Online Appendix to "Do Energy Efficiency Investments Deliver at the Right Time?"

by Judson Boomhower and Lucas Davis

A Electricity Market Data

A.A Wholesale Electricity Prices and Load

Hourly wholesale electricity price and load data are from SNL Financial and are for 2011–2015. For California, we use CAISO day-ahead prices at the SP-15 node. For New England, we use ISO-NE day-ahead prices at the H Internal hub. For Texas, we use ERCOT day-ahead prices at the HB North hub. For New York, we use NYISO day-ahead prices at the J Hub. For PJM, we use day-ahead prices at the Western hub. For MISO, we use day-ahead prices at the Illinois hub. All times in the paper are reported in local prevailing time: Standard Time or Daylight Time according to which is in effect. The load data in each market come from the SNL hourly "Actual Load" series for 2011–2015. Appendix Figure 1 plots hourly average load profiles by month-of-year for each market.

A.B Capacity Prices

Capacity values were calculated under a range of assumptions. For each market, we used auction or regulatory data to infer monthly or annual capacity prices, and allocated those values across hours based on historical load. Capacity market institutions vary across regions, so capacity values are not perfectly comparable across markets. However, we have attempted to use relatively comparable data and methods and to be transparent about our sources and calculations.

ERCOT has no capacity market so capacity values are equal to zero in all hours. In all other markets, generation capacity is procured in advance at the monthly or annual level, and capacity contracts obligate generators to be available every hour during that period. Specifically, California (CAISO) and New York (NYISO) have monthly contracts, whereas the Midwest (MISO), Mid-Atlantic (PJM), and New England (ISONE) have annual contracts. In order to value energy savings in a given hour, we need to allocate these capacity prices across individual hours. We do this several ways and report the results of each. The amount of capacity to be purchased each period is determined by the regulator's forecast of peak demand. If the precise hour of the peak could be predicted with certainty, that one hour would have capacity costs equal to the contract price, and capacity costs for all other hours would be zero. Changing demand in any of these other hours would have no effect on the capacity market. In reality, it is impossible to perfectly predict the day on which the peak will occur because of uncertainty in weather and other factors. The expected capacity value of a one MWh demand reduction in any hour is equal to the capacity price times the probability that that hour will be the peak hour. Our various approaches to allocating capacity value involve different ways of calculating these probabilities.

For markets with monthly capacity contracts, we start by using hourly load data to calculate the hour-of-the-day with the highest average load each month. We then divide the monthly capacity price evenly across all occurrences of that hour-of-day on weekdays. We allocate capacity costs to weekdays only, because weekend and holiday loads are reliably smaller. This approach assigns capacity values to the top 3% of all hours in each month, see column (2) of Table 2. For the alternative approaches, in columns (3) and (4), we divide the capacity contract price evenly over the top two or three hours-of-the-day with the highest load each month. The final approach in Column (5) treats each day of load data as a single observation of daily load shape in a given month. We calculate the likelihood between 2011 and 2015 that each hour-of-the-day was the daily peak hour, and allocate monthly capacity values to hours of the day proportionally according to these probabilities. For example, during

February in the CAISO market, 6:00 p.m. was the highest-demand hour on 92% of days from 2011–2015. Consequently, we assign 92% of the February contract price to the 6:00 - 7:00 p.m. hour.

For markets with annual capacity contracts, our calculations are very similar, except we assign capacity values to the highest- load hours of the year, rather than to the highest-load hours of the month. Specifically, we allocate annual capacity payments to the top 36 hour-of-day by month-of-year pairs, equivalent to about 6% of all hours throughout the year.

We adjust for reserve margins in all calculations. For every unit of forecast peak demand, regulators require more than one unit of forward capacity purchases (the difference being the required reserve margin). California's reserve margin is 15%, and other markets are similar. Therefore, we increase all capacity values by 15% to reflect that each unit of demand reduction reduces capacity requirements by 1.15 units.

A.B.1 Capacity Market Data

California (CAISO)

CAISO differs from the other markets in that capacity is procured through bilateral contracts, rather than through a centralized auction. The California Public Utilities Commission (CPUC) surveys utilities to track capacity contract prices. We use monthly capacity contract prices from the CPUC "2013–2014 Resource Adequacy Report," page 28, Table 13. This document reports average, 85th-percentile, and maximum contract prices for each month. We use the 85th-percentile values, on the reasoning that these provide a conservative estimate of the marginal cost of procuring capacity. We could instead use the maximum, but choose the 85th percentile to limit the influence of potential outlier observations. These reported prices include capacity contracts from 2013 through 2017, though most of the reported transactions are for 2013–2015 (page 29, Figure 9).

New York (NYISO)

Capacity prices for New York come from SNL Financial and are for NYISO's monthly spot capacity auctions for the NYCA region from May 2013 through April 2016. This auction runs two to four days prior to the beginning of the month being transacted for. NYISO also runs auctions for six-month "strips" of summer or winter capacity, as well as additional monthly auctions one to five months in advance.

New England (ISO-NE)

Capacity prices for New England come from SNL Financial and are for ISO-NE's annual forward capacity auctions for 2013 through 2016. We use the simple average of prices across all zones.

Mid-Atlantic (PJM)

Capacity prices for PJM are from SNL Financial and are market clearing prices from the annual Base Residual Auction. We use the simple average across years and geographic zones for 2013–2016.

Midwest (MISO)

Capacity prices for MISO are from SNL Financial and are annual capacity auction prices for 2013 through 2016. We use the simple average of prices across all zones.

