The Incentive to Overinvest in Energy Efficiency: Evidence from Hourly Smart-Meter Data

Kevin Novan, Aaron Smith

Abstract: Many households pay a marginal price for electricity that exceeds the marginal social cost of supplying that electricity. We show evidence that such pricing schemes can create an incentive to overinvest in energy efficiency. Using hourly smart-meter data for households facing time-invariant increasing block prices, we estimate how air conditioner upgrades affect electricity use. We find that the average participating household reduces consumption by 5%, which provides private savings in the form of lower electricity bills and social cost savings by decreasing generation and pollution costs. The private savings exceed the social savings by an average of 140%, so the average household is faced with an incentive to overinvest in energy efficiency. This incentive to overinvest in energy efficiency would be cut in half if consumers faced any one of three alternative pricing plans with lower marginal price but the same average price.

JEL Codes: H23, Q41, Q58

Keywords: Electricity demand, Energy efficiency, Rate structures

PRIVATE INVESTMENT IN ENERGY EFFICIENCY is widely believed to be inefficiently low, which has motivated policy makers to provide generous financial incentives to spur investment in efficiency upgrades.[1] From 2005 through 2012, the US federal government provided tax credits worth $13.7 billion to households making energy efficiency improvements (Borenstein and Davis 2016). Two arguments drive the

Kevin Novan (corresponding author) is at the Department of Agricultural and Resource Economics, University of California, Davis, One Shields Avenue, Davis, CA 95616 (knovan@ucdavis.edu). Aaron Smith is at the Department of Agricultural and Resource Economics, University of California, Davis, One Shields Avenue, Davis, CA 95616 (adsmith@ucdavis.edu).
We gratefully acknowledge financial support from the Giannini Foundation of Agricultural Economics and the Energy Foundation.

1. For example, the often cited McKinsey Report (McKinsey 2009) suggests that there is widespread underinvestment in energy efficiency across a variety of sectors.

Received August 4, 2016; Accepted September 26, 2017; Published online April 18, 2018.

JAERE, volume 5, number 3. 2333-5955/2018/0503-0003$10.00 http://dx.doi.org/10.1086/697050

belief that energy efficiency investment is too low. First, consumers do not bear the full cost of the pollution created by consuming energy. As a result, energy efficiency upgrades provide external benefits that do not accrue to those making the investments. Second, additional market failures (e.g., imperfect information, principal-agent problems) can result in an "energy paradox"—a situation in which investment in energy efficiency falls below the privately optimal level, let alone the socially optimal level.[2]

We focus on an often overlooked third market imperfection: consumers in the residential electricity sector typically face price schedules that differ from the first-best rate structure. In theory, efficient electricity rates consist of a fixed fee combined with a per kilowatt-hour (kWh) charge set equal to the social marginal cost (Coase 1946). In practice, however, per kWh charges are not equal to the social marginal cost. While the social marginal cost of providing electricity varies substantially across hours and across days, most households face per kWh charges that are fixed over time. Moreover, a household's per kWh rate often increases with its aggregate monthly consumption—a practice referred to as increasing block, or tiered, pricing.[3] Under tiered rates, the marginal price may far exceed the social marginal cost—particularly for households with high electricity consumption.

Inefficient retail rates can distort the stream of benefits from an energy efficiency investment and thereby distort the household's decision to invest. Specifically, an energy efficiency upgrade (e.g., installing an energy efficient air conditioning unit) lowers the household's cost of consuming an energy service (e.g., cooling a home). If the change in the cost to the household differs from the resource and environmental costs avoided by reducing energy consumption, then the household's incentive to invest differs from the social optimum.

The distortion may increase or decrease the investment incentive depending on the setting. If an energy efficiency upgrade will reduce a household's electricity consumption during periods when the marginal price exceeds the social marginal cost of supplying electricity, the household has an incentive to overinvest in energy efficiency—that

2. For overviews of the literature studying the energy paradox, see Allcott and Greenstone (2012), Gillingham and Palmer (2014), and Gerarden et al. (2017).

3. Tiered electricity prices are used throughout California as well as in many other locations. In a 2008 survey of 61 US electric utilities (BC-Hydro 2008), 25 employed tiered pricing for residential consumers. To provide more recent evidence of the prevalence of tiered pricing, we narrow our focus on the 15 largest electric utilities in the United States during 2016 (in terms of the number of customers). This information is provided by the US Energy Information Administration (https://www.eia.gov/energyexplained/). During 2016, these 15 utilities supplied over 20% of the total electricity consumed in the United States. Focusing on the standard (default) residential rates, 12 of the 15 utilities charge tiered rates (five of these tiered rates apply specifically during the summer months only).

is, invest beyond the socially optimal level.[4] In contrast, if the electricity savings occur during periods when the marginal price falls below the social marginal cost, the household is faced with an incentive to underinvest in energy efficiency. The incentives operate in the opposite direction if a household expects to increase energy consumption after an upgrade (e.g., if it plans to use substantially more air conditioning after obtaining a new unit). If the marginal electricity price exceeds the social cost, then such a household would be overpaying for additional electricity, creating an incentive to underinvest in energy efficiency.

In this paper, we examine how tiered rate structures affect households' incentives to invest in energy efficiency. We focus on air conditioning (AC), which is the largest source of US residential electricity consumption.[5] We find that, when households are faced with time-invariant tiered prices, the average private savings generated by investing in energy efficient AC units dramatically exceed the average social cost savings. Contrary to the belief that households underinvest in energy efficiency, our results suggest that private investment in residential energy efficiency may in fact exceed the socially optimal level.

We use data from 2,496 households in Sacramento, California, that participated in an energy efficiency rebate program. During 2012 and 2013, the households received a rebate for installing a new, energy efficient AC unit. Rather than focusing on the impact of the rebate program, we estimate how the AC upgrades affect electricity consumption in participating households. We use smart-meter data that records hourly electricity consumption at each residence. Participating households reduced their estimated nonwinter consumption by an average of 1.36 kWh per day—roughly a 5% reduction in total consumption. The impacts are very heterogeneous, however. In the 17% of households with high historical electricity consumption (>50 kWh/day), the AC upgrades reduced consumption by an average of 4.87 kWh/day. In contrast, in the third of households with low historical electricity consumption (<25 kWh/day), the AC upgrades caused consumption to increase by 0.96 kWh/day.

To quantify the resulting changes in the social cost of supplying electricity, it is crucial to know not only how much consumption changes but also when the changes occur. This imperative arises because the social marginal cost of supplying electricity varies substantially over time. Previous studies that estimate the energy savings achieved by residential efficiency upgrades relied on billing data, which record monthly,

4. In many settings, social costs are defined as private costs plus external costs, so social costs must exceed private costs. In our setting, we focus on the private cost to consumers who typically pay a price in excess of the resource cost of the electricity. Thus, the private cost of electricity includes a transfer to the utility in addition to the cost of supplying the electricity.

5. For a breakdown of residential electricity consumption by use, see the US Energy Information and Administration's "Annual Energy Outlook 2015."

household-level consumption.[6] As a result, the previous studies have been unable to determine when, within a given month, the energy savings take place. In contrast, the hourly data enable us to estimate precisely when the changes in energy consumption occur in each house. We find that consumption changes are concentrated during the evening hours of the hottest days—specifically when space cooling demand peaks.

We quantify the social cost savings provided by the new AC units by combining our estimates of the consumption changes during each hour with (1) observed hourly wholesale electricity prices and (2) estimates of the hour-specific marginal external cost of supplying electricity. We find that, by creating marginal prices that regularly exceed the social marginal cost, the current retail rate structure incentivizes some households to underinvest in energy efficiency and most to overinvest. On average across the sample, household bills decline by $6.46/month whereas the social cost of supplying electricity declines by only $2.69 following the AC upgrades. This incentive to overinvest in energy efficiency would be cut in half if consumers did not face a tiered rate structure but instead faced any one of three revenue-neutral pricing policies that are being actively considered by utilities and policy makers.

This paper contributes to a growing literature highlighting the inefficiencies stemming from current residential energy rate structures. For example, Davis and Muehlegger (2010) stress that the combination of low fixed charges and high marginal prices results in inefficiently low natural gas consumption. Borenstein (2012) and Borenstein and Davis (2012) further demonstrate that tiered rate structures can lead to an inefficient distribution of energy consumption across households.

These previous studies focus on how rate structures distort short-run energy consumption, whereas we focus on how the rate structure affects the incentive to invest in long-lived energy-consuming durables. Similarly, Borenstein (2017) demonstrates that California's tiered rate structures provide high-consuming households with an incentive to invest in solar panels. We show that, even without direct state and federal subsidies, the privately optimal level of residential energy efficiency exceeds the socially optimal level on average. Additional market failures (e.g., imperfect information, principal-agent problems) work in the opposite direction and may reduce investment. However, our results provide evidence that any such "energy paradox" would have to be widespread across households to justify the energy efficiency investment incentives created by the current tiered rate structures.

