

NBER WORKING PAPER SERIES

DEMAND AND PRICING IN ELECTRICITY MARKETS:
EVIDENCE FROM SAN DIEGO
DURING CALIFORNIA'S ENERGY CRISIS

Peter C. Reiss
Matthew W. White

Working Paper 9986
<http://www.nber.org/papers/w9986>

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
September 2003

We thank San Diego Gas and Electric Company for providing access to non-public billing data for this research, and staff at the California Public Utilities Commission for assistance and comments. White gratefully acknowledges financial support from the Russell Ackoff Endowment Fund in the Wharton Risk Management and Decision Processes Center. The opinions and conclusions expressed in this paper are solely those of the authors and should not be construed as representing the opinions or policies of the California Public Utilities Commission or the San Diego Gas and Electric Company. The views expressed herein are those of the authors and are not necessarily those of the National Bureau of Economic Research.

©2003 by Peter C. Reiss and Matthew W. White. All rights reserved. Short sections of text, not to exceed two paragraphs, may be quoted without explicit permission provided that full credit, including © notice, is given to the source.

Demand and Pricing in Electricity Markets: Evidence from San Diego During California's Energy Crisis

Peter C. Reiss and Matthew W. White

NBER Working Paper No. 9986

September 2003

JEL No. D1, L5, L9

ABSTRACT

We study the electricity consumption of San Diego-area households following a series of price changes and related events during California's energy crisis in 2000-01. The analysis uses a five-year panel of disaggregate billing and weather data for a random sample of 70,000 households. In contrast to prior work, these data allow us to proceed without behavioral assumptions regarding a consumer's knowledge of energy prices.

We find that after a rapid price increase in summer 2000, consumption fell substantially over about 60 days, averaging 12% per household; consumption then rebounded to within 3% of pre-crisis levels after a price cap was imposed. Under the price cap, public appeals for energy conservation and a remunerative voluntary conservation program had significant, but transitory, effects. Further, a large share of households reduced electricity consumption substantially (over 10%) but saved small monetary amounts (\$10 or less). Overall, the results indicate consumers may be far more responsive to pecuniary and non-pecuniary incentives for altering their energy use than is commonly believed.

Peter C. Reiss
Graduate School of Business
Stanford University
518 Memorial Way
Stanford, CA 94305-5015
and NBER
preiss@stanford.edu

Matthew W. White
The Wharton School
University of Pennsylvania
3620 Locust Walk
Philadelphia, PA 19104-6372
and NBER
mawhite@wharton.upenn.edu

1 Introduction

Households in the United States are billed for their electricity use monthly, based on tariffs that change much less frequently. This practice leaves utilities with few options that they can use to moderate demand during times of tight supply. Traditionally, the industry has dealt with this lack of flexibility by ensuring that utilities have sufficient capacity or import capability to meet all but extreme demands.

In 1998, California implemented an electricity industry restructuring plan that loosened regulatory control over wholesale electricity prices and capacity decisions. Beginning in June 2000, a variety of factors led to a lengthy period of unexpectedly tight supplies. Because most consumers' electric rates remained fixed at or near pre-restructuring levels, the state and utilities had few short-run options they could employ to reduce demand. This inflexibility of retail electricity prices, combined with escalating wholesale costs, pushed several of the state's utilities to the brink of bankruptcy and left a costly financial burden on the state.

In the wake of the California crisis and problems elsewhere, numerous proposals have been put forward to reform the basic structure of retail electricity pricing. Many of these proposals envision making consumer prices much more responsive to supply conditions.¹ Unfortunately, there is limited evidence available on how households might respond in aggregate or individually to the types of price changes seen during the California crisis. Moreover, it remains an open question how quickly—or even whether—California's households are willing and able to reduce their electricity consumption over short-term horizons.

This paper uses a unique household-level data set to assess consumer reactions to tight supply conditions both when prices freely vary and when prices are capped. Our analyses speak not only to energy policy issues, but also more broadly to economic questions about whether and how consumers respond to price versus non-price incentives. Specifically, we examine the responses of San Diego-area households to an unprecedented rise and fall in electricity prices, and their response to voluntary energy conservation incentives when residential prices were subsequently capped. We use San Diego-area households because at the start of the state's electricity crisis in June 2000, a provision of California law allowed the utility serving the San Diego region to raise its electric rates commensurate with escalating wholesale prices. The

¹See, for example, Lafferty et al. (2001), Borenstein (2002), or CPUC (2002).

subsequent retail price increases in San Diego led to an emergency legislative ‘fix’ of price caps, which remained in effect as the state’s crisis deepened.

Our analysis also complements prior electricity demand studies. Despite several decades of research, economists’ understanding of how households respond to energy price changes—both in speed and magnitude—remains far from complete. This shortcoming persists for electricity because regulatory authorities tightly control price changes. Thus, discerning the effect of modest price ‘signals’ amid the wide variation in demand from other sources (such as weather or geography) has long posed a difficult problem for econometric studies, regardless of technique.² The experience of San Diego is valuable because it overcomes this canonical econometric problem in two unique ways: The industry briefly vacated its traditional checks on consumers’ prices and consumers experienced unprecedented price changes.³

A second feature of the paper examines two post-price cap efforts to curb retail consumer demand. The first of these was an extensive media campaign by the state of California to promote voluntary energy conservation during the state-wide crisis in the spring of 2001. The second offered households financial rewards for reducing electricity use during the summer of 2001. We find that San Diego-area household electricity consumption fell substantially during the spring of 2001, the period of extensive public appeals for energy conservation and media attention to the escalating crisis. This occurred even though there were no significant changes in households’ economic incentives to reduce electricity use. The subsequent remunerative voluntary conservation program in the summer of 2001 further reduced demand, and by magnitudes roughly proportionate to households’ responses to the price spike in 2000. Nevertheless, both of these programs effects were transitory, ending promptly after

²Some progress on this problem has been made through formal pricing experiments. See especially Caves and Christensen (1980), Parks and Weitzel (1984), Ham, Mountain and Luke (1997), or CPUC (2002) and references therein. Unfortunately, truly randomized pricing experiments are rare in this industry. Most rely upon voluntary participation and small samples, raising familiar attrition- and selection-bias concerns that impair generalization to the broader population. A recent effort to address the latter is Ham et al. (1997).

³Ours is not the only examination of San Diegans’ experience. Tabulations in Goldman et al. (2002) and contemporaneous work by Bushnell and Mansur (2003) examine aggregate data, reporting smaller results than our study. This is possibly due to compositional effects from different residential and industrial-consumer price changes and demand responses, which are inseparable in the aggregate data. Smaller effects in aggregate data might also be partly attributable to the difficulty of controlling measurement error in aggregate demand models using standard macroeconomic controls. Survey work by Lutzenhiser (2002) provides complementary information on San Diegans’ energy consumption behavior during this time, which we address further below.

the programs did.

Our analysis is based on a panel of disaggregate utility billing data for a random sample of approximately 70,000 households in or near San Diego, California. The data contain among other things each household's electricity consumption and expenditures from 1997 to 2002. To disentangle the confounding effects of weather, we matched each sample household to daily weather data from one of twenty-one weather stations in this region. Our empirical methods are based on household-specific difference-in-difference comparisons. These compare within-household changes in average daily electricity consumption during a $2\frac{1}{2}$ -year pre-crisis period to consumption changes during and after price changes (or related events) occurred. This within-household analysis allows us to identify, net of weather-related effects, the variation in consumption behavior that followed changes in prices and other programs during California's electricity crisis. In contrast to prior work, these billing data allow us to proceed without any specific assumptions about whether consumers knew marginal energy prices or whether the household monitors its energy consumption.

The next section provides background information and describes how retail prices changed. Sections 3 and 4 summarize the data and our methods, respectively. The main results are presented and discussed in Sections 5 and 6. A brief conclusion closes.

2 Background: Electricity Pricing in San Diego

In 1996, California enacted an electric utility restructuring law that initiated major changes to the industry's organization and regulation. This section describes the changes in electricity prices and related events that directly affected San Diego households during our study period. Readers desiring a more detailed treatment of California's electric markets may wish to consult recent and complementary articles by Blumstein et al. (2002), Joskow (2001), Wilson (2002), or Wolak (2003).

The one million households in the greater San Diego metropolitan area receive electric service from a single utility, the San Diego Gas and Electric Company (SDG&E). Figure 1 shows the average electricity price paid by these households from 1998 to 2002.⁴ There is little variation in this price series across households, except that

⁴SDG&E's residential tariff had an increasing block structure through 2001, with a small price

low-income (poverty-level) households are eligible for slightly subsidized rates. The dashed line in the figure shows the average price that households *would have paid* if a (binding) price cap had not been imposed in September 2000. From the figure, we can distinguish three pricing regimes: (1) a period of basically stable prices, from January 1998 to May 2000; (2) the extraordinary price spike of summer 2000; and (3) the period under the price cap beginning September 2000. These changes, along with two policy interventions under the price cap, are the subjects of our analysis.

The Stable Price Period

Until July 1999, electricity prices in San Diego were regulated in much the same way that they had been for decades. Price changes were few and far between; other than a slight downward adjustment in January 1998 (visible at the far left in Figure 1), average residential electricity prices were essentially flat during the 1990's. In July 1999, however, the regulatory system determining these prices changed substantially. Rather than set prices through periodic regulatory hearings, the tariffs determining retail electricity prices were formally indexed to the price of power on regional wholesale markets. Thus, if wholesale market prices increased from one week to the next, this increase would appear in consumers' bills that month.

The shift to this system of indexing residential electricity prices was part of California's long-term vision for reorganizing its electricity industry. The anticipated effect was that competition in wholesale markets would, over time, drive down retail prices below their previously-regulated levels (see White (1996) or Joskow (2001)). San Diego Gas & Electric was the first of California's utilities to implement this "floating prices" system; the state's other regulated utilities were slated to do so later, had the electricity crisis not intervened.

During its first year, the change to this new system of electricity pricing was almost entirely transparent to the residential consumer. Although the wholesale price index was printed on households' monthly statements, its use resulted in little change in electric bills. As Figure 1 shows, average prices continued at basically the same level

difference (two cents per KWh) between the two blocks. The average price series in Figure 1 is the total revenue per KWh sold under this tariff, based on a large random sample of bills (described further below). Our average price calculations include franchise fees (which are city-specific) and excise taxes (which are small), but exclude special rebates (discussed below). There is no state sales tax on electricity in California. The nonlinearity of the residential tariff accounts for the slight wiggles in the average price series in the figure.

over the year starting July 1999 as the year prior, except for two minor ‘blips’ at the peaks of the 1999 summer cooling and winter heating seasons. The 2½-year period from January 1998 to May 2000 was thus one of stable electricity prices, although the regulatory and market system determining these prices changed dramatically in the interim.

The Price Spike and Cap

In June 2000, California and western U.S. wholesale electricity markets experienced an unprecedented run-up in prices.⁵ Because San Diego residential electricity prices were now indexed to the wholesale market, these price increases promptly appeared in households’ monthly bills. Figure 1 shows that San Diego households’ average prices increased from their historical average of approximately 10 cents per kilowatt-hour (KWh) to over 23 cents per KWh in a span of about three months. The magnitude and speed of this price increase substantially exceeds anything previously experienced in the industry’s history.

