RAND Journal of Economics
Vol. 39, No. 3, Autumn 2008
pp. 636–663

What changes energy consumption? Prices and public pressures

Peter C. Reiss*

and

Matthew W. White**

Policymakers often seek to limit energy prices following market shocks, and instead issue public appeals to reduce demand. This article presents new evidence on how price changes and conservation appeals affect energy consumption, using household-level data from California's energy crisis during 2000 and 2001. The evidence indicates that when policymakers cap energy prices following market shocks, they preclude substantial—and quite rapid—reductions in energy use. The data also reveal that conservation appeals and informational programs can produce sustained reductions in energy demand.

1. Introduction

■ Markets for energy often experience shocks that require rapid changes in consumption or prices. Examples abound, from the 1970s oil shocks to California's electricity crisis and recent supply disruptions in gasoline markets. Efficiency requires that consumers quickly respond to higher prices with attenuated demand, thereby averting shortages or rationing. In practice, however, political decision makers are reluctant to let the price mechanism run its course. Instead, they will often hold prices below new market-clearing levels when laws and regulations enable it, and issue public appeals for voluntary conservation in order to mediate demand.[1]

Do these political responses to market shocks exacerbate shortages, or are they a useful strategy to help avoid them? To answer this question requires evidence on whether—and how promptly—consumers react to (i) price shocks and (ii) public pressures. Recent surveys have

---

* Stanford University and NBER; reiss_peter@gsb.stanford.edu.
** University of Pennsylvania and NBER; matthew.white@wharton.upenn.edu.

We thank San Diego Gas and Electric Company for access to nonpublic billing data for this research, and staff at the California Public Utilities Commission for assistance and comments. Two anonymous referees provided valuable suggestions. White gratefully acknowledges financial support from the University of Pennsylvania Research Foundation and the Russell Ackoff Endowment Fund in the Wharton Risk Management and Decision Processes Center. The opinions and conclusions expressed in this article 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.

[1] Typical price-limiting mechanisms include state anti-"price gouging" statutes (used in Midwestern states after pipeline-related gasoline supply limitations), and public utility codes that prevent suppliers from raising prices (such as during California's electricity crisis).

drawn attention to the fact that, despite more than 30 years of market interventions in response to energy price shocks, policymakers still possess limited information regarding their consequences (IEA, 2005a, 2005b). The lack of compelling evidence makes it difficult to weigh the effects of restraining prices and issuing public appeals when market shocks occur.

This article provides unique evidence on how energy consumption responds to both price shocks and public appeals. The analysis is based on seldom-available utility billing data containing the energy consumption histories for tens of thousands of households in the San Diego region during California's energy crisis. This population experienced an unprecedented increase in electricity prices during 2000, and mass public appeals for energy conservation the following year after prices were (legislatively) rolled back.[2]

The data reveal striking changes in households' consumption habits. After electricity prices unexpectedly (and rapidly) increased in 2000, the average household's electricity use fell more than 13% in about 60 days. As we explain, this requires the typical household to invest in new appliances or make substantial behavioral changes in how often they are used. The latter was especially widespread: after the legislature imposed a (binding) cap on residential electricity prices, consumption rapidly rebounded toward former levels. We infer that when policymakers limit price increases following supply shocks, they preclude substantial and fairly rapid changes in energy consumption behavior.

We then consider the effectiveness of public appeals. To curb excess demand under the new price cap and avoid mandatory rationing (or "rolling blackouts"), California initiated a media campaign to promote voluntary conservation. In essence, this amounted to a massive public appeal to forgo consumption. The typical San Diegan's energy use declined steadily by 7% over this six-month period—absent any pecuniary incentive to do so. The magnitude of this response indicates that well-orchestrated mass public appeals can be an effective means to avert rationing when the price mechanism is unable (or, in this instance, not permitted) to equilibrate the market.

These findings provide new and useful information along three lines. First, although the literature on energy demand is voluminous, little work examines how quickly consumption adjusts to price shocks. This gap in the literature makes it difficult to refute conventional wisdom in regulatory and policy arenas, which have long regarded prices as a minimally effective instrument for influencing energy use.[3] The use of accurate (metered) energy consumption data at the household level is invaluable in this regard, as it enables stark conclusions about the speed and magnitude of demand changes by individual consumers.

Second, public appeals by government officials have a long history in times of energy crises and for other social objectives. President Jimmy Carter urged Americans to reduce oil use with Oval Office broadcasts, California issues "Spare the Air" declarations on smoggy days that encourage people to refrain from driving, municipalities ask citizens to conserve water during droughts (yet often leave water prices unchanged), and local utilities broadcast appeals for consumers to shut off air conditioners in hot weather when electricity supplies are tight. Issuing public appeals like these is like soliciting anonymous contributions to a public good: respondents incur private costs individually, yet achieve tangible benefits only if aggregate participation is high. Despite considerable experimental research on public goods problems, there remains little direct evidence on whether public appeals such as these are effective in practice.[4]

---

[2] Related studies of aggregate energy use during California's crisis include Goldman, Eto, and Barbose (2002) and Bushnell and Mansur (2005). Survey work by Lutzenhiser (2002) provides complementary information on households' behavior during this time.

[3] For example, the price-setting process used by utility regulatory commissions in nearly all 50 states formally assumes that electricity demand has zero sensitivity to price (see Friedman, 1991). We relate our findings to prior research in considerable detail below.

[4] Ledyard (1995) surveys experimental research on public good contributions. Existing field evidence on public appeals includes a handful of energy market studies (surveyed in IEA, 2005a, 2005b), a small literature on water-conservation programs (Michelson, McGuckin, and Stumpf, 1998; Nieswiadomy, 1992) and smog-reduction appeals (Cutter and Neidell, 2006), and studies of public appeals by charitable organizations (e.g., List and Lucking-Reiley, 2002; Croson, 2005). Stern (1992) summarizes studies of energy-conservation motives.

Finally, establishing that households can quickly change their consumption habits for electricity is useful because it is harder for consumers to adjust consumption for this service than for most other goods (electricity is considered the most inelastically demanded form of energy, for example). Thus, our findings are suggestive that consumers could do at least as much following market shocks for other goods (such as gasoline, where substitutes such as carpooling and public transportation exist), but where persistent measurement limitations preclude results as precise as the findings documented here (IEA, 2005b).

Our study design is based on five years of disaggregate billing data for a random sample of 70,000 households. We compare within-household changes in average daily electricity consumption during a $2^{1/2}$-year pre-crisis period to consumption changes during the price shock in 2000 and the voluntary conservation campaign in 2001. To disentangle the confounding effects of weather, we carefully matched each sample household to daily weather data using 21 weather stations in the region. The within-household analysis allows us to identify, net of weather-related effects, the variation in consumption behavior that followed changes in prices and public appeals. Because of the staggered design of households' billing periods in the San Diego region, the data reveal a great deal of information about the timing of consumers' responses to these events.

The next section provides background information describing the events we analyze. 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: prices and policy responses

■ 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.[5] There is little variation in this price series across households, except that 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: (i) a period of basically stable prices, from January 1998 to May 2000; (ii) the price spike in summer 2000; and (iii) the period under the price cap beginning September 2000. These changes, along with the 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. Other than a slight downward adjustment in January 1998 (visible at the far left in Figure 1), average prices were essentially flat during the 1990s. 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.[6] The anticipated effect was that competition

---

[5] The price in Figure 1 is total revenue per kilowatt-hour (KWh) sold to residential consumers, based on a large random sample of bills (described 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. SDG&E's residential tariff had a two-tier (or block) structure through 2001, with a small price difference (two cents per KWh) between the two tiers. The nonlinearity of the tariff accounts for the slight wiggles in the average price series.

[6] Readers desiring more detailed background on electricity market restructuring in California may consult Blumstein, Friedman, and Green (2002), Joskow (2001), Wilson (2002), or Wolak (2005).

FIGURE 1

AVERAGE RESIDENTIAL ELECTRICITY PRICE, 1998–2001

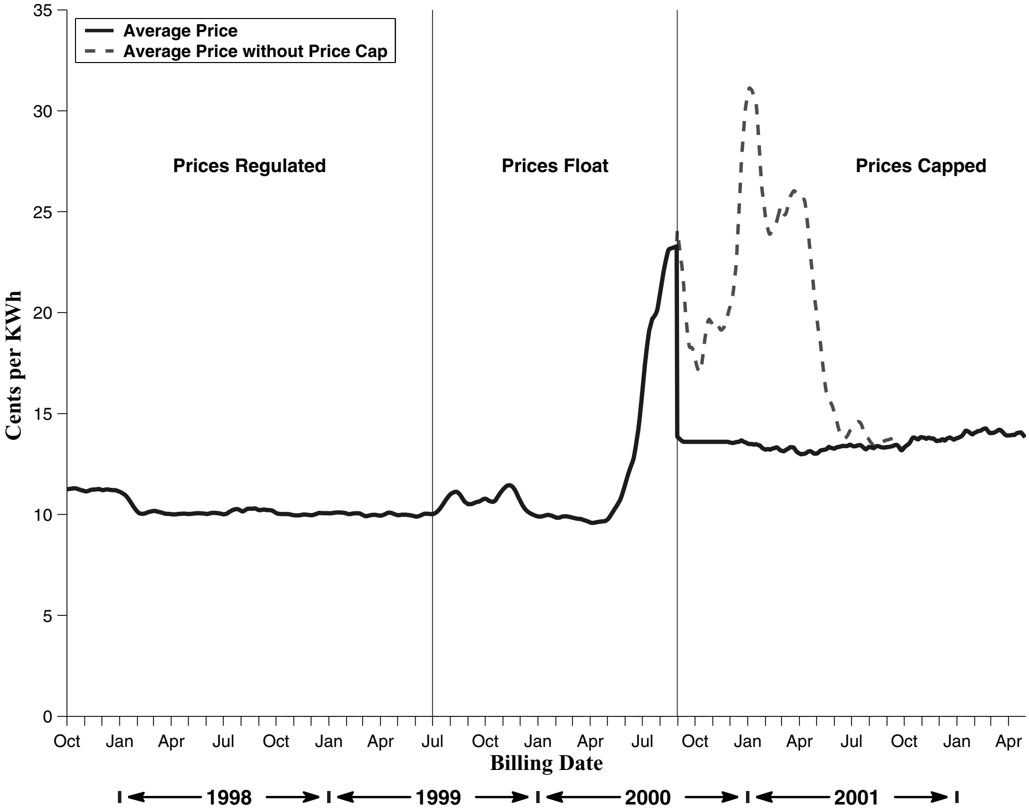

Source: Excludes special rebates (see text).

in wholesale markets would, over time, drive down retail prices below their previously regulated levels (see White, 1996; 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 state's electricity crisis not intervened.