The remainder of the paper proceeds as follows. Section 1 discusses the AC upgrades and consumption data we examine. Section 2 presents estimates of the energy savings achieved by the upgrades, and section 3 explores the private savings. Section 4 compares the estimates of the private savings to estimates of the social costs avoided, and section 5 concludes.

6. For example, see Metcalf and Hassett (1999), Jacobsen and Kotchen (2013), Davis et al. (2014), Graff Zivin and Novan (2016), and Fowlie et al. (forthcoming).

1. DATA SOURCES

1.1. Energy Efficiency Rebate Program

This paper focuses on households purchasing electricity from the Sacramento Municipal Utility District (SMUD). As part of an effort to encourage investment in energy efficiency, SMUD provides customers with rebates for energy efficiency upgrades (e.g., improved insulation, new appliances, etc.). We examine the households that select to participate in SMUD's central air conditioning (AC) program.[7] This program provides residential customers with rebates for installing new central AC units that meet the EPA's Energy Star standards.[8]

We estimate how installing an energy efficient AC unit affects electricity consumption in participating households.[9] Specifically, we estimate the difference in the amount of electricity consumed by a household that has installed a new AC unit relative to the quantity of electricity that would have been consumed had the household not upgraded to the new, energy efficient AC unit.[10] Therefore, our estimates capture the impact of the physical AC upgrades as well as any resulting behavioral responses exhibited by the participating households. Using these estimates, we compute (*a*) the change in households' electricity bills following the upgrades and (*b*) the change in the social cost of supplying electricity. Comparing the changes in households' expenditures and the social cost, we explore whether households face an incentive to over- or underinvest in energy efficiency.

We focus on AC program participants for three reasons. First, the AC program has the most participants among SMUD's energy efficiency rebate programs.[11] Second, in the region we study, cooling is one of the largest sources of residential demand for elec-

---

7. In appendix A, we explore the differences between the households that elect to participate in the AC rebate programs and those that do not participate in the rebate programs.

8. To be classified as an Energy Star central AC unit, an AC unit must have a seasonal energy efficiency ratio (SEER) that exceeds the federal minimum of SEER of 14.

9. We do not estimate the impact of the AC rebate program per se. To evaluate the program, we would need to compare a household's post-upgrade electricity consumption to the consumption that would have occurred in the absence of the AC rebate program. The important distinction is that the AC upgrade may still have occurred even without the rebate program. In this case, while the physical AC upgrade may alter electricity consumption, the rebate program itself would have had no impact on electricity consumption.

10. A similar examination could be conducted in a different setting by estimating how electricity consumption changes differentially across households installing new, relatively energy efficient AC units versus households installing new, relatively inefficient AC units.

11. During 2012 and 2013, the period we examine, 6,142 single-family households received a rebate for installing an energy efficient AC unit. We focus exclusively on single-family premises as opposed to multifamily, rental units. The classifications are provided by Sacramento County Assessor data.

tricity.[12] As a result, improving the efficiency of AC units has the potential to cause large changes in energy consumption—and in particular, during hot days when demand for electricity, and the cost of providing it, peaks. Finally, understanding the effects of new AC units provides insights into how other investments intended to reduce cooling-related electricity demand (e.g., improved insulation) would impact residential electricity consumption. However, installing a new AC unit is a much more homogeneous treatment than improvements such as new insulation, so we are able to explore the extent to which the effects vary across homes.

1.2. Electricity Consumption and Expenditures

To estimate the impact of installing an energy efficient AC unit on electricity consumption, we use household-level smart-meter data. By 2012, each household in the SMUD service territory had a smart meter installed that recorded electricity consumption at the hourly frequency. For the period from January 1, 2012, through December 31, 2013, we observe the hourly consumption from each of the premises participating in the AC rebate program.[13]

In addition to the smart-meter data, we observe the household-level, monthly billing data. The billing data provide two key pieces of information. First, we observe if the SMUD account number at a premise changes. Often, this is a signal that a new owner is residing in the premise. To ensure that any of the observed changes in electricity consumption are not caused by a change in residents, we drop any premises that had multiple account numbers during the 2-year sample. Second, the billing data record the SMUD rate category in which each household is enrolled. Households in the standard rate class—which account for 92% of our sample—paid 9.89 cents per kWh on the first 700 kWh they consumed in a summer month and 18.03 cents per kWh on any additional electricity consumed.[14] During the period we examine (2012–13), the monthly consumption at the participating households placed them in the second price

12. The California Energy Commission provides a summary of electricity consumed by end source at the following location: http://energyalmanac.ca.gov/electricity/electricity_stats/index.html.

13. For security reasons, the consumption data are matched to the program participation data using an anonymized ID. Names, addresses, and account numbers were removed before we received the data.

14. The rates, and the threshold between tier 1 and tier 2 consumption, are not constant across households. Low-income households are eligible to receive a 30% reduction on their rate for all tier 1 and some tier 2 consumption through the Energy Assistance Program (EAPR). Roughly 7% of our sample paid EAPR rates. In addition, households with electric well pumps—roughly 1% of the sample—pay the tier 1 rate on the first 1,000 kWh of consumption during summer months.

tier 58% of the time. We use this information to estimate how the AC upgrades affect each participating household's energy expenditures.

Although the vast majority of residential consumers in the United States pay time-invariant prices—and in many cases, time-invariant prices that follow a tiered rate structure—several alternative pricing strategies are receiving serious consideration. First, a shift toward "flatter" rates that remove the steep, tiered structure has already begun throughout California.[15] In addition, with the increased penetration of smart meters, wider use of dynamic pricing is being considered across the United States.[16] To explore how the private savings provided by the AC upgrades would differ if the households did not pay the current tiered prices, we estimate the private bill savings the households would receive had they been faced with one of three alternative pricing strategies: time of day (TOU), critical peak pricing (CPP), and a flat rate.

2. IMPACT OF ENERGY EFFICIENCY ON CONSUMPTION

In this section, we introduce our empirical strategy and show that our approach accurately uncovers the impact of the AC upgrades. We begin by estimating the average change in daily electricity consumption that occurs after a participating household receives a new AC unit. We then augment the basic model to allow temporal heterogeneity in the impacts of the AC upgrades on electricity consumption. First, we explore how the impacts vary across hours of the day. Second, we examine how the impacts vary not only across hours but also days with different temperatures.

In section 3, we further augment the basic model to estimate how the impacts of the AC upgrades vary across time and across households. Accounting for this heterogeneity is crucial for producing unbiased estimates of the resulting changes in the monthly electricity bills. If the change in electricity consumption for a given household, in a given month, is correlated with the marginal price that household faced during that month, then simply multiplying the average change in electricity consumption by the average marginal price would produce a biased estimate of the change in the monthly bills. For example, suppose that households that pay the higher marginal price, the second tier rate, also conserve more electricity following an AC upgrade. In this case, there would be a positive correlation between the quantity of electricity saved and the marginal price paid for electricity. Multiplying an estimate of the average electricity savings by the average marginal price would understate the reduction in the monthly bills in this case.

---

15. The flattening of California's investor-owned utilities rate structures followed the passage of AB-327.

16. As of the end of 2016, over 50% of US households had smart meters installed (Cooper 2016).

2.1. Econometric Specification

We begin with the following basic model to estimate the average effect:

$$Cons_{i,d} = \alpha_i + \gamma \cdot Post_{i,d} + \boldsymbol{\theta} \cdot \mathbf{W_d} + \epsilon_{i,d}, \tag{1}$$

where i indexes the household and d indexes each day during the 2-year sample. The term $Cons_{i,d}$ represents the total consumption (kWh) for household i on day d. The term $\mathbf{W_d}$, which is discussed in greater detail below, is a flexible spline function controlling for temperature-driven shifts in electricity demand. Finally, $Post_{i,d}$ indicates household i's treatment status; it switches from 0 to 1 beginning on the day household i's AC rebate was sent. The key coefficient of interest, γ, represents the average change in a participating household's daily electricity consumption following the date the AC rebate is sent to the homeowner.

We do not observe the exact dates that the new AC units are installed. Instead, we observe the dates that the rebates were mailed to each of the participating households. Therefore, $Post_{i,d}$ will equal zero on an unknown number of days when premise i had already received a new AC unit. By including these post-upgrade observations during the pretreatment period (i.e., when $Post_{i,d} = 0$), we would expect to understate any energy savings provided by the new AC units.[17] However, we know anecdotally that the lag between the mailing dates and the physical upgrades is typically only 2–3 weeks. Moreover, due to the program requirements, the lag between the installation and the rebate date cannot exceed 90 days. To examine the extent to which our imperfect measure of the treatment date can impact our estimate of the average effect of the AC units, we present several estimates of equation (1), each time dropping between 0 and 90 days worth of observations immediately preceding the observed rebate dates. By accurately separating the pre- and post-upgrade observations, we can interpret the coefficient γ as the average change in a participating household's daily consumption following the AC upgrade.

