, a). These CPP and CPP-IPR rates would raise the same amount of money as the time invariant rate did from the customers on time-invariant rates in California's Statewide Pricing Pilot. California has some of the highest electricity prices in the country.

Table 3.2: Examples of rates. The IP rebate offer here is appropriate for a high use customer with air conditioning in a hot climate.

		Price per kWh			
Price Period	Times in effect	**Time Invariant Rates**	**CPP**	**CPP-IPR**	
				initial price	beyond 450 kWh/mo.
Low	weekdays before 2PM; after 7PM; all day weekends & holidays	14.6 cents	12 Cents	14.5 cents	12 cents
High	Weekdays 2:00PM-6:59PM	14.6 cents	24 Cents	26.5 cents	24 cents
Critical	Announced with a telephone call at least 24 hours in advance	14.6 cents	60 Cents	First 25KWh: usual price for the period, as listed in the two rows above; 36 cent rebate for every kWh you save. Additional kWh after the first 25: 60 cents	

2. The customer has the right to access up to 25 kWh of power during an event at the usual price (typically the high price) for that period. If the customer uses less than 25 kWh during the event that lasts 5 hours or less, he also earns a 36 cent per kWh rebate for the difference between their 25 kWh of rights and their actual use. For example, a customer who used 20 kWh during a critical afternoon would get a rebate on $25kWh\ of\ rights - 20kWh\ used = 5kWh$, for a total of a \$1.80 rebates.

3. CPP puts a ceiling of 1% on the proportion of hours in which the utility can invoke critical prices, while CPP-IPR has a floor that requires the utilities to offer customers rebate opportunities equivalent to declaring 1% of all hours as critical periods. Section 3.4.3 discusses this in more detail.

3.4.3 Strict revenue neutrality: customers cannot profit by becoming eligible for extra rights if they pay a dollar for every dollar of rights that they get

It is important to design IP rebates to provide the appropriate level of rights to each person without creating perverse incentives. If an IP rebate program is strictly revenue

Table 3.3: Bills by time of electricity use. CPP and CPP-IPR generated lower bills than time-invariant rates would have for these California CPP customers. IP rebates rescheduled savings into critical periods. Data: California State Wide Pricing Pilot CPP customers, described in section 3.8.2, CPP-IPR benchmark offers described in section 3.9.3.

	offpeak	peak	critical	annual total bill
Time Invariant (% of bill and total kWh's used)	84.6%	14.1%	1.3%	$939
CPP	71.4%	23.8%	4.8%	$909
CPP-IPR	77.6%	24.9%	-2.4%	$909

neutral because it charges customers a dollar for every dollar of rights they get, the marginal change in bills with respect to a change in rights levels is zero. Thus, customers and utilities are economically indifferent about how many rights they get and face the same marginal incentives as CPP. This section considers how to use strict revenue neutrality to preserve incentives in the assignment of a rights level, in preventing customers from cashing in rights and then exiting the program before they paid for them, and in giving utilities the right incentives to call events.

Customers' electricity use is an important – but manipulable – signal about the quantity of rights a customer needs to avoid bill increases from critical peak events. If our IP rebate system uses customer consumption to adjust the number of rights we sell them, it is important to do so in a way that avoids creating perverse incentives. If a system lets a customer increase the value of the rights he gets in future periods by using more power now, it would implicitly change the price of present power consumption unless it increases the customer's future payments by an amount equal to the increase in rights. Consider a provision that gives the customer an extra $1 worth of rights if they increase consumption by q^*. If the provision increases their contributions by $1 for every $1 of rights that they get, it does not affect their incentive to increase consumption by q^*. If it changes their contributions by less (more) than $1, it creates a perverse incentive to (not to) increase consumption by q^*. Section 3.7 takes up this issue in more detail.

A clever calendar can help ensure person-level revenue neutrality

Sloppy handling of customer exit in mid-year or of under-contribution could break the person-level revenue neutrality. A clever calendar that concentrates events at the end of the fiscal year can, however, ensure that customers buy and pay for property rights before they have a chance to use them. This means that customers cannot profit from the program by strategically entering for just the peak season, claiming rebates, and then exiting the program. A calendar that clusters events at the end of the fiscal year also means it is always possible to reduce credit sizes if customers use too little power during a month and prevents customers from using rights during the summer and then underpaying for them in the fall. This calendar would make customers who leave CPP-IPR before the end of a fiscal year eligible to cash in their unused rights. This is preferable to leaving customers who exit owing the utility money or leaving their neighbors to pay for their rights.[20],[21] A standard fiscal year facilitates making equitable, revenue neutral program revisions since everyone would experience new charges or benefits at the same time.

Ensuring that customers get the rights they paid for

A good CPP-IPR implementation needs to handle year to year fluctuations in the number of times that weather and equipment problems justify critical events, while a CPP-IPR rate is designed to return the funds it raised to customers during a preset number of rebate opportunities per year. An attractive way to deal with this is to return the fixed credits for any unused event days. In other words, the customer would get the rebate they would have received if the critical periods had been called and the customers used zero power during the period.[22]

3.4.4 Psychological criteria suggest offering customers consistent rebates.

If customers dislike paying critical prices or experiencing bill spikes, then we can sell customers enough rights to ensure that most customers receive rebates during each billing cycle containing an event.[23] The IP rebate design means selling customers more credits will not reduce utility revenue or distort incentives. Thus, we aspire to provide most

[20]If customers are reluctant to contribute through declining blocks now to pay for future benefits, we could phase in the program with an abbreviated first fiscal cycle that began in late spring and called proportionally fewer critical events.

[21]The fiscal year approach would require a special rule, like calling new customers for a reduced number of critical events during their first year.

[22]Refunding the unused rights rather than calling an event will slightly reduce the utility's revenue since the utility forgoes the possibility of charging customers the critical price – in a way that is identical to the loss that a utility would take if it failed to call a critical day that it could call under a conventional CPP rate. Thinking about how to get incentives to call critical days right is an important issue, but is beyond the scope of this paper. By contrast, if utilities could simply pocket the fixed credits from unused event days, they would have a fairly strong incentive to not call events.

[23]Analogously, the US income tax system is tuned so that many citizens withhold too much and get refunds when they file rather than writing large checks.

customers with "consistent rebates".

Offering rebates so broadly may have some minor downsides. IP rebates might create incentives for customers to actively manage their air conditioning use at the cost of reducing their awareness of the total cost of their air conditioning. Customers might misinterpret consistent rebates as a sign that they were already managing their peak use well – especially if they were less motivated to find gains than to avoid losses. Ranking each customer's rebate size might avoid this misinterpretation by sending messages like "7 out of 10 of your neighbors earned significantly bigger rebates than you did last month."

We need to the amount of rights that the rate offers each customer carefully to deliver consistent rebates that the customer can fully fund through a declining block that is consistently inframarginal and keeps marked up offpeak power cheaper than the time invariant price. Section 3.8.4 shows that it is not hard to make offers meeting these criteria for most California customers.

3.5 A formal introduction to CPP and CPP-IPR

This section analyzes CPP and CPP-IPR and formally establishes that IP rebates are revenue neutral and preserve marginal incentives.

Consider a CPP rate with three periods: low-priced, offpeak periods (denoted "L"), higher-priced, peak hours ("H"), and the highest-priced, critical hours ("c"). During month m, denote the set of critical hours C_m, higher-priced, peak hours H_m and low-priced, offpeak hours L_m. The quantity of power that the customer uses during period i is $\mathsf{Q_i}$.[24] The rate sets prices for each period, denoted P_H, P_L, and P_c.[25] The total monthly bill, TC_m^{CPP}, under this rate is:

$$TC_m^{CPP} = P_c \sum_{c \in C_m} \mathsf{Q_c} + P_L \sum_{L \in L_m} \mathsf{Q_L} + P_H \sum_{H \in H_m} \mathsf{Q_H}$$

3.5.1 A formal overview of CPP-IPR

IP rebates change the presentation of CPP by adding charges and credits that sum to zero and that retain CPP's marginal incentives.

[24] Characteristics that vary by customer – like quantity consumed, $\mathsf{Q_i}$ – appear in sans serif.

[25] Rate characteristics like P_L and miscellaneous entries appear in the *math* typeface. Rate characteristics reflect local system costs and this document generally takes them as given in designing an IP rebate system. Section O lists the notation used in this document.

Consider a rate that calls N_c critical events per year and a month m in which the utility called $N_m \leq N_c$ events. The rate includes a declining block that imposes a markup of $\mathcal{M}$ on the first $\mathbf{Q_D}$[26] kWh.

3.5.2 Offering the right to units at the regular price and rebates is equivalent to offering an incentive-preserving fixed credit

There are a variety of ways to describe the rights that CPP-IPR customers get during critical events. Explanations for customers like table 3.2 report that CPP-IPR customers get access to up to $\mathbf{q_R}$ of power at the normal price, typically P_H, during a critical event. Section 3.4.1 discusses the rationale for that presentation. Further, if they use $\mathsf{q_e} < \mathbf{q_R}$ they get a rebate of $P_c - P_H$ per unit for any unused rights, $\mathbf{q_R} - \mathsf{q_e}$. Writing this out and multiplying through shows that this is mathematically equivalent to offering each customer a fixed credit that reduces bills by $\mathbf{R} = (P_c - P_H)(\mathbf{q_R})$ during each event.

$$
\begin{aligned}
TC_e^{CPP-IPR} &= P_H\mathsf{q_e} - (P_c - P_H)(\mathbf{q_R} - \mathsf{q_e}) \\
&= P_H\mathsf{q_e} - P_H\mathsf{q_e} + P_c\mathsf{q_e} - (P_c - P_H)(\mathbf{q_R}) \qquad (3.1) \\
&= P_c\mathsf{q_e} - (P_c - P_H)(qr) \\
&= P_c\mathsf{q_e} - \mathbf{R} \qquad (3.2)
\end{aligned}
$$

Equation 3.2 makes it clear that CPP-IPR customers face the same price, P_c, during an event as CPP customers do, but get a lower bill because they receive a fixed credit of $\mathbf{R}$. Equation 3.1 shows that the two formulations are equivalent because the sum of the usual price P_H and the forgone rebate, $P_c - P_H$ is equal to the critical price, P_c.

Customers who use $\mathsf{q_e} > \mathbf{q_R}$ pay P_c for their marginal power use. The calculations to show the equivalence between the two descriptions are analogous and are omitted here for brevity.

The balance of this analysis describes IP rebates as providing a fixed credit of $\mathbf{R}$ to both simplify its notation and draw attention to the fact that the credit of $\mathbf{R}$ does not affect the marginal incentive.

[26] IP rebate designers choose values for the variables listed in **bold**, including $\mathbf{Q_D}$, $\mathbf{R}$, and $\mathbf{q_R}$.

3.5.3 CPP-IPR Total bills and revenue equivalence for customers who buy all the rights the rate offers

This section defines the CPP-IPR monthly bill in the well-behaved case where customers buy $Q_m \geq \mathbf{Q_D}$ kWh in each month m.[27] Thus, they purchase all the rights the rate offers them, namely $\mathcal{M}\mathbf{Q_D}$ worth of rights per month. This sets up the proof that CPP and CPP-IPR generate the same total bill over the course of a year. These contributions provide customers with power rights worth $\mathbf{R}$ during each of N_c events. Revenue neutrality requires that the amount the customer pays through the purchase of marked up units equal the value of the rights the customer gets back, formally that $12\mathcal{M}Q_D = N_c\mathbf{R}$.[28]

Section 3.5.5 considers the analogous general case that maintains revenue neutrality even if customers buy less than the planned $\mathcal{M}\mathbf{Q_D}$ worth of rights in some months.

This customer's total monthly bill, $TC_m^{CPP-IPR}$, will be:

$$TC_m^{CPP-IPR} = \mathcal{M}Q_D - N_m\mathbf{R} + TC_m^{CPP} = \mathcal{M}Q_D - N_m\mathbf{R} + P_c \sum_{c \in C_m} \mathsf{Q_c} + P_L \sum_{L \in L_m} \mathsf{Q_L} + P_H \sum_{H \in H_m} \mathsf{Q_H}$$

3.5.4 CPP-IPR generates the same total annual bill as CPP for each customer

Each customer pays the same amount over the course of a year on a CPP-IPR rate that they would pay on the underlying CPP rate. The total annual CPP-IPR bill, $TC_a^{CPP-IPR}$, is simply the sum of the monthly CPP bills, TC_m^{CPP}, plus exactly offsetting rights and charges. We can see this by computing the total annual CPP-IPR bill, rearranging terms, and recalling that $12\mathcal{M}Q_D = N_c\mathbf{R}$, as follows:

---

[27] IP rebates change seasonal bill patterns by raising CPP bills by up to $\mathcal{M}Q_D$ each month and returning that money during (typically summer) months with events. A section below discusses how impacts on seasonal bill patterns vary by region. Most areas' highest use season coincides with the California summer peak, but usage in other regions peaks during the winter.

[28] I present these examples without calculating interest on the contributions to keep the algebra simple. It is technically correct to equate the two values in net present value, formally: $\sum_{m=1}^{12}(1+r)^{\frac{m}{12}}\mathcal{M}Q_D = \sum_{m=1}^{12}(1+r)^{\frac{m}{12}}N_m\mathbf{R}$ where r is the annual interest rate. The interest-free approximation differs from the net present value by less than the interest rate r, a few percent. Ensuring that the net present value of the rights and charges match exactly would require making minute adjustments to the value of the rights, $\mathbf{R}$, depending on the distribution of event dates, since designers have to set rates before the fact using estimates of when heat waves and equipment problems will cause critical days. Failure to adjust $\mathbf{R}$ or utility revenues to compensate for the timing of events creates tiny incentives for the utility to earn extra interest by calling events later in the year.

$$\begin{aligned} TC_a^{CPP-IPR} &= \sum_{m=1}^{12}[\mathcal{M}Q_D - N_m\mathbf{R} + P_c \sum_{c\in C_m} \mathsf{Q_c} + P_L \sum_{L\in L_m} \mathsf{Q_L} + P_H \sum_{H\in H_m} \mathsf{Q_H}] \\ &= 12\mathcal{M}Q_D - N_c\mathbf{R} + \sum_{m=1}^{12}[TC_m^{CPP}] \\ &= 0 + \sum_{m=1}^{12}[TC_m^{CPP}] \end{aligned}$$

3.5.5 The general case: dealing with customers who do not buy all of the offered rights

This section generalizes the discussion above to allow customers to buy less than $\mathbf{Q_D}$ kWh of power in some months, which means that they did not purchase a full $\mathcal{M}\mathbf{Q_D}$ worth of rights. It maintains revenue neutrality by only selling customers the rights they have paid for, $\hat{\mathbf{R}}$.[29] Selling customers only the number of rights that they pay for is consistent with treating this bill volatility control strategy as a well-defined property right – which section 3.7 discusses further. The customer buys rights worth the markup times the lesser of $\mathbf{Q_D}$ and their actual consumption, $\mathsf{Q_m}$ each month. Formally, they buy rights worth $\mathcal{M}\min\{\mathbf{Q_D}, \mathsf{Q_m}\} \leq \mathcal{M}\mathbf{Q_D}$.

This implies that customers own rights worth $\hat{\mathbf{R}}_c$ during event c. Customer level revenue neutrality requires that the sum of adjusted right values equal the sum of customer contributions, formally:

$$\sum_{c=1}^{N_c} \hat{\mathbf{R}}_c = \mathcal{M} \sum_{m=1}^{12} \min\{\mathbf{Q_D}, \mathsf{Q_m}\}$$

Strategies to restore person-level revenue neutrality reduce bills relative to the full contribution scenario during the months when the customer contributes too little and increases bills during months in which the customers get a reduced credit of $\hat{\mathbf{R}}_c < \mathbf{R}$. Reference dependent people may perceive these adjustments as a gain and a loss of the same size, which they would take as a net loss since they weigh losses more heavily than gains.[30]

The cumulative deficit, δ_m in month m, is a deficit in a customer's purchases of

[29]Equivalently, we could return to the original plan by marking up more than $\mathbf{Q_D}$units of power and buying the missing units in a later month.

[30]Customers may be equally frustrated that issues in the fine print of their offer letter are costing them money.

rights for future events relative to the value of rights that the rate slated for the customer. The deficit grows when monthly consumption is too low, $\mathsf{Q_m} < \mathbf{Q_D}$, and shrinks when the customer gets fewer rights during an event.[31] Formally, the definition is:

$$\delta_{\mathsf{m}} = \min\{0, N_m\mathbf{R} - \delta_{\mathsf{m-1}} - \mathcal{M}(\mathbf{Q_D} - \min\{\mathbf{Q_D}, \mathsf{Q_m}\}\} \tag{3.3}$$

Formula 3.3 defines the cumulative deficit as the sum of that month's deficit and the previous deficit, less any amount of the deficit that can be applied to reduce the value of the rights offered during events in that month. There can never be a positive deficit so the deficit returns to zero after enough rights are applied to it.

We can ensure budget balance by offering rights each month worth

$$N_m\hat{\mathbf{R}} = \max\{0, N_m\mathbf{R} - \delta_{\mathsf{m-1}} - \mathcal{M}(\mathbf{Q_D} - \min\{\mathbf{Q_D}, \mathsf{Q_m}\}\}$$

Notice that this reduces back to the full contribution case considered above if the customer buys at least $\mathbf{Q_D}$ kWh each month, so $\delta_{\mathsf{m-1}} = 0$, $min\{\mathbf{Q_D}, \mathsf{Q_m}\} = \mathbf{Q_D}$, and $\hat{\mathbf{R}} = \mathbf{R}$.

Using a declining block rate to fund rebates never creates a deadweight loss

This rate, unlike most declining block rates, does not create a deadweight loss because every extra dollar that a customer pays for through this rate's markup, $\mathcal{M}$, comes directly back to the customer as an extra dollar of rights.[32]

Customers who buy less than $\mathbf{Q_D}$ units on a typical declining block rate end up paying $\mathcal{M}$ more for their marginal unit. This price increase reduces purchases and creates a deadweight loss.

Formally, consider the marginal incentives for a customer to buy one more unit of power during period $i \in \{c, H, L\}$ for a month m when $\mathsf{Q_m} < \mathbf{Q_D}$. Taking the derivative with respect to the total annual bill shows that the cost of the marginal unit includes an

[31] This algebra, for simplicity, closes the entire deficit at the first available event. There are equivalent, perhaps more palatable, approaches that would spread the reduction over multiple events where possible.

[32] Customers may not notice this subtle connection, but the markup is still unlikely to cause significant distortions because many customers are unaware of whether the quantity they have consumed so far during a month means they are paying a markup on the margin and because demand at the offpeak and peak prices is quite inelastic. All we really need customers to know for CPP-IPR to work well is that there is some economic reason for them to shift power use away from weekday afternoons, and stronger reason to shift power use when the utility notifies them of a critical period.

increase in price of $\mathcal{M}$, but the increases in future rebates of $\mathcal{M}$ exactly offsets the price increase:

$$\frac{\partial TC_a}{\partial \mathsf{Q_m}} = P_i + \mathcal{M} - \frac{\partial \delta_\mathsf{m}}{\partial \mathsf{Q_m}} = P_i + \mathcal{M} - \mathcal{M}$$

3.6 IP Rebates generalize to many pricing challenges

IP rebates generalize to work with a wide variety of dynamic pricing plans that improve incentives over uniform pricing for products which have underlying costs that fluctuate over time. The generalized IP rebate approach will offer each customer:

- rights to buy a block of the product at the usual nominal price during the high priced periods,
- rebates for any unused rights, and
- opportunity costs to purchase the product that are closer to the cost of production during both the high and low cost periods.

Further, an IP rebate implementation will often cause a smaller change to the customer's uniform pricing annual bill patterns than a conventional implementation of dynamic pricing would.[33]

The uniform pricing that we seek to improve can generally only survive in the context of exclusive, long term relationships with a single supplier. Customers who choose frequently among competing suppliers will purchase from firms with dynamic pricing during low cost seasons. Further, cheap storage opportunities will smooth differences in cost over time, so this approach will have the greatest benefits for products that are not cost effective to store. It is common for companies to provide difficult-to-store products through long term, exclusive relationships in utilities, in telecommunications, and in services like shipping, answering phones at call centers, or technical support.

[33] This IP rebate approach will not change a customer's uniform pricing bill patterns at all if the customer has zero demand elasticity and consumes the average ratio of peak to off peak power. In other words, the customer who sees no bill impact is neither gaining nor losing from the price insurance inherent in uniform pricing. By contrast, consider a customer who simply buys the product during the lowest cost season and uniformly spread her demand among seasons under uniform pricing. She would see a larger change under IP rebates than they would under conventional dynamic pricing. Under both pricing schemes, this customer will shift all of their spending to the low cost season. Under an IP rebate scheme that marks up the product during and then returns those markups during the low cost season, her bill increase (decrease) would be larger during the low (high) cost seasons than it would be under the conventional dynamic pricing approach.

I simplify this discussion by assuming that the firm prices at average marginal cost plus a uniform markup. The uniform markup makes the firm indifferent between selling high cost and low cost units of the product.[34],[35] This assumption makes the terms "high cost periods" and "high price periods" interchangeable in the discussion below.

Intuition: Uniform pricing makes customers pay a markup during low cost periods that covers extra production costs during high cost periods. We can set opportunity costs that better reflect production cost during each period. Then we can divert each customer's markup to buy him a credit that he can either either use to buy the product during high cost periods or keep as a rebate. These markups can preserve the nominal prices during low cost seasons and the credits can preserve the nominal prices during high cost seasons for every unit that the customer buys rights to.

Proof: We can decompose the uniform price paid during low priced periods into the low cost and the markup, then applies the total markup paid to buy an equal value of refundable rights to buy a fixed quantity of power at the uniform price during the high priced period. Specifically:

Consider a market with low and high cost sets of time periods with prices $\bar{P}_L$ and $\bar{P}_H$ respectively. The high price can be decomposed into the low price plus a price increase, formally $\bar{P_H} = \bar{P_L} + \Delta P$.

Let f_H be the fraction of all consumption at uniform prices P_u that takes place during high cost hours. Then, setting a uniform price of $P_u = \bar{P_L} + f_H \Delta P$ will raise the same amount of revenue as would selling the same amount of product during each time period, but charging $\bar{P_L}$ and $\bar{P_H}$ for it. Let $\mathcal{M}_u = f_H \Delta P$ be the markup that customers pay during low priced periods to offset the cost of the expensive product. Then customers who buy the population average proportion of the product, f_H, during high cost periods pay in exactly as much in markups as they get back in reductions of the price of the high cost product.

We can move prices closer to costs by setting each customer's peak period opportunity costs to $\bar{P}_H$ and converting each customer i's payment of $Q_L^i \mathcal{M}_u$ into a bill credit of **R** that the customer gets regardless of his critical period use. An IP rebate style description

[34]Adams and Yellen (1976) show that pricing that makes the firm indifferent between selling two different products can be part of an optimal bundling strategy.

[35]Changing to either kind of dynamic pricing will generally change the total quantity that the firm sells and often change the seller's profits. I ignore the profit issue here because utility regulators can adjust rates to ensure that the utility earns its rate of return despite the change in quantity and because any welfare improving change in pricing creates a potential Pareto improvement that can increase firm profits.

would present this as rights to buy $\mathbf{q_R} = \frac{Q_L^i \mathcal{M}_u}{\bar{P}_H - \bar{P}_u}$ units at the uniform price of $\bar{P}_u$ where Q_L^i is the customer's consumption during low priced-periods. Offering a rebate $\bar{P}_H - \bar{P}_u$ for each unused unit to make the opportunity cost $\bar{P}_H$. This approach also implicitly lowers the opportunity cost of consuming off peak to P_L because the customer gets every cent they pay through the markup of $\mathcal{M}_u$ back. This strategy moves opportunity costs closer to true costs while maintaining nominal prices.

That strategy generalizes to pricing schemes that further subdivide the high and low cost periods into any number of subsets. The generalization requires that the customer pay markups during low priced periods that equal the value of the rights the customer gets back during the high priced periods, or:

$$\sum_{i \in L} \mathcal{M}_u^i Q_L^i = \sum_{h \in H} \mathbf{R}^h$$

if $L(H)$ represents the set of low (high) cost periods. It generalizes to a single high and single low price period (e.g. CPP with just low and critical price periods), cases with a few periods per subset (e.g. CPP with low, high, and critical periods), and to cases with a very large number of price periods per subset (e.g. real time pricing that charges the market price every hour).

3.6.1 This two period generalization does not perform as well as the three-period CPP-IPR approach

This two period generalization is harder to explain to customers than the three-period CPP-IPR approach and will not offer consistent rebates.

- Each customer get rights as a function of their usage. The fluctuations in rights levels may be hard to explain to customers.
- This approach leaves the nominal price during low-cost periods at P_u. It may be difficult to explain that customers pay $P_u = P_L + \mathcal{M}_u$, but that the opportunity cost is really P_L since the customer gets the $\mathcal{M}_u$ component of the price back. Customers who do not understand that the opportunity cost has dropped to P_L may inappropriately continue to consume as if the opportunity cost were the higher P_u .
- The two period implementation does not offer consistent rebates to most customers.

> Customer-base wide revenue neutrality implies that, in the absence of demand elasticity, the average customer in the population would get rights to exactly as much power as they use during high priced periods. Customers who use a greater than average proportion of their energy during high priced periods will not get a rebate and will pay the high price on the margin. However, demand elasticity will increase the number of customers getting rebates. Elasticity increases purchases during low priced periods that come bundled with rights and decrease the use of rights to buy expensive power.

Further, this approach creates winners and losers relative to uniform pricing. Customers who used a smaller (larger) than the population average proportion of their power during high-priced periods will see their total annual bills decrease (increase) under the new pricing. For example, if P_u reflects the fact that the average customer buys 75% of his purchases of the product during low cost periods, a customer who buys 80% of her total consumption during low cost periods will see her bill for the same bundle drop because her average unit will now cost $.8P_L + .2P_H$ rather than the uniform cost of $P_u = .75P_L + .25P_H$. The new pricing eliminates the cross subsidy that offering unlimited access to high cost product at $\bar{P}_u$ provides. This flaw is typical of dynamic pricing approaches unless they are designed from the ground with complex, hard to explain and implement, features to preserve existing cross subsidies and deliver a Pareto improvement.

3.6.2 Concentrating credits on selected parts of the high-priced period can address these shortcomings

CPP-IPR works so well because it splits high cost periods into a set of "high" priced periods without rights to buy at the usual price and "critical" periods that get such rights. This is a generally applicable strategy that frees up cash to address some flaws in a two-period IP rebate implementation. Marking up the same number of low priced units while offering credits for fewer purchases creates a surplus of potential credits. This surplus can be used to offer more customers consistent rebates, to reduce nominal offpeak prices, or to implement a declining block and a fixed credit size.[36]

[36] I have made no assumptions about the proportion of unhedged hours in the unhedged high price bin – so it is not clear how much cash is available to fund consistent rebates for more customers or to move to a declining block implementation.

3.7 Comparing CPP-IPR to other rates

CPP-IPR offers better economic incentives than existing, time invariant and baseline-rebate rate designs, while being more compatible with customer decision-making heuristics than CPP. The major existing rate designs are:

- Most customers are on time invariant pricing and seem satisfied. Time invariant pricing gets prices wrong during almost all hours which leads to enormous waste and to significant cross subsidies. Specifically, it charges a uniform price, P_u, during every period. P_u is too high off peak hours and too low during peak and critical periods.
- Dynamic rates including CPP create significantly better incentives than time invariant rates but consumers resist signing up for them.
- Baseline rebate rates, discussed at length below, create dynamic incentives while using behaviorally astute rebate opportunities – but also create perverse incentives for customers to distort their consumption patterns to become eligible for larger rebates.

Figures 3.2 and 3.3 graphically compare their incentives to CPP and CPP-IPR.

The rates that consumers accept include hedges or other features that dampen bill volatility and reduce exposure to high prices by default, while the dynamic rates customers reject generally make risk management optional if it is available at all.[37] There may be both good customer perception and conventional economic risk management reasons to manage bill risks given the volatility of electricity prices. Borenstein (2007) reports that RTP leads to significant increases in bill volatility, but that simple hedges can control that volatility. It makes sense to make risk management that uses well defined property rights part of the default rate.[38]

The existing rates that manage volatility generally fail to manage volatility through well defined property rights and thus create flawed incentives. For example, time invariant

[37] The rights that IP rebate customers buy ahead of time have bill-volatility reduction effects that are akin hedging by buying ahead on the futures market, but differ from conventional hedging in that there is no unknown state of the world that affects the realization of the price of the commodity when the customer uses their forward rights, although there is uncertainty about the number of critical events in each billing period and about factors that affect demand during an event like weather and whether the event falls during the customer's vacations.

[38] It is clear that zero expected cost risk management or month-to-month volatility reduction would make many customers happier for both behavioral and neoclassical reasons. IP rebates are in this category since they provide features that reduce month to month bill volatility at zero cost to the consumer while their dynamic pricing lets customers reduce their overall bill. Rational customers, however, should be willing to pay only a very small risk premium to reduce the risk of a bill spike of a few tens of dollars (Rabin, 2000).

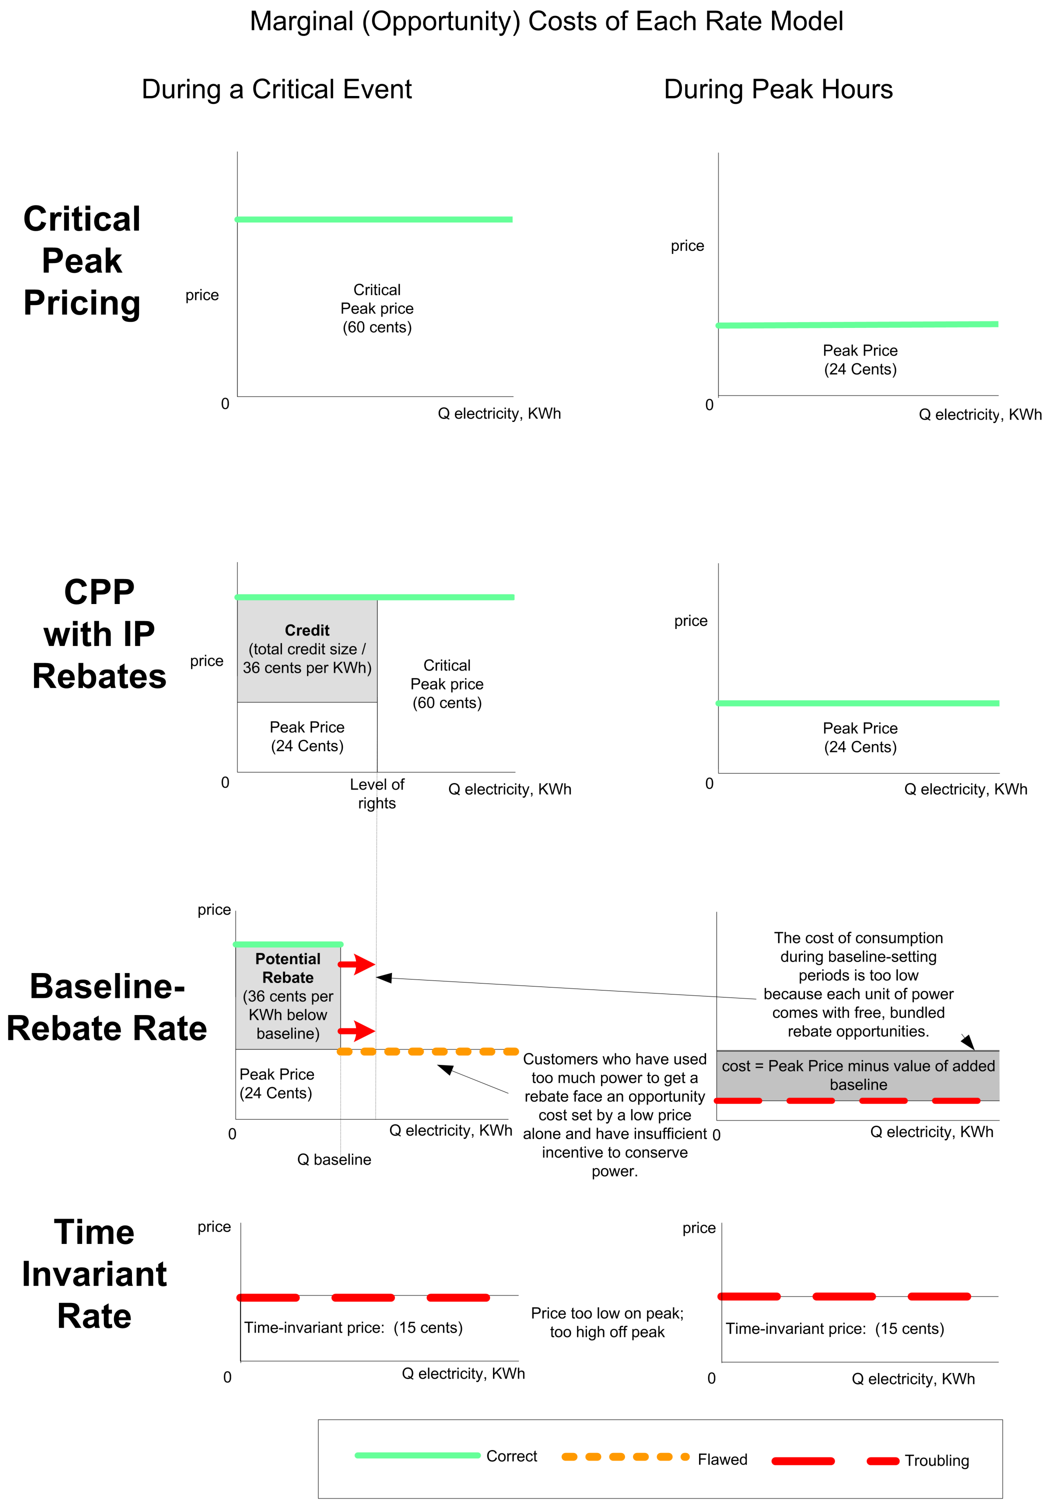

Figure 3.2: *These diagrams compare the plans' marginal incentives during critical events and peak periods, including the peak periods that may set baselines for the critical events.*

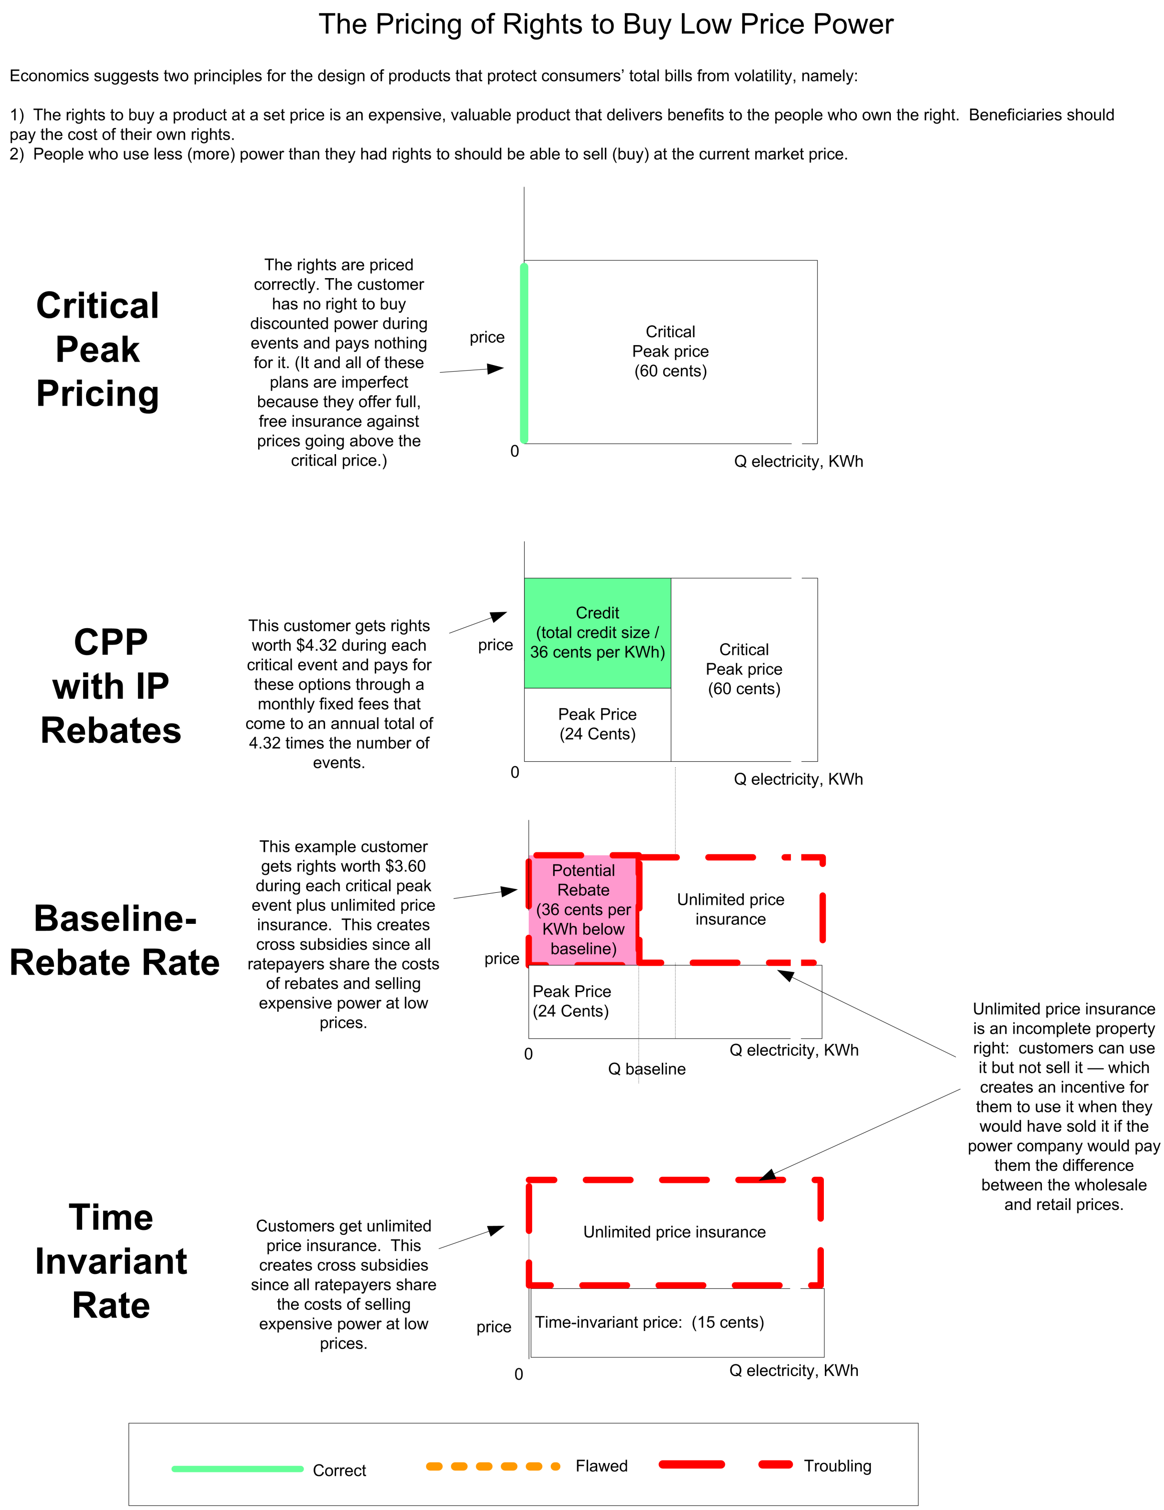

Figure 3.3: *These diagrams compare the plans' pricing of rights to access low cost power during events.*

rates include a built in mandatory hedge that gives customers a use-it-or-lose-it right to as much cheap power as they want. Customers use too much power during periods when wholesale power is expensive and they would prefer to sell their rights if they could sell them for their true cost. Time invariant rates and CPP-IPR ask customers to contribute toward insurance against critical events year round, while plain CPP does not – so IP rebates are a more incremental change from the time invariant status quo than CPP would be. Baseline-rebate rates bundle a "free" hedge with electricity purchases during baseline setting periods – which are typically weekday afternoons during the hottest months – which creates a perverse incentive, discussed at length below, for customers to consume more during the baseline setting periods to get more rebates during events. Baseline rebate rates introduce perverse incentives because they succumb to the temptation to develop ad hoc solutions to provide rebates and create incentives to shift away from critical periods rather than doing so with well defined property rights.

Dynamic rates generally offer the foundations for fairly well defined property rights, but generally include no volatility management by default and – at best – leave it to customers to acquire this kind of hedge separately. Bundling default rights that are priced at marginal cost addresses this omission. Doing so is a major step forward, but the remaining challenges include choosing the level of protection and paying for it in ways that are economically efficient and psychologically attractive.

There are compelling psychological and economic reasons to make a revenue-neutral volatility dampening mechanism or actuarially fair hedge the default.

- Making it a default minimizes transaction costs. The bill shocks that are involved are fairly small, and do not justify customers' spending hours to understand plans and choose a hedge.

- People often refuse to choose when they face too many choices (Iyengar and Lepper, 2000; Dhar, 1997); and are strongly influenced by default offers (Choi et al., 2003).

- Incentive preserving rebate type interventions are designed not to affect marginal incentives or annual per-customer revenues. A revenue neutral bill volatility reduction strategy – like those proposed in IP rebates – has no effect on total annual bills, so economically rational customers who really understand the program should be nearly indifferent between the default IP rebate eligibility-size offers and a menu of alternatives. Similarly, an actuarially fair offer has zero impact on total bills in expectation.

3.7.1 IP rebates avoid the economic flaws in baseline-rebate rates

Some utilities have used baseline-rebate rates that are superficially similar to IP rebates. Baseline-rebate rates calculate a personalized baseline demand level from each customer's consumption history and then offer rebates to customers who use less than their baseline level during critical events. Baseline rebate programs, unlike IP rebates, create troubling cross subsidies and create flawed incentives both during ordinary periods that are used to set the baseline and sometimes during critical events.

Utilities have fielded baseline rebate rates in a variety of contexts. Utilities generally make every customer eligible to earn rebates because baseline-rebate plans are described as containing only rewards for conservation.[39] Thus, baseline rebate programs tend to expose far more customers to improved incentives than opt-in dynamic pricing programs do. Wolak (2006) analyzes an experiment with a baseline-rebate dynamic pricing program in Anaheim, California and reports that consumers reduced use during critical periods but many customers exploited the poor incentives. San Diego Gas and Electric has proposed offering all of its customers a "Peak Time Rebate" rate based on the Anaheim design.[40]

California utilities offered a "20/20" baseline rebate plan during its electricity crisis that offered customers a 20% rebate on their electricity bill if they reduced their total summer electricity use 20% below their use the previous summer. California Utilities offered a "10/20" natural gas baseline rebate program during a price spike in Winter 2005-06 that offered a 20% rebate for reducing gas consumption at least 10% relative to the previous winter. Utility staff and regulators report that they dislike baseline-rebate rates. The section below describes how baseline rebate rates work and then lays out three significant flaws of baseline-rebate rates: perverse incentives during baseline-setting periods; inconsistent incentives during critical periods; and significant revenue impacts.

Baseline rebate mechanics: A baseline-rebate rate customer gets rebates for getting consumption below their baseline usage level, which is a function of their "normal"

[39] Baseline-rebate programs tend to include every customer by default on the claim that they provide only rebate opportunities – carrots without visible sticks. They often quietly recover rebate costs by raising every customers' rate later. This means that some customers would have done better had they been able to opt out of the rebate program and the responsibility to pay for it. PG&E's proposal for its 10/20 baseline-rebate program reads in part "[T]he 10/20 Winter Gas Savings Program is forecasted to pay out $200 million in rebates...PG&E proposes that these costs...be recovered in residential and small commercial customers' transportation rates during the summer gas season..." (Pacific Gas & Electric, b, 4).

[40] Both the Anaheim and San Diego rate designs comply with a California law, AB1X, that limits aspects of utility rates and makes it difficult to implement CPP. AB1X will sunset once a set of obligations from the California electricity crisis are paid off and may get amended or repealed even before that date.

behavior during similar, but non-critical periods.[41] The baseline amount for a baseline-rebate program applied to critical electricity periods, $\bar{\mathsf{Q}}_{\mathsf{bt}}$, is generally the customer's average use, $\mathsf{q}_{\mathsf{t-i,H}}$, during the set of N_b weekdays afternoons or peak periods (hence the subscript H), $t - N_b \cdots t - 1$, before the event at time t.

$$\bar{\mathsf{Q}}_{\mathsf{bt}} = \frac{1}{N}\sum_{i=1}^{n} \mathsf{q}_{\mathsf{t-i,H}}$$

The customer's total bill under a baseline rebate plan is:

$$TC^{baseline} = P_L \sum_{L \in L_m} \mathsf{Q_L} + P_H(\sum_{H \in H_m} \mathsf{Q_H} + \sum_{c \in C_m} \mathsf{Q_c}) - \mathcal{P}_B \sum_{c \in C_m} \max\{0, \bar{\mathsf{Q}}_{\mathsf{bt}} - \mathsf{Q_c}\} \tag{3.4}$$

Baseline-rebate rates create perverse incentives during baseline-setting periods. Baseline-rebate rates bundle free baseline rights with power during baseline setting periods. This makes power artificially cheap and gives customers incentives to increase usage during baseline setting periods. Sometimes the rates offer a negative cost of power during baseline setting periods that pays customers to use power. To see this, substitute the formula for $\bar{\mathsf{Q}}_{\mathsf{bt}}$ into the baseline-rebate bill formula, 3.4, and take the partial derivative with respect to the quantity of power used during the baseline-setting period. The result is as follows, assuming for notational simplicity that the customer is getting a rebate:[42]

$$\frac{\partial TC}{\partial Q_i} = P_H - \mathcal{P}_B \sum_{i \in Bt} \frac{1}{N}$$

This formula also reveals that this distortion becomes small (large) as the baseline setting period gets large (small). Making the baseline-setting period large, however, is likely to include cooler weather in the baseline and thus to make it a less accurate estimate of what people would have been doing on the critical day in the absence of an incentive to conserve.

[41] Situations where the customer knows his true demand for electricity, but the utility can only knows his usage level are asymmetric information games. It is generally difficult to design efficient mechanisms to get customers to reveal their types. In the absence of mechanisms designed to minimize distortions, customers have large incentives to strategically misrepresent their usual consumption during the baseline setting periods.

[42] There is no distortion for customers who will never get rebates. And there is a tedious, unenlightening corner case for customers who switch from no rebates to rebates as they use more during the baseline period.

These bundled baseline rights create significant changes in incentives. For example, San Diego Gas and Electric's proposed baseline-rebate rate offers a 65 cent rebate for every kWh that a customer's period use is below a baseline set by the customer'average use on the five non-event weekdays preceding the event day (Gaines, 2006). This means that San Diego customers get another roughly[43] $\frac{1}{5}$ kWh of baseline rights, worth 13 cents, bundled with every baseline-setting kWh. Residential customers in San Diego pay between 4 and 18 cents per kWh of power (San Diego Gas and Electric).[44] The bundled rights can be worth even more if one day sets the baseline for more than one critical event if there were fewer than five non-event weekdays between events. Anaheim offered a 35 cent rebate and used the average of the consumption during the three highest use non-event weekdays of the summer season as its baseline for every event. Since a single additional kWh consumed on a baseline-setting day increases the baseline by $\frac{1}{3}$ kWh over 12 events, this unit that costs either 6.75 or 11.07 cents comes bundled with rights worth $1.40 to a customer getting rebates (Wolak, 2006, 14).

This distortion is more disturbing because it increases demand for expensive power. Baseline-setting periods are typically moderately hot weekday afternoons when wholesale power is moderately scarce and expensive.

Baseline-rebate rates offer customers who have used too much power to get rebates an unlimited amount of power at the usual price and does nothing to give these customers an extra incentive to reduce usage on the margin during events. The class of customers who generally consume more than their baseline quantity during critical events will quickly learn that they cannot earn rebates and have no incentive to conserve. Baseline-rebate rates, like time invariant rates, implicitly include the cost of mandatory, unlimited critical period price insurance in the price of basic electric service.

These are expensive programs that create cross subsidies and uncomfortable trade offs: Baseline-rebate rates are not revenue neutral for individual customers which means that baseline-rebate programs create unpredictable rebate costs that utilities will need to recover later. Utilities recover the costs of the rebate program by marking

[43]San Diego's proposes to set its baseline by multiplying the average consumption during the baseline-setting period by a scaling factor, namely the ratio between the system wide demand on baseline-setting days and the critical day. This lets them correct for differences in demand – especially in weather-driven air conditioning demand – between the baseline-setting and critical days.

[44]Most California utilities – including San Diego Gas and Electric and Anaheim Public Utilities – use an increasing block rate structure that offers the first few kWh per month at a low price, then increases the marginal price as customers use more.

up power during all the non event periods, which can create inequitable cross-subsidies. Thus baseline levels can be mistakes since they both transfer cash among customers and create incentives, while IP rebates' choice of $\mathbf{q_R}$ neither creates incentives[45] nor transfers cash among consumers. Hence, baseline-rebate designers aspire to calculate baselines that reflect precise predictions of how much each customer would have used in the absence of the rebate opportunity. The limited data available to regulators and significant, normal day-to-day variation in power use mean that many customers' event usage would deviate significantly from their baselines even in the absence of rebate opportunities. Customers who get baselines above what they would have used on the day under the normal incentives get socially expensive "structural" rebates , while customers who would use far more than their baseline levels get no incentives to save.[46] By contrast, IP rebate designers can choose levels of rights that offer almost everyone rebates and are using a rate design that provides a constant opportunity cost of P_c.