During its first year, the change to this new system of electricity pricing was immaterial 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 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 $2^{1/2}$ years 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.[7] Because San Diego residential electricity prices were now indexed to the wholesale market, these price increases promptly appeared in

---

[7] The 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 Borenstein (2002), Borenstein, Bushnell, and Wolak (2002), Joskow and Kahn (2002), or Wolak (2005).

households' monthly bills. San Diego households' 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.

Because price changes under this system were not announced (or even known) in advance, they comprised a true "price shock" when consumers received their June 2000 bills. Continuing increases the next month generated a storm of public protest and, ultimately, 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.[8]

We know a great deal about when information regarding these price changes became available to San Diegans. There is no reason to believe that a typical consumer anticipated the June 2000 price spike. In the third week of June, local media began reporting in business news segments the rapid escalation in wholesale market prices (which now determine monthly bills). By the first week of July, media coverage of the price spike appeared prominently on front pages of regional newspapers and in local newscasts. By then, the increase in prices would have been evident to most any consumer through his or her last utility bill. The legislative price cap first received general media coverage in the third week of August 2000. Its enactment on September 6, 2000 was highly publicized.

Consequently, a household could learn that prices were changing either through its utility bill or through the media. From a research standpoint, we will quantify when households react to price shocks given this concurrent media coverage. That is, we do not attempt to identify what the impact of a large energy price change would be in a "media vacuum" that did not report on it at the time.[9]

□ **Post-cap interventions.** Although imposing a binding price cap quickly 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 policymakers had no authority, now exceeded the capped retail price by a considerable margin (see 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 SDG&E's now-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 2000–2001.

Rather than increase retail prices and face the same political uproar again, state leaders attempted to prevent such rationing by reducing energy consumption with other means. Starting in February 2001, state agencies and utilities undertook a major campaign to urge conservation and educate people about how to reduce energy use (we summarize its implementation in Section 6). These mass public appeals for voluntary energy conservation, and the deepening electricity crisis, received prominent media coverage throughout the spring and early summer of 2001.

California's effort to alter individual consumption behavior through its conservation campaign provoked considerable attention and initial skepticism. 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 policymakers. Yet it is an open question whether promoting (voluntary) electricity conservation has much of an effect if consumers' economic incentives to modify their behavior remain unchanged.

---

[8] 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.

[9] It is difficult to imagine how an energy price shock might occur without media coverage, although smaller unannounced price changes might not draw the same media attention and thus slower consumer responses. See also Peck and Doering (1976).

TABLE 1 Data on Households and Turnover

Sample Condition	Number of Households	Average Kilowatt-Hours Consumed March 2001[a]	Average Consumption Growth (in %)[b] March 1999 to March 2000	Average Consumption Growth (in %)[b] March 2001 to March 2002
Total households sampled,	70,000	473	4.4	−4.6
March 2001 (sample draw date)		(1.4)	(0.1)	(0.1)
Households observed:				
(1) October 1997	37,845	514	3.0	−5.4
(sample start date)		(1.5)	(0.1)	(0.1)
(2) Entering from November 1997	7,363	464	6.7	−4.0
to November 1998		(1.3)	(0.1)	(0.1)
(3) Entering from December 1998	24,792	415	n.m.[c]	−2.65
to March 2001		(1.2)		(0.1)
(4) Exiting from March 2001	13,689	398	4.5	n.m.[c]
to April 2002		(1.1)	(0.1)	
(5) April 2002	56,311	492	4.3	−4.1
(sample end date)		(1.4)	(0.1)	(0.1)

Notes: Standard errors for subsample averages are shown in parentheses.

[a] Average household consumption for the March 2001 billing period, normalized to 30 days.

[b] Growth rates expressed as a percent of average consumption in March 2001.

[c] Not meaningful for this subsample.

To summarize, our agenda is to assess how household electricity consumption responded to three successive events: the price increase during the summer of 2000, the price cap in early fall 2000, and the state's public appeals for conservation and concurrent crisis attention during the spring and early summer of 2001.

3. Data

■ We examine these events using the Household Electricity Research Billing Sample (HERBS), a large, five year panel of San Diego Gas and Electric Company residential utility bills. 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 $3^{1/2}$ 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 1000 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.

One consequence of this sample mobility is that we cannot estimate consumption responses to crisis events for late-stage sample entrants. This group comprises most households 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 have sufficient data to assess year-over-year changes in household consumption.

Sample-entry/attrition issues are a concern for our analysis only 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 in order 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. The principal reason 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 (columns (4) and (5) of Table 1) are small in comparison to typical household consumption responses we estimate to electricity crisis events.[10]

☐ **Timing issues and cohorts.** An important feature of our data is the timing of consumers' bills, which reveal how quickly consumers reacted to price changes and other events. SDG&E uses a cohort-based (or staggered) billing system, wherein each household is permanently 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 the month, with the bill arriving a few days later.[11]

These timing issues affect the way we examine the data in two ways. First, we measure a household's electricity use by its *average daily consumption* in each billing period. This adjusts for variation in the number of days in each billing period, which differs slightly across cohorts each month. Second, we present households' data by billing cohort when describing consumption changes over time. Despite the fact that any single bill averages consumption over the last 30 (or so) days, the staggered billing system enables us to observe changes in households' average consumption (across cohorts) at nearly a daily frequency.

☐ **Weather data.** One limitation of utility billing data is that they do not include household-specific weather information. This might not pose much of a problem if households experienced the same weather. In SDG&E'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 temperatures. Because home electricity use is highly sensitive to weather, it is important to account for this variation to avoid misspecifying the conditions a household actually experienced.

---

[10] This was not our initial expectation. Our concerns about entry and attrition biases motivated considerable effort to address this issue, detailed in Appendix A.

[11] Occasionally, a meter cannot be read on the scheduled cycle-closing day, but is read a few days later. The 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. In addition, because 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. Because 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.

TABLE 2 San Diego Region Summer Weather and Electricity Consumption, 1998–2001

Billing-Month	1998	1999	2000	2001
	Cooling Degree-Days[a]			
June	14	18	69	52
July	133	124	147	147
August	244	137	228	152
September	289	126	201	167
October	72	127	109	110
Annual	778	613	800	690
	Average Daily Electricity Consumption (per household, in kilowatt-hours)			
June	14.2	14.5	15.5	13.3
July	15.7	16.0	16.3	14.1
August	18.9	16.9	17.1	14.5
September	20.1	16.7	16.2	15.2
October	15.4	16.3	15.4	14.8
Annual	16.5	16.4	16.2	15.0

Notes: All figures are entry- and attrition-weighted averages across households (see Appendix A).

[a] Cooling degrees are the number of degrees (F) by which the midpoint of the minimum and maximum daily air temperature exceeds 65°F. Cooling degree-days sum each household's cooling degrees over the billing month, then average across households.

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 in the region. This process matches a household to a local weather station considering both proximity and elevation. 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 daily temperature data and the exact start and end dates for each bill. The weather variables are billing-period heating and cooling degree-days, which are highly predictive measures of home heating and cooling energy demand.[12]

During our study period, weather conditions in San Diego County varied greatly from year to year. For example, the top panel in Table 2 reveals that cooling degree-days in August and September 2000—during the energy crisis—were more than 60% higher than the same months in 1999. This variation is primarily attributable to the El Niño Pacific weather disturbance in 1998 and 1999.[13]

Most importantly, this variation makes direct comparisons of aggregate energy consumption levels between the electricity crisis period and preceding years problematic. For instance, although prices more than doubled between August 1999 and August 2000, the lower panel of Table 2 reveals that average household electricity consumption actually *increased* slightly. The increase in consumption during a price spike is attributable to the offsetting influence of the weather, as indicated in the top panel of Table 2. Thus, in order to identify changes in household behavior due to energy crisis events, it is essential to disentangle them from the confounding effects of weather.

---

[12] A household's cooling degree-days from 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 on day t (in °F) at the household's local (matched) weather station. High values occur in the summer, reflecting high demand for cooling services; cooling degree-days are usually zero in winter. Heating degree-days are computed with the summand replaced by $\max\{65 - A(t), 0\}$. Heating degree-days are high in winter months, and usually zero (or nearly so) in summer. All weather data are from the National Oceanic and Atmospheric Administration's Climate Prediction Center daily temperature archives (cpc.ncep.noaa.gov).

[13] These monthly averages mask considerable variation in cooling and heating degree-days across households. For instance, in August 2000, households located along the coast experienced less than 150 cooling degree-days, while those far inland averaged over 700.

4. Empirical strategy and methods

■ Our empirical methods are designed to overcome two key problems in estimating consumers' responses to the electricity crisis. First, there is the problem that everyone in San Diego experienced the same events at essentially the same time (Figure 1). Thus, from a research design standpoint, there is no way to construct the equivalent of two randomly assigned "treatment" and "control" groups of households.

The second issue relates to price information. We do not know what price each consumer *expected* to pay during the summer and fall of 2000. As prices rose, perceptions of what the floating price would be each billing period (surely) differed across consumers—depending, for instance, on how attuned a household was to media coverage of the unfolding electricity crisis. The upshot is that it is difficult to correctly specify a statistical model that assumes, implicitly or explicitly, what the consumer thought its price would be each month (see also Bushnell and Mansur, 2005).

Our approach addresses these issues 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 and no public appeals occurred. The second is the crisis period, during which (i) prices fluctuated and were subsequently capped (in 2000), and (ii) the public appeals for voluntary conservation occurred (in 2001). This three-part sequence allows us to use variation in consumption during the stable-price, no-appeals period to disentangle the effects of some confounding factors (weather, in particular) during the two later periods. In essence, we are using a household's consumption during the period when prices were held fixed and appeals never occurred as the "control" (and implicit counter-factual), and contrasting this with consumption following each crisis-related event.

□ **Modeling household responses.** For transparency, consider first the price changes in 2000. Inference regarding the subsequent conservation campaign in 2001 requires only minor modifications, noted subsequently.