Because equation (1) includes household fixed effects, we are using within-household variation to identify γ. Therefore, it is important to control for time-varying determinants of electricity demand that could be correlated with the timing of the AC upgrades. In the region we study, the temperature is the key factor driving daily variation in electricity consumption. To control for temperature-driven shifts in electricity demand, equation (1) includes the average daily temperature in Sacramento (T_d).[18]

17. It is also possible that a household installs a new AC unit in response to their existing unit breaking. As a result, consumption during the days immediately preceding $Post_{i,d}$ switching from 0 to 1 could be lower than would be expected if the existing AC unit was still functioning. In this case, we would expect to overestimate the energy savings provided by the new AC unit relative to the previous, functioning AC unit.

18. We use National Oceanic and Atmospheric Administration temperature readings from the Sacramento Airport. We considered several alternative strategies instead of using the aver-

The daily temperature enters the model in piecewise linear form with three knot points (at 63°F, 70°F, and 75°F).[19] Specifically, $\mathbf{W_d}$ represents the following 5×1 vector:

$$\mathbf{W}_d = \begin{bmatrix} 1 \\ \min(T_d, 63) \\ \min(\max(T_d - 63, 0), 70 - 63) \\ \min(\max(T_d - 70, 0), 75 - 70) \\ \max(T_d - 75, 0) \end{bmatrix}. \tag{2}$$

In an effort to focus solely on the impacts of the AC upgrades, we drop households that participate in multiple SMUD rebate programs. Additionally, some households receiving a new AC unit may be replacing a heat pump. In contrast to AC units, which are used only for cooling, heat pumps use electricity for cooling and heating. Therefore, a household that replaces a heat pump with a new AC unit may achieve a reduction in cooling-related electricity consumption as well as a reduction in heating-related electricity consumption—which is not caused by the new AC unit but, rather, by a switch from electric to natural gas heating. We do not observe natural gas consumed by the households in our data set.

To ensure that any electricity savings caused by a switch to gas heating are not attributed to the new AC units, we place two additional restrictions on our sample. First, we drop all households enrolled in SMUD's electric-heat rate plan. The electric-heat plan is open to households using electricity as their primary source for heating. Therefore, dropping these homes removes the premises that are likely to have heat pumps to begin with. Second, given that some eligible households will not enroll in the electric-heat rate program, we drop consumption during the five coldest months (November through March) from the analysis and focus exclusively on the warmer months when cooling-driven demand shifts may occur.[20]

---

age daily temperatures—including using the individual hourly temperatures. However, the average daily temperature is very highly correlated with the hourly temperatures. For each of the 24 hours of the day, the correlation coefficient between the average daily temperature and the hour-specific temperature never falls below 0.81 and, on average, equals 0.92. Appendix E demonstrates that the estimated electricity savings are robust to the alternative temperature controls. In addition, we demonstrate that the inclusion of monthly fixed effects does not meaningfully alter the estimates.

19. The temperatures 63°F, 70°F, and 75°F are approximately the 25th, 50th, and 75th percentiles of average daily temperatures between April and October—the nonwinter months that we ultimately focus on in the empirical analysis.

20. The five excluded months are the only months with average daily temperatures below 60°F.

With these restrictions, we observe 5,423 households that received an AC upgrade. We include in our analysis only the 2,496 households that received a rebate between June 16, 2012, and July 4, 2013. Using only rebates inside this window has a negligible quantitative effect on our estimates of the basic model in equation (1). However, by ensuring that the pre-rebate and post-rebate observations have a common support over the range of temperatures, this restriction enables us to estimate heterogeneous effects in section 3.

2.2. Basic Model Results

Table 1 presents results from estimating equation (1) using the daily observations from April to October of 2012–13 for the 2,496 households in our estimation sample. The standard errors are clustered by household and week by year. On average, a household's daily consumption falls by 1.27 kWh following the AC rebate.[21] Median daily consumption during the months examined is 24 kWh, so this estimate represents about a 5% reduction in consumption.

The average reduction of 1.27 kWh per day likely understates the average energy savings caused by the new AC units because the AC units were installed before the rebate date. To account for this, we estimate equation (1) dropping the 14, 30, 60, or 90 days immediately preceding the observed rebate dates. Columns 2–5 of table 1 present the resulting estimates. Consistent with the fact that AC upgrades occur up to 90 days prior to the rebate date, as the number of dropped pre-rebate observations grows, the estimates of the average change in electricity consumption become more negative. Dropping 90 days immediately preceding the rebate dates, we estimate that, on average, daily consumption falls by 1.40 kWh following the installation of a new AC unit. The results presented in table 1 also reveal that the majority of the jump in the estimate of γ is achieved by dropping the first 30 pre-rebate days. This finding corroborates the anecdotal evidence that the majority of households receive their rebates within the first month following the AC upgrade.

To provide additional evidence regarding when the actual AC upgrades occur, we generalize equation (1) to allow γ to vary flexibly across the months leading up to, and following, the observed AC rebate dates. Specifically, we estimate the following model:

$$Cons_{i,d} = \alpha_i + \sum_{m=-7}^{-1} \gamma_m \cdot Pre_{i,d,m} + \sum_{m=1}^{7} \gamma_m \cdot Post_{i,d,m} + \boldsymbol{\theta} \cdot \mathbf{W_d} + \epsilon_{i,d} \,. \qquad (3)$$

21. Table A1 (tables A1–A3, B1, B2, C1, C2, D1 are available online) presents estimates of equation (1) using the full set of 5,423 households that receive an AC rebate at any point between January 1, 2012, and December 31, 2013. The corresponding point estimate is a 1.24 kWh drop in daily consumption following the AC rebate.

Table 1. Average Change in Daily Electricity Consumption (kWh)

	Number of Days Dropped Prior to Rebate Date				
	0 Days (1)	14 Days (2)	30 Days (3)	60 Days (4)	90 Days (5)
Post	−1.27**	−1.32**	−1.36**	−1.38**	−1.40**
	(.42)	(.44)	(.45)	(.47)	(.48)
Temperature < 63°F	−.01	−.01	−.01	−.01	−.01
	(.03)	(.03)	(.03)	(.03)	(.03)
63°F < Temperature < 70°F	.74**	.74**	.74**	.74**	.75**
	(.08)	(.08)	(.08)	(.08)	(.08)
70°F < Temperature < 75°F	1.48**	1.48**	1.48**	1.48**	1.46**
	(.12)	(.12)	(.12)	(.12)	(.12)
Temperature ≥ 75°F	1.80**	1.80**	1.80**	1.81**	1.80**
	(.07)	(.07)	(.08)	(.07)	(.07)
N	1,091,230	1,067,301	1,040,107	989,737	940,166
Within R^2	.41	.41	.41	.41	.41

Note. Each model is estimated using household fixed effects. Standard errors are robust to heteroskedasticity and clustering at the household and week-by-year level.

* Significant at the 5% level.

** Significant at the 1% level.

The dummy variables $\{Pre_{i,d,m}\}$ separate the days immediately preceding each household's observed rebate date into seven mutually exclusive 30-day pre-rebate windows. Thus, γ_{-1} represents the average change in a household's daily electricity consumption that occurs during the first 30 days preceding the AC rebate dates while γ_{-2} represents the average change in daily consumption during the window 30–60 days prior to the rebate dates. Similarly, the seven post-rebate dummy variables $\{Post_{i,d,m}\}$ separate the observations following each household's rebate date into 30-day windows representing 1, 2, 3, 4, 5, 6, or 7+ months after the rebate date. All of the changes are measured relative to the daily consumption that occurs more than 240 days prior to the household's rebate dates.

Figure 1 presents the point estimates of the pre- and post-rebate γ_m values from equation (3). Focusing on the pre-rebate months, there is evidence that the AC upgrades occur largely during the first 30 days prior to the AC rebate date. On average, households reduce their daily consumption by 1.14 kWh during the first pre-rebate month. This is approximately 75% as large as the estimated energy savings during the first post-rebate month (1.47 kWh/day). This is consistent with approximately 3 weeks of the first pre-rebate month coming from households which had already received a new, energy efficient AC unit.

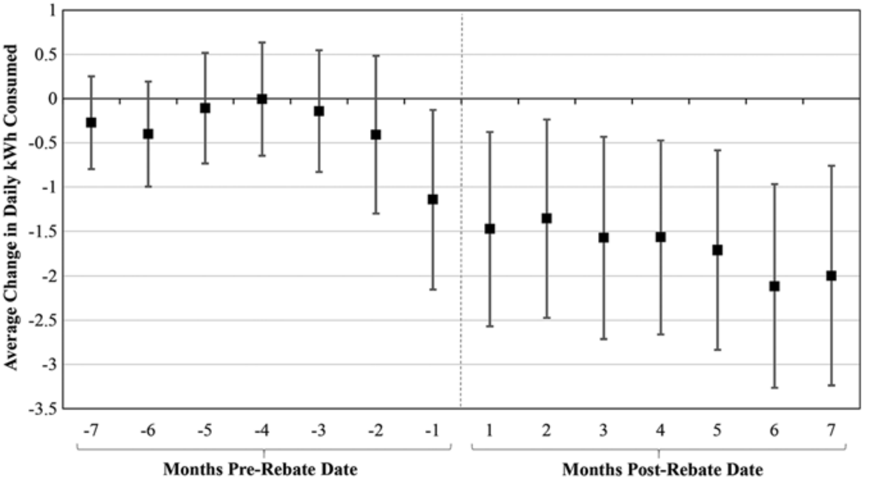

Figure 1. The graph displays the point estimates, and the corresponding 95% confidence intervals, of the average change in the daily electricity consumption during the months before and after the AC rebates are sent. The first "pre-rebate month" (month = −1) includes all observations during the 30 days immediately preceding the date each household's AC rebate is sent. The first "post-rebate month" (month = −1) includes all observations during the first 30 days following each household's AC rebate date. The changes are measured relative to the average daily consumption on days 8 or more months prior to the AC rebate being sent. The confidence intervals are robust to heteroskedasticity and clustering at the household level and across households within each week by year.