In sum, IP rebates avoid flaws in baseline rebate designs. The flaws in baseline-rebate rates reduce efficiency and focus attention on dealing with perverse incentives, baseline estimation challenges, and the redistributive effects of the baselines rather than the real challenge of reducing consumption during critical and peak periods. Baseline-rebate rates offer customers opportunities to reduce their total annual bills through strategic baseline-manipulation without lowering the social cost of electricity provision.

3.8 The political, organizational, and financial feasibility of CPP-IPR: Evidence from California Customers

CPP-IPR implementations need to meet administrative and financial constraints. Data from a California CPP pilot study shows that most customers meet the financial constraints on CPP-IPR. The central, interlocking feasibility issues are:

- **Administrative and political feasibility.** CPP-IPR has to coexist with existing analytic categories and be an incremental change from existing rates. It has to give regulators the flexibility to address local equity concerns and distributional concerns.

[45] The underlying CPP rate creates the incentives in CPP-IPR; IPR's simply presents those incentives in a more palatable way.

[46] San Diego Gas and Electric also reports that more than a quarter of their customers had usage in absence of the critical event that would have either given them a rebate or have given them a baseline that would require them to reduce usage 15% before they got a rebate.

- **Economic feasibility.** CPP-IPR works well if we can assign each customers a revenue neutral pair of a rights size, $\mathbf{q_R}$, and a declining block size, $\mathbf{Q_D}$, that is likely to work well for the customer. An offer that works well has enough rights that it never leaves the customer paying a high nominal price for power during a month that contains an event. And the customer needs to use at least $\mathbf{Q_D}$kWh during each month in order to buy all of the rights that the rate offered the customer. Rebates are only feasible if a customer's demand pattern means that this kind of offer exists. Utilities need to be able to make these offers using only limited information about customers' demand patterns and, probably, a limited amount of flexibility to customize offers.

3.8.1 The central economic feasibility constraints: consistent rebates, inframarginal declining blocks, and revenue neutrality

IP rebates make each customer an **offer**, $(\mathbf{q_R}, \mathbf{Q_D})$, which specifies the quantity of rights that the customer gets during each event, $\mathbf{q_R}$, and the number of kWh the declining block marks up each month, $\mathbf{Q_D}$. It is desirable for offers to meet the following constraints for as many customers as possible:

i. **Consistent rebates:** The offer includes enough kWh at the usual price so that the customer gets a (weakly positive) rebate during each month with an event, or $\mathbf{q_R} \geq \bar{\mathsf{q}}_\mathsf{c}$. In other words, the number of protected kWh, $\mathbf{q_R}$, has to be at least as great as the customer consumed during the average event in the customer's highest average-event-use month $\bar{\mathsf{q}}_\mathsf{c} = \max_{m \in M}\{\mathsf{Q_c}/N_m\}$.

ii. **Consistent purchases through inframarginal declining blocks:** Customers buy all of the rights that the offer promised only if the declining block marks up less power than the customer uses each month[47], or $\mathbf{Q_D} \leq \underline{\mathsf{Q_m}}$.

iii. **Customer-level revenue neutrality:** each customer makes payments for rights equal to the value of the rights they receive. If the customer consistently purchases $\mathbf{Q_D}$ per month (constraint ii) then, this becomes $12\mathcal{M}\mathbf{Q_D} = N_c\mathbf{q_R}(P_c - P_h)$.

[47] I assume that $\mathbf{q_R}$ and $\mathbf{Q_D}$ stay the same year round. Making seasonal changese to the declining block size, $\mathbf{Q_D}$, may be an important way to provide consistent offers. Requiring extra contributions early in the year could provide a reserve fund to cover under contributions later. It may be particularly natural to consider seasonal variations in $\mathbf{Q_D}$ or $\mathcal{M}$ in electricity systems that already seasonally adjust rates. Seasonal adjustments, however, make rates harder for customers to understand.

IP rebate offers must satisfy revenue neutrality constraint[48] iii and aspire to do so while meeting consistent rebate and inframarginal declining block constraints for as many customers as possible. Throughout the discussion below, an offer is **consistent** if it satisfies the consistent rebates constraint i and the consistent rights purchase constraint ii over the course of a year.[49]

Substituting the first constraints i and ii into constraint iii, we discover a criterion that determines whether an offer exists that marks up $\mathbf{Q_D} \leq \underline{\mathsf{Q_m}}$ each month and provides consistent rebates, namely:

$$12\mathcal{M}\underline{\mathsf{Q_m}} \geq 12\mathcal{M}\mathbf{Q_D} = N_c\mathbf{q_R}(P_c - P_h) \geq N_c\bar{\mathsf{q}}_\mathsf{c}(P_c - P_h) \tag{3.5}$$

Dropping out the middle terms that specify a revenue neutral offer creates a feasibility criterion that depends only on customer characteristics and characteristics of the rate that we take as given. The criterion implies that an offer exists only if:

$$12\mathcal{M}\underline{\mathsf{Q_m}} \geq N_c\bar{\mathsf{q}}_\mathsf{c}(P_c - P_h) \tag{3.6}$$

Figure 3.4 visualizes the three constraints and their implications by plotting the use during events on the x-axis and monthly use on the y-axis. The y-axis plots the offer's requirement that the customer use at least $\mathbf{Q_D}$ kWh per month. We can plot the customers' use patterns that make $\underline{\mathsf{Q_m}}$ available on the same axis. The customer's right to buy $\mathbf{q_R}$ kWh per critical event at the usual nominal price can be plotted on the x-axis. The customer's need to get $\bar{\mathsf{q}}_\mathsf{c}$ kWh of rights to get consistent rebates can also be plotted on the x-axis.

IP rebate offers are consistent if they assign each customer a value of $\mathbf{q_R}$ that satisfies feasibility condition 3.5. Specifically, an offer is consistent if it provides **consistent rebates** paid for through a declining block that the customer **consistently purchases**. Providing consistent offers to most customers requires that the distribution of customers have two characteristics:

[48] Any deviation from the revenue neutrality constraint means that the rebate program will sometimes pay a customer more or less in rebates than they contributed to buy their rights, which creates flawed incentives that customers can exploit.

[49] This project focuses on maximizing the probability of getting a consistent offer during each customer year. Since $\underline{\mathsf{Q_m}}$and $\bar{\mathsf{q}}_\mathsf{c}$are a minimum and a maximum, respectively, so the more observations they consider, the more extreme results they will report. From a policy perspective, it is interesting to know that the Park family's annual values of $\underline{\mathsf{Q_m}}$and $\bar{\mathsf{q}}_\mathsf{c}$supported consistent offers in 19 of 20 years. It is less interesting to know that calculating $\underline{\mathsf{Q_m}}$' and $\bar{\mathsf{q}}_\mathsf{c}$' over 20 years picks up outlying values – like an extended vacation and running the dryer during a critical period – and makes it impossible to find a consistent offer.

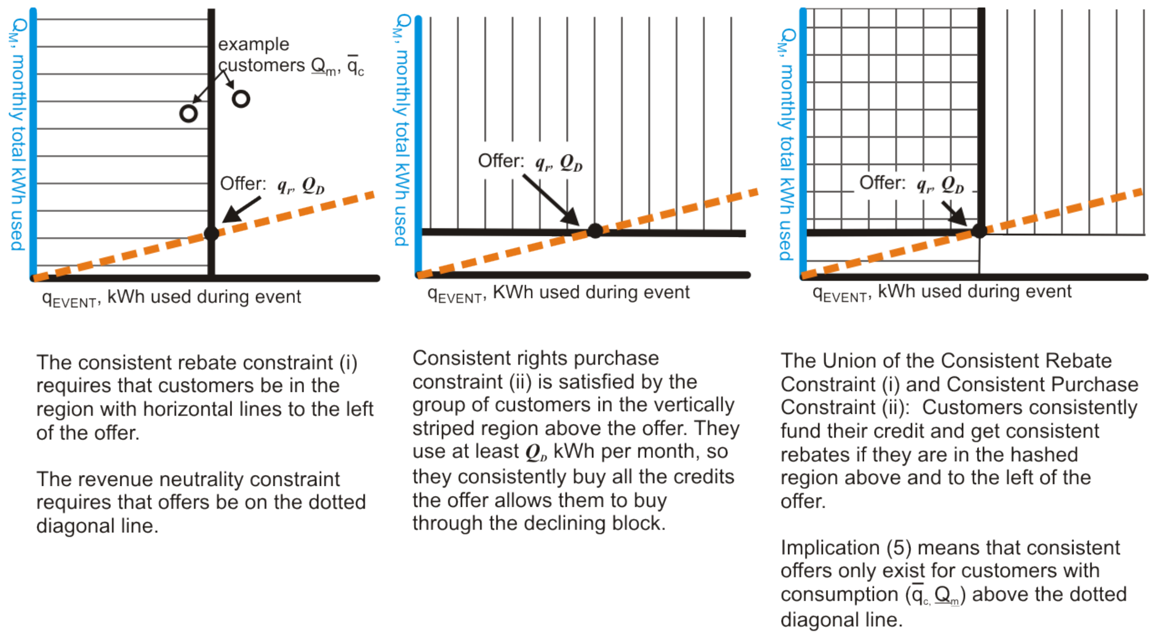

Figure 3.4: *Visualizing constraints 1-3.*

i. Customer-specific rights levels, $\mathbf{q_R}^i$ exist that satisfy feasibility condition 3.5 for most customers.

ii. The organizations making offers have enough data to predict a $\hat{\mathbf{q}}^i_{\mathbf{R}}$ that satisfies feasibility condition 3.5 for customer i.

Answering these questions requires data about the behavior of real customers on CPP-IPR or comparable dynamic pricing. CPP-IPR is yet to be tested on real consumers, but there are data on customers on economically-similar CPP rates.

3.8.2 The California Statewide Pricing Pilot offers important evidence

California's Statewide Pricing Pilot (SPP) exposed about 500 customers to 27 CPP events (Charles River Associates, c, 20) from summer 2003 into fall 2004 while collecting survey data and recording hourly electricity use.[50] This created a 15 month panel of data. The SPP data are a powerful source of evidence about electricity use patterns. SPP data

[50]The SPP was a vast field experiment. The data considered here are from its largest cell, customers on a CPP-Fixed Period ("CPP-F") rate who experienced events that ran from 2-7PM, who were notified of events by telephone the day before, and who did not get thermostats that could respond to price signals automatically.

are particularly relevant to the example CPP-IPR rate in table 3.2 since that rate is adapted from an SPP Welcome kit (Pacific Gas & Electric, a).

3.8.3 Good IP rebate offers exist for most SPP customers

Figure 3.5 plots the constraints in the style of figure 3.4 with real data. Evidence from the SPP suggests that 97% of customers statewide have demand patterns that satisfy feasibility criterion 3.6 for the example rate.

The rectangular region above the diagonal line is the single offer that provides consistent rebates to the largest number of customers. The graph shows that one size does not fit all. The single-optimal offer is not consistent for customers outside of the rectangle. It is not very surprising that the monthly usage of a small apartment in a temperate climate is insufficient to pay for the level of rights required to provide consistent rebates for a big house in the desert.

We can generalize this analysis to a family of CPP-IPR rates by rearranging feasibility constraint 3.6 as a relationship between characteristics of customers and characteristics of rates. This rearrangement yields:

$$\frac{12\mathcal{M}}{N_c(P_c - P_h)} \geq \frac{\bar{\mathsf{q}}_\mathsf{c}}{\underline{\mathsf{Q}_\mathsf{m}}} \tag{3.7}$$

The left side of this equation describes characteristics of the rate, while the right side describes characteristics of the customers. The left side is the ratio of the rate's ability to raise money to the cost of providing each kWh of rights during each event. The right hand side is the ratio of the the number of rights required to offer the customer consistent rebates to the biggest declining block size that they consistently purchase. Figure 3.6 shows the cumulative distribution of the right hand side of the rearranged criterion, 3.7, $\frac{\bar{\mathsf{q}}_\mathsf{c}}{\underline{\mathsf{Q}_\mathsf{m}}}$ and uses it to see the percentage of customers who could get consistent offers under the IP rebates that could be added to a variety of real CPP offers. It suggests that IP rebates work well with three-period rates. IP rebates struggle with a two-period rate proposed by Pepco for Washington DC customers for the reasons outlined in section 3.6.1. The figure also shows that mindlessly implementing this IP rebate approach struggles with Ameren's four period rate and is less than ideal for Gulf Power's four period rate, because they both split the low priced period into a low priced rate and an intermediate, shoulder rate that is quite close (.9 cents in Gulf Power; 0.14 cents for Ameren) to the time invariant rate. Markups

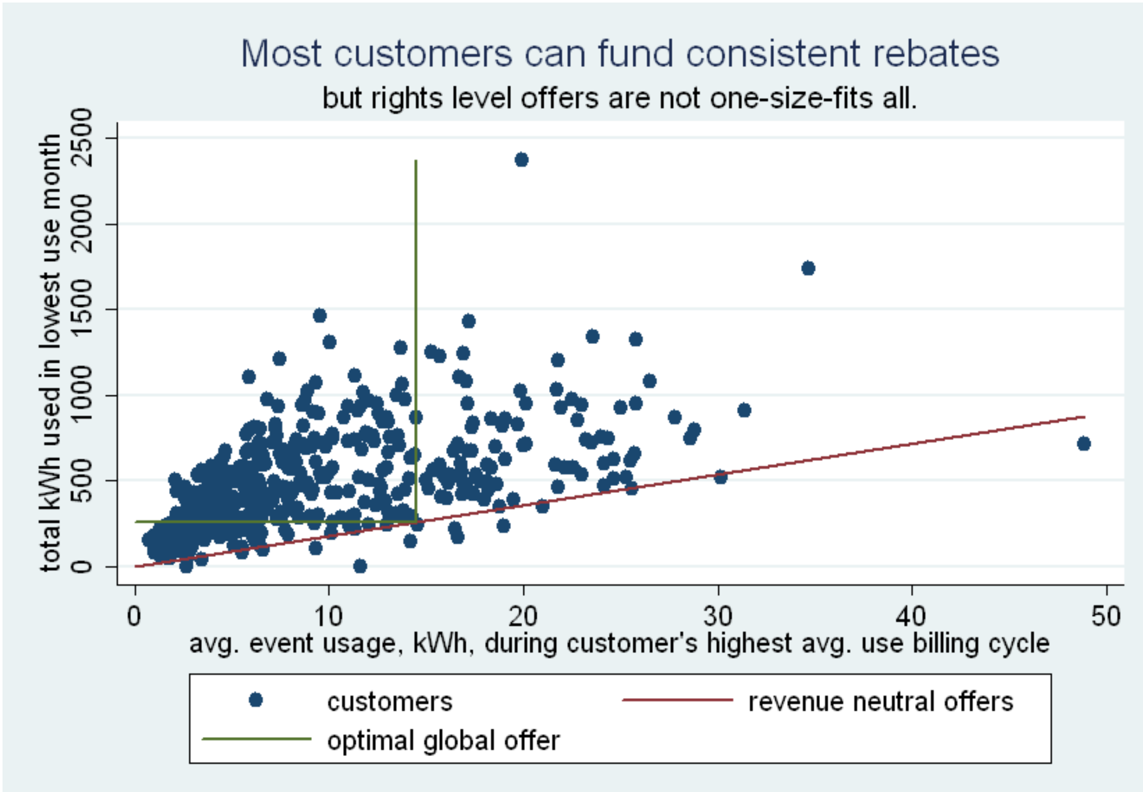

Figure 3.5: *Most California's SPP customers' demand patterns are above the diagonal line defined by feasibility criterion 3.6, so consistent offers exist for them. But the single offer that provides consistent rebates and rights purchases for the greatest number of customers does not perform particularly well, so we should consider more customized offers.*

that keep the shoulder rate less expensive than the time invariant rate often generate no consistentoffers. Either imposing a larger markup (a four cent markup would keep prices lower 64% of the time under Ameren's rate) or sacrificing some economic efficiency by reducing the price during the shoulder period (perhaps by adding more low priced hours to it) could address these problems.

3.8.4 It is easy to predict consistent offers given readily available information

The rate implementers need to be able to identify consistent offers for each customer but will often not have data about how much power the customer used during hot weekday afternoons. This usage level determines the level of rights the customer needs to

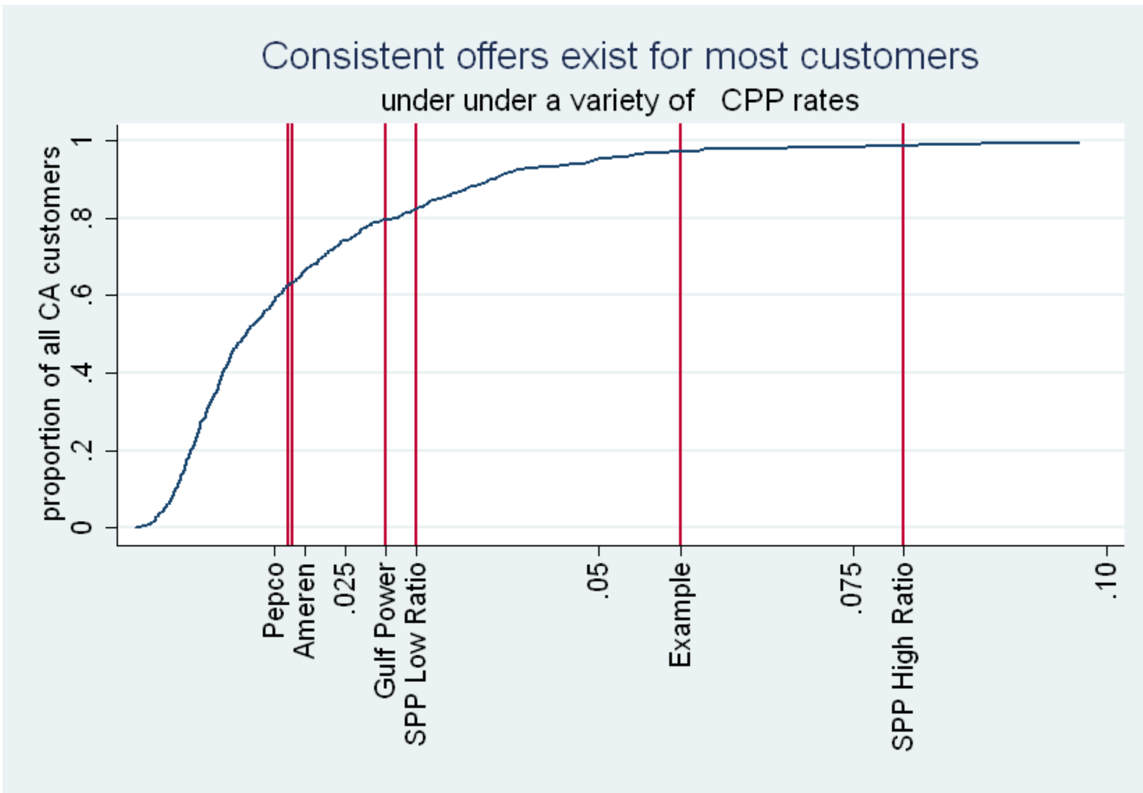

Figure 3.6: *Offer feasibility under a variety of rates: Rearranging criterion 3.6 to* $\frac{12\mathcal{M}}{N_c(P_c - P_h)} \geq \frac{\bar{\mathsf{q}}_\mathsf{c}}{\mathsf{Q}_\mathsf{m}}$ *lets us compute the percentage of California CPP customers for whom consistent offers exist for real CPP-rates. This approximation assumes negligible demand elasticity. Table 3.4 describes the rates pictured here.*

Source of Rate (State)	# rate pds.	crit. hrs $5N_c$	uniform price P_u	peak price P_L	off peak price P_H	critical price P_c	$\frac{12(P_u-P_L)}{N_c(P_c-P_h)}$
Pepco (DC)	2	60	7.92	6.81	6.81	63.98	.019
Ameren (MO)	4	32	7.64	7.5, 4.8	16.75	30.0	.020
Gulf Power (FL)	4	87.6	8	7.1, 5.9	11.7	32.6	.029
SPP Low Ratio (CA)	3	75	P_U	$P_U - 1.2$	$P_U + 9.8$	$P_U + 41.8$	.030
Example, Table 3.2	3	75	14.6	12	24	60	.058
SPP High Ratio (CA)	3	75	P_U	$P_U - 5.09$	$P_U + 11.64$	$P_U + 60.91$	.083

Table 3.4: The table describes the CPP rates plotted in figure 3.6. They have 2-4 rate periods. Two-period rates have just normal and critical periods. The two period rate presented here is designed to generate identical average bills to the time-invariant rate and its IP rebate modification looks like the one proposed in section 3.6 and suffers the shortcomings of taking approach to a two period rate that are described in section 3.6.1. Three period rates have offpeak, peak, and critical rates. Four period rates further subdivide the offpeak period into a low rate and a "shoulder" rate that contains hours during the transition between the peak and offpeak periods. The four-period rates struggle to fund rights because the shoulder rate is quite close to the uniform price. Adjusting the rate to expand the shoulder period and reduce its average price or changing the IP rebate markup structure would improve their performance. For example, basing a markup on Ameren's lowest price of 4.8 cents per kWh would yield prices that are lower during 64% rather than 90% of all hours, but would yield a $\frac{12(P_u-P_L)}{N_c(P_c-P_h)}$ of .39 – which performs so well as to be off this chart. The SPP called 5 hour events, so N_C counts events assuming that they lasts 5 hours. I calculate the number of "5-hour events" that other rates call by dividing number of hours of events that they call by 5. This table reports the summer prices of seasonally varying rates. (Sources: Wilson (2006); Pepco; Voytas (2006); Ameren; Gulf Power; Pacific Gas & Electric (c); San Diego Gas & Electric)

get consistent rebates. Many utilities – including Gulf Power – will lack this information because they only install "interval" electricity meters that provide disaggregated usage data when customers sign up for dynamic pricing. California's three major utilities plan to install interval meters for everyone, but may want to offer dynamic pricing as soon as they install the meters. Either of these scenarios would require the firm to make an initial offer to CPP-IPR customers based on the data they already have, like monthly usage data from old meters that provide only aggregate data, account type, and geographic data. The analysis below shows that the data utilities have predict consistent offers quite well, so it appears to be feasible to implement CPP-IPR.

In order to estimate the desirable offers using a conventional approach, we need to identify an optimal offer from the set of consistent offers. Most customers' use patterns mean that criterion 3.5 defines a range of consistent offers between the smallest $\mathbf{q_R}$ that provides consistent rebates and the largest $\mathbf{Q_D}$ that the customer can buy each month. For the purposes of this analysis, I selected the consistent offer, $(\mathbf{q_R}^*, \mathbf{Q_D}^*)$, that satisfies criterion 3.5 in a way that is robust to the largest number of dollar deviations in total ability to buy rights, $12\underline{\mathsf{Q}_\mathsf{m}}$, and the needs for rights, $N_c\bar{\mathsf{q}}_\mathsf{c}(P_c - P_h)$.[51]

This analysis proceeded in three steps:

i. I constructed an optimal offer $\mathbf{q_{R,'04}}^{i*}$ for each customer. It specified the set of offers that would be consistent for that customer-year, typically the year October '03-September '04.[52]

ii. I ran the following OLS regression: $\mathbf{q_{R,'04}}^{i*} = \alpha + \beta_1 * useSummer02 + \beta_2 * ClimateZone + \beta_3 * apartment + \epsilon$ where *useSummer02* is the customer's average kWh per day during three summer months the year before the experiment began, *apartment* is 1 if the account is in a multifamily building and zero if the account is a single family home, and *ClimateZone* is a set of dummies indicating whether the account is located in each of four mutually exclusive climate zones. Fog-belt zone 1 largely near San Francisco (the

[51] The optimal offer should minimize the likelihood that random variation would prevent the offer from providing consistent rebates or purchases. This requires knowing the within-customer standard deviations of rights needs, $\bar{\mathsf{q}}_\mathsf{c}$, and of the ability to purchase rights, $\underline{\mathsf{Q}_\mathsf{m}}$. The SPP data only tracks 27 events over 15 months, so there are too few years and too little variation in exogenous factors like weather, economic conditions, appliance upgrades, and family configuration changes to calculate meaningful standard deviations.

[52] I focus on the last 12 months of the 15 month sample where possible because the experiment enrolled customers gradually, but ended abruptly, so looking at the initial 12 months would yield different date ranges for different subjects. This would make the results harder to understand, especially because usage is heavily weather driven.

omitted category) is the coolest. The zones get progressively hotter and culminate in desert zone 4. Table 3.5 shows that regressing the optimal offer calculated from a 12 month period in 2003-04 ($\mathbf{q_R}^*, \mathbf{Q_D}^*$) on total summer usage in 2002 explains 76% percent of the variation and that adding readily available variables about the climate and whether the account is at a single or multifamily building improves the fit to explain 78% percent of the variation.

iii. I used the results of that regression to predict a consistent offer, $\hat{\mathbf{q}}^{i*}_{\mathbf{R},'04}$, for each customer and determined whether it was consistent in the sense of satisfying criterion 3.5 for the values of ($\underline{\mathsf{Q}_\mathsf{m}}$, $\bar{\mathsf{q}}_\mathsf{c}$) for that year. The full regression predicts consistent offers that satisfy criterion 3.5 for 80% of all customers for whom a consistent offer exists. When $\hat{\mathbf{q}}^{i*}_{\mathbf{R},'04}$ was not a consistent offer, it was typically substantively fairly close to being a consistent offer. Half were less than 2.4 kWh of rights away from the nearest consistent offer. That size of deviation customers would force customers with too few rights to buy high-priced power costing no more than $1.44 per event.

iv. I tested a model from one summer's ability to predict appropriate offers for another summers which real rate designers need to be able to do. Most California scarcity events take place in the summer months of July through September, and the SPP called 21 of its 27 events during those months. The 15-month experiment contained the important part of 2 years. This allows us to make some preliminary investigations of how well parameters developed from one year predict for a different year. I went out of sample to check whether $\hat{\mathbf{q}}^{i*}_{\mathbf{R},'04}$ calculated from the year containing Summer 2004 (namely October 2003-September 2004) was a consistent offer that satisfied criterion 3.5 using the customers' consumption patterns ($\underline{\mathsf{Q}_{\mathsf{m},03}}$, $\bar{\mathsf{q}}_{\mathsf{c},03}$) for the year including Summer 2003, namely July 2003-June 2004. The out of sample universe contained 61% of customers. These customers had to be in the sample for two summers, and had to have ($\underline{\mathsf{Q}_{\mathsf{m},03}}$, $\bar{\mathsf{q}}_{\mathsf{c},03}$) that was different from ($\underline{\mathsf{Q}_{\mathsf{m},04}}$, $\bar{\mathsf{q}}_{\mathsf{c},04}$). This implies that either or both their their highest event use that set $\bar{\mathsf{q}}_\mathsf{c}$or their minimum consumption that set $\underline{\mathsf{Q}_\mathsf{m}}$had to occur between July and September.[53] The out of sample prediction

[53]The SPP CPP treatment started in July 2003 and ran through September 2004. California's electricity demand (and scarcity) peaks during the summer, so we observe two separate summers but not two separate years. A significant proportion of Zone 1, fog-belt customers used more during the winter events than during any summer events. These customers set their their rights needs, $\bar{\mathsf{q}}_\mathsf{c}$, during the winter and thus get dropped from the out-of-sample analysis which used summer 2004 data to predict summer 2003 needs.

Table 3.5: Using an OLS regression to predict the optimal IP rebate offer in kWh per event works well. Standard Errors in parentheses.

	usage only model	usage, climate, account type model
avg. daily use Summer 2002, kWh	.78*** (.028)	.80*** (.031)
climate zone 2		.84 (.621)
climate zone 3		-.45 (.668)
climate zone 4		-2.95*** (.869)
apartment		-1.47** (.466)
intercept	2.62*** (.452)	2.91*** (.644)
N	482	482
R^2	0.764	0.781

of $\hat{\mathbf{q}}^{i*}_{\mathbf{R},'04}$ was a consistent offer that satisfied criterion 3.5 using $(\underline{\mathsf{Q}}_{\mathsf{m},03}, \bar{\mathsf{q}}_{\mathsf{c},03})$ for 82% of the out of sample universe.

3.9 Implementation concerns

3.9.1 IP rebates are a rate feature, not a whole rate, and leave significant flexibility to rate designers

IP rebates are a revenue neutral feature that can be added to any CPP rate without affecting its marginal incentives. Implementing IP rebates as a flexible feature that coexists with a wide variety of rates preserves CPP rate designers' freedom to meet local needs and their ability to choose rates given limited information. Real CPP rates reflect compromises between pricing near marginal cost, meeting revenue requirements, maintaining simplicity, making incremental changes to the status quo, and treating rate payers equitably. CPP rate designers choose a small number of rate periods and prices for each. These parameters have reasonably transparent implications – unlike, for example, the choice of baseline-rebate parameters. Further, CPP rate designers generally set prices without knowing short term customer demand elasticities that Ramsey pricing would require. Ramsey pricing is most economically efficient approach to meeting a utility's revenue requirement by marking up the products it sells (in the case of CPP they would be offpeak, peak, and critical period power and perhaps connection the electricity system) in a way that minimizes deadweight loss (Hausker, 1986). The designers have to choose without knowing the marginal cost of power in each period. Market power and policies that control market power and prevent shortages make spot market electricity prices diverge from the marginal cost of power.

While IP rebates can be added to any underlying CPP rate, the number of customers who get consistent offers is sensitive to the size of the difference between offpeak and time-invariant prices because that difference is the upper bound on the markup, $\mathcal{M} \leq P_u - P_L$. Often the markup is a tenth of a cent less than the difference, $\mathcal{M} = P_u - P_L - .001$. Larger markups mean that each kWh that the declining block marks up provides more rights, $\mathbf{q_R}$. Thus, increasing the markup $\mathcal{M}$ expands the set of consistent offers by relaxing criterion 3.7 (a form of criterion 3.6) which should allow improvements in the percentage of customers' getting consistent offers.

3.9.2 There are good policy reasons to divide customers into coarse subsets

An IP rebate implementation can either make each customer a customized offer or categorize customers and make an offer to each category. The quantitative analysis below shows that making offers to broad categories of customers defined by use and geography can make consistent offers to a large percentage of customers. Making offers to categories of customers has compelling practical advantages over customizing offers because category-level offers are easy to understand, seem fair, and discourage distortion.

A small number of offers and clear rules are an advantage for analysts, regulators, utilities, and customers.

It is advantageous for a system to be easy for customers, utility staff, advocates, and regulators to understand.

- **Policy makers:** It is easier for regulators and advocates to understand, discuss, and tune a small menu of offers. It is easy to understand and adjust IP rebates' seasonal bill impacts if a large group of customers makes identical monthly contributions and then get identical credits during each event. Categorical offers allow conversations about the precise, category-wide impact rather than about average impacts. A rate will be easiest for regulators to work with if it makes offers to existing categories that the regulators are already familiar with and used to treating as a unit.

- **Customers:** If customers understand why they got their offer and that their neighbors got the same offer, they may be less likely to call their utilities with questions.

- **Customer service:** A simple system will make it easier to train call center staff, reduce the number of questions about offers that the call center receives, and make those questions quicker to answer. Broad categories may simplify the challenge of assigning hedges to new tenants or to new buildings and of explaining this initial decision to the customers.

Customers need to perceive the offers as fair

Every IP rebate offer gives a customer charges and credits that sum to zero over the course of the year, and it is difficult to consider differences in offers unfair while focusing on the zero annual effect bottom line. However, some customers will not know this and may see differences in rebate eligibility as unfair. CPP-IPR should be designed to work well even if some customers understand only that it is advantageous to reduce use of pricey weekday afternoon power and more advantageous to reduce power use during critical events to earn rebates. One step toward this goal is to maximize the number of customers who perceive the program as fair based on superficial knowledge of their own hedges and those of their neighbors and friends. Consumers' lack of knowledge about the program and whether their neighbors use electricity in a similar way makes it harder to maintain the perception of fairness.

There are at least three components of perceived fairness in rate offers: offer equity, process, and justice.

- **Offer equity** requires that (superficially) similar customers get similar rebate opportunities. Customers' electricity consumption patterns determine their need for rights. These patterns – and the equipment efficiency and habits that drive them – are often invisible to neighbors. Thus, assigning rights by consumption patterns may be objectively equitable, but appear inequitable to customers comparing bills over the back fence.[54] By contrast, assigning the same rights level to customers who live in similar

[54] Multiple, mutually exclusive notions of fairness come into play on most policy issues. Stone (1997)[39-41]

buildings in the same geographic area may appear significantly fairer.

- **Process fairness** requires the use of transparent, objective category assignment rules. The policy should articulate simple criteria that explain why two customers received different levels of rights.

- **Justice** requires the rebate program minimize real and perceived opportunities to profit through strategic efforts to exploit the program's rules.

Avoiding consumption distortions by confused customers

One of the central design features of incentive-preserving rebates is that they do not create incentives for rational, well-informed customers to distort their buying patterns to profit by getting more rebates. It takes considerable analysis to convert a rate schedule like table 3.2 and the fine print that would accompany it to uncover this incentive compatibility. A significant literature reports that lab subjects do not respond as intended to incentive compatible mechanisms in part because the incentive compatibility is often not obvious (See Chen (Forthcoming) for a review).

It is important to avoid using mechanisms that set offers in ways that systematically induce customers to believe wrongly that they can benefit by manipulating their demand to get more rights, $\mathbf{q_R}$, or reducing the number of marked up units, $\mathbf{Q_D}$. Using relatively immutable characteristics to set hedge levels can help achieve this goal. Utilities know which climate zone each account is in and whether it is an apartment. Customers can change these characteristics by moving, but even confused customers are unlikely to think that the misunderstood incentives justify the cost of moving unless they were already on the cusp of relocating. Using coarse total annual consumption bins to define categories may be compelling because they predict well and are readily accessible to utilities and because most customers have to engage in a prolonged, costly change in consumption to switch consumption bins.

There are strong practical reasons to assign each customer to one of a small number of categories and to make one offer per category. There is tension between using simple rules based on immutable characteristics to categorize customers in a way that seems fair and that minimizes distortion and the need to match customers with the right hedge-level.

performs a thought experiment about how to equitably decide who can eat a cake and comes up with eight mutually exclusive notions of fairness.

The balance of this section explores whether we can reach an adequate compromise between the goals of categorizing customers and of ensuring an adequate fit.

3.9.3 Making offers using existing categories worked well

Employing categories that rate designers already use would facilitate the implementation of CPP-IPR. This section tests the feasibility of that approach by calculating the optimal offers for each of the SPP's categories of customers. The SPP divided the state into 4 climate zones and each climate zone into three groups: apartments, high use single family houses, and low use single family houses. It classified customers' use levels as high or low using consumption from the summer before the experiment began. The ceiling on the low use category is 16 kWh per day in the coolest climate zones and rises with progressively hotter climate zones to 28 kWh per day in the hottest climates.

The analysis shows that the SPP's raw categories were not optimal because the customers in the low use and apartment categories in the hottest climate zones were too diverse for a single IP rebate offer to fit well. While the raw categories from the SPP performed poorly, a set of categories that preserved most of the SPP's distinctions and further subdivided customers at each category's median use level allowed us to make consistent offers to the vast majority of customers.

The modified set of 16 categories took the SPP's raw categories, discarded the distinction between apartments and single family homes, and categorized all customers using the SPP's high and low use categories.[55] The modified categories subdivided the SPP's high and low use categories at each category's median usage level. This yielded very low use, low use, high use, and very high use categories in each of the four climate zones for a total of 16 groups. It discarded the apartment category because the sample only contains about 500 customers. Retaining and subdividing the apartment category yielded unacceptably small cells. Further, apartment status explains far less variation than does total use.

Calculating the offer that is consistent for the greatest number of people in each group yields a set of offers listed in table 3.6 and visualized in figure 3.8. These are consistent for 86% of all customers statewide regardless of whether a feasible offer exists for that customer. This approach sometimes outperformed the regression approach in part because

[55] The vast majority of apartments were low use.

it used data on each customer's whole range of consistent offers rather than a single point representation of that range. This number reflects a 92% consistent offer rate in the more temperate climates zones 1 and 2 and a 77% rate in the hottest climate zones 3 and 4. Five percent of the customers in zones 3 and 4 have no feasible consistent offers.

The algorithm to determine the optimal set of offers for each category proceeded as follows:

- It calculated the range of offers, $[\mathbf{q_{R,i}}^m in, \mathbf{q_{R,i}}^m ax]$, that satisfies feasibility criterion 3.5 for each customer, i. It ranges from the smallest offer that provides consistent rebates, $\mathbf{q_{R,i}}^m in = \bar{\mathsf{q}}_\mathsf{c}$, to the largest offer that the customer can consistently buy, $\mathbf{q_{R,i}}^m in = \frac{12\mathcal{M}\underline{\mathsf{Q}_\mathsf{m}}}{N_c(P_c - P_h)}$.
- It used each customer's optimal range $[\mathbf{q_{R,i}}^m in, \mathbf{q_{R,i}}^m ax]$ to calculate the proportion of all customers in each group who would get a consistent offer for each value of $\mathbf{q_R}$.[56]

  This yields an objective step function like that pictured in figure 3.7. Table 3.6 summarizes the 16 optimal offers and their performance in providing consistent offers, while figure 3.8 displays identical information about the optimal offers but not about their performance. The balance of this paper will use this 16-offer CPP-IPR rate as a benchmark in calculating the impacts of CPP-IPR.[57] The 16 analysis cells contain both apartment and single family customers. The analysis used weights to make each cell representative of the portion of the statewide population of accounts with its usage and climate zone characteristics.

We can get from 16 to 9 categories without making any compromises by merging groups with similar rights needs. Most often, these involve combining usage categories from the same region or combining the same usage level in neighboring regions.[58]

[56] A future revision will use bootstrap resampling of the population to put confidence intervals on the optimal offer estimates and the estimates of the offer's performance.

[57] These results and the results below modify the example rate in table 3.2 by moving from the 15 events per year that the CPP promised to the 18 events that it in fact called during the 12 months from October '03 through September '04. It reallocates the fixed credits by offering 15/18 of the **R**value per event – reducing the critical price from 60 cents to 54 cents. This is a very small implicit CPP rate reduction

[58] Specifically, we can make offers from the following offer ranges to each of the following sets of categories {3VH, 4VH: offer 31.42-31.46 kWh of rights per event}, {3H, 4H, 1VH, 2VH: 25.83-26.00}, {1L, 2L, 1H: 10.15 - 12.79 }, and {3VL, 4VL: 6.13-6.36}. The four other groups would get their own offers, as listed in 3.6.

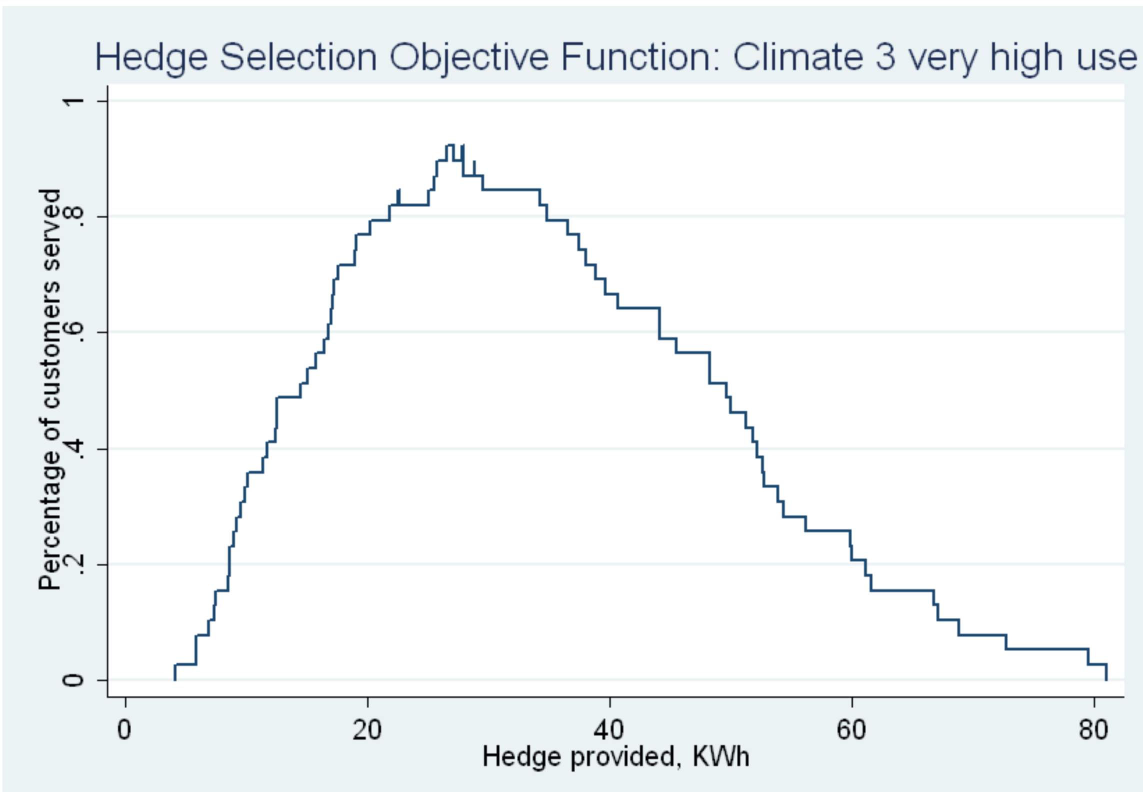

Figure 3.7: *The proportion of very high use customers in climate zone 3 for whom each possible offer is consistent.*

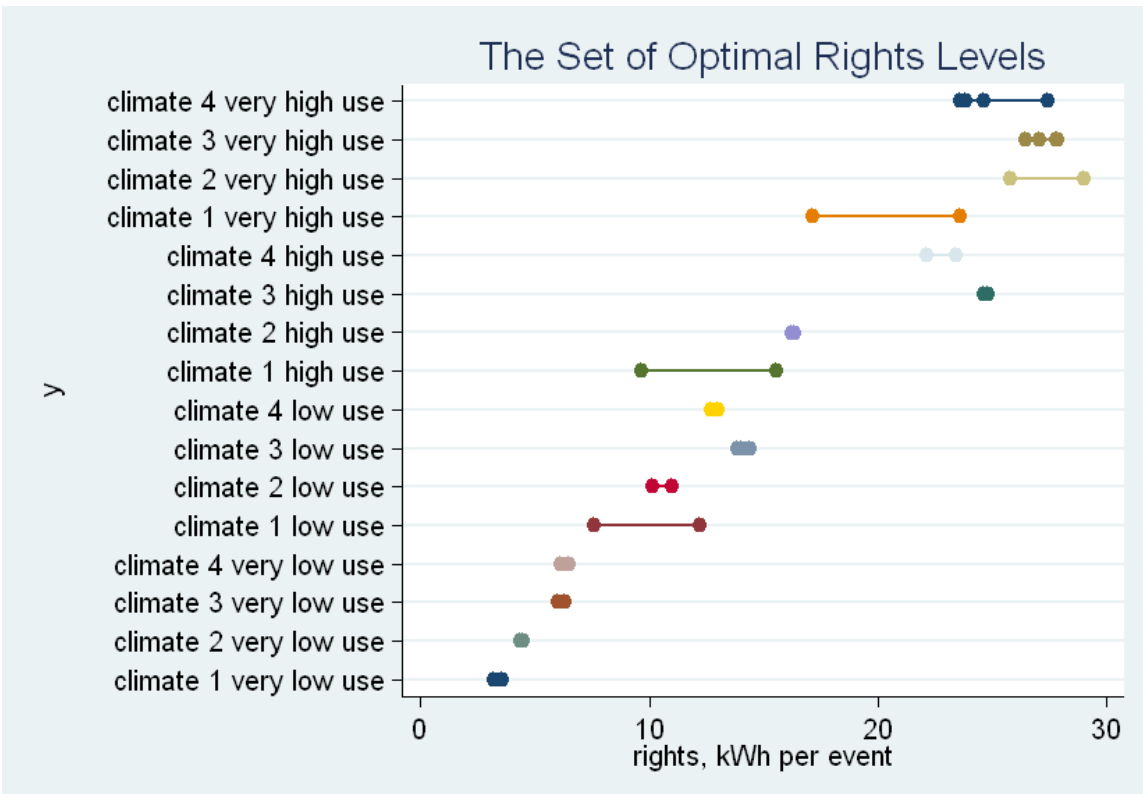

Figure 3.8: *Optimal Offers by Group: consistent offers for higher use customers and customers in hotter climates provide more rights.*

Table 3.6: Optimal group offers: the range of rights values that provides consistent offers to the greatest number of customers. This reports the percentage of customers for whom consistent offers exist who get them.

	proportion getting	optimal range:	
climate zone, use type	consistent offers	optimal range: lower end	upper end
1 very low use	.93	3.19	3.28
2 very low use	.87	4.97	5.17
3 very low use	.84	6.00	6.36
4 very low use	.87	6.13	7.55
1 low use	1.00	7.61	14.21
2 low use	1.00	10.15	12.79
4 low use	.81	14.30	15.13
3 low use	.84	17.98	19.42
1 high use	1.00	9.65	18.04
2 high use	1.00	18.38	18.98
3 high use	.85	24.15	26.00
4 high use	.96	23.49	27.18
1 very high use	1.00	17.13	27.43
2 very high use	1.00	25.83	28.30
3 very high use	.97	28.82	31.46
4 very high use	.88	31.42	31.82

3.9.4 Offers that are not consistent are typically fairly close to being consistent

Offers that are not consistent are generally pretty close to being consistent and expose customers to only a few dollars per year of either exposure to critical pricing or of reduced rebates. Specifically:

- The majority of the 9.3% of customers who did not get consistent rebates paid the high marginal price in just one month. The weighted median (mean) customer who did not get consistent rebates paid for 4.67 (8.23) kWh at the full critical price, which cost \$2.80 (\$4.94).

- The rate marked up more power in at least one month than customers bought for 3.9% of all customers. The mean amount of rights that these customers were supposed to buy, but did not was a total of \$3.65 per year.

- To put this in perspective, this sample of customers spent a weighted average of \$898.71 on power over the course of the year under the example CPP or CPP-IPR rates.

- Customers in hot climates 3 and 4 were roughly twice as likely to have too few rights or to contribute too little as customers in more temperate climates 1 and 2.

3.9.5 Robustness of these offers to changes in weather, economic conditions, and customer characteristics

Good rights offers need to work not only for the summer and customer-base that they were designed for, but also for summers that have differing weather and economic conditions and for an unexpected subset of the customers. A thorough exploration of these issues merits a paper in its own right, but the results from some simple tests suggest that the offers are reasonably robust. One promising way to understand the robustness of the offers is to look for evidence about the engineering and social limits on power consumption. If a customer would get consistent rebates despite running their air conditioner flat out and turning on another major appliance like an oven or dryer, their offer is quite robust. And if the customer is either never home to activate the other major appliance or is paying enough attention to not do so during a critical event, then their offer also seems to be robust.

- 72% of all customers would get consistent rebates even if every event matched their highest use event over the 15 month study. The other customers who got consistent rebates did so by averaging an extreme event with lower-use events over the course of a month.

- 47% of all customers would get consistent rebates even if they equaled their maximum use weekday afternoon over the 15 month study. These customers appear to have engineering or social limits that are likely to prevent them from using more power than they have rights to.

- The median customer uses only 49.1% of their rights to get consistent rebates and gets a declining block that marks up only 57.1% of the power that they use in their lowest use month. Similarly the 75th percentile customer has only 71.3% of their power marked up in their lowest use month and needs only 71.4% of their hedge to get consistent rebates. Most customers get significant cushions that make their IP rebate offers fairly robust to variations in conditions. Once we reach the 90th percentile, however, the cushions largely disappear and customers have 93.7% of their use marked up in their lowest use month and need 99.0% of their hedges to get consistent rebates. It's likely that some of the marginal customers in the SPP experiment who needed the greatest rights levels joined the experiment to contribute to knopwledge and earn $175 but were not responding to price signals and would not opt in CPP-IPR.

3.10 IP rebates smooth seasonal bill variations in regions where peak demand coincides with the system's peak demand

Consumers and policy makers both express a preference for bill levels that are consistent from month to month. Many utilities offer balanced payment plans that send customers bills of a constant size 11 months a year and then adjust for differences between the preset payments and the actual charges in the last billing period of the year. Further, some utilities have sold "flat bill" plans to a significant number of customers. These charge the customer a flat fee that reflects their expected bill plus a risk premium on the order of 10% regardless of the customer's usage.

IP rebates reduce CPP bills during critical periods and increase CPP bills during the first $\mathbf{Q_D}$ hours each month. This shift of bills among hours also drives a bill shift among months since the monthly contributions are spread evenly around the year, while most critical periods take place during the summer months. Thus, IP rebates reduce seasonal variations in bills for customers whose total use peaks during the season with the largest number of critical events. IP rebates can amplify seasonal differences in regions where electricity use peaks in a different season from the majority of the electricity use in their system.

Figure 3.9 shows how these possibilities play out in California's most temperate and hottest climate zones. The top three lines on the graph show that CPP-IPR smooths the air-conditioning driven, summer bill peak in the desert (zone 4). CPP-IPR amplifies the modest bill peak in the zone 1's temperate climate, where electricity demand peaks during the winter in a summer peaking system.