We analyze consumption changes using the difference in a household's average daily consumption (ADC) between one billing period and the same billing period 12 months earlier. 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.[14]

We divide the variation in a household's 12 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.[15] Averaged across all households, things in this first category define the *secular trend* in residential electricity consumption. By contrast, the second category includes nonidiosyncratic factors that induce correlation in consumption changes among households. These include regional weather patterns, common price changes, crisis media attention, and the (subsequent) voluntary conservation program.

We assume that a household's same-month, year-over-year electricity consumption changes can be decomposed into its pre-crisis trend, crisis-period responses, weather, and idiosyncratic

---

[14] From 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 (2005). Because these factors are unobserved in billing data, here we seek to explain variation in consumption within, rather than between, households.

[15] We 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.

factors according to

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

where x_{icm} is the ADC of household i in cohort c for the billing period ending in month m (indexed serially), w_{icm} is a vector of contemporaneous weather-related covariates, and ε_{icm} is an idiosyncratic shock. Here I_m^{pre} is an indicator variable for pre-crisis months (May 2000 or earlier), and I_m^{post} indicates 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 household and by month in order 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 by month m is the difference (or break) from its pre-crisis trend,

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

Ultimate interest lies not in the individual values of D_{icm} but in its population distribution and the time-path of the average

$$\overline{D}_{cm} = \frac{1}{N_{cm}} \sum_{i=1}^{N_{cm}} D_{icm},$$

where N_{cm} is the population cohort size.

Some comments regarding specification are in order. Our use of a linear-in-the-parameters specification in (1) is based upon a large literature (principally in statistics and engineering) on the relationship between weather and electricity sales. A well-known paper in this area is Engle et al. (1986), who examine parametric and nonparametric techniques for modelling 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.[16] For this reason, we prefer the simple and parsimonious specification in (1) rather than less-efficient nonlinear smoothing techniques. We have also estimated more flexible specifications with our data that accommodate nonlinear weather effects on consumption. These yield *de minimus* differences in the results presented below. (For additional details, see Appendix B.)[17]

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

□ **Estimation details.** The specification in (1) implies a household's sensitivity to weather is time invariant. In contrast, a natural response to the rise in prices during the summer of 2000 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. Fortunately, it is straightforward to accommodate such weather-dependent responses empirically. We estimate the difference D_{icm} for each household without imposing this time-invariance restriction using the following procedure.

To estimate the parameters in (1), we first fit to the (up to) 32 months of pre-crisis data the model

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

---

[16] These slopes vary widely across utilities; see especially Engle et al. (1986). Such geographic differences reflect the varying prevalence of residential air conditioning and electric space heating in different parts of the country.

[17] One might also anticipate some degree of serial correlation in ε_{icm}, due to (for example) the lumpiness of durable-appliance replacement timing. In the data, we find little evidence of within-household serial correlation in ε_{icm}. Note that serial correlation poses only an efficiency issue here; a small loss of efficiency is of negligible concern in these data.

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,

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

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

For periods more than 12 months after the crisis began, interest still centers on how consumption differed from its (weather-adjusted) pattern during pre-crisis years. Thus, we modify (3) slightly when m is later than May 2001, computing $\widehat{D}_{icm}$ by replacing $x_{ic,m-12}$ and $w_{ic,m-12}$ with their values at 24 month-lags and doubling $\hat{\delta}_i^{pre}$. This implies $\widehat{D}_{icm}$ for (say) m of August 2001 is relative to the pre-crisis month of August 1999, not the crisis-period month of August 2000.

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 number of observed 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, it is statistically desirable to pool estimates from more than one year of pre-crisis data. Thus, in estimating consumption changes during the first 12 months of the crisis, we use a pooled difference estimator of the form

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

where $\widehat{\overline{D}}{}^{(24)}_{cm}$ 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 $\hat{\delta}_i^{pre}$ doubled).[18] Pooling these two estimates reduces variability attributable to idiosyncratic departures from a household's pre-crisis trend in any one year. The pooled difference estimator $\widehat{\overline{D}}{}^{p}_{cm}$ 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. This is the statistic we report in the time-series results below.

□ **Identification issues.** Identification and interpretation of δ_{icm}^{post} and $\overline{D}_{cm}$ merit brief discussion. Technically, an estimate of $\overline{D}_{cm}$ 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 energy crisis-related events is justified under two assumptions. These are (i) that the population's average pre-crisis trend, $\bar{\delta}^{pre}$, would have continued during 2000–2001 had the crisis not occurred; and (ii) 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 main findings.

Although identifying assumptions are by nature not directly testable, these two can be addressed informatively. Assumption (ii) 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% per annum (p.a.)) for nearly a decade. The effect of a gradual increase of this magnitude on our results is negligible;

---

[18] Similarly, we use 24 and 36 month lags in constructing the pooled difference estimator for any month more than 12 months into the crisis period (that is, after May 2001).

FIGURE 2

PREDICTED AND ACTUAL AVERAGE WITHIN-HOUSEHOLD CONSUMPTION CHANGES DURING THE STABLE-PRICE PERIOD

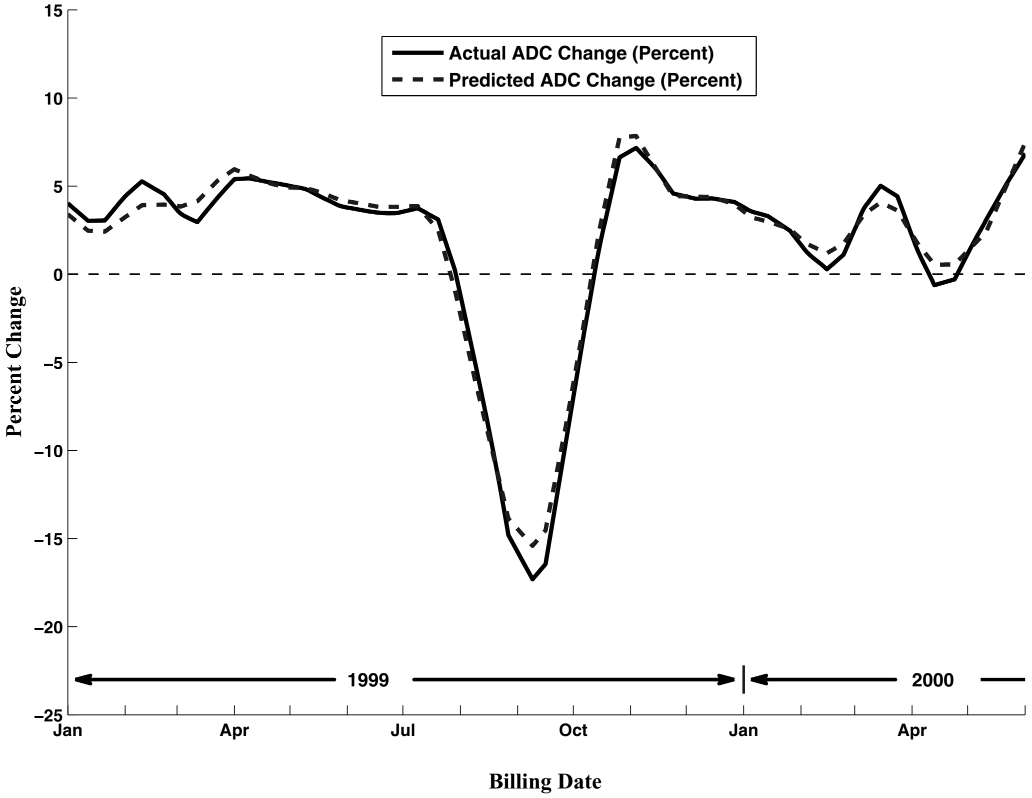

it would take more than a double-digit increase in the percent of homes with air conditioning to discernibly affect our estimate of the time-path of consumption changes, $\overline{D}_{cm}$. An increase of that magnitude over one to two years is not a realistic concern.

At issue with assumption (i) 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 (see Figure 2, discussed next). Any substantial unobserved factor would have to become influential in 2000 or 2001, but not be similarly influential during preceding years. Because the average pre-crisis secular trend in electricity consumption represents the aggregation of hundreds of thousands of individual household decisions about appliance purchases, vacations, remodelling, and so on, it is difficult to conceive that this trend would have changed abruptly in 2000–2001 absent a macroeconomic shock.

One potential point of concern is income. Real personal income per household grew at 3.5% annually from 1997 through 2000 in San Diego, but began to slow in late 2000 and declined slightly in 2001. Because estimates of the short-run income elasticity of residential electricity demand are low, on the order of 0.1 (Hsiao and Mountain, 1985), this change is unlikely to reduce household electricity consumption in 2001 by more than several tenths of 1%. Nevertheless, we address the sensitivity of our 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 pre-crisis consumption history to fit equation (2). The top panel

TABLE 3 Summary Estimation Results for the Stable-Price Period

Distribution of Parameter Estimates across Households

Coefficient	Units	Mean	Standard Error	Percentiles 25th	Percentiles 50th	Percentiles 75th
Household Parameter Estimates						
Constant, pre-10/99	Δ kwh/day	0.63	(0.04)	−0.58	0.36	1.63
Constant, post-10/99	Δ kwh/day	0.60	(0.02)	−0.46	0.34	1.43
Heating, Winter	Δ kwh/Δ hdd	0.21	(0.01)	−0.16	0.10	0.47
Heating, Spring/Fall	Δ kwh/Δ hdd	0.21	(0.00)	−0.14	0.11	0.46
Cooling, Summer	Δ kwh/Δ cdd	0.77	(0.04)	−0.03	0.33	1.17
Cooling, Spring/Fall	Δ kwh/Δ cdd	0.45	(0.03)	−0.09	0.23	0.79
Adjusted R^2		0.34		0.14	0.35	0.55
Entry-/Attrition-Adjusted Estimates[a]						
Constant, pre-10/99	Δ kwh/day	0.62	(0.04)	−0.58	0.36	1.62
Constant, post-10/99	Δ kwh/day	0.64	(0.01)	−0.47	0.35	1.4
Heating, Winter	Δ kwh/Δ hdd	0.21	(0.01)	−0.16	0.10	0.47
Heating, Spring/Fall	Δ kwh/Δ hdd	0.21	(0.00)	−0.14	0.11	0.46
Cooling, Summer	Δ kwh/Δ cdd	0.77	(0.03)	−0.03	0.33	1.15
Cooling, Spring/Fall	Δ kwh/Δ cdd	0.43	(0.02)	−0.10	0.23	0.80
Adjusted R^2		0.33		0.14	0.34	0.55
Interpretable Effects (Using Adjusted Estimates)						
Trend Consumption Growth						
Pre-10/99	%/year	3.54		−3.29	2.08	9.27
Post-10/99	%/year	3.68		−2.71	2.00	6.46
Consumption Elasticity with Respect to Weather						
Summer (per %Δ cdd)		0.39		−0.01	0.11	0.44
Winter (per %Δ hdd)		0.17		−0.09	0.05	0.30

Notes: Summary results for approximately 46,800 individual household regressions. The dependent variable is the household's change in average daily electricity consumption, in kilowatt-hours, for two billing periods 12 months apart. The heating and cooling degree-day variables are changes in average daily degree-days. Not all regressions estimate a pre-10/99 constant or off-season heating or cooling coefficients.