Figure 1 shows no significant changes in the average level of electricity consumption more than 30 days before the observed AC rebate dates. These findings, combined with the estimates presented in table 1, suggest that we can quite accurately separate the daily observations into pre-upgrade and post-upgrade observations by (*a*) continuing to use the $Post_{i,d}$ indicator to reflect whether a household has been treated with a new AC unit and (*b*) dropping the 30 days immediately preceding each individual household's observed rebate date. Throughout the remainder of our analysis, we follow this procedure to estimate the impact of the AC upgrades.

In addition to shedding light on when the AC upgrades occur, the results presented in figure 1 also provide evidence that the estimated changes in consumption are being caused by the AC upgrades as opposed to an alternative, time-varying household characteristic that we cannot observe. If the consumption changes were instead being driven by a confounding variable that is simply correlated—but not perfectly—with the timing of the AC upgrades, then we would expect to see significant changes in consumption occur prior to the AC upgrades—that is, during the period of time more than 30 days before the AC rebate dates. However, there are no significant changes in consumption preceding the first pre-rebate month when the AC upgrades largely occur. Moreover, if the AC upgrades were being performed prior to other household changes that could affect electricity demand (e.g., new household members), then we might expect the estimated consumption changes during the post-rebate months to vary over time as the confounding demand shifts occur. However, average electricity consumption is stable during the post-rebate months.

2.3. Heterogeneity in Energy Savings over Time

A key advantage of the hourly consumption data is that it allows us to estimate not only how much energy is conserved but also when the energy savings occur. This is valuable for two main reasons. First, given that the marginal cost of supplying electricity varies substantially across time, understanding when the reductions in consumption occur plays an important role in determining the social costs that are avoided. Second, by uncovering when the energy savings occur, we can provide additional evidence that our estimation strategy is uncovering the consumption changes caused by the AC upgrade. In particular, if the reductions in consumption are caused by installing energy efficient AC units, then we would expect the energy savings to be concentrated during the late afternoon hours on the hottest days.

To examine how the energy savings vary across hours of the day, we estimate the model in equation (1) separately for each hour of the day. The new dependent variable is now $Cons_{i,h,d}$ the electricity consumed at household i during hour h of day d, and we obtain 24 point estimates of γ_h, one for each hour. Figure 2 displays the estimated γ_h coefficients and reveals that the largest energy savings caused by the AC upgrades occur during the late afternoon and evening hours. In contrast, during the morning hours, the estimated energy savings are small and statistically insignificant. These findings are consistent with the daily temperature patterns in Sacramento. During the months in our sample, the average hourly temperature during the 5 p.m. hour was 90°F—the highest across all hours—and remained above 75°F through the 9 p.m. hour. In contrast, during the morning hours (6 a.m. through 10 a.m.), the average hourly temperature was 64°F.

Importantly, the energy savings are not confined to the period between 4:00 p.m. and 7:00 p.m. when demand on the California grid typically peaks—and consequently, when wholesale electricity prices are at their highest. Instead, the post–7 p.m. hours display the largest average energy savings. This finding suggests that a large share of the energy savings take place during hours when the marginal generation costs are relatively low.

If the AC upgrades are reducing the electricity consumed to cool homes, then we would expect the energy savings to vary not only across hours of the day but also across days with different average temperatures. To directly examine how the quantity of energy saved varies with the daily temperature, we generalize equation (1) further. We

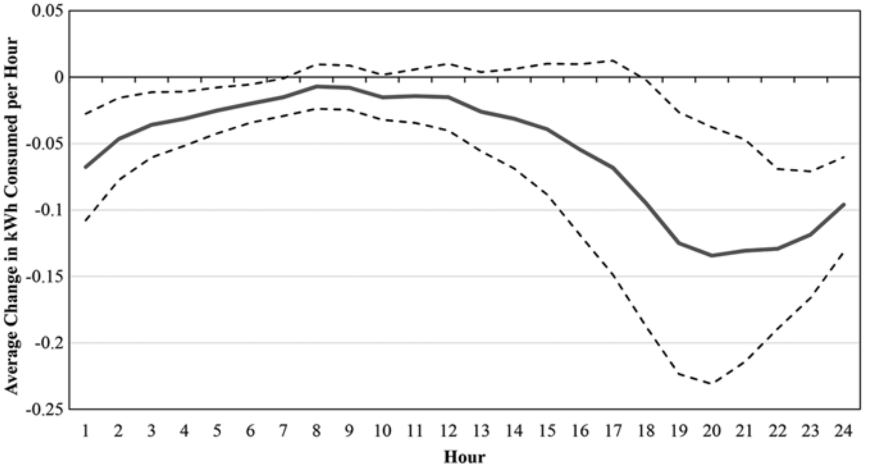

Figure 2. The graph displays the point estimates, and the corresponding 95% confidence intervals, of the average hourly changes in a household's energy consumption following an AC upgrade. The confidence intervals are robust to heteroskedasticity and clustering at the household level and across households within each week by year. To produce the point estimates, the observations from the 30 days preceding each household's AC rebate date are removed from the sample.

add terms that interact $Post_{i,d}$ with the spline function of the average daily temperature ($\mathbf{W_d}$) and reestimate the model separately for each hour of the day. To highlight the heterogeneity in the resulting energy savings, we report two specific hours of the day—8 a.m., when the minimum average energy savings occurs, and 8 p.m., when the maximum average energy savings occurs.

Figure 3 displays the estimates of the average hourly change in electricity consumption, following the receipt of an AC rebate, on days with average temperatures ranging from 52°F to 86°F.[22] The estimates reveal that, regardless of the average daily temperature, the new AC units do not affect the average level of electricity consumed during the 8 a.m. hour. This is consistent with households not using their AC units for cooling during the morning hours. In contrast, during the 8 p.m. hour, we find that the new AC units provide significant energy savings on days when the average temperature exceeds 72°F. Moreover, as the temperature increases beyond 72°F, the average energy savings during the 8 p.m. hour increase.

22. These temperatures represent the 1st and 99th percentiles of the average daily temperatures during April through October of 2012 and 2013.

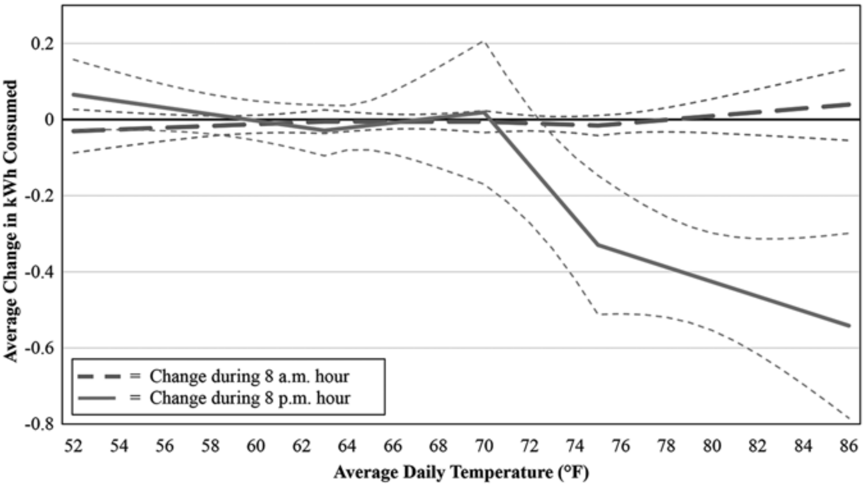

Figure 3. The graph displays the point estimates, and the corresponding 95% confidence intervals, of the average hourly changes in a household's electricity consumption following an AC upgrade during two different hours—8 a.m. and 8 p.m.—as a function of the average daily temperature. The confidence intervals are robust to heteroskedasticity and clustering at the household level and across households within each week by year. To produce the point estimates, the observations from the 30 days preceding each household's AC rebate date are removed from the sample.

In sum, we find that the energy savings are concentrated during the late afternoon hours on the hottest days. These results provide strong evidence that, rather than being driven by omitted variable bias, our estimates are uncovering the impacts of the AC upgrades on the participating household's electricity consumption.

3. PRIVATE COST SAVINGS FROM ENERGY EFFICIENCY

In this section, we present estimates of the private savings households receive on their electricity bills following the installation of the energy efficient AC units. Specifically, we consider the following question. How much less would each household have paid for electricity in the summer of 2013 if their new AC unit had been in place than if they had their old AC unit?