Appendix Figure 1: Load Profiles in Six Major U.S. Electricity Markets

(a) California (CAISO)

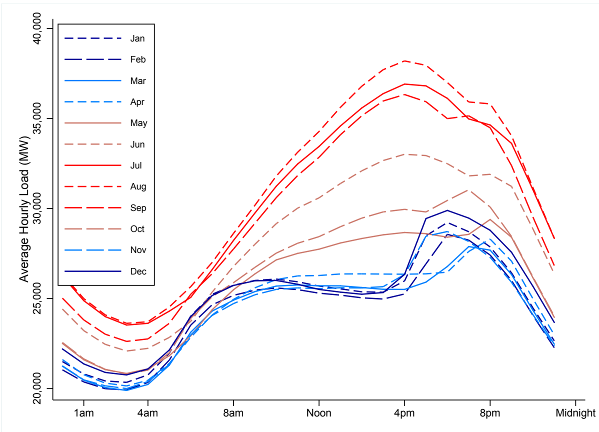

(b) Texas (ERCOT)

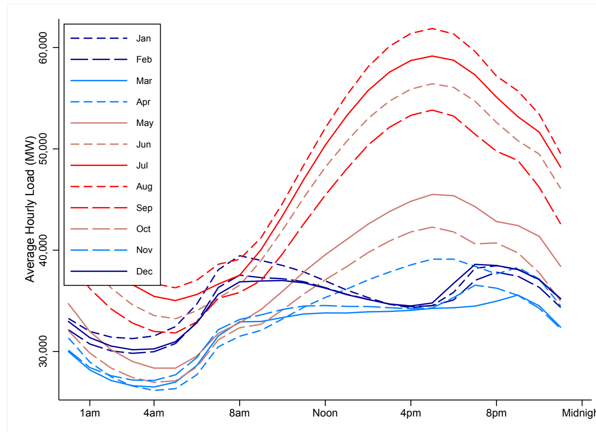

(c) Mid-Atlantic (PJM)

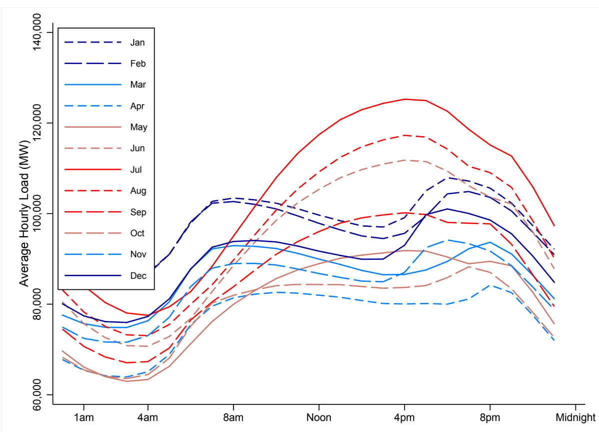

(d) Midwest (MISO)

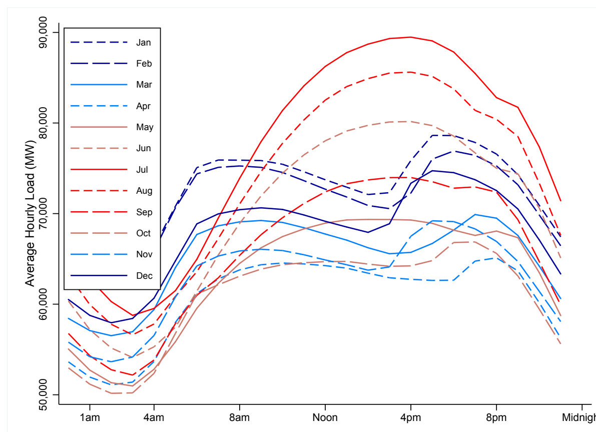

(e) New York (NYISO)

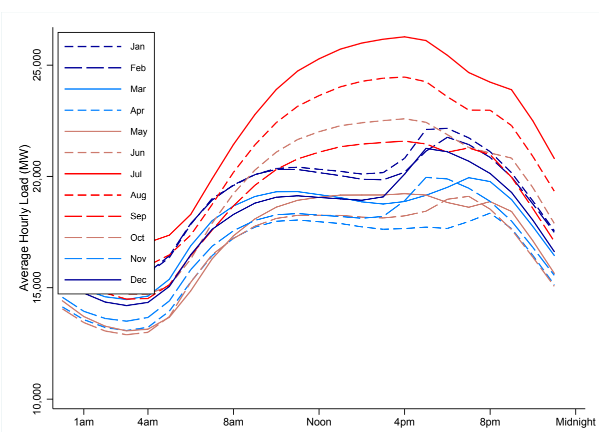

(f) New England (ISONE)

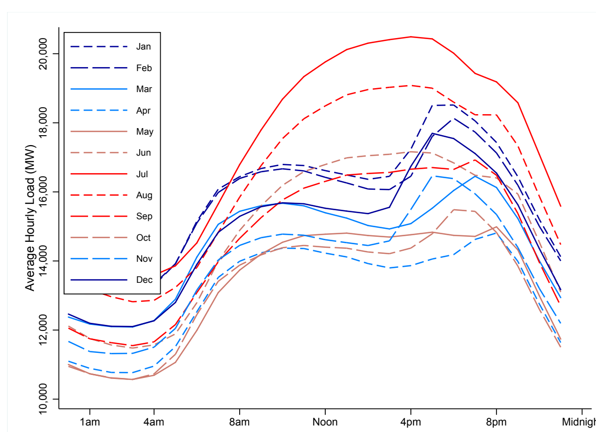

B Additional Data Description

B.A Program Data

The program data describe all 10,848 households who participated in the *Quality Installation Program* program between 2010 and 2015. These data were provided by Southern California Edison. We drop 968 duplicate participant records. These records have the exact same account number as other participant records, so are clear duplicates. We also drop an additional 291 households who installed a new heat pump rather than a new central air conditioner; the expected energy savings for heat pumps follows a very different temporal pattern than the temporal pattern for air conditioning so it does not make sense to include these participants. We further drop 2,431 households who participated before the start of 2012; we use electricity consumption data beginning in 2012, so these early participants would not contribute to our savings estimates. We also drop an additional 757 households who installed rooftop solar at any time during our sample period; rooftop solar dramatically changes household net electricity consumption (we only observe net consumption, not generation and consumption separately) so we drop these households to avoid biasing our savings estimates. In addition, we drop 60 households for whom we do not have a nine-digit zip code; a nine-digit zip code is required for merging with temperature data, and we cluster all standard errors at the nine-digit zip code. We successfully merged 94% of the participant records to the electricity consumption data, so we are left with a total of 5,973 participants in our analysis dataset. Appendix Figure 2 shows the pattern of participation between 2012 and 2015.