Because electricity price changes under this system were not announced (or even known) a month in advance, they comprised a true ‘price shock’ when consumers received their June 2000 bills. Further increases in prices over the next two months (see again Figure 1) generated a storm of public protest and, ultimately, appeals for legislative intervention. By the end of August, the California state legislature responded by imposing a retail price cap on residential and small-commercial electricity rates in the San Diego region. The cap effectively limited residential prices to 13.5 cents per KWh, and has been binding ever since.

Although the price cap first appeared in consumer’s bills in September 2000, it was legislatively made retroactive to June 1, 2000. To carry this out, in the fall of 2000 SDG&E gave each of its residential customers a billing credit equal to the difference between what the household paid for electricity from June through August and what it would have paid under the (retroactively imposed) price cap. These credits appeared on households’ bills in October and November.

There are no systematic data describing San Diegans’ expectations about electric-

⁵The causes and consequences of the wholesale price run-up are many, including increases in production costs (natural gas and air permits, chiefly), flaws in California’s wholesale market design, and the behavior of suppliers. For a more complete discussion, see, e.g., Borenstein (2002), Borenstein, Bushnell, and Wolak (2002), Joskow and Kahn (2002), or Wolak (2003).

ity prices during this period. We do, however, know a great deal about the information available to them. There is little reason to believe that a typical consumer anticipated the June 2000 price spike. But by the end of June 2000, a household attentive to the media's coverage of this market would have been aware that the wholesale prices affecting monthly bills were increasing steadily. By the end of July, the increase in prices should have been apparent to most consumers.

The possibility of a legislative price cap appears to have first received media coverage in the second and third week of August 2000; its enactment in the final week of August was highly-publicized. Consequently, an attentive consumer might have been aware of the price spike as early as June, and the likelihood of its rescission by mid- to late August. An inattentive consumer, by contrast, might not have noticed the electricity price increase until opening its electric bill in late June or early July—and not noticed the price cap's effect until receiving its electric bill for September.

Post-Cap Interventions

Although the price cap expediently dissipated consumer uproar over the price spike, maintaining the cap created other problems. Chief among these was the fact that prices on regional wholesale markets, over which California policy-makers had no authority, now exceeded SDG&E's retail prices by a considerable margin (see again Figure 1). This same situation affected the state's other regulated utilities, who had not yet moved to the floating prices system and sold electricity at regulated rates similar to San Diego's capped price. Essentially, the state's utilities were buying high and obligated to sell low, losing money with each kilowatt-hour consumed. This unstable situation quickly pushed the utilities toward insolvency, turning a financial crisis into one that threatened the state with shortages and electricity rationing (via periodic rotating power outages) by the winter of 2000-01.

Two ways to prevent such rationing are to increase prices, or to lower consumption by other means. Rather than increase retail prices and face the same political uproar that occurred in San Diego the previous summer, state leaders threw considerable resources into promoting energy conservation. For households, these conservation efforts took two forms. First, during the winter and spring of 2001, state agencies and utilities undertook a large media campaign to urge energy conservation and educate people about how to reduce energy use. In essence, this amounted to a massive public appeal for voluntary energy conservation. These public appeals, and the deepening

electricity crisis, received prominent media coverage throughout the spring of 2001.

The second aspect of these efforts was an electricity conservation rebate program. Faced with the possibility of widespread shortages ('rolling blackouts') during the summer of 2001, the state adopted a program of financially rewarding consumers for reducing electricity use. Specifically, beginning in June 2001, residential customers were awarded rebates if their electricity use fell sufficiently from their previous year's level. In the San Diego region, this program's parameters were set so a household received a 20% credit against its monthly electric bill if the household's electricity use fell 15% or more from the same month of the prior year. No adjustments were made for weather, prior conservation efforts, or other factors that might affect a household's electricity use. The program was implemented through utilities' billing systems using households' account histories.

California's effort to alter individual consumption behavior through these programs has provoked considerable attention and skepticism regarding their effectiveness. On the one hand, strategies that successfully reduce energy use without directly taxing it (or otherwise raising price) are of great interest to energy and environmental policy makers. Yet it is an open question whether promoting (voluntary) electricity conservation has any effect if consumers' economic incentives to modify their behavior remain unchanged. This is the situation San Diego households faced during the first six months of 2001.

By contrast, the price-rebate program that commenced in June 2001 changed a household's economic incentives for electricity use in a novel and significant way. While the price spike of summer 2000 and price-rebate program of summer 2001 had different informational and economic structures, they provide complementary evidence on households' willingness to reduce electricity consumption given pecuniary incentives to do so.

In sum, our empirical agenda is to assess how households' electricity consumption responded to four successive events: the spike in consumers' electricity prices during the summer of 2000, the price cap imposed in early fall 2000, the state's public appeals for electricity conservation and concurrent crisis media attention during the winter and spring of 2001, and the electricity-conservation rebate program offered in summer 2001.

3 Data

We examine these events using the *Household Electricity Research Billing Sample* (HERBS), a large panel of San Diego Gas and Electric Company residential utility bills covering the time periods we study. Each observation in the HERBS is a consumer’s monthly bill, including total electricity consumption, exact billing-period dates, the total electricity bill, line-item charges, taxes, any special discounts or rebates, and so on. The data also include information on the location of the residence served, tariff schedule parameters, the billing cohort (described below), and the total bill the household would have paid absent the price cap. We constructed the HERBS with permission of SDG&E from their billing system data archives.

Our sampling design selected 70,000 accounts, at random, from the population of all residential accounts served by SDG&E in March 2001. The number of accounts in this population is the same as the number of households in San Diego County and we therefore use the two terms synonymously. We follow each sample household’s account history backward $3\frac{1}{2}$ years to October 1997 (or to the date it established service, if later) and forward one year to April 2002 (or to the date it ended service, if earlier). Slightly less than 1,000 households enter the HERBS each month prior to March 2001, and a similar number exit each month after March 2001.

This sample entry and exit raises the possibility of bias, if household electricity consumption is associated with mobility. Table 1 provides information on this issue. Households in the sample since October 1997 (row (1)) have higher consumption levels, and lower consumption growth rates, than households that enter the sample later (rows (2) and (3)). These differences are not surprising, inasmuch as households who have moved recently tend to inhabit smaller residences (e.g., apartments) and may spend several years acquiring new appliances. A similar pattern is evident with respect to households exiting the sample, although the difference in consumption growth rates between attritants (row (4)) and non-attritants (row (5)) is negligible.

We note these differences between ‘movers’ and ‘stayers’ for two reasons. First, the fact that entry/attrition may be non-random with respect to electricity use implies that raw statistics from the HERBS for periods far from March 2001 may be unrepresentative of the general population. For example, unweighted averages of sample households’ consumption growth rates are likely to be biased downward as one moves away from March 2001, due to the progressive omission of ‘movers’ who tend to have

higher growth rates (see columns 4 and 5 in Table 1).

The second problem is that we cannot estimate consumption responses to crisis events for late-stage entrants to the sample. This group is identified in row (3) of Table 1. In essence, the statistical models we employ require sufficient consumption observations during the stable price period to identify how a household’s consumption subsequently changed during the crisis period. This leaves us with approximately 46,800 sample households for whom we can measure individual consumption responses to the electricity crisis.

These sample-entry/attrition issues are a concern for our analysis insofar as they create bias when we aggregate consumption responses estimated for individual households to obtain population-level figures. Appendix A describes how we deal with this problem using a sample weighting procedure. These weights are calculated using statistical methods developed to handle missing data (see, for example, Little and Rubin (2002)). In the results that follow, we adjust our estimates (unless indicated otherwise) to account for sample entry/attrition so as to provide better estimates for the population of all San Diego MSA households. In the end, however, the results we report below are largely robust to how we adjust for entry/attrition in the HERBS.⁶ The principal reason for this is that our results are based on within-household consumption changes, and the differences between ‘movers’ and ‘stayers’ in this metric both before and after the crisis (see columns 4 and 5 of Table 1) are small in comparison to typical household consumption responses we estimate to electricity crisis events.

Timing Issues and Cohorts

One feature of the data to clarify is the timing of bills, which we exploit to understand how quickly consumers reacted to price changes and other events. SDG&E uses a cohort-based (or ‘staggered’) billing system, wherein each household is assigned to one of 21 cohorts. Each weekday, one cohort’s meters are read, their current billing period closes, and their next billing period begins. The following weekday the same events occur for the next cohort, and so on. The “monthly” bill for a household in the 10th cohort, for instance, covers a period ending within two days of the 14th of

⁶This was not our initial expectation, however. Our concerns about more serious entry/attrition biases motivated our considerable efforts to address with problem, as detailed in Appendix A.

the month, with the bill arriving a few days later.⁷

These timing issues affect the way we handle the data in two important ways. First, to track changes in consumption through time, we group and compare the consumption data by billing cohort. This enables us to track almost daily changes in consumption behavior. Second, because of differences in the length of monthly billing periods, we compare households' bills based on their average daily consumption (the household's total consumption for the billing period (in kilowatt-hours) divided by the number of days in the billing period).⁸

Weather Data

One limitation of utility billing data is that they do not include household-specific weather information. In practice, this might not be much of a problem if most households experienced the same weather. In San Diego's service territory, however, the weather can and does differ dramatically across households' locations. This territory includes temperate coastal areas, near inland areas with considerable summer heat, mountain areas that receive snow in winter and remain cool in summer, and (lightly populated) regions of the Sonora desert with extreme hot and cold weather. Since home electricity use is highly sensitive to weather, it is important to account for this variation to avoid mis-specifying the conditions a household actually experienced.

To this end, we used the nine-digit zip code information in the HERBS to map each sample household to one of 21 National Weather Service (NWS) stations. This process matches a household to a local weather station considering both proximity and elevation.⁹ The 21 weather stations are located throughout SDG&E's service territory, but are concentrated near population centers.

We construct a household's weather variables using the matched weather station's

⁷Occasionally, a meter cannot be read on the scheduled cycle-closing day, but is read a few days later. Actual read dates are recorded in our data. In rare cases, SDG&E postpones a bill until the following month because it cannot read a meter (billing the household for two months). In our analyses, we treat the middle month(s) as missing. Because it is willing to postpone billing, SDG&E estimates very few bills.

⁸Because of the cohort system, SDG&E may send a shorter 'start-up' or 'closing' bill when a household moves into or out of the area. Since these periods commonly reflect highly irregular use (such as a dwelling temporarily empty), we exclude a small number of bills in the data that cover start-up or closing periods of less than 15 days. We also drop a very small number of irregular bills that cover more than 60 days.