As explained in section 2, estimating bill savings requires us to further generalize the basic model to allow heterogeneity across households. Thus, we estimate the savings separately for each household. Using these estimates, we compute how each household's expenditure on electricity is affected by the new AC units and then summarize the predicted private savings across households. This approach allows households that experience larger energy savings following an AC upgrade to be in a different pricing tier from those with smaller savings.

3.1. Household-Level Consumption Changes

We estimate the following model separately for each premise and for each hour of the day:

$$Cons_{i,h,d} = \alpha_{i,h} + \boldsymbol{\beta}_{i,h} \cdot \mathbf{W_d} \cdot Post_{i,d} + \boldsymbol{\theta}_{i,h} \cdot \mathbf{W_d} + \epsilon_{i,h,d}, \tag{4}$$

where $Cons_{i,h,d}$ represents the electricity consumed at household i during hour h of day d and $\mathbf{W}_d$ is again specified by equation (2).[23] As in section 2, we include the 2,496 households that received an AC rebate between June 16, 2012, and July 4, 2013, and we use electricity consumption from April through October of 2012 and 2013, excluding the 30 days immediately preceding each household's observed rebate date.

Combining the estimates $\{\hat{\alpha}_{i,h}, \hat{\boldsymbol{\beta}}_{i,h}, \hat{\boldsymbol{\theta}}_{i,h}\}$ with the daily temperatures from June through September 2013, we can predict the expected consumption with and without the new AC unit for each household during the summer of 2013. Specifically, the estimated pre-upgrade and post-upgrade consumption levels for household i during hour h of day d are given by:

$$\text{Pre-upgrade consumption}_{i,h,d} = \hat{\alpha}_{i,h} + \hat{\boldsymbol{\theta}}_{i,h} \cdot \mathbf{W_d} \tag{5}$$

$$\text{Post-upgrade consumption}_{i,h,d} = \hat{\alpha}_{i,h} + \hat{\boldsymbol{\theta}}_{i,h} \cdot \mathbf{W_d} + \hat{\boldsymbol{\beta}}_{i,h} \cdot \mathbf{W_d}\,. \tag{6}$$

The first column of table 2 presents the estimates of the average pre- and post-upgrade electricity consumption and shows that the energy efficient AC units reduce summer consumption by 1.13 kWh per day on average.[24] There is, however, considerable heterogeneity in the estimated energy savings across households. To highlight this heterogeneity, we separate the households into three groups—low, medium, and high electricity users. To create these groups, we first aggregate each household's consumption during the summer (June through September) of 2011—prior to the 2012 and 2013 period when we observe the household's hourly consumption.[25] The low-

23. In related work, Reiss and White (2008) examine how households respond to electricity price shocks and calls for conservation by estimating household-specific temperature response functions.

24. This estimate differs from the pooled fixed effects (FE) estimates presented in table 1 because, by solving for the simple mean of the household-specific, average energy savings, we are now placing an equal weight on the average change in electricity consumption that occurs at each household. In contrast, the pooled FE model in equation (1) places a greater weight on the energy savings achieved by households that receive their AC upgrade in the middle of the sample period.

25. To calculate the 2011 summer consumption, we utilize the monthly billing data. While all of the households we examine reside in their home for all of 2012 and 2013, not all were in the home during the summer of 2011. As a result, our sample shrinks to 2,437 households. For households that have billing periods that do not align with the calendar months, we uniformly

Table 2. Pre- and Post-upgrade Consumption under Tiered Prices

		Consumption Group		
Average Consumption (kWh/day)	All Households	Low	Medium	High
Pre-upgrade	32.56	17.86	34.14	59.35
	(16.68)	(6.29)	(8.39)	(16.16)
Post-upgrade	31.43	18.82	32.79	54.48
	(15.22)	(6.79)	(8.27)	(16.61)
Change	−1.13	.96	−1.35	−4.87
	(7.02)	(4.95)	(6.34)	(10.21)
Number of households	2,496	826	1,220	391

Note. The point estimates represent the simple average of the household-level, mean predicted pre- and post-upgrade daily consumption levels from June 1, 2013, through September 30, 2013. The Change point estimates provide the simple average of the household-level, mean changes in daily electricity consumption following the upgrades. The standard deviations represent the standard deviation among the household-level mean daily consumption levels and changes.

consumption group includes the households that, on average, use less than 25 kWh/day during the summer of 2011. These low users account for approximately 33% of the households. The medium-consumption group, approximately 50% of the homes, uses between 25 kWh and 50 kWh per day. Finally, the high-consumption group, the remaining 17% of households, uses more than 50 kWh/day.

The last three columns of table 2 summarize the estimates of the pre- and post-upgrade electricity consumption across the three different consumption groups.[26] On average, the new AC units save households in the high-consumption group 4.87 kWh per day. Households in the medium-consumption group save an average of 1.35 kWh per day.[27] Finally, households in the low-consumption group increase electricity usage

allocate consumption to a calendar month based on the share of the billing period occurring in each given month.

26. Figure D2 (figs. B1, D1–D3, E1 are available online) plots the average consumption changes by consumption group and by hour of day.

27. Imposing the assumption that the estimated consumption changes are independent across households, the information in the last three columns of table 2—that is the mean consumption changes by group, the sample standard deviation by group, and the number of homes in each consumption group—can be used to conduct a difference in means test between the average consumption changes for any two consumption groups. Doing so reveals that the average consumption changes differ significantly across low-, medium-, and high-consumption groups at the 5% significance level.

by an average of 0.96 kWh per day.[28] In appendix B (appendix with sections A–E available online), we explore potential causes of the consumption increase among the low-consumption households. We demonstrate that almost all these households used AC prior to participating in the rebate program, suggesting that the increase in their electricity cannot be explained by the addition of a functioning AC unit to a house that previously did not have one. Rather, we present evidence that the increase in their electricity use is consistent with a sizable rebound in their use of space cooling.

3.2. Private Cost Savings under Tiered Pricing

We assume that customers are billed by the calendar month. Then, for June, July, August, and September of 2013, we estimate the monthly aggregate pre-upgrade and post-upgrade consumption by each household by summing the hourly estimates in equation (5) and equation (6). Next, we use these estimates to calculate household electricity bills with and without the new AC units. To calculate the bills, we use the actual rates paid by each household.[29]

The first column of table 3 summarizes the estimates of the monthly expenditure changes. The average bill dropped by an estimated $6.46 per month. Without the new AC units, the households would pay an average summer bill of $125.58 per month, which falls to $119.12 with the new AC unit. Within the low-consumption group, the new AC units increase the average monthly bill by $3.95, whereas monthly bills fall by $7.03 and $26.56 for the medium- and high-consumption groups, respectively.

The variation in the average expenditure changes across the low, medium, and high users is largely driven by the differences in the energy conserved by the new AC units. On average, high-consuming households save the most energy, and therefore, save the most on their bills. However, some of the variation in the private savings is driven by the tiered rate structure. Given that high-consuming households are almost exclusively

---

28. To examine whether the pattern across groups is being driven in part by mean reversion in household-level electricity consumption, we use an alternative approach to classify households as low, medium, and high consumers. We use billing data to calculate the average daily consumption during three summers (2009, 2010, and 2011). low-consuming households are defined as those that have an average consumption below 25 kWh/day during each of the three historical summers. high-consuming households are all households that have an average daily consumption above 50 kWh during each of the three summers. Medium-consuming households are all other households. Table A2 presents the average consumption changes across the three consumption groups. The average changes are very similar to the results presented in table 2.

29. Households also pay a fixed charge of $12 per month. Given that the fixed charge is unaffected by the energy efficient AC units, we do not include the $12 charge in our predicted electricity expenditures.

Table 3. Change in Average Monthly Expenditure and Social Cost

	Monthly Expenditure Change				Monthly Social Cost Change	
Consumption Group	Tier	TOU	CPP	Flat	\$38/ton CO_2	\$100/ton CO_2
All households	–\$6.46	–\$4.37	–\$5.81	–\$4.39	–\$2.69	–\$4.22
	(35)	(29)	(29)	(27)	(17)	(26)
Low group	\$3.95	\$4.16	\$2.96	\$3.56	\$2.41	\$3.70
	(21)	(20)	(22)	(19)	(12)	(19)
Medium group	–\$7.03	–\$5.48	–\$7.50	–\$5.18	–\$3.34	–\$5.15
	(33)	(26)	(27)	(24)	(15)	(24)
High group	–\$26.56	–\$19.01	–\$19.35	–\$18.73	–\$11.44	–\$18.01
	(55)	(40)	(39)	(39)	(24)	(38)

Note. The estimates represent the mean of the household-level average change in the monthly expenditure, during June through September of 2013, that is caused by installing a new AC unit. The standard deviations of the electricity expenditures represent the standard deviations of the household-level average monthly expenditures changes. Under TOU, customers pay \$0.27 per kWh for electricity used between 4:00 p.m. and 7:00 p.m. and \$0.093 per kWh otherwise. Under CPP, customers pay \$0.75 per kWh between 4:00 p.m. and 7:00 p.m. on the 13 hottest days during summer 2013 and \$0.089 per kWh otherwise. Under flat pricing, customers pay \$0.128 per kWh across all hours and all days. We set the off-peak and flat rates to ensure that the pre-upgrade bills would, on average, equal \$125.58 per month as they do under the tiered prices under the assumption that the pre-upgrade consumption would be unchanged, regardless of how the prices change. TOU = time of use; CPP = critical peak pricing.

in the top price tier, conserving a kWh of electricity results in the maximum potential private savings. Dividing the estimated average monthly expenditure changes by the average monthly consumption changes, we find that households in the high-consumption group save an average of 17.95 cents per kWh conserved and households in the medium-consumption group save an average of 17.3 cents per kWh conserved. Low-consumption-group households spend an average of 12.39 cents per additional kWh used. The following subsection examines how removing the tiered price structure would affect the private savings provided by the AC investments.