B.B Electricity Consumption Data

The electricity consumption data describe hourly electricity consumption for all program participants. We were provided with the complete history of hourly consumption for these households beginning when each household received a smart meter and continuing until August 2015, or, in some cases, somewhat before August 2015. Most Southern California Edison customers received a smart meter for the first time in either 2011 or 2012. Appendix Figure 3 shows the number of participants with smart meter billing data during each week of the sample.

B.C Engineering-Based Savings Profiles

Appendix Figure 4 plots savings profiles for eight additional energy-efficiency investments. These figures are constructed in exactly the same way as Figure 6, and describe five residential investments and three commercial/industrial investments. As described in the paper, these engineering-based savings profiles come from the Database for Energy Efficient Resources (DEER), maintained by the California Public Utilities Commission. We use values developed in 2013/2014 for DEER 2011, reported in the file DEER2011-HrlyProfiles-SCE.xls. For each energy-efficiency investment the DEER reports 8,760 numbers, one for each hour of the year. We use these data to construct average hourly profiles by month. These savings profiles are intended to reflect average impacts in Southern California Edison territory.

The underlying model that generates the DEER hourly profiles does not account for daylight savings time. Building occupants are assumed to observe Standard Time for the full year. As a result, the model inputs for physical phenomena such as solar angle and temperature are correct, but inputs related to human schedules, like building opening times, are "off" by one hour. Some analysts adjust for daylight savings after the fact by "shifting" the DEER profile one hour: that is, replacing predicted savings for all hours during Daylight Time with predicted savings one hour later. This corrects building schedules but introduces error in physical factors. Whether such a shift helps or hurts accuracy depends on whether building schedules or physical factors are more important in determining hourly savings. We do not make any adjustments to the DEER profiles in our main specifications. If we do impose a "shift" during Daylight Time, the estimated timing premiums for DEER investments change only slightly.