[a] Distribution of parameter estimates after re-weighting to correct for sample entry and attrition. See text.

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 for entry/attrition, yielding only minor changes. Although the weather parameters' degree-day units are difficult to interpret intuitively, their means and medians are in accord with prior work (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 cooling degree-days between one summer month and the same month of the previous year implies a 39% increase in household electricity consumption. The analogous consumption sensitivity in winter is about one half as large.[19]

The secular trend in household consumption averages between 3.5% and 3.7% per annum during the pre-crisis period. Table 3 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, which suggested a change about this time. The

---

[19] Table 3 also reveals that there is considerable variation in the estimated responses to the weather variables, with some households' estimated effects being small or negative. Although it is possible that some households reduce consumption in response to increases in heating or cooling degree-days, our preferred interpretation is that the response of many households is truly zero and that negative coefficients reflect idiosyncratic factors (e.g., travel, a major 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, 2005) indicates that households differ substantially in their response to weather, and a nontrivial fraction of households exhibit no consumption response to temperature.

estimated difference in the mean trend before versus after October 1999 is small, however, and the medians are nearly identical. We use the post-October 1999 trend parameters in the estimates of $\overline{D}_{cm}$ for June through December 2000 reported below.

An important issue is how well our parsimonious model predicts consumption in the absence of price changes and other crisis events. Figure 2 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 household's change by its average daily consumption during this period). The actual and predicted series are close, despite the dramatic variation in average consumption and weather over this period. The overall R^2 for the average actual and predicted data summarized in Figure 2 is 0.98. Importantly, the model accurately predicts the substantial 20% drop in average consumption between August/September of 1998 and 1999. It does equally well with the 10% increase in consumption in October 1999. The weather to come during the crisis period (2000–2001) falls well within the range of variation observed during these pre-crisis years.

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 2 is that year-over-year changes in average household electricity consumption are basically *all* weather induced, absent a change in prices or a major conservation campaign. It is the stability of this tight relationship over time that we rely upon to draw inferences about how households responded to the crisis-related events.

5. The price spike and price cap

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

□ **Results.** Figure 3 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 is the pooled difference estimator $\widehat{\overline{D}}^{p}_{cm}$ in (4), calculated relative to households' 1998 and 1999 consumption levels and smoothed nonparametrically for clarity.[20] This 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.

During the stable-price period, the line shown in Figure 3 is effectively zero. The absolute error represented by the small fluctuations before June 2000 averages less than 1%; this much is to be expected from Figure 2. By contrast, Figure 3 shows that during the price spike in summer 2000, average residential electricity consumption declined precipitously. By the end of August, San Diego's million households consumed 13% less electricity than in prior years, net the effect of weather and relative to their pre-crisis trend. Given that bills reflect consumption over the preceding 30 or so days, Figure 3 implies this 13% drop in average consumption occurred in a short span of about 60 days.

Figure 3 further reveals that the decline in average consumption ends nearly as abruptly as it began. This change is first observed in bills closing in late September, indicating that households began to reverse course in their consumption behavior by the middle of that month. Figure 3 reveals a rapid rise in average consumption from late September through October. We refer to this as the *rebound effect.* This effect is coincident with the imposition of the (retroactive) price cap in early September, discussed below.

---

[20] Here and throughout, we use a loess (locally estimated scatterplot smoothing) procedure to present time-series results. The standard errors for these averages are extremely small, as each billing cohort averages over 2000 households.

FIGURE 3

AVERAGE WITHIN-HOUSEHOLD CONSUMPTION CHANGES DURING THE 2000 PRICE SPIKE AND PRICE CAP

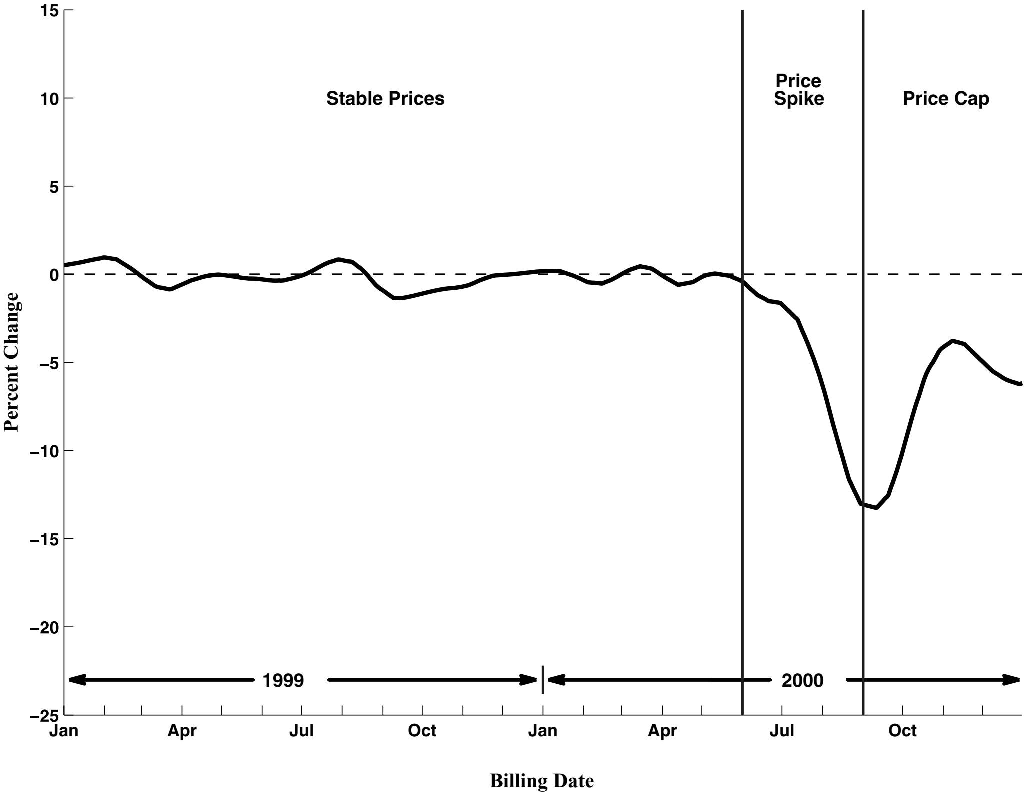

Source: Changes are relative to the same month of prior years, with weather and trend removed.

Interpreting response magnitudes. Figure 3 presents a compelling case that San Diego's one million households responded to the abrupt rise and subsequent fall in electricity prices during 2000. Moreover, they did so fairly rapidly. To grasp the implications of these results, it is valuable to first interpret the magnitudes: does a 13% drop in electricity consumption require a big change in behavior, or a small one?

A 13% reduction in electricity use requires considerable and deliberate efforts 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 for the San Diego region indicate that air conditioning accounts for about 15% of aggregate residential electricity consumption during the summer.[21] Only steep reductions in air conditioner use could therefore account for the observed 13% drop in aggregate residential electricity consumption—and this during an exceptionally warm summer. Alternatively (or in combination), 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 clotheslines), 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

---

[21] Derived from EPRI (1989), which is based on appliance sub-metering studies of San Diego households performed by SDG&E.

television time by one half, then *combined* these five changes could amount to a 10% reduction in average electricity consumption.[22] One can continue with similar such examples for less electricity-intensive appliances, but the picture here is clear. Because a large share of 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 effects 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 resulting electricity savings at San Diegans' summer 2000 electricity prices.

Interpreting response timing. It bears emphasis that the speed of these changes reflects both how quickly consumers are willing to change consumption habits and how quickly they notice prices are changing. Specifically, the fact that consumption shows a much slower rate of decline before versus after mid-July suggests that consumers were not widely aware of the rapid rise in prices before then. The steep consumption decline after mid-July is consistent with consumers (in all billing cohorts) making significant consumption adjustments by then, and continuing to do so through early September. The rate of decline in Figure 3 should thus be viewed as a result of consumer learning over time about the price shock, and progressively greater changes in consumption habits as prices continued to rise.

The rebound effect of the price cap. The rapid rebound of consumption in the fall of 2000 might not be considered surprising, given the rate at which consumers reduced demand when prices rose during the summer. Figure 3 shows that after the price cap was imposed, average household consumption (adjusted for weather) increased approximately 8% from mid-September to mid-October. Just as they did following the earlier price increases, consumers reacted—here with less than a one month billing lag—to decreases in their electricity prices.

One feature to note here is that the response to the price cap in the fall of 2000 may have been influenced 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 from mid-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 responses, one might expect a complete rebound in the fall. The demand rebound following the price cap leaves average consumption about 4% below the pre-spike level, however, and it never exceeds that level (on a weather-adjusted basis) subsequently.

One interpretation of this fact is that it suggests one third of consumers' aggregate response to the price spike was realized through changes in appliance stocks, dwelling improvements,

---

[22] Based on appliance-level electricity consumption and utilization data from the U.S. Department of Energy (EIA, 1997).