3.3. Private Cost Savings under Alternative Pricing Strategies

We examine how three alternative pricing policies would affect the private savings provided by the new AC units. First, we simulate a time of use (TOU) plan under which customers pay \$0.27 per kWh for electricity used between 4:00 p.m. and 7:00 p.m. and a constant off-peak rate otherwise. Second, we simulate a critical peak pricing (CPP) program under which customers pay \$0.75 per kWh during peak hours (4:00 p.m. to 7:00 p.m.) on the 13 hottest days during summer 2013 and a constant off-peak rate on all other hours and days. Finally, we simulate is a flat price (\$0.128 per kWh) that is constant across all hours and all days. We set the off-peak and flat rates to ensure that the pre-upgrade bills would, on average, equal \$125.58 per month as they do under the tiered prices.[30]

The second through fourth columns of table 3 present the estimates of the average impact the new AC units would have on the household's monthly bills during summer 2013 under TOU, CPP, and flat retail rates.[31] Assuming demand is perfectly inelastic, households would save an average of \$4.37/month under the TOU plan, \$5.81/month under the CPP plan, and \$4.39/month under flat retail rates. In appendix C, we present estimates of the expenditure changes relaxing the assumption that demand is perfectly inelastic. The estimates reveal that the private savings are largely unaffected by the assumed elasticity of demand.

Overall, aggregating across all households, the private savings provided by the new AC units are the greatest under the current tiered pricing structure. This outcome is driven by two factors. First, the majority of energy savings comes from the high-consuming households. Under tiered pricing, these households typically pay the higher tiered rate, even during the off-peak hours. The second factor is that the majority of energy savings occur after 7:00 p.m.. As a result, with tiered prices, the off-peak energy savings in the high-consuming households provides a relatively large reduction in their bills. In contrast, these off-peak energy savings will become less valuable privately under the TOU, CPP, and flat rate policies.

Table 3 also highlights how the private savings vary across consumption groups. Regardless of the policy, the high-consuming households receive the largest average private savings from the new AC units. Comparing the expenditure changes across the different pricing policies, the gap between the private savings under tiered prices and the alternative policies is most pronounced for the high-consumption group. The new AC units reduce the monthly expenditures within the high-consuming households by an average of \$26.56 per month—roughly 40% more than the private savings under the alternative policies (≈\$19/month).

---

30. The TOU and CPP plans follow the plans SMUD recently explored during a Smart Pricing pilot study (SMUD 2014). The actual Smart Pricing pilot study only increased the peak hour prices on nonholiday weekdays. Rather than estimating how each individual household's demand for electricity differs by weekdays and weekends, we instead assume that, under our TOU plan, the peak hour prices increase each day of the week.

31. In table C2, we present the point estimates of the average difference between the change in monthly expenditure under each counterfactual pricing policy relative to the observed tiered rate structure. To do so, we calculate the difference in the estimated savings under a given counterfactual policy and the tiered rates for each individual household. We then solve for the simple mean of this difference across all households as well as for the low-, medium-, and high-consumption groups. In addition, we present estimates of the standard errors of the differences in the mean savings—which enables us to highlight whether the differences in the estimated expenditure savings are statistically significant. To calculate these standard errors, we assume the estimated expenditure savings are independent across households.

4. SOCIAL COST SAVINGS FROM ENERGY EFFICIENCY

The preceding estimates quantify the private cost savings that households received by installing new, energy efficient AC units. In this section, we compare the private cost savings to estimates of the social cost savings. Specifically, we quantify the reduction in the generation, external pollution, and capacity investment costs that occur during the summer of 2013 as the result of a household installing an energy efficient AC unit.

4.1. Estimating the Avoided Social Costs

Changes in electricity consumption affect social costs in two broad ways. First, the cost of producing energy (i.e., private generation costs and external costs) can change. Second, if consumption changes during the period(s) when electricity demand peaks in the region, the cost of the capacity required to meet the maximum demand also changes. Here, we outline our approach to estimating the change in private generation and external costs, as well as the change in the capacity cost. Appendix D describes the estimation procedures in detail.

We focus first on estimating the change in the social cost of producing electricity. From equation (4), the expected impact of a new, energy efficient AC unit on household i's consumption during hour h on day d is equal to $\boldsymbol{\beta}_{i,h} \cdot \mathbf{W_d}$, where $\mathbf{W_d}$ is specified by equation (2). Therefore, installing an energy efficient AC unit in household i will cause the social cost of consuming electricity during hour h of day d to change by the following amount:

$$\Delta\text{Social cost}_{i,h,d} = (\rho_{h,d} + \mu_{h,d}) \cdot \boldsymbol{\beta}_{i,h} \cdot \mathbf{W_d}, \tag{7}$$

where $\rho_{h,d}$ denotes the marginal private cost of supplying electricity during hour h of day d and $\mu_{h,d}$ is the marginal external cost of electricity.

To estimate the hourly marginal private costs ($\rho_{h,d}$), we use the average hourly locational marginal price (LMP) paid for electricity in SMUD's service region in the day-ahead market.[32] During the nonwinter months of 2013 (April through October), the average hourly price never exceeded 8 cents per kWh ($80 per MWh) outside of 2 p.m. through 8 p.m. In contrast, from 4 p.m. through 6 p.m., the hourly prices often increased above 13 cents per kWh on the highest demand days.[33] The magnitude of the temporal variation in marginal prices highlights why it is crucial to determine when the energy savings occur.

Quantifying the avoided external costs is complicated by existing environmental policies. In particular, if the cap in California's CO_2 cap-and-trade program were binding, then a marginal decrease in electricity demand among the SMUD customers would not affect the aggregate emissions of CO_2. Given the uncertainty surrounding

32. Table D1 provides estimates using the average hourly prices in the real time market.
33. Figure D1 summarizes the distribution of the hourly wholesale prices during this period.

California's business-as-usual CO_2 emissions (Borenstein et al. 2015), we take a conservative approach and assume the cap is not binding and therefore that reductions in electricity use also reduce CO_2 emissions.

To produce estimates of the marginal external cost of supplying electricity to the SMUD region ($\mu_{h,d}$), we follow the approach used in several recent studies (Siler-Evans et al. 2012; Carson and Novan 2013; Jacobsen 2014; Callaway et al. 2018) using hourly data from the EPA's Continuous Emissions Monitoring Systems. The approach abstracts from variation in the marginal external cost across days, but it does capture the variation across hours in the average day (see app. D for details).

In addition to the avoided generation and pollution costs, the energy efficient AC units can also provide social cost savings by reducing, or deferring, the required investment in generation capacity.[34] To approximate these additional social cost savings, we assume that the avoided social cost from reducing peak demand by a kilowatt (kW) is equal to $2.66 per month. This is the average monthly contracted price for capacity from 2013 through 2017 in the Northern Zone of California under California's Resource Adequacy Program (CPUC 2015a).[35] To estimate the change in the cost of the required capacity, we multiply the average changes in 5 p.m. summer consumption by the average monthly cost of capacity ($2.66/kW). This adjustment raises the social marginal cost by about 6%.

4.2. Comparison of Private and Social Cost Savings

The last two columns of table 3 show the average change in the monthly social cost of supplying electricity to a participating household. We report estimates for two values of the social cost of carbon (SCC): $38 per ton and $100 per ton.[36] The estimates of the social cost changes apply specifically to the case where the current tiered rate structure is maintained.[37]

Using a $38 per ton SCC, the new AC units reduce the estimated total social cost of providing electricity to a participating household by an average of $2.69 per month—of

34. There can also be reductions, or deferments, in the required investment in distribution infrastructure. However, Cohen et al. (2016) estimate that, in the vast majority of locations, these cost savings are likely to be negligible.

35. Capacity contract prices are presented in table 11 of the CPUC's 2013–14 Resource Adequacy Report.

36. To estimate the monthly household-level social costs avoided, we sum the estimates of the hourly avoided generation and pollution costs over each calendar month and add the estimates of the change in monthly capacity costs. The simple means of the household-level average monthly avoided generation, pollution, and capacity costs are individually presented in table D1. Figure D3 presents the average change in the monthly social cost of supplying electricity to a participating household alongside the private cost savings under each pricing policy.