⁹The weather data are from the National Oceanic and Atmospheric Administration (NOAA) daily temperature files available at www.noaa.gov.

data. The weather variables are billing-period heating and cooling degree-days; these are standard measures of electric heating and cooling demands. Cooling degree-days are (a discrete approximation to) the positive part of the temperature time path less 65°F, integrated over time.¹⁰ High values occur in the summer, reflecting high demand for cooling services; cooling degree-days are usually zero in winter. Heating degree-days approximate the negative part of the temperature time path less 65°F over the billing period. They are large in winter months, reflecting high demand for home heating, and usually zero (or nearly so) in summer. We calculate a household’s heating and cooling degree-day values using the exact start and end dates for each bill, standardized by the number of billing days.

Table 2 summarizes how cooling degree days and electricity consumption varied from 1998 to 2001. The top panel reveals that San Diegan’s cooling degree-days vary considerably over time. Most importantly, 1999 was substantially cooler than 2000; cooling degree-days in August and September 2000 increased more than 60 percent from the same period in 1999.¹¹ An increase of this magnitude would customarily produce a much greater demand for electricity, *ceteris paribus*.

This large variation in weather makes direct casual inferences about changes in consumption between the electricity-crisis period (2000-01) and the preceding years of 1998-99 problematic. For instance, although prices more than doubled from August 1999 to August 2000, the lower panel of Table 2 reveals that average household electricity consumption *slightly increased*. The obvious explanation for this increase is the offsetting effect of weather changes between August 2000 and August 1999. Our chief task is thus to develop a method for disentangling the effects of weather and prices on consumption.

¹⁰A household’s cooling degree-days for the billing period running from calendar date d_1 to d_2 are $CDD = \sum_{t=d_1}^{d_2-1} \max\{A(t) - 65, 0\}$, where $A(t)$ is the average of the high and low temperatures in °F at the household’s local (matched) weather station on date t . Heating degree days (HDD) are computed with the summand replaced by $\max\{65 - A(t), 0\}$.

¹¹The large variation across years is due in part to an El Niño Pacific weather disturbance during 1998-99. As noted earlier, the monthly averages also mask considerable variation in cooling degree days across households. For instance, in August 2000 households located along the coast experienced less than 150 cooling degree days, while those in the inland desert areas averaged over 700.

4 Empirical Strategy and Methods

This section motivates and summarizes our methods. These are designed to overcome two key problems in the estimation of consumers' responses to the electricity crisis. First, from a research design standpoint, there is the problem that everyone in San Diego experienced the same price changes at about the same time (see again the discussion of Figure 1). Thus, there is no way to construct the equivalent of two randomly assigned 'treatment' and 'control' groups. All households in the population experienced the same price changes in 2000, public appeals in 2001, and so on.

The second difficulty is that it is far from clear what consumers knew about prices at the time. As noted earlier, the price indexing system in place between July 1999 and September 2000 determined the price consumers would pay when bills were issued. Consequently, during the summer of 2000 a consumer might have behaved as if its prior month's price would apply, or behaved as if a much higher marginal price would apply, or anything in between. The same uncertainty applies in the fall with respect to the price cap. Further complicating matters is the fact that the perception of what price would be in effect each month surely differed across consumers—depending, for example, on how attuned the household was to media coverage of the unfolding electricity crisis. The upshot of these issues is that it is difficult to correctly specify statistical models that assume, implicitly or explicitly, what the consumer knew about prices at the time consumption occurred.

Our approach handles these problems by exploiting the fact that we observe consumption in two distinct 'regimes': The first is the pre-crisis period, during which prices did not vary. The second is the crisis period, during which prices fluctuated and related events transpired. This sequence allows us to use variation in consumption during the stable-price period to disentangle the effects of some confounding factors (weather, in particular) during the subsequent crisis period. In essence, we are using a household's consumption during the period when prices were held fixed as the 'control' (and implicit counter-factual), and contrasting this with consumption following crisis-related events.

Our basic unit of analysis is the change in a household's average daily consumption (ADC) between one billing period and the same billing period twelve months earlier (*i.e.*, twelve-month differences). Twelve-month differencing removes seasonal effects of no immediate interest, such as the number of daylight hours per day. These effects

can also be household specific, as with a family that vacations each August or that has enormous outdoor light displays during the December holidays.¹²

We divide the variation in a household’s twelve-month consumption differences into two distinct categories. The first are factors idiosyncratic to the household. These include remodeling the kitchen during the intervening year, having a teenager leave home for college, changing the month of a family vacation between one year and the next, replacing a failing major appliance (for reasons unrelated to the electricity crisis), and so on.¹³ Averaged across all households, things in this first category define the *secular trend* in residential electricity consumption. By contrast, the second category includes non-idiosyncratic factors that induce correlation in consumption changes among households. These include common price changes, crisis media attention and the public appeals noted earlier, and regional weather patterns.

As is well known from the electricity demand literature, much of the variation in electricity consumption comes from changes in the weather. This is particularly true in California. (See, for example, studies by Acton, Mitchell, and Mowill (1976), Parti and Parti (1980), and Reiss and White (2001).) The close correspondence between weather and electricity consumption in the San Diego area is summarized in Figure 2. The solid line in Figure 2 shows average daily consumption by billing cohort. The dashed and dotted lines similarly depict heating and cooling degrees experienced by these billing cohorts during their billing months.

Figure 2 shows that average daily consumption varies considerably over the course of a year. Most of this variation closely coincides with the variation in the weather. For example, during 1998, consumption is at its highest in January when heating degree-days peak. Electricity usage peaks again in the summer as cooling degree-days peak. From peaks to troughs, consumption typically varies 20 to 30 percent. Thus, it is important to account for this variation before we can begin to address the influence of crisis-related events on consumption.

We assume that a household’s same-month, year-over-year electricity consumption

¹²From prior work we also know that electricity consumption depends greatly on appliance stocks, dwelling structure attributes, and demographics (e.g., the number of children in the home); see Reiss and White (2001). Because these factors are unobserved in our billing data, here we seek to explain variation in consumption only within, rather than between, households.

¹³We term these factors idiosyncratic in the statistical sense that even if we knew such an event occurred for a household, that information would be of no value for predicting the contemporaneous consumption change of another randomly-sampled household.

changes can be decomposed into its pre-crisis trend, crisis-period responses, weather, and idiosyncratic factors according to

$$x_{icm} - x_{ic,m-12} = \delta_i^{pre} \cdot I_m^{pre} + \delta_{icm}^{post} \cdot I_m^{post} + (w_{icm} - w_{ic,m-12})' \beta_i + \varepsilon_{icm} \quad (1)$$

where x_{icm} is the ADC of household i in cohort c for the billing period ending in month m (indexed serially), w_{icm} a vector of contemporaneous weather-related covariates, and ε_{icm} an idiosyncratic shock. Here I_m^{pre} is an indicator variable for pre-crisis months (May 2000 or earlier), and I_m^{post} similarly for months after the crisis begins (June 2000 and later). The δ 's and β 's are household-specific parameters to be estimated. We allow δ_{icm}^{post} to vary by month, to accommodate heterogeneity in how rapidly households react to crisis events.

Although we compute things slightly differently, a standard measure of how a household's consumption changed from pre-crisis behavior by month m during the crisis is the difference-in-differences,

$$D_{icm} = \delta_{icm}^{post} - \delta_i^{pre}.$$

Ultimate interest lies not in the individual values of D_{icm} , but in its population distribution and the time-path of the population average $\bar{D}_{cm} = \sum_{i=1}^{N_{cm}} D_{icm} / N_{cm}$, where N_{cm} is the cohort population size.

The specification in (1) implies a household's sensitivity to weather is time-invariant. In reality, this is likely to be implausible for many households. A natural response to the rise in prices during the summer of 2000, for instance, is to curtail home air conditioning during hot weather. This means that β_i during pre-crisis periods may differ systematically from its value after the crisis begins. Since it is the change in consumption behavior relative to pre-crisis use patterns that is of interest counter-factually, we estimate the difference D_{icm} for each household without imposing this time-invariance restriction.

A Few Details

To be specific, we estimate the parameters of (1) the following way. First, we fit to the (up to) 32 months of pre-crisis data the regression model:

$$x_{icm} - x_{ic,m-12} = \delta_i^{pre} + (w_{icm} - w_{ic,m-12})' \beta_i + \varepsilon_{icm} \quad (2)$$

where ε_{icm} is assumed mean zero given the weather. This equation is fitted separately for each household via ordinary least squares. We then compute the change in consumption during crisis-period months as actual minus predicted,

$$\widehat{D}_{icm} = [x_{icm} - x_{ic,m-12}] - [\widehat{\delta}_i^{pre} + (w_{icm} - w_{ic,m-12})'\widehat{\beta}_i] \quad (3)$$

with $\widehat{\delta}_i^{pre}, \widehat{\beta}_i$ estimated from pre-crisis data. We expect \widehat{D}_{icm} to be close to zero for any month m prior to the crisis, and non-zero for any month after.¹⁴

Two wrinkles remain in estimating the average difference \overline{D}_{cm} from the individual values \widehat{D}_{icm} . First, as noted earlier there is the matter of sample entry and attrition. We estimate \overline{D}_{cm} using weighted sample averages of the form

$$\widehat{\overline{D}}_{cm} = \sum_{i=1}^{n_{cm}} \omega_{icm} \widehat{D}_{icm}$$

where n_{cm} is the observed number of households in cohort c in month m , and ω_{icm} is a time-dependent sampling weight that diverges from $1/n_{cm}$ to account for entry and attrition (see Appendix A).

Second, in practice it is statistically desirable to pool estimates from more than one year of pre-crisis data. Thus, in estimating consumption changes relative to the same month of pre-crisis years, we use a pooled difference estimator of the form

$$\widehat{\overline{D}}_{cm}^p = [\widehat{\overline{D}}_{cm} + \widehat{\overline{D}}_{cm}^{(24)}] / 2 \quad (4)$$

where $\widehat{\overline{D}}_{cm}^{(24)}$ differs from $\widehat{\overline{D}}_{cm}$ only in that the values of x and w at 12-month lags in (3) are replaced by their values at 24-month lags (with $\widehat{\delta}_i^{pre}$ doubled).¹⁵ Pooling two estimates in this way reduces variability attributable to idiosyncratic departures from a household's pre-crisis trend in any one year during the pre-crisis period—a source of

¹⁴ Because the electricity crisis continued into the summer of 2001, the 12-month lag structure in (1-3) is not universally appropriate. For the months of June 2001 and following, interest still centers on how consumption differed from its (weather-adjusted) pattern during pre-crisis years. Consequently, to compute a value for \widehat{D}_{icm} when m is later than May 2001, it is necessary to replace $x_{ic,m-12}$ and $w_{ic,m-12}$ in (3) with their values at 24-month-lags and double $\widehat{\delta}_i^{pre}$. This implies \widehat{D}_{icm} for m of August 2001 is relative to the pre-crisis month of August 1999, not the crisis-period month of August 2000.

¹⁵In constructing the 24-lag analog for the pooled difference estimator, the coefficient estimates in (3) are still computed using (2) as written.

statistical ‘noise’ from our perspective. The pooled difference estimator \widehat{D}_{cm}^p is thus the average change in within-household consumption between month m and the same months of pre-crisis years, less the change we would expect on the basis of weather and pre-crisis trend alone.¹⁶ This is the statistic we report in the time-series results below.