FIGURE 4

CROSS-SECTIONAL DISTRIBUTION OF HOUSEHOLDS' CONSUMPTION CHANGES

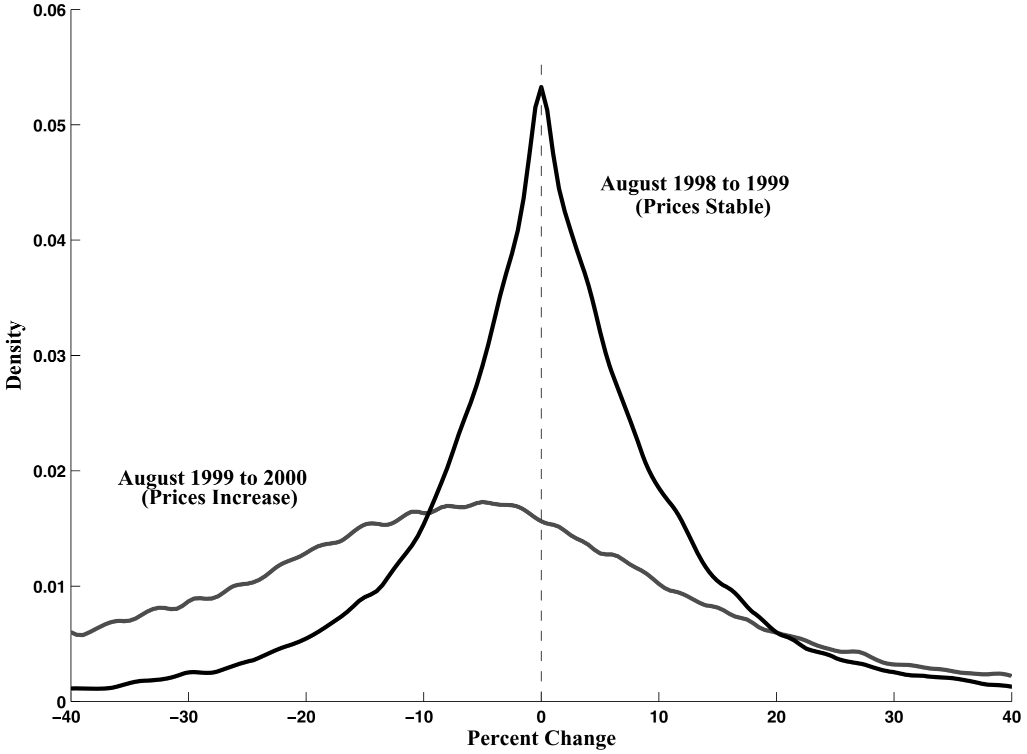

Source: Changes are from August to August of each year, with weather and trend removed.

or persistent changes in utilization decisions. This gives some quantitative credence to the self-reported evidence on appliance replacement 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 longer-term impacts on energy consumption.

□ **Response heterogeneity.** The household-level data reveal considerable heterogeneity in households' consumption responses. This heterogeneity has noteworthy implications for conservation incentives and ongoing pricing debates.

Are consumption changes widespread? Figure 4 shows the cross-sectional distribution of individual households' consumption changes between August 1999 and August 2000, which spans the period when prices rose. For comparison, this figure also shows the analogous distribution for August 1998 to August 1999, when prices did not change. Both distributions are based on each household's consumption change adjusted for weather and relative to pre-crisis trend.

Figure 4 reveals that the distribution of consumption changes between August 1998 and 1999, the stable price period, is nearly symmetrically distributed about zero. In contrast, the distribution of changes between August 1999 and 2000 is quite skewed. Interestingly, the 1999-to-2000 distribution shows a considerable fraction of households made large consumption changes, in *both* directions. On the right-hand side, 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 35% of households, which is down from 50% during the prior

FIGURE 5

SCATTERPLOT OF HOUSEHOLDS' CONSUMPTION LEVELS VERSUS

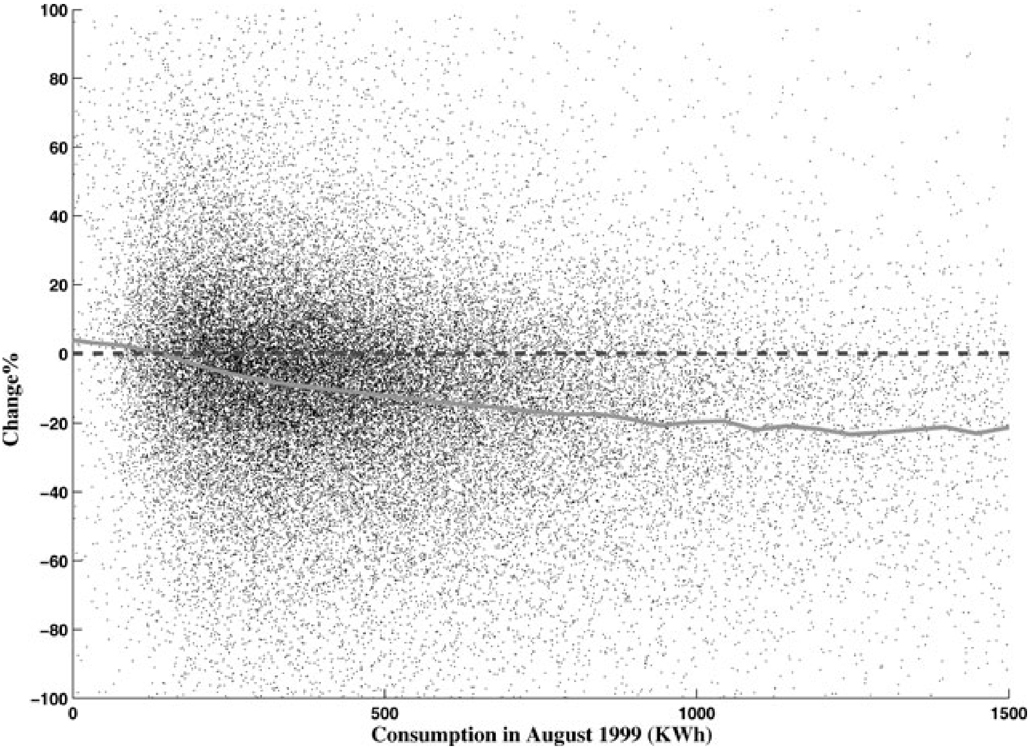

Source: Levels are for August 1999. Changes are between August 2000 and August of pre-crisis years, with weather and trend removed.

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 (Reiss and White, 2005). 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.

At the opposite end, there is a large share of households that show extraordinary reductions in electricity consumption. Between August of 1999 and 2000, approximately 1 in 3 households reduced electricity use by 20% or more (relative to their pre-crisis trend and adjusted for weather). For the prior year's period of August 1998 to 1999 when no price changes occurred, the directly comparable value is only 1 in 14 households. Thus, a substantial share of the decline in average (and aggregate) consumption following the price shock is attributable to a minority of the population who reduced their electricity use dramatically. Consumption reductions of 20% or more require careful and attentive efforts to lower electricity use, major investments in new appliances, or substitution behavior that one suspects would pose considerable inconvenience within the household.

High- versus low-demand consumers. Figure 5 shows the percent change in households' electricity consumption between August 2000 and August of prior years, as a function of households' consumption *levels* (kilowatt-hours in August 1999). Each point in the scatter is an individual household. (The figure omits a small number of outliers with consumption over 1500 KWh per month.) The change shown in the figure is the household's actual change less the change predicted by pre-crisis trend and weather alone; in statistical terms, this is the household's pooled difference estimate $\widehat{D}_{icm}^{p}$ for August 2000. The solid line in the figure is a (nonparametric) line of averages, indicating the mean consumption response at each level of consumption.

Figure 5 reveals the enormous unexplained variance—recall it is already net of weather's influence—in year-over-year consumption changes across households at all consumption levels. Despite this dispersion, the change in consumption is clearly negative at the mean (of approximately 500 KWh). Further, it slopes downward fairly steeply. The mean consumption change is approximately -25% for households using 1000 KWh or more per month, which is typical for a large single family dwelling with air conditioning. In contrast, the mean consumption response to the dramatic rise in prices is zero for households that use 200 KWh or less. That level is typical of small households (one or two members) residing in apartments without major electricity-using appliances beyond a refrigerator.[23]

To our knowledge, there is no prior evidence on this relationship. We note it here because the empirical relation between consumption levels and price sensitivity informs a number of ongoing welfare and regulatory policy questions.[24] First, this is essential information for evaluating the consumption and expenditure consequences of new tariff designs, which vary prices based on household consumption. This proved to be important in California, where policymakers changed tariff designs for 10 million households after the energy crisis. Because the state's regulatory authorities did not have any information on the relationship shown in Figure 5 at the time—and errantly assumed it to be flat—they significantly mis-estimated the ensuing changes in consumption and utilities' revenues. Reiss and White (2005) analyze this in detail.

Second, over the past decade, economists have pointedly emphasized the value of enhancing the "demand response" of energy markets—that is, increasing the ability of consumers to adjust consumption quickly to varying electricity prices. As a practical matter, this requires a significant investment in new pricing programs, metering hardware, systems to communicate price information, and the like. The data in Figure 5 suggest that the most cost-effective way to pursue this goal is to target these programs selectively to the subset of consumers with high consumption levels. These households are only a small share of the population, are easily identified in utilities' billing data, and account for most of the aggregate residential demand response when prices change.

□ **On price elasticities and prior research.** It is perhaps not surprising to many economists that the typical household would adjust its energy consumption in response to price changes even over short-term durations. But this fact has not been widely recognized by policymakers, 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 energy use of their household appliances, and will modify that use substantially in response to pecuniary incentives.

One reason these debates persist is that few (if any) prior studies document how quickly households respond to energy price shocks (IEA, 2005a, 2005b). Instead, the literature obtains energy demand elasticities using other forms of price variation. For household electricity consumption, there are two: cross-sectional survey data and utilities' field experiments with tariff designs. The cross-sectional estimates vary widely, with typical values (conditional on existing appliance stocks) near -0.3.[25] Purely cross-sectional based estimates have limited appeal in many prospective applications, however, because they convey no information about the speed with which demand adjusts to price changes.

---

[23] By comparison, a similar plot (not shown here) of August 1998 to 1999 changes versus August 1998 levels—during the stable price period—shows mean weather- and trend-adjusted consumption changes at all consumption levels are nearly perfectly symmetrically distributed about zero.

[24] We thank an anonymous referee for emphasizing this point.

[25] Noteworthy studies include Parti and Parti (1980), Dubin and McFadden (1984), and Dubin (1985). More precise estimates can be obtained by using repeated cross-sections with intervening price changes and nonlinear price schedules; see Reiss and White (2005).

