Resource and Energy Economics 33 (2011) 820–842

Contents lists available at ScienceDirect

Resource and Energy Economics

journal homepage: www.elsevier.com/locate/ree

Rethinking real-time electricity pricing

Hunt Allcott

Department of Economics, New York University, 19 W. 4th St., 6th Floor, New York, NY 10012, United States

ARTICLE INFO

Article history:
Received 15 March 2010
Received in revised form 10 June 2011
Accepted 14 June 2011
Available online 23 June 2011

JEL classification:
C93
L51
L94
Q41

Keywords:
Real time electricity pricing
Energy demand
Randomized field experiments

ABSTRACT

Most US consumers are charged a near-constant retail price for electricity, despite substantial hourly variation in the wholesale market price. This paper evaluates the first program to expose residential consumers to hourly real-time pricing (RTP). I find that enrolled households are statistically significantly price elastic and that consumers responded by conserving energy during peak hours, but remarkably did not increase average consumption during off-peak times. The program increased consumer surplus by $10 per household per year. While this is only one to two percent of electricity costs, it illustrates a potential additional benefit from investment in retail Smart Grid applications, including the advanced electricity meters required to observe a household's hourly consumption.

1. Introduction

Because electricity is very costly to store, wholesale prices vary from day to day and often fluctuate by an order of magnitude between low-demand nighttime hours and high-demand afternoons. Nearly all retail consumers, however, are charged some average price that does not reflect the wholesale price at the time of consumption. In theory, economists have long recognized that this creates allocative inefficiencies, and there is a long literature[1] on "peak load pricing" and "real-time pricing." In practice, the welfare implications of correcting this inefficiency fundamentally depend on how price elastic consumers are.

E-mail addresses: hunt.allcott@nyu.edu, allcott@mit.edu.

[1] The earliest peak load pricing discussion dates to Houthakker (1951), Steiner (1957), and Williamson (1966). Recent theoretical and simulation analyses include Borenstein (2005, 2007a,b), Borenstein and Holland (2005), and Holland and Mansur (2006, 2008).

0928-7655/$ – see front matter
doi:10.1016/j.reseneeco.2011.06.003

Recent advances in "Smart Grid" consumer energy management technologies can increase households' price elasticities and reduce the cost of the advanced electricity meters required to record hourly consumption. This has increased the likelihood that real-time pricing (RTP) would have positive net welfare effects, magnifying the business and policy interest in RTP. While electric utilities have experimented substantially with other price structures,[2] however, no utility has taken the seemingly natural step of exposing residential consumers to the real-time variation in wholesale market prices. This paper analyzes the program that has come closest, by providing hourly varying retail prices based on day-ahead wholesale prices.

Specifically, I examine the Energy-Smart Pricing Plan, which has operated in Chicago since 2003. From the households that opted into the pilot program, program managers randomly assigned a control group to be kept on the standard flat rate tariff, allowing an unbiased estimate of the treatment effects of real-time pricing on experimental households. To construct demand equations useful for welfare and policy analysis, I derive and estimate structural demand functions from indirect utility, while also highlighting the connections to a "reduced-form" treatment effect framework.

The demand estimation results bring to light three features of household electricity demand under RTP. First, the households that selected into the experiment have statistically significant elasticities: the overall reduced-form price elasticity of demand is about -0.1. Second, households' behavioral changes take the net form of energy conservation in peak price hours, with no net increase in consumption during low price hours. The idea that RTP could cause peak energy conservation with no net load shifting has important implications for the effects on energy costs, consumer welfare, the equilibrium entry of different power generation technologies, and the carbon emissions from electricity generation. The third central finding is that because the variation in hourly prices is small, reducing the cost of observing and responding to energy prices can substantially affect households' behavior. In this program, this was achieved by distributing Pricelights, glowing plastic orbs that change colors to indicate the current electricity price, to a set of households randomly selected from the group that had requested the device.

The demand system results are then used to estimate the program's effects on consumer surplus, under the simplifying assumption that wholesale prices and retailer profits are fixed. The calculation shows that this RTP program resulted in an annualized consumer surplus gain of approximately \$10. This is 1–2 percent of the average household's electricity expenditures. As the welfare gains from RTP increase in the amount of wholesale price variation, however, this number would likely be higher in other years and regions, where electricity prices tend to vary much more. I then extend the welfare analysis by constructing a simple model of the greenhouse gas emissions from the marginal electricity suppliers in each hour. This shows that the program reduced carbon dioxide emissions by just over four percent.

Because the demand parameters are estimated from an experiment into which households had self-selected, they are not informative about the price elasticity of the general population. Analysis of optional residential RTP, however, is of great policy interest. Although the regulators that approve electric utilities' pricing structures typically share economists' intuition about the benefits of RTP, they are often concerned that consumers perceive RTP as complicated or risky relative to average cost pricing (Faruqui and Sergici, 2009). A further political economy problem is that a mandatory shift to real-time pricing could increase electricity bills for many consumers that tend to use more electricity than average at times when market prices are high (Borenstein, 2007b). As a result, many real-time pricing programs introduced over the medium-term may be optional instead of mandatory, and the demand parameter estimates from this program may be of particular interest in understanding the early phases of these future optional programs.

Although this discussion is specific to a particular setting, real-time electricity pricing is a manifestation of an extremely general economic problem. Monitoring costs cause consumers to take

[2] There is empirical evidence on the response of larger commercial and industrial customers to RTP, including Patrick and Wolak (2001), Boisvert et al. (2007), Herriges et al. (1993), and Taylor et al. (2005). Barbose et al. (2004) provides a comprehensive overview on real time pricing programs operated by US utilities. There is also a substantial literature on residential electricity demand under other price structures, such as "critical peak pricing" and "time of use pricing." See Faruqui and Sergici (2008) for an overview of recent work.

unobserved actions that differ from the perfect information optimum. Observing actions, however, only improves welfare if consumers' behavior under the new contract (real-time pricing) changes sufficiently to justify the monitoring or contracting cost. Somewhat more specifically, there are many settings when prices do not fully reflect how input costs or the shadow price of capacity vary over time, including cellular phone and internet service contracts, most restaurants, and bridge and highway tolls. In theory, the firm or regulator in these settings determines the profit-maximizing or socially optimal contract using price elasticities and consumer surplus analyses similar to those carried out here.

The introduction proceeds first by motivating why now is an important time to study real-time pricing. Section 2 details the experimental design, the marketing and recruitment process, and baseline household characteristics. In Sections 3 and 4, demand functions are derived from indirect utility and estimated, exploiting the randomization to purge the estimator of simultaneity bias. Section 5 presents empirical results, and Section 6 discusses policy implications.

1.1. Motivation: why study real-time pricing?

Real-time pricing is one of the central issues in an important industry. In 2007, the electric power sector accounted for 2.5 percent of United States GDP, or $326 billion in retail sales per year (US Energy Information Administration, 2008). Broader discussions of wholesale and retail electricity market design, such as Bandt et al. (2003), Borenstein (2002), Joskow and Tirole (2006, 2007), and Wolak (2007), often center on the importance of real-time pricing for market efficiency.

The most commonly discussed channel of efficiency gains from moving consumers from the standard flat rate tariff to real-time pricing is static allocative efficiency improvements: conditional on a capital stock of power plants, there are gains from shifting consumption from peak periods when the marginal cost of production is high to off-peak periods when marginal cost is low (Holland and Mansur, 2006). As discussed in Allcott (2009a), Borenstein and Bushnell (1999), and Borenstein (2005), however, there are other important channels. The inelastic demand that results from the lack of retail RTP is one of the central challenges in designing electricity markets: inelastic short-term demand gives producer firms market power in wholesale markets, and markups above marginal cost can cause an inefficient allocation of production between firms. Inelastic short-term demand, combined with the extremely high cost of blackouts, also means that electricity market operators must ensure that sufficient generating capacity is in the ground to satisfy extreme realizations of the demand shock. In an industry where capacity is a substantial part of the cost structure, the capacity reduction that could result from implementing RTP is a substantial potential source of welfare gains. These potential gains from real-time pricing, of course, depend on the magnitude of consumers' price elasticity.

This is a particularly interesting time to be studying real-time pricing. Most US households currently have electricity meters that simply record the total consumption of electricity since installation, meaning that the consumer cannot be charged prices that vary from hour to hour. Furthermore, the only way for the electric utility to observe households' consumption is to actually send a worker to read the meter, a costly and potentially error-prone process. The "Smart Grid" is a set of emerging electric power information technologies that include, among other things, household energy management devices and technologies that facilitate communication between electricity retailers and consumers. From the utility's perspective, improvements in these technologies offer reduced meter reading and administrative costs and the potential for real-time metering of electricity use. Furthermore, by allowing households to more easily observe prices and consumption, and even to automate how air conditioners and other appliances turn on and off in response to real-time prices, Smart Grid technologies can increase consumers' price elasticity of demand.

Substantial business and policy interest has been associated with these technological changes. Electric utilities in California, Colorado, Florida, Illinois, Indiana, Ohio, Texas, Washington, and other states are introducing Smart Grid technologies to large groups of customers. The Energy Independence and Security Act of 2007 provides $100 million annually in research and development funding and provides incentives for utilities to invest in Smart Grid infrastructure. Furthermore, the 2009 US economic stimulus package includes $3.9 billion in funds for wholesale and retail-level Smart Grid projects. Now that a larger share of households will have meters capable of administering real-time pricing, utilities, regulators, and policymakers will want to understand the efficiency gains from that pricing structure relative to the remaining perceived costs.

There is a long empirical literature on different forms of household electricity pricing. Reiss and White (2005), for example, examine *increasing block pricing*, under which the marginal price increases by the total quantity consumed. Wolak (2006) estimates consumers' response to *critical peak pricing*, where consumers pay higher prices or receive rebates for conservation on occasional high price days, but pay their standard rate at all other times. Train and Mehrez (1994), the analyses in Aigner (1984), and many others focus on *time of use pricing*, in which customers pay different fixed prices in peak and off-peak hours. While responses to these other price structures can suggest a reasonable range of responses to RTP, these distinct structures provide a distinct set of short run and long run incentives. Despite this experimentation, the vast majority of US households are currently on a very simple *flat rate tariff*, in which the marginal price per kilowatt-hour consumed does not vary other than year-to-year or perhaps season-to-season.

A number of utilities have some large commercial and industrial customers on RTP, and such programs have been analyzed in Patrick and Wolak (2001), Boisvert et al. (2007), Herriges et al. (1993), Taylor et al. (2005), and other work. These larger customers, and in particular firms for which electricity represents a large share of input costs, could respond in very different ways. Because the fixed costs of advanced meters are a greater share of total electricity costs, and because of the perceived complexity and risk to these smaller consumers, real-time pricing has historically been a lower priority for residential relative to commercial and industrial customers. Indeed, after the early time of use pricing experiments in Aigner (1984), conventional wisdom held that although households had statistically non-zero price elasticities, these were not large enough to justify the expense of installing advanced electricity meters (Faruqui and Sergici, 2008). In recent years, the prior about the net benefits of residential RTP has begun to change, but no direct evidence existed until recently.