Appendix Figure 2: Histogram of Installation Dates

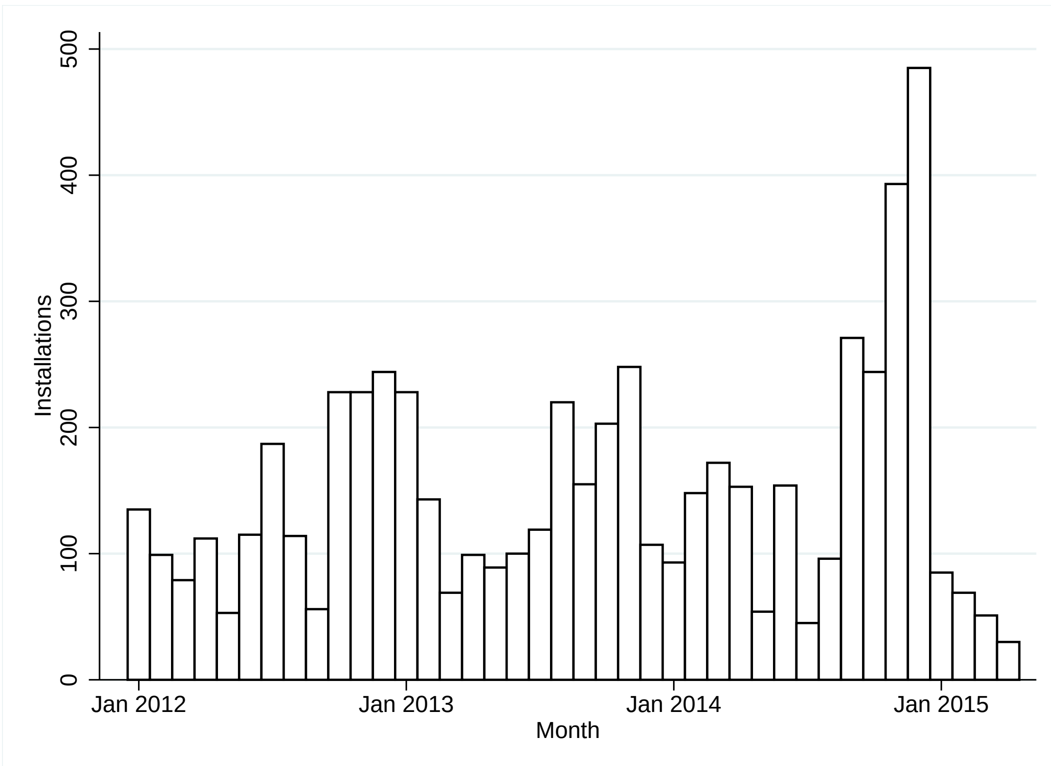

Appendix Figure 3: Number of Participants with Smart Meter Data

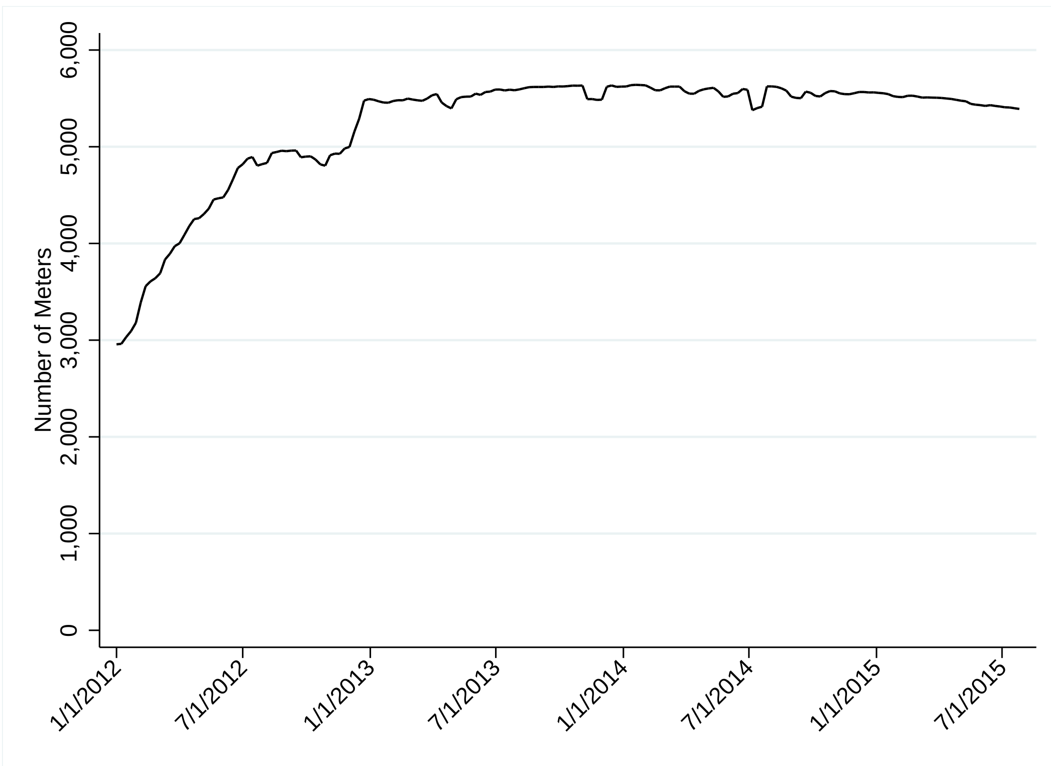

Appendix Figure 4: Savings Profiles for Additional Energy-Efficiency Investments

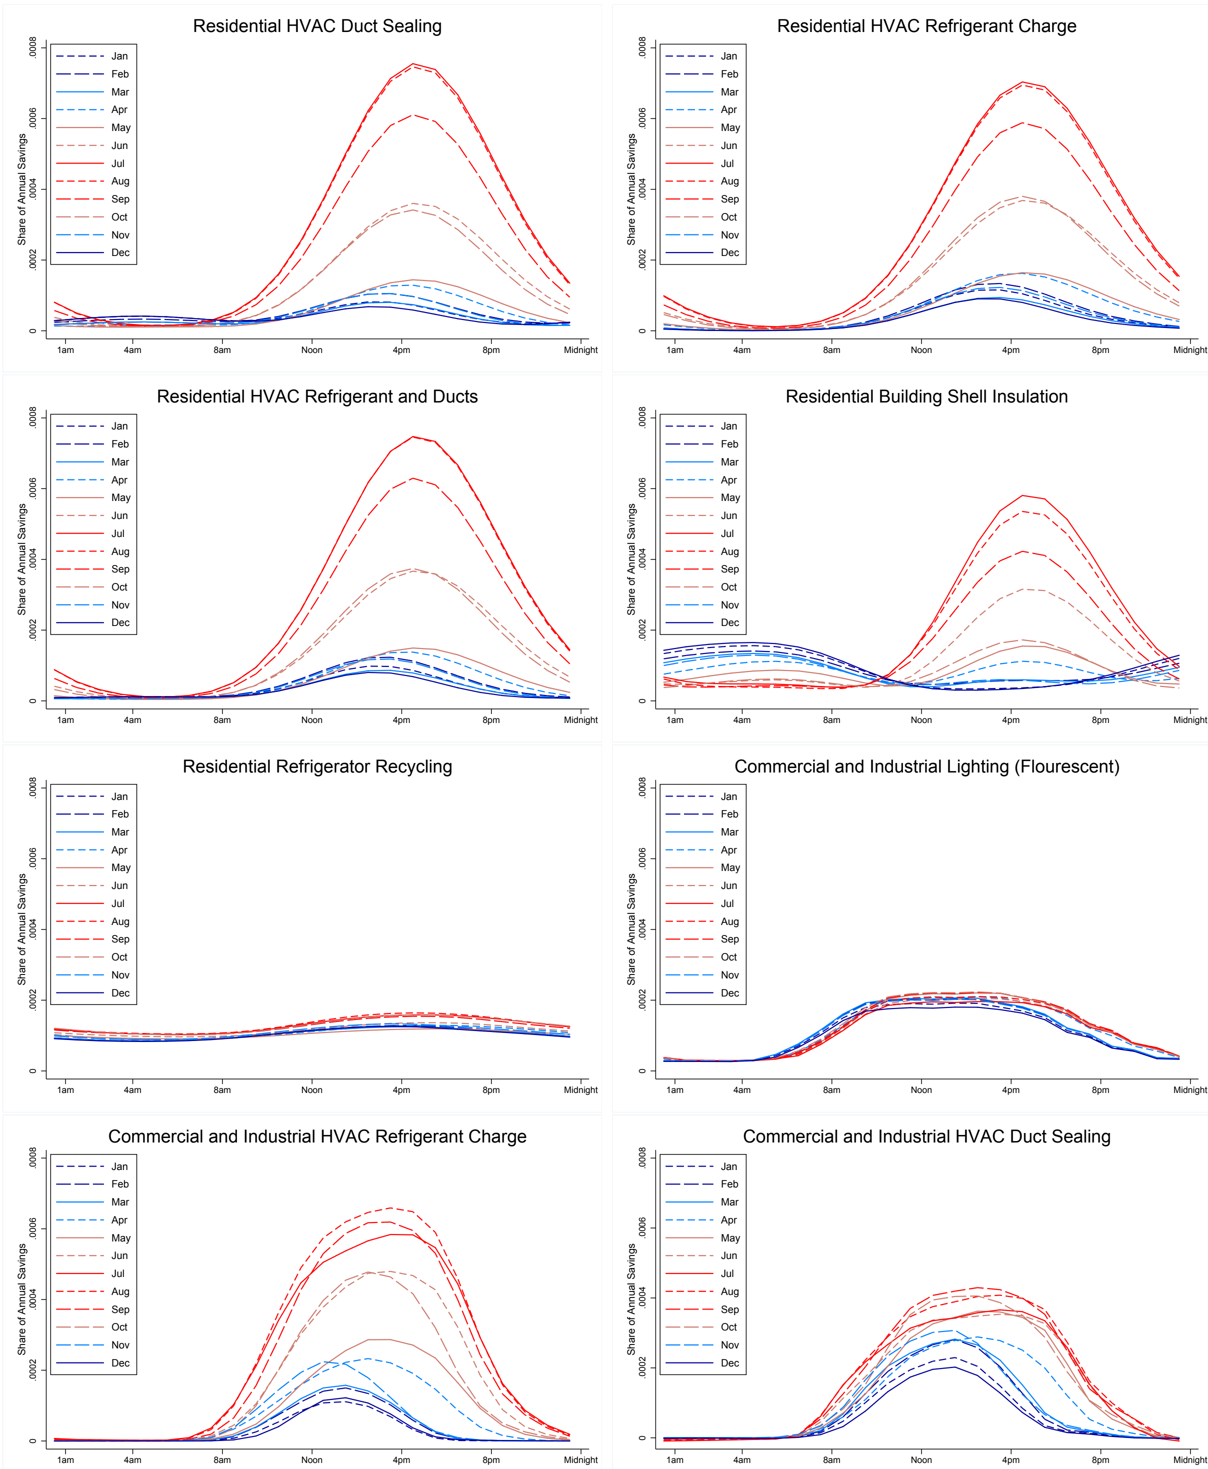

C Alternative Specifications Using Data from Non-Participants

This section presents estimates from alternative specifications which incorporate electricity consumption data from non-participating households. Overall, these alternative estimates are quite similar to the main estimates in the paper.

The key challenge in our empirical analysis is to construct a counterfactual for how much electricity the participants would have consumed had they not installed a new air conditioner. The analyses in the paper construct this counterfactual using data from participants only, exploiting the natural variation in the timing of program participation to control for trends in electricity consumption, weather, and other time-varying factors. An alternative approach, however, is to estimate the model using data from both participants and non-participants.

There are advantages and disadvantages with this alternative approach. The potential advantage of including non-participant data is that these data may help better control for trends in electricity consumption, weather, and other time effects, while also potentially improving the precision of the estimates. The disadvantage is that non-participants tend to be quite different from participants, making them potentially a less valid counterfactual. Without any *ex ante* reason to prefer one approach over the other, it makes sense to report estimates from both approaches.

Appendix Table 1 provides descriptive statistics. The columns refer to three different samples. The first column describes the 5,973 participants used for the main estimates in the paper. The second column describes a random sample of non-participants. We were provided with data from a 5% random sample of Southern California Edison residential customers who did not participate in the program, and this is a random subset of 5,973 households from that sample. Finally, the third column describes a matched sample of non-participants. For the matched sample we selected non-participants based on zip codes. In particular, for each participant, we randomly selected a non-participant from the same nine-digit zip code, or five-digit zip code when nine-digit zip code is not available. Weather is a major determinant of electricity consumption so this matching ensures that comparison households are experiencing approximately the same