In contrast, energy demand elasticities at monthly or daily frequencies are sometimes extrapolated from tariff design experiments. These estimates are commonly half as large as those from cross-sectional studies.[26] However, these experiments have important limitations: most are conducted to inform alternative pricing systems (such as the introduction of time-of-use pricing), not to evaluate the effects of large changes in price levels. Acton (1982) points out that such experiments are generally unreliable for the latter purpose. Additionally, nearly all such experiments carefully inform subjects of the tariff changes they will face well before the experiment begins (an exception is CEC, 2005). This intensive consumer education process makes the informational context of pricing experiments somewhat special, relative to how consumers belatedly learn (from their bills or the media) of price changes in practice.

We note these attributes here to emphasize that contrasting the consumption responses from San Diego with most price elasticity estimates in the literature makes for apples-to-oranges comparisons. Each type of data and study design informs a different aspect of energy consumption behavior (with respect to time horizons, price change magnitudes, initial awareness of the price change, and so on).[27] Alternatively stated, the data from San Diego offer information about consumers' responses to price shocks that are not available in prior work. This information appears considerably more meaningful for assessing potential policy interventions following energy supply shocks.

A related study is Bushnell and Mansur (2005), who analyze the time series of aggregate electricity consumption in San Diego and Los Angeles during 2000. The Bushnell-Mansur study reports much smaller consumption changes than we find. Although the reasons for the difference are not entirely clear, two explanations are salient. One is that some attenuation bias should perhaps be expected when using imperfectly measured macroeconomic data to model aggregate electricity demand variation. A second likely explanation is compositional: smaller changes are apt to result from pooling residential and industrial consumers, who have different price sensitivities (and experienced different price changes) but are inseparable in aggregate consumption data.

Because we are studying the effect of a price shock here, we offer an additional word of caution against 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. Rather, one should interpret the results here as informing a large population's *60 day elasticity* (or so) of residential electricity demand with respect to an unanticipated price shock that (roughly) doubles price.

---

[26] An extensive survey can be found in the 1984 *Journal of Econometrics* (vol. 26, no. 1-2) special issues on electricity pricing experiments; see also Hausman, Kinnucan, and McFadden (1979), Acton and Mitchell (1980), Caves and Christensen (1980), and Parks and Weitzel (1984). Unfortunately, most of these experimental studies are now quite dated. Moreover, most use small samples and rely upon voluntary participation without randomization, raising attrition- and selection-bias concerns that impair generalization (see Aigner, 1984, and references therein for a critique). Ham, Mountain, and Chan (1997) suggest statistical corrections.

[27] One of the only experimental studies that offers an apples-to-apples comparison to our results is described in Acton and Mitchell (1980). In one treatment, there was a 150% increase in the (flat rate) electricity price for approximately 100 households in the Los Angeles area. The average decline in electricity consumption in ensuing months was 6.2%, relative to a control group with no price increase. San Diego households experienced a proportionally similar price increase in 2000 (of about 130%), with a consumption decline twice as large (13%). Note, however, that the prevalence of home air conditioning at the time of Mitchell and Acton's study was half that of San Diego in 2000. Given this, these consumption response magnitudes seem reasonably consistent.

6. Post-cap interventions

■ 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.

The deepening crisis was routinely aired on nightly news programs, particularly following the realization of rolling blackouts (electricity rationing) 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'. Residential prices in San Diego were held at the capped level shown in Figure 1 from the fall of 2000 through 2001.

To forestall excess demand and further electricity shortages under the price cap, the state undertook a major campaign to promote energy conservation. This campaign involved a $65 million program of advertising and outreach through television, radio, newspapers, billboards, and other outlets (e.g., public schools). The advertising had both persuasive and informational components, appealing to consumers to voluntarily conserve energy and simultaneously educating consumers about how to do so through simple changes in behavior (e.g., changing thermostat settings, unplugging second refrigerators, and so on). The campaign's major ad buys ran for about six weeks in February and March of 2001, and again in May and June. State officials also made dramatic televised appeals for consumers to voluntarily conserve electricity. Details regarding the campaign's design and implementation can be found in CDCA (2003) or Bender et al. (2002).

Did these events—the state's public appeals for voluntary electricity conservation and concurrent media coverage—alter households' consumption habits? Because prices were capped at a constant level, it would seem that consumers had little incentive individually to reduce their energy consumption. The empirical evidence suggests that the opposite was the case. Figure 6 presents the main evidence on this point. This figure extends the time series shown in Figure 3 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.[28]

Figure 6 shows a steady and sizable continuous reduction in household electricity consumption from January through June 2001. This decline amounts to approximately 7% of average pre-crisis consumption, on a weather- and trend-adjusted basis. The large drop evident in mid-March to early-April bills followed a limited, but well-publicized, rolling blackout that affected approximately 41,000 customers in the San Diego area.

□ **Public appeals and public goods.** That households' electricity consumption fell quite substantially when prices were not changing is a striking and informative empirical result. Public appeals for energy conservation by government officials have a long history in times of energy crises, to questionable effect. In economic terms, voluntary conservation is a collective-action problem subject to extreme free-rider difficulties. Tangible benefits (viz., averting a blackout or lowering the market price) accrue only if aggregate consumer participation is high, and no one can be excluded from the benefits of group success. Yet reducing energy consumption by the magnitudes here is likely to pose considerable inconvenience within the household, as noted earlier. In addition, the chance that any single household's effort might be pivotal in averting a

---

[28] We make one adjustment to our prior methods here. Because personal income growth had begun to slow by the beginning of 2001, 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 6 assuming that there was *zero* secular consumption growth after January 2001. This yields a conservative estimate of the changes in consumption during 2001. Adjusting the mean secular trend for the decline in mean personal income growth using conventional short-run income elasticity estimates for residential electricity consumption (Hsaio and Mountain, 1985) yields greater consumption declines; estimated consumption in Figure 6 would be 2–3 percentage points lower in mid-2001.

FIGURE 6

AVERAGE WITHIN-HOUSEHOLD CONSUMPTION CHANGES, 2000–2002

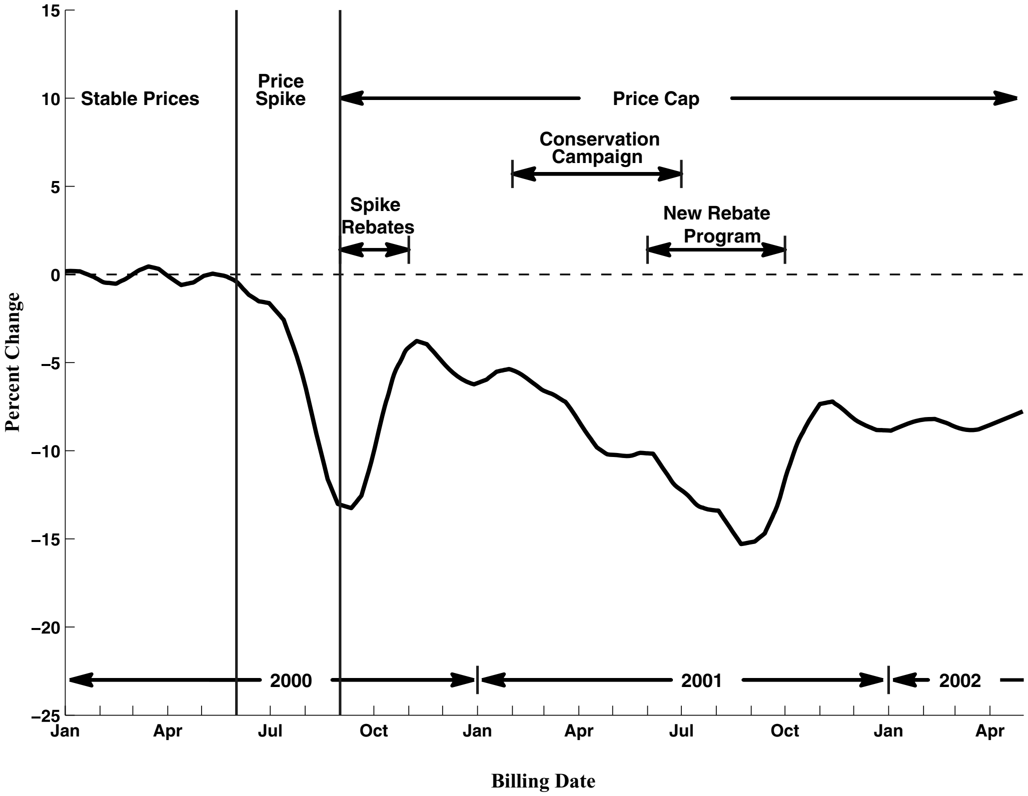

Source: Changes are relative to the same months during pre-crisis years, with weather and trend removed.

blackout (given an electrical grid the size of SDG&E's) is vanishingly small. The bottom line: each household faced private costs of reducing consumption in response to the public appeals, a virtually zero possibility of bringing about any tangible benefit with respect to the crisis through individual effort, and a considerable incentive to free-ride on whatever efforts are made by others.

In these respects, voluntarily conserving energy in response to public appeals is much like contributing anonymously to a public good. Anonymity arises because—as California households quickly discovered—there is no real way 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, pool pump operating hours, refrigerator efficiency, or the other major appliances that account for the vast majority of residential electricity consumption.[29] Thus, explanations that rely upon either social sanctions or a household's intrinsic desire not to be seen as a wasteful energy user during a crisis face a basic problem: a household's true effort to reduce electricity consumption—or not—is (essentially) visible only to that household. It is true (rather than self-reported) conservation effort that is captured in our data, and this effort is substantial.

---

[29] The one form of electricity consumption that is readily visible to others, outdoor lighting, is a negligible share of the typical household's electric bill. Even if many households were unaware of this fact, this does not help explain the actual consumption changes shown in Figure 6; as discussed in Section 5, changes of this magnitude necessitate major reductions in the use of appliances that are not observable outside the home.

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, and so forth.[30] Although economists tend to be dismissive of public appeals that 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 6. The empirical facts here indicate that consumers do respond to voluntary appeals to modify their consumption behavior, provided (i) the costs of a collective-action failure are tangible (here, involuntary blackouts for some), and (ii) the public is well aware of it.

□ **Identification issues, redux.** As discussed earlier, the observed decline during the spring and early summer of 2001 can be attributed to energy crisis events under quite credible assumptions. A more subtle concern is whether there are other possible *crisis-related* mechanisms, distinct from the voluntary conservation campaign, that might account for some of the spring 2001 consumption decline in Figure 6. Two other possible mechanisms merit brief discussion here: pecuniary incentives via appliance subsidies and other rebate programs, and the possible effect of natural gas prices on household electricity use.