The two intermediate zones have seasonal patterns between these two extremes. Bills in Central Valley, climate zone 3, peak during the summer. CPP-IPR dampens this peak, much as it does in climate zone 4. Climate zone 2 has the largest population of any of

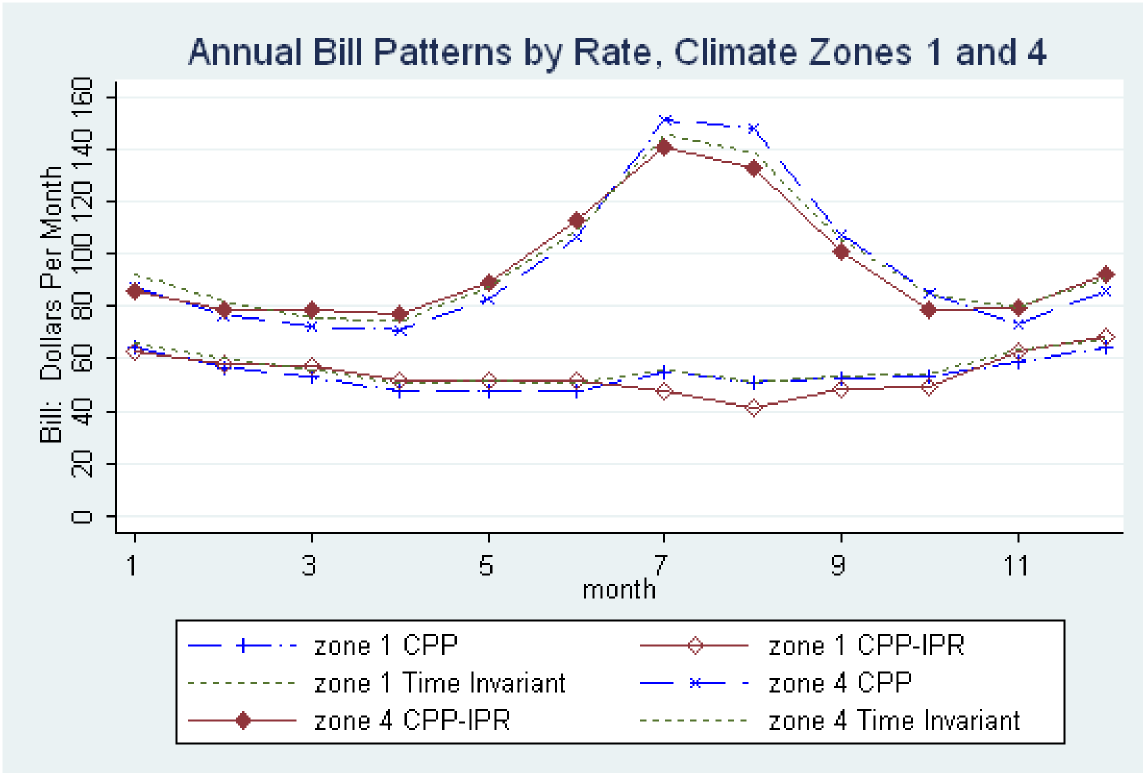

Figure 3.9: *IP rebate rates (diamond) smooth a high summer peak in the climate zone 4 (desert) at the top of this graph but exacerbate modestly winter peaking bills in climate zone 1 (the temperate fog belt; largely the San Francisco Bay Area) at the bottom of the graph.*

the four zones and includes much of the Los Angeles and San Diego areas, and some inland parts of the San Francisco area. It has modest summer and winter bill peaks that increase average bills from about \$70 in the fall and spring to about \$80 in the winter and summer under both CPP and under time invariant rates. The example CPP-IPR rate eliminates the summer peak, leaving the average customer with a modest winter peak.

Looking at customer-level bill volatility – as Borenstein (2007) did – yields qualitatively similar results: IP rebates reduce each customer's month-to-month bill volatility relative to CPP in climates that hit their peak consumption season when the system does and increase each customer's month-to-month bill volatility in regions where residential use peaks in a different season than the statewide system does.

3.11 Conclusion

Efficient incentives can be an important part of improving public policies. But the obvious, natural implementations of efficient policies often repel customers who use flawed behavioral decision making heuristics.

The evidence from this project and from a line of behavioral field experiments suggests that behavioral insights can provide insights about the source of behavior that serves neither individuals nor society well and suggest levers for interventions that address this behavior. Using economics and psychology to guide the design of improved incentives can often yield implementations that preserve the important economic properties of a policy while helping some consumers make significantly better choices.

Incentive Preserving Rebates are an example of an intervention that preserves incentives and revenues while changing the presentation of incentives to address a significant set of psychological and implementation concerns. The behavioral considerations drive five constraints, while implementation challenges add requirements like using a small number of existing categories.

The evidence suggests that IP rebates are administratively feasible and that coarse, categorical offers can meet the needs of most customers. Most customers will fully fund rights that deliver them consistent rebates. Those who do not fully fund their rights or who do not get get consistent rebates experience deviations in the form of payments at the critical price or reductions in rights from the promised level that are typically less than 1% of their total annual bill.

IP rebates smooth bills and reduce volatility relative to conventional CPP. IP rebates are somewhat sensitive to customer diversity, but not nearly as sensitive as baseline rebate rates are.

Changing the framing of prices is a well recognized tool in marketing. This application to public policy may be the first of many important potential applications to help people make better choices in areas like the purchase of energy efficient appliances and vehicles or choices between owning a car and using public transportation.

Chapter 4

Optimal Deployment of a simple menu of Incentive Preserving Rebates for CPP Rates with Heterogeneous Customers

4.1 Overview and Background

This project articulates criteria for an optimal implementation of Incentive Preserving (IP) Rebates under realistic constraints. The paper's analysis uses electricity consumption profiles from the California Statewide Pricing Pilot (SPP) to assess the performance of simple, realistic implementations of an IP rebate rate.

Chapter 3proposes using Incentive Preserving (IP) Rebates to present critical peak pricing in a way that is more attractive to consumers. IP rebates work by selling each customer rights to either a rebate or a fixed quantity of power at the usual price during a critical peak event. Customers buy these rights by paying a markup on the first few units of power that they buy each month. An IP rebate implementation has to choose an **offer** specifying

i. the number of kWh of rights, $\mathbf{q_R}$, that the customer gets per event and

ii. the number of kWh per month, $\mathbf{Q_D}$, that the rate marks up to pay for these rights.

Good rates choose offers ($\mathbf{q_R}$, $\mathbf{Q_D}$) so that as many customers as possible:

i. get "consistent rebates", because they purchase at least as many rights, $\mathbf{q_R}$, as they use in the average event in the month with the highest average event use,

ii. make consistent rights purchases because they consume at least $\mathbf{Q_D}$, the amount of power that their offer bundled with rights, each month.

Offers that meet both constraints are "consistent". Design constraints and budget balance requirements fix a ratio between the quantity of rights that each customer gets during each event and the number of units the rate marks up to pay for them (Chapter 3).[1] Thus, there is a tradeoff between providing enough rights ($\mathbf{q_R}$) so that larger customers get consistent rebates and marking up few enough kWh per month ($\mathbf{Q_D}$) that smaller and more seasonal customers pay for all their rights each month.

Customers are unlikely to understand all the details of IP rebate pricing. Customers with an incomplete understanding often know enough to respond rationally in the way the IP rebate program designers desired. Instead, many will understand only that their CPP-IPR rate typically offers rebates as incentives to save during critical period opportunities and a predictable schedule of prices and rebate opportunities except in extraordinary situations. Consistent offers fit this understanding. They create incentives to conserve during critical periods through rebates, not surcharges, and they do so without activating budget balancing surcharge or rights reduction provisions from the 'fine print.' Inconsistent offers which fail to meet these expectations may confuse and annoy customers, but will have no effect on a customer's total annual bills, the utility's revenue from that customer, or the program's marginal incentives unless they spur an unexpected change in customer behavior.[2] Thus, this paper addresses an optimization problem that is solely about being able to offer accurate assurances to customers that the program is likely to work "as promised" and then delivering on those pledges. Good dynamic pricing programs deploy limited marketing resources and use limited rate flexibility to attract customers who offer the greatest

[1](Chapter 3)discusses this design as a "declining block" rate. The industrial organization literature describes the potential economic efficiency of declining block rates, but declining block rates are currently banned in the state of California and perceived as encouraging the wasteful use of energy.

[2]Some customers may respond to paying a high price on the margin by reducing their use further during future critical events. A neoclassical model would suggest that unexpectedly paying the high marginal price would create only a tiny income effect, but loss averse customers and customers who have a mental accounting rule of thumb that it is never OK to buy at the high marginal price might increase their response to critical prices.

reductions in peak and critical period demand.[3] Desirable menus of offers assign consistent offers to a large proportion of the customer population. Customers who perceive that the program will not or did not "work as promised", because they got an inconsistent offer may be less likely to sign up and more likely to leave. Thus, this project explores the feasibility and implications of offering category-level offers tuned to work particularly well for the most desirable customers.

Judging IP Rebate offers by their ability to recruit customers makes sense because IP rebates have very narrow impacts. IP rebate offers are a tool to make critical peak pricing more attractive to consumers. They incidentally reduce month-to-month bill variation where most critical events take place during months with the highest residential electricity demand.[4] Adding IP rebates to a CPP rate does not affect a customer's total annual bill, the utility's revenue or the distributional implications of the rate. This revenue and annual-bill neutrality requires that customers' consumption under CPP with IP Rebates be identical to the incentive-equivalent CPP rate. Their narrowness allows them to coexist well with policies designed to affect revenue, incentives, and equity.

Rates should incorporate additional complexity only when its benefits in tailoring offers to customers' use patterns outweigh its costs. It is costly for regulators, utilities, and consumers to develop, update, and understand complex rate features. Using existing categories of customers that are familiar to regulators and utilities can reduce the cost of tailoring rates. We can maximize the probability that each customer will get a consistent offer via custom offers, this adds complexity. Using historical consumption to customize offers can lead confused customers to increase their consumption in the hope of getting a new rate that offers more rebates and profit opportunities. Using immutable characteristics to make offers to broad categories of customers can avoid this problem.[5]

This project takes a preexisting assignment of customers to 16 groups[6] and shows that collapsing these groups into three categories based on consumption and geography

[3]Chapter 2finds that customers in hotter climates and customers with high total electricity respond more to dynamic price signals.

[4]For clarity and brevity, this document's wording often assumes that the critical events and the customer's highest bills take place in the summer because the electricity system is summer peaking with air-conditioning-driven peaks. These assumptions are correct in most regions of the United States. However, almost all of the statements about summer impacts can be generalized to winter- and dual-peaking systems by replacing "summer" with "the season(s) with high residential demand and most of the critical events."

[5]See Chapter 3 for an extended discussion of this.

[6]These 16 groups come from a modification of the consumption and climate zone categories that the Statewide Pricing Pilot designers used. Chapter 3describes them in more detail.

can perform 96% as well as does making the 16 optimal group-level offers. Using three optimal categories far outperforms one- and two-category offers.[7] Four and five category offers perform modestly better than the three category offer. The optimal, three-category offer assigns at least one group from each of the four climate zones to the low, medium, and high use categories.

Finding that three offers can perform 96% as well as 16 optimal offers is very good news. There are, however, two concerns about these offers:

- The high-performing rates get down to three offers, but are unable to achieve some simplicity goals. These high performing offers categorize customers by overall energy use and geography. Further simplifications that eliminate the need for data about customer energy consumption would would simplify implementation. However, attempts to assign customers to groups on coarse, purely geographic criteria perform much worse. Making a single offer for everyone in the state also performs poorly. The 500-customer SPP data set lacks size to evaluate the effectiveness of making offers by ZIP code or neighborhood – but creating a system that assigned every ZIP code or neighborhood to even one of three offers would be complicated.

- Most customers get consistent offers under the optimal 3, 4, or 5 category approaches, but some do not. The optimal three-category offer makes "consistent" offers to between 75% and 90% of the customers in most groups.[8] The offer inconsistency reflects both the limited predictability of each customer's set of consistent offers, and the compromises required to make a single optimal offer to a category of customers with diverse needs.

Thus, IP rebate implementation in California would assign customers to categories

[7] This paper identifies and reports the characteristics of the optimal one-size-fits-all offer. We cannot rule out the possibility that there is a two category offer that uses better customer-assignment data to perform nearly as well as the best three category offer that collapses the 16 groups. The results here make that seem unlikely. The two category offer is roughly the optimal two-category offer with its two smallest categories collapsed into a single category which gets roughly the average of their offers of 6.5 and 14 kWh per event and makes inconsistent offers on both its large and small customer fringes. Thus, an improved 2 category offer that outperforms the best offer found here would have to selectively reassign the biggest of the small group customers to the large category, so that it could make a smaller offer to the small category. There is reason for concern that, even if such a high performance 2-offer rate exists, that many of its customers would be forcefit into offers that were barely consistent for them, leaving its performance vulnerable to weather shocks and that making changes to accomodate implementation concerns might have a large impact on its performance.

[8] No consistent offer exists for about 3% of all customers. Even an omniscient system would be unable to make these customers a consistent offer, so they are not in the denominator here.

based on energy consumption or a proxy for it and would have to be comfortable with some customers getting offers that are not quite consistent.

Brown and Sibley (1986, 7.3) is a similar intellectual exercise that uses simulations to compare the benefits of a variety of rate structures, some of which are optimized for a diverse group of customers. Orhun (2006) develops menus of products designed to maximize profits from a diverse population of loss averse customers, but assumes that the customer will choose from a menu of products rather than having the firm assign rates using visible customer characteristics. Borenstein (2007) simulates the impact of dynamic pricing and demonstrates the effectiveness of simple hedges using customer-specific characteristics to set the quantity hedged.

The optimization in this paper that attempts to assign 16 geographic and consumption based groups of customers to broader categories can be articulated as a mixed integer linear programming problem. In particular, its structure is equivalent to the facilities' siting problem. There are extensive literatures (Wolsey, 1998; Velle, 1993; ReVelle and LaPorte, 1996) and active research in developing more efficient algorithms to solve these problems which are known for their computational difficulty. This paper develops and uses a specialized algorithm that uses the problem's structure to create a tractable, simplified representation of it. The appendix describes the algorithm.

This paper proceeds as follows:

i. It describes the 16 groups of customers and the objectives and constraints on the menu of offers.

ii. It then estimates the peak-period, energy use impact of exposing each of the 16 groups of customers to peak and critical prices.

iii. The paper then converts these two energy use impacts into a single estimate of each category's summer season, deadweight loss reduction value per customer.

iv. The paper describes the assumptions about the relationship probability that a customer's offer is consistent to the probability that he signs up for dynamic pricing used in the optimization.

v. It uses these estimates and assumptions to calculate value maximizing menus of offers. We explore the implications of restrictions on menu sizes and on the nature of offers.

vi. It compares the customization challenges that this optimization posed to the customization challenges that other dynamic rates pose.

Sections 4.3 and 4.4 attempt to make reasonable approximations of the relative value of each group's response, but 4.6 finds that there are categories of groups of customers with such similar optimal offers that the offers chosen at each size are relatively insensitive to the exact value of each groups response.[9]

4.2 Objectives and Constraints

This project develops an optimal set of offers to customers given psychological, economic and practical constraints.

4.2.1 Economic Constraints

This section uses notation introduced in Chapter 3.

IP rebates make each customer an **offer**, $(\mathbf{q_R}, \mathbf{Q_D})$, which specifies the quantity of rights that the customer gets during each event, $\mathbf{q_R}$, and the number of kWh the declining block marks up each month, $\mathbf{Q_D}$. It is desirable for offers to meet the following constraints for as many customers as possible:

i. **Consistent rebates:** The offer includes enough kWh at the usual price so that the customer gets a (weakly positive) rebate during each month with an event, or $\mathbf{q_R} \geq \bar{\mathsf{q}}_\mathsf{c}$. In other words, the number of protected kWh, $\mathbf{q_R}$, has to be at least as great as the customer consumed during the average event in the customer's highest average-event-use month $\bar{\mathsf{q}}_\mathsf{c} = \max_{m \in M}\{\mathsf{Q_c}/N_m\}$.

ii. **Consistent purchases through inframarginal declining blocks:** Customers buy all of the rights that the offer promised only if the declining block marks up less power than the customer uses each month[10], or $\mathbf{Q_D} \leq \underline{\mathsf{Q_m}}$.

---

[9]The situation is analogous to siting distribution centers to serve a set of clients. Determining the optimal sites for distribution centers relative to the clients depends on measuring the number and cost of each kind of shipment. If the clients are clustered in a few metropolitan areas, so that they are close together within each city but there are large distances between cities, then a wide variety of shipment costs and frequencies will drive us toward building one distribution facility roughly in the center of each cluster of clients.

[10]I assume that $\mathbf{q_R}$ and $\mathbf{Q_D}$ stay the same year round. Making seasonal changese to the declining block size, $\mathbf{Q_D}$, may be an important way to provide consistent offers. Requiring extra contributions early in the year could provide a reserve fund to cover under contributions later. It may be particularly natural to

iii. **Customer-level revenue neutrality:** each customer makes payments for rights equal to the value of the rights they receive. If the customer consistently purchases $\mathbf{Q_D}$ per month (constraint ii) then, this becomes $12\mathcal{M}\mathbf{Q_D} = N_c\mathbf{q_R}(P_c - P_h)$.

4.2.2 Political and Implementation Constraints

It seems to be a political reality that CPP-IPR will be deployed first as an opt-in program and that customers will have to actively choose to leave their status-quo, time-invariant rate to join the CPP-IPR program. This analysis assumes CPP-IPR will be offered to all customers statewide.

This project operationalizes the goals of using existing categories by making offers to categories formed by further aggregating 16 existing climate/usage level groups of customers.[11] Each customer is categorized by its membership in one of four SPP climate zones. Each climate zone contains four power usage levels based on each customer's pre-experiment, summer 2002 average daily consumption. The SPP designated low and high use categories for each climate zone using a climate-zone-specific threshold for membership in the high use group, and this analysis subdivides each high or low use group at the median. Chapter 3 describe the construction of these groups in detail. This is one of many ways to operationalize these simple menu from existing category constraints. It is hopefully typical of the kinds of performance and challenges that we might expect from such an effort. The current effort does not attempt to categorize customers in a way that maximizes the similarity of the customers within each of the initial groups or to solve the exact problem that a real California utility might, in part, because the data set is too small. Clustering data into groups of similar customers is an active research area in operations research (Jain and Murty, 1999; Verma and Meila, 2003).

consider seasonal variations in $\mathbf{Q_D}$ or $\mathcal{M}$ in electricity systems that already seasonally adjust rates. Seasonal adjustments, however, make rates harder for customers to understand.

[11] Chapter 3reports that assigning each group to one of 9 optimally chosen offers can perform as well as choosing one optimal offer per group. The finding that reducing from 16 offers to 9 yields zero loss of performance is likely to be an artifact of the small sample size, but the loss of performance from going from 16 to 9 offers would probably be small as we increase the sample size.

4.2.3 Objective Function

I assume that the utility is trying to choose offers ($\mathbf{q_R}$, $\mathbf{Q_D}$) to attract the customers who will offer the greatest reduction in peak load[12] given that customers' propensity to sign up for CPP-IPR increases with the probability that they get a consistent offer that meets the three economic constraints above. Thus, the offer's benefits are roughly the sum of the peak and critical demand reductions that customers in each group provide when they sign up for CPP times the fraction of customers in that group who sign up.

Formally, the objective function is:

$$\max_{G_{gk},\mathbf{Q}} \sum_{g} A_g V_g P(F(\mathbf{Q}))$$

where:

G_{gk}	is a matrix of 1's and 0's recording the assignment of groups of customers, g, to a smaller number of broader categories, k.
$\mathbf{Q}$	is a matrix specifying ($\mathbf{q_R}$, $\mathbf{Q_D}$) for each category c
A_g	is the number of accounts statewide that group g represents[13]
V_g	is the summer period deadweight loss reduction value of customers in group g. Specifically, $V_g = (\Delta_c^g DWL_c + \Delta_H^g DWL_H)$ where Δ_i^k is group $g's$ kW response to price $i \in \{c, H\}$ and DWL_i is the summer season deadweight loss reduction value of changing consumption by 1 kW during all hours in price period i. Section 4.3 estimates Δ_i^g. Section 4.4 approximates the relative sizes of the two DWL_i values.
$F(\mathbf{Q_g})$	is the probability that $\mathbf{Q_g}$ is a consistent offer for a customer in group g.
P()	is the probability that a customer signs up for CPP-IPR as a function of the probability that they got a consistent offer.

Subject to the following constraints:

[12] I assume that the utility is indifferent between customers reducing load altogether or shifting it other time periods.

	kW impact		weighted average benefit ratio	
	critical	TOU Peak impact	2.2x	6.7x
1 very low	0.016	0.133	0.052	0.031
1 low	-0.013	0.025	-0.009*	-0.011*
1 high	-0.053	-0.149	-0.083	-0.066
1 very high	-0.085	-0.395	-0.182	-0.125
2 very low	-0.012	0.120	-0.008*	-0.010*
2 low	-0.048	-0.052	-0.049	-0.048
2 high	-0.090	-0.220	-0.130	-0.107
2 very high	-0.117	-0.429	-0.215	-0.158
3 very low	-0.070	-0.062	-0.068	-0.069
3 low	-0.113	-0.230	-0.149	-0.128
3 high	-0.132	-0.393	-0.213	-0.166
3 very high	-0.205	-0.740	-0.372	-0.274
4 very low	-0.136	-0.256	-0.174	-0.152
4 low	-0.179	-0.471	-0.270	-0.217
4 high	-0.202	-0.570	-0.317	-0.250
4 very high	-0.286	-0.869	-0.468	-0.361
*indicates a negative value was rounded to zero before computing the wtd. avg.				

$\sum_k G_{gk} = 1$	each group is in a total of exactly one category.
$.36 * 15\mathbf{R}_k = 12 * .025 * \mathbf{Q}_{\mathbf{D}k}$	**Customer level revenue equivalence.** the number of units, $\mathbf{Q}_{\mathbf{D}k}$ that are marked up for category k exactly pay the costs of the category's rights
$\mathbf{Q_i} = \mathbf{Q_j}$ if $\exists c$ such that $G_{ic} = G_{jc} = 1$	subject to the constraint that if groups i and j are in the same category c, they get the same offer.

4.3 The average magnitude of response of customers in each group

This section estimates the size of customer response to CPP's peak and critical prices during the summer rate season. This section uses the same data, faces the same estimation problems, and uses a strategy quite similar to that described in Chapter 2. This section estimates critical impacts using the average conditions from the 15 critical events that the SPP called on days with statewide loads above the 80th percentile of the

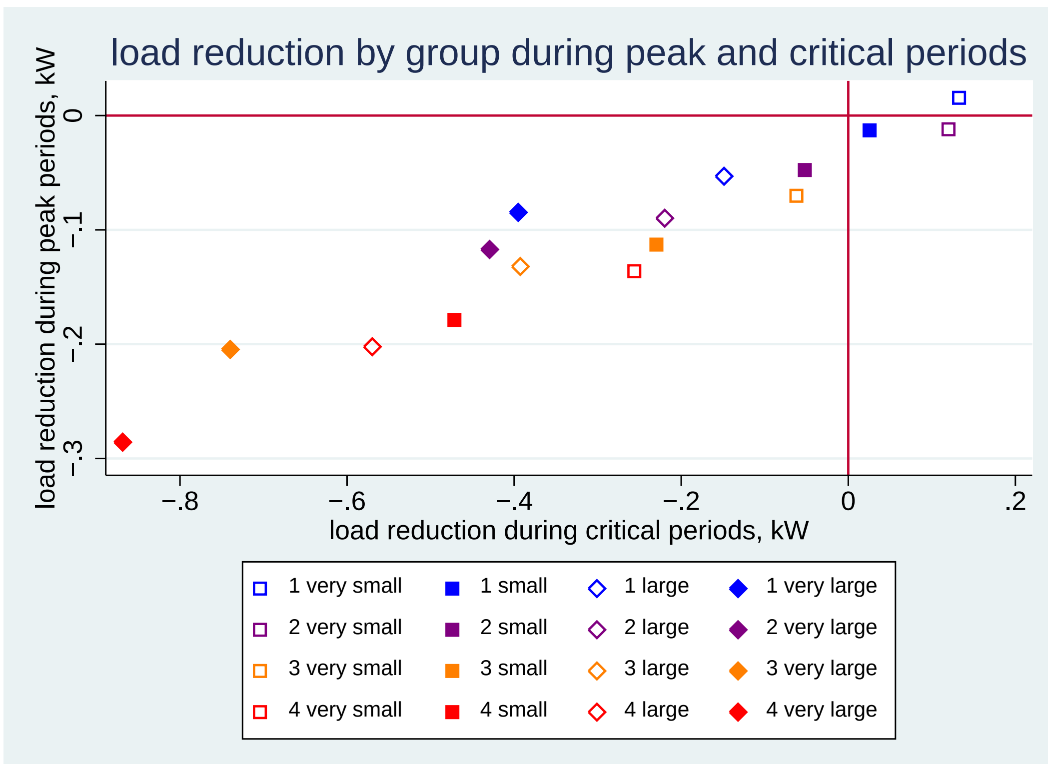

Figure 4.1: Difference regressions show that big customers in the hottest climates responded the most; and responded significantly more during critical periods than during ordinary peak periods.

distribution of 2003-04 summer season loads. Very high demand is typically a component of the scarcity scenarios that would prompt critical prices and make reductions in electricity use quite socially valuable. Generation and transmission problems are also often important components of scarcity, but high demand makes the system far more vulnerable to these. The ability to call 12 critical days allows operators to impose a critical price on nearly 10% of days during the average summer season which contains 128.4, non-holiday weekdays.

The difference-in-difference and single difference versions of the three regression specifications that use splines for temperature impacts in Chapter 2provide a variety of estimates that tell qualitatively similar stories. The results reported in table 4.3 are calculated with methodology identical to regression 5 in Chapter 2except that they use just data from the "treatment" period and drop all variables identified from pretreatment data alone. Figure 4.3 and table 4.3 shows that high use customers and customers in hot climates respond significantly more to CPP than do customers in more temperate climates.

4.4 The value of response:

Section 4.3 estimates impacts on critical and peak period use separately. The optimization in section 4.6 requires a single measure of each kind of customer's value. A TOU peak period every weekday creates incentives to conserve during far more hours than do sporadic critical periods, but critical-period consumption has a much higher social cost. Unless customers' changes in use in response to peak and critical prices are perfectly correlated,[14] then it is important to consider the relative social value of responses to the critical and peak price changes.

This section makes the impacts commensurable by approximating the two price's relative deadweight-loss-reduction benefits per kW. This allows us to minimize the deadweight loss from underpricing power during summer season critical and peak periods.[15]

Precise deadweight loss calculations require accurate demand curves and hour-by-hour, weather-dependent marginal social costs. We have quite limited knowledge of

[14]If, for example, each group's average customer reduced his use by exactly twice as much during critical hours as during peak hours, then Section 4.6's optimal offers would not depend on the relative importance of peak and critical prices calculated in this section.

[15]A logical extension of this criterion would be to minimize deadweight loss round the clock and during both winter and summer rate seasons. Quirks in the SPP rate make analyzing the winter season challenging. Including round-the-clock summer impacts requires running a straightforward set of regressions and then dealing with the likely fact that the results are likely to be substantively and significantly indistinguishable from zero and to show a lot of small-magnitude point estimates that contradict economic theory.

both. Most Regional Transmission Organization (RTO) electricity spot markets have weak demand-side abilities to react to high-price and manage scarcity. They thus cap spot market prices and often use side payments to ensure that there will be adequate of generation. This approach causes electricity spot prices prices to understate the true social costs of scarcity. Despite these interventions, wholesale prices that are well under \$0.10/kWh in a typical afternoon hour sometimes hit price caps that are typically \$1.00/kWh (\$1,000/MWh) and a few very high prices often represent a significant part of the cost of summer energy. [16] The available estimates of demand curves are imperfect and Chapter 2finds that the impact of price changes on quantity demanded is quite sensitive to weather conditions which suggests that the demand curve both shifts and changes slope as temperatures change.

A variety of utilities have deployed CPP rates. These rates attempt to offer customer a Pareto-improving deal by moving prices closer to marginal cost, subject to a variety of constraints.

There are, however, a variety of sources of evidence about the relative deadweight-loss reduction value of a kW response during peak and critical periods.[17],[18],[19] Table 4.1

[16]We could expand the analysis in this section with the simulated hourly prices from Borenstein (2005a) and with data from regional transmission organization (RTO) markets. Indeed, the gold standard for the present analysis would be to use daily weather conditions to link market prices to estimated change in quantity demanded. Most RTO markets' spot-market price series understate the costs of scarcity and managing extreme demand. Dynamic pricing is likely to have the greatest value in markets experiencing scarcity. Price-responsive customers can be a substitute for RTO markets' more distortionary and costly efforts to manage scarcity through non-price means. It might be an imperfect, but interesting analogy to look at spot prices from an RTO market that is experiencing some scarcities that did not lead to the kind of institutional crisis that California saw in 2000-01. Data from the Delmarva region (Delaware and the eastern shores of Virginia and Maryland) of the PJM market might be a reasonable choice for this task as might be ISO-New England's Southern Connecticut region. The appropriateness of these data depends strongly on exactly how the RTO's strategies for managing scarcity affect the visibility of scarcity rents in its prices.

[17]See Borenstein (2005a) for an extended discussion of this and for an effort to simulate equilibrium wholesale prices in the absence of these interventions. Borenstein and Holland (2005) prove that adding a marginal customer to dynamic pricing offers the largest benefits when most customers are on time invariant pricing and that the marginal benefits decrease in the number of customers enrolled.

[18]PG&E has filed a Business Case for its Advanced Metering Deployment that calculates the costs and benefits of a critical-period only CPP rate. Ameren planned to conduct a cost-benefit analysis using evidence from the trial with the rate discussed in table 4.1.

[19]Gulf Power's Good Cents Select CPP program is a full scale, opt-in CPP program that may illustrate the potential of the approach. Gulf Power's claims that its 7200 customers' response to a critical price signal is equivalent to an 80 MW combustion turbine plant (White, 2005). Energy Information Administration (EIA) estimates suggest that plant would cost \$32,000,000 to build and that its fixed operating and maintenance costs would be \$880,000 a year (Conti et al., 2006). This comes out to averting a one-time investment of \$4400 and annual maintenance costs of \$120 per customer from the critical price portion of the program alone – before we factor in fuel and variable operating and maintenance costs. A careful engineering cost-benefit study of its hour-by-hour social benefits and impacts on the utilities' revenues might be useful to policy makers in understanding both what to expect from dynamic pricing and in understanding how to make dynamic pricing a win-win proposition for participating customers, utilities, and other customers.

shows examples of these rates and the DWL ratios they would imply if the rates peak and critical prices, P_H and P_c, were the marginal cost of power during the peak and critical periods respectively.[20] This exercise attempts to glean significant insights by using evidence from a variety of imperfect data points. These may be useful approximations of the bounds on the relative size of the deadweight loss reduction during each period and to see whether adjusting the periods' relative importance within these bounds has much effect on the optimal offer. If the optimal offer is highly sensitive to the choice of bounds, it would suggest that the existing evidence is insufficient to optimize.

We find evidence that the value of a kW of reduction during TOU peak periods is worth between 0 and 6.7 times as much in DWL as a kW of reduction during critical periods. Thus, we will run the analysis below using the DWL ratio of 2.2 from the example rate; and then run robustness checks with ratios of 0 and 6.7.

This analysis calculates the deadweight loss reduction values of each kW, a flow variable, rather than kWh, a unit of energy, because its regressions estimate impacts on per-hour flows (or, as other papers call them, kilowatt-hours-per-hour). Five of the six rates considered here have daily afternoon TOU peak periods. These five rates have between 8.5 and 10.7 times as many TOU peak hours as critical hours. Thus each TOU-peak kW (or, equivalently, kWh/h) of flow translates into 8.5 to 10.7 times as many units of energy –

Unfortunately, Gulf Power is not part of an organized energy spot market which limits the availability of data about the kind of scarcity rents, market power and capacity market prices that CPP will affect in other regions. Further, these are customers in a hot, humid climate who received automated thermostats.

A similar style of engineering analysis could calculate the social costs avoided per kW, or benefits B_i, of peak and critical response as follows:

$$B_H = \Delta_H(AnnualFixedCost/kW\frac{N_H}{N_H+N_c} + (N_H)(VariableCost/kWh))$$

$$B_c = (\Delta_c - \Delta_H)(AnnualFixedCost/kW + \Delta_H(AnnualFixedCost/kW\frac{N_c}{N_H+N_c} + \Delta_c(N_c)(VariableCost/kWh)$$

$$B_L = (\Delta_L)(AnnualFixedCost_L/kW + \Delta_L(N_c)(VariableCost_L/kWh)$$

where Δ_i is the change in electricity consumption in response to price i and N_i is the number of hours in period i. The nature of the offpeak change in costs is likely to be qualitatively different from the peak change in costs – since it is likely to reflect the cost of the increased use of fairly efficient plants that are on the margin offpeak and perhaps the additional capital cost of substituting a few more efficient plants that can now run more hours for less efficient plants.

Reasonable estimates of these costs are readily available from natural gas spot market prices, utility's costs of capital, and EIA data about one time capital costs and annual fixed costs of a combustion turbine plant. Borenstein (2005a) makes this kind of calculation of the costs of building and operating all the plants in a competitive electricity market in equilibrium with varying percentages of customers on real time pricing.

[20] The calculations here further assume that the cost of energy is uniform across the utilities' service area, but these calculations are easy to extend to incorporate the kind of locational cost differences that are common in electricity markets.

kWh's – as does a critical period kW. This far greater number of kWh's per kW flow allows ratios that find that a TOU peak kW has more value than a critical kW, despite the fact that there are far greater deadweight losses per unit of energy (kWh) during critical periods.

4.4.1 Approximating the size of reductions in deadweight losses

PEPCO's rate has only critical and time-invariant periods. If this reduces deadweight loss as much as any rate that also raised prices every weekday afternoon, this would imply that raising prices every weekday afternoon has zero benefit.[21] The PJM spot market prices from PEPCO's service area are higher during weekday afternoons, which means that PEPCO's rate failed to address some deadweight loss. It is possible, however, that PEPCO's implication that only critical periods matter is closer to reality than any of the other rates. Estimates of the cost of blackouts ("the value of lost load") or of the highest market clearing price in a perfectly competitive spot market (Borenstein, 2005a) are roughly 100 times higher than the average afternoon marginal cost of power. Thus, it is worth considering the case in which (nearly) all the benefits come from critical period power.

Estimating the size of deadweight loss requires making an assumption about the demand function. This paper experiments with two demand functions. One assumes that demand is linear but varies between peak and critical periods. The other assumes that demand is a step function.

4.4.2 Benefit Ratios with a Linear Demand Function

If demand is linear, with different slopes on peak-priced and critical days, then the deadweight loss reduction from charging P_H instead of P_u for power that has a social cost of P_H is $.5S(P_H - P_u)^2$ where S is the slope of the demand curve. The critical period slope may differ on critical days if a program phones customers to inform them of the high price of power, as the SPP did. Further, the slope may differ because critical days tend to fall on days when factors like heat and a relatively small number of people on vacation combine to create extreme demand. We can back out the slope S given using the available kW impacts of the peak and critical prices as $S = \frac{\Delta kW}{\Delta P}$. This section's approximation of the relative value of a kW of demand reduction normalizes the ΔkW component of S to 1.

[21] PG&E is rolling out the same style of opt-in rate that adds critical periods to otherwise time-invariant prices (Pacific Gas & Electric News Department).

						ratio of the value of a kW response, peak to critical	
Source of Rate (State)	summer season critical hours	summer peak hours	ratio: peak to critical hrs.	critical price diff, ρ_c	TOU peak price diff, ρ_H	upper bound, step function	linear
Pepco (DC)	48	0.0	0.0	56.1	-1.1	0.0	0.0
Ameren (MO)	40	341.7	8.5	22.4	9.1	3.5	6.7
Gulf Power (FL)	87.6	642.1	7.3	24.6	3.7	1.1	0.8
SPP Low Ratio (CA)	60	642.1	10.7	41.0	9.8	2.6	2.9
Example, Table 2 in Chapter 3	60	642.1	10.7	45.4	9.4	2.2	2.2
SPP High Ratio (CA)	60	642.1	10.7	60.0	11.6	2.1	1.9

Table 4.1: The relative sizes of summer season, peak and critical period deadweight losses under a variety of CPP rates. Notes: Table 4 in Chapter 3presents more details about these rates. The Ameren rate allows for the 10 critical periods reported in this table, but only called eight during its experiment as reported in Chapter 3. The number of peak hours is the average number of non-holiday weekday hours reflecting the fact that the first two days in any partial week contain $\frac{10}{7}$ weekdays in the average year and that any additional days will be weekdays. Ameren has daily peak periods during a four month summer season, while Gulf Power and the California SPP used a six month summer season. The Gulf Power rate specifies that no more than 1% of all hours year round will be critical, without commiting to any seasonal schedule of events. Gulf Power does call winter critical events during unusually cold weather to deal with electric-heat-driven load. This table thus presents the upper bound on the importance of summer critical periods in the Gulf Power region. (Sources: Wilson (2006); Pepco; Voytas (2006); Ameren; Gulf Power; Pacific Gas & Electric (c); San Diego Gas & Electric)

The final optimization uses the ΔkW_g value for each group from Table 4.3. Since the high and low ratio rates' impacts were statistically indistinguishable and the high- and low-ratio customer pools had roughly the same size, Table 4.1 uses the simple average of the high and low ratio price change as ΔP.[22] Table 4.1 computes the ratio of the value of a kW conserved during each TOU peak hour to a kW conserved during each critical hour as:

$$\frac{N_H \frac{1}{\Delta P_H}(P_H - P_u)^2}{N_c \frac{1}{\Delta P_H}(P_c - P_u)^2}$$

Where:

N_H is the number of peak priced hours during the summer season.

N_c is the number of critically priced hours during the summer season.

ΔP_i is the average price change during period i that California Statewide Pricing Pilot customers experienced. Linear demand implies that customers who experience different price changes in other markets will respond by different amounts than the SPP customers did.

4.4.3 Benefit Ratios with a Step Demand Function

The finding that customer responses to high and low ratio rates are statistically indistinguishable suggests that linear demand curves may be an inappropriate model. Step-function demand predicts that customers' demand would be constant within price intervals, which is consistent with our regression findings. Further, holding temperature constant, a step demand function approximates the behavior of dynamic pricing customers who program their thermostats to run their climate control systems less when prices hit threshold levels.[23] However, finding a step demand function that predicts response to both the SPP and Gulf Power's rates is tricky. Gulf Power's summer weekday afternoon price of 11.7 cents is lower than the California rates' time-invariant and offpeak prices for high use customers. Consider CPP price, $p_i, i \in \{H, c\}$, and the time-invariant price, p_u. A family of step functions that specify demand as a function of the price difference $\rho = p_i - p_u$ can, however, model this

[22]Specifically, it uses average price increases of 50.5 cents during critical periods and 10.72 cents for TOU.

[23]Indeed, we could imagine a field experiment that would expose customers with these thermostats to several price schedules; record their thermostat settings about how much they respond to each price level; and use this to bound a step demand function for each customer. Exploring how customers set their thermostats to respond to dynamic pricing and how the quantity demanded emerges from the interaction of set point and temperature has significant value in understanding how putting large numbers of customers on dynamic pricing changes the investments required to handle a variety of extreme demand scenarios.

behavior. This family has the form:

$$\begin{aligned} D(\rho) &= Q_L \textit{ if } \rho < \rho_1 \\ & \quad Q_H \textit{ if } \rho_1 \leq \rho < \rho_2 \\ & \quad Q_c \textit{ if } \rho_2 \leq \rho \end{aligned}$$

Members of this family of functions predict the expected response to a set of rates r if $0 < \rho_1 \leq \min_r\{\rho_{H,r}\}$ and $\max_r\{\rho_{H,r}\} \leq \rho_2 < \min_r\{\rho_c\}$. In other words, ρ_1, the price increase sufficient to cause people to reduce to peak period consumption levels must be positive, but weakly smaller than the smallest of the peak-period price increases. Similarly, ρ_2, the price increase sufficient to get customers to shift from a peak to a critical conservation strategy must be larger than any peak period markup, but weakly smaller than any critical markup.

Such a demand function could come from by a variety of combinations of substitution-driven, cross-price elasticity, reference-dependent own-price elasticity relative to a reference point of p_u, and conventional own-price elasticity cannot explain. Reference dependent customers might conserve power during periods when it is above their reference price – leading to conservation rather than substitution in patterns that conventional own price elasticity. Neoclassical own-price elasticity alone would predict that California customers who pay more offpeak than Gulf Power customers pay during afternoon peaks would use set their thermostats higher around the clock than their Gulf Power counterparts would on peak.[24]

The TOU peak deadweight loss rectangle from a step function is the dark purple region in the lower right corner of figure 4.2 or:

$$(D(\rho_H) - D(0))(\rho_H - \rho_1)$$

The critical price deadweight loss region is the sum of the three shaded rectangles in figure 4.2. It is the sum of the rectangles formed by the TOU peak and critical responses to the critical price:

[24] We do not have sufficient data either to refute this strange implication or to develop a four step function that would describe behavior consistent with this implication.

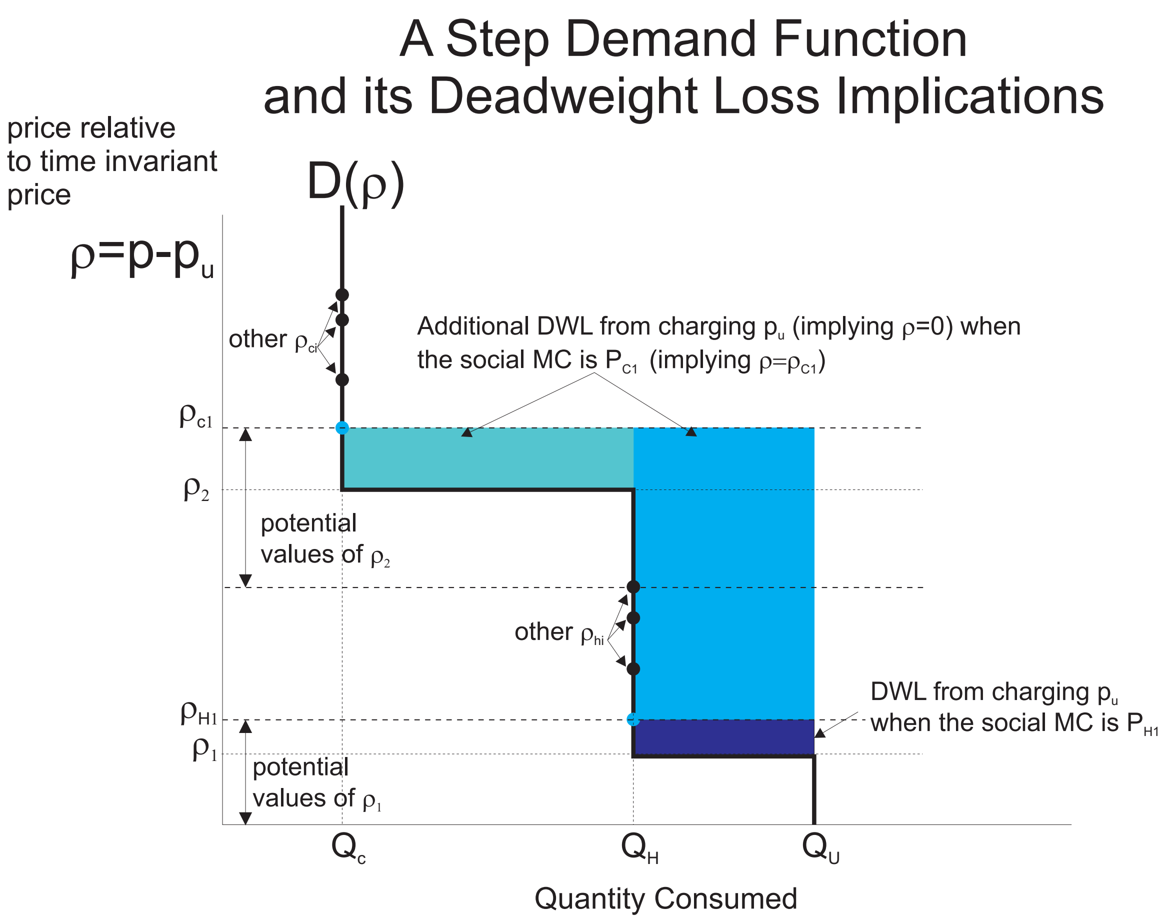

Figure 4.2: This illustrates the family of step function demand curves that consistent with the rates considered in this essay.

$$(D(\rho_c) - D(\rho_H))(\rho_c - \rho_2) + (D(\rho_H) - D(0))(\rho_c - \rho_1)$$

This step function requires further assumptions or data to calculate the deadweight loss reduction ratios. However, we show that zero is the lower bound on the step demand function deadweight loss ratio for any set of rates and report an upper bound in Table 4.1. The ratio of deadweight losses from peak and critical prices is sensitive to the unobserved values of ρ_1 and ρ_2. The ratio is often sensitive to the relative sizes of the responses to peak and critical prices $D(\rho_c) - D(\rho_H)$ and $D(\rho_H) - D(0)$. Thus, this family thus supports a range of deadweight loss ratios, α, for rate j ranging from:

$$\frac{(D(\rho_H) - D(0))(\rho_{H,j} - \min_r\{\rho_{H,r}\}}{(D(\rho_c) - D(\rho_H))(\rho_{c,j} - \min_r\{\rho_{c,r}\}) + (D(\rho_H) - D(0))(\rho_{c,j} - \min_r\{\rho_{H,r}\})} \leq \alpha <$$
$$\frac{(D(\rho_H) - D(0))(\rho_{H,j} - 0)}{(D(\rho_c) - D(\rho_H))(\rho_{c,j} - \max_r\{\rho_{H,r}\}) + (D(\rho_H) - D(0))(\rho_{c,j} - 0)}$$

Cases where the customer is indifferent about buying creates two interesting bounding cases:

- One such case implies that this family of step functions includes members that have zero peak-price-period deadweight loss. There is at least one rate for which the customer's maximum willingness to pay for that power, ρ_1, equals the peak price increase, ρ_H.[25] The PEPCO rate provides further evidence that it is worth considering a lower bound of 0.

- The directly analogous case where the customer is indifferent between buying $D(\rho_H)$ and $D(\rho_c)$ at the critical price because $\rho_c = \rho_2$ makes critical period deadweight loss reduction benefits as small as possible. It provides an upper bound on the ratio of peak to critical benefits. Graphically, the top left rectangle in figure 4.2 would have zero

[25]There is an analogous, but far weaker, result which observes that this ratio has an upper bound where critical and peak kWh reductions have the same deadweight loss reduction value. In other words, this ratio can achieve an upper bound of $\frac{N_H}{N_c}$. In practice, deadweight losses will be well below this bound. Peak and critical prices approach having the same deadweight loss implications in the limit only if all of the peak and critical prices of the rates under consideration approach being the same. Specifically, this requires that all the rates we observe have $\rho_{c,i}$ arbitrarily close to $\rho_{H,i}$ and to $\rho_{H,j}, \forall j$.

height. The upper bound on the deadweight loss ratios is, conveniently, independent of the quantity responses to these two price changes and has the functional form:

$$\alpha \leq \frac{N_H(D(\rho_H) - D(0))(\rho_{H,j} - 0)}{N_c(D(\rho_H) - D(0))(\rho_{c,j} - 0)} \tag{4.1}$$

Table 4.1 reports the upper bounds calculated using formula 4.1 to yield the relative DWL reduction values of a 1 kW reduction during each hour under each rate. The largest of the upper bounds identified by this calculation happens to come from the rate with the smallest ρ_c which means there are step demand functions consistent with the constraints above with a ρ_2 that makes this approximation exact.[26]

4.4.4 Dealing with Negative Point Estimates of Benefits

The optimization approach taken below attempts to provide the maximum possible sum of savings. Groups of customers with point estimates that indicate that they generated negative benefits pose challenges for this method.

i. I chose a set of estimated results that yielded positive point estimates for the benefits from a larger number of groups than other estimates yielded. The single-difference estimates that compared the two groups' outcomes during just the treatment period yielded more positive point estimates than did the difference-in-difference estimates. The simplest estimate, regression specification 5, yielded more positive and more tightly correlated point estimates (and also more tightly correlated estimates) than did the more complicated estimates 6 and 8.

ii. The next section describes a method of calculating the benefits of enrolling each customer as a weighted average of its TOU peak and critical benefit point estimates. For each pair with one negative and one positive benefit point estimate, I round the

[26]This, in general, need not be the case. To see this, we can take a rate with the smallest critical markup, $\underline{\rho_c}$ and assume that this rate's TOU peak markup obeys: $\underline{\rho_H} < \frac{2}{3}\underline{\rho_c}$. (Its TOU peak rate need not be the lowest of the set seen.) Then choose any $\epsilon > 0$ such that $\underline{\rho_H} + 3\epsilon < \underline{\rho_c}$. We can now construct a rate with a critical price-difference strictly higher than $\underline{\rho_c}$, namely $\underline{\rho_c} + \epsilon$ that outperforms $\underline{\rho_c}$ on this metric by increasing ρ_H by an amount greater than the ratio $\frac{\underline{\rho_H}}{\underline{\rho_c}}$. Specifically, choose $\underline{\rho_h} + 2\epsilon$ and then notice that:

$$\frac{\underline{\rho_H}}{\underline{\rho_c}} < \frac{\underline{\rho_H} + 2\epsilon}{\underline{\rho_c} + 3\epsilon}$$

negative benefits to zero before calculating the weighted average.[27] This rule affects the weighted average benefits for the Climate Zone 1 low use and the climate zone 2 very low use group which had negative point estimates of critical period benefits[28], while the climate zone 1 very low use group had negative point estimates of both benefits. This change simply increased the Climate Zone 1 low use benefits while it allowed us to include the climate zone 2 very low use group in the optimization that would have otherwise had to been dropped. This change does not affect the rank ordering of the benefits and even after rounding negative benefits upward, the two groups with potentially negative benefits still have very small positive benefits.

iii. I drop the groups that have negative point estimates of total benefits from the optimal offer calculation. I then assign them to best fitting offer that was chosen to maximize benefits for all other groups. This means that the optimization problem becomes easier and eventually trivial as the number of groups (pigeons) shrinks toward the number of amalgamated categories (pigeonholes).