Some comments regarding specification are in order. Our use of a linear-in-the-parameters specification for the model in (1) and (2) is driven by the results of a large literature (principally in statistics and engineering) on the relationship between weather and electricity sales. A key paper in this literature is Engle et al. (1986), who examine parametric and non-parametric techniques for modeling this relationship with residential billing data from various utilities. One conclusion from this work is that the relationship between temperature and electricity demand is linear (or very close to it) for temperatures above 65°F, and similarly (with a different slope) below 65°F.¹⁷ For this reason we use the simple and parsimonious specification in (2), rather than less-efficient nonlinear smoothing techniques.¹⁸

A second conclusion from this literature is that when measured carefully to account for billing-cycle timing, local weather, et cetera, the relationship between weather and residential electricity consumption is tight. When predicting average (or aggregate) household electricity consumption, parametric linear models of the form in (2) commonly have adjusted R-square values of 0.95 or better using degree-day weather data (see Engle et al. (1986) or EPRI (1983)). This high level of predictive accuracy is confirmed in the HERBS with our models as well, as shown presently.¹⁹

¹⁶Apropos note (14), we use 24- and 36-month lags in constructing the pooled difference estimator for any month after May 2001.

¹⁷These slopes vary widely across utilities; see especially Engle et al. (1986), Figs. 1–4. Such geographic differences reflect the prevalence of air conditioning and electric space heating among homes in different parts of the country.

¹⁸The specification in (2), which uses heating and cooling degree-days at base 65°F in w , is a linear transform of the piecewise-linear relationship (about 65°F) between average daily temperature and electricity consumption.

¹⁹One might also anticipate some degree of serial correlation in ε_{icm} , on the basis of various behavioral considerations (e.g., the lumpiness of durable-appliance replacement timing). In the data, we find little evidence of this. In any event, serial correlation in ε_{icm} is only an efficiency issue here, and a small loss of efficiency is tolerable given the large sample size.

Identification Issues

Identification and interpretation of δ_{icm}^{post} and \bar{D}_{cm} merit brief discussion. Technically, an estimate of \bar{D}_{cm} will measure the combined effect of everything, other than weather, that moved average household consumption away from its pre-crisis trend at time m . Its interpretation as consumers' responses to prices and other crisis-related events is justified under two assumptions. These are (1) that the average pre-crisis secular trend, $\bar{\delta}^{pre}$, would have continued during 2000-01 had the crisis not occurred; and (2) that households' weather-sensitivity parameters (the β_i 's) are stable over time 'but-for' the crisis. Both of these assumptions are plausible. Moreover, as will become clear from the results in Sections 5 and 6, small violations of these identifying assumptions will not imperil the major findings.

Although identifying assumptions are by nature not directly testable, these two can be addressed informatively. Assumption (2) is essentially one about the rate at which air conditioning is installed in residential dwellings. Roughly one-third of homes in San Diego County have air conditioning, and its prevalence has increased at a low rate (approximately 1% p.a.) for nearly a decade. The effect of a gradual increase of this magnitude on our results is negligible; it would take more than double-digit increase in the percent of homes with air conditioning to discernibly affect our estimate of the time-path of consumption changes, \bar{D}_{cm} . An increase of that magnitude over one to two years is not a realistic concern.

At issue with assumption (1) is whether there are any unobserved factors that would have shifted average consumption growth, net of weather, absent the crisis. The principal quantitative evidence against this possibility is the close fit between fluctuations in weather and electricity consumption in the data during the pre-crisis years. Any substantial unobserved factor would have to become influential in 2000 or 2001, but not be similarly influential during preceding years. Since the average pre-crisis secular trend in electricity consumption represents the aggregation of hundreds of thousands of individual household decisions about appliance purchases, vacations, remodeling, and so on, it is difficult to conceive that this trend would have changed abruptly in 2000-01 absent a macroeconomic shock.

One potential point of concern is income. Real personal income per household grew at 3.5% p.a. from 1997 through 2000 in San Diego, but began to slow in late

2000 and declined slightly in 2001.²⁰ Since estimates of the short-run income elasticity of residential electricity demand are low, on the order of 0.1, this change is unlikely to reduce household electricity consumption in 2001 by more than a few tenths of one percent. Nevertheless, we address the sensitivity of our 2001 results to this issue directly in Section 6.

Pre-Crisis Period Fit

Table 3 summarizes the coefficient estimates for the approximately 46,800 households with sufficient data to estimate the model (2) using pre-crisis data. The top panel reports the means and selected percentiles of the raw coefficient estimates across households; the middle panel presents the same figures after re-weighting to account entry/attrition, yielding only minor changes. Although the weather-related parameters' units make their values difficult to interpret intuitively, the means and medians are in accord with prior work in this area (*cf.*, Engle et al. (1986), EPRI (1983)). Note we estimate separate heating- and cooling-sensitivity parameters for in-season and out-of-season use. The bottom panel converts the weather parameters to elasticities for interpretation. On average, a doubling of temperatures (as measured by cooling degree-days) between one summer month and the same month of the previous year implies about a 30% increase in household electricity consumption. (Somewhat smaller changes occurred between 1998, 1999, and 2000; see Table 2). The analogous consumption sensitivity to winter temperature changes is about one-third as large.²¹

The secular trend in household consumption averages between 2.5 and 3.3% per annum during the pre-crisis period. The table reports separate pre-crisis trend coefficient estimates before versus after October 1999. Our motivation for this arose from income and aggregate electricity consumption data from elsewhere in California,

²⁰Calculated from the BEA's personal income series for the San Diego MSA, local CPI-U, and Census Bureau annual household estimates.

²¹Table 3 also reveals that there is considerable variation in the estimated responses to the weather variables, with some estimated effects being small or negative. While it is possible that some households reduce consumption in response to small increases in heating and cooling degree-days, our preferred interpretation is that the response of many households is truly zero and that the negative coefficients reflect idiosyncratic factors (e.g., travel, an appliance purchase) that happened to occur in a month of notable weather change from the previous year. Our prior work using different data for California (Reiss and White (2001)) indicates that households differ substantially in their response to weather, and a non-trivial fraction of households exhibit no measurable response to temperature.

which suggested a change about this time. The estimated difference in the mean trend before versus after this point is small and the medians are nearly identical. Additional investigation revealed no indication of a change in the average trend at other times during the pre-crisis period. We use the post-October 1999 trend parameters in our estimates of \overline{D}_{cm} during June through December 2000, reported below.

A critical question that needs to be addressed for later analyses is whether the parsimonious model in (2) predicts consumption accurately (when prices remain fixed). Figure 3 presents the average actual and predicted consumption changes during the stable-price period. To facilitate interpretation, we have normalized the vertical units from ΔKWh to percentages (dividing each data point by the population average daily consumption during this period, a constant equal to 16.04 KWh). The actual and predicted series are close, despite the large variation in average consumption and weather during this period. The overall R-square for the average actual and predicted data summarized in Figure 3 is 0.97. Notably, the models pick up the substantial 20% drop in average consumption between August and September of 1998 and 1999; they do equally well with the 10% increase in consumption in October 1999.²² The most significant departure between actual and predicted occurs March-April 2000. After some investigation, this appears attributable to an unusual pattern of unseasonably-cold snaps in late spring that year; the time-path of temperature suggests that heating-degree days may not fully capture the demand for heating at the time. Since this pattern is not observed in the weather data previously, the models fit this time less than perfectly. Such a pattern did not recur during 2000 or 2001.

The close agreement between average actual and predicted electricity consumption in our data is not unexpected, given prior work in this area. The important consequence that comes through in Figure 3 is that year-over-year changes in average household electricity consumption are basically all weather, absent a change in prices. It is the stability of this tight relationship over time that we rely upon to draw inferences about how households responded to subsequent crisis-related events.

²²Note that weather during the two crisis-period years (2000-01) falls within the range of variation observed during the fitted stable-price period (1998-99); see Table 2.

5 The Price Spike and Price Cap During 2000

We now describe our empirical findings on households' responses to the summer 2000 price spike and subsequent price cap. Our discussion proceeds in two parts: the first summarizes and interprets the main statistical results, and the second addresses their implications.

5.1 Main Results

Figure 4 describes households' average consumption responses to the crisis in the summer and fall of 2000. For comparison, we also provide the same information for 1999 and the spring of 2000, prior to the crisis. The solid line shown is the pooled difference estimator \widehat{D}_{cm}^p in (4), calculated relative to households' 1998 and 1999 consumption levels. This statistic, which varies (nearly) daily, is non-parametrically smoothed slightly for clarity.²³ The line represents the average within-household change in electricity consumption from the same month of prior years, after subtracting the change in consumption predicted by pre-crisis trend and weather alone.

During the stable-price period, the line shown in Figure 4 is effectively zero. This much is to be expected from Figure 3. The small fluctuations of the line during 1999 and the first-half of 2000 are (primarily) attributable to unmeasured factors that influence same-month changes in average electricity demand from one year to the next (such as unusual weather patterns as in March 2000, noted earlier). The absolute error represented by these fluctuations is typically about one percent, which is small compared to the changes to come in mid-2000.

The data in Figure 4 reveal an abrupt change in household electricity use during the summer of 2000. From mid-July on, average residential electricity consumption begins a precipitous decline. Table 4 provides a numerical summary of these changes along with information on average prices. By the end of summer 2000, San Diego's million households consumed 12 to 13 % less electricity than in prior years, net of the effects of weather and adjusted for trend. Table 4 also shows that although consumers faced dramatically higher prices in July, their consumption over the prior

²³Here and throughout we use a standard loess (locally-estimated scatter plot smoothing) procedure with a narrow bandwidth to present time-series results. This procedure reduces sampling variability and variation due to the composition of bills that close on a particular day (i.e., billing-cohort effects). Billing-cohort variation is due to SDG&E's geographically-rotating meter-reading schedule.

month did not appear to anticipate these higher prices. The decline in consumption is first evident toward the end of July – a full month after prices began their rapid rise. Given that bills reflect consumption over the preceding 30 or so days, Table 4 implies a 12 to 13 % drop in consumption occurred in a short span of about 60 days.

Figure 4 further reveals that the decline in average consumption ends nearly as abruptly as it began. This change is first observed in bills closing in September, which implies that households began to reverse course in their consumption behavior conceivably as early as late August. Figure 4 and Table 4 then show a rapid rise in average consumption from late September through October. This effect is coincident with the imposition of the (retroactive) price cap in early September. We refer to this as the *rebound effect*. The slight decline that follows the fall 2000 consumption rebound, evident for December in Figure 4, fore-shadows another change in households' consumption behavior as the electricity crisis escalated during winter 2000 and early 2001.

Interpreting Response Magnitudes

The electricity consumption changes shown in Figure 4 and Table 4 present a reasonably compelling case that San Diego-area households responded to the abrupt rise and subsequent fall in electricity prices during 2000. Moreover, they did so fairly rapidly. In order to grasp the implications of these results, it is important to address first a question of interpretation: Is a 12 to 13% drop in electricity consumption a big response, or a small one?