The ESPP pilot program analyzed here is of particular interest because it exposes residential consumers to hour-to-hour price variation, which has never been done before in the United States. This is as close as utilities have come to residential real-time pricing.[3]

2. Experimental design and data

This section begins by presenting substantial background on the program and participants, which allows more insight into the nature of households' self-selection. I then detail the experiment itself and present descriptive statistics.

2.1. Setting: the energy-smart pricing plan

At the beginning of the decade, demand growth had pushed parts of Chicago's local electricity distribution network near their capacity limits. A large electric utility called Commonwealth Edison (ComEd) owns this local network and provides retail electricity service to residential customers at prices regulated by the state. Temporarily prohibited by state electricity restructuring rules from financing infrastructure investments through higher electricity rates, ComEd was interested in low cost strategies to reduce demand during peak periods.

A Chicago NGO called the Center for Neighborhood Technology (CNT) helped ComEd to design and operate air conditioner replacement programs, in which a household with an inefficient but functional window air conditioner could receive a rebate for trading in for a new, energy efficient model. These programs were targeted to several specific neighborhoods where both ComEd's infrastructure was

[3] I note that this experiment has been previously studied in a consulting report (Summit Blue Consulting, 2004). Remarkably, the report did not exploit the randomized control group. The authors recognized that a "comparison approach" between treatment and control groups was possible, but worry that "the relationship between hourly prices is complicated, and therefore it is difficult to determine from this approach" whether treatment effects are correlated with prices. They conclude that if the control group were used in the analysis, "it would be difficult to directly determine the level and significance of the relationship between hourly prices and usage." In Section 5.1, I detail how non-experimental approaches appear to produce biased parameter estimates.

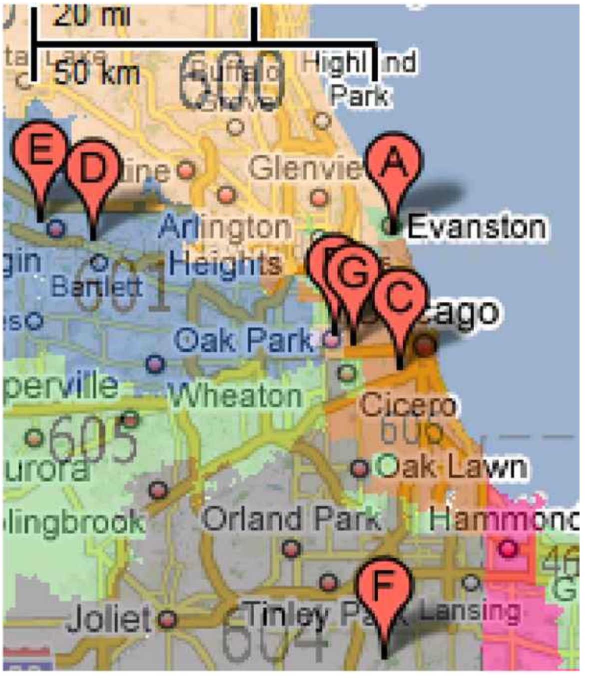

Fig. 1. ESPP geographic areas. CNT neighborhoods: A: Evanston; B: Austin; C: Pilsen; D, E: Elgin; F: Park Forest; G: Near West Side.

under stress and CNT believed it could operate effectively. In 2003, CNT initiated the Energy-Smart Pricing Plan (ESPP) to test whether real-time pricing could incentivize significant reductions in peak electricity demand. This was a convenient partnership for ComEd, because state law also prevented them from promoting new products or rates[4] (Isaacson et al., 2006).

2.2. Recruitment and baseline characteristics

CNT's outreach targeted its existing areas of operation, shown in Fig. 1, and in particular focused on households that had chosen to participate in the previous air conditioner replacement programs. Beginning in late 2002, CNT mailed marketing materials to their 7000 existing participants, organized community meetings in the areas where they operated, and publicized the new program via word of mouth. By the end of April 2003, 693 households had opted in. Although the program was open to all ComEd customers in and around Chicago, over two-thirds of the initial households had been participants in previous CNT programs. Of the 225 that had no previous affiliation with CNT, most lived in the existing areas of operation and had found out via word of mouth or the community meetings.

CNT and a consultant carried out a survey of ESPP program households and of other CNT affiliate households that had received direct mail marketing materials but did not sign up for the program. The survey indicated that saving money was by far the primary factor driving enrollment, followed by "environmental benefits." As shown in Fig. 2, the most common way that enrollees found out about the program was through direct mail. Among CNT households that had not enrolled, half did not recall

[4] Restrictions against raising rates and promoting new rate structures were common features of state electricity market restructuring law passed in the late 1990s. The rate freeze was typically used to garner political support from consumer groups, while the prohibition against the incumbent utility offering new products or rates was intended to encourage competing retail electricity suppliers to enter the market.

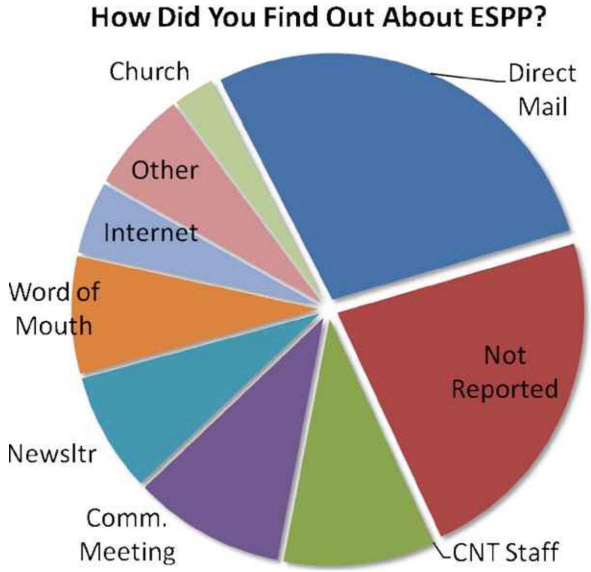

Fig. 2. Recruitment.

hearing about the program, while others did not expect sufficient savings or thought that the program was too risky or complex (Summit Blue Consulting, 2004).

Since many program households had participated in CNT's earlier air conditioner replacement programs, they were probably more interested in energy conservation or attentive to electricity use than the Chicago population average. The fact that these households had recently purchased an energy efficient air conditioner, however, means that their price response could also be understated relative to the general population. This is because treatment group households would have had one less inefficient air conditioner to replace, and also because turning down an efficient air conditioner reduces consumption less than turning down an inefficient model.

Of the 693 initial households that enrolled in the pilot, program managers randomly assigned 103 to a control group. Control group households received a letter saying that they were not on real-time pricing, complete with three $15 gift certificates for groceries as a consolation. These households received no further information during the 2003 experiment and were only allowed to enter real-time pricing at the beginning of 2004.

I observe each household's Census tract and use this to incorporate tract-level information from the 2000 Census Long Form on housing stock (median year of construction and percent of dwellings that are multifamily vs. single family) and demographics (percent of individuals not in the labor force and percent with a college diploma). The 693 treatment and control households were in 255 different tracts. At the household level, I also observe income (in six buckets) and number of household members, although some missing data is imputed from census tract medians. Finally, I use the monthly electricity bills for May through December 2002 to construct pre-program baseline electricity consumption.

Columns I and II of Table 1 presents the means and standard deviations of the treatment and control groups' baseline observable characteristics. Column V presents the difference in mean characteristics between treatment and control group households, with standard errors in parenthesis. As Column V shows, treatment group households are slightly but not statistically significantly larger and more concentrated in the higher income buckets. They are also similar on average household size and on the tract-level demographic and housing stock variables. An F-test of a regression of a treatment group indicator on observable characteristics fails to reject that

Table 1
Baseline household characteristics.

	(I) Treatment	(II) Control	(III) CNT	(IV) Chicago	(V) T-C	(VI) ESPP-CNT
Pre-program usage	954	879	1050	1050	75	−135
SD (and SE):	(497)	(436)	(629)	(629)	(47.5)	(40.8)**
Household size	2.56	2.58	2.65	2.64	−0.024	−0.059
	(1.02)	(1.09)	(1.66)	(1.55)	(0.110)	(0.100)
Income 10k–25k	0.20	0.26	0.19	0.14	−0.058	0.034
	(0.27)	(0.36)	(0.39)	(0.35)	(0.037)	(0.024)
Income 25k–50k	0.29	0.32	0.29	0.24	−0.029	0.022
	(0.31)	(0.37)	(0.45)	(0.43)	(0.038)	(0.029)
Income 50k–75k	0.23	0.21	0.19	0.22	0.023	0.026
	(0.29)	(0.31)	(0.39)	(0.42)	(0.033)	(0.026)
Income 75k–150k	0.18	0.13	0.18	0.24	0.053	−0.015
	(0.27)	(0.23)	(0.38)	(0.42)	(0.025)**	(0.025)
Income >150k	0.04	0.03	0.04	0.08	0.016	−0.014
	(0.14)	(0.10)	(0.20)	(0.26)	(0.011)	(0.014)
Median year constructed	1950	1949	1948	1962	1.31	0.60
	(11.4)	(9.00)	(8.38)	(14.6)	(1.00)	(0.67)
Pct multifamily housing	0.49	0.48	0.63	0.44	0.017	−0.051
	(0.29)	(0.29)	(0.48)	(0.50)	(0.031)	(0.031)*
Pct not in labor force	0.40	0.41	0.40	0.39	−0.005	0.015
	(0.10)	(0.09)	(0.49)	(0.49)	(0.010)	(0.028)
Pct college graduates	0.25	0.23	0.19	0.31	0.023	0.041
	(0.19)	(0.17)	(0.39)	(0.46)	(0.018)	(0.025)*
Previous CNT affiliates	421	450	252	26.0	−29.1	187
	(574)	(603)	(480)	(170)	(63.7)	(33.4)**
F-Test *p*-Value					0.263	0.000

The first four columns present mean characteristics, with standard deviations below in parenthesis. Columns V and VI present the difference between mean characteristics in Treatment and Control and ESPP vs. CNT neighborhood households, respectively. For Columns V and VI, standard deviations are below in parenthesis. *, **Statistically different from zero with 90 and 95 percent confidence, respectively. CNT neighborhoods are the 30 zip codes that had more than 10 CNT affiliates before the ESPP program began. ESPP Treatment Group: 590 households. ESPP Control Group: 103 households. CNT neighborhoods: 631 thousand households. Chicago four-county metropolitan area: 3.14 million households. For ESPP households, Pre-Program Electricity Usage (W), Household Size, and Income buckets (Real 2003 $/year) are observed at the household level. Previous CNT Affiliates is observed by zip code. All other variables are at the Census tract level; from 2000 US Census Long Form data. CNT and Chicago electricity usage distribution are from microdata for the East North Central division from the 2001 Residential Energy Consumption Survey US Energy Information Administration, 2005. For CNT-area and Chicago households, all other data are from the 2000 Census demographic distributions.

the treatment and control groups are the identical, as we would expect in a randomized experiment.