weather as the treatment households. In addition, households in close geographic proximity tend to have similar income and other demographics. Some non-participants matched to more than one participant, yielding 5,643 unique households in the matched sample of non-participants. For both random and matched samples we excluded households with rooftop solar or a missing nine-digit zip code, just as we did for participants.

Across all households, mean hourly electricity consumption is about one kWh per hour. Participants tend to consume more than non-participants, especially during summer months. But this appears to be largely a question of geography and the pattern of consumption in the matched sample is much more similar to participants. More generally, the characteristics of the matched sample are more similar but not identical to the characteristics of participants. Among participants, 13% are on the low-income tariff, compared to 30% in the random sample and 25% in the matched sample. Similarly, only 2% of participants are on the all-electric tariff, compared to 10% in the random sample and 6.1% in the matched sample.

We used these alternative samples to construct alternative estimates of several of our main results. We begin in Appendix Figure 5 by showing an event study for winter months using an identical sample of households as in Figure 1. This event study figure was constructed in exactly the same way as Figure 1 but using data from January and February, and excluding data from installations that occurred during February, March, or April. As expected, winter consumption is essentially unchanged after the new air conditioner is installed. This suggests that the sharp drop in electricity consumption during summer is indeed due to the new air conditioner and not some other unrelated change in household appliances or behavior.

Appendix Figure 6 moves on to show event study estimates using alternative household samples. Whereas the event study figure in the paper is estimated using data from participants only, these are estimated using data from both participants and non-participants. The plots on the top include the random sample of non-participants while the plots on the bottom include the matched sample. These alternative event studies follow a very similar pattern to the event study figures that include only participants. Summer consumption drops sharply in the year that the new air conditioners are installed and the magnitude of this decrease is 0.2 kWh/hour, identical to the decrease in the event study figure in the paper. Moreover, the pattern for winter is very similar, with no change when the new air conditioners are installed.