A 12 to 13% percent reduction in electricity use requires a non-trivial behavioral change for most households. It involves not turning on air conditioners or raising temperature settings substantially, re-setting pool filters to operate far fewer hours, dramatically reducing television time, or other such actions. To have a sense of the magnitudes involved, appliance-level electricity consumption data from San Diego indicate that air conditioning accounts for about 15% of aggregate residential electricity consumption during the summer.²⁴ Only steep reductions in air conditioner use could therefore account for the observed 12 to 13% drop in aggregate residential electricity consumption—and this during a relatively warm summer. Alternatively (or in combi-

²⁴Derived from EPRI (1989) Table 3-4, which is based on appliance sub-metering studies of San Diego households performed by SDG&E.

nation), changes in other appliances' use would have to be substantial. For example, if households reduced their aggregate use of electric clothes dryers by one-half (by substituting to clothes lines), electric stoves and ovens by one-half (by changing the foods they prepare or dining elsewhere), dishwashers by one-half (ibid), home lighting by one-third, and additionally cut television time by one-half, then *combined* these five changes could amount to a 10% reduction in average electricity consumption.²⁵ One can continue with similar such examples for less electricity-intensive appliances, but the picture here is clear. Because a large share of San Diegans' aggregate residential electricity consumption is attributable to appliances for which consumption cannot be easily changed in the short-run (most importantly refrigerators, freezers, and electric water heaters), substantial changes must have taken place in how often the other major appliances are used.

Additionally, households could also respond by replacing aging and inefficient appliances with newer, more-energy efficient ones. There is evidence of this from contemporaneous secondary sources (appliance sales data, press accounts, and surveys). Telephone interviews with a sample of 400 SDG&E-area households by Lutzenhiser (2002) indicate that 1 in 8 report purchasing and replacing a major appliance with a more efficient model during the electricity crisis. The net effect of these longer-term actions with respect to electricity consumption are discussed further below. Here we note that these actions represent a substantial behavioral response of a different sort, inasmuch as the expenditures involved may represent primarily foregone consumption of other goods and services. Replacing any electricity-intensive durable appliance in the home (e.g., air conditioner, refrigerator, water heater, dishwasher, or clothes dryer) is an expensive outlay, even in comparison to the electricity savings at San Diegans' summer 2000 electricity prices.²⁶

Heterogeneity of Responses

There is, of course, a great deal of heterogeneity in households' consumption responses to what was (essentially) an identical price change for all consumers. Figure 5 shows

²⁵Based on appliance-level electricity consumption and utilization data from the U.S. Department of Energy (EIA (1997), Table 3.1).

²⁶There is also the issue of to what extent the price increases simply accelerated replacement of durables that would have soon been replaced anyway. Unfortunately, this question is not easily assessed.

the cross-sectional distribution of individual households' consumption changes between August 1999 and August 2000. For comparison the figure also shows the distribution of households' consumption changes between August 1998 and August 1999, when no appreciable price changes occurred. The two distributions are based on each household's adjusted (net) consumption change, which is the actual change less the change predicted by pre-crisis trend and weather alone.

The distribution for 1999-to-2000 shows extraordinary variance, with a substantial number of households exhibiting consumption changes at the extremes in *both* directions. On the right, there is a substantial share of the population who actually increased consumption from August 1999 to August 2000, net of weather and (pre-crisis) trend. This group comprises approximately 36% of households, which is down from 50% during the prior years' comparison period. This magnitude is in line with prior econometric work using different data that indicates approximately 2/5ths of the California population are completely price-inelastic electricity users (see Reiss and White (2001)). As a practical matter, these findings suggest that roughly one-third of the population would exhibit no observable consumption decrease in response to even large changes in electricity prices.

On the opposite end, there is a large share of households that show dramatic reductions in electricity consumption. Between August 1999 and August 2000, approximately 1 in 3 households reduced electricity use by 20% or more (again, relative to pre-crisis trend and adjusted for weather). For the prior years' period of August 1998 to 1999, the directly comparable figure is only 1 in 14 households. Thus, a substantial share of the decline in average (and aggregate) consumption is attributable to a minority of the population who reduced their electricity use dramatically during the summer of 2000. Consumption reductions of 20% or more require careful and attentive efforts to lower electricity use, major investments in new appliances, or substitution behaviors that one suspects would pose considerable inconvenience within the household.

Figure 6 provides additional information on the heterogeneity in consumers' behaviors during California's electricity crisis. It shows the percent change in households' electricity consumption between August 2000 and August of prior years, again adjusted for weather and relative to pre-crisis trend, versus households' consumption *levels* (in kilowatt-hours) for August 1999. Each point in the scatter is an individual

household.²⁷ The solid line in the figure is the (non-parametric) line of averages, which shows the mean consumption response at each level of consumption. Average household consumption for August 1999 is 498 KWh.

A striking aspect of Figure 6 is the enormous variance in year-over-year consumption changes across households, at all consumption levels. This variance—recall it is already net of weather’s influence—is precisely why it is difficult to estimate the sensitivity of electricity demand to small price changes. Such price changes are the nature of the historical data upon which most econometric studies are based. Despite this dispersion in consumption responses, however, it is clear that the change at the mean is negative here (as summarized in Figure 4). By comparison, a similar scatter plot (not shown here) of August 1998 to 1999 changes versus August 1998 levels—an interval when no appreciable price changes occurred—shows no such change. In that case, the weather- and trend-adjusted consumption changes are nearly perfectly symmetrically distributed about zero at all consumption levels.

5.2 Discussion

The results in Figures 4 through 6 speak to issues of broader significance for energy markets and regulatory policies. It is perhaps not surprising to many economists that households would adjust their electricity consumption in response to price changes even over short-term durations. But this fact has not been widely recognized by policy-makers, in part because clear and unambiguous evidence of such behavior has been heretofore lacking. The results above indicate that consumers have considerable control over the short-run electricity use of their household appliances, and are willing to modify that use in response to pecuniary incentives.

In post-mortem analyses of electricity market problems in California and elsewhere, the sensitivity of demand to price signals has become a central focus of reforms. Numerous observers have noted the insensitivity of electricity demand to the actual cost of producing power at any point in time (Borenstein (2002), Joskow (2001), Wolak (2003)). As these authors point out, however, the diagnosis that demand is nearly perfectly inelastic with respect to the wholesale price of power (San Diego being a brief exception) need not reflect consumer preferences. Rather, it more likely reflects the tenuous link between the naturally volatile price of electricity on

²⁷Figure 6 clips a small fraction of households whose consumption exceeds 1500 KWh per month.

wholesale commodity markets and the average-cost-based retail electricity prices employed by state utility regulators. This observation has led to a variety of proposals for reforming retail pricing mechanisms to improve demand responsiveness and thus market efficiency (such as real-time pricing, critical-peak pricing, or more widespread use of time-of-day pricing; see Lafferty et al. (2001) or CPUC (2002)). The principal unknown confronting these proposals is whether consumers—households, especially—are actually willing and able to forgo electricity consumption over short-term horizons. The data from San Diego lend considerable credence to the view that consumers—or at least a large share of them—would react to such programs by changing their consumption behavior substantially whenever prices do.

Nevertheless, we caution against directly interpreting these results as classical demand elasticities. It is tempting to divide the August 2000 drop in average demand by the contemporaneous percent change in prices from Figure 1, and declare the ratio a price elasticity of demand. We have not done so, as this is likely to be misleading if applied uncritically in other circumstances. Electricity consumption was still falling precipitously at the time the price cap was imposed, ending the unintended ‘experiment’ in demand sensitivity before consumer adjustment to these prices had stabilized at a new consumption level. This limitation means we cannot construe the consumption drop here as a complete adjustment to the price change, which is the assumption implicit in economists’ customary use and interpretation of price elasticities. At best, one can interpret the results here as informing the *60-day elasticity* (or so) of residential electricity demand with respect to an *unanticipated* increase (roughly doubling) in price.

Marginal Valuations

An important related observation about consumers’ willingness-to-pay can be derived from these results. The magnitude of San Diego households’ average electricity consumption response is all the more remarkable when one considers that the typical monetary savings achieved is rather small. The median household reduction in electricity consumption during August 2000, after adjustment for weather and relative to pre-crisis trend, is 44 kilowatt-hours in physical terms. This is about 9% of average pre-crisis consumption. At then-prevailing electricity prices, reducing demand by 44 KWh would lower a household’s total expenditures by approximately \$10.25.

This amount seems a small income gain for implementing a 9% decrease in house-

hold electricity use, particularly given our earlier discussion. Since this magnitude of an income-electricity use trade-off was in fact made—by half of a population of 1.1 million San Diego households—it strikes us that the data show consumers are remarkably willing to forgo electricity consumption for comparatively paltry monetary amounts. In other words, the marginal value of the last 5% or so of aggregate residential electricity consumption each month must be low—on the order of $\$10.25 \times \frac{1}{2}$ 1.1 million in the aggregate, or about \$5.12 per household on average.

Effects of the Price Cap

Figure 4 shows that after the price cap was imposed in the fall of 2000, average electricity consumption quickly rebounded upward. This rapid reversal might not be considered surprising, given the rate at which consumers reduced demand when prices rose during the summer. Average household consumption increased approximately 7% between mid-September and mid-October, on a weather-adjusted basis. As the evidence from earlier in 2000 showed, consumers clearly react—here with less than a one month lag—to changes in their electricity prices.

One feature to note here is that the response to the price cap in the fall of 2000 was surely influenced, to some extent, by its retroactive provision. Billing credits issued in the fall of 2000 refunded summer electricity charges in excess of the new price cap. For many households, these credits made electricity bills in September through November net to zero. An inattentive consumer might have observed and reacted to only this bill total, rather than the positive marginal price (capped at roughly 13.5 cents per KWh) during this period.

It is also noteworthy that the consumption rebound amounts to only about two-thirds of the decline that occurred during the price spike. If simply foregoing consumption (*i.e.*, not turning on appliances) accounted for all of households' summer 2000 behavioral responses, one might expect a complete rebound in the fall. The demand rebound following the price cap leaves average consumption about four percent below its pre-spike level, however, and it remains at that level (on a weather-adjusted basis) throughout that fall.

One interpretation of this fact is that it suggests that approximately one-third of consumers' aggregate response to the price spike was realized through changes in appliance stocks, dwelling improvements, or persistent changes in utilization decisions. This gives some quantitative credence to the self-reported evidence on appliance re-

placement from Lutzenhiser (2002), noted above. Moreover, during all of 2001 and into 2002, demand (net weather effects) never rises above the level evident immediately following the price cap's imposition. The implication is even transitory price spikes, if not pre-announced (as would generally be the case) appear to have measurable long-term impacts on energy demand.

