

Do Bill Shocks Induce Energy Efficiency Investments?

Corey Lang, Kevin Nakolan, David S. Rapson, and Reid Taylor *

December 13, 2023

Abstract

Electricity bill shocks can draw attention to the benefits of home energy efficiency investments. Our novel identification strategy builds on the fact that prolonged extreme weather events (which raise electricity costs for many customers) fall within a single billing cycle for some customers but are split across cycles for others. We find that households exposed to bill shocks are 22 percent more likely to invest in energy efficiency than households for whom the same weather shock was split across two bills. This result suggests that an opportunity exists for electric utilities to overcome consumer inattention around energy efficiency investments by targeting information to households who have experienced bill shocks.

* (Lang) University of Rhode Island. clang@uri.edu. (Nakolan) Energy Information Administration, Department of Energy. kevin.nakolan@eia.gov. (Rapson) University of California, Davis and Federal Reserve Bank of Dallas. dsrapson@ucdavis.edu. (Taylor) University of California, Davis. rbtaylor@ucdavis.edu. We gratefully acknowledge research funding from the Sloan Foundation and data from United Illuminating Company. All opinions and errors are our own. The views expressed here are those of the authors and do not necessarily reflect those of the Federal Reserve Bank of Dallas or the Federal Reserve System.

1 Introduction

Human beings have limited capacity for attention (Kahneman (2003), Miller (1956)). We actively attend to the most urgent or important decisions in our lives while leaving the rest to heuristics or habits. This has implications for societal welfare. When ignored decisions have external benefits, understanding and remedying the source of inattention can improve social welfare (Allcott et al. (2023)). This paper investigates an opportunity for such an improvement in the context of residential energy efficiency investments.

Energy efficiency continues to be a central element to climate change mitigation plans worldwide (Cabeza et al. (2022)). Yet, the reality of energy efficiency has consistently fallen far short of the aspiration, with a large and growing body of evidence showing that most energy efficiency programs fail to yield energy reduction benefits that meet (or even approach) expectations (Allcott and Greenstone (2017), Fowlie et al. (2018), Burlig et al. (2020)). Economists have sought to understand why energy efficiency programs tend to underperform (Christensen et al. (2023), Boomhower and Davis (2020), Gilbert et al. (2022), Zivin and Novan (2016)) so that we may be able to inform how to better target these programs in the future. While the reasons for energy efficiency underperformance are many, this literature often returns to a common theme: consumers appear to be inattentive to energy efficiency investment opportunities, even when they may be privately net-beneficial.

In this paper we test the hypothesis that electricity customers exposed to exogenous bill shocks will subsequently invest more in home energy efficiency upgrades. Doing so credibly requires overcoming a standard identification challenge: high bills and the choice of durable good attributes are jointly determined by heterogeneous consumer preferences. We develop a novel identification strategy to address this based on severe weather events and idiosyncrasies of electricity billing patterns across customers. Prolonged periods of abnormal heat or cold typically increase electricity use, and while these periods will fall entirely within a single billing month for some customers, the exact same weather event is split across two billing months for other customers by chance. The heterogeneous timing of billing dates relative to exogenous weather events thus forms the basis for testing whether bill shocks induce energy efficiency investments.

Our empirical setting is a utility district in Connecticut from 2008-2017. We observe household-level monthly electric bills, with information on the usage and total billed amount. We focus on people living in single-family homes with prolonged (more than 3 years of) continuous service from their electric utility company, United Illuminating (UI). During this period, UI par-

ticipated in an energy efficiency program that offered UI customers reduced prices on home audits and energy efficiency investments, the records from which we match to the billing data for 300,000 residential electricity customers.

We use the timing of extreme weather events relative to billing month cutoff dates to designate treatment households as those who experienced a single bill encompassing the seasonal peak weather event as an instrument for the endogenous bill shock variable, measured as percent deviation from the average bill over the prior 12 months. Households that had the peak weather event nearly evenly split across two bills serve as control units. We show supporting evidence that the required identifying assumptions for an instrumental variables identification strategy are satisfied in our setting. Results are consistent with our hypothesis. Households exposed to these bill shocks are 22 percent more likely to invest in energy efficiency upgrades in the following months than those that are not. While the relative change is large, the absolute change is small due to the low baseline level of energy efficiency program participation.