4.5 The Model of Consumer Choice Underlying this Optimization

The optimization below uses the follow model of consumer choice.

- Customers decide whether to sign up for dynamic pricing or to remain on their existing, time invariant pricing.
- Customers do not know the details of how their consumption patterns relate to the consistent offer thresholds, but get a signal from the utility about the fraction, f of customers in their class will get consistent offers.
- Customers base their decision about whether to sign up, in part, on f the fraction of customers in their class who get consistent offers.

[27] Notice that this has some troubling implications. For example, consider two groups of customers with positive, identical benefit point estimates during TOU peak hours. If one has a negative point estimate of critical benefits, while the other has a zero point estimate of critical benefits, there is reason to think that the group with zero benefits is more valuable, but this rounding technique will render them indistinguishable.

[28] There are far fewer critical days than TOU peak days, so it is likely that critical impacts are less precisely estimated.

- Customers vary in the weight they place on the behavioral component of their preferences – i.e. on how much disutility they expect to experience from a bill surprise and thus on the likelihood that they will get a consistent offer.[29]
- Any increase in the probability that a customer gets a consistent offer leads to some increase in the probability that they sign up[30]
- Customers have the same sensitivity to changes in the probability that they get a consistent offer regardless of their climate zone and the offer that they are getting.

It is clear how to do this optimization if we slightly strengthen that assumption to require that the probability of each person signing up for dynamic pricing is a linear function of the probability that customers in his category get a consistent offer. In other words, the assumption that an increase of 1% in the probability that a customer will get a consistent offer has the same effect on the probability that the customer signs up[31], then it is straightforward to rank offers' desirability. More formally, model the probability that a customer in category c given an offer $\mathbf{Q}$ signs up as $P_c = af_c(\mathbf{Q}) + b$ where P_c is the probability that a customer signs up given f_c, the fraction of customers in category c who got a consistent offer. If we can make these assumptions, then our preference orderings of offerings will remain the same for any values of a and b that satisfy $a > 0$, $aP_c(1.0) + b \leq 1$, and $b \geq 0$ because $\frac{\partial P_c^s}{\partial P_c^c(\mathbf{Q})} = a$. If we relax the assumption that $\frac{\partial P_c^s}{\partial P_c^c(\mathbf{Q})} = a$ where a

[29] We can operationalize this as a distribution of the η parameter in Koszegi and Rabin's (2005a) "A Model of Reference-Dependent Preferences". I am hopeful that careful choices of the distribution of the distribution of the η parameter exist that can yield the linear behavior described below for a variety of well-behaved demand functions. Utility from consuming bundle x while having a reference point r has two components: $U(x|r) = M(x) + \eta N(x|r)$ where $M(x)$ is conventional economic utility from consuming bundle x and $\eta N(x|r)$ is gain-loss utility. It is likely that the formal models would predict a feedback effect where decreasing the likelihood of losses from a purchase of power would increase their utility directly, by decreasing their expected losses, and indirectly, by increasing marginal utility which prompts them to buy more power.

[30] Having the number of customers accepting the offer increase for all f between 0% and 100% requires that some customers expect to get so much (dis)utility from other aspects of dynamic pricing that they will sign up for (decline) dynamic pricing if the probability that they get consistent rebates is 0%(100%). The number of accepting (declining) customers could hit zero in the limit as f goes to 0 (100%) but would have to be positive for any $\epsilon > 0$ ($< 100\%$). If the number of accepting or declining customers were to hit zero anywhere in the interior of the range from 0 to 100% then the marginal benefit of increasing f would be zero at that point. Some of the customers who are unaffected by the probability of getting a consistent offer could be neo-classically economically rational and thus indifferent to the performance of IP rebates' efforts to reframe prices.

[31] This assumption, for mathematical tractibility, ignores Kahneman and Tversky's well-documented finding that while the difference between a 50% and a 51% chance has expected utility implications identical to the benefits from going from 99% to 100%, they seem very different. A future revision could incorporate probability weighting.

is the same for all customer classes and for all $\mathbf{Q}$, then optimization will become quite sensitive to the values of a specific to their situations. Relaxing the assumptions would require the evidence about a from field experiments or market research before performing the optimization.

4.6 Choosing an optimal, small set of CPP-IPR offers:

Chapter 3 shows that splitting the state into 16 bins and making an offer to each bin can yield very good CPP-IPR performance. It constructed objective functions for each group that mapped the number of kWh of rights offered to the number of customers who got consistent offers. Its analysis shows that we can collapse the 16 group-level offers to 9 categories without making any compromise in the number of customers getting consistent offers, but that the optimal one-size fits all offer would perform poorly. This section finds that optimally assigning the 16 groups to 3 categories performs only very slightly worse than the 16 group (or equivalent 9 category) rate. The optimal one and two offer rates perform far worse, while adding more categories beyond three adds a small amount of value. Further, we find that the groups' objective functions within each optimal category are similar enough that revising the criteria described in sections 4.3-4.5 is unlikely to have an important impact on the offers made so long as the goal remains the provision of consistent offers to a large number of customers in many regions.[32]

However, the optimization finds that all optimal rates split customers from each climate zone into at least high and low use-level groups that get different offers. Rates that make the same offer to every customer in a broad geographic area are attractive in their administrative simplicity, fair appearance, and clarity to customers. This section, however, provides strong evidence that making a single offer to broad geographic areas that contain a diverse housing stock cannot perform well. The SPP sample is too small to test whether making offers at the ZIP code or neighborhood level could work well.

This section uses the empirical results and assumptions about consumer choice

[32]The solution is probably less vulnerable to poorly chosen weightings of groups than to weather conditions that drive up critical period use beyond customers' levels of rights within a month or that depress use during low use months. This poses the greatest likelihood of giving inconsistent offers to identifiable groups of customers whose offers were near an endpoint of the range from the smallest number of rights that we predict that they need to get consistent rebates to the largest number of rights we predict that they can buy. If customers' propensity to exit rises significantly when they get inconsistent offers, knowing which customers are most likely to get inconsistent offers could facilitate policy responses to extreme weather patterns.

above to characterize the performance of the optimal 1-5 category groups and two climate-based categorizations. The formal model implies that customers sign up for CPP in direct proportion to the probability that it makes them a consistent offer.

Figure 4.3 gives the flavor of the results:

- Making the optimal single statewide offer performs only about 74% as well as does making the optimal offer to each of the 16 groups of customers.

- Increasing the complexity of offers by defining a second and third category of customers and making different offers to each category can have significant benefits. Optimally assigning the 16 groups to 2 (3) categories and then making the optimal offer to each of those categories performs 91% (96%) as well as does making an optimal offer to each of the 16 groups.

- The marginal benefits of adding more categories are much smaller. Adding fourth and fifth categories increases performance to 98 and 99% of the 16 group optima. Chapter 3 showed that 9 a category offer could make consistent offers to every customer who would get one under the 16 group optima. If the costs of adding more categories are low, adding them is worthwhile. But the modest benefits of adding fourth and fifth categories only justify modest outlays to do so.

- Requiring that the offer make a single offer to every customer in each climate zone significantly reduces performance. Making an optimal offer to customers in "hot climates" (zones 3 and 4) and an optimal offer to customers in "cool climates" (zones 1 and 2) performs 78% as well as does making 16 optimal offers. Further dividing these categories to make an optimal offer to each climate zone increases performance by less than a percentage point. The tables in section 4.6.2 show that this approach performs poorly because the optimal one-size-per-climate zone offers work well for the customers in the middle of the size distribution, but perform poorly for both high and low use customers.

4.6.1 Categories Bring Together Groups with Similar Objective Functions

Figures 4.4, 4.5, and 4.6 show the group level objective functions grouped using the optimal 3 category offer. Each group's objective function is the percentage of customers in

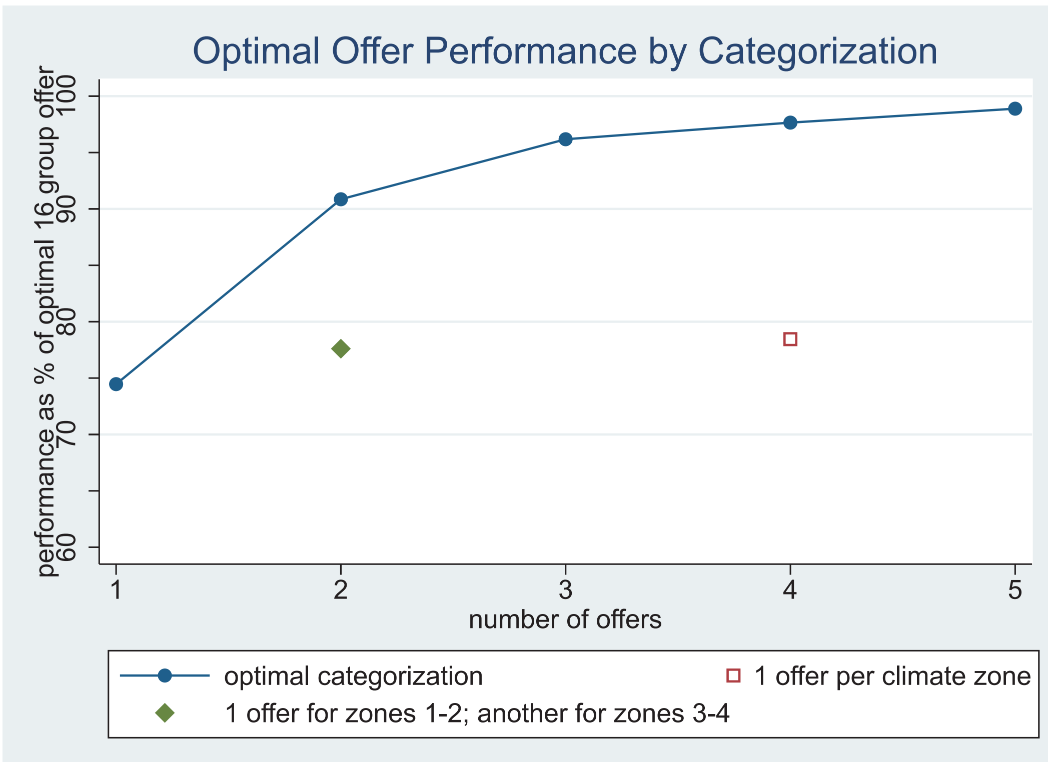

Figure 4.3: The performance of the optimally chosen 1 to 5 category offers, plus the performance of the optimal hot and cold climate zone offers and the optimal offer to each of the four climate zones. Customer heterogeneity within and across climate zones turns out to be quite important. Hence, there are large performance benefits from 1) making at least three categories and 2) allowing similarly located, large and small consumers to get different offers.

each group who would have gotten a consistent offer under the group-level optimal offer who got consistent offers at each level of in kWh per event rights for each group. The vertical lines show the optimal offer for each category. The optimization attempted to choose the categorization and offer (vertical line) that maximizes the weighted sum of the group-level objective functions.

Figures 4.5 and 4.6 make it particularly clear that the optimal categories contain groups with objective functions with similar, overlapping optima. Category 1 in figure 4.4 appears slightly more diverse.[33] There is a broad harmony of objectives within homogeneous categories like these since delivering the maximum level of benefits for one group's objective function or weighted sum of group objective functions will generally perform quite well for the groups with similar objective functions.

Testing the sensitivity of the optimal offer to changes in the metric used to measure the value of each customer shows the implications of the combination of similar objective functions and the chosen estimate's very high correlation between critical and peak period benefits.[34] The optima presented in this section use a metric that puts 2.2 times as much importance on the kW response to the TOU peak price as it puts on response to the critical price. Changing to a metric that uses a TOU:critical ratio of 6.7 does not change the optimal 2, 3, and 5 category offers. The optimal, four-category offer changed from one that achieved 97.7% of the 16 category offer's performance to one that achieved 97.4% of its performance according to the 2.2 ratio metric.

4.6.2 The Optimal, Group-Level Performance of a Variety of Categorizations

This section presents the performance of the each of the optimized offers for each group. Its tables report 1) the category to which the optimization assigns each of the 16 groups and 2) this categorical offer's benefits as a percentage of the benefits that the group would deliver if we made 16 optimal, group-level, offers.[35] The optimization uses group-level

[33]Bootstrapping the set of customers would let us test the hypothesis that the optima or regions near the optima are statistically indistinguishable.

[34]The high correlation between the two kinds of impacts could be sufficient to drive the insensitivity findings here. Eyeballing the objective functions suggests that the similarity among objective functions may also be sufficient to get this result. Future robustness tests can check whether the similarity of objective functions is sufficient to preserve these findings when we use other impact estimates, that happen to be much less correlated.

[35]The algorithm always assigns zone 4, very high use customers to category 1. Otherwise, the category numbering is arbitrary.

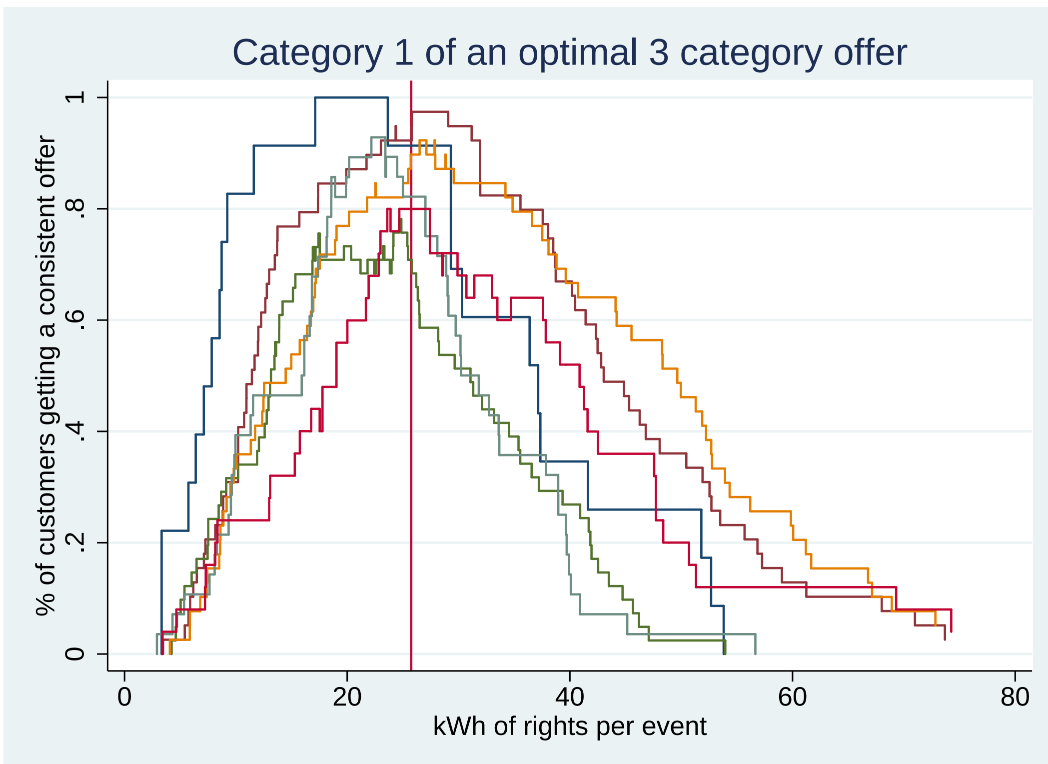

Figure 4.4: The objective step functions for the groups in category 1.

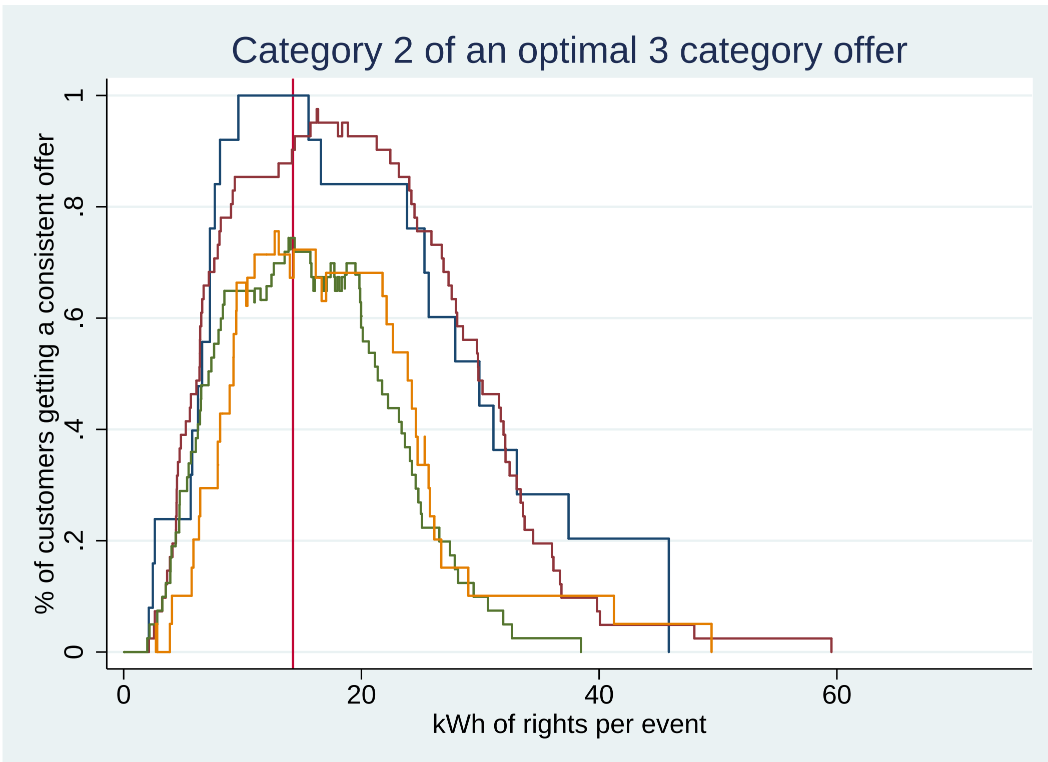

Figure 4.5: The objective step functions for the groups in category 3.

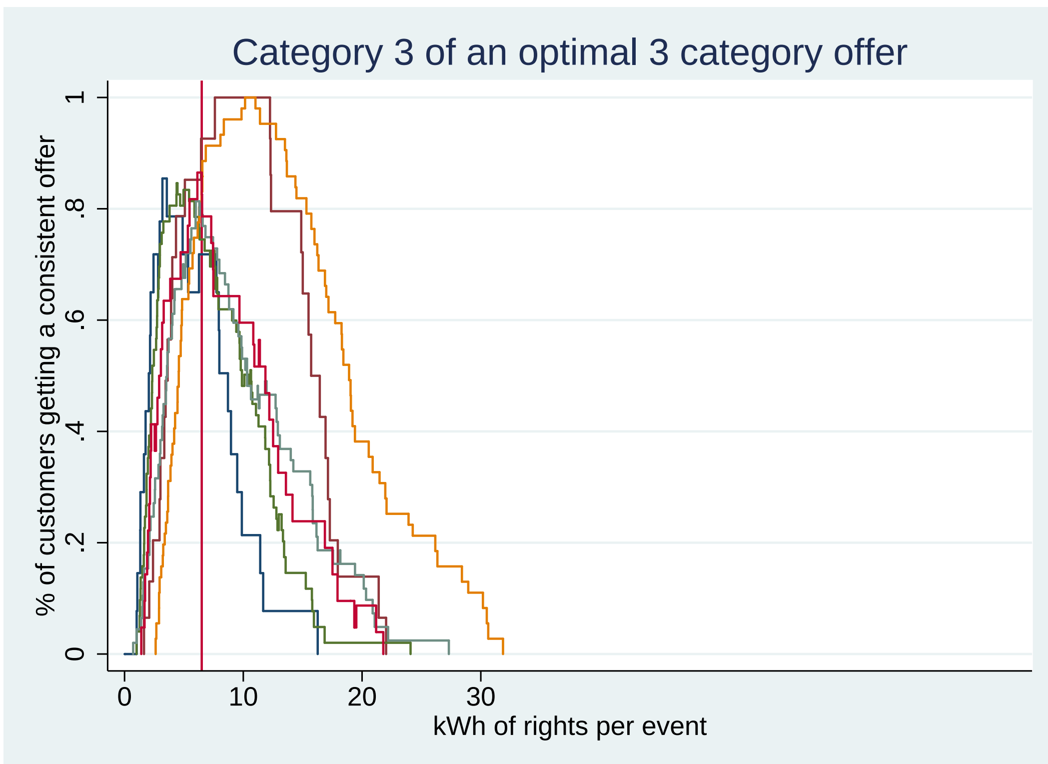

Figure 4.6: The objective step functions for the groups in category 3.

average benefits per customer in its benefits calculations. This means that the group-level percentage of potential benefits delivered is identical to the percentage of people in the group getting consistent offers.[36] The overall results reported in Figure 4.3 are a weighted average of the group level benefits.

The one and two category offers forced the algorithm to choose which groups of customers would get good offers. The algorithm's choices reflect the fact that groups in hotter climates and groups with higher use were more valuable because those groups responded more to price incentives.

- The optimal 1 category rate makes consistent offers to a large majority of the high and very-high use category customers. It performs quite poorly (probably bundling rights with too many kWh per month) for very low use customers – fewer than 20% of whom get consistent offers in any climate zone. Low use customers in zones 1 and 2 also fare poorly. Similar patterns are apparent in the optimal one-offer-per-climate zone results.

- The preferential allocation of flexibility to high use, hot climate customers is also evident in the 2 category offer. This offer's low use category serves 10 groups, while its high use category serves just 6 groups. All of the groups in the high use category make consistent offers to at least 89% of their customers. The lower use category provides consistent offers to at least 84% of members of each low and high use group, but to as few as 34% of customers in the very low use category.

- These stark differences largely disappear in the 3, 4, and 5 category offers. There is enough flexibility, enough similarity among low use groups, and enough response value in some low-use groups to justify deployments of offers that make consistent offers to more than 80% of all customers for whom the 16 optimal group level offers would be consistent.

[36]The graphs calculate the "percentage of people" by using weights that map each data point in the sample to the number of customers it represents. This process puts different weights on single family and apartment customers in each of the 16 cells.

Optimal 1 category offer					
size class	**zone 1**	**zone 2**	**zone 3**	**zone 4**	**offer**
very low	**0.0%** **Cat. 1**	**2.4%** **Cat. 1**	**19.9%** **Cat. 1**	**11.0%** **Cat. 1**	**Category 1** **19.0 kWh/event**
low	**13.9%** **Cat. 1**	**54.7%** **Cat. 1**	**94.1%** **Cat. 1**	**89.1%** **Cat. 1**	
high	**84.1%** **Cat. 1**	**97.5%** **Cat. 1**	**85.3%** **Cat. 1**	**88.8%** **Cat. 1**	
very high	**100.0%** **Cat. 1**	**86.8%** **Cat. 1**	**81.1%** **Cat. 1**	**69.7%** **Cat. 1**	

Optimal 2 category offer					
size class	**zone 1**	**zone 2**	**zone 3**	**zone 4**	**offer**
very low	**34.0%** **Cat. 2**	**68.4%** **Cat. 2**	**73.2%** **Cat. 2**	**74.3%** **Cat. 2**	**Category 2** **9.75 kWh/event**
low	**100.0%** **Cat. 2**	**98.0%** **Cat. 2**	**84.9%** **Cat. 2**	**86.8%** **Cat. 2**	**Category 1** **25.75 kWh/event**
high	**100.0%** **Cat. 2**	**87.5%** **Cat. 2**	**97.1%** **Cat. 1**	**89.2%** **Cat. 1**	
very high	**91.3%** **Cat. 1**	**100.0%** **Cat. 1**	**94.6%** **Cat. 1**	**100.0%** **Cat. 1**	

Optimal 3 category offer					
size class	**zone 1**	**zone 2**	**zone 3**	**zone 4**	**offer**
very low	**84.0%** **Cat. 3**	**88.0%** **Cat. 3**	**99.5%** **Cat. 3**	**100.0%** **Cat. 3**	**Category 3** **6.5 kWh/event**
low	**92.6%** **Cat. 3**	**88.6%** **Cat. 3**	**100.0%** **Cat. 2**	**100.0%** **Cat. 2**	**Category 2** **14.25 kWh/event**
high	**100.0%** **Cat. 2**	**92.5%** **Cat. 2**	**97.1%** **Cat. 1**	**89.2%** **Cat. 1**	**Category 1** **25.75 kWh/event**
very high	**91.3%** **Cat. 1**	**100.0%** **Cat. 1**	**94.6%** **Cat. 1**	**100.0%** **Cat. 1**	

Optimal 4 category offer					
size class	**zone 1**	**zone 2**	**zone 3**	**zone 4**	**offer**
very low	**84.0%** **Cat. 3**	**88.0%** **Cat. 3**	**99.5%** **Cat. 3**	**100.0%** **Cat. 3**	**Category 4** **27.75 kWh/event**
low	**92.6%** **Cat. 3**	**88.6%** **Cat. 3**	**100.0%** **Cat. 2**	**100.0%** **Cat. 2**	**Category 3** **6.5 kWh/event**
high	**100.0%** **Cat. 2**	**92.5%** **Cat. 2**	**94.2%** **Cat. 1**	**100.0%** **Cat. 1**	**Category 2** **14.25 kWh/event**
very high	**100.0%** **Cat. 1**	**100.0%** **Cat. 4**	**100.0%** **Cat. 4**	**100.0%** **Cat. 1**	**Category 1** **24.0 kWh/event**

Optimal 5 category offer					
size class	**zone 1**	**zone 2**	**zone 3**	**zone 4**	**offer**
very low	**84.0%** **Cat. 4**	**92.8%** **Cat. 4**	**100.0%** **Cat. 4**	**100.0%** **Cat. 4**	**Category 5** **10.25 kWh/event**
low	**100.0%** **Cat. 5**	**100.0%** **Cat. 5**	**100.0%** **Cat. 3**	**100.0%** **Cat. 3**	**Category 4** **6.25 kWh/event**
high	**100.0%** **Cat. 3**	**92.5%** **Cat. 3**	**94.2%** **Cat. 2**	**100.0%** **Cat. 2**	**Category 3** **14.25 kWh/event**
very high	**100.0%** **Cat. 2**	**100.0%** **Cat. 1**	**100.0%** **Cat. 1**	**100.0%** **Cat. 1**	**Category 2** **24.0 kWh/event**
					Category 1 **27.75 kWh/event**

One offer for climate zones 1-2 and one for zones 3-4					
size class	**zone 1**	**zone 2**	**zone 3**	**zone 4**	**offer**
very low	**9.0%** **Cat. 1**	**23.9%** **Cat. 1**	**17.4%** **Cat. 2**	**10.1%** **Cat. 2**	**Category 2** **20.25 kWh/event**
low	**79.6%** **Cat. 1**	**92.5%** **Cat. 1**	**88.7%** **Cat. 2**	**89.1%** **Cat. 2**	**Category 1** **13.75 kWh/event**
high	**100.0%** **Cat. 1**	**90.0%** **Cat. 1**	**88.3%** **Cat. 2**	**92.5%** **Cat. 2**	
very high	**91.3%** **Cat. 1**	**78.9%** **Cat. 1**	**83.8%** **Cat. 2**	**74.8%** **Cat. 2**	

One offer per climate zone					
size class	**zone 1**	**zone 2**	**zone 3**	**zone 4**	**offer**
very low	**25.0%** **Cat. 1**	**23.9%** **Cat. 2**	**19.9%** **Cat. 3**	**4.6%** **Cat. 4**	**Category 4** **22.25 kWh/event**
low	**100.0%** **Cat. 1**	**92.5%** **Cat. 2**	**94.1%** **Cat. 3**	**89.1%** **Cat. 4**	**Category 3** **19.0 kWh/event**
high	**100.0%** **Cat. 1**	**90.0%** **Cat. 2**	**85.3%** **Cat. 3**	**96.3%** **Cat. 4**	**Category 2** **13.75 kWh/event**
very high	**91.3%** **Cat. 1**	**78.9%** **Cat. 2**	**81.1%** **Cat. 3**	**84.9%** **Cat. 4**	**Category 1** **11.5 kWh/event**

4.6.3 Performance Relative to Total Number of Feasible Customers in the Group

The performance figures above describe performance as a percentage of what making 16 group level offers would achieve. Their denominators omit the customers who did not get consistent offers under the optimal group-level offer. Table 4.7.1 shows the performance of the optimal 3 category offer as a percentage of the total number of customers for whom a consistent offer exists. This is the product of the percentages reported above and the percentage of customers for whom the optimal group-level offer is consistent, given that a consistent offer exists for the customer.

There are four reasons why a customer might not get a consistent offer:

i. **No budget balanced, consistent offers exist:** it is impossible to make a consistent offer to a handful of customers given the constraints and criteria in Chapter 3. It goes on to report that budget balanced, consistent offers exist for 97% of all customers (Chapter 3). Making consistent offers to the remaining customers requires either seasonally adjusting the number of kWh that rights come bundled with or abandoning a desirable property of the offers. Even an ideal, omniscient system using the current ground rules would be unable to make consistent offers to these customers, so I do not count them in the current analysis that aims to evaluate performance relative to that ideal.

ii. **Compromise across groups:** Getting from 16 offers to 3 offers requires making some outlying customers in outlying groups offers that are not quite consistent.

iii. **Compromise within groups:** Making just 16 offers to 500 customers means that some outlying customers within each group may get offers that are not quite consistent.

iv. **Inability to predict the right offer:** Some customers are outliers in the relationship between the consistent offer range, climate zone, and previous summer usage. The budget balance constraints are close to binding for others, leaving them with a very narrow range of consistent offers. If such a customer is larger or smaller than the typical customer in their group, they might not get a consistent offer.

Table 6 in Chapter 3reports the percentage of customers in each group for whom the optimal group level offer is consistent.[37]

Table 4.7.1 is generally hopeful. The group-level offers, however, did not fit perfectly in the hotter climate zones and that making category level offers fits slightly worse. Half of all categories make consistent offers to more than 80% of the customers for whom a consistent offer exists. The other offers make consistent offers to more than two thirds of the feasible customers. Chapter 3found that inconsistent offers often deviated from consistency by well under 1% of total annual electricity bils.

4.6.4 Categorizing Customers by Usage Level Alone

The findings above suggest that a household's energy consumption level alone might be a good predictor of the appropriate offer, while the household's climate alone cannot. Table 4.6.4 suggests that the 2 and 3 category optimal assignments offers reported elsewhere appear to be consistent with assignment by energy use alone. There is, however, reason to think that customers who use air conditioning to deal with hot climates use more of their power during hot afternoons than do customers with identical average daily use levels in cooler climates. Indeed, the four and five category offers make assignments that average daily usage alone would not predict. A modest revision of the algorithm used here could find the optimal statewide categorizations by average 2002 summer daily usage. An optimal categorization by previous summer usage could outperform the sixteen category optimization if benefits from increased consumption-level resolution outweighed the loss of

[37]These calculations should be identical to multiplying the values in the tables above with the number of customers getting consistent offers from optimal group level offers in Chapter 3, Table 6. The present analysis and Chapter 3, Table 6 weight customers slightly differently and thus yield slightly different answers. Figures 4.4-4.6 use the same weighting as Chapter 3Table 6. A future revision will fix this discrepancy.

upper bounds on previous summer average kWh/day that define cells				
size class	zone 1	zone 2	zone 3	zone 4
very low	8.5	10.7	14.0	16.0
low	16.0	19.0	24.0	28.0
high	22.7	24.0	32.0	36.4

Table 4.2: Upper bounds on usage levels that define cells.

climate zone distinctions.[38],[39]

Categorizing customers by consumption alone would be modestly simpler than the current approach. It would create unified, state wide category assignment criteria rather than climate zone specific criteria. But unified statewide consumption-level categories lack the compelling advantages of climate zone-level offers. Customers can manipulate their consumption levels. Consumption levels do not designate an obvious category for new buildings or new accounts, will not always treat neighbors alike, and do not guarantee that an account will be in the same offer category from year to year. IP rebate implementations will have to address these modest problems regardless of whether they use consumption alone or in combination with other criteria.

4.6.5 Policy Implications of Inconsistent Offers

It is important that policy makers have an accurate appraisal of likely performance and that they do market research to see how customers react to inconsistent offers given that realistic offers will expose a non-trivial minority of customers to inconsistent offers.[40]

Careful implementation could reduce damage from inconsistent offers. Creating a modest reserve fund early in the program's fiscal year (or early in the lifetime of an account) could reduce the number of bill surprises by covering the shortfalls that make offers inconsistent. Chapter 3finds that most offer inconsistencies are small. Thus, there may be opportunities to design bills that emphasize important things like opportunities to

[38]Insofar as implementation concerns rule out the possibility of using finer grained distinctions, pointing out their potential benefits of more precise division points is not particularly constructive or interesting.

[39]California's increasing block rates already apply usage-level tiers to categorize customers' monthly consumption. Tiers vary geographically. It might be reasonable to assign customers to offer categories using the existing tier system. Further analysis could explore the effectiveness of a variety of functions that map each customer's customer monthly usage tier time series to an offer category.

[40]All of the optimization in this project assumes that correcting for undercontribution disturbs customers exactly as much as exposure to high prices on the margin. A more nuanced understanding of customer preferences from market research might suggest adjustments to the optimization.

reduce bills by conserving during high priced periods, while playing down the deviations from the normal sources of charges.

Appendix M presents these results for other categories.

4.7 These Optimization Challenges Compared to Other Dynamic Rates

4.7.1 Optimization requirements of a variety of rate designs

An IP rebate implementation in California would probably use customer energy consumption or a proxy for it to assign customers to categories. It would have to be comfortable with some customers getting offers that are not quite consistent. This finding places IP rebate rates' implementation challenges between critical peak and real time pricing and baseline rebate rates. Real time pricing requires no customer-specific optimization, but is often so complex and exposes customers to enough bill-spikes that customers find it unattractive. CPP's challenges are similar but CPP would benefit from optimization to develop a rate that approximates spot energy prices well.

Conventional CPP and real time pricing rates offer implementers a paucity of tools to address customers' decision making heuristics like reference-dependent loss aversion.[41] Both IP rebates and baseline-rebate rates offer potentially important tools to reframe high prices as gains, but require some tailoring of offers to meet customer characteristics. CPP with IP rebates adds an additional layer of customer-class specific optimization challenges to deliver consistent offers. The offers affect how well the rate attracts customers and how likely it is to work as promised for those customers. Deviations from optimal behavior do not affect incentives or revenue streams. Baseline-rebate rates pose analogous optimization problems that affect incentives and revenue streams. IP rebate rates can solve these problems at the customer category level, while baseline-rebate rates almost always customize offers for each customer.

Baseline-rebate rates have been repeatedly deployed, but they struggle with an optimization problem that is analogous to the IP rebate optimization, but is both fundamentally harder and involves higher stakes. Table 4.7.1 compares the IP rebate and

[41] CPP designers might explore whether customers have loss averse reactions to high, evening rates. A conventional CPP rate could end its TOU peak price period early in the evening to address this loss aversion.

	avoid offering too few rights because	avoid offering too many rights because
IP rebate	expose customers to high marginal prices	customer may be unable to buy all offered rights
baseline-rebate	breaks marginal incentive to save	means that the customer gets an automatic bill reduction payed for through a cross subsidy from other customers or a rate increase

Table 4.3: Comparing the optimization problems posed by baseline-rebate and IP rebate rates.

baseline-rebate optimization problems. Both baseline rebate and IP rebate rates seek to offer each customer a level of rights that is neither too small nor too big. The ideal baseline rights level would be the customers' use level in the counterfactual in which the were on conventional, time-invariant pricing. A smaller baseline would give customers too weak an incentive to conserve power. Baselines above this counterfactual level pay customers "structural rebates" for not using power that the customers were never would have used. Other customers often end up cross subsidizing customers who get structural rebates. The customer's counterfactual usage is unobserved. Analysis of billing time series data shows that factors like weather fluctuations and the installation of new appliances cause fluctuations in counterfactual use. Chapter 3describes how trying to use customers' own behavior to calculate a baseline can create perverse incentives and is an example of a class of asymmetric information games that have no first-best solutions.

While the optimal baseline-rebate offer is the exact counterfactual, the IP rebate optimum offer is generally any choice from a broad range of analogous, consistent offers.[42] The existence of a broad range of consistent offers makes the optimization easier than those that baseline rebate rates pose, but the analysis here finds that even the simplified problem poses some significant challenges. Further, the IP rebate stakes are lower because, unlike a poor choice of baseline, IP rebate offer inconsistency does not affect firm revenues, total annual bills, or incentives. A significant number of baseline-rebate programs have been deployed despite the difficult and more important optimization problems they face. IP rebates appear to provide practitioners with attractive, practical tools to help customers who

[42]Chapter 3does some analysis that observes that the most desirable offers are not only consistent under the most likely scenario, but are also robust to a variety of foreseeable deviations from that scenario. So a more accurate statement is that IP rebates create a range of offers that will be optimal with high probability.

Optimal 3 category offer				
size class	**zone 1**	**zone 2**	**zone 3**	**zone 4**
very low	**67.2%**	**71.0%**	**77.5%**	**79.0%**
low	**92.6%**	**88.6%**	**76.5%**	**76.5%**
high	**100.0%**	**90.2%**	**76.8%**	**86.0%**
very high	**91.3%**	**97.4%**	**89.7%**	**79.8%**

use decision making heuristics like reference-dependent loss aversion make good decisions about participating in dynamic pricing.

4.8 Conclusion

The IP Rebate optimization problem has characteristics that facilitate its implementation.

- Assigning all customers statewide to three or more categories can perform very well. Optimally aggregating 16 groups of customers into three categories performs nearly as well as making each of the 16 groups an optimal offer.
- Utilities already have the geographic and billing data to develop these categories.
- The group level objective functions that report the number of customers who would get consistent offers at each level of offers are fairly similar within each of the three optimal categories. This means that there is likely to be broad agreement about the optimal offers among interest groups that have divergent priorities for the program and divergent attitudes about whether to target groups of customers.
- Most of the group-level objective functions have large, fairly flat regions around the optima.[43] Thus, implementers have flexibility to address implementation concerns like offering rights in round numbers of kWh or having the offers correspond to usage thresholds in existing rates.
- These performance statistics come from the majority of the state of California which has diverse climates, housing stock, and socio-economic conditions. Many utilities,

[43] The significant change from the three category offer to the two category offer only leads to a 6 percentage point change in overall performance. This suggests that the category-level objective functions have fairly flat peaks as well.

including two of the three that participated in the California SPP, have far more homogeneous customer bases than does California as a whole.[44]

IP Rebate deployments cannot, however, be as simple as we might like. Achieving high performance appears to require at least three categories that divide customers by usage or a proxy for usage. Every IP rebate offer provides the right incentives and the right total annual bills. Even the best offers that have no seasonal variation in contributions, however, expose a minority of customers to offers that include high marginal payments or corrections for under purchasing of rights.

Thus, answers to hard questions about the implications of inconsistent offers would be useful. There is little evidence about whether customers who get inconsistent offers will exit. It is unclear under what circumstances customers might prefer complex provisions that make offers more likely to be consistent to exposure to inconsistent offers. These are crucial open questions. We can hope that even customers who get occasional, modestly inconsistent offers will exhibit the fairly low customer attrition rates seen in conventional dynamic pricing programs. If, however, customers who get inconsistent offers tend to exit even when they are coming out ahead under dynamic pricing, it is worth investigating the effectiveness of bills that draw attention to information about total costs and consequences rather than offer inconsistencies.

IP rebates rates, unlike conventional CPP and real time pricing, have features that address customer loss aversion that require some offer customization. IP rebates pose an offer customization problem that is both easier and far less important than the optimization problem that baseline-rebate rates pose. The IP rebate design means that every customer faces the incentives and pays the total annual bill that the rate designers intended. Between 67% and 100% of customers in each category will get rates with features designed to address loss aversion that work in a desirable way. Thus, IP rebates offer every customer the right incentives and a large majority of customers reframing of high prices as opportunities that works as intended.

[44]Only PG&E has customers in climate zone 1. SDG&E has neither a fog belt nor a desert region.

Chapter 5

Conclusion

Demand for air conditioning on the hottest summer weekday afternoons drives extreme electricity demand in many electricity markets. Meeting this extreme demand is very expensive. It is especially expensive to deal with when it coincides with problems in the electricity supply system that exacerbate the scarcity. Dynamic pricing has been proposed as a way to manage these costs.

These essays fill in important aspects of the picture of the implications of opt-in, residential critical peak pricing (CPP), and the challenge of making such a program work well. Taken together, they have several important policy implications that cut across the chapters:

5.0.1 Customer Heterogeneity is Important

- It is propitious, but perhaps unsurprising, that dynamic pricing has its greatest benefits on hot weekday afternoons for customers in climates hot enough to justify air conditioning. It estimates that the benefits of dynamic pricing range from zero in cooler climates on cooler days to .3 (.4) kW every hour for ordinary afternoon TOU peak (extreme scarcity, "critical peak") prices on the hottest days in hot climates. The notion of a typical customer on a typical day turns out to be fairly unhelpful in a place with climates and buildings as diverse as California.

- Using at least three different rate makes consistent offers to far more customers than does a one-size-fits-all approach. Making the optimal single statewide offer performs only about 74% as well as does making the optimal offer to each of the 16 groups of customers. Optimally assigning the 16 groups to 3 categories and then making the optimal offer to each of those categories performs 96% as well as does making an optimal offer to each of the 16 groups. In this context, heterogeneity within climate zones appears more important than heterogeneity between similarly-sized customers across climate zones. Making an optimal offer to each climate zone performs 78% as well as does making 16 optimal offers.

5.0.2 Recruiting and Retaining the Right Customers is a Significant Challenge

- A variety of customer decision making heuristics could drive resistance to signing up for dynamic pricing. There is reason to think that these heuristics cause a significant number of customers to decline to participate despite the fact that they would be happy and save money on the program. Incentive preserving rebates are an approach to work around these heuristics that has desirable economic properties. At least in this case, rates can incorporate desirable microeconomic incentives and be presented in a way that works around troubling heuristics.

- Neoclassical explanations for resistance are also important. The customers who are most able to respond are also often the customers who would have to give up the greatest cross subsidies to participate. A rate that is revenue neutral for the statewide average load profile, who used an average of 22.3% of their summer-season power during peak hours, would increase bills for more than 55% (45%) of treatment-group customers in climate zone 4 (3) even after they shifted their usage patterns in response to the new prices. Developing a rate that charges revenue-neutral prices within each climate zone, however, would reduce this problem and leave only between 25% and 40% (40% and 45%) of treatment group customers coming out behind despite shifting load in zone 3 (in zone 4).

- In the longer term, there is evidence that customers who experienced bill increases or who were frustrated with critical events left the program. Thus, retaining customers appears to be a separate challenge from attracting customers, although retaining customers appears to be an easier challenge.

There are significant open questions here:

- This dissertation shows that it is possible to design incentive preserving rebates that have desirable economic properties, while addressing decision making heuristics that could be driving customer resistance to signing up for CPP. This dissertation presents significant evidence from the literature, from focus groups, and from field experience that suggests that these heuristics may be driving resistance, but does not directly explore customers' thought processes to confirm the origin of resistance. Future research should explore whether incentive preserving rebates are an effective way to help customers make better choices. In particular, additional evidence would be useful about how people react to efforts that address loss aversion with a more complex reframing.

- The SPP generated a data set about customer decisions whether to continue on dynamic pricing. Differences between customers on high and low ratio rates and differences in bill spikes might allow future analysis to explore the behavioral and neoclassical economic explanations of attrition. It could also explore the demographic correlates of attrition and whether the program was keeping its most responsive customers.

5.0.3 Targeting may offer compelling benefits, but is not required

Pricing programs will be judged, at least in large part, by their ability to reduce peak period demand, especially during scarcity periods. Some customers respond far more than others – and thus provide far more social benefit. This project suggests that carefully designed programs can address inter-regional cross subsidies by designing rates that are revenue neutral for the average consumption profile within each region. Simple menus of IP rebate offers can perform well for large and small customers in a variety of climates. Indeed, we get good performance for low benefit customers even when the algorithm is tuned to maximize the expected total kWh response, which puts roughly 50 times more weight per customer on making a consistent offer to the most responsive kind of customer than to the least. Thus, deploying 3 or more offers that use customer-consumption levels do not appear to require selecting which customers would get offers that are consistent. A scheme of climate-specific offers can maintain cross subsidies. By contrast, a one-size-fits-all offer, a two statewide offer, or a one offer per climate zone offer all force designers to choose which customers will get inconsistent offers.

While targeting may be optional, addressing heterogeneity is not. A one-size-fits-all approach that makes a single dynamic pricing offer in a place with climates and air conditioning adoption rates as diverse as those in California will struggle with the fact that the most responsive customers will have to give up cross subsidies to participate and, if it uses IP rebates, is likely to make many customers inconsistent offers. Many utilities have much less diverse customer bases than does the state of California as a whole and may find that they can achieve high performance with a single rate. Thus, policies that understand customer diversity and, as appropriate, address differences among customers are quite important.

It is up to the policy designers to decide whether to take this one step further to target the most responsive customers with extra marketing and consumer education activities, IP rebate offers that are more likely to be consistent, or even incentives to try dynamic pricing or rates that increase the percentage of customers from highly responsive categories who would save money if they exhibited at least modest response.

Factors that might lead to mixed views of the wisdom of targeting include:

- Equity which suggests that all customers should be treated the same. This may suggest against targeting, but also urges policy makers to get prices as close to cost as possible to stop subsidizing people who habitually run big air conditioners during scarcity periods.

- Planning for growth. Limiting an initial deployment to a small category of customers or concentrating its resources too narrowly may leave the program without evidence that it can work for other kinds of customers. Viewing initial deployments as steps in a transition toward improving pricing for most consumers suggest both including customers who represent the set of consumers as a whole, and using techniques like targeting to ensure that early programs deliver significant social benefits.

- There are deadweight losses from exposing any price-elastic customer to prices that diverge from marginal cost, so improving prices for any customer should offer social benefits. The research here, suggests, however that some customers deliver far more social benefits than others.

One crucial open question here is to approximate the set of Pareto-improving deals. Such an exercise would let us understand how much surplus putting each type of customer on dynamic pricing creates, so we could decide how to allocate those benefits among the responsive-customer, the utility, and other customers. If putting a certain class of customers on dynamic pricing saved \$50 per year in the fixed cost of building and maintaining peaking capacity beyond the marginal costs incorporated into their rates, we could imagine a rate that "targets" these customers with an offer of \$40 to sign up and still leaves \$10 in benefits to split between other customers and the utility.[1]

Another crucial unknown is the cost and effectiveness of recruiting. If fixed-cost efforts like designing and testing marketing materials are central in determining effectiveness and if there is little need to customize for different regions and types of customers, then there is relatively little reason to target. If recruiting has large marginal costs, then targeting a scarce recruiting budget at the most responsive customers makes sense. This might happen if the firm faced high costs to contact customers, to offer them incentives to participate, to install new meters, and to answer new customers' questions. If it were feasible, it might be particularly compelling for the regulator to offer the firm a reasonably good set of incentives and give the utility latitude to set and allocate marketing and sign-up incentive budgets, as well as to incrementally refine some aspects of the program and its marketing materials.[2]

5.0.4 Utilities have enough information to make offers and predict response. Good Offers Require Some but Limited Flexibility

Climate zone and a customer's overall energy consumption level are powerful predictors of the appropriate offer for a customer. These factors plus the temperature on a particular day are powerful predictors of customers' response. Utilities know where their customers live and have billing time series data for all but the newest accounts and buildings.

A menu of offers that makes small, medium, and large offers and differentiates among customers by historical use and geography can perform quite well.

Incentive preserving rebates have the potential to be part of a carefully designed residential opt-in dynamic pricing program that delivers significant benefits and that paves the way for future expansion of dynamic pricing.

[1] Before we rush to the conclusion that we can increase participation rates with modest amounts of cash to sign up, it is worth noting that the statewide pricing pilot used exactly this kind of scheme. Market researchers talked to customers who refused to sign up and discovered that many suspicious customers were ready to interpret any increase in the cash payment as being compensation for a larger expected bill increase.

[2] Getting incentives right in this arena is a difficult, important question that is well beyond the scope of this dissertation.

Bibliography

WJ Adams and JL Yellen. Commodity bundling and the burden of monopoly. *Quarterly Journal of Economics*, 1976.