6 Post-Cap Interventions

6.1 Public Appeals and Crisis Attention

Although San Diego consumers' electricity prices were capped throughout 2001, prices on California's wholesale electricity market continued to escalate dramatically in December 2000 and January 2001 (see again Figure 1). The fact that wholesale electricity prices exceeded retail prices by a substantial margin led California's regulated utilities to the brink of insolvency. An ensuing credit crisis further exacerbated the inefficiency problems in the wholesale market, leading the state into a true (physical) electricity supply crisis during the winter of 2000 and spring of 2001. In fact, a limited number of consumers experienced involuntary 'rolling blackouts' (electricity rationing, in effect) during January, March, and again in May of 2001.

During these winter and early spring months, state officials made dramatic televised appeals for consumers to voluntarily conserve electricity. The deepening crisis was routinely aired on nightly news programs in California (and nationally), particularly following the realization of rolling blackouts to homes and businesses. Nevertheless, state leaders feared a widespread—and well-founded—political backlash if consumers' electricity prices were to increase commensurate with wholesale markets'. State utility regulators refused to allow essentially any change in retail electricity prices throughout this period.²⁸ Residential prices in San Diego, in particular, continued at the capped level shown in Figure 1 throughout the spring, summer, and early fall of 2001. A variety of conservation programs and incentives also were put in place in early 2001, but these programs were not primarily aimed at residential

²⁸The state's public utilities commission approved a one cent per KWh increase in prices for two of the states' utilities (serving areas other than San Diego) in January of 2001. This 'emergency' measure was little more than a token political maneuver by regulators to signal awareness of the looming cash-flow crisis to utilities' creditors; substantive changes in prices did not occur until June of 2001 for these two utilities, and not until November 2001 for the utility serving San Diego.

customers.

Did these (non-price) events—media attention to the crisis and officials’ public appeals for voluntary electricity conservation—alter households’ consumption? Standard economic theory suggests that given prices were capped at a constant level, consumers would have little incentive to reduce individual consumption. The empirical evidence suggests that the opposite was the case: there was, in fact, a substantial response by California households to these events.

Figure 7 presents the main evidence on this point. This figure extends the time-series shown in Figure 4 into 2001. As earlier, the line shows the average change in consumption from corresponding months of the pre-crisis period, net the effects of weather and relative to pre-crisis trend.²⁹ Figure 7 shows a steady and sizable continuous reduction in household electricity consumption from January through May 2001. This decline amounts to approximately 6% of average pre-crisis consumption, on a weather-adjusted basis. Interestingly, the two largest drops (in late-January to February bills and mid-March to early-April bills) occur following two sets of limited, but well-publicized, rolling-blackouts in mid-January and mid-March. The second set affected approximately 41 thousand customers in the San Diego region.

That consumption dropped substantially when prices were not changing is important for several reasons. Public appeals for energy conservation by government officials have a long history in times of energy crises, to questionable effect. In economic terms, voluntary energy conservation is a collective-action problem that is subject to seemingly extreme free-rider difficulties. Factually, the probability an individual household could affect the aggregate supply-demand electricity balance in a grid the size of a major California city is infinitesimally small. Moreover, there is no real way for others to know whether a particular household is conserving electricity or not—even among immediate neighbors. One cannot readily check a neighboring household’s indoor air temperature, daily hot-water use, or other major appliances. Thus, explanations that rely upon a household’s desire not to be seen as a profligate

²⁹We make one important adjust to our prior methods here. Because personal income growth had begun to slow by the beginning of 2001 (see again Section 4), it is possible that households’ secular trend in electricity consumption would have slowed from the pre-crisis rates estimated in Table 3. Consequently, we constructed Figure 7 assuming that there was *zero* secular consumption growth after January 2001. We view this assumption as highly conservative, since assuming positive growth (the historical norm) would result in greater estimated consumption decreases and electricity consumption is highly income-inelastic. One can visually adjust the line in Figure 7 proportionately for an alternative trend growth-rate assumption.

user by others face a basic problem: one’s effort to reduce electricity consumption—or not—is (essentially) visible only to that household.

Despite these features and the absence of any change in pecuniary incentives, however, households collectively responded dramatically to the media attention and public appeals for electricity conservation during this period. The nature of individuals’ free-rider problem here and the lack of private incentives for electricity conservation leave largely ‘moral suasion’-type arguments to explain their behavior: Consumers individually wanting to ‘do their part’ to mitigate the electricity crisis, et cetera.³⁰ Although economists tend to be dismissive of public appeals than run contrary to private incentives when free-rider problems are great, one cannot dismiss the striking reduction in average household electricity consumption shown in Figure 7. The empirical facts here indicate that consumers do respond to voluntary appeals to modify their consumption behavior, provided (1) the costs of a collective-action failure are tangible (here, involuntary blackouts for some consumers), and (2) media attention to the problem is considerable.

6.2 The Conservation Price-Rebate Program

In the spring of 2001, California faced the prospect of more frequent electricity shortages and continuing wholesale market problems during the coming summer. One major policy response was a novel voluntary conservation rebate program for the summer of 2001. In the San Diego region, this program offered each household a 20% rebate against its monthly electricity bill (via a line-item bill credit) if the household reduced its electricity use by 15% or more from the same month of 2000. This program was promoted as a “cash for energy savings” program, with the program’s cost paid from the state’s general revenue. Although similar rebate programs had been tried on smaller scales, this program appears to be the first large-scale effort to pay households for reducing their consumption. As such, it was followed by industry participants with great interest.

³⁰It is also possible consumers believed that by reducing their consumption individually, they could reduce the chance of future involuntary blackouts. Although this is difficult to assess directly, Lutzenhiser (2002, Table 18) reports that among households surveyed in the Southern California Edison service area during 2001, 30% claimed to have experienced a blackout ordered by the electricity system operators. The actual number is closer to five percent. Californians evidently viewed their blackouts as widespread, although they directly affected relatively few.

Before the program was implemented, there was some controversy about whether it would alter consumers' behavior. First, industry experts worried that the rebate threshold was not sensitive to potential weather differences between 2000 and 2001. Thus a cooler summer in 2001, as actually occurred, would make it possible for a household to receive a rebate without changing its behavior from what would have occurred without the program.³¹ Second, there is the issue that large reductions in average household electricity use had already occurred, absent any price changes, during the spring of 2001 (as shown in Figure 7, above). This fact makes aggregate tabulations of year-over-year consumption changes between the summers of 2000 to 2001, as reported in Goldman et al. (2002) and others, difficult to interpret.

Table 5 and Figure 7 present information that disentangles these factors. The first column of Table 5 shows the fraction of all San Diego-area households that qualified for a rebate by reducing consumption 15% or more from the same month of 2000. For comparison, we also provide the corresponding figures for pre- and post-program months when no rebates were offered. The data indicate that on average about 39% of residential bills qualified for a rebate during the program months of June through September.³² The average rebate amounted to \$11. When we examine the distribution of consumption changes, there is no clustering of accounts on either side of 15%—that is, we do not observe a disproportionate number of accounts just qualifying for the rebate, nor just missing the rebate threshold.

The second column of Table 5 shows the fraction of households that would have qualified for a rebate in the absence of consumption decreases attributable to weather alone. We obtain these figures by calculating, separately for each household and each month, the actual consumption change between 2000 and 2001 less the change attributable to weather and trend alone. The June through September average is four percent below the unadjusted average in the first column, suggesting that the weather had a modest impact on the fraction of households that qualified. Only in August does the weather appear to have greatly aided households' chances of qualifying.

³¹One manager for Southern California Edison noted in late 2001, "I got a 20 percent rebate during the summer and I did absolutely nothing. I just never had to turn on my air conditioner." (Chan (2001)).

³²According to CPUC (2002b), approximately 36% of SDG&E residential customers qualified for a rebate. The report is unclear as to how this figure is computed. The same report states that total rebate expenditures in the SDG&E service area were \$16 million. Our data imply aggregate SDG&E-issued rebates of \$18 million.

The final two columns of Table 5 provide information about the aggregate effect of the program on household consumption. Column (3) shows households that qualified for a rebate reduced consumption by 35% relative to the same months of pre-crisis years, adjusted for weather and net of trend. Overall, column (4) shows the aggregate effect for all households was (understandably) smaller, with average household consumption during program months typically 11 to 14% below what we would expect on the basis of consumption behavior prior to the crisis.

To gauge the program's direct effect, however, it is important to make comparisons across different months in the rows of Table 5. These show that in April and May 2001, prior to the program's commencement, average household consumption ran 8 to 9% below comparable pre-crisis levels; this increased to 11% in June, then 13 to 14% for July through September. Thus we view the best estimate of this program's overall consumption impact as the decrement from spring pre-program levels, or approximately 4 to 6% of average pre-crisis household consumption.

The important point to note here is that, despite wide (and ultimately legitimate) concerns about this program's potential free-rider problem, it evidently did reduce average household consumption significantly. A consumption reduction of 4 to 6% is a big effect, particularly given it was incremental to an 8 to 9% decrease from pre-crisis consumption levels prior to the program's start. Much as we have suggested that the 12 to 13% decrease in average consumption during the price spike of 2000 required substantial practical changes in how households use their appliances, it appears that households were again willing and able to make similar-magnitude changes in response to a combination of pecuniary and non-pecuniary incentives during 2001.

In contrast to the price spike, however, the effects of these voluntary conservation programs appear transitory. Figure 7 shows that average household consumption increased dramatically in late September and October 2001, after the conservation rebate program ended. (The state's broader media campaign had effectively ended by this time as well). Consumption in October 2001 through March 2002, which is the end of our sample data, averages approximately 6% below pre-crisis levels. This is far above the average 13% below pre-crisis level during the rebate program in Table 5. The consumption increase in the fall of 2001 exceeds the magnitude of the program's effect, reversing much of the consumption decline in the spring of 2001 during the height of media attention to the crisis. In short, while the unannounced price spike in 2000 had a persistent but small effect after the price spike was rescinded (see again

the price cap effect discussion in Section 5.2), it seems clear the effects of the two voluntary energy conservation programs were largely transitory.

Program Cost Estimates

Another perspective on consumption behavior during the conservation rebate program is from the viewpoint of program cost effectiveness. Table 6 estimates how much the state effectively “paid” for reductions in household electricity consumption during summer 2001. We use the household-level estimates of electricity consumption changes underlying the figures the final column of Table 5 to calculate the program’s cost effectiveness.

Column (1) in Table 6 shows the average value of the rebate paid to households that qualified for it, by program month. Column (2) reports the fraction of households that qualified, as shown in Table 5. We then report two measures of the change in household electricity consumption during program months. Column (3) is the average consumption change in KWh for all households, relative to the same months of pre-crisis years (1998 and 1999) adjusted for weather and trend. This statistic reveals the cumulative reduction in average household electricity consumption relative to pre-crisis behavior. (The time-series behavior of this statistic is shown in more detail in Figure 7). The program-effect column (4) is calculated beginning with the value in column (3) but then deducting the average reduction in consumption (relative to the pre-crisis period) already achieved in the neighboring non-rebate months of April, May, October and November. Column (4) is interpretable as the incremental reduction in average consumption during program months.