Columns III and IV of Table 1 present the means and standard deviations of characteristics of households in the 30 zip codes where CNT had ten or more existing affiliates and in the Chicago metropolitan area, respectively. Column VI presents the difference in mean characteristics between the ESPP participants and the population in the CNT zip codes, with standard errors in parenthesis. On household size and income, ESPP enrollees are not statistically different from households in the 30 CNT zip codes. Their pre-program energy consumption is lower, likely resulting from their previous participation in CNT's air conditioner replacement programs. Program participants are disproportionately drawn from the better-educated of the areas where CNT had ten or more affiliates, although CNT itself had focused its operations in areas with lower average education than the rest of Chicago.

In sum, while program participants are not highly unusual on observable characteristics, they are a self-selected group. As a result, this experiment would not be very useful in understanding the effects of a mandatory, population-wide real-time pricing program. It does, however, allow us to consistently estimate internally valid demand parameters for the experimental population, which should be indicative of the first few percent of households that would opt in to a first wave of optional residential real-time pricing programs in similar parts of the country.

2.3. The experiment

Households assigned to the control group were forced to remain on the standard ComEd residential tariff, which is 8.275 cents per kilowatt-hour (kWh) in the summer and 6.208 cents/kWh in other seasons. For the treatment group, prices were set such that expected total electricity bills would be slightly lower under RTP compared to the flat rate tariff, for a household with a typical load shape. This reduction in total expected costs was achieved through a Participation Incentive, which was designed to compensate households for an increase in perceived price risk. ComEd fixed each day's hourly retail prices by 4 pm the day before, according to the following formula:

$$p_{hd} = P_{hd}^{DA} + D - PI \tag{1}$$

p_{hd} = retail price for hour h of day d; P_{hd}^{DA} = Day-Ahead[5] wholesale market price; D = distribution charge (5.0 cents/kWh); PI = Participation Incentive (1.4 cents/kWh).

On the evenings before days when the wholesale component of price was to exceed 10 cents/kWh, treatment group households received a special e-mail or telephone High Price Alert. This happened on nine summer days in 2003. Program managers also made available programmable thermostats to the treatment group, which allowed automated temperature control by time of day.

Survey results, website login data, and discussions with program managers and participants indicate that although prices were available via telephone and internet, only rarely did households actively check prices. Treatment group households could, however, form reasonably precise beliefs about the joint distribution of prices with season, hour, and temperature. To help inform these beliefs, households were sent quarterly ''ESPP Updates'' and a ''Summer Readiness Kit,'' which explained that prices are typically higher during the afternoon, on hot days, and in the summer.

In an effort to increase the ease with which households could observe, and thus respond to, hourly price fluctuations, program managers introduced a device called the Pricelight. This is a small plastic globe that changes colors in real-time on a continuous spectrum from blue to red, indicating low to high electricity prices. CNT offered Pricelights to all ESPP participants in 2006. Of the 223 households that submitted requests, 47 were randomly selected to receive the device. These households form the treatment and control groups for a separate experiment that allows an estimate of the treatment effect of owning a Pricelight conditional on being on real-time pricing and having requested the device.

After a household had opted in to the program, and regardless of whether the household was assigned to treatment or control, ComEd installed a new electricity meter to record hourly consumption. I observe hourly electricity consumption for all program households between 2003 and 2006. This includes 3.98 million hourly observations from the randomized RTP experiment from May through December 2003, as well as 814,000 observations from the Pricelight experiment between June and October 2006.

Table 2 shows descriptive statistics for the 2003 RTP experiment. Compared with the same months of the previous year, which had higher temperatures, treatment group households reduced consumption by 90.6 W,[6] or about 10 percent, while control group households reduced consumption by about 5 percent. Table 3 displays similar information for the Pricelights experiment in 2006.

An example investment decision may help put these price and consumption figures in context. A household deciding to purchase a window air conditioner might choose between a standard model, for

[5] Specifically, the prices between 6 am and 10 pm were on-peak Day-Ahead prices from a nearby region, as reported in Platt's Energy Trader, applied to the shape of hourly Locational Marginal Prices at the PJM West hub. Between 10 pm and 6 am, hourly wholesale prices have little variance and were based on seasonal historical averages. As of 2008, the program uses the PJM ComEd zone Balancing Market Locational Marginal Price for all hours.

[6] I will sometimes refer to treatment effects in watts, which is a measure of power, the flow of energy. In this setting, energy consumption is actually measured on an hourly basis, in total watt-hours. Treatment effects could also be thought of in ''average watts over an hour.''

Table 2
Descriptive statistics: RTP experiment.

	Obs	Mean	SD	Min	Max
RTP retail price (summer)	2208	6.85	2.12	4.62	16.0
RTP retail price (summer peak)	384	9.73	2.46	5.27	16.0
RTP retail price (non-summer)	3672	6.20	1.07	4.78	11.0
Control retail price (summer)	1	8.275	0	8.275	8.275
Control retail price (non-summer)	1	6.208	0	6.208	6.208
1 (high price hour)	5880	.0052	.072	0	1
1 (high price day)	245	.037	.19	0	1
Treatment quantity	3,396,511	865	810	0	14,750
Control quantity	581,520	830	900	0	21,880
Quantity – baseline use (T)	3,396,511	−90.6	694	−3860	12,300
Quantity – baseline use (C)	581,520	−47.5	787	−2380	19,890
N (total households)	5880	677	9.40	658	689
N (treatment)	5880	578	6.67	563	586
N (control)	5880	98.9	3.16	95	103

Includes the RTP experimental period, May–December 2003. "Summer peak" includes noon to 6 pm on all workdays from June to August. Observations column reflects distinct observations. Quantities are in watts; prices are in cents/kWh.

Table 3
Descriptive statistics: Pricelight experiment.

	Obs	Mean	SD	Min	Max
RTP retail price (summer)	2208	9.24	3.50	3.71	40.1
RTP retail price (non-summer)	1464	7.83	1.55	4.29	14.9
1 (high price hour)	3672	.013	0.11	0	1
1 (high price day)	153	0.065	0.25	0	1
Treatment quantity	171,867	1190	1180	0	16,700
Control quantity	642,095	1090	1040	0	16,600
Quantity – baseline use (T)	171,867	−91.1	872	−4790	11,100
Quantity – baseline use (C)	642,095	−44.3	876	−2640	15,300
N (total households)	3672	222	0.693	219	223
N (treatment)	3672	46.8	0.396	46	47
N (control)	3672	175	0.369	173	176

Includes the Pricelights experimental period, June–October 2006.

$270, and an efficient "Energy Star" model, for $300. Air conditioning can represent a substantial portion of household electricity consumption: when turned to its highest setting, the standard model uses a flow of 1000 W, while the energy efficient model requires about 100 W less. At standard usage[7], a household on ComEd's flat rate tariff saves $5.34 per year with the Energy Star model and chooses that model if it discounts these future cash flows at less than 5.8 percent per year. A household on real-time pricing would save $6.87 per year on the prices observed during the first four years of the experiment and would purchase the Energy Star model under a discount rate of 13.4 percent or less. Note that program households did have long run incentives: CNT and ComEd promoted RTP as a permanent program, and as of 2008, there were approximately 5000 households enrolled.

The retail prices during the experiment ranged from 4.62 to 16.0 cents/kWh, and there were 30 h on nine days in which the wholesale component of price exceeded the High Price Alert cutoff of 10 cents/kWh. Both the level and variability of wholesale prices are lower during the experiment than they are

[7] The calculation assumes that the air conditioner lasts seven years and is used 683 hours per year, as suggested by the Energy Star website (US Department of Energy, 2008). Air conditioner capital costs are from the same source. The RTP savings assumes that this 683 h of usage is spread equally across the summer months between 10 am and 6 pm. Note that the Participation Incentive lowered electricity prices for RTP customers by 1.4 cents per kilowatt-hour and thus somewhat reduces their estimated savings. Note further that even these relatively small savings may be overstated if air conditioners are used during evening hours, which in Chicago are still quite hot but have lower prices than 10 am to 6 pm.

in most power markets today. There are a number of ways of quantifying this difference; one simple measure is the mean and standard deviation of daily average prices. During the RTP experiment, these are 2.8 cents/kWh and 0.85 cents/kWh, respectively. By comparison, the mean and standard deviation of daily average prices in 2010 from a selection of major wholesale hubs across the United States are 4.7 cents/kWh and 1.4 cents/kWh, respectively.[8]

Five percent of households attritted from the sample over the eight month experiment. The attrition was reportedly due to customers closing accounts as they moved, and ComEd has confirmed that this is consistent with the rate at which their general residential customer base closes accounts. There is no statistical difference between the attrition rates of the treatment and control groups ($p = 0.251$). Furthermore, within each group, attrition does not appear to be correlated with observable characteristics. The first two columns of Table 4 present regressions of an attritter indicator variable on household characteristics in each group, and F-tests cannot reject that attrition is uncorrelated with observables. This is consistent with the proposition that attrition is random and thus independent of demand parameters, which would be sufficient for the parameter estimates to be unbiased by attrition. Only four households attritted from the sample during the Pricelights experiment.

3. Indirect utility and demand functions

In this section, I derive demand functions from indirect utility. This structural approach is useful for welfare analysis and counterfactual simulations. The results section will discuss reduced form analyses as well as estimates of these structural demand parameters.

Household i has utility $V(\mathbf{P}, w_i)$, which depends on wealth w and the vector $\mathbf{P}$ of electricity prices and High Price Alerts in all future days. The Gorman form is used, because it will give simple linear[9] demand functions and has zero wealth effects, which is reasonable given that electricity is a small share of the household budget. Assuming negligible time discounting over the experimental period, we have:

$$V_i(\mathbf{P}, w_i) = w_i - \sum_d \left(\sum_{h=1}^{24} p_{hd} \cdot \left(\frac{1}{2} \overset{D}{\underset{ih}{\eta}} p_{hd} + (\overset{A}{\underset{is}{\eta}} - \overset{D}{\underset{ih}{\eta}})\, \bar{p}_{hs} + \overset{HP}{\underset{ig}{\eta}} HP_d + \overset{T}{\underset{is}{\eta}} T_i + \xi_{ihd} \right) \right) \tag{2}$$

$\{\eta_{ih}^D, \eta_{is}^A, \eta_{ig}^{HP}, \eta_{is}^T\}$ = demand parameters for household i, where g denotes a group of hours; ξ_{hd}= demand shifter in hour h of day d. (Note that ξ need not have zero mean.); $\bar{p}_{hs}$ = average price for hour h in season s; HP_d= High Price Alert day indicator function; T_i= indicator for whether household i is enrolled in RTP.