Ameren. Service classification no. 1(m) residential service rate. Missouri Public Service Commission, Schedule Number 5, Sheet Number 28, 34th revised, August 2002. https://www2.ameren.com/ACMSContent/Rates/Rates_umbe28rt1M.pdf.

Galen Barbose, Charles Goldman, and Bernie Neenan. A survey of utility experience with real time pricing. Technical report, Lawrence Berkeley National Laboratory, http://eetd.lbl.gov/ea/ems/reports/54238.pdf, December 2004.

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), 2005a.

Severin Borenstein. Time-varying retail electricity prices: Theory and practice. In Griffin and Puller, editors, *Electricity Deregulation: Choices and Challenges*. University of Chicago Press, 2005b.

Severin Borenstein. Wealth transfers from implementing real-time retail electricity pricing. *Center for the Study of Energy Markets Working Paper 156* *http://www.ucei.berkeley.edu/PDF/csemwp156.pdf*, July 2006. University of California Energy Institute.

Severin Borenstein. Customer risk from real-time retail electricity pricing: Bill volatility and hedgability. *The Energy Journal*, April 2007. Also: Center for the Study of Energy Markets Working Paper 155 http://www.ucei.berkeley.edu/PDF/csemwp155.pdf.

Severin Borenstein and Stephen Holland. On the efficiency of competitive electricity markets with time-invariant retail prices. *Rand Journal of Economics*, 36, Autumn 2005.

Stephen J. Brown and David S. Sibley. *The theory of public utility pricing*. Cambridge University Press, 1986.

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.

Colin Camerer, Samuel Issacharoff, George Loewenstein, Ted O'Donoghue, and Matthew Rabin. Regulation for conservatives: Behavioral economics and the case for 'asymmetric paternalism'. *University of Pennsylvania Law Review*, 151:1211–1254, 2003.

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.

www.energy.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://www.energy.ca.gov/demandresponse/documents/group3_final_reports/2005-03-24_SPP_APPENDICES.PDF.

Gary Charness and Matthew Rabin. Understanding social preferences with simple tests. *Quarterly Journal of Economics*, 2002.

Yan Chen. *The Handbook of Experimental Economics Results*, chapter Incentive compatible mechanisms for pure public goods: A survey of experimental research. Amsterdam, Forthcoming.

JJ Choi, D Laibson, B Madrian, and A Metrick. Optimal defaults. *American Economic Review*, 2003.

John J. Conti et al., editors. *Annual Energy Outlook 2006.* Energy Information Administration, 2006. "Conv. Combustion Turbineëntry from Table 38. Cost and Performance Characteristtics of New Central Station Electricity Generationg Technologies http://www.eia.doe.gov/oiaf/aeo/assumption/pdf/electricity_tables.pdf.

Ravi Dhar. Consumer preference for a no-choice option. *The Journal of Consumer Research*, 1997.

Alexis Diamond and Jasjeet S. Sekhon. "genetic matching for estimating causal effects: A general multivariate matching method for achieving balance in observational studies". *Institute of Governmental Studies. Paper WP2006-35.*, September 18 2006. http://repositories.cdlib.org/igs/WP2006-35.

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.

Mark F. Gaines. *Application of San Diego Gas & Electric Company (U-902-E) for Adoption of an Advanced Metering Infrastructure Deployment Scenario and Associated Cost Recovery and Rate Design: Chapter 24 Rebuttal Testimony.* San Diego Gas & Electric Company, 2006. http://www.sdge.com/regulatory/tariff/05_03_015_RT_Gaines.PDF.

Uri Gneezy and Jan Potters. An experiment on risk taking and evaluation periods. *The Quarterly Journal of Economics*, 1997.

Uri Gneezy, Arie Kapteyn, and Jan Potters. Evaluation periods and asset prices in a market experiment. *Journal of Finance*, 58(2):821–838, 2003.

Gulf Power. Goodcents select. Gulf Power letter and brochure. Pensacola, FL., 2005. request form available online at: http://www.southerncompany.com/gulfpower/residential/requestinfo.asp?mnuOpco=gulf&mnuType=res&mnuItem=ps.

Karl Hausker. The politics and economics of auction design in the market for sulfur dioxide pollution. *Journal of Policy Analysis and Management*, 11(4):553–572, 1992.

Karl Hausker. *Polar Ramsey Pricing: a third-best approach to efficient pricing*. PhD thesis, Goldman School of Public Policy, UC Berkeley, 1986.

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.

Teck-Hua Ho and Juanjuan Zhang. Does the format of pricing contracts matter? *XLab Working Paper XL05-002 University of California, Berkeley*, September 28 2004.

Stephen P. Holland and Erin T. Mansur. The short-run effects of time-varying prices in competitive electricity markets. *Center for Study of Energy Markets Working Paper 143R*, 2006.

Stephen P. Holland and Erin T. Mansur. Is real-time pricing green?: The environmental impacts of electricity demand variance. *Center for Study of Energy Markets Working Paper 136*, 2005.

SS Iyengar and MR Lepper. When choice is demotivating: Can one desire too much of a good thing? *Journal of Personality and Social Psychology*, 2000.

A K Jain and M N Murty. Data clustering: A review. *http://citeseer.ist.psu.edu/jain99data.html*, 1999.

Paul Joskow. Deregulation and regulatory reform in the US electric power sector. In S Peltzman and C Winston, editors, *Deregulation of Network Industries: What's Next?*, chapter 4, pages 113–188. AEI-Brookings Joint Center for Regulatory Studies, 2000.

Daniel Kahneman and Amos Tversky. Prospect theory: An analysis of decision under risk. *Econometrica*, 47(2):263–292, March 1979. Reprinted in Choices, Values, and Frames 2000.

Daniel Kahneman, Jack Knetsch, and Richard Thaler. Fairness as a constraint on profit seeking: Entitlements in the market. *The American Economic Review*, 1986.

Chris King and Patti Harper-Slaboszewicz. Pepco smart meter pilot program inc. focus group. Audio Tape, March 30 2006. Customers from Washington, DC.

Botond Koszegi and Matthew Rabin. A model of reference-dependent preferences. *Quarterly Journal of Economics*, forthcoming. Revised July 2005. available at: http://emlab.berkeley.edu/users/botond/refdep.pdf.

David Lineweber. Post-spp rate choice assessment research: Final report. *Momentum Market Intelligence*, June 2005.

George Loewenstein, Ted O'Donoghue, and Matthew Rabin. Projection bias in predicting future utility. *The Quarterly Journal of Economics*, 118(4):1209–1248, 2003.

LL Lopes. *Advances in experimental social psychology*, chapter Between hope and fear: the psychology of risk. Academic Press, 1987.

JG March and Z Shapira. Managerial perspectives on risk and risk taking. *Management Science*, 1987.

Momentum Market Intelligence. *Customer Preferences Market Research (CPMR): A Market Assessment of Time-Differentiated Rates Among Residential Customers in California*. Momentum Market Intelligence, December 2003. http://www.energy.ca.gov/demandresponse/documents/group3_final_reports/2004-10-29_CPMR_RES_REPORT.PDF.

A. Yesim Orhun. *Optimal Product Line Design When Consumers Exhibit Choice-Set-Dependent Preferences*. PhD thesis, Haas School of Business, UC Berkeley, 2006.

Pacific Gas & Electric. Welcome package. Pamphlet Mailed to Statewide Pricing Pilot Customers on the CPP Low Ratio Rate in the PG&E Service Territory, June 2006a. filename: "06.06.03 6.9 CPP-F PG&E_Rate B_Welcome Package fmt.pdf".

Pacific Gas & Electric. Request for expedited implementation of PG&E's natural gas 2005-2006 winter gas savings program. California Public Utilities Commission Advice Letter 2675-G, November 3 2005b. http://www.pge.com/tariffs/advice/adviceletters/2675-G.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 2003c.

Pacific Gas & Electric News Department. Pacific Gas and Electric Company's SmartMeter proposal approved by California Public Utilities Commission. Press Release, July 20 2006. http://www.pge.com/news/news_releases/q3_2006/060720a.html.

Pepco. District of columbia residential service schedule r updated billing month of july 2006. http://www.pepco.com/_res/documents/dc_schedule-r-2006.pdf, July 2006.

Steve Quinn. Hot temperatures, power demand force blackouts. *Dallas Morning News*, April 18 2006. http://www.dallasnews.com/sharedcontent/APStories/stories/D8H281MG9.html.

Matthew Rabin. Diminishing marginal utility of wealth cannot explain risk aversion. In Daniel Kahneman and Amos Tversky, editors, *Choices Values and Frames*, pages 202–208. Russell Sage, 2000.

Matthew Rabin. Incorporating fairness into game theory and economics. *American Economic Review*, 83(5):1281–1302, 1993.

Daniel Read, George Loewenstein, and Matthew Rabin. Choice bracketing. *Journal of Risk and Uncertainty*, 19(1-3):171–97, 1999. available at http://ideas.repec.org/a/kap/jrisku/v19y1999i1-3p171-97.html.

Joseph P. Redden and Stephen J. Hoch. The psychology of two-part tariffs. *Wharton School Working Paper*, January 2005.

CS ReVelle and G. LaPorte. The plant location problem: New models and research prospects. *Operations Research*, 1996.

Dorothy Robyn. *Braking the Special Interests: Trucking Deregulation and the Politics of Policy Reform*. University of Chicago Press, Chicago, 1987.

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.

San Diego Gas and Electric. *Schedule DR: Residential Service*. San Diego Gas and Electric, October 5 2006. CPUP Sheet number: 19621-E http://www.sdge.com/tm2/pdf/DR.pdf.

Dean Schultz and David Lineweber. Real mass market customers react to real time-differentiated rates: What choices do they make and why? *Association of Energy Services Professionals Conference*, 2006.

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.

Debra Stone. *Policy paradox: the art of political decision making*. Norton, 1997.

Cass R. Sunstein and Richard H. Thaler. Libertarian paternalism is not an oxymoron. *University of Chicago Law Review*, 2003. http://ssrn.com/abstract=405940.

Richard H Thaler. Mental accounting matters. *Journal of Behavioral Decision Making*, 12 (3):183 – 206, 1999.

Richard H. Thaler and Shlomo Benartzi. Save more tomorrow: Using behavioral economics to increase employee saving. *Journal of Political Economy*, 112:S164–S187, 2004.

Richard H Thaler and et. al. The effect of myopia and loss aversion on risk taking: An experimental test. *The Quarterly Journal of Economics*, 112(2):647–61, 1997.

R Velle. Facility siting and integer-friendly programming. *European Journal of Operational Research*, 1993.

D Verma and M Meila. A comparison of spectral clustering algorithms. *http://citeseer.ist.psu.edu/587964.html*, 2003.

Richard Voytas. AmerenUE critical peak pricing pilot. Presentation to the National Demand Response Town Meeting http://drrc.lbl.gov/pubs/drtown-pricing-voytas.pdf, June 26 2006.

Brian White. Gulf power. Personal Interview, 2006.

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

Anthony C. Wilson. Application of the Potomac Electric Power Company for approval of the tariff and meter to be used in a small customer smart meter project. Formal Case 1002 Public Service Commission of the District of Columbia, June 2006.

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.

Laurence A. Wolsey. *Integer Programming.* 1998.

Lisa Wood. The new vanilla: Why making time-of-use the default rate for residential customers makes sense. *Fortnightly's Energy Customer Management*, July/August 2002a.

Lisa Wood. Effective demand response programs for mass market customers. Presented at the NYSERDA Time Sensitive Electricity Pricing Workshop. Albany NY, October 3 2002b.

Paul Wright et al. Demand-responsive electrical appliance manager website. http://dr.berkeley.edu/dream/index.htm.

Appendix A

Customer Recruitment Process

Figure A.1: The recruiting process that generated the control and CPP groups.

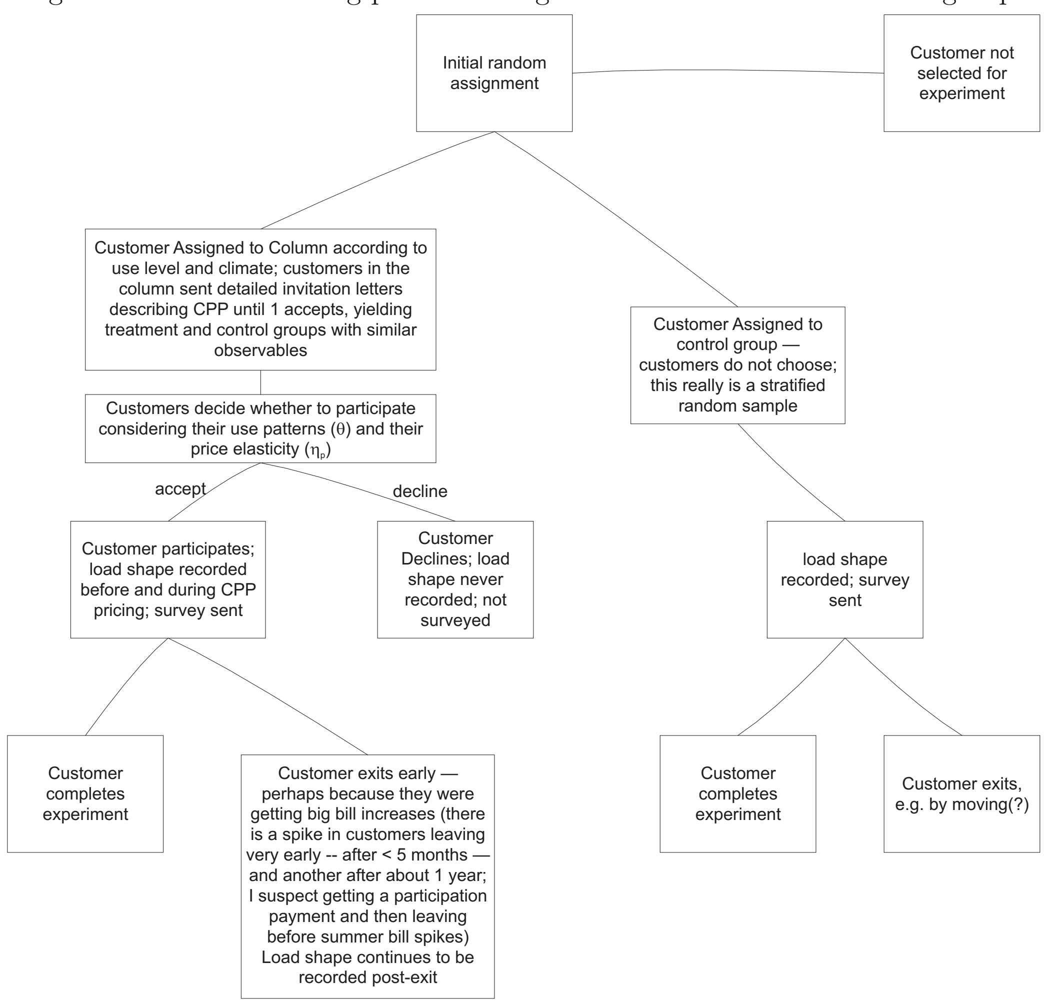

Appendix B

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

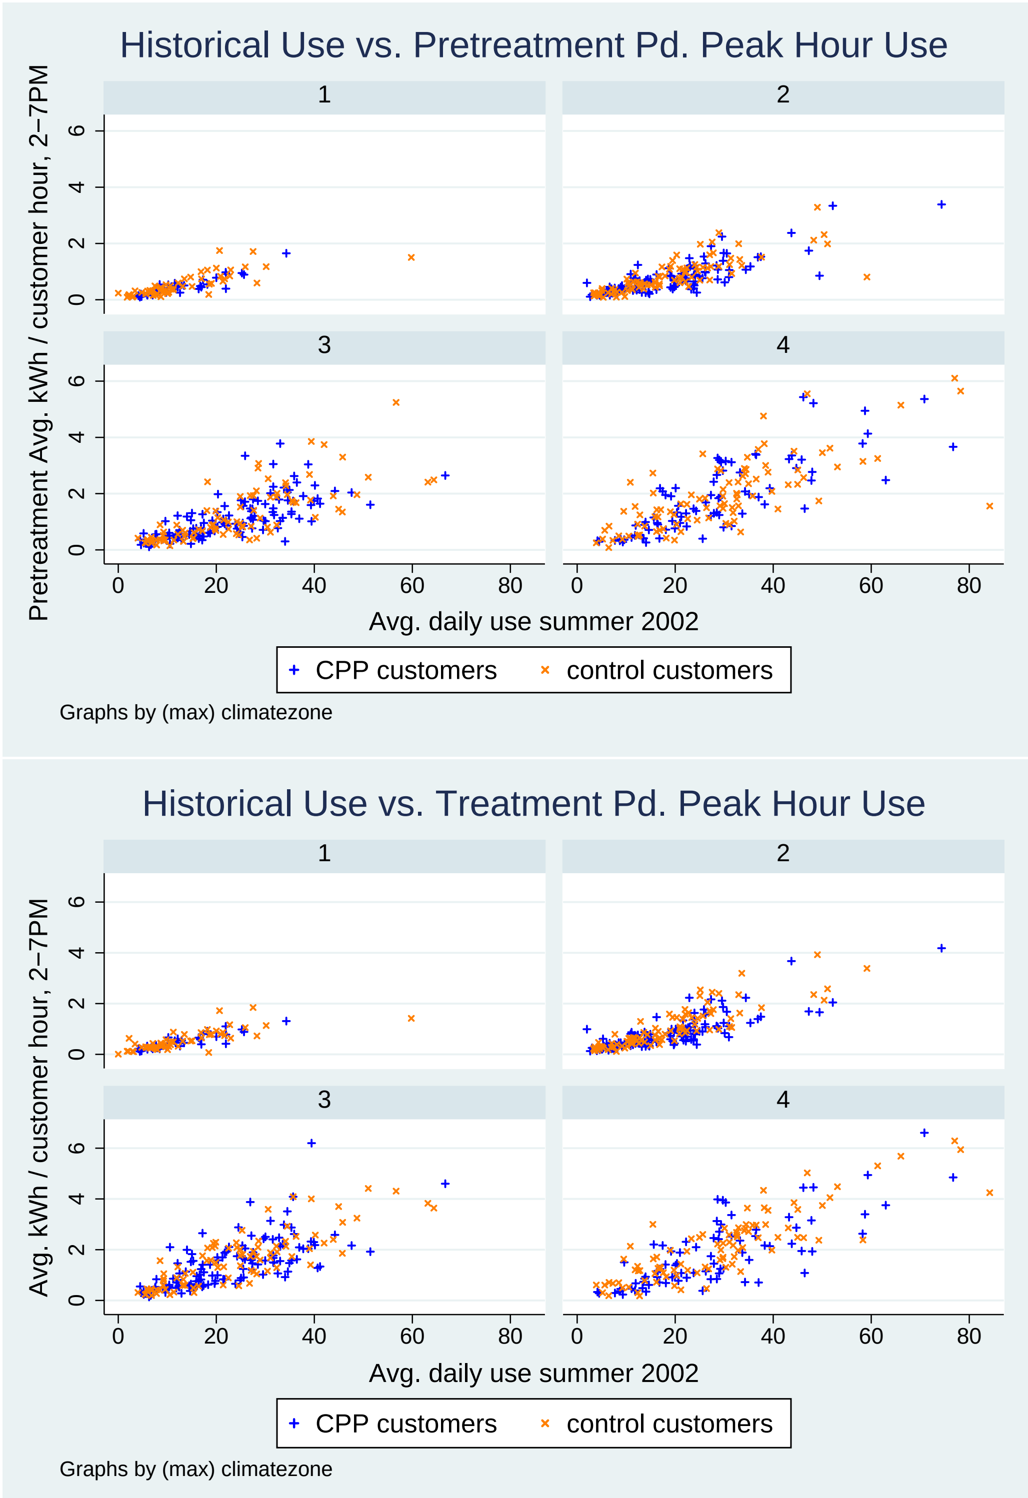

Figure B.1: This graph shows raw average daily use plotted against raw average daily afternoon use before and after the experiment began. This illustrates the identification strategy for the relationship between customer historical usage level and customer use during weekday peak hours. The simple, graphical results are not nearly as striking as those that in figures 2.4, 2.3, and 2.2 which support the regression finding that dynamic pricing has greater impact on hotter days.

Appendix C

Means of the whole sample for comparison to other papers

The selection problems described in section 2.2.4 extend to other slices of the data, including those used in the previous papers that use the same data set like Faruqui and George (2005); Charles River Associates (c); Herter et al. (2007); Herter (2006b,a) and Chapter 3. I illustrate this point by reporting the sample averages for every customer who reports any data and every customer who has at least 105 days of participation in the experiment. A substantial number of customers completed the survey, but have no detailed electricity use data. Some subjects who did participate in the experiment never completed the survey. This problem is more serious among control customers. It is especially serious among control customers who use a fairly small amount of power.

	whole sample			>4 months of usage data		
	Control Gp.	CPP Gp.	P-value	Control Gp.	CPP Gp.	P-value
number of customers with at least one of these variables	778	1,146		331	355	
avg. daily use, kWh, summer 2002	22.6	21.8	0.420	22.9	21.5	0.230
avg. use, kWh, weekdays 2-7PM, June 1-15 ’03	5.89	5.25	0.083	6.02	5.1	0.015
avg. daily use offpeak usage, kWh, June 1-15 ’03	13.7	13.6	0.933	13.8	13.5	0.673
avg. 4PM temperature, June 1-15 ’03	77.3	77.3	0.967	77.4	77.2	0.756

	whole sample			>4 months of usage data		
	Control Gp.	CPP Gp.	P-value	Control Gp.	CPP Gp.	P-value
# children 0 to 5	.387	.299	0.078	.396	.311	0.148
# children 6-18	.702	.701	0.984	.714	.695	0.829
# people over 65	.298	.301	0.950	.292	.323	0.559
everyone in household is > 65	.094	.109	0.460	.093	.117	0.333
home built after 1979	.44	.479	0.244	.455	.442	0.753
% work from home part/full time	.164	.172	0.752	.187	.132	0.061
agrees w/ "everyone should pay a little ...[for] a cleaner environment"	.525	.649	0.000	.504	.666	0.000
agrees that "a cleaner environment will mean fewer jobs"	.2	.231	0.269	.224	.241	0.638
agree/strongly agree that 'global warming is a threat...'	.668	.664	0.910	.66	.634	0.513
1=rates utility performance good or excellent	.781	.786	0.857	.76	.797	0.261
household head is a college graduate	.439	.468	0.370	.433	.434	0.991
has central air conditioning	.517	.531	0.673	.544	.519	0.544
has 1+ room air conditioners	.139	.164	0.321	.143	.152	0.766
electric well pump	.059	.054	0.749	.056	.041	0.376
# refrigerators + freezers	1.45	1.45	0.946	1.46	1.46	0.958
electric hot water	.127	.125	0.937	.107	.11	0.925
electric range	.383	.375	0.818	.376	.317	0.118
electric oven	.483	.504	0.572	.467	.445	0.611
electric dryer	.378	.416	0.237	.377	.362	0.692
programmable thermostat for Central AC	.246	.292	0.115	.272	.286	0.710
swimming pool	.133	.14	0.768	.142	.121	0.451
electric spa	.066	.071	0.773	.075	.086	0.644

	whole sample			>4 months of usage data		
	Control Gp.	CPP Gp.	P-value	Control Gp.	CPP Gp.	P-value
number of customers contacted before one accepted	2.33	4.14	0.000	1.19	3.16	0.000
total annual household income	69977	65464	0.112^m	70852	62400	0.026^m
sq. feet of living space	1652.4	1658.4	0.891^m	1681.3	1621.6	0.371^m
significance: *=10% ** =5% ***=1%						

Appendix D

CPP Impacts: Complete regression results

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 FE's; 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 before critical price	0.001 (0.012)	0.001 (0.013)	0.003 (0.013)	0.001 (0.013)
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 * apartment	-0.050	0.010	-0.002	-0.013

	Specification 1: Simplest Diff in Diff	Specification 2: Adding Survey Variables	Specification 3: Adding CAC*CDH interactions	Specification 4: Adds person FE's; controls
	(0.062)	(0.087)	(0.087)	(0.104)
TOU Peak Price in Effect * climate zone 2	-0.027 (0.051)	-0.060 (0.062)	-0.065 (0.060)	-0.042 (0.074)
TOU Peak Price in Effect * climate zone 3	-0.063 (0.072)	-0.065 (0.092)	-0.070 (0.093)	-0.084 (0.110)
TOU Peak Price in Effect * climate zone 4	-0.228* (0.131)	-0.174 (0.154)	-0.138 (0.149)	-0.096 (0.176)
TOU Peak Price in Effect * cooling degree hours 2-7pm	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), 2-7pm	-0.102*** (0.033)	-0.078** (0.035)	-0.013 (0.135)	-0.106*** (0.038)
TOU Peak Price in Effect * heating degree hours 2-7pm	-0.00010 (0.002)	-0.000049 (0.002)	-0.00035 (0.002)	-0.007 (0.006)
TOU Peak Price in Effect * heating degree hours 2-7pm squared (1000's)	. .	. .	. .	0.141 (0.123)
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 * number of bedrooms	. .	0.059 (0.041)	0.050 (0.039)	0.043 (0.043)
TOU Peak Price in Effect * # people in the household	. .	0.010 (0.022)	0.010 (0.022)	0.049 (0.038)
TOU Peak Price in Effect * cooling degree hours 2-7pm * central AC	. .	. .	-0.000022 (0.002)	. .
TOU Peak Price in Effect * cooling degree hours squared * central AC	.	.	-0.00044	.

	Specification 1: Simplest Diff in Diff	Specification 2: Adding Survey Variables	Specification 3: Adding CAC*CDH interactions	Specification 4: Adds person FE's; controls
	.	.	(0.00084)	.
TOU Peak Price in Effect * cooling degree hours 2-7PM, previous day	. .	. .	. .	-0.001* (0.00059)
TOU Peak Price in Effect * cooling degree hours 2-7PM, two days before	. .	. .	. .	-0.0000072 (0.00045)
TOU Peak Price in Effect * cooling degree hours 2-7PM, three days before	. .	. .	. .	-0.001** (0.00045)
TOU Peak Price in Effect * work from home 11-30 hrs/wk	. .	. .	. .	0.055 (0.106)
TOU Peak Price in Effect * work from home >30 hrs/wk	. .	. .	. .	-0.319 (0.267)
TOU Peak Price in Effect * swimming pool	. .	. .	. .	-0.279* (0.148)
TOU Peak Price in Effect * spa	. .	. .	. .	0.062 (0.139)
TOU Peak Price in Effect * cooling degree hours 2-7pm * room AC	. .	. .	. .	0.010*** (0.003)
TOU Peak Price in Effect * heating degree hours 2-7PM * electric heat	. .	. .	. .	0.010*** (0.003)
TOU Peak Price in Effect * electric heat	. .	. .	. .	-0.161* (0.095)
TOU Peak Price in Effect * # kids under 5 in household	. .	. .	. .	-0.106 (0.072)
TOU Peak Price in Effect * # kids over 5 in household	. .	. .	. .	-0.054 (0.052)
TOU Peak Price in Effect * # people over 65 in household	.	.	.	-0.117**

	Specification 1: Simplest Diff in Diff	Specification 2: Adding Survey Variables	Specification 3: Adding CAC*CDH interactions	Specification 4: Adds person FE's; controls
	.	.	.	(0.057)
TOU Peak Price in Effect * work from home 0-10 hrs/wk	. .	. .	. .	-0.026 (0.125)
TOU Peak Price in Effect * electric cooktop	. .	. .	. .	0.184 (0.145)
TOU Peak Price in Effect * electric oven	. .	. .	. .	-0.139 (0.140)
TOU Peak Price in Effect * number of refrigerators and freezers	. .	. .	. .	-0.121 (0.085)
TOU Peak Price in Effect * customer stayed in expt. < 4.5 months	. .	. .	. .	-0.079 (0.126)
TOU Peak Price in Effect * customer stayed in expt. throughout expt.	. .	. .	. .	-0.193** (0.089)
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 before critical price	0.081*** (0.025)	0.081*** (0.027)	0.064** (0.027)	0.047 (0.033)
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 * apartment	-0.011 (0.090)	0.017 (0.123)	0.028 (0.122)	-0.034 (0.147)
Critical Price in Effect * climate zone 2	0.003	-0.040	-0.052	-0.058

	Specification 1: Simplest Diff in Diff	Specification 2: Adding Survey Variables	Specification 3: Adding CAC*CDH interactions	Specification 4: Adds person FE's; controls
	(0.064)	(0.075)	(0.072)	(0.099)
Critical Price in Effect * climate zone 3	-0.074 (0.110)	-0.010 (0.133)	-0.005 (0.136)	-0.032 (0.156)
Critical Price in Effect * climate zone 4	-0.171 (0.169)	-0.062 (0.197)	-0.023 (0.191)	0.019 (0.217)
Critical Price in Effect * cooling degree hours 2-7pm	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), 2-7pm	-0.110*** (0.038)	-0.065 (0.041)	-0.074 (0.161)	-0.107** (0.044)
Critical Price in Effect * heating degree hours 2-7pm	0.007 (0.008)	0.003 (0.010)	0.00042 (0.008)	-0.029* (0.015)
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 * number of bedrooms	. .	0.033 (0.059)	0.031 (0.058)	0.025 (0.059)
Critical Price in Effect * # people in the household	. .	0.030 (0.027)	0.035 (0.027)	0.080* (0.048)
Critical Price in Effect * cooling degree hours 2-7pm * central AC	. .	. .	0.000017 (0.003)	. .
Critical Price in Effect * 2-7pm squared * central AC	. .	. .	0.0000065 (0.00093)	. .
Critical Price in Effect * heating degree hours 2-7pm squared (1000's)	. .	. .	. .	1.194** (0.592)
Critical Price in Effect * cooling degree hours 2-7PM, previous day	.	.	.	-0.002**

	Specification 1: Simplest Diff in Diff	Specification 2: Adding Survey Variables	Specification 3: Adding CAC*CDH interactions	Specification 4: Adds person FE's; controls
	.	.	.	(0.001)
Critical Price in Effect * cooling degree hours 2-7PM, two days before	. .	. .	. .	0.002 (0.001)
Critical Price in Effect * cooling degree hours 2-7PM, three days before	. .	. .	. .	-0.002 (0.001)
Critical Price in Effect * work from home 11-30 hrs/wk	. .	. .	. .	0.018 (0.187)
Critical Price in Effect * work from home >30 hrs/wk	. .	. .	. .	-0.243 (0.291)
Critical Price in Effect * swimming pool	. .	. .	. .	-0.289 (0.196)
Critical Price in Effect * spa	. .	. .	. .	0.124 (0.193)
Critical Price in Effect * cooling degree hours 2-7pm * room AC	. .	. .	. .	0.010*** (0.003)
Critical Price in Effect * heating degree hours 2-7PM* electric heat	. .	. .	. .	0.042** (0.021)
Critical Price in Effect * electric heat	. .	. .	. .	-0.178 (0.140)
Critical Price in Effect * # kids under 5 in household	. .	. .	. .	-0.221** (0.093)
Critical Price in Effect * # kids over 5 in household	. .	. .	. .	-0.051 (0.068)
Critical Price in Effect * # people over 65 in household	. .	. .	. .	-0.223*** (0.085)
Critical Price in Effect * work from home 0-10 hrs/wk	.	.	.	0.001

	Specification 1: Simplest Diff in Diff	Specification 2: Adding Survey Variables	Specification 3: Adding CAC*CDH interactions	Specification 4: Adds person FE's; controls
	.	.	.	(0.176)
Critical Price in Effect * electric cooktop	. .	. .	. .	0.377* (0.195)
Critical Price in Effect * electric oven	. .	. .	. .	-0.249 (0.185)
Critical Price in Effect * number of refrigerators and freezers	. .	. .	. .	-0.245** (0.106)
Critical Price in Effect * customer stayed in expt. < 4.5 months	. .	. .	. .	0.161 (0.164)
Critical Price in Effect * customer stayed in expt. throughout expt.	. .	. .	. .	-0.352** (0.138)
Treatment Customer	-0.062 (0.077)	0.220 (0.160)	0.118 (0.154)	. .
Treatment Customer * electric use, kWh / day, summer 2002	0.000034 (0.004)	0.002 (0.005)	0.001 (0.005)	. .
Treatment Customer * apartment	0.102* (0.058)	-0.022 (0.092)	0.017 (0.089)	. .
Treatment Customer * climate zone 2	0.027 (0.056)	-0.013 (0.072)	-0.023 (0.068)	. .
Treatment Customer * climate zone 3	0.061 (0.074)	0.021 (0.095)	0.045 (0.093)	. .
Treatment Customer * climate zone 4	0.271* (0.156)	0.248 (0.173)	0.237 (0.163)	. .
Treatment Customer * cooling degree hours 2-7pm	-0.013*** (0.004)	-0.010** (0.004)	-0.007 (0.005)	-0.012** (0.005)
Treatment Customer * cooling degree hours squared (1000's), 2-7pm	0.098***	0.062	0.027	0.096**

	Specification 1: Simplest Diff in Diff	Specification 2: Adding Survey Variables	Specification 3: Adding CAC*CDH interactions	Specification 4: Adds person FE's; controls
	(0.037)	(0.039)	(0.121)	(0.041)
Treatment Customer * heating degree hours 2-7pm	-0.00018 (0.002)	0.000074 (0.002)	-0.00030 (0.002)	0.004 (0.005)
Treatment Customer * central AC	. .	-0.042 (0.074)	-0.026 (0.075)	. .
Treatment Customer * room AC	. .	0.099 (0.083)	0.098 (0.081)	. .
Treatment Customer * number of bedrooms	. .	-0.104** (0.044)	-0.080* (0.042)	. .
Treatment Customer * # people in the household	. .	0.013 (0.021)	0.020 (0.021)	. .
Treatment Customer * cooling degree hours 2-7pm * central AC	. .	. .	-0.003 (0.003)	. .
Treatment Customer * cooling degree hours 2-7pm squared * central AC	. .	. .	0.00029 (0.00076)	. .
Treatment Customer * heating degree hours 2-7pm squared (1000's)	. .	. .	. .	-0.160 (0.125)
Treatment Period (after 7/1/2003)	-0.035 (0.054)	0.036 (0.109)	0.020 (0.105)	-0.267* (0.145)
Treatment Period * electric use, kWh / day, summer 2002	0.007** (0.003)	0.006 (0.004)	0.005 (0.004)	-0.00040 (0.004)
Treatment Period * apartment	0.037 (0.041)	0.008 (0.061)	0.013 (0.062)	0.026 (0.070)
Treatment Period * climate zone 2	-0.013 (0.040)	-0.016 (0.046)	0.006 (0.045)	0.00013 (0.059)
Treatment Period * climate zone 3	0.035	-0.013	0.004	-0.015

	Specification 1: Simplest Diff in Diff	Specification 2: Adding Survey Variables	Specification 3: Adding CAC*CDH interactions	Specification 4: Adds person FE's; controls
	(0.055)	(0.069)	(0.069)	(0.079)
Treatment Period * climate zone 4	0.055 (0.097)	-0.045 (0.111)	-0.021 (0.106)	-0.120 (0.142)
Treatment Period * cooling degree hours 2-7pm	-0.003 (0.003)	-0.002 (0.003)	0.005 (0.004)	-0.004 (0.004)
Treatment Period * cooling degree hours squared (1000's), 2-7pm	0.025 (0.029)	0.019 (0.031)	-0.155* (0.094)	0.018 (0.032)
Treatment Period * heating degree hours 2-7pm	-0.003 (0.003)	-0.002 (0.003)	-0.001 (0.003)	0.010 (0.006)
Treatment Period * central AC	. .	0.138** (0.060)	0.075 (0.063)	0.144** (0.067)
Treatment Period * room AC	. .	-0.052 (0.060)	-0.046 (0.060)	0.021 (0.090)
Treatment Period * number of bedrooms	. .	-0.027 (0.031)	-0.025 (0.029)	-0.031 (0.033)
Treatment Period * # people in the household	. .	0.00042 (0.018)	0.003 (0.018)	-0.019 (0.028)
Treatment Period * cooling degree hours 2-7pm * central AC	. .	. .	-0.001 (0.002)	. .
Treatment Period * cooling degree hours 2-7pm squared * central AC	. .	. .	0.001** (0.00057)	. .
Treatment Period * heating degree hours 2-7pm squared (1000's)	. .	. .	. .	-0.062 (0.156)
Treatment Period * cooling degree hours 2-7PM, previous day	. .	. .	. .	0.003*** (0.00046)
Treatment Period * cooling degree hours 2-7PM, two days before	.	.	.	0.00095**

	Specification 1: Simplest Diff in Diff	Specification 2: Adding Survey Variables	Specification 3: Adding CAC*CDH interactions	Specification 4: Adds person FE's; controls
	.	.	.	(0.00040)
Treatment Period * cooling degree hours 2-7PM, three days before	. .	. .	. .	0.00100*** (0.00034)
Treatment Period * work from home 11-30 hrs/wk	. .	. .	. .	-0.020 (0.080)
Treatment Period * work from home >30 hrs/wk	. .	. .	. .	0.085 (0.213)
Treatment Period * swimming pool	. .	. .	. .	0.251** (0.118)
Treatment Period * spa	. .	. .	. .	-0.058 (0.107)
Treatment Period * cooling degree hours 2-7pm * room AC	. .	. .	. .	-0.008*** (0.002)
Treatment Period * heating degree hours 2-7PM* electric heat	. .	. .	. .	-0.007*** (0.002)
Treatment Period * electric heat	. .	. .	. .	0.151** (0.072)
Treatment Period * # kids under 5 in household	. .	. .	. .	0.106* (0.055)
Treatment Period * # kids over 5 in household	. .	. .	. .	0.062* (0.033)
Treatment Period * # people over 65 in household	. .	. .	. .	0.155*** (0.047)
Treatment Period * work from home 0-10 hrs/wk	. .	. .	. .	0.169* (0.089)
Treatment Period * electric cooktop	.	.	.	-0.130

	Specification 1: Simplest Diff in Diff	Specification 2: Adding Survey Variables	Specification 3: Adding CAC*CDH interactions	Specification 4: Adds person FE's; controls
	.	.	.	(0.118)
Treatment Period * electric oven	. .	. .	. .	-0.003 (0.108)
Treatment Period * number of refrigerators and freezers	. .	. .	. .	0.011 (0.064)
Treatment Period * customer stayed in expt. < 4.5 months	. .	. .	. .	0.062 (0.101)
Treatment Period * customer stayed in expt. throughout expt.	. .	. .	. .	0.184*** (0.072)
Critical Period	-0.250*** (0.043)	-0.314*** (0.085)	-0.253*** (0.079)	-0.187 (0.130)
Critical Period * electric use, kWh / day, summer 2002	0.017*** (0.002)	0.015*** (0.003)	0.015*** (0.003)	0.017*** (0.003)
Critical Period * high ratio rate customer.	-0.209* (0.121)	-0.226* (0.133)	-0.201 (0.125)	-0.063 (0.071)
Critical Period * apartment	-0.004 (0.041)	0.045 (0.062)	0.022 (0.058)	0.042 (0.080)
Critical Period * climate zone 2	0.042 (0.030)	-0.005 (0.032)	0.00040 (0.030)	0.042 (0.049)
Critical Period * climate zone 3	0.217*** (0.054)	0.100 (0.064)	0.070 (0.066)	0.084 (0.075)
Critical Period * climate zone 4	0.183** (0.084)	0.046 (0.087)	-0.00070 (0.085)	-0.011 (0.106)
Critical Period * cooling degree hours 2-7pm	-0.003* (0.001)	-0.002 (0.001)	0.002 (0.003)	-0.004*** (0.001)
Critical Period * cooling degree hours squared (1000's), 2-7pm	-0.009	-0.017	-0.020	0.004

	Specification 1: Simplest Diff in Diff	Specification 2: Adding Survey Variables	Specification 3: Adding CAC*CDH interactions	Specification 4: Adds person FE's; controls
	(0.012)	(0.013)	(0.055)	(0.013)
Critical Period * heating degree hours 2-7pm	-0.004 (0.005)	-0.007 (0.008)	-0.005 (0.005)	0.013 (0.011)
Critical Period * central AC	. .	0.267*** (0.052)	0.114** (0.057)	0.274*** (0.060)
Critical Period * room AC	. .	-0.069 (0.062)	-0.062 (0.069)	0.072 (0.093)
Critical Period * number of bedrooms	. .	0.026 (0.024)	0.020 (0.024)	0.007 (0.031)
Critical Period * # people in the household	. .	-0.010 (0.012)	-0.011 (0.012)	-0.003 (0.019)
Critical Period * cooling degree hours 2-7pm * central AC	. .	. .	-0.002** (0.00093)	. .
Critical Period * cooling degree hours 2-7pm squared * central AC	. .	. .	-0.000073 (0.00031)	. .
Critical Period * heating degree hours 2-7pm squared (1000's)	. .	. .	. .	-0.697 (0.442)
Critical Period * cooling degree hours 2-7PM, previous day	. .	. .	. .	-0.001 (0.001)
Critical Period * cooling degree hours 2-7PM, two days before	. .	. .	. .	-0.00079 (0.001)
Critical Period * cooling degree hours 2-7PM, three days before	. .	. .	. .	0.002* (0.001)
Critical Period * work from home 11-30 hrs/wk	. .	. .	. .	0.083 (0.113)
Critical Period * work from home >30 hrs/wk	.	.	.	0.049

	Specification 1: Simplest Diff in Diff	Specification 2: Adding Survey Variables	Specification 3: Adding CAC*CDH interactions	Specification 4: Adds person FE's; controls
	.	.	.	(0.136)
Critical Period * swimming pool	. .	. .	. .	-0.009 (0.093)
Critical Period * spa	. .	. .	. .	-0.081 (0.080)
Critical Period * cooling degree hours 2-7pm * room AC	. .	. .	. .	0.00018 (0.001)
Critical Period * heating degree hours 2-7PM* electric heat	. .	. .	. .	0.005 (0.009)
Critical Period * electric heat	. .	. .	. .	-0.065 (0.071)
Critical Period * # kids under 5 in household	. .	. .	. .	0.027 (0.048)
Critical Period * # kids over 5 in household	. .	. .	. .	-0.021 (0.028)
Critical Period * # people over 65 in household	. .	. .	. .	0.106** (0.049)
Critical Period * work from home 0-10 hrs/wk	. .	. .	. .	-0.004 (0.071)
Critical Period * electric cooktop	. .	. .	. .	-0.147* (0.081)
Critical Period * electric oven	. .	. .	. .	0.091 (0.072)
Critical Period * number of refrigerators and freezers	. .	. .	. .	0.017 (0.045)
Critical Period * customer stayed in expt. < 4.5 months	.	.	.	-0.329***

	Specification 1: Simplest Diff in Diff	Specification 2: Adding Survey Variables	Specification 3: Adding CAC*CDH interactions	Specification 4: Adds person FE's; controls
	.	.	.	(0.110)
Critical Period * customer stayed in expt. throughout expt.	. .	. .	. .	0.161** (0.069)
	.	.		
electric use, kWh / day, summer 2002	0.047*** (0.003)	0.046*** (0.004)	0.045*** (0.004)	. .
trt. customer on high-ratio rate	-0.024 (0.036)	0.006 (0.039)	0.007 (0.038)	. .
apartment	-0.054 (0.044)	0.056 (0.081)	0.006 (0.078)	. .
climate zone 2	0.013 (0.042)	0.002 (0.059)	0.039 (0.056)	. .
climate zone 3	0.020 (0.057)	-0.004 (0.078)	0.032 (0.076)	. .
climate zone 4	-0.207 (0.129)	-0.239* (0.137)	-0.238* (0.131)	. .
cooling degree hours 2-7PM, base 78	0.017*** (0.004)	0.017*** (0.004)	0.008* (0.005)	0.019*** (0.005)
heating degree hours squared (1000's), 2-7pm	-0.054 (0.035)	-0.033 (0.038)	. .	. .
cooling degree hours squared (1000's), 2-7pm	-0.054 (0.035)	-0.033 (0.038)	0.009 (0.094)	-0.040 (0.040)
heating degree hours squared (1000's), 2-7pm	-0.173 (0.063)	-0.204 (0.076)	-0.131 (0.066)	0.108 (0.107)
heating degree hours 2-7pm	0.009*** (0.003)	0.010** (0.004)	0.006* (0.003)	-0.003 (0.005)

	Specification 1: Simplest Diff in Diff	Specification 2: Adding Survey Variables	Specification 3: Adding CAC*CDH interactions	Specification 4: Adds person FE's; controls
Tuesday	-0.007 (0.006)	-0.005 (0.007)	-0.004 (0.007)	-0.003 (0.009)
Wednesday	-0.006 (0.006)	-0.005 (0.007)	-0.005 (0.007)	-0.004 (0.008)
Thursday	-0.019** (0.008)	-0.017* (0.010)	-0.015 (0.009)	-0.008 (0.010)
Friday	-0.003 (0.008)	0.00030 (0.010)	-0.00056 (0.010)	0.002 (0.010)
year 2004	-0.029 (0.023)	-0.027 (0.026)	-0.021 (0.027)	0.008 (0.032)
June	0.071*** (0.013)	0.076*** (0.014)	0.065*** (0.014)	0.042*** (0.016)
July	0.130*** (0.019)	0.151*** (0.021)	0.126*** (0.021)	0.112*** (0.022)
August	0.157*** (0.021)	0.177*** (0.024)	0.148*** (0.024)	0.131*** (0.026)
September	0.101*** (0.017)	0.113*** (0.019)	0.096*** (0.019)	0.078*** (0.022)
October	0.040 (0.024)	0.050* (0.028)	0.045 (0.028)	0.042 (0.032)
heating degree hours 2-7pm squared (1000's)	-0.173*** (0.063)	-0.204*** (0.076)	-0.131** (0.066)	. .
Tue * cooling degree hours 2-7pm	-0.00073 (0.00061)	-0.00062 (0.00064)	-0.00094 (0.00060)	-0.001 (0.00069)
Tue * cooling degree hours 2-7pm squared (1000's)	0.00076 (0.007)	-0.002 (0.007)	0.002 (0.006)	0.003 (0.007)

	Specification 1: Simplest Diff in Diff	Specification 2: Adding Survey Variables	Specification 3: Adding CAC*CDH interactions	Specification 4: Adds person FE's; controls
Tue * heating degree hours 2-7pm	-0.001 (0.001)	-0.002 (0.002)	-0.001 (0.001)	-0.001 (0.001)
Tue * heating degree hours 2-7pm squared (1000's)	0.044* (0.027)	0.048 (0.029)	0.038 (0.025)	0.023 (0.023)
Wed * cooling degree hours 2-7pm	-0.002** (0.00069)	-0.001 (0.00074)	-0.002** (0.00072)	-0.002** (0.00082)
Wed * cooling degree hours 2-7pm squared (1000's)	0.015** (0.007)	0.011 (0.008)	0.016** (0.007)	0.012 (0.008)
Wed * heating degree hours 2-7pm	0.001 (0.002)	0.00078 (0.002)	0.001 (0.002)	0.001 (0.002)
Wed * heating degree hours 2-7pm squared (1000's)	0.009 (0.030)	0.010 (0.034)	-0.001 (0.031)	-0.015 (0.032)
Thu * cooling degree hours 2-7pm	-0.00091 (0.00071)	-0.00077 (0.00079)	-0.002** (0.00073)	-0.001 (0.00073)
Thu * cooling degree hours 2-7pm squared (1000's)	0.005 (0.007)	0.004 (0.008)	0.012 (0.008)	0.006 (0.008)
Thu * heating degree hours 2-7pm	-0.002 (0.002)	-0.003 (0.002)	-0.002 (0.002)	-0.002 (0.002)
Thu * heating degree hours 2-7pm squared (1000's)	0.082* (0.045)	0.091* (0.052)	0.072 (0.044)	0.049 (0.034)
Fri * cooling degree hours 2-7pm	-0.002*** (0.00072)	-0.002** (0.00082)	-0.003*** (0.00088)	-0.00086 (0.00090)
Fri * cooling degree hours 2-7pm squared (1000's)	0.012* (0.007)	0.011 (0.008)	0.023** (0.010)	0.004 (0.009)
Fri * heating degree hours 2-7pm	-0.005*** (0.002)	-0.006*** (0.002)	-0.005** (0.002)	-0.005** (0.002)

	Specification 1: Simplest Diff in Diff	Specification 2: Adding Survey Variables	Specification 3: Adding CAC*CDH interactions	Specification 4: Adds person FE's; controls
Fri * heating degree hours 2-7pm squared (1000's)	0.108*** (0.035)	0.128*** (0.044)	0.104*** (0.037)	0.084** (0.033)
year_2004 * cooling degree hours 2-7pm	-0.004** (0.002)	-0.004** (0.002)	-0.005*** (0.002)	-0.003* (0.002)
year_2004 * cooling degree hours 2-7pm squared (1000's)	0.037** (0.016)	0.037** (0.018)	0.043** (0.018)	0.033* (0.018)
year_2004 * heating degree hours 2-7pm	0.002 (0.002)	0.002 (0.003)	0.002 (0.003)	0.003 (0.005)
year_2004 * heating degree hours 2-7pm squared (1000's)	0.029 (0.054)	0.032 (0.064)	-0.003 (0.060)	-0.129 (0.133)
June * cooling degree hours 2-7pm	-0.002 (0.002)	-0.003 (0.002)	-0.003 (0.002)	0.002 (0.002)
June * cooling degree hours 2-7pm squared (1000's)	0.026 (0.028)	0.037 (0.028)	0.027 (0.025)	-0.010 (0.023)
June * heating degree hours 2-7pm	-0.005** (0.002)	-0.005** (0.002)	-0.004* (0.002)	-0.002 (0.002)
June * heating degree hours 2-7pm squared (1000's)	0.069** (0.031)	0.079** (0.036)	0.069* (0.035)	0.062 (0.039)
July * cooling degree hours 2-7pm	-0.00049 (0.002)	-0.002 (0.002)	0.00010 (0.002)	-0.00038 (0.002)
July * cooling degree hours 2-7pm squared (1000's)	0.017 (0.021)	0.023 (0.022)	-0.002 (0.020)	0.012 (0.023)
July * heating degree hours 2-7pm	-0.015*** (0.004)	-0.016*** (0.004)	-0.013*** (0.004)	-0.013*** (0.003)
July * heating degree hours 2-7pm squared (1000's)	0.398*** (0.122)	0.351*** (0.126)	0.309*** (0.120)	0.285*** (0.094)

	Specification 1: Simplest Diff in Diff	Specification 2: Adding Survey Variables	Specification 3: Adding CAC*CDH interactions	Specification 4: Adds person FE's; controls
Aug * cooling degree hours 2-7pm	-0.002 (0.002)	-0.003 (0.002)	-0.001 (0.002)	-0.00099 (0.002)
Aug * cooling degree hours 2-7pm squared (1000's)	0.027 (0.020)	0.034 (0.021)	0.010 (0.019)	0.012 (0.021)
Aug * heating degree hours 2-7pm	-0.035*** (0.010)	-0.037*** (0.011)	-0.030*** (0.009)	-0.018** (0.007)
Aug * heating degree hours 2-7pm squared (1000's)	1.038** (0.420)	1.064** (0.444)	0.898** (0.388)	0.438 (0.280)
Sept * cooling degree hours 2-7pm	-0.006*** (0.002)	-0.007*** (0.002)	-0.005*** (0.002)	-0.005** (0.002)
Sept * cooling degree hours 2-7pm squared (1000's)	0.053*** (0.017)	0.058*** (0.018)	0.034* (0.018)	0.036* (0.019)
Sept * heating degree hours 2-7pm	-0.004** (0.002)	-0.005* (0.003)	-0.003 (0.003)	-0.002 (0.002)
Sept * heating degree hours 2-7pm squared (1000's)	0.056* (0.031)	0.067* (0.040)	0.044 (0.035)	0.032 (0.028)
Oct * cooling degree hours 2-7pm	-0.014*** (0.002)	-0.015*** (0.002)	-0.015*** (0.003)	-0.010*** (0.002)
Oct * cooling degree hours 2-7pm squared (1000's)	0.109*** (0.025)	0.115*** (0.027)	0.111*** (0.029)	0.056** (0.027)
Oct * heating degree hours 2-7pm	-0.00019 (0.003)	-0.00099 (0.003)	0.001 (0.003)	0.002 (0.005)
Oct * heating degree hours 2-7pm squared (1000's)	0.063 (0.052)	0.070 (0.062)	0.020 (0.058)	-0.106 (0.131)
constant	-0.150*** (0.057)	-0.434*** (0.142)	-0.309** (0.137)	0.627*** (0.030)

	Specification 1: Simplest Diff in Diff	Specification 2: Adding Survey Variables	Specification 3: Adding CAC*CDH interactions	Specification 4: Adds person FE's; controls
central AC	.	0.082	-0.063	.
	.	(0.056)	(0.058)	.
room AC	.	-0.036	-0.026	.
	.	(0.066)	(0.065)	.
number of bedrooms	.	0.072**	0.053	.
	.	(0.035)	(0.033)	.
# people in the household	.	0.014	0.012	.
	.	(0.017)	(0.017)	.
cooling degree hours 2-7pm * central AC	.	.	0.013***	.
	.	.	(0.002)	.
cooling degree hours 2-7pm squared * central AC	.	.	-0.00048	.
	.	.	(0.00055)	.
one * cooling degree hours 2-7pm squared (1000's)	.	.	0.009	-0.041
	.	.	(0.094)	(0.040)
heating degree hours 2-7pm squared, 1000's	.	.	.	0.108
	.	.	.	(0.108)
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%

Cooling degree hours are base 78^o F. Heating degree hours are base 65^o F.