Next, Appendix Table 2 reports regression estimates of total energy savings from new air conditioner installation. This table is constructed in exactly the same way as Table 1, but estimated using data from both participants and non-participants. Including data from non-participants has little overall effect. The estimates are slightly larger, but the pattern across specifications is similar, increasing when dropping eight weeks pre-installation in Column (3).[18]

Finally, Appendix Figure 7 plots estimates of energy savings by month-of-year and hour-of-day. These figures are constructed in exactly the same way as Figure 5, but are estimated using data from both participants and non-participants. Overall, including data from non-participants has very little effect on the temporal pattern of savings. Electricity savings still tend to occur disproportionately during July and August, and during the hours 3 p.m. to

[18]An alternative specification for both Table 1 and Appendix Table 2 would be to use a single new appliance indicator variable to measure average savings across all hours, and then multiply by the number of hours in a year. However, this approach incorrectly weights hours of the year according to the composition of post-installation observations for each household. For example, since our data end in April 2015, a household that installed in late 2013 would be observed for two winters and one summer. This uneven weighting could potentially be addressed by re-weighting or restricting the sample to include exactly one year of post-installation data for each household and throwing out installations after April 2014; or by re-weighting across the sample to equalize the effective number of post-installation observations. We prefer to simply estimate average savings for each hour-of-day by month-of-year pair and sum up to annual savings. Moreover, we need these 288 separate estimates for the analyses elsewhere in the paper.

9 p.m. In addition, during winter months the estimates remain very close to zero during all hours of the day. Moreover, the random and matched samples yield virtually identical estimates across hours and months.

Appendix Figure 5: Event Study Figure, Winter Electricity Consumption

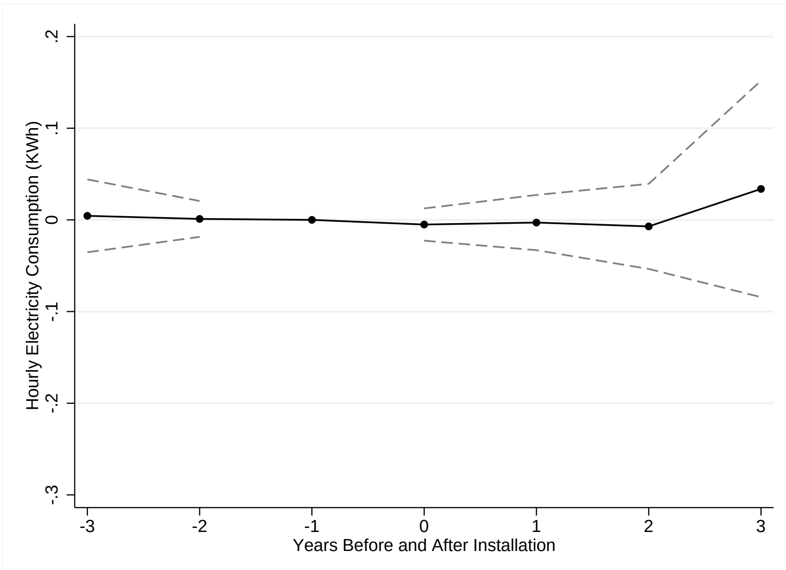

Notes: This event study figure plots estimated coefficients and 95% confidence intervals from a least squares regression. The dependent variable is average hourly electricity consumption during January and February at the household by year level. Time is normalized relative to the year of installation ($t = 0$) and the excluded category is $t = -1$. The regression includes year by climate zone fixed effects. Standard errors are clustered by nine-digit zip code.

Appendix Figure 6: Event Study Figures, Alternative Specifications

Random Sample of Non-Participants

Summer

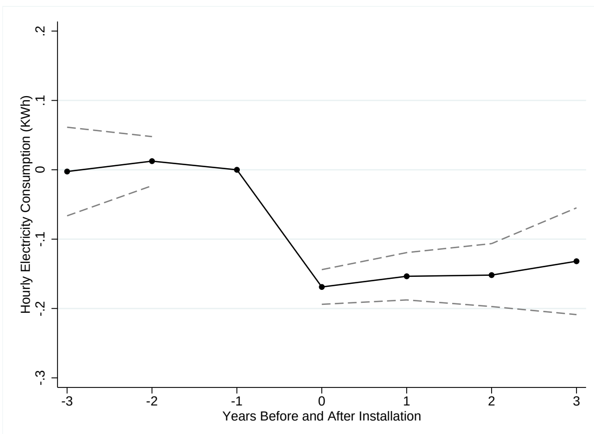

Winter

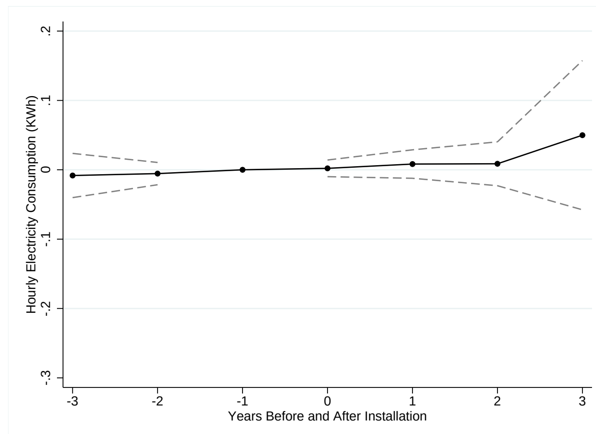

Matched Sample of Non-Participants

Summer

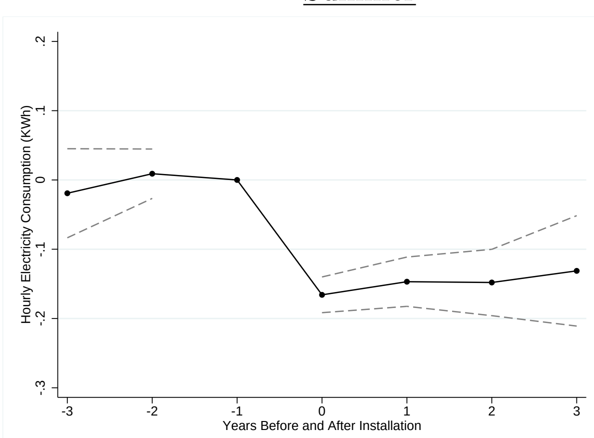

Winter

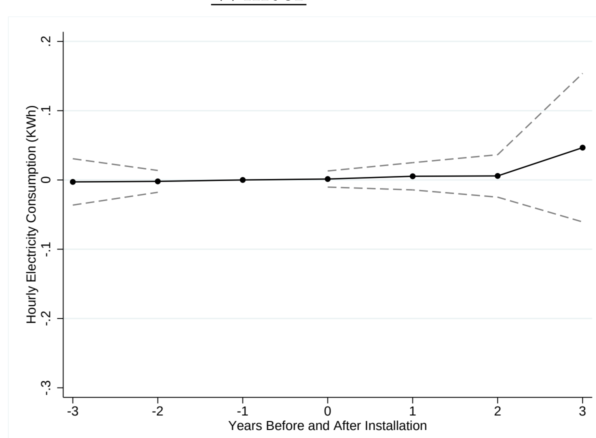

Appendix Figure 7: Econometric Estimates of Electricity Savings, Alternative Specifications