Rebate programs. Did San Diego households face new pecuniary incentives to reduce energy use through changes in appliance promotion programs, and the like? A careful review of these programs indicates that this was not the case for households in San Diego during the winter of 2000 and spring of 2001. To be clear, the state (and, to a lesser extent, SDG&E) launched a large number of conservation programs that targeted energy use by businesses, institutions (e.g., universities), and government agencies in 2001.[31] But, because of various implementation delays, the state's expanded subsidy programs for household purchases of energy-efficient appliances were not launched until August 2001 (CDCA, 2003). Additionally, SDG&E did not significantly change its preexisting appliance rebates and residential energy-efficiency programs until then (SDG&E, 2002). In sum, from the time the price cap took effect in September 2000 until the end of May 2001, San Diego households' *pecuniary* incentives to reduce electricity use remained unchanged.

The first major change in San Diego households' pecuniary incentives to reduce electricity use came in June 2001, under a state conservation price-rebate program. For June through September 2001, Figure 6 reveals the combined effects of the ongoing voluntary conservation campaign *and* this new conservation price rebate program.[32] Consumption begins another sharp rebound upward in the fall of 2001, when both the media conservation campaign and the price-rebate programs were effectively over.

The "combination bill" conjecture. During the winter and early spring of 2001, natural gas prices increased significantly in California. These increases compounded the state's energy crisis and raised households' total energy bills significantly. Conceivably, households who were billed for both natural gas and electricity on the same billing statement might not have clearly distinguished between the two. A popular conjecture is that households might have cut *electricity* use in response to the increase in *natural gas* prices in the spring of 2001, rather than as a response to the state's public appeals for voluntary electricity conservation (IEA, 2005a). Such

---

[30] Survey results of self-reported motivations for conserving energy corroborate this; see Lutzenhiser (2002) and Stern (1992).

[31] CALMAC (2003) contains a comprehensive summary of the 218 energy-conservation programs run by California state agencies or utilities during 2001.

[32] For an analysis that separates these two effects, see our NBER working paper (Reiss and White, 2003).

behavior, if common, would complicate interpretation of the observed electricity consumption decline.[33]

Actual (metered) consumption data are not consistent with this conjecture, however. In San Diego County, almost one third of households do not use natural gas and thus do not receive a combined gas and electric bill (all-electric home heating is relatively common in the region). The other two thirds of households receive a combined bill. If households mistakenly reduced electricity use in response to the increase in *natural gas* prices (instead of the media campaign to reduce electricity use), then electricity use should decline differently for (i) households with combined gas and electric bills, versus (ii) households without combined bills (that is, without natural gas service). The latter group clearly should not be affected by natural gas prices.

In the data, we find little evidence of differential consumption changes between these groups. During the spring of 2001, households with combined-bill electric and natural gas service and those with electric but no natural gas service reduced their electricity consumption at the same speed, and by the same magnitude (approximately 10% each by April 2001, relative to pre-crisis levels and net the effect of weather). Our conclusion is that anecdotes from newspapers (and small focus groups) about people confusing electricity and gas prices are atypical. In metered consumption data for a large representative sample of households, there is no evidence that consumers mistook natural gas for electricity prices during California's energy crisis.

7. Final remarks

■ This article analyzes household-level energy consumption during a period when households saw unprecedented price changes and public appeals for conservation. The results are relevant to a number of energy policy discussions, and to broader questions about whether (and how) a large population reacts to pecuniary and nonpecuniary incentives to modify consumption behavior.

Our findings are particularly germane to policies that envision making consumers' electricity prices more responsive to shifts in supply and demand. The benefits of such changes depend crucially on consumers' willingness to curtail electricity consumption in response to high prices. As noted in Section 5, policymakers have limited direct evidence on how a large and diverse population of consumers might respond under such proposals (see also FERC, 2001; CPUC, 2002). Unfortunately, this makes it difficult to dissuade a prevailing view among policymakers that consumers do not (or cannot) quickly respond to price changes. Two years after San Diego's price spike, economist and former Federal Energy Regulatory Commission Commissioner Charles Stalon still asserted, "[We] can't say whether we saw any significant response on the demand side or not" (Hand, 2002). The data examined in this article present compelling evidence to the contrary. Average household consumption fell 13% over a short span of approximately 60 days, in response to an *unannounced* price increase. Because these results are obtained from actual (metered) energy use at the consumer level for a large representative sample of households, they leave little room to dispute the magnitude or timing of consumer behavior. The evidence here should be sufficient to dissipate any lingering views that consumers cannot (or will not) respond quickly to energy price changes.

Our findings regarding households' voluntary conservation responses are also significant. California's effort to alter energy consumption in this manner 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 policymakers (see esp. IEA, 2005a). The data reveal that the average household reduced its consumption significantly during the state's public appeals for energy conservation, even though it faced no pecuniary incentive to do so. The fact that public pressure of this sort works at all raises the important question of whether prices or non-price mechanisms should be the preferred means of mediating energy consumption

---

[33] Note that this combination bill conjecture does not arise because, say, only aggregated information appears on the consumer's monthly bill. The prices, monthly usage, and total bills for each form of utility service (electricity and natural gas) are printed separately, in tabular and in graphical form, on each billing statement from SDG&E.

when a market shock requires it. Although these results suggest policymakers' proclivity for the latter can be effective, a complete treatment of that issue awaits further research.

Appendix A

■ **Entry and attrition bias corrections.** Our sample of $n = 70{,}000$ residential billing accounts was drawn by simple random sampling from the approximately 1 million SDG&E residential electric accounts in March 2001. The entry or attrition of approximately 1000 sample accounts per month before and after that date poses both standard and nonstandard 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 nonuniform 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

$$\omega_{icm} = \frac{1/\phi_{icm}}{\sum_{i=1}^{n^*_{cm}} 1/\phi_{icm}},$$

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 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's nine-digit zip-code from the HERBS data.[34] A 2000 Census tract in the San Diego MSA contains approximately 1600 households.

We model an observation as missing or nonmissing 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 is 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 Census tract, the fraction of housing units that are detached single-family dwellings units in the account's Census tract, the fraction of 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 modelling is quite flexible in that we allow the effect of a household's covariates to vary by cohort and by calendar month.[35]

Although it is not easy to summarize the many-thousand probit coefficient estimates, the results primarily accord with intuition. Most of the aforementioned 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 weights, we estimate separate models for the availability of data on levels and the availability of data in 12 month differences used in the regressions. 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).

Appendix B

■ **Alternative model specifications.** This Appendix summarizes a more flexible specification of the econometric model in Section 4. The specification of weather's influence in equations (1)–(3) makes two implicit assumptions of

---

[34] For 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.

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

possible concern: (i) it treats two days at 75° F the same as a day each at 65° and 85°; and (ii) it treats the effect of a change in temperature from 65° to 80° the same as a change from 80° to 95°. We have checked the validity of both of these restrictions in the data using the following procedure.[36]

First, rewrite the specification for the weather terms in (1). In level form,

$$w'_{icm}\beta_i = w_{0icm} + \beta_i^C\, C_{icm} + \beta_i^H\, H_{icm}, \tag{A1}$$

where C_{icm} is the *average daily* cooling degree-days in (billing-) month m for household i in cohort c, and H_{icm} analogously for heating degree-days. The term w_{0icm} is a "baseline" contribution of weather to consumption when C_{icm} and H_{icm} are zero.

Let $f^C_{k,icm}$ denote the fraction of days in billing-month m for which household i's cooling degrees equals k ($k = 1, 2, \dots, K$). Define $f^H_{j,icm}$ analogously for heating degrees. Then

$$C_{icm} = \textstyle\sum_k k\, f_k^C \qquad \text{and} \qquad H_{icm} = \textstyle\sum_j j\, f_j^H.$$

The baseline term w_{0icm} can similarly be represented as a household-specific constant, β_i^0, times the fraction of days, $f_{0,icm}$, for which the average temperature is exactly 65°F. Inserting these expressions into (A1) yields

$$w'_{icm}\beta_i = \beta_i^0 f_{0,icm} + \left(\sum_{k=1}^{K} \beta_i^C\, k\, f^C_{k,icm} + \sum_{j=1}^{J} \beta_i^H\, j\, f^H_{j,icm}\right).$$

Note that the fractions f_k^C and f_j^H must satisfy the adding-up condition $\sum_{j=1}^{J} f_j + \sum_{k=1}^{K} f_k = 1 - f_0$. Substituting this for $f_{0,icm}$ gives

$$w'_{icm}\beta_i = \beta_i^0 + \left(\sum_{k=1}^{K} \left(\beta_i^C\, k - \beta_i^0\right) f^C_{k,icm} + \sum_{j=1}^{J} \left(\beta_i^H\, j - \beta_i^0\right) f^H_{j,icm}\right).$$

This expression makes explicit the two restrictions noted above, that, net of the baseline effect, two days with cooling (heating) degrees at $65 + k$ ($65 - k$) have the same effect on consumption as one day at $65 + 2k$ ($65 - 2k$).

This expression also reveals a flexible way to relax such restrictions and allow each possible degree-day outcome to have a different effect on consumption. This amounts to replacing the single cooling-weather sensitivity parameter β_i^C with a vector of temperature-specific cooling-weather parameters $\beta^C_{k,i}$, $k = 1, 2, \dots, K$, and similarly for the heating-weather sensitivity coefficients. That is, we can model the effect of weather with the richer specification

$$\begin{aligned} w'_{icm}\beta_i &= \beta_i^0 + \left(\sum_{k=1}^{K} \left(\beta^C_{k,i}\, k - \beta^0\right) f^C_{k,icm} + \sum_{j=1}^{J} \left(\beta^H_{j,i}\, j - \beta^0\right) f^H_{j,icm}\right) \\ &\equiv \beta_i^0 + f'_{icm}\tilde{\beta}_i. \end{aligned} \tag{A2}$$

This specification can be estimated by projecting average daily electricity consumption (or yearly changes in it) on the fractions of the billing period (or changes in these fractions) at various degree-day values, f'. We can then examine whether the estimated coefficients for $\tilde{\beta}_i$ satisfy the joint restrictions $\tilde{\beta}^H_{j,i} = \beta_i^H$ and $\tilde{\beta}^C_{k,i} = \beta_i^C$ for all j, k, with β_i^H and β_i^C estimated from the restricted model using A1.