Appendix E

CPP Impacts in Just Climates with Hot Summers and Lots of Central Air Conditioning

California has regions with hot summers and air conditioning in the majority of residences. The SPP designated these areas as climate zones 3 (Central Valley) and 4 (Desert)). California also has more temperate coastal regions. Zones 3 and 4 are more comparable to conditions in much of the rest of the United States and in California regions like Sacramento and Imperial County where there are large, public utilities. Hence, this table reproduces the SPP's impacts for just climate zones 3 and 4. Qualitatively, the main results are quite similar to and often slightly stronger than those found above.

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 FE's; controls
TOU Peak Price in Effect	0.027	-0.228	-0.281	0.422

	Specification 1: Simplest Diff in Diff	Specification 2: Adding Survey Variables	Specification 3: Adding CAC*CDH interactions	Specification 4: Adds person FE's; controls
	(0.128)	(0.267)	(0.265)	(0.437)
TOU peak price in effect * day before critical price	0.025 (0.024)	0.027 (0.026)	0.028 (0.026)	0.017 (0.023)
TOU peak price in effect * day after critical price	0.046* (0.027)	0.057** (0.029)	0.057** (0.028)	0.043* (0.024)
TOU Peak Price in Effect * electric use, kWh / day , summer 2002	-0.002 (0.005)	-0.002 (0.007)	-0.00057 (0.007)	0.013* (0.008)
TOU Peak Price in Effect * high ratio rate customer.	-0.040 (0.078)	-0.016 (0.089)	-0.003 (0.088)	0.007 (0.082)
TOU Peak Price in Effect * apartment	-0.004 (0.129)	0.081 (0.178)	0.055 (0.177)	0.095 (0.200)
TOU Peak Price in Effect * climate zone 4	-0.193 (0.162)	-0.162 (0.172)	-0.119 (0.169)	-0.009 (0.169)
TOU Peak Price in Effect * cooling degree hours 2-7pm	0.010** (0.004)	0.009** (0.004)	0.009 (0.007)	0.009** (0.004)
TOU Pk Price in Effect * cooling degree hours squared (1000's), 2-7pm	-0.108*** (0.038)	-0.089** (0.039)	-0.096 (0.160)	-0.109*** (0.039)
TOU Peak Price in Effect * heating degree hours 2-7pm	-0.020* (0.010)	-0.025** (0.012)	-0.023** (0.012)	0.010 (0.021)
TOU Peak Price in Effect * central AC	. .	-0.106 (0.122)	-0.065 (0.125)	-0.242 (0.149)
TOU Peak Price in Effect * room AC	. .	0.127 (0.159)	0.152 (0.159)	-0.333 (0.235)
TOU Peak Price in Effect * number of bedrooms	. .	0.071 (0.078)	0.066 (0.075)	0.074 (0.090)
TOU Peak Price in Effect * # people in the household	.	0.023	0.023	0.068

	Specification 1: Simplest Diff in Diff	Specification 2: Adding Survey Variables	Specification 3: Adding CAC*CDH interactions	Specification 4: Adds person FE's; controls
	.	(0.040)	(0.040)	(0.091)
TOU Peak Price in Effect * cooling degree hours 2-7pm * central AC	. .	. .	-0.00086 (0.003)	. .
TOU Peak Price in Effect * cooling degree hours 2-7pm squared * central AC	. .	. .	0.000048 (0.00096)	. .
TOU Peak Price in Effect * heating degree hours 2-7pm squared (1000's)	. .	. .	. .	-0.237 (1.241)
TOU Peak Price in Effect * cooling degree hours 2-7PM, previous day	. .	. .	. .	-0.001* (0.00072)
TOU Peak Price in Effect * cooling degree hours 2-7PM, two days before	. .	. .	. .	-0.00011 (0.00064)
TOU Peak Price in Effect * cooling degree hours 2-7PM, three days before	. .	. .	. .	-0.001** (0.00062)
TOU Peak Price in Effect * work from home 11-30 hrs/wk	. .	. .	. .	0.283 (0.338)
TOU Peak Price in Effect * work from home >30 hrs/wk	. .	. .	. .	-0.234 (0.492)
TOU Peak Price in Effect * swimming pool	. .	. .	. .	0.081 (0.217)
TOU Peak Price in Effect * spa	. .	. .	. .	-0.407 (0.247)
TOU Peak Price in Effect * cooling degree hours 2-7pm * room AC	. .	. .	. .	0.009** (0.004)
TOU Peak Price in Effect * heating degree hours 2-7PM* electric heat	. .	. .	. .	0.009* (0.006)
TOU Peak Price in Effect * electric heat	.	.	.	-0.085

	Specification 1: Simplest Diff in Diff	Specification 2: Adding Survey Variables	Specification 3: Adding CAC*CDH interactions	Specification 4: Adds person FE's; controls
	.	.	.	(0.168)
TOU Peak Price in Effect * # kids under 5 in household	. .	. .	. .	-0.117 (0.139)
TOU Peak Price in Effect * # kids over 5 in household	. .	. .	. .	-0.138 (0.124)
TOU Peak Price in Effect * # people over 65 in household	. .	. .	. .	-0.176 (0.124)
TOU Peak Price in Effect * work from home 0-10 hrs/wk	. .	. .	. .	-0.027 (0.239)
TOU Peak Price in Effect * electric cooktop	. .	. .	. .	0.030 (0.281)
TOU Peak Price in Effect * electric oven	. .	. .	. .	-0.076 (0.267)
TOU Peak Price in Effect * number of refrigerators and freezers	. .	. .	. .	-0.190 (0.131)
TOU Peak Price in Effect * customer stayed in expt. < 4.5 months	. .	. .	. .	0.029 (0.327)
TOU Peak Price in Effect * customer stayed in expt. throughout expt.	. .	. .	. .	-0.507** (0.216)
Critical Price in Effect	0.076 (0.206)	-0.133 (0.370)	-0.302 (0.390)	0.865 (0.603)
Critical Price in Effect * day before critical price	0.069 (0.053)	0.086 (0.056)	0.080 (0.056)	0.050 (0.063)
Critical Price in Effect * day after critical price	0.097* (0.052)	0.118** (0.056)	0.113** (0.056)	0.037 (0.065)
Critical Price in Effect * electric use, kWh / day , summer 2002	-0.015*	-0.018**	-0.017*	0.000069

	Specification 1: Simplest Diff in Diff	Specification 2: Adding Survey Variables	Specification 3: Adding CAC*CDH interactions	Specification 4: Adds person FE's; controls
	(0.008)	(0.009)	(0.009)	(0.011)
Critical Price in Effect * high ratio rate customer.	0.150 (0.279)	0.395 (0.294)	0.366 (0.283)	0.046 (0.204)
Critical Price in Effect * apartment	0.248 (0.197)	0.316 (0.261)	0.329 (0.257)	0.439 (0.292)
Critical Price in Effect * climate zone 4	-0.109 (0.203)	-0.083 (0.217)	-0.063 (0.213)	0.021 (0.205)
Critical Price in Effect * cooling degree hours 2-7pm	0.010 (0.006)	0.007 (0.006)	0.019* (0.011)	0.004 (0.006)
Critical Price in Effect * cooling degree hours squared (1000's), 2-7pm (base 78	-0.118** (0.047)	-0.083* (0.049)	-0.317 (0.217)	-0.089* (0.048)
Critical Price in Effect * heating degree hours 2-7pm	-0.119 (0.117)	-0.092 (0.113)	-0.031 (0.121)	-0.051 (0.358)
Critical Price in Effect * central AC	. .	-0.226 (0.167)	-0.178 (0.234)	-0.245 (0.222)
Critical Price in Effect * room AC	. .	0.298 (0.214)	0.290 (0.212)	-0.807** (0.370)
Critical Price in Effect * number of bedrooms	. .	0.056 (0.114)	0.064 (0.111)	0.052 (0.121)
Critical Price in Effect * # people in the household	. .	0.053 (0.047)	0.056 (0.047)	0.159 (0.105)
Critical Price in Effect * cooling degree hours 2-7pm * central AC	. .	. .	-0.001 (0.003)	. .
Critical Price in Effect * cooling degree hours 2-7pm squared * central AC	. .	. .	0.001 (0.001)	. .
Critical Price in Effect * heating degree hours 2-7pm squared (1000's)	.	.	.	-111.179

	Specification 1: Simplest Diff in Diff	Specification 2: Adding Survey Variables	Specification 3: Adding CAC*CDH interactions	Specification 4: Adds person FE's; controls
	.	.	.	(112.521)
Critical Price in Effect * cooling degree hours 2-7PM, previous day	. .	. .	. .	-0.002 (0.002)
Critical Price in Effect * cooling degree hours 2-7PM, two days before	. .	. .	. .	0.00049 (0.002)
Critical Price in Effect * cooling degree hours 2-7PM, three days before	. .	. .	. .	-0.00046 (0.002)
Critical Price in Effect * work from home 11-30 hrs/wk	. .	. .	. .	-0.813* (0.468)
Critical Price in Effect * work from home >30 hrs/wk	. .	. .	. .	0.059 (0.506)
Critical Price in Effect * swimming pool	. .	. .	. .	0.264 (0.301)
Critical Price in Effect * spa	. .	. .	. .	-0.619* (0.329)
Critical Price in Effect * cooling degree hours 2-7pm * room AC	. .	. .	. .	0.015*** (0.004)
Critical Price in Effect * heating degree hours 2-7PM* electric heat	. .	. .	. .	0.108 (0.472)
Critical Price in Effect * electric heat	. .	. .	. .	-0.180 (0.252)
Critical Price in Effect * # kids under 5 in household	. .	. .	. .	-0.329** (0.164)
Critical Price in Effect * # kids over 5 in household	. .	. .	. .	-0.209 (0.152)
Critical Price in Effect * # people over 65 in household	.	.	.	-0.352*

	Specification 1: Simplest Diff in Diff	Specification 2: Adding Survey Variables	Specification 3: Adding CAC*CDH interactions	Specification 4: Adds person FE's; controls
	.	.	.	(0.191)
Critical Price in Effect * work from home 0-10 hrs/wk	. .	. .	. .	-0.071 (0.327)
Critical Price in Effect * electric cooktop	. .	. .	. .	0.222 (0.369)
Critical Price in Effect * electric oven	. .	. .	. .	-0.141 (0.345)
Critical Price in Effect * number of refrigerators and freezers	. .	. .	. .	-0.322** (0.156)
Critical Price in Effect * customer stayed in expt. < 4.5 months	. .	. .	. .	0.052 (0.427)
Critical Price in Effect * customer stayed in expt. throughout expt.	. .	. .	. .	-0.717** (0.347)
Treatment Customer	-0.134 (0.118)	0.262 (0.279)	0.297 (0.263)	. .
Treatment Customer * electric use, kWh / day , summer 2002	0.004 (0.006)	0.002 (0.007)	0.001 (0.007)	. .
Treatment Customer * apartment	0.350*** (0.110)	0.152 (0.169)	0.186 (0.165)	. .
Treatment Customer * climate zone 4	0.264 (0.173)	0.300* (0.178)	0.242 (0.177)	. .
Treatment Customer * cooling degree hours 2-7pm	-0.016*** (0.005)	-0.014*** (0.005)	-0.014* (0.008)	-0.014*** (0.005)
Treatment Customer * cooling degree hours squared (1000's), 2-7pm	0.121*** (0.043)	0.089** (0.045)	0.119 (0.148)	0.108** (0.043)
Treatment Customer * heating degree hours 2-7pm	0.018*	0.023*	0.022*	-0.015

	Specification 1: Simplest Diff in Diff	Specification 2: Adding Survey Variables	Specification 3: Adding CAC*CDH interactions	Specification 4: Adds person FE's; controls
	(0.010)	(0.012)	(0.011)	(0.022)
Treatment Customer * central AC	. .	0.055 (0.123)	-0.013 (0.125)	. .
Treatment Customer * room AC	. .	0.076 (0.173)	0.034 (0.167)	. .
Treatment Customer * number of bedrooms	. .	-0.138 (0.088)	-0.124 (0.084)	. .
Treatment Customer * # people in the household	. .	0.00079 (0.035)	0.006 (0.035)	. .
Treatment Customer * cooling degree hours 2-7pm * central AC	. .	. .	0.00012 (0.003)	. .
Treatment Customer * cooling degree hours 2-7pm squared * central AC	. .	. .	-0.00019 (0.00087)	. .
Treatment Customer * heating degree hours 2-7pm squared (1000's)	. .	. .	. .	0.258 (1.260)
Treatment Period (after 7/1/2003)	0.056 (0.098)	0.029 (0.192)	0.080 (0.186)	-0.417 (0.306)
Treatment Period * electric use, kWh / day , summer 2002	0.004 (0.003)	0.00053 (0.004)	0.00023 (0.004)	0.000032 (0.005)
Treatment Period * apartment	-0.047 (0.078)	-0.048 (0.114)	-0.027 (0.112)	0.049 (0.138)
Treatment Period * climate zone 4	0.098 (0.124)	0.069 (0.132)	0.045 (0.131)	-0.024 (0.140)
Treatment Period * cooling degree hours 2-7pm	-0.004 (0.004)	-0.005 (0.004)	0.002 (0.006)	-0.007 (0.004)
Treatment Period * cooling degree hours squared (1000's), 2-7pm	0.043	0.043	-0.113	0.042

	Specification 1: Simplest Diff in Diff	Specification 2: Adding Survey Variables	Specification 3: Adding CAC*CDH interactions	Specification 4: Adds person FE's; controls
	(0.035)	(0.035)	(0.116)	(0.032)
Treatment Period * heating degree hours 2-7pm	0.017 (0.049)	0.014 (0.054)	-0.029 (0.048)	-0.030 (0.114)
Treatment Period * central AC	. .	0.184** (0.091)	0.065 (0.092)	0.227* (0.125)
Treatment Period * room AC	. .	-0.015 (0.102)	-0.018 (0.100)	0.210 (0.193)
Treatment Period * number of bedrooms	. .	-0.00072 (0.061)	-0.003 (0.058)	-0.038 (0.062)
Treatment Period * # people in the household	. .	-0.005 (0.028)	-0.004 (0.028)	-0.052 (0.072)
Treatment Period * cooling degree hours 2-7pm * central AC	. .	. .	0.00063 (0.002)	. .
Treatment Period * cooling degree hours 2-7pm squared * central AC	. .	. .	0.00091 (0.00069)	. .
Treatment Period * heating degree hours 2-7pm squared (1000's)	. .	. .	. .	13.851 (25.589)
Treatment Period * cooling degree hours 2-7PM, previous day	. .	. .	. .	0.003*** (0.00059)
Treatment Period * cooling degree hours 2-7PM, two days before	. .	. .	. .	0.00095* (0.00051)
Treatment Period * cooling degree hours 2-7PM, three days before	. .	. .	. .	0.001** (0.00043)
Treatment Period * work from home 11-30 hrs/wk	. .	. .	. .	0.021 (0.245)
Treatment Period * work from home >30 hrs/wk	.	.	.	-0.168

	Specification 1: Simplest Diff in Diff	Specification 2: Adding Survey Variables	Specification 3: Adding CAC*CDH interactions	Specification 4: Adds person FE's; controls
	.	.	.	(0.366)
Treatment Period * swimming pool	. .	. .	. .	-0.052 (0.161)
Treatment Period * spa	. .	. .	. .	0.321 (0.200)
Treatment Period * cooling degree hours 2-7pm * room AC	. .	. .	. .	-0.009*** (0.003)
Treatment Period * heating degree hours 2-7PM* electric heat	. .	. .	. .	-0.005 (0.004)
Treatment Period * electric heat	. .	. .	. .	0.088 (0.130)
Treatment Period * # kids under 5 in household	. .	. .	. .	0.089 (0.111)
Treatment Period * # kids over 5 in household	. .	. .	. .	0.140 (0.092)
Treatment Period * # people over 65 in household	. .	. .	. .	0.195** (0.090)
Treatment Period * work from home 0-10 hrs/wk	. .	. .	. .	0.078 (0.145)
Treatment Period * electric cooktop	. .	. .	. .	-0.081 (0.228)
Treatment Period * electric oven	. .	. .	. .	-0.048 (0.211)
Treatment Period * number of refrigerators and freezers	. .	. .	. .	-0.006 (0.099)
Treatment Period * customer stayed in expt. < 4.5 months	.	.	.	-0.058

	Specification 1: Simplest Diff in Diff	Specification 2: Adding Survey Variables	Specification 3: Adding CAC*CDH interactions	Specification 4: Adds person FE's; controls
	.	.	.	(0.264)
Treatment Period * customer stayed in expt. throughout expt.	. .	. .	. .	0.467** (0.184)
Critical Period	-0.152 (0.093)	-0.491*** (0.148)	-0.384** (0.174)	-0.552 (0.362)
Critical Period * electric use, kWh / day , summer 2002	0.017*** (0.003)	0.016*** (0.003)	0.016*** (0.003)	0.015*** (0.004)
Critical Period * high ratio rate customer.	-0.164 (0.239)	-0.308 (0.241)	-0.262 (0.233)	0.050 (0.166)
Critical Period * apartment	-0.172** (0.073)	-0.083 (0.103)	-0.116 (0.102)	-0.258** (0.123)
Critical Period * climate zone 4	-0.082 (0.078)	-0.086 (0.078)	-0.082 (0.076)	-0.082 (0.087)
Critical Period * cooling degree hours 2-7pm	-0.00027 (0.002)	-0.00088 (0.002)	-0.00016 (0.005)	0.002 (0.002)
Critical Period * cooling degree hours squared (1000's), 2-7pm	-0.012 (0.017)	-0.015 (0.018)	0.027 (0.086)	-0.018 (0.017)
Critical Period * heating degree hours 2-7pm	0.004 (0.048)	0.050 (0.044)	-0.013 (0.042)	0.058 (0.288)
Critical Period * central AC	. .	0.257*** (0.064)	0.088 (0.116)	0.123 (0.095)
Critical Period * room AC	. .	0.072 (0.094)	0.084 (0.093)	0.569*** (0.204)
Critical Period * number of bedrooms	. .	0.064 (0.044)	0.059 (0.044)	0.069 (0.044)
Critical Period * # people in the household	.	-0.004	-0.004	-0.017

	Specification 1: Simplest Diff in Diff	Specification 2: Adding Survey Variables	Specification 3: Adding CAC*CDH interactions	Specification 4: Adds person FE's; controls
	.	(0.019)	(0.019)	(0.034)
Critical Period * cooling degree hours 2-7pm * central AC	. .	. .	-0.002 (0.002)	. .
Critical Period * cooling degree hours 2-7pm squared * central AC	. .	. .	-0.00030 (0.00043)	. .
Critical Period * heating degree hours 2-7pm squared (1000's)	. .	. .	. .	1.806 (54.306)
Critical Period * cooling degree hours 2-7PM, previous day	. .	. .	. .	-0.001 (0.002)
Critical Period * cooling degree hours 2-7PM, two days before	. .	. .	. .	-0.00031 (0.002)
Critical Period * cooling degree hours 2-7PM, three days before	. .	. .	. .	0.001 (0.002)
Critical Period * work from home 11-30 hrs/wk	. .	. .	. .	0.587*** (0.198)
Critical Period * work from home >30 hrs/wk	. .	. .	. .	-0.001 (0.166)
Critical Period * swimming pool	. .	. .	. .	-0.116 (0.119)
Critical Period * spa	. .	. .	. .	-0.029 (0.108)
Critical Period * cooling degree hours 2-7pm * room AC	. .	. .	. .	-0.004* (0.003)
Critical Period * heating degree hours 2-7PM* electric heat	. .	. .	. .	0.208 (0.275)
Critical Period * electric heat	.	.	.	0.002

	Specification 1: Simplest Diff in Diff	Specification 2: Adding Survey Variables	Specification 3: Adding CAC*CDH interactions	Specification 4: Adds person FE's; controls
	.	.	.	(0.121)
Critical Period * # kids under 5 in household	. .	. .	. .	0.051 (0.068)
Critical Period * # kids over 5 in household	. .	. .	. .	-0.020 (0.053)
Critical Period * # people over 65 in household	. .	. .	. .	0.159* (0.083)
Critical Period * work from home 0-10 hrs/wk	. .	. .	. .	-0.061 (0.115)
Critical Period * electric cooktop	. .	. .	. .	-0.113 (0.143)
Critical Period * electric oven	. .	. .	. .	-0.011 (0.137)
Critical Period * number of refrigerators and freezers	. .	. .	. .	-0.034 (0.056)
Critical Period * customer stayed in expt. < 4.5 months	. .	. .	. .	-0.064 (0.312)
Critical Period * customer stayed in expt. throughout expt.	. .	. .	. .	0.236 (0.148)
electric use, kWh / day, summer 2002	0.050*** (0.004)	0.051*** (0.005)	0.050*** (0.005)	. .
trt. customer on high-ratio rate	-0.016 (0.067)	0.025 (0.073)	0.008 (0.072)	. .
apartment	-0.188** (0.088)	-0.029 (0.143)	-0.088 (0.142)	. .
climate zone 4	-0.365**	-0.370**	-0.334**	.

	Specification 1: Simplest Diff in Diff	Specification 2: Adding Survey Variables	Specification 3: Adding CAC*CDH interactions	Specification 4: Adds person FE's; controls
	(0.143)	(0.144)	(0.144)	.
cooling degree hours 2-7PM, base 78	0.024*** (0.005)	0.024*** (0.005)	0.014** (0.006)	0.025*** (0.005)
cooling degree hours squared (1000's), 2-7pm	-0.100** (0.041)	-0.085** (0.043)	-0.037 (0.115)	-0.090** (0.043)
heating degree hours 2-7pm	0.037 (0.051)	0.040 (0.054)	0.056 (0.050)	0.081 (0.082)
Tuesday	-0.014 (0.012)	-0.017 (0.013)	-0.017 (0.013)	-0.016 (0.014)
Wednesday	-0.024* (0.013)	-0.030** (0.014)	-0.033** (0.014)	-0.037*** (0.014)
Thursday	-0.041*** (0.015)	-0.052*** (0.016)	-0.049*** (0.016)	-0.034** (0.017)
Friday	-0.017 (0.016)	-0.018 (0.018)	-0.020 (0.018)	-0.015 (0.018)
year 2004	-0.049 (0.045)	-0.047 (0.047)	-0.028 (0.048)	0.021 (0.053)
June	0.148*** (0.029)	0.171*** (0.032)	0.151*** (0.031)	0.095*** (0.031)
July	0.272*** (0.047)	0.310*** (0.049)	0.255*** (0.049)	0.196*** (0.043)
August	0.299*** (0.057)	0.343*** (0.060)	0.279*** (0.061)	0.235*** (0.050)
September	0.145*** (0.031)	0.173*** (0.034)	0.147*** (0.033)	0.095*** (0.033)
October	0.035	0.040	0.048	0.001

	Specification 1: Simplest Diff in Diff	Specification 2: Adding Survey Variables	Specification 3: Adding CAC*CDH interactions	Specification 4: Adds person FE's; controls
	(0.048)	(0.052)	(0.052)	(0.060)
one * heating degree hours 2-7pm squared (1000's)	-5.367 (7.473)	-3.700 (8.744)	-7.054 (7.889)	. .
Tue * cooling degree hours 2-7pm	-0.00077 (0.00083)	-0.00032 (0.00084)	-0.00040 (0.00079)	-0.001 (0.00082)
Tue * cooling degree hours 2-7pm squared (1000's)	0.003 (0.008)	-0.003 (0.008)	-0.002 (0.008)	0.006 (0.008)
Tue * heating degree hours 2-7pm	-0.013* (0.008)	-0.011 (0.008)	-0.010 (0.008)	-0.022** (0.011)
Tue * heating degree hours 2-7pm squared (1000's)	0.798* (0.437)	0.697 (0.429)	0.625 (0.408)	1.159** (0.588)
Wed * cooling degree hours 2-7pm	-0.001 (0.00084)	-0.00046 (0.00084)	-0.00055 (0.00081)	-0.00086 (0.00089)
Wed * cooling degree hours 2-7pm squared (1000's)	0.013 (0.008)	0.004 (0.008)	0.005 (0.008)	0.008 (0.009)
Wed * heating degree hours 2-7pm	0.007 (0.018)	0.005 (0.019)	0.008 (0.018)	0.00028 (0.015)
Wed * heating degree hours 2-7pm squared (1000's)	-0.525 (1.166)	-0.177 (1.089)	-0.251 (1.013)	0.070 (0.791)
Thu * cooling degree hours 2-7pm	-0.00044 (0.00094)	0.00050 (0.001)	0.000045 (0.00092)	-0.00052 (0.00089)
Thu * cooling degree hours 2-7pm squared (1000's)	0.002 (0.009)	-0.006 (0.010)	-0.001 (0.009)	0.004 (0.009)
Thu * heating degree hours 2-7pm	-0.034*** (0.010)	-0.031*** (0.011)	-0.030*** (0.011)	-0.039*** (0.011)
Thu * heating degree hours 2-7pm squared (1000's)	2.118***	1.907**	2.105***	2.936***

	Specification 1: Simplest Diff in Diff	Specification 2: Adding Survey Variables	Specification 3: Adding CAC*CDH interactions	Specification 4: Adds person FE's; controls
	(0.742)	(0.836)	(0.750)	(0.663)
Fri * cooling degree hours 2-7pm	-0.002** (0.00097)	-0.001 (0.001)	-0.002* (0.001)	-0.000079 (0.00097)
Fri * cooling degree hours 2-7pm squared (1000's)	0.015* (0.009)	0.011 (0.010)	0.018 (0.012)	-0.002 (0.010)
Fri * heating degree hours 2-7pm	-0.029*** (0.010)	-0.028** (0.011)	-0.027** (0.011)	-0.034*** (0.010)
Fri * heating degree hours 2-7pm squared (1000's)	2.036*** (0.740)	1.851** (0.833)	2.017*** (0.742)	2.818*** (0.646)
year_2004 * cooling degree hours 2-7pm	-0.004 (0.002)	-0.004* (0.002)	-0.005** (0.003)	-0.004 (0.002)
year_2004 * cooling degree hours 2-7pm squared (1000's)	0.035* (0.021)	0.038* (0.023)	0.047** (0.023)	0.038* (0.022)
year_2004 * heating degree hours 2-7pm	-0.029 (0.064)	-0.017 (0.065)	-0.00080 (0.062)	-0.039 (0.067)
year_2004 * heating degree hours 2-7pm squared (1000's)	4.084 (7.353)	1.189 (8.695)	5.045 (7.886)	-3.607 (17.346)
June * cooling degree hours 2-7pm	-0.002 (0.003)	-0.004* (0.003)	-0.003 (0.002)	0.002 (0.002)
June * cooling degree hours 2-7pm squared (1000's)	0.019 (0.031)	0.043 (0.031)	0.028 (0.028)	-0.013 (0.025)
June * heating degree hours 2-7pm	-0.033 (0.049)	-0.040 (0.053)	-0.063 (0.048)	-0.089 (0.084)
June * heating degree hours 2-7pm squared (1000's)	4.385 (7.399)	2.873 (8.675)	6.234 (7.791)	10.724 (12.594)
July * cooling degree hours 2-7pm	-0.003	-0.004	-0.001	-0.001

	Specification 1: Simplest Diff in Diff	Specification 2: Adding Survey Variables	Specification 3: Adding CAC*CDH interactions	Specification 4: Adds person FE's; controls
	(0.003)	(0.003)	(0.003)	(0.002)
July * cooling degree hours 2-7pm squared (1000's)	0.026 (0.026)	0.035 (0.027)	0.004 (0.024)	0.018 (0.026)
July * heating degree hours 2-7pm	0.017 (0.071)	0.080 (0.089)	0.00052 (0.079)	-0.182*** (0.050)
July * heating degree hours 2-7pm squared (1000's)	-9.070 (10.990)	-16.535 (13.126)	-1.704 (11.569)	3.710 (14.075)
Aug * cooling degree hours 2-7pm	-0.005* (0.003)	-0.007** (0.003)	-0.004 (0.003)	-0.003 (0.002)
Aug * cooling degree hours 2-7pm squared (1000's)	0.039 (0.025)	0.051** (0.026)	0.022 (0.024)	0.024 (0.023)
Aug * heating degree hours 2-7pm	-0.710*** (0.260)	-0.861*** (0.279)	-1.033*** (0.244)	-0.582*** (0.172)
Aug * heating degree hours 2-7pm squared (1000's)	343.860** (149.653)	447.452*** (159.862)	512.971*** (142.571)	152.020 (92.895)
Sept * cooling degree hours 2-7pm	-0.006*** (0.002)	-0.008*** (0.002)	-0.006*** (0.002)	-0.003 (0.002)
Sept * cooling degree hours 2-7pm squared (1000's)	0.052** (0.022)	0.063*** (0.023)	0.044** (0.022)	0.028 (0.022)
Sept * heating degree hours 2-7pm	-0.054 (0.074)	-0.120 (0.087)	-0.116 (0.075)	-0.134 (0.117)
Sept * heating degree hours 2-7pm squared (1000's)	-0.570 (11.572)	8.352 (13.015)	10.971 (11.202)	20.700 (16.934)
Oct * cooling degree hours 2-7pm	-0.016*** (0.003)	-0.017*** (0.003)	-0.018*** (0.003)	-0.010*** (0.003)
Oct * cooling degree hours 2-7pm squared (1000's)	0.121***	0.138***	0.147***	0.055*

	Specification 1: Simplest Diff in Diff	Specification 2: Adding Survey Variables	Specification 3: Adding CAC*CDH interactions	Specification 4: Adds person FE's; controls
	(0.030)	(0.032)	(0.034)	(0.031)
Oct * heating degree hours 2-7pm	-0.025 (0.068)	-0.027 (0.068)	-0.00098 (0.064)	-0.015 (0.069)
Oct * heating degree hours 2-7pm squared (1000's)	3.333 (7.406)	1.845 (8.492)	5.036 (7.695)	-5.471 (16.987)
constant	-0.259*** (0.092)	-0.766*** (0.220)	-0.545** (0.212)	0.694*** (0.055)
central AC	. .	0.041 (0.098)	-0.148 (0.097)	. .
room AC	. .	0.021 (0.136)	0.017 (0.133)	. .
number of bedrooms	. .	0.116* (0.069)	0.104 (0.067)	. .
# people in the household	. .	0.018 (0.024)	0.017 (0.025)	. .
cooling degree hours 2-7pm * central AC	. .	. .	0.010*** (0.002)	. .
cooling degree hours 2-7pm squared * central AC	. .	. .	-0.00026 (0.00064)	. .
heating degree hours 2-7pm squared, 1000's	. .	. .	. .	-11.161 (12.514)
N	68372	59686	59686	46191
R^2	0.4559	0.4636	0.4773	0.6289

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.

	Specification 1: Simplest Diff in Diff	Specification 2: Adding Survey Variables	Specification 3: Adding CAC*CDH interactions	Specification 4: Adds person FE's; controls

Appendix F

CPP Impacts Split into High and Low / Apartment Customers

F.1 Specification 2: Survey Variables and Specification 3: Adding CAC*CDH interactions

The table below provides complete results from the regressions presented above in tables 2.8 and 2.9.

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

	Specification 2: Adding Survey Variables		Specification 3: Adding CAC*CDH interactions	
	Low Use / Apt. Customers	High Use Customers	Low Use / Apt. Customers	High Use Customers
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 before critical price	0.001 (0.013)	0.008 (0.034)	0.003 (0.013)	0.013 (0.034)
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)

	Specification 2: Adding Survey Variables		Specification 3: Adding CAC*CDH interactions	
	Low Use / Apt. Customers	High Use Customers	Low Use / Apt. Customers	High Use Customers
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 * apartment	0.005 (0.089)	. .	0.004 (0.091)	. .
TOU Peak Price in Effect * climate zone 2	-0.016 (0.068)	-0.354** (0.140)	-0.022 (0.067)	-0.333** (0.132)
TOU Peak Price in Effect * climate zone 3	-0.177* (0.097)	-0.030 (0.244)	-0.179* (0.097)	-0.005 (0.236)
TOU Peak Price in Effect * climate zone 4	-0.039 (0.180)	-0.607* (0.361)	0.022 (0.172)	-0.632* (0.360)
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), 2-7pm	-0.071* (0.038)	-0.092 (0.064)	0.070 (0.133)	-0.204 (0.292)
TOU Peak Price in Effect * heating degree hours 2-7pm	-0.002 (0.002)	0.005 (0.006)	-0.002 (0.002)	0.004 (0.005)
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 * number of bedrooms	0.072 (0.045)	0.071 (0.075)	0.068 (0.044)	0.044 (0.070)
TOU Peak Price in Effect * # people in the household	-0.015 (0.026)	0.081* (0.045)	-0.013 (0.026)	0.069 (0.044)
TOU Peak Price in Effect * cooling degree hours 2-7pm * central AC	. .	. .	-0.00013 (0.003)	-0.003 (0.005)

	Specification 2: Adding Survey Variables		Specification 3: Adding CAC*CDH interactions	
	Low Use / Apt. Customers	High Use Customers	Low Use / Apt. Customers	High Use Customers
TOU Pk Price in Effect * cooling degree hrs 2-7pm squared * central AC	.	.	-0.00089	0.00067
	.	.	(0.00081)	(0.002)
Critical Price in Effect	-0.093	0.389	-0.189	0.470
	(0.210)	(0.467)	(0.206)	(0.451)
Critical Price in Effect * day before critical price	0.068**	0.117*	0.055*	0.100
	(0.029)	(0.062)	(0.029)	(0.063)
Critical Price in Effect * day after critical price	0.048	0.071	0.033	0.055
	(0.031)	(0.058)	(0.031)	(0.058)
Crit. Price in Effect * elec. use, kWh / day summer 2002	-0.015	-0.020	-0.012	-0.020
	(0.011)	(0.012)	(0.011)	(0.012)
Critical Price in Effect * high ratio rate customer.	0.280	0.190	0.238	0.275
	(0.173)	(0.216)	(0.162)	(0.209)
Critical Price in Effect * apartment	0.014	.	0.038	.
	(0.131)	.	(0.131)	.
Critical Price in Effect * climate zone 2	0.033	-0.265	0.017	-0.322*
	(0.083)	(0.187)	(0.080)	(0.179)
Critical Price in Effect * climate zone 3	-0.108	0.260	-0.116	0.258
	(0.141)	(0.297)	(0.144)	(0.287)
Critical Price in Effect * climate zone 4	0.019	-0.080	0.057	-0.103
	(0.236)	(0.436)	(0.224)	(0.432)
Critical Price in Effect * cooling degree hours 2-7pm	0.004	0.012	0.002	0.020
	(0.005)	(0.007)	(0.007)	(0.014)
crit. price in effect * cooling degree hours squared (1000's)	-0.058	-0.092	-0.010	-0.251
	(0.046)	(0.074)	(0.163)	(0.330)
crit. price in Effect * heating degree hours 2-7pm	0.008	-0.007	0.005	-0.005
	(0.010)	(0.017)	(0.007)	(0.015)
Critical Price in Effect * central AC	-0.102	-0.545**	-0.040	-0.272
	(0.127)	(0.227)	(0.139)	(0.250)

	Specification 2: Adding Survey Variables		Specification 3: Adding CAC*CDH interactions	
	Low Use / Apt. Customers	High Use Customers	Low Use / Apt. Customers	High Use Customers
Critical Price in Effect * room AC	0.219* (0.128)	0.534** (0.248)	0.224 (0.136)	0.520** (0.248)
Critical Price in Effect * number of bedrooms	0.044 (0.065)	0.013 (0.108)	0.051 (0.064)	-0.012 (0.105)
Critical Price in Effect * # people in the household	0.019 (0.030)	0.028 (0.052)	0.026 (0.030)	0.009 (0.053)
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)
Treatment Customer	0.181 (0.180)	0.536 (0.406)	0.104 (0.173)	0.328 (0.393)
Trt. Customer * elec. use, kWh / day summer '02	0.003 (0.008)	0.007 (0.011)	0.004 (0.008)	0.008 (0.011)
Treatment Customer * apartment	-0.039 (0.088)	. .	-0.015 (0.086)	. .
Treatment Customer * climate zone 2	-0.059 (0.069)	0.276 (0.200)	-0.059 (0.066)	0.180 (0.191)
Treatment Customer * climate zone 3	0.044 (0.086)	0.231 (0.307)	0.064 (0.083)	0.184 (0.298)
Treatment Customer * climate zone 4	0.012 (0.183)	0.923** (0.430)	-0.023 (0.173)	0.949** (0.427)
Treatment Customer * cooling degree hours 2-7pm	-0.007 (0.005)	-0.018** (0.007)	-0.00098 (0.005)	-0.021* (0.012)
Treatment Customer * cooling degree hours squared (1000's)	0.058 (0.042)	0.088 (0.072)	-0.076 (0.115)	0.291 (0.276)
Treatment Customer * heating degree hours 2-7pm	0.001 (0.002)	-0.004 (0.006)	0.00088 (0.002)	-0.003 (0.005)

	Specification 2: Adding Survey Variables		Specification 3: Adding CAC*CDH interactions	
	Low Use / Apt. Customers	High Use Customers	Low Use / Apt. Customers	High Use Customers
Treatment Customer * central AC	-0.055 (0.076)	-0.281 (0.190)	-0.060 (0.081)	-0.199 (0.181)
Treatment Customer * room AC	0.034 (0.085)	0.150 (0.218)	0.035 (0.082)	0.183 (0.220)
Treatment Customer * number of bedrooms	-0.103** (0.045)	-0.169* (0.088)	-0.085** (0.043)	-0.132 (0.086)
Treatment Customer * # people in the household	0.034 (0.024)	-0.076* (0.045)	0.040 (0.024)	-0.064 (0.045)
Trt. Cust. * cooling degree hours 2-7pm * central AC	. .	. .	-0.001 (0.003)	-0.007 (0.006)
Trt. Cust. * cooling degree hours 2-7pm squared * central AC	. .	. .	0.00086 (0.00074)	-0.001 (0.002)
Treatment Period (after 7/1/2003)	0.042 (0.116)	0.116 (0.294)	0.056 (0.115)	0.006 (0.276)
Treatment Period * electricity use, kWh / day summer 2002	0.004 (0.006)	0.008 (0.007)	0.002 (0.006)	0.008 (0.007)
Treatment Period * apartment	0.016 (0.066)	. .	0.013 (0.068)	. .
Treatment Period * climate zone 2	-0.056 (0.050)	0.273** (0.113)	-0.031 (0.048)	0.261** (0.106)
Treatment Period * climate zone 3	0.077 (0.068)	-0.063 (0.199)	0.098 (0.068)	-0.065 (0.190)
Treatment Period * climate zone 4	-0.118 (0.132)	0.163 (0.256)	-0.099 (0.127)	0.239 (0.250)
Treatment Period * cooling degree hours 2-7pm	0.002 (0.004)	-0.008 (0.006)	0.006 (0.004)	-0.003 (0.009)
Treatment Pd. * cooling degree hours squared (1000's), 2-7pm	-0.019 (0.032)	0.059 (0.055)	-0.147 (0.094)	-0.034 (0.196)

	Specification 2: Adding Survey Variables		Specification 3: Adding CAC*CDH interactions	
	Low Use / Apt. Customers	High Use Customers	Low Use / Apt. Customers	High Use Customers
Treatment Period * heating degree hours 2-7pm	-0.001 (0.003)	-0.001 (0.009)	-0.00051 (0.003)	-0.00080 (0.008)
Treatment Period * central AC	0.097 (0.060)	0.138 (0.146)	0.028 (0.067)	0.099 (0.146)
Treatment Period * room AC	-0.068 (0.067)	-0.104 (0.107)	-0.062 (0.068)	-0.112 (0.104)
Treatment Period * number of bedrooms	-0.032 (0.035)	-0.059 (0.054)	-0.033 (0.035)	-0.040 (0.047)
Treatment Period * # people in the household	0.011 (0.021)	-0.034 (0.033)	0.013 (0.020)	-0.023 (0.032)
Treatment Pd. * cooling degree hours 2-7pm * central AC	. .	. .	0.00052 (0.002)	-0.002 (0.004)
Trt. Pd. * cooling degree hours 2-7pm squared * central AC	. .	. .	0.00090 (0.00057)	0.00061 (0.001)
Critical Period	-0.187* (0.108)	-0.844*** (0.203)	-0.107 (0.097)	-0.792*** (0.201)
Critical Period * electricity use, kWh / day summer 2002	0.015*** (0.005)	0.013** (0.005)	0.013*** (0.005)	0.012** (0.005)
Critical Period * high ratio rate customer.	-0.204 (0.148)	-0.305 (0.190)	-0.157 (0.136)	-0.387** (0.180)
Critical Period * apartment	0.030 (0.066)	. .	0.006 (0.063)	. .
Critical Period * climate zone 2	-0.025 (0.032)	-0.080 (0.082)	-0.019 (0.030)	-0.022 (0.076)
Critical Period * climate zone 3	0.114 (0.070)	-0.160 (0.160)	0.097 (0.072)	-0.165 (0.150)
Critical Period * climate zone 4	0.061 (0.106)	-0.191 (0.207)	0.038 (0.101)	-0.231 (0.197)

	Specification 2: Adding Survey Variables		Specification 3: Adding CAC*CDH interactions	
	Low Use / Apt. Customers	High Use Customers	Low Use / Apt. Customers	High Use Customers
Critical Period * cooling degree hours 2-7pm	-0.001 (0.002)	-0.004 (0.003)	0.00069 (0.003)	0.003 (0.005)
Critical Period * cooling degree hours squared (1000's), 2-7pm	-0.018 (0.014)	-0.00013 (0.027)	-0.011 (0.067)	-0.030 (0.096)
Critical Period * heating degree hours 2-7pm	-0.009 (0.007)	0.008 (0.015)	-0.006 (0.005)	0.004 (0.014)
Critical Period * central AC	0.211*** (0.057)	0.576*** (0.100)	0.072 (0.060)	0.329*** (0.126)
Critical Period * room AC	-0.051 (0.067)	-0.016 (0.130)	-0.048 (0.073)	-0.024 (0.130)
Critical Period * number of bedrooms	0.018 (0.028)	0.074* (0.043)	0.010 (0.027)	0.070 (0.043)
Critical Period * # people in the household	-0.022 (0.015)	0.044** (0.021)	-0.023 (0.014)	0.052** (0.022)
Critical Period * cooling degree hours 2-7pm * central AC	. .	. .	-0.001 (0.00100)	-0.002 (0.002)
Critical Period * cooling degree hours 2-7pm squared * central AC	. .	. .	-0.000085 (0.00036)	-0.000012 (0.00051)
electricity use, kWh / day summer 2002	0.041*** (0.006)	0.041*** (0.009)	0.040*** (0.006)	0.041*** (0.009)
trt. customer on high-ratio rate	-0.004 (0.040)	0.080 (0.090)	-0.006 (0.039)	0.081 (0.090)
apartment	0.079 (0.077)	. .	0.045 (0.076)	. .
climate zone 2	0.043 (0.056)	-0.257 (0.179)	0.060 (0.053)	-0.124 (0.171)
climate zone 3	-0.00030 (0.066)	-0.176 (0.275)	0.023 (0.064)	-0.092 (0.266)

	Specification 2: Adding Survey Variables		Specification 3: Adding CAC*CDH interactions	
	Low Use / Apt. Customers	High Use Customers	Low Use / Apt. Customers	High Use Customers
climate zone 4	-0.046 (0.150)	-0.627* (0.358)	-0.032 (0.143)	-0.696** (0.351)
cooling degree hours 2-7PM,	0.010** (0.004)	0.032*** (0.007)	0.003 (0.004)	0.024** (0.010)
heating degree hours 2-7pm	0.007* (0.004)	0.005 (0.010)	0.005 (0.004)	-0.001 (0.009)
central AC	0.077 (0.057)	0.284* (0.164)	-0.022 (0.061)	0.083 (0.154)
room AC	-0.011 (0.070)	0.040 (0.169)	-0.005 (0.069)	0.027 (0.173)
number of bedrooms	0.087** (0.037)	0.117* (0.067)	0.072** (0.034)	0.086 (0.065)
# people in the household	0.010 (0.022)	0.059* (0.033)	0.008 (0.021)	0.054* (0.032)
Tuesday	-0.006 (0.008)	-0.001 (0.016)	-0.006 (0.008)	0.002 (0.016)
Wednesday	-0.003 (0.008)	-0.011 (0.018)	-0.004 (0.008)	-0.009 (0.018)
Thursday	-0.008 (0.011)	-0.044** (0.018)	-0.008 (0.011)	-0.037** (0.018)
Friday	0.003 (0.010)	-0.008 (0.024)	0.001 (0.010)	-0.006 (0.024)
year 2004	-0.018 (0.028)	-0.051 (0.060)	-0.016 (0.029)	-0.043 (0.060)
June	0.041*** (0.014)	0.172*** (0.037)	0.034** (0.014)	0.153*** (0.037)
July	0.086*** (0.020)	0.370*** (0.054)	0.072*** (0.021)	0.324*** (0.053)

	Specification 2: Adding Survey Variables		Specification 3: Adding CAC*CDH interactions	
	Low Use / Apt. Customers	High Use Customers	Low Use / Apt. Customers	High Use Customers
August	0.103*** (0.023)	0.420*** (0.066)	0.087*** (0.023)	0.367*** (0.066)
September	0.080*** (0.020)	0.203*** (0.047)	0.069*** (0.021)	0.174*** (0.047)
October	0.043 (0.029)	0.045 (0.072)	0.039 (0.029)	0.037 (0.072)
heating degree hours 2-7pm squared (1000's)	-0.181** (0.081)	0.023 (0.192)	-0.136* (0.074)	0.138 (0.178)
Tue * cooling degree hours 2-7pm	-0.00045 (0.00076)	-0.001 (0.001)	-0.00053 (0.00071)	-0.002* (0.001)
Tue * cooling degree hours 2-7pm squared (1000's)	-0.004 (0.008)	0.009 (0.012)	-0.003 (0.007)	0.017 (0.012)
Tue * heating degree hours 2-7pm	0.00058 (0.002)	-0.009** (0.004)	0.00096 (0.001)	-0.009** (0.004)
Tue * heating degree hours 2-7pm squared (1000's)	0.015 (0.027)	0.130* (0.073)	0.006 (0.024)	0.121* (0.070)
Wed * cooling degree hours 2-7pm	-0.00067 (0.00079)	-0.003** (0.001)	-0.00100 (0.00078)	-0.004*** (0.001)
Wed * cooling degree hours 2-7pm squared (1000's)	0.003 (0.008)	0.032** (0.015)	0.006 (0.008)	0.038** (0.015)
Wed * heating degree hours 2-7pm	0.00058 (0.002)	0.002 (0.005)	0.00086 (0.002)	0.003 (0.005)
Wed * heating degree hours 2-7pm squared (1000's)	0.008 (0.035)	-0.022 (0.101)	-0.00027 (0.032)	-0.038 (0.099)
Thu * cooling degree hours 2-7pm	-0.00075 (0.00089)	-0.002 (0.001)	-0.001 (0.00084)	-0.003** (0.001)
Thu * cooling degree hours 2-7pm squared (1000's)	0.001 (0.009)	0.021 (0.014)	0.005 (0.009)	0.035** (0.014)

	Specification 2: Adding Survey Variables		Specification 3: Adding CAC*CDH interactions	
	Low Use / Apt. Customers	High Use Customers	Low Use / Apt. Customers	High Use Customers
Thu * heating degree hours 2-7pm	-0.001 (0.002)	-0.009 (0.005)	-0.00088 (0.002)	-0.008 (0.005)
Thu * heating degree hours 2-7pm squared (1000's)	0.050 (0.045)	0.223** (0.102)	0.036 (0.037)	0.191** (0.093)
Fri * cooling degree hours 2-7pm	-0.002* (0.00098)	-0.002* (0.002)	-0.002** (0.00097)	-0.004** (0.002)
Fri * cooling degree hours 2-7pm squared (1000's)	0.013 (0.010)	0.017 (0.016)	0.018* (0.010)	0.035* (0.021)
Fri * heating degree hours 2-7pm	-0.004* (0.002)	-0.013** (0.006)	-0.003 (0.002)	-0.011** (0.005)
Fri * heating degree hours 2-7pm squared (1000's)	0.091** (0.043)	0.211** (0.099)	0.072* (0.037)	0.173* (0.093)
year 2004 * cooling degree hours 2-7pm	-0.005*** (0.002)	-0.004 (0.004)	-0.005*** (0.002)	-0.005 (0.004)
June * cooling degree hours 2-7pm	-0.002 (0.002)	-0.004 (0.004)	-0.003 (0.002)	-0.002 (0.004)
June * cooling degree hours 2-7pm squared (1000's)	0.026 (0.031)	0.043 (0.052)	0.025 (0.027)	0.019 (0.048)
June * heating degree hours 2-7pm	-0.004 (0.002)	-0.002 (0.007)	-0.003 (0.002)	-0.00062 (0.007)
June * heating degree hours 2-7pm squared (1000's)	0.065* (0.036)	0.014 (0.094)	0.057* (0.035)	-0.007 (0.096)
July * cooling degree hours 2-7pm	-0.003 (0.002)	-0.002 (0.005)	-0.002 (0.002)	0.002 (0.004)
July * cooling degree hours 2-7pm squared (1000's)	0.030 (0.024)	0.009 (0.046)	0.018 (0.022)	-0.030 (0.041)
July * heating degree hours 2-7pm	-0.009** (0.004)	-0.035*** (0.012)	-0.007* (0.004)	-0.031*** (0.011)

	Specification 2: Adding Survey Variables		Specification 3: Adding CAC*CDH interactions	
	Low Use / Apt. Customers	High Use Customers	Low Use / Apt. Customers	High Use Customers
July * heating degree hours 2-7pm squared (1000's)	0.141 (0.123)	1.100** (0.436)	0.124 (0.113)	1.055** (0.422)
Aug * cooling degree hours 2-7pm	-0.003 (0.002)	-0.004 (0.004)	-0.002 (0.002)	-0.00086 (0.004)
Aug * cooling degree hours 2-7pm squared (1000's)	0.032 (0.023)	0.021 (0.041)	0.021 (0.022)	-0.015 (0.038)
Aug * heating degree hours 2-7pm	-0.024** (0.010)	-0.070*** (0.025)	-0.020** (0.008)	-0.064*** (0.025)
Aug * heating degree hours 2-7pm squared (1000's)	0.680* (0.410)	2.047* (1.063)	0.562 (0.344)	1.984* (1.032)
Sept * cooling degree hours 2-7pm	-0.007*** (0.002)	-0.007* (0.003)	-0.005*** (0.002)	-0.004 (0.003)
Sept * cooling degree hours 2-7pm squared (1000's)	0.059*** (0.020)	0.040 (0.036)	0.042** (0.020)	0.016 (0.035)
Sept * heating degree hours 2-7pm	-0.004 (0.003)	-0.00023 (0.007)	-0.003 (0.003)	0.003 (0.007)
Sept * heating degree hours 2-7pm squared (1000's)	0.069 (0.042)	-0.073 (0.095)	0.052 (0.036)	-0.111 (0.095)
Oct * cooling degree hours 2-7pm	-0.013*** (0.003)	-0.020*** (0.005)	-0.012*** (0.003)	-0.021*** (0.005)
Oct * cooling degree hours 2-7pm squared (1000's)	0.103*** (0.029)	0.139*** (0.053)	0.090*** (0.031)	0.149** (0.058)
Oct * heating degree hours 2-7pm	-0.002 (0.003)	0.010 (0.008)	-0.00056 (0.003)	0.015* (0.008)
Oct * heating degree hours 2-7pm squared (1000's)	0.090 (0.064)	-0.253 (0.167)	0.061 (0.059)	-0.333** (0.161)
constant	-0.369** (0.166)	-0.819** (0.349)	-0.278* (0.160)	-0.606* (0.335)

	Specification 2: Adding Survey Variables		Specification 3: Adding CAC*CDH interactions	
	Low Use / Apt. Customers	High Use Customers	Low Use / Apt. Customers	High Use Customers
cooling degree hours squared (1000's), 2-7pm	-0.015 (0.037)	-0.076 (0.073)	0.068 (0.091)	-0.222 (0.201)
cooling degree hours 2-7pm * central AC	. .	. .	0.009*** (0.003)	0.019*** (0.004)
cooling degree hours 2-7pm squared * central AC	. .	. .	-0.00068 (0.00056)	0.00065 (0.001)
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.