The final column in Table 6 reports program cost estimates, calculated as the ratio of the total rebates paid to the incremental reduction in average consumption during program months. These are probably the most accurate estimates available for the average cost to the state of California of its conservation rebate program (excluding administrative costs) on a dollars-per-“KWh saved” basis.³³ The figures here indicate that the program’s average cost (in cents per KWH) was above the average price of electricity to consumers at the time, as expected. Nevertheless, these figures are below the average price of power on wholesale markets during the spring of 2001,

³³These calculations ignore the administrative costs of running the program. We have been unable to obtain program administrative cost data.

when the program was developed.

Of broader interest, however, is the fact that the magnitude of consumers' responses to this incentive program are not dissimilar to their behavior during the price spike one year earlier. In the aggregate, the incremental reduction in monthly consumption during the rebate program amounted to approximately 24 KWh per household. This reduction lowered average household electricity expenditures by approximately \$4.30 per household. By comparison, the 12% average household consumption reduction during the price spike corresponds to approximately 62 KWh, for an average expenditure savings of \$13.70. While comparisons of these figures reflect consumer average, rather than marginal, willingness to pay, the ratio of aggregate expenditure "savings" to electricity consumption foregone are in line with one another: \$0.18 per KWh for the rebate program and \$0.22 per KWh during the price spike. This is notable because both programs, although different in economic structure, provide similar measures of consumers' aggregate willingness to substitute away from electricity use given a pecuniary incentive to do so.

7 Conclusion

This paper analyzes unique household-level electricity consumption data during a period when households saw unprecedented retail price changes and were subject to notable conservation programs. Our results are relevant to energy policy discussions and also to economic questions about whether and how a large population of consumers might react to short-term price and non-price incentives to modify consumption behavior.

Our findings are of particular interest to debates about the merits of reforming retail electricity pricing—especially those that envision making prices more responsive to the short-term balance of electricity supply and demand. These proposals vary from having "real-time" prices, where a consumer's price might vary hourly, to proposals that change consumers' prices over longer windows (on a day's notice, at set times each day, et cetera). The benefits of such flexible pricing systems depend critically on consumers' willingness to curtail electricity consumption in response to high prices; absent a significant consumer behavioral response, such proposals incur costs with little efficiency gains. At present, policy-makers have limited direct evidence on how a large and diverse population of residential consumers might respond under such

proposals (see CPUC (2002)). Unfortunately, this makes it difficult to dissuade the view that consumers do not (or cannot) respond to electricity price changes at all. As recently as April 2002, economist and former Federal Energy Regulatory Commission commissioner Charles Stalon, speaking about the San Diego summer 2000 price spike, remarked “[We still] can’t say whether we saw any significant response on the demand side or not.”³⁴

The data from San Diego examined in this paper present compelling evidence to the contrary: Average household consumption fell twelve percent over a short span of approximately sixty days, in response to an *unannounced* price increase. To be clear, our results are only suggestive (at best) of what might happen if price changes of this magnitude were promulgated or scheduled differently, or changed with greater frequency. Nevertheless, the evidence from San Diego shown here should be sufficient to dissipate (at least in principle) any lingering views among policy-makers that consumers do not respond to short-run electricity price changes.

Two aspects of our results also carry noteworthy policy implications. The extraordinary heterogeneity in consumers’ responses to the price spike surely reflects both variation in households’ marginal valuations of electricity, and heterogeneity in how much attention people pay to electricity prices, their bills, and energy use. The data suggest most of the aggregate consumption decline we observed was attributable to a small share of all households. This fact suggests a cost-effective way of introducing new pricing mechanisms for making supply and demand adjust more efficiently is to identify these price-responsive consumers from the data, and target the new pricing system accordingly.

The findings regarding California’s initiatives to reduce consumption during 2001 are also of interest. California’s effort to alter individual consumption behavior, particularly through media attention, has provoked considerable attention because strategies that successfully reduce electricity use without directly taxing it (or otherwise raising price) are of great interest to energy and environmental policy makers. We find that San Diego-area households did reduce their consumption significantly during the height of the state’s energy crisis and public pleas for energy conservation. In contrast to consumer’s responses to price changes, however, the effect of such appeals proved transitory.

³⁴Hand (2002), p. 16.

Appendix A. Entry and Attrition Bias Corrections

Our sample of $n = 70,000$ residential billing accounts was drawn by simple random sampling from the roughly 1 million SDG&E residential electric accounts in March 2001. The entry or attrition of approximately one thousand sample accounts per month before and after that date poses both standard and non-standard statistical issues. Little and Rubin (2002) provide a detailed discussion of statistical methods designed to correct for attrition in longitudinal data.

A conventional approach to correcting for the effect of attrition is to adjust the uniform random-sample weights for non-uniform attrition. For example, suppose there were n_{cm} households in cohort c in the original sample of 70,000 households. If we had n_{cm} households in every month m , we would estimate cohort c 's month m population average electricity consumption by $\sum_{i=1}^{n_{cm}} \omega_{icm} y_{icm}$, where y_{icm} is household i 's consumption in month m and $\omega_{icm} = 1/n_{cm}$. In our case, with the exceptions of February and March 2001, we observe consumption for $n_{cm}^* < n_{cm}$ households. Because of this, we replace $\omega_{icm} = 1/n_{cm}$ with

$$\omega_{icm} = \frac{\frac{1}{\phi_{icm}}}{\sum_{i=1}^{n_{cm}^*} \frac{1}{\phi_{icm}}}$$

where ϕ_{icm} is the probability household i in cohort c is not missing in month m . To compute these sampling weights, one needs to estimate the attrition probabilities ϕ_{icm} . This is typically done by estimating a discrete dependent variable model for the probability that an observation is missing. We follow this approach.

We first matched our sample households to Census 2000 census tracts so that we could obtain demographic and other information that would help predict when we would lose a household from the sample. We match each account to its census tract using the account service address' nine-digit zip-code from the HERBS data.³⁵ A 2000 Census tract in the San Diego MSA contains approximately 1,600 households.

We model an observation as missing or non-missing using a probit specification. The covariates that predict the likelihood of an observation being missing are: indicator variables for the decile into which the household's average electricity consumption falls in February 2001, an indicator for whether the household has natural gas service, an indicator whether the household is on a low-income electricity tariff, an indicator for whether the account-holder is an SDG&E employee, an indicator for whether the household's electric service as billed by a third-party when that option was available, the fraction of households renting in the account's census tract, the median age of household heads in the account's census tract, the median household income in the account's census tract, the median number of rooms per housing unit in the account's

³⁵For a few, mostly ex-urban accounts, we do not have the last four zip code digits and instead map the account to a central town Census tract.

census tract, the fraction of housing units that are detached single-family dwellings units in the account's census tract, the fraction housing units with natural gas heating in the account's census tract, the median structure age in the account's census tract, the fraction of population that is urban in the account's census tract, and the median house value in the account's census tract.

Our estimated probit coefficients vary by time. In practice, this means we estimate a separate missing-observation probit model for each cohort-month (i.e., each end billing date) in our sample. This yields a total of nearly 1000 probit models. Thus, the entry/attrition modeling is quite flexible in that we allow the effect of a household's covariates to vary by cohort and by calendar month.³⁶

While it is not easy to summarize the many-thousand probit coefficient estimates, the results primarily accord with intuition. Most of the afore-mentioned covariates are statistically significant predictors of missing observations. For example, a greater fraction of renters in the census tract is associated with a higher chance of missing data. Similarly, accounts in lower consumption deciles are significantly more likely to have missing account information as one moves backward and forward from March 2001. As noted in the text, we find that adjusting for these effects significantly impacts inferences about consumption levels but does not appreciably alter average year-over-year consumption change measures.

In constructing our entry-/attrition-bias correction weightings, we estimate separate models for the availability of data on levels and the availability of data in twelve-month differences used in the regression analyses. Thus, we calculate separate weights for reporting averages of consumption levels and averages of consumption changes. The need for these weights arises because late-stage entrants into the sample (row (3) in Table 1) are treated as missing for purposes of the regression analyses, but are not missing when computing average consumption levels (reported in Table 2).

³⁶Because the probits are estimated by day, each probit is for approximately 2600-4000 households whose bills closed (or, absent attrition, would have closed) on a particular day.

References

- [1] Acton, J. P., B. M. Mitchell, and R. S. Mowill (1976). Residential Demand for Electricity in Los Angeles: An Econometric Study of Disaggregated Data. Santa Monica, CA: Rand Corporation, Report R-1899-NSF.
- [2] Blumstein, Carl, Lee S. Friedman, and Richard J. Green (2002). "The History of Electricity Restructuring in California." University of California Energy Institute, Center for the Study of Energy Markets Working Paper #103.
- [3] Borenstein, Severin (2002). "The Trouble with Electricity Markets: Understanding California's Restructuring Disaster." *Journal of Economic Perspectives* 16(1): 191-211.
- [4] Borenstein, Severin, James Bushnell, and Frank Wolak (2002). "Measuring Market Inefficiencies in California's Restructured Wholesale Electricity Market." *American Economic Review* 92(5): 1376-1405.
- [5] Bushnell, James and Erin Mansur (2002). "The Impact of Retail Rate Deregulation on Electricity Consumption in San Diego." University of California Energy Institute, Program on Workable Energy Regulation Working Paper #82.
- [6] Caves, Douglas W., and Laurits Christensen (1980). "Econometric Analysis of Residential Time-of-Use Electricity Pricing Experiments." *Journal of Econometrics* 14: 287-306.
- [7] Chan, Terry (2001). *20:20 Vision: California's Cash for Conservation Strategy*. Address to the Association of Energy Service Professionals, October 15, 2001. Available at <http://www.aesp.org/public/articles/20paper10.15.01.htm> (accessed July 1, 2003).
- [8] CPUC (2002). *Dynamic Tariffs for Residential and Small Commercial Customers*. Report of Working Group 3, Final Version 5 (December 10, 2002), Demand Response Rulemaking (R.02-06-001). San Francisco, CA: California Public Utilities Commission.
- [9] CPUC (2002b). *Resolution E-3770*. San Francisco, CA: California Public Utilities Commission.
- [10] Engle, Robert F., C. W. J. Granger, John Rice, and Andrew Weiss (1986). "Semiparametric Estimates of the Relationship Between Weather and Electricity Sales." *Journal of the American Statistical Association* 81(394): 310-20.
- [11] EPRI (1983). *Weather Normalization of Electricity Sales*. Palo Alto, CA: Electric Power Research Institute, Report EA-3143.
- [12] EPRI (1989). *Residential End-Use Energy Consumption: A Survey of Conditional Demand Estimates*. Palo Alto, CA: Electric Power Research Institute, Report CU-6487.