Random Sample of Non-Participants

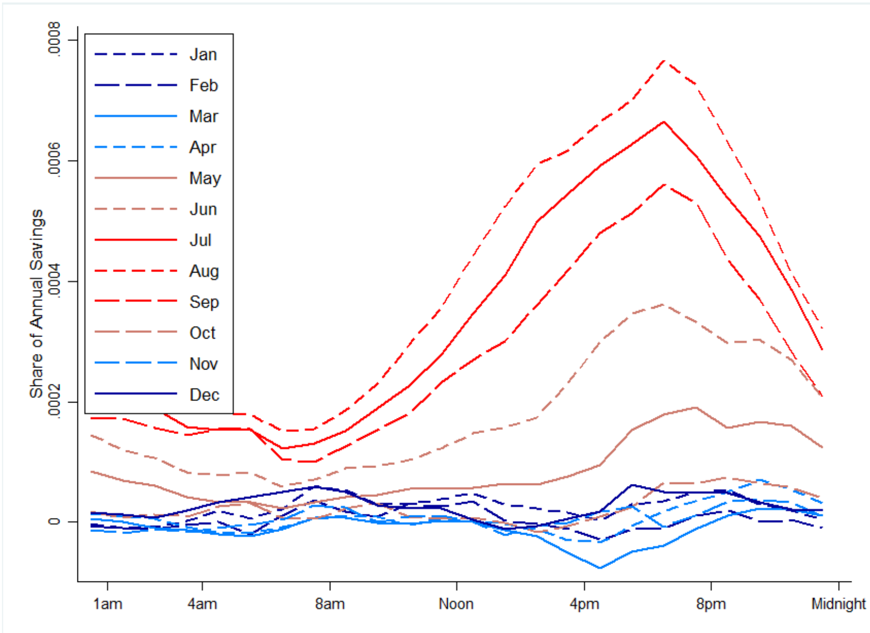

Matched Sample of Non-Participants

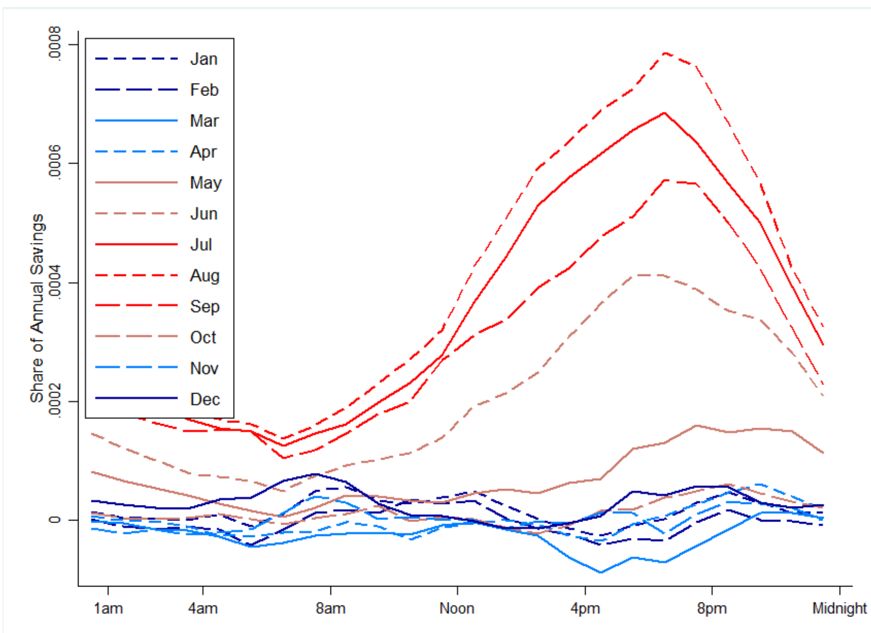

Appendix Table 1: Smart Meter Data, Descriptive Statistics

	Participants (1)	Random Sample of Non-Participants (2)	Matched Sample of Non-Participants (3)	p-Value: Column 1 vs Column 2 (4)	p-Value: Column 1 vs Column 3 (5)
Mean Hourly Electricity Consumption					
All Months	1.076	0.878	0.998	0.000	0.000
Summer Months (July and August)	1.521	1.205	1.435	0.000	0.000
Winter Months (January and February)	0.852	0.729	0.789	0.000	0.000
Type of Electricity Tariff					
Proportion on Low-Income Tariff	0.128	0.303	0.247	0.000	0.000
Proportion on All-Electric Tariff	0.020	0.101	0.061	0.000	0.000

Notes: Columns (1), (2), and (3) report the variables listed in the row headings for the group listed at the top of the column. There are a total of 5,973 participants and an equal number of non-participating households in the random and matched samples. Columns (4) and (5) report p-values from tests that the means in the subsamples are equal.