This paper makes two main contributions. First, we add to a literature that estimates the effect of bill shocks on consumer decisions. Grubb and Osborne (2015) demonstrate that consumers are inattentive to cellular phone usage, and that bill shock alerts affect subsequent behavior. Jessoe et al. (2014) find that unexpected (and quasi-random) changes in electricity price tariffs appear to shock customers into attention, leading to behavior that is consistent with intermittent attentiveness to electricity prices. Second, our results suggest that a pro-social opportunity exists for electric utilities to target information about energy efficiency programs to customers who have recently experienced bill shocks. The information is readily available and the cost to do so small.

Section 2 reviews the empirical setting and describes the data. Section 3 explains the empirical approach, the results of which are presented in Section 4. Section 5 concludes.


2 Data

We combine data from several sources to construct a panel of customer-level monthly electricity consumption, monthly electricity bill amount, and energy efficiency investments. The primary data source comes from United Illuminating Company (UI), an electric utility company focused on retail transmission and distribution in southwestern Connecticut. The billing dataset contains electricity usage and bill amounts at the customer billing-month level for 302,046 unique customers in the UI service territory for years 2008 to 2017. Due to our focus on durable investments, it is desirable to observe households with long occupancy durations.

For this reason, we restrict the dataset to customers who had continuous service at a particular address with UI for at least three years between 2008 and 2017. We further restrict the sample to exclude observations with implausibly low electricity usage (less than 10kWh total for the month).

We observe the meter read date and billing cycle for each customer-billing month (“billing month”) combination. UI has created 17 distinct billing cycles into which customers are sorted upon enrolling for service. Each customer in a given billing cycle has the same billing period and are billed on the same day. The different billing period end dates for the different cycles are spread relatively evenly through each month. The billing cycle designation is a vestige of the analog era, when meters were read manually on-site by utility employees. Customers in a given billing cycle thus reside in close geographic proximity. Figure 1 shows the distribution of households across billing cycles, weighted by the number of seasons present in our sample.

Figure 1: Density of Households in Main Sample by Billing Cycle



Notes: Households included had continuous service with UI for at least 3 years during the sample period of 2008-2017, were classified as a single-family residence, and are classified as either treatment or control for at least one season.

We match the billing data to a dataset of energy efficiency (EE) investments that was made available through the UI home energy audit and rebate program. These data contain a unique customer identifier, the date of installation, and a category and sub-category for the investment. We observe each line-item investment made by the customer, with a single installation

visit often encompassing multiple investments from different categories.¹ The most frequent line item investment category is “site visit”, which contains audits, tests, and survey items. Most often, these administrative items are associated with other investments, but for a small number of observations, the only investment made is a site visit. Our empirical investigations focus primarily on investments that are most likely to have tangible impacts on energy efficiency, which leads us to drop investment activities that only contain a site visit from our main specifications. We include results from the full sample as a robustness check, with nearly identical results. This allows us to identify which households make investments in energy efficiency and when. We construct our main outcome variable from the item-level data for each household: an indicator variable reflecting whether an investment was made during the months following each season’s peak weather event.

We collect 2016 Connecticut tax assessor data for the UI service territory to identify single-family homes. Households that rent or live in multi-unit structures may lack the ability to make alterations to their dwellings, either due to feasibility or contractual obligation. Even when these households are exposed to high electricity bills, and may wish to invest in more efficient home energy services, they may not have the incentives to do so (Gillingham et al. (2012), Davis (2012)). The assessor data identifies single-family dwellings, and we merge these data with customer billing data by address. The match rate is 43%, and we drop all non-matching customers from our sample. We acknowledge that we are likely discarding some owner-occupied units, but do not have the means for more precise cuts.² Regardless, given that our sample is restricted to customers living in single-family homes that have had prolonged service with UI, we feel confident that this sample is primarily owner-occupied residences, and thus we have identified a sample in which energy investments are most likely. However, to the extent that our sample still contains some renters, our estimates of investment responses to bill shocks may be slightly attenuated. Lastly, the assessor data also include whether a housing unit uses electric heat, which we include in our data as a binary variable.