37. With nonzero demand elasticities, the social cost changes could differ under alternative rates.

which $1.62 is avoided private generation costs, $0.96 in avoided external pollution costs, and $0.11 is avoided generation capacity costs. If CO_2 imposes an external cost of $100 per ton, the AC units will reduce the monthly social costs by an average of $4.23—the avoided external pollution costs increase to an average of $2.50/month. In contrast, under the existing tiered electricity prices, households save an average of $6.46 per summer month.

For the high-consumption group, energy efficient AC units reduce the social cost of providing electricity by an average of $11.44 per month at a $38/ton SCC. Setting the SCC to $100/ton increases monthly cost savings to $18.01, which is still substantially smaller than the private savings of $26.56 per month. The 50% of households in the medium-consumption group display the same pattern. The average monthly private savings exceed average social savings by 110%, assuming an SCC of $38/ton. In the lowest consuming households, the average monthly private expenditures increase by $3.95 under the observed tiered prices, whereas the energy efficient AC units only increase the average monthly social cost by $2.41 assuming SCC = 38/ton.

The differences between the private and social cost changes arise because the marginal price paid for electricity exceeds the social marginal cost of supplying electricity during the vast majority of hours. As a result, reductions in energy consumption caused by the new AC units reduce the private expenditures by more than the social costs. Similarly, increases in energy consumption observed for the low-consumption group lead to larger increases in private expenditures than in the social costs of supplying the additional electricity.

The differences between the private and social cost changes are caused by charging relatively high average per kWh rates, but the tiered price structure exaggerates the gap, especially among the high-consuming households. Assuming the SCC is $38/ton, the average private savings among the high-consuming households is 132% larger than the average social savings under the tiered rates. In contrast, under the TOU, CPP, and flat rate policies, private savings are only 66%, 69%, and 64% larger than the avoided social costs.

4.3. Discussion

Upgrading residential energy efficiency (e.g., installing energy efficient AC units) requires sizable upfront investments.[38] In return, the upgrades provide a variety of future expected benefits. From the perspective of a homeowner making the investment, these benefits include increased welfare ($\Delta U \geq 0$) from additional consumption of energy services (e.g., making their home cooler during the summer). Depending on the sign of the change in the household's electricity consumption (Δq), there could be an addi-

38. For example, a 2.5 ton AC unit with a SEER rating of 15—which would be eligible to receive a $400 rebate from SMUD—would cost around $1,600, and upward of $3,000 with installation included.

tional private benefit or cost. For households that reduce energy use ($\Delta q < 0$), there is an additional private benefit equal to the reduction in their monthly energy bill ($-p \cdot \Delta q$, where p represents the marginal retail price paid by the household).[39] In contrast, households that increase energy use face an additional private cost in the form of increased energy bills. In sum, the private benefit a household receives is equal to the utility gain less the expenditure changes, $\Delta U - p \cdot \Delta q$.

From society's perspective, the stream of potential benefits includes the same welfare increase from the increased consumption of energy services ($\Delta U \geq 0$). However, the magnitude of the social benefits, or costs, does not depend on the change in a household's energy bills. Instead, the social benefit, or cost, depends on the change in the resource and pollution costs incurred by altering the level of energy consumed. The social benefit can therefore be expressed as $\Delta U - c \cdot \Delta q$ where c is equal to the marginal social cost of supplying electricity.

Assuming the upfront private cost of installing an AC unit is equal to the fixed cost for society, a household will have an incentive to overinvest in energy efficiency (e.g., purchase an AC unit that is more energy efficient than is socially optimal) if the following is true:

$$\begin{aligned} 0 < (\text{Private benefit}) - (\text{Social benefit}) &= (\Delta U - p \cdot \Delta q) - (\Delta U - c \cdot \Delta q) \\ &= (c - p) \cdot \Delta q \,. \end{aligned} \tag{8}$$

Whether a household faces an incentive to over- or underinvest in energy efficiency depends not only on the difference between the marginal social cost and the marginal retail price but also on the expected change in energy consumption. Table 2 shows that the average low-consuming household increases its energy consumption ($\Delta q > 0$) following the AC upgrades. These households pay an average of 12.39 cents for each additional kWh of electricity they consume. In contrast, during each hour of the day, the median hourly wholesale prices during 2013 are below 6 cents per kWh. When factoring in the additional external costs from generating electricity, we still find that the marginal price paid by the low-consuming households exceeds the marginal social cost of supplying the additional electricity demanded ($c - p < 0$). Therefore, on average, the inequality in equation (8) does not hold for the average low-consuming household, that is, it faces an incentive to underinvest in energy efficiency.

---

39. To highlight how rates can incentivize over- or underinvestment in energy efficiency, we abstract from heterogeneity over time. In our empirical analysis, the change in consumption (Δq), the marginal social cost (c), and the marginal price (p) all vary across time for a given household. In such a setting, the expression $p \cdot \Delta q$ can be thought of as representing the sum across hours of the hourly price times the hourly consumption change. Similarly, $c \cdot \Delta q$ would represent the sum across hours of the hourly social marginal cost and the hourly change in consumption, and ΔU represents the total change in welfare resulting from the additional energy services consumed across all hours.

In contrast, the medium- and high-consuming households, which combined account for the majority of our sample, experience an average reduction in their electricity consumption ($\Delta q < 0$) following the AC upgrades. These households also face marginal retail prices that exceed the social marginal cost ($c - p < 0$). Medium-consuming households save an average of 17.3 cents for each kWh of consumption they reduce while high-consuming households save an average of 17.95 cents per kWh conserved. Thus, the inequality in equation (8) holds on average for medium- and high-consuming households. The time-invariant tiered rates are creating a clear incentive for the majority of households to overinvest in energy efficiency.

Our estimates of the social cost savings depend on the California market conditions during the sample period. As the stock of generating capacity changes, or as fuel prices change, the social marginal cost of producing electricity can change—which may in turn alter the social benefits provided by energy efficiency upgrades.[40] Nonetheless, if households continue to face time-invariant tiered rate structures that regularly charge marginal prices in excess of the social marginal cost, the privately optimal level of residential energy efficiency will exceed the socially optimal level.

By itself, our results do not necessarily imply that the actual level of private investment will be inefficiently high. If additional market failures contribute to the existence of an energy paradox, then investment in energy efficiency will lag behind the privately optimal level. However, even in the presence of an energy paradox, it is difficult to justify the incentives generated by the tiered rate structure on the grounds of economic efficiency. The available evidence reveals that the investment inefficiencies stemming from an energy paradox are unlikely to be large (Allcott and Greenstone 2012). In contrast to claims of widespread underinvestment in energy efficiency (e.g., McKinsey 2009), the observed low investment levels can instead be largely explained by factors including overstated potential energy savings (Metcalf and Hassett 1999; Graff Zivin and Novan 2016; Fowlie et al., forthcoming) as well as unaccounted for, hidden upgrade costs (Fowlie et al. 2015). These findings suggest that, in the setting we examine, the incentive to overinvest in energy efficiency generated by the current rate structure—which leads to private savings that exceed the social savings by over 130%—is inefficiently large.

Even if we were to assume that a sufficiently large energy paradox exists, the incentives created by the tiered rate structure will not efficiently address the cause of the underinvestment. The energy paradox literature attributes the cause to market failures such as imperfect information about potential savings, principal-agent problems, and

40. Relative to 2013, the 2014 CAISO (California Independent System Operator) day ahead prices increased on average slightly during the peak hours (http://www.caiso.com/Documents/2014AnnualReport_MarketIssues_Performance.pdf). However, from 2015 and on, the wholesale prices have been falling—driven in part by a large increase in solar capacity in the region. This suggests that the social benefits provided by the previous energy efficiency investments have fallen.

liquidity constraints (e.g., Gillingham and Palmer 2014). Rather than directly addressing these potential market failures, however, high marginal prices do not necessarily target the customers underinvesting in energy efficiency.[41] For example, we may expect that an energy paradox results in the greatest underinvestment among low-income or credit-constrained households. However, the incentive to invest in energy efficiency created by the high marginal prices does not target this set of consumers—it is provided to all households that reduce consumption and is largest for high-consuming households, which, on average, have higher incomes (Borenstein 2012).

Although our results reveal that SMUD's standard rate structure—a fixed monthly fee of $12 combined with a two-tiered, increasing block per kWh charge—creates an incentive for the majority of households to overinvest in energy efficiency, customers throughout the rest of California have faced even more extreme tiered rate schedules with higher marginal prices. For example, in 2015, Pacific Gas and Electric (PG&E)—the nation's largest investor-owned electric utility in terms of number of customers—charged a marginal price of $0.33/kWh for households in the top consumption tier.[42] As a result, California households outside of SMUD's territory have historically received even stronger incentives to overinvest in energy efficiency.

The residential rate structures charged by California's investor-owned utilities (IOUs) are undergoing changes. Following a recent ruling by the California Public Utilities Commission (CPUC 2015b), the IOUs are transitioning toward residential rate structures very similar to the SMUD rate structure examined in this paper.[43] While these changes can be expected to reduce the incentive to overinvest in energy efficiency, our analysis using the SMUD rate structure reveals that the general incentive to overinvest will remain.