Appendix Table 2: Average Energy Savings, Alternative Specifications

	(1)	(2)	(3)
Random Sample of Non-Participants			
Energy Savings Per Household (kWh/year)	502.9	440.6	512.2
	(45.4)	(46.5)	(51.6)
Number of observations	27.0 M	27.0 M	26.4 M
Number of households	5,976	5,976	5,975
Matched Sample of Non-Participants			
Energy Savings Per Household (kWh/year)	399.7	387.9	453.8
	(47.3)	(47.7)	(52.7)
Number of observations	27.0 M	27.0 M	26.4 M
Number of households	5,887	5,887	5,886
Household by hour-of-day by month-of-year fixed effects	Y	Y	Y
Week-of-sample by hour-of-day fixed effects	Y		
Week-of-sample by hour-of-day by climate zone fixed effects		Y	Y
Drop 8 weeks pre-installation			Y

Notes: This table reports results from six separate regressions and is identical to Table 1 in the paper except for the sample includes data on non-participating households. In particular, Panel A includes a random sample of non-participating households and Panel B includes a matched sample of non-participating households in which the non-participating households are drawn from the same nine digit zip code as participating households. For computational reasons, we restrict these regressions to a 50% random sample of participating households along with an equal number of non-participating households.

D Visualizing the Correlation between Savings and Value

As a graphical complement to Table 2, Appendix Figure 8 shows the correlation between energy savings and the value of energy. Panel A compares hourly average energy savings to energy prices only. Panel B compares the same savings estimates to the sum of energy and capacity values. Each marker in each plot corresponds to an hour-of-day by month-of-year pair (for example, 1:00–2:00 p.m. during November). The vertical axes show average hourly energy savings. These are the 288 coefficients from the regression described in Section II.E. In Panel A, the horizontal axis shows average wholesale energy prices from California for 2011–2015. In Panel B, the horizontal axis shows energy and capacity values, using the probabilistic allocation method for capacity prices described in Section III.A.

Several facts are apparent in Panel A. First, the summer months include many more high-price realizations than the winter months. We use dark markers to indicate April through September, and the number of intervals with energy prices above $40/MWh is much higher during these summer months. Second, this energy-efficiency investment delivers much larger savings in the summer. We saw this before in Figure 3, with average savings in excess of 0.1 kWh/hour in most summer hours.

The figure also includes least-squares fitted lines. The fitted line for April–September slopes steeply upward. In Panel A, predicted savings when energy prices are $55/MWh are twice as large as predicted savings at $35/MWh. The fitted line for winter, in contrast, is essentially flat. Savings are near zero in all winter hours, so there is little correlation between savings and price.

The same patterns are apparent in Panel B, but this panel emphasizes the importance of accounting for capacity values. During a few ultra-peak hours in the summer, generation capacity is extremely valuable and the value of energy surges to above $200/MWh. Air conditioner investments deliver above-average savings in these hours, so the correlation is again strongly positive.

Appendix Figure 8: Correlation Between Savings and Prices, By Season

Panel A. Energy Prices Only

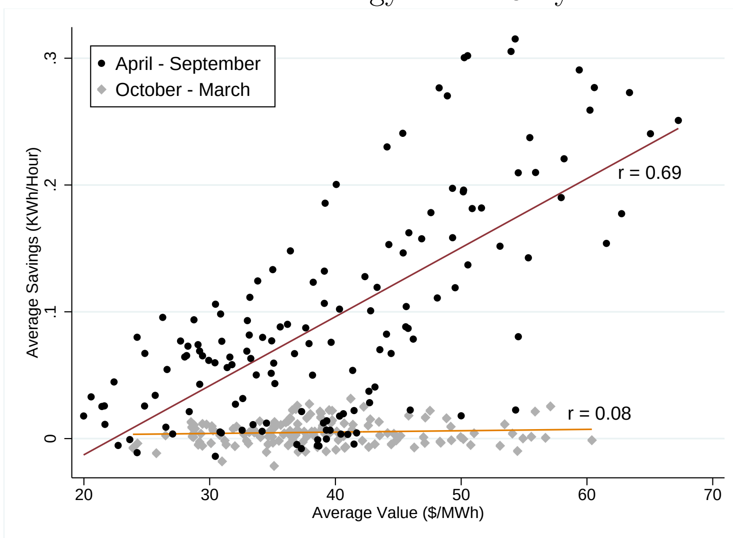

Panel B. Energy and Capacity Prices

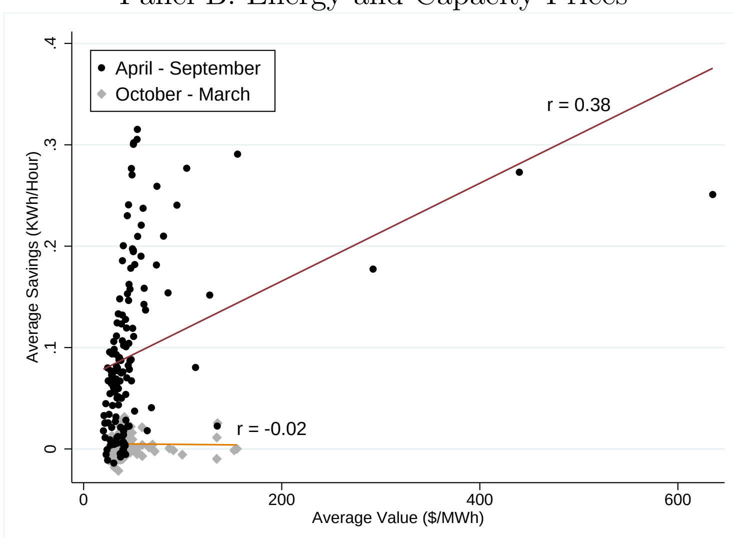

Notes: These scatterplots show the correlation between electricity savings and the value of electricity. Each observation is an hour-of-day by month-of-year pair (e.g. 1–2 p.m. during November). Electricity savings are estimated using a regression which controls for household by hour-of-day by month-of-year and week-of-sample by climate zone fixed effects. Electricity savings are identical in Panels A and B. Panel A uses wholesale electricity prices only, while Panel B also includes hourly capacity values. Energy and capacity price data are from the California electricity market during 2011–2015. See text for details. The figure also includes least squares fitted lines for April-September and October-March observations with the correlation indicated in text above.