The advantage of (A2) is that a sufficiently dense partition of degree-day outcomes can accurately approximate (essentially) any temperature-consumption relation as m tends large. The disadvantage of (A2) is that its implementation in our data requires restricting the heterogeneity in the coefficients across households to preserve degrees of freedom. This is the standard parameter homogeneity versus functional flexibility tradeoff with nonparametric regression. We have estimated the article's results using (1)–(3) with (A2) under several alternative homogeneity restrictions: households have the same weather-sensitivity coefficients ($\tilde{\beta}$'s) if they are in the same billing cohort, if they have the same weather station, or if both conditions hold. The results under any implementation yield little evidence to suggest nonlinear weather effects except in the extremes (degree-days exceeding 25 per day); there are very few households subject to these extremes in our sample, however. None of these alternative specifications result in material changes to the mean or median coefficient estimates reported in Table 3 or the main results in Figures 4–6.

We regard this as an informative outcome for future researchers, insofar as it suggests (or rather, confirms the existing literature's observation) that household-level weather and electricity consumption relationships can be satisfactorily modelled using the simpler degree-day metrics and specifications employed in Section 4 of this article.

References

Acton, J. "An Evaluation of Economists' Influence on Electric Utility Rate Reforms." *American Economic Review*, Vol. 72 (1982), pp. 114–119.

---

[36] Thanks to an anonymous referee for suggesting these extensions.

——— AND MITCHELL, B. *The Effect of Time-of-Use Rates in the Los Angeles Electricity Study*. Santa Monica, CA: RAND Corporation, Report no. N-1533-DWP/HF, 1980.

AIGNER, D. "The Welfare Econometrics of Peak-Load Pricing for Electricity: Introduction." *Journal of Econometrics*, Vol. 26 (1984), pp. 1–15.

BENDER, S., MOEZZI, M., GOSSARD, M., AND LUTZENHISER, L. "Using Mass Media to Influence Energy Consumption Behavior: California's 2001 'Flex Your Power' Campaign as a Case Study." In *Proceedings of the 2002 ACEEE Summer Study*, Vol. 8. Washington, DC: American Council for an Energy-Efficient Economy, 2002.

BLUMSTEIN, C., FRIEDMAN, L.S., AND GREEN, R. "The History of Electricity Restructuring in California." Working Paper no. 103, Center for the Study of Energy Markets, University of California Energy Institute, 2002.

BORENSTEIN, S. "The Trouble with Electricity Markets: Understanding California's Restructuring Disaster." *Journal of Economic Perspectives*, Vol. 16 (2002), pp. 191–211.

———, BUSHNELL, J., AND WOLAK, F. "Measuring Market Inefficiencies in California's Restructured Wholesale Electricity Market." *American Economic Review*, Vol. 92 (2002), pp. 1376–1405.

BUSHNELL, J. AND MANSUR, E. "Consumption under Noisy Price Signals: A Study of Electricity Retail Rate Deregulation in San Diego." *Journal of Industrial Economics*, Vol. 53 (2005), pp. 493–513.

CALMAC. *California Summary Study of 2001 Energy Efficiency Programs.* Sacramento: California Measurement Advisory Council, Report no. 02-1099, 2003.

CAVES, D.W. AND CHRISTENSEN, L. "Econometric Analysis of Residential Time-of-Use Electricity Pricing Experiments." *Journal of Econometrics*, Vol. 14 (1980), pp. 287–306.

CDCA. *Flex Your Power Annual Report, 2001–2002.* Sacramento: California Department of Consumer Affairs, 2003.

CEC. *Impact Evaluation of the California Statewide Pricing Pilot.* Sacramento: California Energy Commission, March 16, 2005.

CPUC. *Dynamic Tariffs for Residential and Small Commercial Customers*. Report of Working Group 3, Final Version 5, Demand Response Rulemaking R.02-06-001. San Francisco: California Public Utilities Commission, December 10, 2002.

CROSON, R. "Theories of Commitment, Altruism, and Reciprocity: Evidence from Linear Public Goods Games." Working Paper, The Wharton School, University of Pennsylvania, 2005.

CROSON, R. "Theories of Commitment, Altruism, and Reciprocity: Evidence from Linear Public Goods Games." *Economic Inquiry*, Vol. 45 (2007), pp. 199–218.

CUTTER, W.B. AND NEIDELL, M. "Voluntary Information Programs and Environmental Regulation: Evidence from 'Spare the Air.' " Working Paper, Department of Economics, Columbia University, 2006.

DUBIN, J.A. *Consumer Durable Choice and the Demand for Electricity.* Amsterdam: North Holland, 1985.

DUBIN, J.A. AND MCFADDEN, D.L. "An Econometric Analysis of Residential Appliance Holdings and Consumption." *Econometrica*, Vol. 52 (1984), pp. 345–362.

ENGLE, R.F., GRANGER, C., RICE, J., AND WEISS, A. "Semiparametric Estimates of the Relationship between Weather and Electricity Sales." *Journal of the American Statistical Association*, Vol. 81 (1986), pp. 310–320.

EPRI. *Weather Normalization of Electricity Sales*. Palo Alto, CA: Electric Power Research Institute, Report no. EA-3143, 1983.

———. *Residential End-Use Energy Consumption: A Survey of Conditional Demand Estimates.* Palo Alto, CA: Electric Power Research Institute, Report no. CU-6487, 1989.

FERC. *Demand Responsiveness in Electricity Markets*. Washington, DC: Federal Energy Regulatory Commission, Office of Markets, Tariffs and Rates, 2001.

FRIEDMAN, L. "Energy Utility Pricing and Customer Response." In R.J. Gilbert, ed., *Regulatory Choices: A Perspective on Developments in Energy Policy*. Berkeley: University of California Press, 1991.

GOLDMAN, C., ETO, J., AND BARBOSE, G. *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 no. LBNL-49733, 2002.

HAM, J., MOUNTAIN, D., AND CHAN, M. "Time-of-Use Prices and Electricity Demand: Allowing for Selection Bias in Experimental Data." *RAND Journal of Economics*, Vol. 28 (1997), pp. S113–S141.

HAND, M. "The Economists: On the Future of Energy Markets." *Public Utilities Fortnightly*, Vol. 140 (2002), pp. 12–18.

HAUSMAN, J., KINNUCAN, M., AND MCFADDEN, D. "A Two-Level Electricity Demand Model: Evaluation of the Connecticut Time-of-Day Pricing Test." *Journal of Econometrics*, Vol. 10 (1979), pp. 263–289.

HSIAO, C. AND MOUNTAIN, D. "Estimating the Short-Run Income Elasticity of Demand for Electricity." *Journal of the American Statistical Association*, Vol. 80 (1985), pp. 259–265.

IEA. *Saving Electricity in a Hurry: Strategies to Deal with Temporary Shortfalls in Electricity Supplies.* Paris: International Energy Agency, 2005a.

———. *Saving Oil in a Hurry: Measures for Rapid Demand Restraint in Transport.* Paris: International Energy Agency, 2005b.

JOSKOW, P. "California's Electricity Crisis." *Oxford Review of Economic Policy,* Vol. 17 (2001), pp. 365–388.

——— AND KAHN, E. "A Quantitative Analysis of Pricing Behavior in California's Wholesale Electricity Market during Summer 2000: The Final Word." *Energy Journal*, Vol. 23 (2002), pp. 1–35.

Ledyard, J. "Public Goods: A Survey of Experimental Research." In J. Kagel and A.E. Roth, eds., *Handbook of Experimental Economics.* Princeton, NJ: Princeton University Press, 1995.

List, J. and Lucking-Reiley, D. "The Effects of Seed Money and Refunds on Charitable Giving: Experimental Evidence from a University Capital Campaign." *Journal of Political Economy*, Vol. 110 (2002), pp. 215–233.

Little, R. and Rubin, D. *Statistical Analysis with Missing Data*, 2d ed. Hoboken, NJ: John Wiley & Sons, 2002.

Lutzenhiser, L. *An Exploratory Analysis of Residential Consumption Survey and Billing Data: Southern California Edison, Summer 2001.* Sacramento: California Energy Commission, Report no. 400-02-006F, 2002.

Michelson, A., McGuckin, T., and Stumpf, D. "Effectiveness of Residential Water Conservation Price and Nonprice Programs." Denver American Water Works Research Foundation, Report no. 90747, 1998.

Nieswiadomy, M. "Estimating Urban Residential Water Demand: Effects of Price Structure, Conservation, and Education." *Water Resources Research,* Vol. 28 (1992), pp. 609–615.

Parks, R. and Weitzel, D. "Measuring the Consumer Welfare Effects of Time-Differentiated Prices." *Journal of Econometrics*, Vol. 26 (1984), pp. 35–64.

Parti, M. and Parti, C. "The Total and Appliance-Specific Conditional Demand for Electricity in the Household Sector." *Bell Journal of Economics*, Vol. 11 (1980), pp. 309–321.

Peck, A. and Doering, O. "Voluntarism and Price Response: Consumer Reaction to the Energy Shortage." *Bell Journal of Economics*, Vol. 7 (1976), pp. 287–292.

Reiss, P. and White, M. "Demand and Pricing in Electricity Markets: Evidence from San Diego during California's Energy Crisis." Working Paper no. 9986, NBER, 2003.

——— and White, M. "Household Electricity Demand, Revisited." *Review of Economic Studies*, Vol. 72 (2005), pp. 853–883.

SDG&E. *Energy Efficiency Programs Annual Summary and Technical Appendix–2001 Results.* San Diego: San Diego Gas and Electric Corporation, 2002.

Stern, P. "What Psychology Knows about Energy Conservation." *American Psychologist*, Vol. 47 (1992), pp. 1224–1232.

White, M. "Power Struggles: Explaining Deregulatory Reforms in Electricity Markets." *Brookings Papers on Economic Activity: Microeconomics*. Washington, DC: The Brookings Institution, 1996.

Wilson, R. "Architecture of Power Markets." *Econometrica*, Vol. 70 (2002), pp. 1299–1340.

Wolak, F. "Lessons from the California Electricity Crisis." In J. Griffin and S. Puller, eds., *Electricity Deregulation: Choices and Challenges.* Chicago: University of Chicago Press, 2005.