Lastly, we collect daily temperature data from the National Oceanic and Atmospheric Administration (NOAA) for our entire sample timeframe. We use daily readings from the 10 weather stations located within the UI service territory to calculate average daily temperature as the mean of daily high and low observations, as well as daily heating and cooling degree days during the sample period. Daily heating degree days (HDD) and cooling degree days

¹Table A.1 reports the number of investments made during the sample period by investment category as reported by UI. Table A.2 reports the sub-category for the site-visit category.

²While we do not have estimates of the owner-occupied rate for the whole UI service territory, the owner-occupied rate for New Haven, CT is 62%, according to the Census Bureau’s 2015 American Community Survey (ACS).

(CDD) are defined as (65 – average daily temperature) in Fahrenheit, with positive values representing CDD and negative values HDD (though both measures are recorded as positive values). As described in the methods section, we use these data as a treatment intensity measure to identify periods of anomalous heat or cold events which can lead to large electricity bill increases, the incidence of which will depend on the customer’s monthly billing cycle.

Table 1: Summary Statistics

| | Statistic | Source |
|-----------------------|--------------------|---------------|
| Households | 120,030 | Utility Data |
| Monthly Observations | 11,520,232 | Utility Data |
| HHs that Ever Invest | 23,330 | Utility Data |
| Electric Heating | 0.01 | Assessor Data |
| Years Present | 8.03 (2.38) | Utility Data |
| Monthly KWh | 818.02 (815.50) | Utility Data |
| Monthly Bill Amount | 181.96 (117.74) | Utility Data |
| Investments per Month | 0.26 (0.62) | Utility Data |

Notes: Utility data is sourced from customer-level data provided by United Illuminating for years 2008-2017. Assessor data is collected from municipalities in the territory covered by United Illuminating. Means are reported with standard errors below in parentheses.

Table 1 presents simple summary statistics and data sources that describe our sample. Our data consist of 120,030 unique customers and a total of 11,520,232 customer-billing month observations. The average customer is present for 8 years during our sample period, and the average monthly electricity consumption is 818 kWh, yielding an average bill of \$182. Over our entire sample period, 19.5% of households make energy efficiency investments, but the number of household investments per month is quite low at .26. Only 1% of customers have electric heat (the bulk of customer use natural gas, propane or heating oil). In the following section, we detail our identification strategy and how we use these data to build the panel for our analysis.

3 Methods

When people are generally inattentive about their energy use, as is widely believed in the electricity demand setting, their attention can be drawn by an event such as an abnormal shock to their household’s monthly electricity bill. Our main hypothesis is that investments in energy efficiency are likely to be made in the period following such a shock. In our setting, a household experiences a bill shock when their monthly electricity bill is high and outside of the range of what is normal for them. Our empirical challenge is to assess how a household responds to such a shock in the context of energy efficiency investments. We are interested in the following relationship:

$$Invest_{it} = \beta_1 \Delta Bill_{it} + \alpha_i + \delta_t + \epsilon_{it} \quad (1)$$

where $\Delta Bill_{it}$ is customer i ’s percent change in electricity bill in month t relative to their average bill over the previous twelve months, $Invest_{it}$ is a binary variable equal to 1 if customer i invests in energy efficiency during the six months following month t , α_i are customer fixed effects, which control for unobservable, time-invariant, customer-specific determinants of investment, and δ_t are time fixed effects, which control for macro-level shocks to both bills and investments in a given time period, such as weather or economic conditions. We hypothesize that $\beta_1 > 0$ because anomalously large bills will draw attention to electricity usage and lead households to invest in otherwise profitable energy efficiency measures.

While Equation (1) presents our intuition, estimating it via OLS would likely yield biased estimates due the endogeneity of bill shocks. Electricity bills fluctuate for many reasons: seasonality, adding household members, shifting to remote work, and home renovations, among others. Some of these factors are likely correlated with underlying investment decisions, leading to omitted variable bias.