Household i enrolls in real-time pricing if its expected utility from entering and participating in the program is greater than its expected utility on the standard flat rate tariff. Econometric unobservables influence selection in two ways. First, households with stronger demand parameters are more likely to enroll. Second, households that naturally have lower demand during high priced hours, i.e. whose demand shifters ξ_i are less correlated with prices, are more likely to enroll.[10] This analysis focuses on

[8] The underlying data are from the U.S. Energy Information Administration's website, at http://www.eia.gov/cneaf/electricity/wholesale/wholesale.html. I used all hubs where data were available for 2010, which were NEPOOL, AEP-Dayton Peak, Entergy Peak, ERCOT South, PJM West Peak, and SP-15 EZ.

[9] Any applied specification of the functional form for indirect utility and demand functions will be ad-hoc, and there is little evidence on the appropriate functional form for electricity demand curves. Constant elasticity of substitution demand curves would be a natural alternative to the linear specification used here. In this setting, however, non-parametric examinations of average treatment effects as a function of prices suggest that treatment effects only become larger in magnitude for retail prices above 10 cents per kilowatt-hour. There is little apparent elasticity below that level, suggesting that the CES specification would overstate elasticity at low prices. While linear demand functions would presumably overstate price elasticity at higher prices, this is not a concern here because demand functions will only be fitted to the relatively moderate prices in sample.

[10] As suggested by Borenstein (2007b), for most commercial and industrial consumers, the savings from RTP due to their price elasticity are outweighed by the transfers implied by how their natural load shapes (represented by ξ_i) interact with the price changes implied by moving from the flat rate tariff to RTP. If this holds for residential customers, this implies that the selection of consumers with lower demand during high price hours ("flatter load shapes") is more powerful than selection of consumers with higher unconditional elasticities. These consumers with flat load shapes, furthermore, may also actually have lower unconditional elasticities. As a result, although it seems plausible that those who enrolled in the program are more price elastic, this is not an unambiguous theoretical result.

Table 4
Attrition.

	Treatment	Control
	(1)	(2)
Pre-program usage	8.49E−06	−1.47E−06
	(0.00002)	(0.00009)
Household size	0.0001	0.0007
	(0.007)	(0.009)
Income 10–25k	−.028	−.009
	(0.068)	(0.051)
Income 25–50k	−.005	0.002
	(0.071)	(0.044)
Income 50–75k	−.026	0.027
	(0.069)	(0.047)
Income 75–150k	−.050	0.07
	(0.07)	(0.074)
Income >150k	−.068	0.117
	(0.075)	(0.197)
Med const year	0.0008	0.002
	(0.0008)	(0.002)
Pct multifamily housing	0.112	0.037
	(0.038)**	(0.062)
Pct not in labor force	0.028	0.064
	(0.179)	(0.167)
Pct college graduates	0.018	0.024
	(0.072)	(0.08)
Const.	−1.662	−3.615
	(1.657)	(3.436)
Obs.	590	103
R^2	0.024	0.009
F statistic	1.459	0.308
F test *p*-value	**0.983**	**0.143**

Dependent variable is an indicator for whether the household attritted. Standard errors in parenthesis. *, **Statistically different from zero with 90 and 95 percent confidence, respectively. Pre-program electricity use is in watts. Income buckets are real 2003 dollars per year.

estimating the (internally valid) average demand parameters for the population that enrolled in the experiment. I thus drop the individual subscripts on demand parameters, and it is understood that the η parameters are the averages for the group that enrolled.

The demand function for household *i* on hour *h* of day *d* is derived from Roy's Identity and a minor re-arrangement:

$$q_{ihd} = \eta_s^A \bar{p}_{hs} + \eta_h^D (p_{hd} - \bar{p}_{hs}) + \eta_g^{HP} HP_d + \eta_s^T T_i + \xi_{ihd} \quad (3)$$

The η parameters were conceived to represent three different responses to time-varying prices. First, the parameter η^A captures responses to the average hourly price shape, such as habitually using a washing machine during low price evening hours instead of high price afternoon hours. The year is broken into two seasons, summer and non-summer, and separate parameters allowed for each of the two average price shapes. The η_s^T parameters, one for each season, are the intercepts for the response to hourly price shape that was parameterized by η^A. These capture the level shift in average consumption for the real-time pricing treatment group.

The second response to time varying prices is captured by the η_h^D parameters, which are responses to deviations from the average price shape. The two parameters η^A and η_h^D are separately specified because large scale RTP would likely affect both the typical price shape and the magnitudes of deviations from that shape, and households could respond differently to the two types of variation. A third type of price response is captured by η_g^{HP}, which reflects additional consumption changes associated with High Price Alert days. Separate parameters are estimated for four hour groups g: Early Morning (6–10 am), Morning (10 am–2 pm), Afternoon (2 pm–6 pm), and Evening (6 pm–10 pm).

These demand functions do not include intra-day substitution parameters, through which consumption in hour h could be affected by price in another hour of the same day. While this restriction is undesirable, it is necessary because price variation is principally movement of an entire day's prices up or down, affecting relative hourly prices by a fairly constant proportion. Prices in different hours of the same day have correlation coefficients of roughly 0.9, and this collinearity makes it impossible to separately estimate both η_h^D and intra-day substitution parameters.[11] This means that η_h^D should be thought of as the association between consumption in hour h and the deviation of the day's prices from average, which may include response to price in hour h as well as some intra-day substitution.[12] Also, notice that while price responsiveness might differ between weekends and weekdays, all day types are pooled because there is not sufficient weekend data to separately estimate these parameters with sufficient precision.

4. Identification and estimation

As in the typical demand estimation, the demand shifter ξ is unobserved. Attempting to estimate the demand functions using only treatment group usage data would be biased by simultaneity, because the same unobservable factors affecting RTP households' demand shifters ξ also shift the aggregate market-level demand curve and thus affect equilibrium prices. The randomized control group, however, has in expectation the same ξ as the treatment group, meaning that any difference in demand is the effect of response to real-time prices. Intuitively, the variation in average treatment effects across hours with different prices can identify the η parameters.

This can be formalized using the "potential outcomes" framework of Rubin (1974) and the program evaluation literature that builds on his approach. Upon enrollment, each household has two possible states of the world: one in which it is randomized into the real-time pricing treatment, and one in which it is randomized into the control group. Define $q_{ihd}(T_i=1)$ as household i's potential consumption in hour h of day d in the treated state, while $q_{ihd}(T_i=0)$ is the potential consumption if it were assigned to the control. For each household, consumption q_{ihd} is observed only for the state to which it was actually assigned:

$$q_{ihd} = q_{ihd}(T_i) = \begin{Bmatrix} q_{ihd}(T_i = 0) & \text{if } T_i = 0 \\ q_{ihd}(T_i = 1) & \text{if } T_i = 1 \end{Bmatrix} \tag{4}$$

The average treatment effect (ATE) is the expected effect on program households' consumption in hour h of day d caused by being on real-time pricing instead of the flat rate tariff:

$$\tau_{hd} = E[q_{ihd}(T_i = 1) - q_{ihd}(T_i = 0)|h, d] \tag{5}$$

The control group paid ComEd's seasonal flat rate residential tariff $\bar{p}_s^{T=0}$. For both the summer and non-summer days, substituting this into the demand function gives:

$$q_{ihd}(T_i = 0) = \eta_s^A \bar{p}_s^{T=0} + \xi_{ihd} \tag{6}$$

[11] Some analyses of real time pricing for larger commercial and industrial customers (e.g. Patrick and Wolak (2001) and Taylor et al. (2005)) do not have this collinearity problem because they analyze programs that offered each day's Balancing Market prices. These prices covary less than the Day-Ahead market prices that the ESPP program used at the time of the experiment.

[12] A possible behavior that the specification will thus misrepresent is "pre-cooling," in which a consumer air conditions the house in the morning hours to reduce the need for cooling in a relatively high price afternoon. Based on the energy use changes reported by program households, both anecdotally and in small surveys conducted by the program managers, it appears that intra-day substitution may not be large in magnitude. The demand functions do capture what seem *a priori* to be two more likely forms of substitution. First, the strongest intra-day substitution should be observed on High Price Alert days, and the net effect of this is captured through the η_g^{HP} parameters. Second, the η^A parameters measure substitution between hours of the average day based on the average price shape.

The randomization of the N households into treatment and control groups allows the identifying assumption that the groups' average demand shifters are equal as $N \rightarrow \infty$:

$$E[\xi_{ihd}|T = 1, h, d] = E[\xi_{ihd}|T = 0, h, d] \tag{7}$$

Subtracting this and the treatment and control groups' demand functions, each hour's average treatment effect can be parameterized as:

$$\tau_{hd} = E\left[\eta_s^A(\bar{p}_{hs} - \bar{p}_s^{T=0}) + \eta^D(p_{hd} - \bar{p}_{hs}) + \eta_g^{HP} HP_d + \eta_s^T | h, d\right] \tag{8}$$

The demand parameters can therefore be consistently estimated by pooling across all hours of the experiment and adding a fixed effect for each of the 5880 h observed. Including control variables, the estimating equation is:

$$\begin{aligned} q_{ihd} &= \left\{\eta_s^A(\bar{p}_{hs} - \bar{p}_s^{T=0}) + \eta^D(p_{hd} - \bar{p}_{hs}) + \eta_g^{HP} HP_d + \eta_s^T\right\} \cdot T_i \\ &\quad + \left\{\alpha_1(\bar{p}_{hs} - \bar{p}_s^{T=0}) + \alpha_2(p_{hd} - \bar{p}_{hs}) + \alpha_3 HP_d^{aft} + \alpha_4 HP_d + \alpha_5\right\} X_i + \zeta_{hd} + \varepsilon_{ihd} \end{aligned} \tag{9}$$

X={Pre-Program Average Hourly Consumption, Household Size, log(Income)}; HP_d^{aft} = indicator for an afternoon hour of a High Price Alert day; ζ_{hd} = fixed effect for hour h of day d; ε_{ihd} = econometric error.

The estimation uses the standard fixed effects estimator. The data are demeaned to remove fixed effects ζ, and ordinary least squares is applied to the demeaned data. Standard errors are Newey-West, allowing a 1-h lag. In a structural sense, the "econometric error" ε_{ihd} is part of the household's demand shock ξ_{ihd} that is residual of the other covariates.