Looking farther down the road, the CPUC has also mandated that California's IOUs begin charging a default time-of-use (TOU) rate plan for residential consumers by 2019.[44] Our results highlight that TOU rates, which remove the tiered structure currently in place, will substantially reduce the incentive to overinvest in energy efficiency

41. Allcott et al. (2015) highlight this inefficiency, which is inherent in direct energy efficiency subsidies as well.

42. Similar to SMUD, California's big three investor-owned utilities (IOUs)—that is PG&E, SCE (Southern California Edison Company), and SDG&E (San Diego Gas and Electric Company)—each receive the majority of their revenue from the per kWh charges as opposed to fixed fees. At the beginning of 2015, PG&E customers on the standard tiered rate schedule paid a minimum of $4.44/month while SDG&E customers paid a minimum of $5.10/month. If not binding, the minimum charges serve as fixed fees. SCE customers on the standard residential tiered rate plan paid a fixed charge of $0.93/month.

43. Specifically, the CPUC has mandated that the IOUs move to a two-tiered rate schedule with fixed charges of $10 per month.

44. Moreover, on August 15, 2013, the SMUD board of directors voted to phase out the residential tiers by 2017 and to begin preparing for time-of-use rates.

upgrades. However, if the TOU rate plans continue to include low fixed monthly fees, and therefore generate revenue mainly from the per kWh charges, then our estimates highlight that, for most households, a financial incentive to overinvest in energy efficiency will remain, on average.

Inefficient electric rate structures are not unique to California. Tiered prices have become commonplace across the United States, and utilities regularly rely on variable (per kWh) charges to generate a large share of the revenue required to cover their costs. In 2015, the median fixed charge in the United States was $10 per month, similar to the low fixed fees charged by California utilities.[45] With such little revenue generated by the fixed monthly fees, utilities are forced to charge inefficiently high variable rates. Our results provide support for moving toward residential energy rate structures that bring the marginal prices down, aligning them with the social marginal costs. Doing so requires not only removing tiered rates but also, in many cases, increasing the share of revenue generated by fixed fees.

5. CONCLUSION

In this paper, we compare the private and social cost savings achieved by investing in residential energy efficiency. Using hourly, household-level smart-meter data from Sacramento, we estimate not only how much electricity households save by installing new, energy efficient AC units but also when the energy savings occur. We predict how much the new AC units reduce participating households' monthly electricity bills, which we compare to the reduction in resource and environmental cost savings from reduced electricity generation.

After installing an energy efficient AC unit, a participating household's nonwinter electricity consumption fell by an average of 1.36 kWh per day—approximately a 5% reduction in total consumption. Under the current residential electricity rate structure, which combines low fixed monthly fees with tiered prices, the participating households saved an average of $6.46 per month. In contrast, the estimated social cost savings were only $2.69 per month. The gap between the private and social cost savings is even more pronounced among the households with high historical levels of electricity consumption—the homes that account for the vast majority of the energy savings. On average, a high-consuming household's monthly electricity bill fell by $26.56 while the corresponding social costs fell only by $11.44 per month, on average. Thus, by setting the marginal price above the marginal social cost of supplying electricity, the current residential rate structure leads to an outcome in which the privately optimal level of investment in energy efficiency exceeds the socially optimal level.

45. Information on the rate codes is available from the OpenEI project (http://en.openei.org/apps/USURDB/). We focus on the single phase, residential rates as of January 1, 2015. The individual rates are not weighted by the share of customers on each pricing plan—information that is not available.

In many regions, regulators and policy makers are actively debating making changes to existing residential energy rate structures. In California in particular, the Public Utilities Commission has recently mandated that the investor-owned utilities begin to charge flatter tiered electricity prices with slightly larger fixed fees. However, the results presented in this paper demonstrate that, even after these changes, most homeowners will still have an incentive to overinvest in energy efficiency.

REFERENCES

Allcott, Hunt, and Michael Greenstone. 2012. Is there an energy efficiency gap? *Journal of Economic Perspectives* 26 (1): 3–28.

Allcott, Hunt, Christopher Knittel, and Dmitry Taubinsky. 2015. Tagging and targeting of energy efficiency subsidies. *American Economic Review* 105 (5): 187–91.

BC-Hydro. 2008. 2008 residential inclining block rate application. Technical report, British Columbia Utilities Commission.

Borenstein, Severin. 2012. The redistributional impact of nonlinear electricity pricing. *American Economic Journal: Economic Policy* 4 (3): 56–90.

———. 2017. Private net benefits of residential solar PV: The role of electricity tariffs, tax incentives, and rebates. *Journal of the Association of Environmental and Resource Economists* 4 (S1): S85–S122.

Borenstein, Severin, James Bushnell, Frank A. Wolak, and Matthew Zaragoza-Watkins. 2015. Expecting the unexpected: Emissions uncertainty and environmental market design. NBER Working Paper, National Bureau of Economic Research, Cambridge, MA.

Borenstein, Severin, and Lucas W. Davis. 2012. The equity and efficiency of two-part tariffs in US natural gas markets. *Journal of Law and Economics* 55 (1): 75–128.

———. 2016. The distributional effects of US clean energy tax credits. *Tax Policy and the Economy* 30 (1): 191–234.

Callaway, Duncan, Meredith Fowlie, and Gavin McCormick. 2018. Location, location, location: The variable value of renewable energy and demand-side efficiency resources. *Journal of the Association of Environmental and Resource Economists* 5 (1): 39–75.

Carson, Richard T., and Kevin Novan. 2013. The private and social economics of bulk electricity storage. *Journal of Environmental Economics and Management* 66 (3): 404–23.

Coase, Ronald H. 1946. The marginal cost controversy. *Economica* 13 (51): 169–82.

Cohen, Michael, Paul Kauzmann, and Duncan Callaway. 2016. Effects of distributed PV generation on California's distribution system, part 2: Economic analysis. *Solar Energy* 128:139–52.

Cooper, Adam. 2016. Electric company smart meter deployments: Foundation for a smart grid. Technical report, Institute for Electric Innovation, Edison Foundation, Washington, DC.

CPUC. 2015a. The 2013–2014 resource adequacy report. Technical report, Public Utilities Commission of the State of California.

———. 2015b. Decision on residential rate reform for Pacific Gas and Electric Company, Southern California Edison Company, and San Diego Gas and Electric Company and transition to time-of-use rates. Technical report, Public Utilities Commission of the State of California.

Davis, Lucas W., Alan Fuchs, and Paul Gertler. 2014. Cash for coolers: Evaluating a large-scale appliance replacement program in Mexico. *American Economic Journal: Economic Policy* 6 (4): 207–38.

Davis, Lucas W., and Erich Muehlegger. 2010. Do Americans consume too little natural gas? An empirical test of marginal cost pricing. *RAND Journal of Economics* 41 (4): 791–810.

Fowlie, Meredith, Michael Greenstone, and Catherine Wolfram. 2015. Are the non-monetary costs of energy efficiency investments large? Understanding low take-up of a free energy efficiency program. *American Economic Review* 105 (5): 201–4.

———. Forthcoming. Do energy efficiency investments deliver? Evidence from the Weatherization Assistance Program. *Quarterly Journal of Economics*.

Gerarden, Todd D., Richard G. Newell, and Robert N. Stavins. 2017. Assessing the energy-efficiency gap. *Journal of Economic Literature* 55 (4): 1486–1525.

Gillingham, Kenneth, and Karen Palmer. 2014. Bridging the energy efficiency gap: Policy insights from economic theory and empirical evidence. *Review of Environmental Economics and Policy* 8 (1): 18–38.

Graff Zivin, Joshua, and Kevin Novan. 2016. Upgrading efficiency and behavior: Electricity savings from residential weatherization programs. *Energy Journal* 37 (4): 1–23.

Jacobsen, Grant D. 2014. Estimating end-use emissions factors for policy analysis: The case of space cooling and heating. *Environmental Science and Technology* 48 (12): 6544–52.

Jacobsen, Grant D., and Mathew J. Kotchen. 2013. Are building codes effective at saving energy? Evidence from residential billing data in Florida. *Review of Economics and Statistics* 95 (1): 34–49.

McKinsey. 2009. Unlocking energy efficiency in the US economy. http://www.mckinsey.com/client_service/electric_power_and_natural_gas/latest_thinking/unlocking_energy_efficiency_in_the_us_economy.

Metcalf, Gilbert E., and Kevin A. Hassett. 1999. Measuring the energy savings from home improvement investments: Evidence from monthly billing data. *Review of Economics and Statistics* 81 (3): 516–28.

Reiss, Peter C., and Matthew W. White. 2008. What changes energy consumption? Prices and public pressures. *RAND Journal of Economics* 39 (3): 636–63.

Siler-Evans, Kyle, Ines Lima Azevedo, and M. Granger Morgan. 2012. Marginal emissions factors for the US electricity system. *Environmental Science and Technology* 46 (9): 4742–48.

SMUD. 2014. Smart pricing options final evaluation. Technical report, Sacramento Municipal Utility District.