To address this endogeneity, we implement an instrumental variable (IV) strategy that is based on simple intuition. When a heatwave (or cold snap) occurs, this will increase electricity use and thus the amount on electricity bills. Within a given season, the exact timing of these weather events is random. Due to the staggered nature of billing windows across customers, when these shocks occur relative to a household’s billing cycle will determine the extent to which the “shock” will impact a household’s next bill. If a single bill cycle encompasses the entire heatwave, then that customer’s bill will be anomalously large. However, a customer on a different bill cycle that splits the heat wave evenly between two monthly bills will not receive

the same shock on a single bill, despite having been exposed to identical weather. Since this relationship – when heat waves or cold snaps occur relative to the household’s billing cycle – is as good as random, it forms the basis of an identification strategy that can recover unbiased estimates of the causal effect.


We operationalize this intuition by first defining “winter” and “summer” seasons during which the cold snaps and heat waves can occur. “Winter” is December 1 through March 31 and “Summer” is June 1 through September 30. A “heat wave” of window length “ W ” is defined as the W consecutive days during which time the average daily temperature is higher than during any other group of W consecutive days in that season. “Cold snaps” are calculated analogously for the coldest stretch of W days during a winter season. Our main results use a window length of 20-days, which allows for relative balance between the number of treatment and control households.

We define the outcome variable of interest, $Invest_{it}$, as any energy investment made during the six month window that begins the day after the final billing window containing part of the given season’s heatwave or coldwave has closed. This allows for some lag between the weather event that caused the bill shock and the investment itself. This lag is likely to occur for two main reasons. First, several weeks may pass between the time when the abnormal increase in electricity usage occurs and the moment when the bill is received and paid. Secondly, after a household decides to invest in an energy efficiency upgrade, some weeks or months may pass before the upgrade is installed in their home.

As discussed in Section 2, utility customers are divided into 17 different billing cycles, each with different start and stop dates staggered over the course of a calendar month. We define the instrument, $treat_{it} = 1$, if customer i is on a bill cycle such that the entirety of the heatwave or coldwave is contained within a single billing month. For our main specification, we define control households as having at most 70 percent of the weather shock occurring in a single billing month (i.e. the shock is split relatively evenly across two bills). Further, if a heatwave or coldwave is split across two bills, but the proportion on one bill ranges between 70 and 99%, then that household is excluded from both treatment and control for that season. Figure 2 shows the proportion of seasons each billing wave is considered treatment, control, or omitted during the span of our sample. Importantly, every bill cycle is at some point part of the treatment group and at some point part of the control group (and sometimes omitted). Since identification is coming from all parts of the sample, our estimates are internally valid and this increases the likelihood of being externally valid. Figure 3 shows specifically which bill cycles

fall into which treatment group for each season in our sample.

Figure 2: Treatment Status Across Seasons by Billing Cycle




Notes: This figure presents the proportion of seasons designated as treatment, control, or omitted for each of the 17 billing cycles. Data is provided by United Illuminating for years 2008-2017.

Figure 3 presents a visual representation of the IV setup for the summer 2015 season using a window length of $W = 20$ days. The y-axis represents the 17 different billing cycles, ranging from 1 to 17. The x-axis represents time (in days), with the length of each horizontal line representing the days that are included in a given bill for each billing cycle. For this season, the 20 consecutive hottest days occurred between July 17 to August 5, with this time span indicated by the vertical lines. Billing cycles depicted in dashed lines are the treated group, as their billing dates span the entire heat wave. Billing cycles depicted in dash-dot lines are the control group because the heatwave is relatively evenly between two billing months for these customers. Finally, billing cycles depicted in solid lines are omitted from that season's observations. Above, we stated that $Invest_{it}$ includes investments in a subsequent six month window. In Figure 3, the start of the six month investment period would be August 28 and would end 180 days later, and this would be true for all bill cycles.

We organize our data such that the unit of observation is a customer-season. For each season, a customer can be treated, control, or omitted, but those classifications will change from season

Figure 3: Treatment Status by Billing Cycle During the 2015 Summer Heat Wave



Notes: This figure presents as an example the timing of the various billing cycles during the Summer 2015 20-day peak temperature event, represented by the vertical bars. Cycles 1 and 11-13 are omitted, cycles 2-10 are designated treated, and cycles 14-17 are designated as control.

to season. Empirically, we estimate how treatment impacts bill changes, and then in turn, how do the exogenous changes in bills affect energy efficiency investments in the six months following the shock.