Although simultaneity bias has been removed via the randomization, there are remaining limitations. One concern is that although the causal effect of being on real-time pricing is identified for each hour, the demand functions are an *a priori* imposition of functional form. This means that the parameters η are only causal in the (unlikely) event that the demand functions are correctly specified. For example, unless consumers' true demand functions are linear in prices, $\hat{\eta}^{HP}$ is not a consistent estimate of the causal impact of a High Price Alert on consumption.[13] As an illustration of this issue, recall that purchasing an energy efficient air conditioner reduces consumption most on hot days, which are more likely to be High Price Alert days. The estimated coefficient $\hat{\eta}^{HP}$ could therefore be nonzero even if households had no incremental response to the Alerts. There is an analogous problem in discrete choice demand estimation in characteristics space: analysts typically impose some simple functional form for indirect utility as a function of characteristics, but these characteristics are often correlated with each other and are not randomly assigned to products.

A second and related limitation is that we do not observe the behaviors that underlie the treatment effects or when those behaviors occurred. This procedure simply estimates the correlation between average treatment effects and prices. It does not identify the frequency at which behavioral changes occurred, i.e. whether the responses were day-to-day, short-term adjustments to air conditioner settings or long-term adjustments to thermostats each season. It similarly does not identify whether the effects were produced by long run changes to energy-using capital stock versus short run changes to the usage of that capital stock.[14]

[13] A natural way to estimate responses to High Price Alerts would be a regression discontinuity framework, in which demand on days with highest wholesale price just below \$0.10 per kWh is compared to demand on days with highest price just above that cutoff. Unfortunately, there were not enough High Price Alert days near the cutoff to carry this out, even when including additional summers of non-experimental data from 2004 to 2006.

[14] Short-term price response could be identified by instrumenting for hourly prices with short-term supply shocks that are exogenous to the unobservable demand shifters. One suggestive test of the validity of the exclusion restriction is whether the control group, which faced the flat rate tariff, appears to be responsive to prices instrumented by the short-term supply shocks. Using this sort of "placebo test," I ruled out potential instruments for short-term price variation such as deviations in daily natural gas spot prices from trend and variation in relative temperatures in nearby regions.

Structurally estimating short run and long run elasticities would require data on households' stock of energy using durable goods, but these data are not available for ESPP households.

5. Results

This section presents estimation results in the form of three key conclusions that can be drawn from the data. Some reduced form results will be presented alongside the demand parameter estimates, with an eye toward the connections between the two approaches.

The first conclusion is that program households are statistically significantly price elastic. Specification I of Table 5 presents the demand parameters estimated using Eq. (9). Responses to average price $\hat{\eta}_s^A$ are −17.4 and −21.8 W/(cents/kWh), respectively. These are statistically stronger than the responses to deviations from average prices: the deviation coefficients $\hat{\eta}_h^D$ average -12 W/(cents/kWh). This implies that a one standard deviation increase in an hour's price (in the afternoon, 1.5 cents/kWh) is associated with the equivalent of just under one in four households turning off a 75 W lightbulb. Because there is little price variation in the nighttime hours (from 10 pm to 6 am), these deviation parameters are not estimated. In the next section, we will examine whether these statistically significant price responses are also economically significant, in the sense of generating substantial welfare benefits.

Specification I takes up the first three columns, with the parameter estimates presented in the middle column. These three columns present different adjustments for serially correlated errors, which could be a concern given that consumption is measured by the hour. The leftmost column presents robust standard errors, while the second column uses Newey-West standard errors with one lag, and the third uses Newey-West with up to 24 h lag. Going from no serial correlation to allowing for up to 24 lags changes the standard errors only slightly: the median standard error estimate increases by just under eight percent, although a couple of the parameters appear to be much less precisely estimated.

The second conclusion is that households respond to RTP through energy conservation, with no net load shifting from high to low price hours. Fig. 3 illustrates this in reduced form by showing the mean average treatment effect for each hour of the day, across all non-High Price Alert summer days. The relationship of these hourly average ATEs and the average summer price shape is in essence the variation that identifies the parameters η_s^A and the intercept η_s^T. Average prices increase from 5 cents/kWh at night to 9 cents in the mid-afternoon, and this is associated with lower treatment group consumption by an average of 50–80 W. On average, the afternoon consumption is not being shifted to the nighttime: only between 10 pm and 11 pm is the point estimate for treatment group average consumption higher than in the control, and this increase does not come close to offsetting the conservation that occurs during the rest of the day.

This finding of no increase in off-peak consumption does not necessarily imply that there is zero substitution of electricity consumption across hours. As intra-day substitution parameters could not be estimated, I do not rule out this form of substitution. Furthermore, different combinations of conservation and shifting could be consistent with this aggregate finding on the level and shape of average consumption by hour. For example, households could make some investment that conserves energy in all hours by some constant amount, which if combined with a second change that shifts some consumption from afternoon to nighttime hours would make it appear as though no change had occurred at night. Alternatively, the treatment group could leave nighttime consumption unchanged and conserve more substantially in the afternoons. What can be concluded is that *on net*, RTP causes households to significantly reduce consumption on the average afternoon and does not cause significant increases in average consumption at night. As a result, the treatment group's consumption during the experiment dropped approximately five percent more than the control group's consumption during the experiment relative to pre-program baseline.

Consumers' behavior exhibits this same pattern on High Price Alert days. Fig. 4 illustrates the hourly shape of mean ATEs for the nine summer High Price days, showing that the treatment group reduces consumption by an additional 100–200 W during daytime hours, or about five to 14 percent. Only in 4 h is the treatment group estimated to have increased consumption, all by less than 10 W and statistically indistinguishable from zero. This is the reduced form illustration of the variation that identifies the η_g^{HP} parameters. Other specifications, not reported, show that there is no increase in consumption associated with the evening before a High Price Alert day, or the day after.

The energy conservation result can be explained by the technologies available to households to respond to RTP. In a small survey in the fall of 2003, treatment group households reported the changes

Table 5
Demand system estimates.

		I (T and C)		II (T Only)	III (T Only)	IV (T Only)
Average price η^A: summer		−17.4		40.2	N/A	N/A
	(1.0)**	(1.3)**	(2.0)**	(0.8)**	N/A	N/A
Non-summer		−21.8		119.8	N/A	N/A
	(1.9)**	(2.4)**	(2.1)**	(0.9)**	N/A	N/A
Price deviation η^D: Hour 6		−10.5		35.2	32.6	29.9
	(6.9)	(6.9)	(7.1)	(1.9)**	(1.9)**	(7.8)**
7		−14.8		27.7	22.7	25.2
	(6.3)**	(6.3)**	(6.5)**	(1.9)**	(2.0)**	(7.6)**
8		−16.7		1.7	−3.0	−4.7
	(10.0)*	(10.0)*	(10.4)	(2.8)	(3.1)	(12.5)
9		−15.6		−4.4	−17.0	−12.1
	(7.9)**	(8.0)*	(8.3)*	(2.5)*	(2.8)**	(10.1)
10		−12.2		−3.3	−22.1	−21.4
	(5.4)**	(5.4)**	(5.8)**	(2.0)*	(2.2)**	(7.3)**
11		−10.6		6.7	−22.6	−23.8
	(5.0)**	(5.1)**	(5.5)*	(2.1)**	(2.4)**	(7.0)**
12		−13.1		13.7	−21.5	−19.2
	(4.3)**	(4.3)**	(4.5)**	(1.9)**	(2.2)**	(5.9)**
13		−13.6		12.6	−18.0	−12.6
	(3.4)**	(3.5)**	(3.6)**	(1.6)**	(1.9)**	(4.8)**
14		−12.7		20.9	2.5	9.4
	(3.7)**	(3.8)**	(4.0)**	(1.8)**	(2.2)	(5.4)*
15		−11.6		34.5	12.3	19.0
	(3.9)**	(4.0)**	(4.2)**	(1.9)**	(2.3)**	(5.7)**
16		−9.1		43.7	23.1	24.0
	(3.7)**	(3.6)**	(3.8)**	(1.8)**	(2.2)**	(5.2)**
17		−10.7		64.2	39.7	41.3
	(4.3)**	(4.4)**	(4.6)**	(2.1)**	(2.5)**	(6.2)**
18		−17.8		44.0	38.9	39.3
	(6.2)**	(6.3)**	(6.6)**	(2.6)**	(3.0)**	(8.5)**
19		−5.7		56.5	54.8	44.5
	(7.6)	(7.7)	(8.1)	(2.9)**	(3.4)**	(10.4)**
20		−5.6		42.4	52.5	43.0
	(5.5)	(5.5)	(5.9)	(2.3)**	(2.7)**	(7.6)**
21		−9.9		54.4	69.7	62.7
	(8.5)	(8.6)	(9.2)	(3.1)**	(3.6)**	(11.5)**
High price day η^{HP}: Morning		−62.9		−8.9	1109.4	1388.5
	(13.5)**	(16.5)**	(22.2)**	(8.0)	(82.2)**	(198.5)**
Mid-Day		−115.8		−58.0	1437.0	2093.3
	(18.4)**	(23.1)**	(32.9)**	(11.5)**	(118.3)**	(290.2)**
Afternoon		−133.1		30.8	52.5	147.0
	(23.3)**	(29.5)**	(42.1)**	(14.7)**	(15.6)**	(38.0)**
Evening		−53.2		607.8	277.4	340.1
	(20.9)**	(26.6)**	(37.7)	(13.7)**	(14.2)**	(32.4)**
Constant η^T: summer		−45.3		N/A	N/A	N/A
	(2.6)**	(3.5)**	(6.9)**	N/A	N/A	N/A
Non-summer		−36.6		N/A	N/A	N/A
	(1.4)**	(1.9)**	(5.5)**	N/A	N/A	N/A

Dependent variable: electricity use (W). Price variables are in cents/kWh. Standard errors in parenthesis. *, **Statistically different from zero with 90 and 95 percent confidence, respectively. Specification I: Left column has robust standard errors. Middle column has Newey-West standard errors with up to 1 h lag. Right column has Newey-West standard errors with up to 24 h lag. Specifications II–IV, Newey-West SE's with up to 1 h lag. Regression I: 3.98 million observations. Regressions II–IV: 3.40 million observations.

they thought they had made after entering the program: turning off lights, using fans instead of air conditioners, turning down or replacing air conditioners, and washing clothes during low price hours instead of during the afternoon. Only this last activity involves substitution of electricity consumption from 1 h to another; the others entail substitution toward more energy efficient capital stock or substitution away from household energy services such as comfort.