F.2 Specification 1: Basic Difference-in-Difference and Specification 4 Fixed Effects and Additional Controls

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

	Specification 1: Simplest Diff in Diff		Specification 4: Adds person FE's; controls	
	Low Use / Apt. Customers	High Use Customers	Low Use / Apt. Customers	High Use Customers
TOU Peak Price in Effect	-0.012 (0.093)	0.341 (0.244)	0.318 (0.201)	-0.347 (0.583)
TOU peak price in effect * day before critical price	0.002 (0.012)	0.005 (0.033)	0.002 (0.013)	0.003 (0.031)
TOU peak price in effect * day after critical price	0.035** (0.014)	0.004 (0.029)	0.026 (0.016)	-0.010 (0.026)
TOU Peak Price in Effect * elec. use, kWh / day, summer '02	-0.002	-0.005	0.003	-0.002

	Specification 1: Simplest Diff in Diff		Specification 4: Adds person FE's; controls	
	Low Use / Apt. Customers	High Use Customers	Low Use / Apt. Customers	High Use Customers
	(0.007)	(0.010)	(0.009)	(0.011)
TOU Peak Price in Effect * high ratio rate customer	0.058 (0.040)	-0.206** (0.092)	0.058 (0.056)	-0.105 (0.105)
TOU Peak Price in Effect * apartment	-0.049 (0.062)	. .	0.075 (0.104)	. .
TOU Peak Price in Effect * climate zone 2	0.031 (0.054)	-0.298** (0.121)	-0.031 (0.082)	-0.179 (0.251)
TOU Peak Price in Effect * climate zone 3	-0.114 (0.074)	-0.038 (0.188)	-0.208* (0.112)	0.101 (0.295)
TOU Peak Price in Effect * climate zone 4	0.006 (0.160)	-0.783*** (0.295)	0.039 (0.208)	-0.027 (0.416)
TOU Peak Price in Effect * cooling degree hours 2-7pm	0.008** (0.004)	0.012** (0.006)	0.006 (0.005)	0.011* (0.007)
TOU Peak Price in Effect * cooling degree hours squared (1000's), 2-7pm	-0.096*** (0.036)	-0.100 (0.063)	-0.102** (0.041)	-0.080 (0.066)
TOU Peak Price in Effect * heating degree hours 2-7pm	-0.002 (0.002)	0.003 (0.005)	-0.001 (0.006)	-0.039** (0.016)
TOU Pk. Price in Effect * heating deg. hrs 2-7pm squared (1000's)	. .	. .	-0.020 (0.123)	0.869* (0.453)
TOU Peak Price in Effect * cooling degree hrs 2-7PM, previous day	. .	. .	-0.001* (0.00066)	-0.00092 (0.001)
TOU Pk Price in Effect * cooling degree hrs 2-7PM, two days before	. .	. .	0.00048 (0.00051)	-0.001 (0.00090)
TOU Pk Price in Effect * cooling deg hrs 2-7PM, three days before	. .	. .	-0.00074 (0.00051)	-0.002** (0.00087)
TOU Peak Price in Effect * central AC	. .	. .	0.015 (0.084)	-0.073 (0.178)
TOU Peak Price in Effect * room AC	.	.	-0.048	-0.046

	Specification 1: Simplest Diff in Diff		Specification 4: Adds person FE's; controls	
	Low Use / Apt. Customers	High Use Customers	Low Use / Apt. Customers	High Use Customers
	.	.	(0.118)	(0.251)
TOU Peak Price in Effect * number of bedrooms	. .	. .	0.092** (0.039)	-0.024 (0.090)
TOU Peak Price in Effect * # people in the household	. .	. .	0.005 (0.044)	0.191*** (0.066)
TOU Peak Price in Effect * work from home 11-30 hrs/wk	. .	. .	0.191* (0.109)	-0.273 (0.204)
TOU Peak Price in Effect * work from home >30 hrs/wk	. .	. .	-0.646*** (0.205)	-0.020 (0.319)
TOU Peak Price in Effect * swimming pool	. .	. .	-0.105 (0.148)	-0.224 (0.232)
TOU Peak Price in Effect * spa	. .	. .	0.396*** (0.150)	-0.207 (0.229)
TOU Pk Price in Effect * cooling degree hours 2-7pm * room AC	. .	. .	0.011*** (0.003)	0.001 (0.006)
TOU Pk Price in Effect * heating deg. hrs 2-7PM * elec. heat	. .	. .	0.010*** (0.003)	0.006 (0.006)
TOU Peak Price in Effect * electric heat	. .	. .	-0.202* (0.104)	-0.254 (0.279)
TOU Peak Price in Effect * # kids under 5 in household	. .	. .	-0.081 (0.092)	-0.156 (0.118)
TOU Peak Price in Effect * # kids over 5 in household	. .	. .	-0.023 (0.060)	-0.152* (0.087)
TOU Peak Price in Effect * # people over 65 in household	. .	. .	-0.023 (0.062)	-0.319*** (0.116)
TOU Peak Price in Effect * work from home 0-10 hrs/wk	. .	. .	-0.084 (0.104)	0.451 (0.291)
TOU Peak Price in Effect * electric cooktop	.	.	0.190	0.245

	Specification 1: Simplest Diff in Diff		Specification 4: Adds person FE's; controls	
	Low Use / Apt. Customers	High Use Customers	Low Use / Apt. Customers	High Use Customers
	.	.	(0.134)	(0.267)
TOU Peak Price in Effect * electric oven	. .	. .	-0.149 (0.125)	-0.112 (0.243)
TOU Peak Price in Effect * # of refrigerators and freezers	. .	. .	-0.228** (0.101)	-0.031 (0.121)
TOU Peak Price in Effect * customer in expt. < 4.5 months	. .	. .	-0.136 (0.115)	0.764* (0.424)
TOU Peak Price in Effect * customer stayed in expt. throughout expt.	. .	. .	-0.213** (0.086)	-0.277 (0.254)
Critical Price in Effect	0.017 (0.129)	0.478 (0.306)	0.447* (0.266)	0.984 (0.756)
Critical Price in Effect * day before critical price	0.075*** (0.026)	0.087 (0.061)	0.039 (0.037)	0.076 (0.071)
Critical Price in Effect * day after critical price	0.047* (0.027)	0.061 (0.055)	-0.004 (0.036)	0.056 (0.061)
Crit. Price in Effect * elec. use, kWh / day, summer '02	-0.014 (0.010)	-0.018 (0.013)	-0.015 (0.013)	-0.020 (0.015)
Critical Price in Effect * high ratio rate customer.	0.244 (0.157)	0.276 (0.195)	0.263** (0.118)	-0.283 (0.236)
Critical Price in Effect * apartment	-0.025 (0.091)	. .	0.088 (0.153)	. .
Critical Price in Effect * climate zone 2	0.089 (0.066)	-0.350* (0.184)	0.004 (0.116)	-0.003 (0.349)
Critical Price in Effect * climate zone 3	-0.094 (0.118)	-0.108 (0.274)	-0.134 (0.167)	0.596 (0.408)
Critical Price in Effect * climate zone 4	0.006 (0.208)	-0.635 (0.386)	0.003 (0.247)	0.900 (0.553)
Critical Price in Effect * cooling degree hours 2-7pm	0.008*	0.011	0.007	0.009

	Specification 1: Simplest Diff in Diff		Specification 4: Adds person FE's; controls	
	Low Use / Apt. Customers	High Use Customers	Low Use / Apt. Customers	High Use Customers
	(0.004)	(0.007)	(0.005)	(0.008)
Crit. Price in Effect * cooling degree hours squared 1000's, 2-7pm	-0.106** (0.044)	-0.100 (0.071)	-0.102** (0.050)	-0.077 (0.078)
Critical Price in Effect * heating degree hours 2-7pm	0.006 (0.006)	-0.006 (0.011)	-0.036* (0.018)	-0.017 (0.045)
Crit. Price in Effect * heating degree hours 2-7pm squared 1000's	. .	. .	1.572*** (0.608)	-0.027 (1.752)
Critical Price in Effect * cooling degree hours 2-7PM, previous day	. .	. .	-0.002 (0.001)	-0.004 (0.003)
Crit. Price in Effect * cooling degree hours 2-7PM, two days before	. .	. .	0.002 (0.001)	-0.000070 (0.003)
Crit. Price in Effect * cooling degree hours 2-7PM, three days before	. .	. .	-0.001 (0.001)	-0.002 (0.002)
Critical Price in Effect * central AC	. .	. .	-0.120 (0.128)	-0.668** (0.269)
Critical Price in Effect * room AC	. .	. .	-0.173 (0.164)	0.052 (0.369)
Critical Price in Effect * number of bedrooms	. .	. .	0.050 (0.057)	-0.077 (0.120)
Critical Price in Effect * # people in the household	. .	. .	0.055 (0.051)	0.184** (0.094)
Critical Price in Effect * work from home 11-30 hrs/wk	. .	. .	0.386* (0.229)	-0.473 (0.314)
Critical Price in Effect * work from home >30 hrs/wk	. .	. .	-0.137 (0.293)	0.168 (0.385)
Critical Price in Effect * swimming pool	. .	. .	-0.269 (0.269)	-0.195 (0.284)
Critical Price in Effect * spa	.	.	0.652**	-0.313

	Specification 1: Simplest Diff in Diff		Specification 4: Adds person FE's; controls	
	Low Use / Apt. Customers	High Use Customers	Low Use / Apt. Customers	High Use Customers
	.	.	(0.266)	(0.302)
Critical Price in Effect * cooling degree hours 2-7pm * room AC	. .	. .	0.010*** (0.003)	0.005 (0.007)
Critical Price in Effect * heating degree hours 2-7PM* electric heat	. .	. .	0.056** (0.027)	-0.006 (0.044)
Critical Price in Effect * electric heat	. .	. .	-0.347** (0.164)	0.037 (0.328)
Critical Price in Effect * # kids under 5 in household	. .	. .	-0.189* (0.111)	-0.309* (0.168)
Critical Price in Effect * # kids over 5 in household	. .	. .	-0.009 (0.073)	-0.166 (0.125)
Critical Price in Effect * # people over 65 in household	. .	. .	-0.107 (0.094)	-0.552*** (0.166)
Critical Price in Effect * work from home 0-10 hrs/wk	. .	. .	-0.162 (0.169)	0.637* (0.384)
Critical Price in Effect * electric cooktop	. .	. .	0.317 (0.213)	0.244 (0.377)
Critical Price in Effect * electric oven	. .	. .	-0.131 (0.193)	-0.247 (0.326)
Critical Price in Effect * number of refrigerators and freezers	. .	. .	-0.340** (0.133)	-0.190 (0.160)
Critical Price in Effect * customer stayed in expt. < 4.5 months	. .	. .	0.036 (0.155)	0.857 (0.546)
Critical Price in Effect * customer stayed in expt. throughout expt.	. .	. .	-0.311** (0.121)	-0.728* (0.395)
Treatment Customer	-0.038 (0.085)	-0.247 (0.257)	. .	. .
Treatment Customer * electricity use, kWh / day, summer 2002	0.00073	0.005	.	.

	Specification 1: Simplest Diff in Diff		Specification 4: Adds person FE's; controls	
	Low Use / Apt. Customers	High Use Customers	Low Use / Apt. Customers	High Use Customers
	(0.007)	(0.009)	.	.
Treatment Customer * apartment	0.074 (0.056)	. .	. .	. .
Treatment Customer * climate zone 2	0.012 (0.051)	0.048 (0.166)	. .	. .
Treatment Customer * climate zone 3	0.085 (0.066)	-0.055 (0.225)	. .	. .
Treatment Customer * climate zone 4	0.029 (0.182)	0.640* (0.340)	. .	. .
Treatment Customer * cooling degree hours 2-7pm	-0.010** (0.005)	-0.017** (0.007)	-0.008 (0.006)	-0.019** (0.008)
Trt. Customer * cooling degree hours squared (1000's), 2-7pm	0.093** (0.040)	0.102 (0.069)	0.081* (0.044)	0.105 (0.069)
Treatment Customer * heating degree hours 2-7pm	0.00039 (0.002)	-0.002 (0.005)	-0.003 (0.005)	0.038** (0.016)
Trt. Cust. * heating degree hours 2-7pm squared (1000's)	. .	. .	0.005 (0.122)	-0.873* (0.454)
Treatment Period (after 7/1/2003)	-0.007 (0.058)	-0.198 (0.182)	-0.255* (0.136)	-0.053 (0.407)
Treatment Period * electricity use, kWh / day, summer 2002	0.005 (0.005)	0.009 (0.007)	0.009 (0.007)	0.004 (0.007)
Treatment Period * apartment	0.035 (0.043)	. .	-0.015 (0.072)	. .
Treatment Period * climate zone 2	-0.065 (0.043)	0.254*** (0.094)	-0.026 (0.062)	0.195 (0.182)
Treatment Period * climate zone 3	0.074 (0.055)	0.022 (0.136)	0.052 (0.079)	-0.142 (0.205)
Treatment Period * climate zone 4	-0.079	0.360*	-0.396**	0.049

	Specification 1: Simplest Diff in Diff		Specification 4: Adds person FE's; controls	
	Low Use / Apt. Customers	High Use Customers	Low Use / Apt. Customers	High Use Customers
	(0.125)	(0.189)	(0.170)	(0.317)
Treatment Period * cooling degree hours 2-7pm	0.001 (0.003)	-0.007 (0.005)	0.002 (0.005)	-0.008 (0.006)
Treatment Period * cooling degree hours squared (1000's), 2-7pm	-0.007 (0.031)	0.052 (0.052)	-0.017 (0.038)	0.018 (0.055)
Treatment Period * heating degree hours 2-7pm	-0.001 (0.003)	-0.004 (0.007)	0.007 (0.006)	0.021 (0.021)
Trt. Pd. * heating degree hours 2-7pm squared (1000's)	. .	. .	-0.031 (0.161)	0.315 (0.816)
Treatment Period * cooling degree hours 2-7PM, previous day	. .	. .	0.002*** (0.00050)	0.005*** (0.001)
Treatment Period * cooling degree hours 2-7PM, two days before	. .	. .	0.00061 (0.00047)	0.002*** (0.00074)
Treatment Period * cooling degree hours 2-7PM, three days before	. .	. .	0.00062 (0.00038)	0.002*** (0.00065)
Treatment Period * central AC	. .	. .	0.125** (0.060)	0.134 (0.150)
Treatment Period * room AC	. .	. .	-0.022 (0.096)	-0.099 (0.177)
Treatment Period * number of bedrooms	. .	. .	-0.058** (0.028)	0.007 (0.074)
Treatment Period * # people in the household	. .	. .	-0.011 (0.032)	-0.107** (0.049)
Treatment Period * work from home 11-30 hrs/wk	. .	. .	-0.157*** (0.056)	0.222 (0.177)
Treatment Period * work from home >30 hrs/wk	. .	. .	0.616*** (0.175)	-0.137 (0.181)
Treatment Period * swimming pool	.	.	0.144	0.099

	Specification 1: Simplest Diff in Diff		Specification 4: Adds person FE's; controls	
	Low Use / Apt. Customers	High Use Customers	Low Use / Apt. Customers	High Use Customers
	.	.	(0.112)	(0.187)
Treatment Period * spa	. .	. .	-0.311*** (0.099)	0.145 (0.194)
Treatment Period * cooling degree hours 2-7pm * room AC	. .	. .	-0.007*** (0.002)	-0.00017 (0.004)
Treatment Period * heating degree hours 2-7PM* electric heat	. .	. .	-0.008*** (0.002)	-0.001 (0.004)
Treatment Period * electric heat	. .	. .	0.196** (0.078)	0.031 (0.152)
Treatment Period * # kids under 5 in household	. .	. .	0.124* (0.068)	0.134 (0.088)
Treatment Period * # kids over 5 in household	. .	. .	0.070* (0.038)	0.092 (0.063)
Treatment Period * # people over 65 in household	. .	. .	0.124** (0.051)	0.303*** (0.079)
Treatment Period * work from home 0-10 hrs/wk	. .	. .	0.177** (0.070)	-0.147 (0.189)
Treatment Period * electric cooktop	. .	. .	-0.057 (0.108)	-0.356* (0.209)
Treatment Period * electric oven	. .	. .	-0.094 (0.098)	0.050 (0.178)
Treatment Period * number of refrigerators and freezers	. .	. .	0.092 (0.065)	-0.043 (0.092)
Treatment Period * customer stayed in expt. $<$ 4.5 months	. .	. .	0.027 (0.094)	-0.439 (0.310)
Treatment Period * customer stayed in expt. throughout expt.	. .	. .	0.124* (0.068)	0.437** (0.204)
Critical Period	-0.192***	-0.396***	-0.068	-1.280***

	Specification 1: Simplest Diff in Diff		Specification 4: Adds person FE's; controls	
	Low Use / Apt. Customers	High Use Customers	Low Use / Apt. Customers	High Use Customers
	(0.056)	(0.128)	(0.123)	(0.393)
Critical Period * electricity use, kWh / day, summer 2002	0.017*** (0.004)	0.014*** (0.005)	0.022*** (0.007)	0.016*** (0.004)
Critical Period * high ratio rate customer.	-0.154 (0.137)	-0.478*** (0.154)	-0.109 (0.075)	0.248 (0.198)
Critical Period * apartment	0.005 (0.041)	. .	0.012 (0.086)	. .
Critical Period * climate zone 2	0.010 (0.028)	0.154* (0.091)	0.003 (0.057)	-0.018 (0.163)
Critical Period * climate zone 3	0.200*** (0.059)	0.273** (0.136)	0.101 (0.084)	-0.317* (0.184)
Critical Period * climate zone 4	0.176* (0.106)	0.280 (0.183)	0.061 (0.143)	-0.481** (0.237)
Critical Period * cooling degree hours 2-7pm	-0.002 (0.001)	-0.002 (0.003)	-0.003** (0.002)	-0.004 (0.003)
Critical Period * cooling degree hours squared (1000's), 2-7pm	-0.007 (0.014)	-0.009 (0.023)	0.00099 (0.015)	0.015 (0.028)
Critical Period * heating degree hours 2-7pm	-0.002 (0.004)	0.002 (0.007)	0.012 (0.013)	0.024 (0.029)
Critical Period * heating degree hours 2-7pm squared (1000's)	. .	. .	-0.583 (0.440)	-0.935 (1.120)
Critical Period * cooling degree hours 2-7PM, previous day	. .	. .	-0.002* (0.001)	-0.00019 (0.002)
Critical Period * cooling degree hours 2-7PM, two days before	. .	. .	-0.00048 (0.001)	-0.00053 (0.003)
Critical Period * cooling degree hours 2-7PM, three days before	. .	. .	0.002* (0.001)	0.001 (0.002)
Critical Period * central AC	.	.	0.207***	0.648***

	Specification 1: Simplest Diff in Diff		Specification 4: Adds person FE's; controls	
	Low Use / Apt. Customers	High Use Customers	Low Use / Apt. Customers	High Use Customers
	.	.	(0.066)	(0.122)
Critical Period * room AC	.	.	0.133	0.128
	.	.	(0.109)	(0.149)
Critical Period * number of bedrooms	.	.	-0.00025	0.062
	.	.	(0.036)	(0.056)
Critical Period * # people in the household	.	.	-0.030	0.036
	.	.	(0.019)	(0.040)
Critical Period * work from home 11-30 hrs/wk	.	.	0.034	0.195
	.	.	(0.131)	(0.159)
Critical Period * work from home >30 hrs/wk	.	.	-0.371**	0.060
	.	.	(0.162)	(0.116)
Critical Period * swimming pool	.	.	0.193	-0.031
	.	.	(0.214)	(0.097)
Critical Period * spa	.	.	-0.137	-0.020
	.	.	(0.106)	(0.119)
Critical Period * cooling degree hours 2-7pm * room AC	.	.	-0.000063	-0.002
	.	.	(0.002)	(0.003)
Critical Period * heating degree hours 2-7PM* electric heat	.	.	-0.00010	.
	.	.	(0.010)	.
Critical Period * electric heat	.	.	-0.046	-0.148
	.	.	(0.087)	(0.102)
Critical Period * # kids under 5 in household	.	.	0.042	0.046
	.	.	(0.050)	(0.077)
Critical Period * # kids over 5 in household	.	.	-0.020	-0.028
	.	.	(0.025)	(0.062)
Critical Period * # people over 65 in household	.	.	0.120*	0.149*
	.	.	(0.064)	(0.082)
Critical Period * work from home 0-10 hrs/wk	.	.	0.052	-0.110

	Specification 1: Simplest Diff in Diff		Specification 4: Adds person FE's; controls	
	Low Use / Apt. Customers	High Use Customers	Low Use / Apt. Customers	High Use Customers
	.	.	(0.078)	(0.119)
Critical Period * electric cooktop	. .	. .	-0.036 (0.086)	-0.158 (0.167)
Critical Period * electric oven	. .	. .	-0.031 (0.074)	0.108 (0.136)
Critical Period * number of refrigerators and freezers	. .	. .	0.007 (0.062)	0.024 (0.059)
Critical Period * customer stayed in expt. < 4.5 months	. .	. .	-0.226** (0.098)	-0.141 (0.316)
Critical Period * customer stayed in expt. throughout expt.	. .	. .	0.093 (0.063)	0.430** (0.195)
electricity use, kWh / day, summer 2002	0.045*** (0.005)	0.042*** (0.008)	. .	. .
trt. customer on high-ratio rate	-0.044 (0.036)	0.058 (0.086)	. .	. .
apartment	-0.044 (0.043)	. .	. .	. .
climate zone 2	0.029 (0.035)	-0.010 (0.151)	. .	. .
climate zone 3	0.015 (0.043)	0.125 (0.194)	. .	. .
climate zone 4	-0.028 (0.158)	-0.354 (0.272)	. .	. .
cooling degree hours 2-7PM, base 78	0.012*** (0.004)	0.029*** (0.006)	0.010* (0.005)	0.034*** (0.006)
heating degree hours 2-7pm	0.007** (0.003)	0.005 (0.007)	0.002 (0.005)	-0.041*** (0.015)
Tuesday	-0.009	-0.002	-0.003	-0.010

	Specification 1: Simplest Diff in Diff		Specification 4: Adds person FE's; controls	
	Low Use / Apt. Customers	High Use Customers	Low Use / Apt. Customers	High Use Customers
	(0.007)	(0.014)	(0.010)	(0.017)
Wednesday	-0.006 (0.006)	-0.006 (0.016)	-0.003 (0.009)	-0.018 (0.020)
Thursday	-0.012 (0.009)	-0.038** (0.016)	-0.002 (0.012)	-0.033* (0.019)
Friday	-0.002 (0.009)	-0.003 (0.021)	0.004 (0.011)	-0.012 (0.026)
year 2004	-0.013 (0.025)	-0.088 (0.057)	0.013 (0.037)	0.002 (0.058)
June	0.038*** (0.013)	0.167*** (0.032)	0.022 (0.018)	0.085** (0.038)
July	0.072*** (0.018)	0.338*** (0.051)	0.065*** (0.024)	0.262*** (0.050)
August	0.090*** (0.020)	0.389*** (0.060)	0.073*** (0.027)	0.314*** (0.063)
September	0.074*** (0.018)	0.172*** (0.043)	0.058** (0.025)	0.126*** (0.045)
October	0.040 (0.025)	0.007 (0.066)	0.039 (0.035)	0.044 (0.073)
heating degree hours 2-7pm squared (1000's)	-0.172*** (0.066)	0.001 (0.149)	. .	. .
Tue * cooling degree hours 2-7pm	-0.00050 (0.00070)	-0.002 (0.001)	-0.001 (0.00085)	-0.002 (0.001)
Tue * cooling degree hours 2-7pm squared (1000's)	-0.003 (0.007)	0.013 (0.012)	0.003 (0.009)	0.011 (0.012)
Tue * heating degree hours 2-7pm	0.00072 (0.001)	-0.006 (0.004)	0.00052 (0.001)	-0.007 (0.006)
Tue * heating degree hours 2-7pm squared (1000's)	0.012	0.110*	0.006	-0.016

	Specification 1: Simplest Diff in Diff		Specification 4: Adds person FE's; controls	
	Low Use / Apt. Customers	High Use Customers	Low Use / Apt. Customers	High Use Customers
	(0.023)	(0.060)	(0.022)	(0.245)
Wed * cooling degree hours 2-7pm	-0.00069 (0.00071)	-0.004*** (0.001)	-0.00097 (0.00095)	-0.003** (0.001)
Wed * cooling degree hours 2-7pm squared (1000's)	0.003 (0.007)	0.045*** (0.015)	0.004 (0.010)	0.029** (0.015)
Wed * heating degree hours 2-7pm	0.00071 (0.002)	0.004 (0.004)	0.00087 (0.002)	0.007 (0.008)
Wed * heating degree hours 2-7pm squared (1000's)	0.006 (0.030)	-0.041 (0.081)	-0.004 (0.027)	-0.348 (0.296)
Thu * cooling degree hours 2-7pm	-0.00071 (0.00079)	-0.002* (0.001)	-0.001 (0.00090)	-0.001 (0.001)
Thu * cooling degree hours 2-7pm squared (1000's)	0.001 (0.008)	0.025* (0.014)	0.005 (0.010)	0.015 (0.012)
Thu * heating degree hours 2-7pm	-0.001 (0.002)	-0.006 (0.004)	-0.001 (0.002)	-0.002 (0.008)
Thu * heating degree hours 2-7pm squared (1000's)	0.047 (0.036)	0.179** (0.081)	0.040 (0.032)	-0.067 (0.275)
Fri * cooling degree hours 2-7pm	-0.002* (0.00085)	-0.004** (0.001)	-0.001 (0.001)	-0.00098 (0.002)
Fri * cooling degree hours 2-7pm squared (1000's)	0.010 (0.008)	0.027* (0.015)	0.007 (0.011)	0.004 (0.017)
Fri * heating degree hours 2-7pm	-0.004** (0.002)	-0.009** (0.004)	-0.004* (0.002)	-0.006 (0.008)
Fri * heating degree hours 2-7pm squared (1000's)	0.085** (0.034)	0.143* (0.083)	0.072** (0.033)	-0.039 (0.255)
year_2004 * cooling degree hours 2-7pm	-0.005*** (0.002)	-0.003 (0.003)	-0.004** (0.002)	-0.004 (0.004)
year_2004 * cooling degree hours 2-7pm squared (1000's)	0.049***	0.028	0.040**	0.032

	Specification 1: Simplest Diff in Diff		Specification 4: Adds person FE's; controls	
	Low Use / Apt. Customers	High Use Customers	Low Use / Apt. Customers	High Use Customers
	(0.015)	(0.036)	(0.016)	(0.039)
year_2004 * heating degree hours 2-7pm	-0.00084 (0.002)	0.012* (0.007)	-0.001 (0.005)	0.024 (0.016)
year_2004 * heating degree hours 2-7pm squared (1000's)	0.063 (0.054)	-0.211 (0.143)	-0.004 (0.143)	-0.950* (0.497)
June * cooling degree hours 2-7pm	-0.00088 (0.002)	-0.004 (0.004)	0.002 (0.002)	0.001 (0.003)
June * cooling degree hours 2-7pm squared (1000's)	0.010 (0.030)	0.041 (0.051)	-0.029 (0.027)	0.017 (0.040)
June * heating degree hours 2-7pm	-0.003 (0.002)	-0.005 (0.005)	-0.002 (0.002)	-0.003 (0.010)
June * heating degree hours 2-7pm squared (1000's)	0.064** (0.030)	0.021 (0.076)	0.049 (0.035)	0.338 (0.449)
July * cooling degree hours 2-7pm	-0.001 (0.002)	-0.00038 (0.004)	-0.001 (0.002)	-0.00043 (0.004)
July * cooling degree hours 2-7pm squared (1000's)	0.021 (0.022)	0.009 (0.043)	0.005 (0.025)	0.032 (0.043)
July * heating degree hours 2-7pm	-0.009*** (0.003)	-0.035*** (0.011)	-0.010*** (0.003)	-0.018 (0.011)
July * heating degree hours 2-7pm squared (1000's)	0.225** (0.110)	1.027*** (0.363)	0.212** (0.105)	0.357 (0.397)
Aug * cooling degree hours 2-7pm	-0.002 (0.002)	-0.002 (0.004)	-0.002 (0.002)	-0.001 (0.004)
Aug * cooling degree hours 2-7pm squared (1000's)	0.027 (0.021)	0.011 (0.039)	0.009 (0.026)	0.020 (0.036)
Aug * heating degree hours 2-7pm	-0.019** (0.008)	-0.066*** (0.023)	-0.011 (0.007)	-0.053** (0.022)
Aug * heating degree hours 2-7pm squared (1000's)	0.525	2.078**	0.230	1.918*

	Specification 1: Simplest Diff in Diff		Specification 4: Adds person FE's; controls	
	Low Use / Apt. Customers	High Use Customers	Low Use / Apt. Customers	High Use Customers
	(0.368)	(0.932)	(0.261)	(0.993)
Sept * cooling degree hours 2-7pm	-0.006*** (0.002)	-0.005 (0.003)	-0.005** (0.002)	-0.005* (0.003)
Sept * cooling degree hours 2-7pm squared (1000's)	0.055*** (0.019)	0.033 (0.034)	0.032 (0.023)	0.046 (0.032)
Sept * heating degree hours 2-7pm	-0.004* (0.002)	0.002 (0.006)	-0.002 (0.002)	0.019 (0.016)
Sept * heating degree hours 2-7pm squared (1000's)	0.065** (0.029)	-0.095 (0.083)	0.037 (0.028)	-0.733 (0.578)
Oct * cooling degree hours 2-7pm	-0.012*** (0.003)	-0.018*** (0.004)	-0.009*** (0.003)	-0.016*** (0.004)
Oct * cooling degree hours 2-7pm squared (1000's)	0.101*** (0.028)	0.123** (0.050)	0.049 (0.031)	0.097** (0.046)
Oct * heating degree hours 2-7pm	-0.002 (0.003)	0.010 (0.007)	-0.002 (0.005)	0.026 (0.018)
Oct * heating degree hours 2-7pm squared (1000's)	0.089* (0.053)	-0.173 (0.142)	0.012 (0.139)	-1.051* (0.569)
cooling degree hours squared (1000's), 2-7pm	-0.034 (0.035)	-0.086 (0.069)	0.001 (0.043)	. .
heating degree hours 2-7pm squared, 1000's	. .	. .	-0.035 (0.104)	0.825 (0.555)
constant	-0.082 (0.061)	-0.213 (0.224)	0.469*** (0.030)	1.219*** (0.061)
N	66832	54576	39483	38177
R-squared	0.3459	0.4228	0.5411	0.5993

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.

	Specification 1: Simplest Diff in Diff		Specification 4: Adds person FE's; controls	
	Low Use / Apt. Customers	High Use Customers	Low Use / Apt. Customers	High Use Customers

Appendix G

Methodology used to Calculate Population Weighted Temperatures

This section describes the methodology used to calculate the population weighted temperatures that I use in Section 2.5.

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.[1] 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:

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

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:

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

[1]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.

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:

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

Appendix H

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

Appendix I

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

Appendix J

Population Weighted Cooling Degree Hours and Cooling Degree Hours Squared

TOU 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[1] 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.

[1] 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

percentiles of peak load distribution	N	Peak	Zone 1		Zone 2		Zone 3		Zone 4		State	
		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

Appendix K

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 single quadratic relationship between temperature and response. There are troubling signs that 1) the rigid nature of 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 K.4 and figures K.4 and K.5 simply substitute the piecewise linear function of temperature for the quadratic function of temperature in the difference-in-difference framework used above. Section K.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 K.4 and K.5 need to display the estimates' standard errors.

K.1 The Kernel Estimates

The kernel estimate in figures K.1, K.2, and K.3 make estimates as follows:

i. 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.

ii. 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[1] and simplify the section 2.3 regression[2] 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.[3] I use no interactions between calendar days and temperature.

iii. 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

[1] 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.

[2] 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 K.1, K.2, and K.3 are unweighted.

[3] The difference-in-difference splines estimate the slope of fewer segments because there are too few very hot days during the pretreatment period.

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

iv. The analysis then uses the control group's coefficients to predict each treatment customer's energy use each day out-of-sample. 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.

K.2 The Difference-in-Difference Spline Estimates

Section 2.4 describes the difference-in-difference spline estimates as an extension of the general econometric model described in Section 2.3.

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 K.4 and K.5 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 2.3, the difference lines report $(\hat{\delta}+\hat{\beta})\mathbf{X}^*$ for TOU peak impacts and $(\hat{\delta}+\hat{\psi})\mathbf{X}^*$ for the impact of the critical price. By contrast, the difference-in-difference estimates are $\hat{\beta}\mathbf{X}^*$ and $\hat{\psi}\mathbf{X}^*$

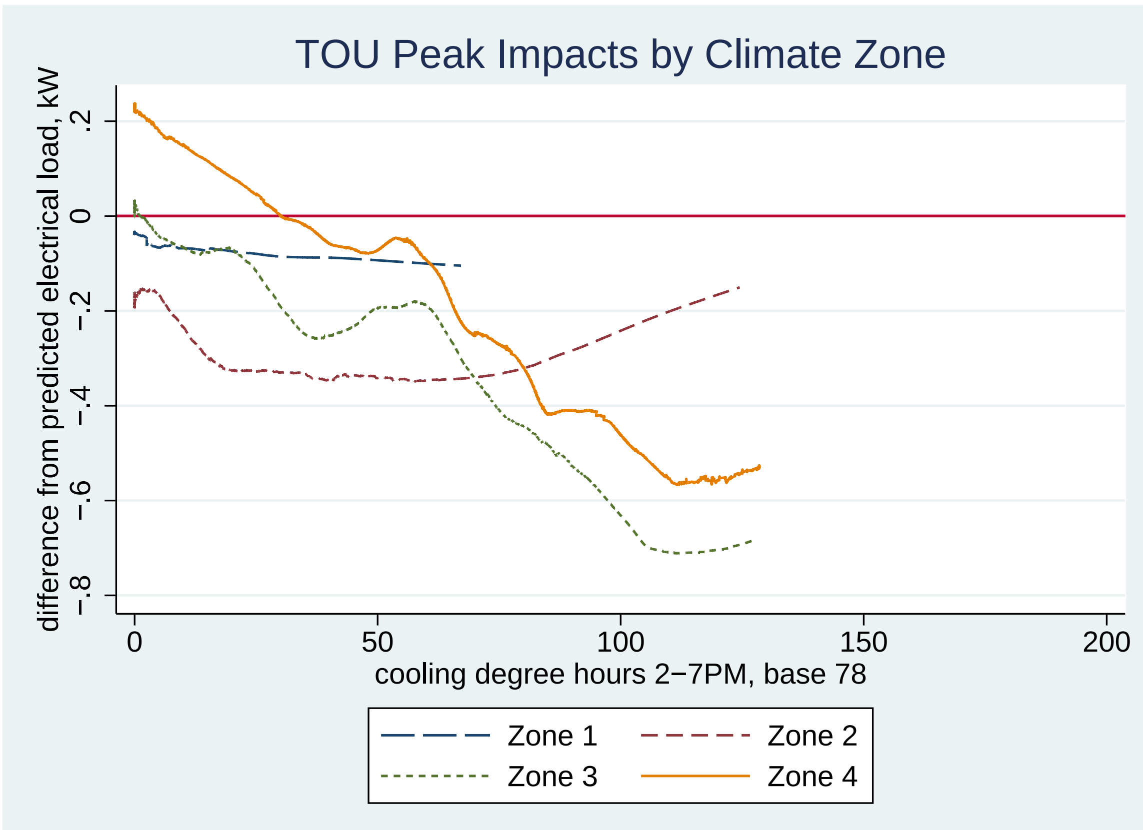

Figure K.1: *Kernel estimates of the relationship between the impact of TOU peak prices and temperature by climate zone. Plotted for temperatures between the 1st and 99th percentile of the temperature range for each climate zone.*

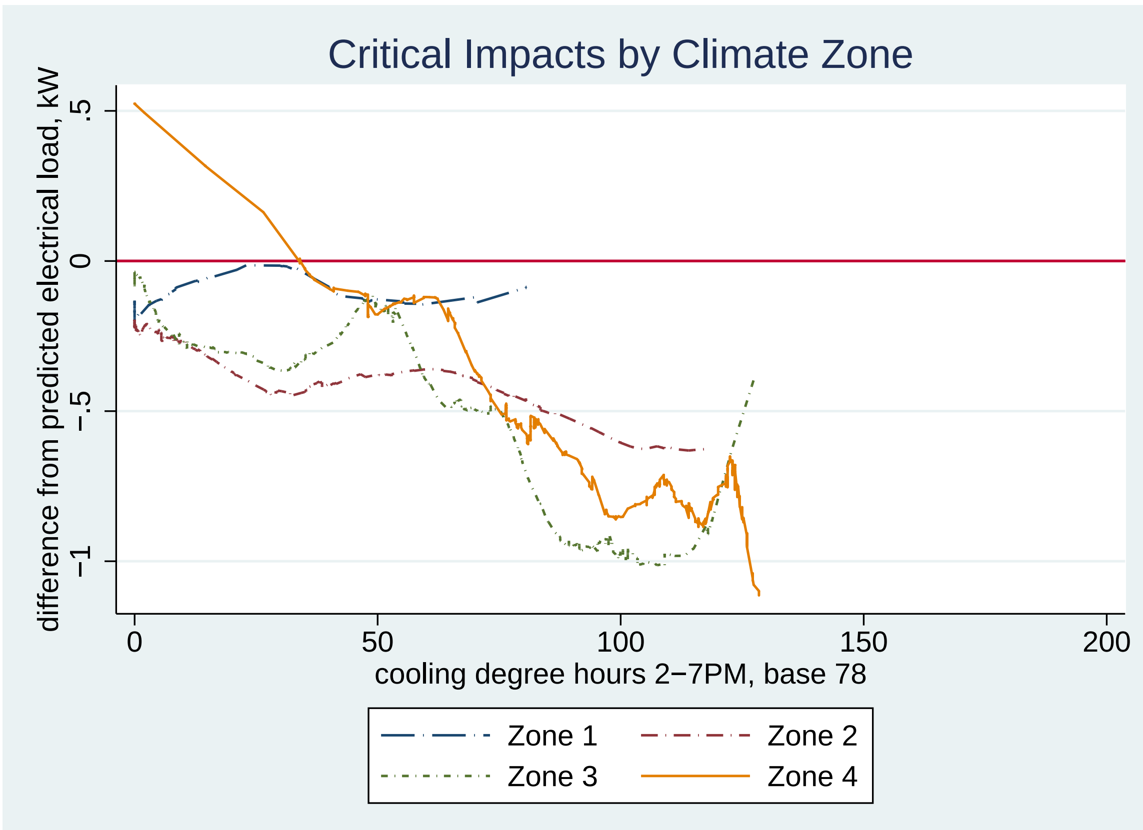

Figure K.2: *Kernel estimates of the relationship between the impact of critical prices and temperature by climate zone. Plotted for temperatures between the 1st and 99th percentile of the temperature range for each climate zone.*

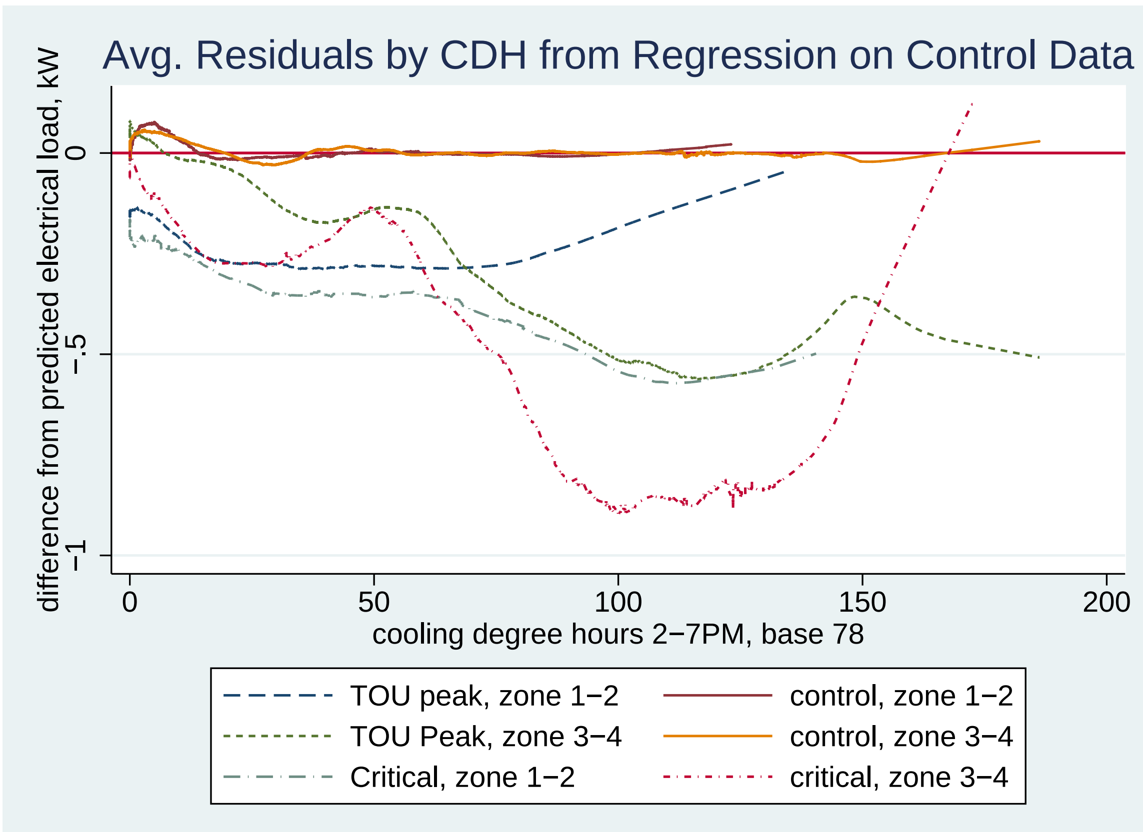

Figure K.3: *Nonparametric estimates of the impacts of dynamic pricing, by climate zone and the price that is in effect. The two control group lines that stay quite close to zero everywhere suggest that the functional form captures the average temperature-driven variation in the control group's electric use. The other four lines approximate the impacts of critical and TOU peak prices by temperature. Plotted for all temperatures. Less than 1% of all zone 2 (4) data is above 95 (140) CDH and climate zone 1 (3) is even cooler, so impacts beyond those temperatures may not be very reliable.*

respectively. The "difference" estimators yield more intuitively appealing results, especially in conditions that were fairly hot for June.