- [13] Ham, John C., Dean C. Mountain, and M.W.L. Chan (1997). "Time-of-Use Prices and Electricity Demand: Allowing for Selection Bias in Experimental Data." *Rand Journal of Economics* 28(0): S113-41.
- [14] Goldman, Charles A., Joseph H. Eto, and Galen L. Barbose (2002). *California Customer Load Reductions During the Electricity Crisis: Did They Help Keep the Lights On?* Berkeley, CA: Lawrence Berkeley National Laboratory, Energy Analysis Department, Report LBNL-49733.
- [15] Joskow, Paul L. (2001). "California's Electricity Crisis." NBER Working Paper #8442.
- [16] Joskow, Paul L. and Edward Kahn (2002). "A Quantitative Analysis of Pricing Behavior in California's Wholesale Electricity Market During Summer 2000: The Final Word." *Energy Journal*, 23(4): 1-35.
- [17] Lafferty, Ronald, David Hunger, James Ballard, Gary Mahrenholz, David Mead, and Derek Bandera (2001). *Demand Responsiveness in Electricity Markets*. Washington, D.C.: Federal Energy Regulatory Commission, Office of Markets, Tariffs and Rates.
- [18] Little, Roderick and Donald Rubin (2003). *Statistical Analysis with Missing Data*, 2nd Edition. John Wiley & Sons: New Jersey.
- [19] Lutzenhiser, Loren (2002). *An Exploratory Analysis of Residential Consumption Survey and Billing Data: Southern California Edison, Summer 2001*. Sacramento, CA: California Energy Commission, Report 400-02-006F.
- [20] Parks, Richard W., and David Weitzel (1984). "Measuring the Consumer Welfare Effects of Time-Differentiated Prices." *Journal of Econometrics* 26(1-2): 35-64.
- [21] Parti, M. and C. Parti (1980). "The Total and Appliance-Specific Conditional Demand for Electricity in the Household Sector." *Bell Journal of Economics* 11: 309-321.
- [22] Reiss, Peter C., and Matthew W. White (2001). "Household Electricity Demand, Revisited." NBER Working Paper #8687.
- [23] White, Matthew W. (1996). "Power Struggles: Explaining Deregulatory Reforms in Electricity Markets." *Brookings Papers on Economic Activity, Microeconomics*. 201-250.
- [24] Wilson, Robert A. (2002). "Architecture of Power Markets." *Econometrica* 70(4): 1299-1340.
- [25] Wolak, Frank (2003). "Lessons from the California Electricity Crisis." University of California Energy Institute, Center for the Study of Energy Markets Working Paper #110.

TABLE 1
SAMPLE INFORMATION, ENTRY, AND ATTRITION
(Standard deviations in parentheses)

Sample Condition	Number of Households	Average Kilowatt-hours Consumed, March 2001 ^a	Average Consumption Growth (in percent) ^b	
			March 1999 to March 2000	March 2001 to March 2002
Total households sampled, March 2001 (sample draw date)	70,000	473 (67)	4.3 (6.2)	-4.4 (5.4)
<i>Households observed beginning:</i>				
(1) October 1997 (sample start date)	37,872	514 (73)	3.0 (5.6)	-5.2 (5.6)
(2) Between Oct. 1997 and Dec. 1998 (entrants)	7,428	460 (60)	6.5 (6.1)	-3.5 (6.3)
(3) Between Dec. 1998 and March 2001 (entrants)	24,700	415 (58)	n.a. ^c	-2.3 (3.6)
<i>Households observed until:</i>				
(4) Between March 2001 and April 2002 (attrition)	14,034	399 (56)	4.5 (6.8)	n.a. ^c
(5) April 2002 (sample end date)	55,966	492 (70)	4.3 (6.1)	-4.0 (5.2)

Notes. The sample extends from October 1997 to April 2002. Available billing histories start at October 1997 if the household began electric service before this date, and end at April 2002 if the household continued service after this date.

^aAverage household consumption for the March 2001 billing period, normalized to 30 days.

^bGrowth rates expressed as a percent of average consumption in March 2001.

^cNot applicable. This figure cannot be calculated for these households.

TABLE 2
 SAN DIEGO REGION SUMMER WEATHER
 AND ELECTRICITY CONSUMPTION, 1998-2001

Billing-Month	1998	1999	2000	2001
Monthly Cooling Degree-Days ^a				
June	15	20	76	59
July	137	127	153	158
August	261	153	250	168
September	277	126	207	164
October	82	142	123	121
<i>All 12 months</i>	797	654	857	735
Average Daily Electricity Consumption (per household, in kilowatt-hours)				
June	13.7	14.2	15.4	13.3
July	15.1	15.6	16.2	14.2
August	18.2	16.5	17.0	14.5
September	19.3	16.3	16.2	15.1
October	14.9	16.0	15.4	14.7
<i>All 12 months</i>	16.0	16.1	16.1	15.0

Notes: Figures reported are averages over households for bills ending in the month shown (see text).

^aMonthly cooling degree-days measure the number of degrees by which the average daily air temperature exceeds 65°F, summed over the course of a month.

TABLE 3
SUMMARY ESTIMATION RESULTS FOR THE STABLE-PRICE PERIOD

<i>Distribution of Parameter Estimates Across Households</i>						
Coefficient	Units	Mean	Std. Error	<u>Percentiles</u>		
				25th	50th	75th
<i>Raw Parameter Estimates</i>						
Constant, Pre-10/99	Δ kwh / day	0.69	(0.13)	-0.70	0.38	1.78
Constant, Post-10/99	Δ kwh / day	0.50	(0.06)	-0.98	0.35	1.93
Heating, Winter	Δ kwh / Δ hdd	0.21	(0.01)	-0.19	0.10	0.50
Heating, Spring/Fall	Δ kwh / Δ hdd	0.19	(0.01)	-0.18	0.10	0.46
Cooling, Summer	Δ kwh / Δ cdd	0.76	(0.04)	-0.04	0.34	1.22
Cooling, Spring/Fall	Δ kwh / Δ cdd	0.47	(0.04)	-0.22	0.21	0.92
Adjusted R^2		0.35		0.10	0.39	0.64
<i>Entry-/Attrition-Adjusted Estimates^a</i>						
Constant, Pre-10/99	Δ kwh / day	0.64	(0.12)	-0.66	0.35	1.65
Constant, Post-10/99	Δ kwh / day	0.45	(0.05)	-0.92	0.31	1.78
Heating, Winter	Δ kwh / Δ hdd	0.20	(0.01)	-0.18	0.09	0.48
Heating, Spring/Fall	Δ kwh / Δ hdd	0.18	(0.00)	-0.17	0.09	0.43
Cooling, Summer	Δ kwh / Δ cdd	0.71	(0.04)	-0.05	0.31	1.10
Cooling, Spring/Fall	Δ kwh / Δ cdd	0.44	(0.03)	-0.22	0.19	0.85
<i>Interpretable Effects (Using Adjusted Estimates)</i>						
Trend Consumption Growth						
Pre-10/99	% / year	3.27		-5.48	2.99	12.11
Post-10/99	% / year	2.50		-7.58	2.77	13.34
Consumption Elasticity with Respect to Weather						
In Summer (per % Δ cdd)		0.31		-0.03	0.18	0.63
In Winter (per % Δ hdd)		0.11		-0.10	0.05	0.26

Notes. Summary results for 46,783 individual household regressions. The dependent variable is the household's change in average daily electricity consumption, in kilowatt-hours, for the same billing-month between two consecutive years (*i.e.*, 12-month differences). Not all regressions estimate a pre-10/99 constant or off-season heating or cooling coefficients.

^aDistribution of parameter estimates after re-weighting to correct for sample entry and attrition. See text.

TABLE 4
HOUSEHOLD CONSUMPTION CHANGES
BETWEEN SUMMER 2000 AND PRE-CRISIS YEARS,
BY BILLING COHORT
(Adjusted for weather and trend)

Billing Month	Cohorts	Consumption Change (Percent) ^a	Average Price (Cents/KWh)
June	1 - 7	-0.9	12.0
June	8 - 14	-1.4	12.5
June	15 - 21	-0.9	14.2
July	1 - 7	-0.5	17.9
July	8 - 14	-0.4	19.6
July	15 - 21	-2.9	20.0
August	1 - 7	-6.5	22.0
August	8 - 14	-9.7	22.9
August	15 - 21	-11.8	23.2
September	1 - 7	-12.7	13.7
September	8 - 14	-12.4	13.3
September	15 - 21	-10.1	13.0
October	1 - 7	-7.5	13.3
October	8 - 14	-5.3	13.3
October	15 - 21	-4.2	13.3

Notes: Figures reported are averages over households in the cohorts listed for bills ending in the month shown (see text).

^aAverage change in consumption from the same months of pre-crisis years, adjusted for weather and trend (see text). Changes expressed as a percentage of average pre-crisis consumption.

TABLE 5
 THE 2001 CONSERVATION REBATE PROGRAM:
 QUALIFYING HOUSEHOLDS AND CONSUMPTION CHANGES

Month	<i>Fraction of Households Qualifying</i>		<i>Average Consumption Change from Pre-Crisis Period^b</i>	
	Actual (%)	Adjusted ^a (%)	Qual. Households (%)	All Households (%)
<i>Pre-Program:</i>				
April	32	38	-32	- 9
May	36	40	-30	- 9
<i>Program Months:</i>				
June	42	39	-31	-11
July	41	43	-36	-13
August	41	30	-34	-13
September	31	27	-41	-14
June-September	39	35	-35	-13
<i>Post-Program:</i>				
October	28	31	-31	- 8
November	30	32	-26	- 7

Notes.

^aAdjusted entries show the fraction of households that *would still have qualified* in the absence of consumption changes attributable to weather alone.

^bAverage change in consumption from the same months of pre-crisis years, adjusted for weather and trend (see text). Changes expressed as a percentage of average pre-crisis consumption.

TABLE 6
 THE 2001 CONSERVATION REBATE PROGRAM:
 COST ESTIMATES FOR ENERGY “SAVED”

Program Month	Average Rebate ^a (dollars) (1)	Households Qualifying (percent) (2)	<i>Average Consumption Reduction, All Households (KWh)^a</i>		Program Cost Estimate (cents per KWh) (1) × (2) ÷ (4)
			Cumulative ^b (3)	Program ^c (4)	
June	10.43	42	54	15	29
July	11.41	41	66	26	18
August	12.06	41	64	24	20
September	10.95	31	69	30	11
June-September	11.03	39	63	24	18

Notes.

^aFigures standardized to a 30-day billing month in this table for comparability.

^bAverage change in consumption from the same months of pre-crisis years (1998 and 1999), adjusted for weather and trend.

^cColumn (3) less the weather- and trend-adjusted reduction in average consumption (relative to the pre-crisis period) already achieved in the adjacent non-program months of April, May, October and November.

**Fig. 1. Average Residential Electricity Price, 1998–2001
(Excludes Taxes and Rebates)**




Fig. 2. Electricity Consumption and Weather, 1998–2001




Fig. 3. Predicted and Actual Average Within-Household Consumption Changes During the Stable-Price Period




Fig. 4. Average Within-Household Consumption Changes Relative to the Same Months During Pre-Crisis Years, Weather and Trend Removed




Fig. 5. Cross-Section Distribution of Households' Year-Over-Year Consumption Changes (Weather and Trend Removed)




Fig. 6. Households' Consumption Changes and Levels
(Changes are Aug. '00 v. Aug. '98 and '99, Weather and Trend Removed)




Fig. 7. Average Within-Household Consumption Changes Relative to the Same Months During Pre-Crisis Years, Weather and Trend Removed