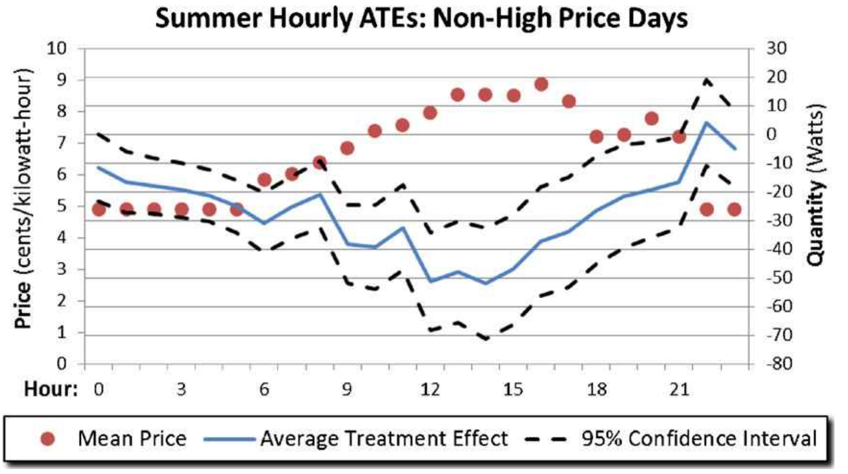

Fig. 3. Average hourly prices and reductions. Newey-West standard errors.

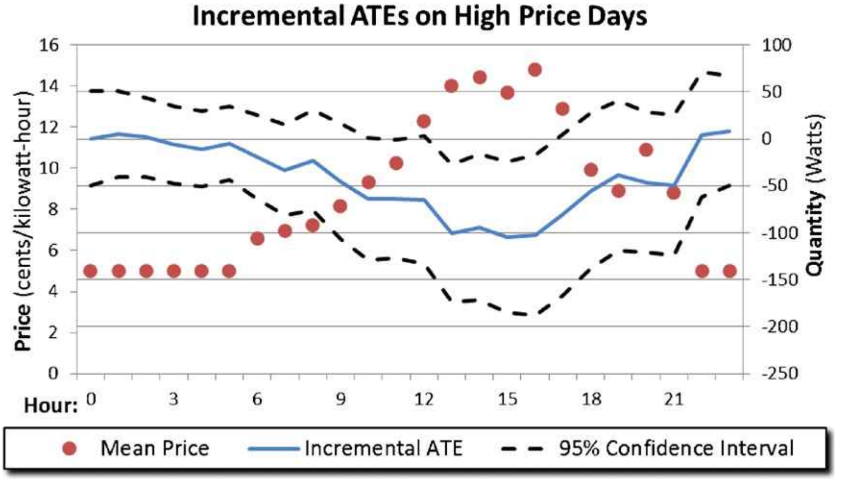

Fig. 4. Prices and incremental reductions on high price days. Newey-West standard errors.

Several factors could change this result in the future. Intra-day substitution elasticities may be increased by energy management devices that automatically allocate activities such as clothes washing to low price hours. Furthermore, potential new sources of electricity demand such as plug-in electric vehicles may also be able to automatically charge when prices are low. Also note that "energy conservation and no net shifting" does not necessarily generalize to industrial facilities, which can respond to RTP through flexibly scheduling production processes across hours. Patrick and Wolak (2001), for example, find statistically significant intra-day substitution parameters for large industrial consumers on RTP.

A third conclusion from the demand results, which is fundamental but perhaps unsurprising, is that energy management and information technology can significantly increase households' price elasticity. Fig. 5 plots the reduced-form ATE for each hour of the Pricelight experiment against the hour's price. For the price range below 10 cents/kWh, the Pricelight does not substantially affect consumption. At 15 cents/kWh, however, the estimated effect of having a Pricelight is about 150 W, and this reaches over 200 W during the highest price hours. This is about two-thirds of the conservation by the RTP treatment group during the highest-price hours of the 2003 experiment,

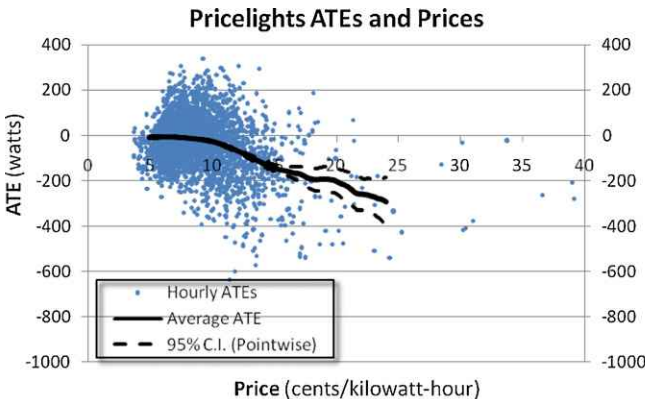

Fig. 5. Pricelights: average treatment effects by price.

although the comparison is made cautiously given that the Pricelights group is self-selected out of the (already self-selected) RTP population.

For context, a short term change that could produce a 200W effect would be if every fifth household turned off a window air conditioner. The Pricelights' importance suggests that because the hourly variance in households' electricity costs induced by real-time pricing is small, devices that lower the cost of price discovery – or simply increase attention paid to electricity use – can substantially affect energy use. Fig. 5 also illustrates that the Pricelight does not cause consumption to increase during low price hours, further reinforcing the net energy conservation result.

An eight month experiment is not ideal evidence on how households would respond to real-time pricing over a period of years. Consumers might learn more about typical price shapes, lose an initial interest in energy conservation, or have time to make additional changes to energy-using capital stock. One way of providing additional evidence on this issue is to compare the treatment and control groups in 2004, when both are on real-time pricing but the treatment group is in its second year and the control group is in its first. For additional statistical power, this exercise departs from the structural formulation and simply estimates a reduced form coefficient on hourly price. Table 6 details the results. The first column is the reduced form analysis of the RTP experiment from May through December 2003. Across all hours, if price was higher by one cent/kWh, treatment group consumption was lower by 18.6W. The latter two columns show that in 2004, the original treatment group consumed six to eight watts less on hours when price was higher by one cent. By this measure, price responsiveness in the second year is about one third greater than in the first year.

5.1. The benefits of randomized experimentation

Randomized experiments are not new to the utilities industry: British utilities used randomized trials to test alternative pricing programs in the late 1960s (Levitt and List, forthcoming), and many of the American time of use pricing experiments from the late 1970s and early 1980s included randomized control groups (Aigner, 1984). Nearly all current energy efficiency programs and many recent real-time pricing programs have been non-experimental,[15] however, despite the fact that the causal effects of these programs are under continuing debate. How useful are non-experimental data in this application?

Lacking a randomized control group or a valid instrument for prices in this setting, one way to address simultaneity bias would be to use observables to soak up the demand shifter ξ and assume that price variation is conditionally exogenous. As a trial of the conditional exogeneity assumption, I parameterize ξ as a function of each hour h's observable characteristics, including a polynomial series

[15] To my knowledge, Herriges et al. (1993) is the only randomized evaluation of a real time pricing program. Many other pricing structures have been evaluated with randomized treatment and control groups, however, and Allcott (2009b) and Davis (2008) evaluate recent experimentally randomized energy efficiency programs.

Table 6
Second year vs. first year of RTP experiment.

	RTP	All	Summer
	(1)	(2)	(3)
P × T	−18.5	−6.1	−4.8
	(0.9)**	(0.7)**	(1.1)**
T × P × Pre-program usage	−14.8	−9.7	−9.1
	(1.5)**	(2.7)**	(2.0)**
T × P × household size	6.7	0.9	−2.1
	(0.8)**	(0.9)	(1.4)
T × P × log(Income)	7.2	1.6	4.7
	(0.7)**	(0.9)*	(1.2)**
P × Pre-program usage	63.7	33.6	51.5
	(1.5)**	(2.7)**	(1.8)**
P × household size	−7.6	2.5	4.0
	(0.8)**	(0.9)**	(1.2)**
P × log(Income)	1.4	7.3	2.2
	(0.7)**	(0.8)**	(1.0)**
Pre-program usage	92.9	330.7	62.0
	(9.5)**	(20.3)**	(12.6)**
Household size	38.2	−20.3	0.5
	(4.9)**	(6.5)**	(8.6)
log(Income)	11.7	−64.3	−14.8
	(4.2)**	(6.0)**	(7.3)**
T × Pre-program usage	1.8	−92.3	72.2
	(9.9)	(20.7)**	(13.6)**
T × household size	6.3	25.9	17.1
	(5.3)	(6.9)**	(9.4)*
T × log(Income)	−69.9	−4.8	−37.9
	(4.6)**	(6.5)	(8.2)**
T	69.2	−20.6	48.5
	(5.9)**	(5.8)**	(7.5)**
Obs.	3,978,019	3,762,972	1,396,648
F statistic	64776.9	29708.4	14750.4

Dependent variable: electricity use (W). Price variables are in cents/kWh. Newey-West standard errors with up to 1 h lag in parenthesis. *, **Statistically different from zero with 90 and 95 percent confidence, respectively. Pre-program usage, household size, and log(Income) are normalized to mean 0, standard deviation 1.

of weather variables and month and workday indicators. I then estimate the demand function from Eq. (3) with only the treatment group data. Specification II (the fourth column) in Table 5 shows the results of this regression: the positive η parameters give the apparently upward-sloping demand indicative of simultaneity bias.

I then omit η^A and include hour dummies, in an attempt to soak up the natural correlation between a typical hourly price shape and households' electricity consumption. As Specification III (the fifth column) of Table 5 shows, this also gives upward-sloping demand functions for most hours. Specification IV repeats this regression for the control group, showing that it apparently was responsive to real-time prices that it did not face. Although the failure of conditional exogeneity in some sense means that the econometrician did not collect enough controls, these results are consistent with a number of other specifications attempted, with all of the control variables an analyst would typically have available.[16].

While the benefits of randomized experiments are well-understood by economists, this exercise highlights the unique opportunity that this experiment provides to understand how consumers respond to real-time electricity prices. This also suggests that a further shift toward randomized evaluations of energy pricing and energy efficiency programs could be useful in allocating policy

[16] Note further that even if an elasticity to price variation conditional on observables could somehow be estimated, that approach would not capture any price response correlated with these observables. For example, recall that Figs. 3 and 4 illustrated substantial energy conservation on average in afternoon hours. Including hour dummy variables in the set of controls would absorb that form of price response into the estimated demand shifter ξ.

Table 7
Welfare calculation and CO_2 effects.