K.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 finding makes sense. The control group's air conditioning load will stop growing in temperature when it gets so hot that air conditioners start running continuously. Meanwhile, dynamic pricing customers' air conditioning loads will increase with temperature at these levels if they responded to the higher price by increasing their thermostat set points a few degrees. Thus the treatment group's use may grow when the control groups' use is not because it is more work to keep a house at 80^oF when the outside temperature is 100^oF rather than 90^oF, but an air conditioner may be running flat out to keep a house as close as possible to 70^oF regardless of whether it is 95^oF or 100^oF.

- Temperature sensitivity appears to be particularly high between roughly 70 and 100 to 110 CDH (i.e. the range between averaging 92 and 100^oF for the afternoon).

- 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. Impacts in zone 1 and, to a lesser extent, zone 2 are fairly temperature insensitive. Thus, zone dummies will capture the difference between zones 1 and 2 reasonably well.

- The difference estimators make very few point estimates suggesting that dynamic pricing was counterproductive and raised energy use.[4] 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

[4]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.

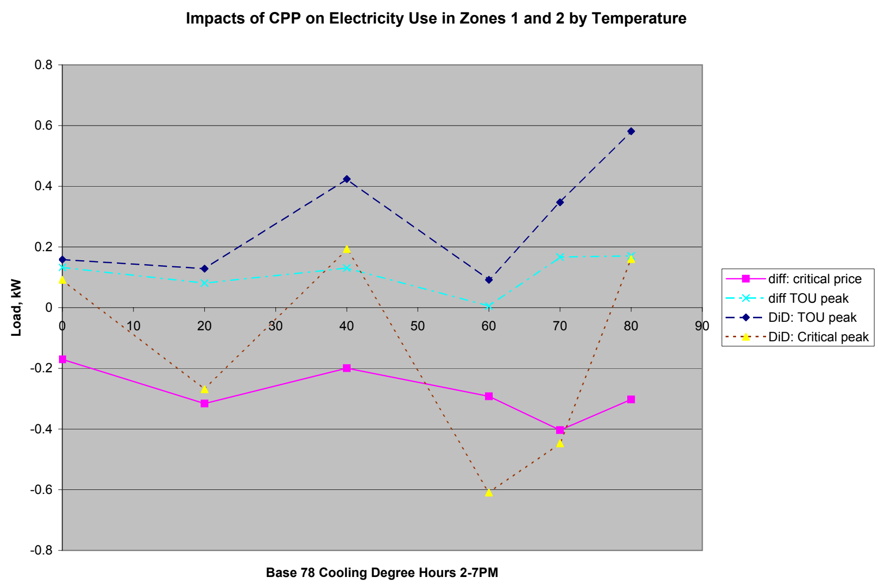

Figure K.4: Splines fit to the temperature-energy use relationship for temperate climate zones 1 and 2. The Difference-in-Difference (DiD) lines measure impact relative to the "preexisting differences" measured from behavior during the quite brief, relatively cool pretreatment period. The difference 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.

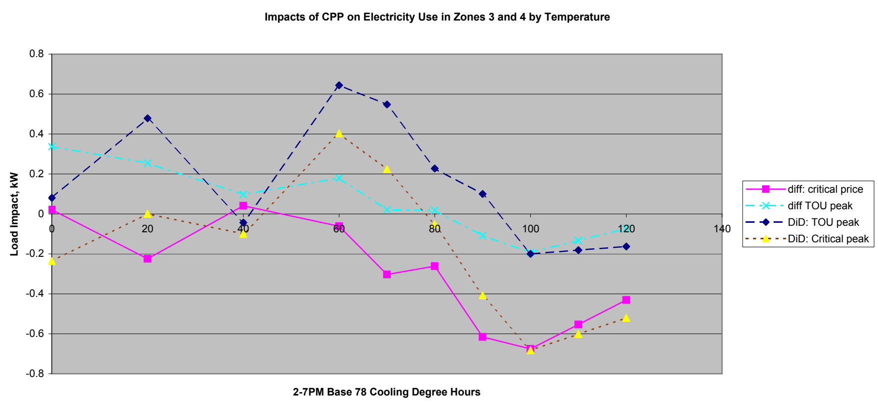

Figure K.5: Splines fit to the temperature-energy use relationship for temperate climate zones 3 and 4. This uses the same techniques and conventions as Graph K.4. The Difference-in-Difference regression finds that the treatment group used far less power than the control group during (rare) very hot 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.

to thin and idiosyncratic high temperature, pretreatment data. When we change them to difference estimators, we notice that the point estimates showing dynamic pricing to be counterproductive exist to almost exactly negate strange, imprecisely estimated pretreatment relationships.

K.4 Regressions with Splines

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 5: Simplest Diff in Diff	Specification 6: Adding Survey Variables	Specification 8: Adds person FE's; controls
TOU Peak Price in Effect	-0.178 (0.177)	-0.330 (0.227)	0.126 (0.256)
TOU Peak Price in Effect * electric use, kWh / day , summer 2002	-0.004 (0.004)	-0.006 (0.005)	0.004 (0.005)
TOU Peak Price in Effect * high ratio rate customer.	-0.013 (0.039)	-0.017 (0.043)	0.017 (0.054)
TOU Peak Price in Effect * apartment	-0.056 (0.061)	0.012 (0.086)	-0.011 (0.100)
TOU Peak Price in Effect * climate zone 1	0.268* (0.161)	0.236 (0.180)	0.121 (0.187)
TOU Peak Price in Effect * climate zone 2	0.219 (0.158)	0.159 (0.178)	0.080 (0.175)
TOU Peak Price in Effect * climate zone 3	0.208 (0.162)	0.170 (0.175)	0.067 (0.169)
TOU Peak Price in Effect * cooling degree hours 2-7pm	0.020*** (0.007)	0.021*** (0.008)	0.024*** (0.009)
TOU Peak Price in Effect * impact of a CDH beyond 20	-0.046*** (0.015)	-0.048*** (0.016)	-0.059*** (0.018)

	Specification 5: Simplest Diff in Diff	Specification 6: Adding Survey Variables	Specification 8: Adds person FE's; controls
TOU Peak Price in Effect * impact of a CDH beyond 40	0.061*** (0.022)	0.058** (0.023)	0.065*** (0.022)
TOU Peak Price in Effect * impact of a CDH beyond 60	-0.044 (0.043)	-0.033 (0.045)	-0.004 (0.044)
TOU Peak Price in Effect * impact of a CDH beyond 70	-0.022 (0.062)	-0.019 (0.063)	-0.079 (0.071)
TOU Peak Price in Effect * impact of a CDH beyond 80	0.019 (0.060)	-0.007 (0.060)	0.035 (0.067)
TOU Peak Price in Effect * impact of a CDH beyond 90	-0.017 (0.054)	0.018 (0.055)	-0.018 (0.057)
TOU Peak Price in Effect * impact of a CDH beyond 100	0.032 (0.044)	0.008 (0.046)	0.042 (0.044)
TOU Peak Price in Effect * CDH * zone is 1 or 2	-0.021** (0.009)	-0.024** (0.010)	-0.027** (0.012)
TOU Peak Price in Effect * zone 1 or 2 * impact of a CDH beyond 20	0.062*** (0.021)	0.066*** (0.022)	0.093*** (0.027)
TOU Peak Price in Effect * zone 1 or 2 * impact of a CDH beyond 40	-0.092*** (0.035)	-0.087** (0.037)	-0.148*** (0.043)
TOU Peak Price in Effect * zone 1 or 2 * impact of a CDH beyond 60	0.086 (0.094)	0.024 (0.101)	0.116 (0.111)
TOU Peak Price in Effect * zone 1 or 2 * impact of a CDH beyond 70	0.020 (0.103)	0.099 (0.107)	0.079 (0.115)
TOU Peak Price in Effect * heating degree hours 2-7pm	-0.002 (0.002)	-0.002 (0.002)	-0.008 (0.006)
TOU peak price in effect * day before critical price	-0.006 (0.012)	-0.006 (0.013)	-0.004 (0.013)
TOU peak price in effect * day after critical price	0.027**	0.034**	0.016

	Specification 5: Simplest Diff in Diff	Specification 6: Adding Survey Variables	Specification 8: Adds person FE's; controls
	(0.013)	(0.014)	(0.014)
TOU Peak Price in Effect * central AC	. .	-0.038 (0.079)	-0.054 (0.085)
TOU Peak Price in Effect * room AC	. .	0.090 (0.082)	-0.091 (0.114)
TOU Peak Price in Effect * number of bedrooms	. .	0.063 (0.040)	0.048 (0.042)
TOU Peak Price in Effect * # people in the household	. .	0.010 (0.022)	0.054 (0.038)
TOU Peak Price in Effect * heating degree hours 2-7pm squared (1000's)	. .	. .	0.148 (0.125)
TOU Peak Price in Effect * cooling degree hours 2-7PM, previous day	. .	. .	-0.00098 (0.00061)
TOU Peak Price in Effect * cooling degree hours 2-7PM, two days before	. .	. .	-0.00012 (0.00043)
TOU Peak Price in Effect * cooling degree hours 2-7PM, three days before	. .	. .	-0.001** (0.00043)
TOU Peak Price in Effect * swimming pool	. .	. .	-0.268* (0.146)
TOU Peak Price in Effect * # kids under 5 in household	. .	. .	-0.118* (0.070)
TOU Peak Price in Effect * # people over 65 in household	. .	. .	-0.114** (0.054)
TOU Peak Price in Effect * electric cooktop	. .	. .	0.176 (0.144)
TOU Peak Price in Effect * customer stayed in expt. throughout expt.	. .	. .	-0.190** (0.087)

	Specification 5: Simplest Diff in Diff	Specification 6: Adding Survey Variables	Specification 8: Adds person FE's; controls
TOU Peak Price in Effect * cooling degree hours 2-7pm * room AC	. .	. .	0.009*** (0.003)
TOU Peak Price in Effect * heating degree hours 2-7PM*electric heat	. .	. .	0.011*** (0.003)
TOU Peak Price in Effect * electric heat	. .	. .	-0.139 (0.093)
TOU Peak Price in Effect * # kids over 5 in household	. .	. .	-0.062 (0.050)
TOU Peak Price in Effect * work from home 0-10 hrs/wk	. .	. .	-0.057 (0.117)
TOU Peak Price in Effect * electric oven	. .	. .	-0.125 (0.139)
TOU Peak Price in Effect * number of refrigerators and freezers	. .	. .	-0.120 (0.083)
TOU Peak Price in Effect * work from home 11-30 hrs/wk	. .	. .	0.063 (0.102)
TOU Peak Price in Effect * work from home >30 hrs/wk	. .	. .	-0.310 (0.255)
TOU Peak Price in Effect * spa	. .	. .	0.070 (0.136)
TOU Peak Price in Effect * customer stayed in expt. < 4.5 months	. .	. .	-0.074 (0.130)
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	0.023	-0.007

	Specification 5: Simplest Diff in Diff	Specification 6: Adding Survey Variables	Specification 8: Adds person FE's; controls
	(0.025)	(0.028)	(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)
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 2-7pm	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)

	Specification 5: Simplest Diff in Diff	Specification 6: Adding Survey Variables	Specification 8: Adds person FE's; controls
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 2-7pm	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** (0.119)	-0.145 (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 2-7pm squared (1000's)	. .	. .	1.132 (0.692)
Critical Price in Effect * cooling degree hours 2-7PM, previous day	. .	. .	-0.002 (0.001)
Critical Price in Effect * cooling degree hours 2-7PM, two days before	.	.	0.001

	Specification 5: Simplest Diff in Diff	Specification 6: Adding Survey Variables	Specification 8: Adds person FE's; controls
	.	.	(0.001)
Critical Price in Effect * cooling degree hours 2-7PM, 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 2-7pm * room AC	. .	. .	0.009*** (0.003)
Critical Price in Effect * heating degree hours 2-7PM*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 FE's; 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)
Treatment Customer	0.309 (0.193)	0.588** (0.247)	. .
Treatment Customer * electric use, kWh / day , summer 2002	0.00021 (0.004)	0.002 (0.005)	. .
Treatment Customer * apartment	0.106* (0.058)	-0.015 (0.090)	. .
Treatment Customer * climate zone 1	-0.382** (0.180)	-0.404** (0.196)	. .
Treatment Customer * climate zone 2	-0.365** (0.175)	-0.429** (0.190)	. .
Treatment Customer * climate zone 3	-0.295* (0.176)	-0.344* (0.184)	. .
Treatment Customer * cooling degree hours 2-7pm	-0.024*** (0.007)	-0.025*** (0.007)	-0.026*** (0.009)
Treatment Customer * impact of a CDH beyond 20	0.042*** (0.014)	0.043*** (0.015)	0.052*** (0.018)
Treatment Customer * impact of a CDH beyond 40	-0.049** (0.021)	-0.044** (0.022)	-0.054** (0.022)
Treatment Customer * impact of a CDH beyond 60	0.024	0.012	-0.008

	Specification 5: Simplest Diff in Diff	Specification 6: Adding Survey Variables	Specification 8: Adds person FE's; controls
	(0.041)	(0.043)	(0.044)
Treatment Customer * impact of a CDH beyond 70	0.038 (0.059)	0.034 (0.060)	0.089 (0.069)
Treatment Customer * impact of a CDH beyond 80	-0.032 (0.054)	-0.008 (0.054)	-0.053 (0.064)
Treatment Customer * impact of a CDH beyond 90	0.021 (0.048)	-0.013 (0.048)	0.033 (0.054)
Treatment Customer * impact of a CDH beyond 100	-0.017 (0.037)	0.006 (0.038)	-0.035 (0.042)
Treatment Customer * Cooling Degree Hours * Zone 1 or 2	0.023** (0.009)	0.025** (0.010)	. .
Treatment Customer * zone 1 or 2 * impact of a CDH beyond 20	-0.054*** (0.020)	-0.056*** (0.022)	-0.078*** (0.027)
Treatment Customer * zone 1 or 2 * impact of a CDH beyond 40	0.071** (0.033)	0.064* (0.036)	0.118*** (0.041)
Treatment Customer * zone 1 or 2 * impact of a CDH beyond 60	-0.044 (0.093)	0.011 (0.101)	-0.079 (0.112)
Treatment Customer * zone 1 or 2 * impact of a CDH beyond 70	-0.052 (0.100)	-0.115 (0.106)	-0.093 (0.112)
Treatment Customer * heating degree hours 2-7pm	0.00020 (0.002)	0.00054 (0.002)	0.004 (0.005)
Treatment Customer * central AC	. .	-0.032 (0.073)	. .
Treatment Customer * room AC	. .	0.095 (0.084)	. .
Treatment Customer * number of bedrooms	. .	-0.100** (0.043)	. .

	Specification 5: Simplest Diff in Diff	Specification 6: Adding Survey Variables	Specification 8: Adds person FE's; controls
Treatment Customer * # people in the household	. .	0.017 (0.020)	. .
Treatment Customer * heating degree hours 2-7pm squared (1000's)	. .	. .	-0.175 (0.123)
Treatment Period (after 7/1/2003)	0.075 (0.127)	0.054 (0.162)	-0.218 (0.193)
Treatment Period * electric use, kWh / day , summer 2002	0.007** (0.003)	0.006 (0.004)	0.00099 (0.004)
Treatment Period * apartment	0.039 (0.041)	0.002 (0.060)	0.016 (0.064)
Treatment Period * climate zone 1	-0.098 (0.122)	0.015 (0.135)	0.005 (0.154)
Treatment Period * climate zone 2	-0.079 (0.121)	0.021 (0.137)	0.022 (0.144)
Treatment Period * climate zone 3	-0.085 (0.125)	-0.034 (0.135)	-0.035 (0.140)
Treatment Period * cooling degree hours 2-7pm	-0.008 (0.006)	-0.009 (0.006)	-0.025*** (0.007)
Treatment Period * impact of a CDH beyond 20	0.015 (0.012)	0.017 (0.012)	0.048*** (0.014)
Treatment Period * impact of a CDH beyond 40	-0.028* (0.016)	-0.027 (0.017)	-0.054*** (0.018)
Treatment Period * impact of a CDH beyond 60	0.033 (0.031)	0.036 (0.031)	0.031 (0.033)
Treatment Period * impact of a CDH beyond 70	-0.013 (0.047)	-0.040 (0.048)	-0.00092 (0.057)
Treatment Period * impact of a CDH beyond 80	0.035	0.078*	0.044

	Specification 5: Simplest Diff in Diff	Specification 6: Adding Survey Variables	Specification 8: Adds person FE's; controls
	(0.046)	(0.046)	(0.053)
Treatment Period * impact of a CDH beyond 90	-0.010 (0.037)	-0.057 (0.037)	-0.037 (0.038)
Treatment Period * impact of a CDH beyond 100	-0.029 (0.026)	0.003 (0.028)	-0.011 (0.022)
Treatment Period * CDH * zone is 1 or 2	0.011 (0.008)	0.011 (0.008)	0.025** (0.010)
Treatment Period * zone 1 or 2 * impact of a CDH beyond 20	-0.028* (0.016)	-0.030* (0.017)	-0.065*** (0.022)
Treatment Period * zone 1 or 2 * impact of a CDH beyond 40	0.060** (0.026)	0.062** (0.027)	0.114*** (0.034)
Treatment Period * zone 1 or 2 * impact of a CDH beyond 60	-0.129* (0.068)	-0.127* (0.066)	-0.184** (0.084)
Treatment Period * zone 1 or 2 * impact of a CDH beyond 70	0.066 (0.076)	0.071 (0.072)	0.060 (0.090)
Treatment Period * heating degree hours 2-7pm	0.002* (0.001)	0.002 (0.002)	0.016*** (0.004)
Treatment Period * central AC	. .	0.149** (0.061)	0.169** (0.067)
Treatment Period * room AC	. .	-0.032 (0.058)	0.028 (0.098)
Treatment Period * number of bedrooms	. .	-0.032 (0.030)	-0.039 (0.031)
Treatment Period * # people in the household	. .	0.00092 (0.017)	-0.023 (0.027)
Treatment Period * heating degree hours 2-7pm squared (1000's)	. .	. .	-0.262** (0.105)

	Specification 5: Simplest Diff in Diff	Specification 6: Adding Survey Variables	Specification 8: Adds person FE's; controls
Treatment Period * cooling degree hours 2-7PM, previous day	. .	. .	0.003*** (0.00047)
Treatment Period * cooling degree hours 2-7PM, two days before	. .	. .	0.001*** (0.00037)
Treatment Period * cooling degree hours 2-7PM, three days before	. .	. .	0.00089*** (0.00034)
Treatment Period * swimming pool	. .	. .	0.229** (0.116)
Treatment Period * # kids under 5 in household	. .	. .	0.110** (0.053)
Treatment Period * # people over 65 in household	. .	. .	0.151*** (0.044)
Treatment Period * electric cooktop	. .	. .	-0.125 (0.117)
Treatment Period * customer stayed in expt. throughout expt.	. .	. .	0.177** (0.070)
Treatment Period * cooling degree hours 2-7pm * room AC	. .	. .	-0.007** (0.003)
Treatment Period * heating degree hours 2-7PM*electric heat	. .	. .	-0.007*** (0.002)
Treatment Period * electric heat	. .	. .	0.140** (0.069)
Treatment Period * # kids over 5 in household	. .	. .	0.069** (0.031)
Treatment Period * work from home 0-10 hrs/wk	. .	. .	0.198** (0.082)
Treatment Period * electric oven	.	.	-0.018

	Specification 5: Simplest Diff in Diff	Specification 6: Adding Survey Variables	Specification 8: Adds person FE's; controls
	.	.	(0.108)
Treatment Period * number of refrigerators and freezers	. .	. .	0.014 (0.061)
Treatment Period * work from home 11-30 hrs/wk	. .	. .	-0.017 (0.077)
Treatment Period * work from home >30 hrs/wk	. .	. .	0.078 (0.203)
Treatment Period * spa	. .	. .	-0.073 (0.105)
Treatment Period * customer stayed in expt. < 4.5 months	. .	. .	0.061 (0.108)
Critical Period	-0.174 (0.123)	-0.385*** (0.141)	-0.316* (0.170)
Critical Period * electric use, kWh / day , summer 2002	0.017*** (0.002)	0.015*** (0.003)	0.017*** (0.003)
Critical Period * high ratio rate customer.	-0.214* (0.118)	-0.224* (0.131)	-0.044 (0.074)
Critical Period * apartment	-0.008 (0.043)	0.047 (0.065)	0.029 (0.083)
Critical Period * climate zone 1	-0.049 (0.114)	0.086 (0.115)	0.135 (0.123)
Critical Period * climate zone 2	-0.027 (0.114)	0.059 (0.117)	0.164 (0.124)
Critical Period * climate zone 3	0.051 (0.081)	0.087 (0.082)	0.091 (0.098)
Critical Period * cooling degree hours 2-7pm	-0.000025 (0.005)	-0.00020 (0.005)	-0.001 (0.006)

	Specification 5: Simplest Diff in Diff	Specification 6: Adding Survey Variables	Specification 8: Adds person FE's; controls
Critical Period * impact of a CDH beyond 20	-0.014 (0.011)	-0.013 (0.012)	-0.009 (0.015)
Critical Period * impact of a CDH beyond 40	0.023* (0.013)	0.022 (0.014)	0.014 (0.017)
Critical Period * impact of a CDH beyond 60	-0.005 (0.016)	-0.006 (0.018)	0.017 (0.020)
Critical Period * impact of a CDH beyond 70	-0.013 (0.025)	-0.009 (0.026)	-0.033 (0.026)
Critical Period * impact of a CDH beyond 80	0.00062 (0.026)	-0.009 (0.026)	-0.00086 (0.027)
Critical Period * impact of a CDH beyond 90	0.016 (0.027)	0.015 (0.028)	0.019 (0.029)
Critical Period * impact of a CDH beyond 100	-0.011 (0.018)	0.00029 (0.019)	-0.006 (0.019)
Critical Period * CDH * zone is 1 or 2	0.003 (0.006)	0.007 (0.006)	0.003 (0.007)
Critical Period * zone 1 or 2 * impact of a CDH beyond 20	0.008 (0.012)	-0.003 (0.015)	-0.001 (0.018)
Critical Period * zone 1 or 2 * impact of a CDH beyond 40	-0.013 (0.016)	0.00010 (0.018)	0.002 (0.022)
Critical Period * zone 1 or 2 * impact of a CDH beyond 60	-0.005 (0.024)	-0.021 (0.027)	-0.042 (0.035)
Critical Period * zone 1 or 2 * impact of a CDH beyond 70	0.005 (0.020)	0.020 (0.022)	0.042 (0.029)
Critical Period * heating degree hours 2-7pm	-0.008 (0.005)	-0.010 (0.008)	0.013 (0.013)
Critical Period * central AC	.	0.261***	0.277***

	Specification 5: Simplest Diff in Diff	Specification 6: Adding Survey Variables	Specification 8: Adds person FE's; controls
	.	(0.053)	(0.062)
Critical Period * room AC	. .	-0.055 (0.061)	0.107 (0.095)
Critical Period * number of bedrooms	. .	0.028 (0.024)	0.006 (0.031)
Critical Period * # people in the household	. .	-0.012 (0.012)	-0.005 (0.018)
Critical Period * heating degree hours 2-7pm squared (1000's)	. .	. .	-0.833* (0.504)
Critical Period * cooling degree hours 2-7PM, previous day	. .	. .	-0.00067 (0.00099)
Critical Period * cooling degree hours 2-7PM, two days before	. .	. .	-0.002** (0.001)
Critical Period * cooling degree hours 2-7PM, three days before	. .	. .	0.002** (0.001)
Critical Period * swimming pool	. .	. .	-0.026 (0.091)
Critical Period * # kids under 5 in household	. .	. .	0.027 (0.048)
Critical Period * # people over 65 in household	. .	. .	0.115** (0.053)
Critical Period * electric cooktop	. .	. .	-0.134* (0.079)
Critical Period * customer stayed in expt. throughout expt.	. .	. .	0.149** (0.068)
Critical Period * cooling degree hours 2-7pm * room AC	. .	. .	-0.00060 (0.002)

	Specification 5: Simplest Diff in Diff	Specification 6: Adding Survey Variables	Specification 8: Adds person FE's; controls
Critical Period * heating degree hours 2-7PM*electric heat	. .	. .	0.004 (0.008)
Critical Period * electric heat	. .	. .	-0.056 (0.071)
Critical Period * # kids over 5 in household	. .	. .	-0.020 (0.027)
Critical Period * work from home 0-10 hrs/wk	. .	. .	-0.003 (0.068)
Critical Period * electric oven	. .	. .	0.093 (0.069)
Critical Period * number of refrigerators and freezers	. .	. .	0.021 (0.045)
Critical Period * work from home 11-30 hrs/wk	. .	. .	0.106 (0.103)
Critical Period * work from home >30 hrs/wk	. .	. .	0.044 (0.143)
Critical Period * spa	. .	. .	-0.055 (0.074)
Critical Period * customer stayed in expt. < 4.5 months	. .	. .	-0.295*** (0.111)
electric use, kWh / day, summer 2002	0.047*** (0.003)	0.045*** (0.004)	. .
trt. customer on high-ratio rate	-0.020 (0.035)	0.009 (0.038)	. .
apartment	-0.059 (0.044)	0.051 (0.078)	. .
climate zone 1	0.403***	0.444***	.

	Specification 5: Simplest Diff in Diff	Specification 6: Adding Survey Variables	Specification 8: Adds person FE's; controls
	(0.147)	(0.155)	.
climate zone 2	0.425*** (0.142)	0.463*** (0.152)	. .
climate zone 3	0.362** (0.144)	0.382*** (0.149)	. .
cooling degree hours 2-7PM, base 78	0.024*** (0.006)	0.025*** (0.006)	0.043*** (0.007)
impact of a CDH beyond 20	-0.021* (0.011)	-0.022* (0.012)	-0.056*** (0.014)
impact of a CDH beyond 40	0.028* (0.015)	0.026 (0.016)	0.053*** (0.017)
impact of a CDH beyond 60	-0.019 (0.029)	-0.022 (0.030)	-0.025 (0.032)
impact of a CDH beyond 70	-0.00051 (0.045)	0.026 (0.045)	-0.004 (0.056)
impact of a CDH beyond 80	-0.012 (0.041)	-0.051 (0.041)	-0.021 (0.050)
impact of a CDH beyond 90	-0.010 (0.032)	0.033 (0.031)	0.011 (0.034)
impact of a CDH beyond 100	0.025 (0.021)	-0.003 (0.021)	0.012 (0.019)
CDH * zone is 1 or 2	-0.020*** (0.008)	-0.019** (0.008)	-0.034*** (0.010)
zone 1-2: impact of a CDH beyond 20	0.032** (0.016)	0.033** (0.017)	0.066*** (0.021)
zone 1-2: impact of a CDH beyond 40	-0.065*** (0.025)	-0.065** (0.026)	-0.107*** (0.032)

	Specification 5: Simplest Diff in Diff	Specification 6: Adding Survey Variables	Specification 8: Adds person FE's; controls
zone 1-2: impact of a CDH beyond 60	0.121* (0.070)	0.121* (0.068)	0.181** (0.088)
zone 1-2: impact of a CDH beyond 70	-0.054 (0.075)	-0.065 (0.072)	-0.061 (0.090)
heating degree hours 2-7pm	0.00033 (0.001)	0.00022 (0.002)	-0.009** (0.004)
Tuesday	-0.019*** (0.005)	-0.018*** (0.006)	-0.020*** (0.007)
Wednesday	-0.011* (0.006)	-0.008 (0.007)	-0.014 (0.008)
Thursday	-0.023*** (0.007)	-0.020*** (0.008)	-0.015* (0.009)
Friday	-0.019*** (0.007)	-0.015* (0.008)	-0.007 (0.009)
year 2004	-0.064*** (0.024)	-0.068*** (0.026)	-0.026 (0.031)
June	0.058*** (0.012)	0.056*** (0.014)	0.049*** (0.016)
July	0.132*** (0.019)	0.136*** (0.020)	0.111*** (0.022)
August	0.135*** (0.019)	0.143*** (0.021)	0.117*** (0.024)
September	0.042** (0.017)	0.045** (0.018)	0.030 (0.021)
October	-0.092*** (0.026)	-0.094*** (0.028)	-0.079** (0.033)
central AC	.	0.068	.

	Specification 5: Simplest Diff in Diff	Specification 6: Adding Survey Variables	Specification 8: Adds person FE's; controls
	.	(0.057)	.
room AC	. .	-0.039 (0.067)	. .
number of bedrooms	. .	0.071** (0.034)	. .
# people in the household	. .	0.010 (0.016)	. .
heating degree hours 2-7pm squared, 1000's	. .	. .	0.256** (0.101)
constant	-0.505*** (0.159)	-0.814*** (0.202)	0.606*** (0.030)

Robust standard errors, clustered by customer in parentheses.

Significance: *=10% ** =5% ***=1%

Abbreviations: AC: air conditioning CAC: central air conditioning FE's: fixed effects

Cooling degree hours (CDH) are base 78^o F. Heating degree hours are base 65^o F.

Splines for impact of a CDH beyond K are defined as: $SplineCDH_k = \max\{0, CDH - K\}$

Appendix L

Optimization Algorithm

The optimization in section can be formulated as a mixed integer linear programming problem. Operations researchers often take a first pass at modeling problems like this one in a mixed integer framework, only to discover that the formulation is intractable because it requires searching an enormous state space. The analysis here collided with exactly that problem. Analysts often respond by changing their modeling methodology or by finding a more tractable, equivalent set of constraints and variables.

This paper changed methodologies to use a specialized algorithm written in Scientific Python that exploits more of problem's structure solves it quite effectively. A naive approach to this problem would attack the challenge of finding the arrangement that maximizes the number of customers get consistent offers by visiting up to 2^N possibilities (where $N \approx 500$ customers) and would repeat this process for each possible grouping of the 16 groups into the smaller number of categories. This algorithm both finds a simpler way to compute the set of customers who would be getting consistent offers and finds a way to do most of that work just once, before beginning to search the few million ways to combine groups into categories. The approach begins by realizing that there is a well defined range of consistent offers for each customer, which Figure L.1 shows.

Graphically, we can use figure L.2 and see the central intuition behind the algorithm algorithm as:

- Try a combination of groups into categories. For this example, ,merge three groups into two categories. One categorization might combine the "x" group and the "Δ" group, leaving the "o" group separate. Another might combine the "x" group and the

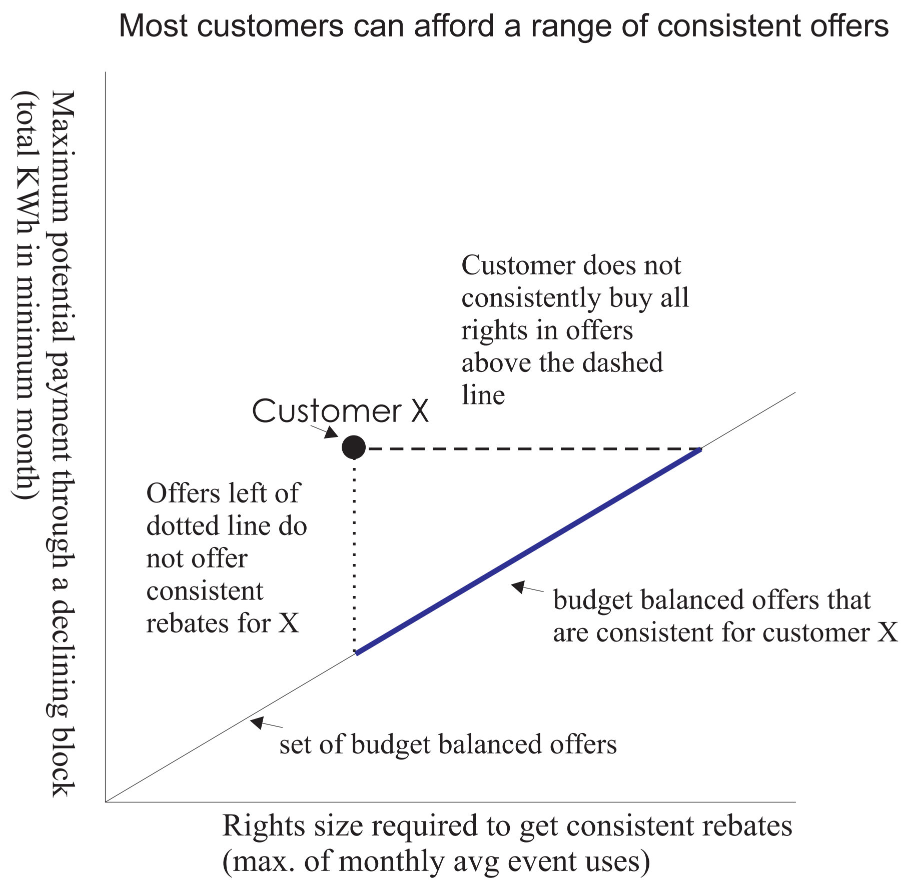

Figure L.1: Visualizing the range of consistent offers for each customer.

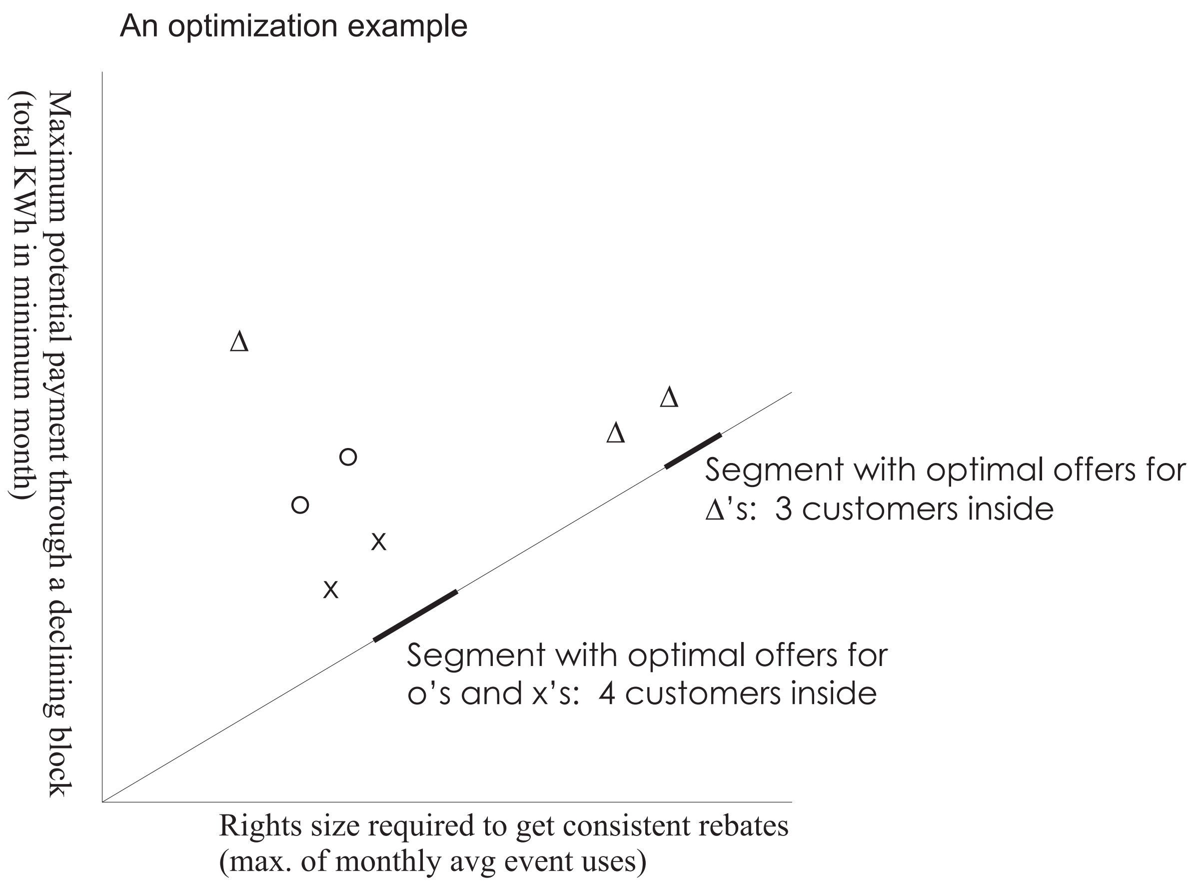

Figure L.2: A simple example of the challenge of grouping 3 groups into two categories.

"o" group, leaving the "Δ" group separate.

- Walk along the budget balanced offer line once for each category.
- Sum the weighted value of the customers in the current category who are (weakly) above and (weakly) to the left of each point on the budget-balanced offer line and would thus get consistent offers if we made that offer to that category.
- Find the maximum point on the budget-balanced offer line by comparing the value at the current point to the best performance achieved at any of the locations previously visited.
- Add up the performance of the optimal offer to each category, getting the optimal performance of the current categorization.
- Repeat this process for all the other possible categorizations, remembering the best categorization seen so far and its performance.

In this simple example, the optimal solution is to combine the "x" group and the "o" group into one category, leaving the group of Δ's as a separate category. This solution makes consistent offers to all seven customers. No other amalgamation of these three groups into two or one categories can make consistent offers to more than 5 customers.

The algorithm that the paper uses a few additional insights about the nature of the optimization problem to reduce the amount of computation required to find the solution. It works as follows.

- Drop any customers for whom no consistent offers exist from the calculations.
- Identify all of the threshold points where a customer in the universe goes from having a consistent offer to an inconsistent offer. Black squares and gray diamonds denote these threshold points on the diagonal line of budget balanced offers in figure L.3. These threshold points divide the continuum of budget balanced offers into a finite set of line segments.[1]

[1] I assign threshold points to .25 kWh bins to reduce the number of line segments, which reduces the size of the vectors listing the value of the objective function for each line segment, simplifying step iii below. There are scenarios in which this approximation could cause very large changes in the location of the optimal value and the percentage of customers reported as getting optimal offers. These scenarios require that many customers get their minimal consistent offer very close to many other customers' maximum abilities to pay

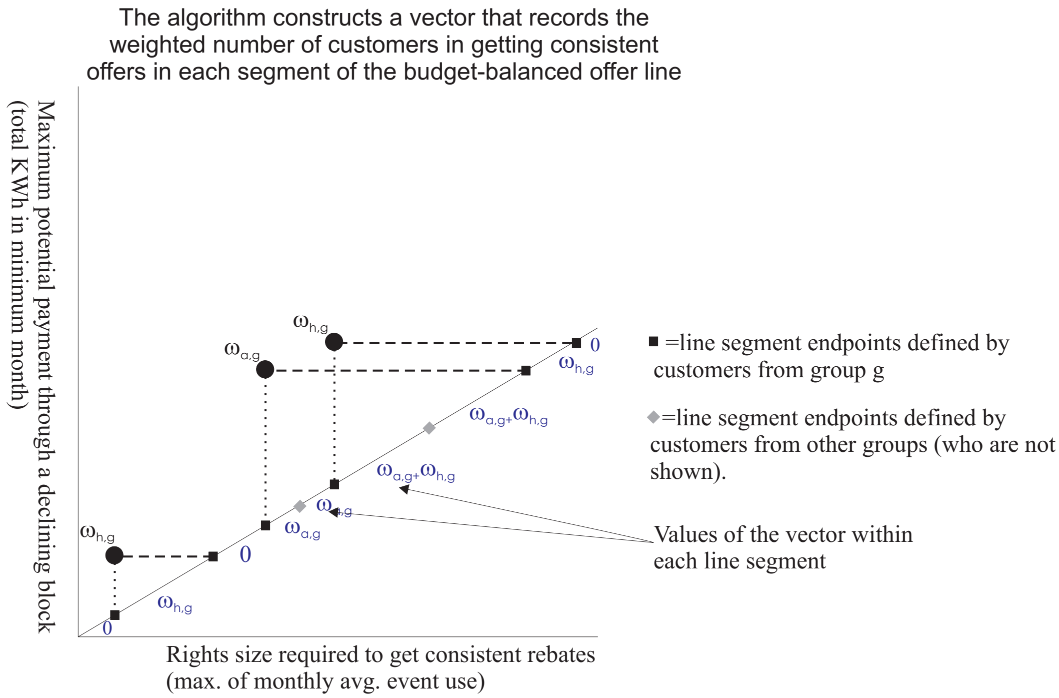

Figure L.3: Calculating the vector listing the number of customers who get consistent offers for each segment of the budget-balanced offer line.

- For each group, create a vector listing the weighted number of customers who would get consistent offers for each budget-balanced-offer line segment.[2] The algorithm creates this vector by creating a list with an entry recording the location of each threshold point and the change in the objective function there (i.e. the weighted total number of customers who start getting consistent offers at the threshold minus the customers who stop getting consistent offers at the threshold). It then sorts the list by threshold value and creates the vector as the running sum of the changes in the number of people getting consistent offers. The objective function is the sum of these vectors.

- Try every possible categorization of groups into categories.[3] The algorithm iterates

for their offer. These scenarios can only happen when the global maximum is not robust to real world noise in customers' energy use. The graphs of the objective function shapes in Chapter 3suggest that this is not an important issue in this data set. It is straightforward to add code to the algorithm to flag any such situations.

[2]The construction of the sixteen categories by preserving climate zones, but collapsing the apartments and single family homes into unified high and low use customer bins means that customers within each group g will have either an apartment (a) or a single family home (h) weight.

[3]We reduce the size of the problem by assuming, without loss of generality, that category 1 always

over each possible assignment of the 16 groups to K categories and finds the optimal arrangement as follows:

i. Increment the assignments of groups to categories. Typically, this will reassign one group to a new category.[4]

ii. Use a branch and bound approach to determine whether the new configuration might have outperformed the optimal configuration found so far, making it worthwhile to run the computationally-intensive, optimum-calculation step iii. The algorithm remembers the best performance achieved so far by any grouping and the configuration's performance the last time the configuration was fully evaluated. It adds the the greatest possible benefits from all of the reconfigurations that have taken place after the most recent computation to the performance of the configuration most recently evaluated. Specifically, the upper bound on the performance of the reconfigured categorization is that each reassigned group vector will contribute its maximum benefits to its new category without taking any benefits away from the category it left. Continue to reassign vectors (step i) and update the best possible performance (step ii), only searching for the exact optimal values (step iii if the upper bound on performance reports that the new configuration could outperform the best performance seen so far.

iii. If this configuration might be a new optimum, calculate the maximum possible performance. First, create category-level benefit vectors by adding each category's constituent group-vectors. Then search every entry in each category-level vector to identify its maximal element. For example, if we have vectors $A = [0, 1], B = [1, 2]$, and $C = [1, 0]$ and groups $i = \{A, B\}$ and $j = \{C\}$, then the category vectors are $i = [1, 3]$ which yields optimal performance of 3 given the offer corresponding to its second segment and $j = [1, 0]$ which yields optimal performance of 1 using its first offer.

---

contains group 1, that group 2 is either in category 1 or 2, and that group g is assigned to category c only if $g \leq c$. We can also require that group $g > 1$ be assigned to category c only if some other group has been assigned to either c or $c - 1$. This assumption eliminates the need to spend computer time evaluating redundant ways to write out the same categorization. A simple example illustrates this: consider trying to assign groups A, B, and C to two nonempty categories. There are three unique categorizations: $(A, B), (C); (A), (B, C); (A, C), (B)$. Unless we assume that group A is in category 1, we can write each unique categorization two ways by shuffling the numbers assigned to the underlying categories: numbering scheme i (ii) has category 1 (2) $=(A, B)$ and category 2 (1) $= (C)$.

[4] This operation is directly akin to counting, and we sometimes reach cases that are analogous to that of adding 001 to 099, which require "carrying" and requires changes to three categorizations.

Appendix M

Performance of More Offers Calculated Relative to the Total Number of Feasible Customers in the Group

Optimal 1 category offer				
size class	**zone 1**	**zone 2**	**zone 3**	**zone 4**
very low	**0.0%**	**2.0%**	**16.2%**	**9.5%**
low	**13.9%**	**54.7%**	**71.9%**	**68.1%**
high	**84.1%**	**95.1%**	**67.5%**	**85.6%**
very high	**100.0%**	**84.5%**	**76.9%**	**55.7%**

Optimal 2 category offer				
size class	**zone 1**	**zone 2**	**zone 3**	**zone 4**
very low	**27.2%**	**55.2%**	**57.0%**	**58.7%**
low	**100.0%**	**98.0%**	**63.6%**	**57.6%**
high	**100.0%**	**85.4%**	**76.8%**	**86.0%**
very high	**91.3%**	**97.4%**	**89.7%**	**71.2%**

Optimal 4 category offer				
size class	**zone 1**	**zone 2**	**zone 3**	**zone 4**
very low	**71.8%**	**74.5%**	**80.9%**	**86.5%**
low	**92.6%**	**88.6%**	**76.5%**	**76.5%**
high	**100.0%**	**90.2%**	**74.5%**	**96.4%**
very high	**100.0%**	**97.4%**	**94.9%**	**79.8%**

Optimal 5 category offer				
size class	**zone 1**	**zone 2**	**zone 3**	**zone 4**
very low	**71.8%**	**78.5%**	**81.4%**	**86.5%**
low	**100.0%**	**100.0%**	**76.5%**	**76.5%**
high	**100.0%**	**90.2%**	**74.5%**	**96.4%**
very high	**100.0%**	**97.4%**	**94.9%**	**79.8%**

One offer for climate zones 1-2 and one for zones 3-4				
size class	**zone 1**	**zone 2**	**zone 3**	**zone 4**
very low	**7.7%**	**20.2%**	**14.2%**	**8.7%**
low	**79.6%**	**92.5%**	**67.8%**	**68.1%**
high	**100.0%**	**87.8%**	**69.8%**	**89.2%**
very high	**91.3%**	**76.8%**	**79.5%**	**59.7%**

One offer per climate zone				
size class	**zone 1**	**zone 2**	**zone 3**	**zone 4**
very low	**21.4%**	**20.2%**	**16.2%**	**3.9%**
low	**100.0%**	**92.5%**	**71.9%**	**68.1%**
high	**100.0%**	**87.8%**	**67.5%**	**92.8%**
very high	**91.3%**	**76.8%**	**76.9%**	**67.8%**

Appendix N

Percentage of Customers Statewide in each cell

<table>
<tr><td></td><td></td><td colspan="2">Statewide Percent of Population in Cell</td></tr>
<tr><td>Climate zone</td><td>Cells</td><td>Single Family Homes</td><td>Apartments</td></tr>
<tr><td>1</td><td>high / very high</td><td>2.3%</td><td rowspan="2">4.9%</td></tr>
<tr><td>1</td><td>low / very low</td><td>5.2%</td></tr>
<tr><td>2</td><td>high / very high</td><td>9.8%</td><td rowspan="2">15.2%</td></tr>
<tr><td>2</td><td>low / very low</td><td>22.3%</td></tr>
<tr><td>3</td><td>high / very high</td><td>8.2%</td><td rowspan="2">6.4%</td></tr>
<tr><td>3</td><td>low / very low</td><td>15.1%</td></tr>
<tr><td>4</td><td>high / very high</td><td>3.1%</td><td rowspan="2">2.1%</td></tr>
<tr><td>4</td><td>low / very low</td><td>5.2%</td></tr>
</table>

Table N.1: The percentage of customers statewide that each cell represents

Appendix O

Notation

- Characteristics that vary by customer – like quantity consumed, $\mathsf{Q_i}$ – appear in sans serif.
- Rate characteritics like P_L and miscellaneous entries appear in the *math* typeface. Rate characteristics reflect local system costs and this document generally takes them as given in designing an IP rebate system.
- Variables that IP rebate designers choose – most notably $\mathbf{Q_D}$, $\mathbf{R}$, and $\mathbf{q_R}$– appear in **bold**.

object	source	notation
basic objects		
Price	exogenous	P
quantity per month	varies	Q_i
quantity per critical or baseline-setting peak period	varies	q_i
time periods – subscripts and sets		
months	exogenous	m
set of months	exogenous	M
critical peak events – e.g. use during an event	exogenous	c, set is C
baseline setting pd	exogenous	b, B
rate of interest	exogenous	r
	CPP	
counts		
number of critical events, per year	exogenous rate design	N_c
number of critical events, this month	exogenous rate design	N_m
number of days in baseline setting period	exogenous	N_b
rate period subscripts		
offpeak - low	exogenous rate design	L
peak - high	exogenous rate design	H
critical	exogenous rate design	c
time invariant, uniform	exogenous rate design	u

what?	source	notation
IPR features		
value of rights / credit	IPR	$\mathbf{R}$
number of kWh per event protected by rights / credit	IPR	$\mathbf{q_R}$
IPR rebate rate	IPR	$\mathcal{P}_R$
baseline-rebate rebate rate	baseline-rebate	$\mathcal{P}_B$
declining block markup	IPR/exogenous	$\mathcal{M}$
number of kWh marked up by the declining block	IPR	$\mathbf{Q_D}$
bills	Note that B abbreviates baseline, not bill.	
CPP Bill	implication	TC^{CPP}
CPP-IPR Bill	implication	$TC^{CPP-IPR}$
Customer Characteristics		
minimum monthly consumption / ability to buy a hedge, shorthand for $\min_{m \in M}\{Q_m\}$	customer level	$\underline{\mathsf{Q}}_\mathsf{m}$
maximum consumption during an event / hedge need, shorthand for $\max_{m \in M}\{\frac{\sum_{cinC_M} q_c}{N_m}\}$	customer level	$\bar{\mathsf{q}}_\mathsf{c}$
deficit / cumulative undercontribution	customer level	δ_m