	Summer	Non-summer
Effects by season		
Hours observed	2208	3672
Comparable flat rate tariff (cents/kWh)	7.17	6.15
Fitted baseline usage (W)	986	804
Fitted usage reduction (W)	45.7	37.7
Electricity cost reduction (cents/h)	0.244	0.118
Compensating variation (cents/h)	0.195	0.088
Electricity generation CO_2 emissions (lbs/household-hour)	1.4	1.53
CO_2 emission reductions (lbs/household-hour)	0.0616	0.0674
Annual effects		
Annual compensating variation ($)	10.05	
Baseline annual electricity costs ($)	480	
Annual electricity cost reduction ($)	13.1	
Percent savings	2.73	
CO_2 emission reductions (short tons/household)	0.289	
Baseline annual CO_2 emissions (short tons)	6.54	
Percent CO_2 reduction	4.42	

Baseline annual electricity costs is the control predicted group's electricity bill on the comparable flat rate tariff. Note that this is lower than ComEd's typical annual electricity bill because the Participation Incentive and the mild summer reduce the comparable flat rate tariff. To compute annual effects, the summer and non-summer observations are re-weighted such that summer is 2208 of the 8760 h in a year and non-summer is 6552 of the hours.

attention and investment dollars to the most effective projects. Allcott and Mullainathan (2010) describe one model for how this shift could occur and discuss lessons for energy program design from field experiments in other areas such as development microeconomics and finance.

6. Consumer surplus

The welfare effects of real-time pricing flow through four primary channels: producer profits, retailer profits, meter costs, and consumer welfare.[17] This analysis computes only the consumer surplus effects, under the simplifying assumption that retail prices are fixed. Modeling producer profits and welfare effects in market equilibrium is well beyond the scope of this paper, and Allcott (2009a) focuses on those issues.

The representative consumer's compensating variation is:

$$CV_{hd} = V(\mathbf{P}^{T=1}, w)|(\hat{\eta}, \hat{\xi}, T = 1) - V(\mathbf{P}^{T=0}, w|\hat{\eta}, \hat{\xi}, T = 0) \tag{10}$$

From Eq. (1), each element in the vector $\mathbf{P}^{T=1}$ of RTP treatment group prices was the wholesale price plus a flat payment per kilowatt-hour that covers the retailer's costs, which are essentially fixed. For welfare analysis, we want the flat rate tariff $\mathbf{P}^{T=0}$ that is "comparable" in equilibrium with $\mathbf{P}^{T=1}$, in the sense that it keeps the retailer at zero profits by covering wholesale electricity costs and holding constraint the *ex-post* net revenues from the distribution charge and the ESPP program's "Participation Incentive." In this case, the $\mathbf{P}^{T=0}$ used for welfare analysis is less than the actual price that the Control group received, because the Participation Incentive and mild weather during the pilot program kept $\mathbf{P}^{T=1}$ relatively low. To compute annualized values, the summer and non-summer observations are re-weighted such that summer is 2208 of the 8760 h in a year and non-summer is 6552 of the hours.

This calculation, detailed in Table 7, gives a Compensating Variation of $10 per year per program household. By comparison, advanced metering infrastructure is expected to cost between $100 and $150 per household (Allcott, 2009a), although the infrastructure's benefits other than the potential to

[17] The indirect utility function from Section 3 brings out an additional channel of welfare effects when RTP is optional. Households that select into RTP should theoretically have underlying demand patterns – as captured by demand shifters ξ – that are less correlated with hourly prices than in the rest of the population. This form of positive selection implies that the average cost of electricity for the group that remains on the flat rate tariff will rise. I abstract away from this because my data provide no insight into demand parameters for the rest of the population. See Borenstein (2007b) for an analysis of this issue.

offer RTP are substantial. Although in electricity pricing, congestion pricing, and many other settings, economists' intuition is that prices should be aligned with marginal costs, this residential RTP program may provide an important real-world example of a situation where this is not currently welfare-enhancing given contracting or information costs. It should again be emphasized that because of the selected set of households in this program and the small amount of price variability in this program compared to in other settings, this result may not be highly generalizable.

The pilot program's effects can also be framed as a reduction in energy costs. To do this, I first compute each household's fitted quantity demanded in each hour under RTP and under the "comparable" flat rate tariff, $\hat{q}(\mathbf{P}^{T=1}|(\hat{\eta}, \hat{\hat{\xi}}_{ihd}))$ and $\hat{q}(\mathbf{P}^{T=0}|\hat{\eta}, \hat{\hat{\xi}}_{ihd})$, respectively. Using this predicted consumption, I estimate that the average RTP household saves $13 per year on electricity bills. As shown in Table 7, this is 2.7 percent of total electricity bill at the calculated $\mathbf{P}^{T=0}$. A different way of putting this is that, even in this selected group, a household that did nothing in response to the offer to enroll in the program would have foregone just $13 annually.

This calculation suggests that some program participants may have gone to great lengths to be price elastic, with small pecuniary returns.[18] This result is consistent with some other studies, such as Reiss and White (2008), that show that populations or subpopulations are in some cases highly motivated to respond to energy prices.[19] These studies are a puzzling counterpoint to a larger literature on the "Energy Paradox" (Jaffe and Stavins, 1994), which suggests that consumers often appear to be less responsive to energy costs than theory would predict that they should be.

6.1. Carbon emissions

While real-time pricing does affect air pollution emissions, RTP programs have historically been motivated by aligning electricity prices with marginal costs, with little consideration of climate change or other environmental issues. The energy conservation result from the empirical analysis, however, suggests that RTP programs could also reduce carbon emissions from electricity generation. How relevant is this factor in the ESPP program?

Carbon emission abatement is the product of the treatment effect of RTP with the carbon dioxide emission rate of the marginal electricity generator. Intuitively, if the marginal power plant during high price hours, when the ATEs are largest in magnitude, has a higher emission rate than the marginal power plant when prices are low, RTP will have a more beneficial impact on emissions. Electricity generators of different technologies, including coal, nuclear, hydro, natural gas, oil, wind, and others, have different carbon emission rates, and the market shares and generation profiles of each technology vary between electricity markets in different regions. Results in a nationwide analysis, such as Holland and Mansur (2008), may therefore be different from those for this specific program.

I construct a simple model[20] that gives the marginal carbon emission rate $\hat{\mathrm{E}}_{hd}(P_{hd})$ as a function of the price that would be set by the marginal generator in the ComEd region in 2003. For each hour of the

[18] One particularly motivated program participant told me that on the afternoons of High Price Alert days, they reduce energy consumption by leaving the house and taking their family to the park. I asked them how much they saved by doing that on a High Price Alert day relative to a normal day; they estimated 25 cents. Even that was an overestimate.

[19] Reiss and White (2008) examine electricity consumption by San Diego households during the California electricity crisis. They find that total consumption dropped 13 percent in two months in response to an average price increase from 10 to 23 cents/kWh, then rebounded immediately by 8 percent when prices dropped to 13 cents/kWh. After the initial increase, one in three households reduced by 20 percent or more relative to the previous year. The authors show that this large fraction of the population would have needed to make significant behavioral or capital stock changes in order to achieve reductions of this magnitude.

[20] The model is a simple merit-order dispatch model constructed from Continuous Emissions Monitoring System data from the US Environmental Protection Agency (2009) for the Reliability First region of the North American Electric Reliability Council, of which ComEd is a member. These data give unit-level fuel input, electricity output, and carbon dioxide emissions, which are then combined with 2003 average prices for coal, oil, and natural gas from the US Energy Information Administration (2009a,b). The marginal generator at a particular price is the next unit in a merit order determined by the sum of fuel costs and other non-fuel variable operating costs.

This model correctly represents marginal CO_2 emissions under either of two sets of assumptions. First, it would be valid assuming fixed capital stock and no ramping constraints or other dynamic considerations. Second, it would be valid in the long run where the generation technology at a particular price is constant. Both sets of assumptions also require marginal cost bidding and abstract away from inter-regional electricity transfers.

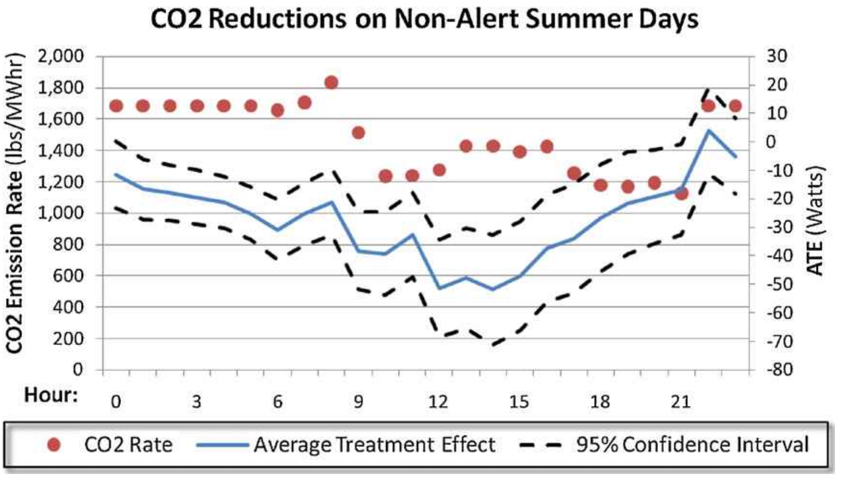

Fig. 6. Carbon dioxide reductions. Newey-West standard errors.

ESPP experiment, this model provides the emission rate of the marginal generator, which can be multiplied by the predicted effect of RTP on quantity demanded at "comparable" retail prices. The average effect per household on over any period of days d and hours h is given by the following sum, where $\hat{q}$ is given by Eq. (3):

$$\sum_{d}\sum_{h=1}^{24}\hat{E}_{hd}(P_{hd})\cdot\frac{1}{N}\sum_{i=1}^{N}\left(\hat{q}(\mathbf{P}^{T=1}|(\hat{\eta},\hat{\xi}_{ihd}))-\hat{q}(\mathbf{P}^{T=0}|(\hat{\eta},\hat{\xi}_{ihd}))\right) \tag{11}$$

In this region, the marginal generator is typically small high-emission coal-fired plants in the off-peak periods and lower-emission natural gas plants during the peak periods. As shown in Fig. 6, this means that the energy conservation from RTP covaries negatively with the marginal emission rate, slightly attenuating RTP's effects on carbon dioxide emissions. As shown in Table 7, the predicted annual carbon reduction from the program is approximately 0.29 short tons per household, which is 4.4 percent of total emissions from the average program household's electricity consumption. This is slightly less than the predicted percent reduction in electricity demanded under RTP. What fundamentally drives this result is the daytime energy conservation parameterized by η_s^T, which is visible in Fig. 3.

To convert these into monetary terms for comparison with the welfare calculation requires placing a value on carbon emissions. Because marginal damages are difficult to estimate, I instead use the predicted price of a carbon emission allowance in 2020 under the recently proposed Waxman-Markey carbon cap-and-trade bill, which is \$29 per short ton (US Energy Information Administration, 2009c). Multiplied by the annual carbon emission reductions, this gives \$8.41 per year. In the specific case of an RTP program that generates net energy conservation, in an electric power market where the marginal off-peak generator is not substantially more carbon-intensive than the marginal peak-hour generator, the gains from reduced carbon emissions could play an important role in the welfare calculation.[21]

7. Conclusion

This paper exploits a randomized field experiment to evaluate the first ever hourly real-time electricity pricing program for residential consumers, Chicago's Energy-Smart Pricing Plan. A central

[21] If and when a carbon cap-and-trade bill passes, this calculation would be different because the costs of carbon emissions would in principle be internalized into the electricity price. At that point, we would simply compute the welfare gains from RTP at the new equilibrium electricity prices, and the marginal damages of carbon dioxide emissions would not enter the calculation separately.

result of the demand estimates is that residential RTP should perhaps be thought of as a peak energy conservation program, instead of a mechanism to shift consumption from peak to off-peak. This means that demand for off-peak power under residential RTP, and thus air pollution emissions, energy costs, and equilibrium entry of baseload power plants, may be lower than analysts had previously expected. A second central result highlights the importance of complementary technologies, in this case the "Pricelight," which help residential consumers respond to time-varying prices.

For several reasons, the precise magnitude of the elasticity and welfare results from this experiment may not generalize to other settings and time periods. First, in this program were a selected group who are potentially more price elastic than the general population. Second, future advances in household Smart Grid technologies could increase price elasticity and therefore directly change this result. Third, the variance in hourly electricity prices during this pilot is lower than that experienced in other regions of the country and in different years. This is important both because higher price variance increases the potential gains from RTP given any particular price elasticity and because higher volatility could induce consumers to rationally devote more attention to prices, which would increase their price elasticity.

The welfare calculation with these parameter estimates shows that moving to RTP from the flat rate tariff improved consumer surplus for program households by an annualized value of $10. This comprises 1-2 percent of households' total electricity expenditures. These benefits do not appear substantially outweigh the gross costs of advanced metering infrastructure required to observe hourly consumption, which are estimated to be around $150 per household. Many utilities are installing these meters and other retail Smart Grid infrastructure, however, because they offer substantial cost savings. The option to offer RTP is thus for many utilities an additional potential benefit of advanced metering infrastructure.

Acknowledgments

I thank, without implicating, Alberto Abadie, Severin Borenstein, Ken Gillingham, Michael Greenstone, Bill Hogan, Erich Muehlegger, Larry Katz, Erin Mansur, Sendhil Mullainathan, Paul Niehaus, Chris Nosko, Ariel Pakes, Dave Rapson, Rob Stavins, Frank Wolak, and seminar participants at Duke, Georgetown, Harvard, IBS Hyderabad, Michigan, NYU Department of Economics, NYU Stern, Notre Dame, Resources for the Future, Stanford, Tufts, and UC Davis. Thanks also to the Editor and to two anonymous referees. Anne Evens, Marjorie Isaacson, Larry Kotewa, and Anthony Star of the Center for Neighborhood Technology facilitated my access to the household electricity demand data. Financial support is acknowledged from the Harvard University Center for the Environment and from the Joseph Crump Fellowship, which is granted by the Harvard Environmental and Natural Resources Program and the Harvard Center for Business and Government.

References

Aigner, D., 1984. The welfare econometrics of peak-load pricing for electricity. Journal of Econometrics 26 (1–2), 1–15.

Allcott, H., 2009a. The Smart Grid, Imperfect Competition, and Entry in Electricity Supply. Working Paper, Massachusetts Institute of Technology (September).

Allcott, H., 2009b. Social Norms and Energy Conservation. Working Paper, Massachusetts Institute of Technology (July).

Allcott, H., Mullainathan, S., 2010. Behavior and energy policy. Science 327 (March (5970)).

Bandt, W., Campbell, T., Danner, C., Demsetz, H., Faruqui, A., Kleindorfer, P., Lawrence, R., Levine, D., McLeod, P., Michaels, R., Oren, S., Ratliff, J., Riley, J., Rumelt, R., Smith, V., Spiller, P., Sweeney, J., Teece, D., Verleger, P., Wilk, M., Williamson, O., 2003. 2003 Manifesto on the California Electricity Crisis. http://www.aei-brookings.org/publications/abstract.php?pid=341. May 2003.

Barbose, G., Goldman, C., Neenan, B., 2004. A Survey of Utility Experience with real-time pricing. Working Paper, Lawrence Berkeley National Laboratory, December.

Boisvert, R.N., Cappers, P., Goldman, C., Neenan, B., Hopper, N., 2007. Customer response to RTP in competitive markets: A Study of Niagara Mohawk's Standard Offer Tariff. The Energy Journal 28 (January (1)), 53–74.

Borenstein, S., 2002. The trouble with electricity markets: understanding California's restructuring disaster. Journal of Economic Perspectives 16 (Winter (1)), 191–211.

Borenstein, S., 2005. The long-run efficiency of real-time electricity pricing. The Energy Journal 26 (April (3)), 93–116.

Borenstein, S., 2007a. Customer risk from real-time retail electricity pricing: bill volatility and hedgeability. The Energy Journal 28 (April (2)), 111–130.

Borenstein, S., 2007b. Wealth transfers among large customers from implementing real-time retail electricity pricing. The Energy Journal 28 (April (2)), 131–149.
Borenstein, S., Bushnell, J., 1999. An empirical analysis of the potential for market power in California's electricity market. Journal of Industrial Economics 47 (September (3)), 285–323.
Borenstein, S., Holland, S., 2005. On the efficiency of competitive electricity markets with time-invariant retail prices. RAND Journal of Economics 36 (3), 469–493.
Davis, L., 2008. Durable goods and residential demand for energy and water: evidence from a field trial. RAND Journal of Economics 39 (Summer (2)), 530–546.
Faruqui, A., Sergici, S., 2008. The Power of Experimentation: New Evidence on Residential Demand Response. Working Paper, Brattle Group (May).
Faruqui, A., Sergici, S., 2009. Household Response to Dynamic Pricing of Electricity: A Survey of the Experimental Evidence. Working Paper, Brattle Group (January).
Herriges, J.A., Mostafa Baladi, S., Caves, D.W., Neenan, B.F., 1993. The response of industrial customers to electric rates based upon dynamic marginal costs. Review of Economics and Statistics 75 (August (3)), 446–454.
Holland, S., Mansur, E., 2006. The short-run effects of time-varying prices in competitive electricity markets. The Energy Journal 27 (October (4)), 127–155.
Holland, S.P., Mansur, E., 2008. Is real-time pricing green? The environmental impacts of electricity demand variance. Review of Economics and Statistics 90 (August (3)), 550–561.
Houthakker, H., 1951. Electricity tariffs in theory and practice. The Economic Journal 61 (March (241)), 1–25.
Isaacson, M., Kotewa, L., Ozog, M., Star, A., 2006. Changing How People Think About Energy. Working Paper, ACEEE Summer Study on Energy Efficiency in Buildings.
Jaffe, A., Stavins, R., 1994. The energy paradox and the diffusion of conservation technology. Resource and Energy Economics 16, 91–122.
Joskow, P., Tirole, J., 2006. Retail electricity competition. RAND Journal of Economics 37 (Winter (4)), 799–815.
Joskow, P., Tirole, J., 2007. Reliability and competitive electricity markets. RAND Journal of Economics 38 (Spring (1)), 60–84.
Levitt, S.D., List, J.A., Field experiments in economics: the past, the present, and the future. European Economic Review, forthcoming, http://www.fieldexperiments.com/uploads/127.pdf.
Patrick, R., Wolak, F., 2001. Estimating the Customer-Level Demand for Electricity Under Real-Time Market Prices. NBER Working Paper No. 8213, April.
Reiss, P., White, M., 2005. Household electricity demand, revisited. Review of Economic Studies 72 (July (3)), 853–883.
Reiss, P., White, M., 2008. What changes energy consumption? Prices and public pressure. RAND Journal of Economics 39 (Autumn (3)), 636–663.
Rubin, D., 1974. Estimating causal effects of treatments in randomized and non-randomized studies. Journal of Educational Psychology 66, 688–701.
Steiner, P., 1957. Peak loads and efficient pricing. Quarterly Journal of Economics 71 (November (4)), 585–610.
Summit Blue Consulting, 2004. Evaluation of the 2003 Energy Smart Pricing Plan. http://www.cnt.org/repository/ESPP-2003-Evaluation.pdf.
Taylor, T.N., Schwartz, P.M., Cochell, J.E., 2005. 24/7 hourly response to electricity real-time pricing with up to eight summers of experience. Journal of Regulatory Economics 27 (May (3)), 235–262.
Train, K., Mehrez, G., 1994. Optional time-of-use prices for electricity: econometric analysis of surplus and pareto impacts. RAND Journal of Economics 25 (Summer (2)), 263–283.
US Department of Energy, 2008. Room Air Conditioners. http://www.energystar.gov.
US Energy Information Administration, 2005. Residential Energy Consumption Survey. http://www.eia.doe.gov/emeu/recs/recs2001/publicuse2001.html.
US Energy Information Administration, 2008. Annual Energy Outlook. http://www.eia.doe.gov/oiaf/aeo/.
US Energy Information Administration, 2009a. Natural Gas Prices. http://tonto.eia.doe.gov/dnav/ng/ng_pri_sum_dcu_nus_a.htm.
US Energy Information Administration, 2009b. Receipts, Average Cost and Quality of Fossil Fuels. http://www.eia.doe.gov/cneaf/electricity/epm/table4_1.html.
US Energy Information Administration, 2009c. Energy Market and Economic Impacts of H.R. 2454, the American Clean Energy and Security Act of 2009. http://www.eia.doe.gov/oiaf/servicerpt/hr2454/execsummary.html.
US Environmental Protection Agency, 2009. Clean Air Markets Data and Maps. http://camddataandmaps.epa.gov/gdm/index.cfm.
Williamson, O., 1966. Peak load pricing and optimal capacity under indivisibility constraints. American Economic Review 56 (September (4)), 810–827.
Wolak, F., 2006. Residential Customer Response to Real-time Pricing: The Anaheim Critical Peak Pricing Experiment. Center for the Study of Energy Markets Working Paper 151, May.
Wolak, F., 2007. Managing Demand-Side Economic and Political Constraints on Electricity Industry Restructuring Processes. Working paper, Stanford University, June.

Further reading

Aubin, C., Fougere, D., Husson, E., Ivaldi, M., 1995. Real-time pricing of electricity for residential customers: econometric analysis of an experiment. Journal of Applied Econometrics 10 (December), S171–S191.
Borenstein, S., Jaske, M., Rosenfeld, A.H., 2002. Dynamic Pricing, Advanced Metering, and Demand Response in Electricity Markets. Center for the Study of Energy Markets Working Paper 105, October.
Ham, J., Mountain, D., Chan, L., 1997. Time-of-use prices and electricity demand: allowing for selection bias in experimental data. RAND Journal of Economics 28, S113–S141.
Wilson, R., 2002. Architecture of power markets. Econometrica 70 (July (4)), 1299–1340.