Equations (2) and (3) represent the first and second stages of the IV model, respectively.

$$\Delta Bill_{it} = \gamma_1 treat_{it} + \alpha_i + \delta_t + v_{it} \quad (2)$$

$$Invest_{it} = \beta_1 \Delta \hat{Bill}_{it} + \alpha_i + \delta_t + \epsilon_{it} \quad (3)$$

We additionally include one more variable in our models, which is not displayed in Equations (2) and (3) for simplicity. We include a binary variable, $PastInvest_{it}$, to account for whether the customer has invested in energy efficiency in the preceeding two years. This variable captures the effect that households recently investing in energy upgrades are unlikely to invest again, regardless of subsequent exposure to a bill shock.

The key assumption for identification of causal estimates using the IV estimator is the exclusion restriction, requiring the instrument to only affect energy investments through its impact

on electricity bill amounts. This concern is nullified by the fact that treatment and control customers both experience the same weather event, only differing on how their billing cycles align with the weather event, which is plausibly random after including temporal fixed effects. One potential threat to identification is that bill cycles are not randomly assigned, instead they are based on a property’s location, with entire neighborhoods being on the same cycle. However, our estimates use within household variation through the inclusion of customer fixed effects controlling for unobserved customer differences, and hence unobserved neighborhood differences. Further, as discussed above, every bill cycle is at some point treated and at some point control (and sometimes omitted), so over time there is balance in which neighborhoods are treated. From Figure 2, there are differences in the proportion of seasons spent in different categories across cycles, however all cycles experience all three categories at least 15% of the time. This means that a few bill cycles, which may be different in unobservable ways, are not driving results.

Table 2: Summary Statistics

| | Treatment | Control | Difference |
|------------------------|--------------------|--------------------|--------------------|
| Observations | 515,789 | 528,863 | |
| Delta Bill | 0.17 (0.39) | 0.14 (0.35) | 0.06 (82.51) |
| Qualifying Investments | 0.015 (0.120) | 0.012 (0.110) | 0.003 (11.940) |
| Prior Investments | 0.026 (0.158) | 0.027 (0.162) | -0.003 (-7.886) |
| Monthly KWh | 1,017 (716) | 960 (639) | 58 (71.92) |
| Montly Bill Amount | 226.19 (150.48) | 213.45 (134.63) | 11.98 (68.28) |
| CDD | 225.93 (37.42) | 242.47 (36.45) | -0.00 (-0.00) |
| HDD | 821.32 (86.27) | 823.35 (98.77) | -0.00 (-0.11) |

Notes: Columns 1 and 2 report the mean and standard deviation for treatment and control observations in our main IV sample. Differences shown in column 3 are calculated from a regression of the variable on a binary indicator for treatment. As in our main regression specifications, we include household and season fixed effects. The t-statistic for the coefficient is shown below in parentheses.

Table 2 presents summary statistics on the customer-season sample closely related to our identification strategy. Columns 1 and 2 present means and standard deviations for the treated and control groups, respectively. In Column 3, we report the estimated coefficient from a re-

gression of the respective variable on the treatment indicator, conditioning on household and season fixed effects. We observe a statistically significant difference in the bill shock amount for the treatment group, $\Delta Bill$, as well as a small increase, yet significant increase, in qualifying investments in the post period. At the beginning of Section 4, we further study the reduced form relationship between treatment and investments by estimating an event study model. As expected, the treatment group has higher energy use and monthly bill amounts during the treated seasons, compared to control households. For prior investments, we observe a small negative coefficient on treatment households. While the difference is small, one potential explanation is that control groups have previously made investments and that treatment results in a “catching-up” effect. We take this difference into account in our regressions by conditioning on $PastInvest_{it}$ in our IV regressions.

4 Results

4.1 Reduced Form

We begin the discussion of our results by estimating the reduced form relationship of our instrument, the binary treatment indicator, with our dependent variable of interest, energy efficiency investments. The hypothesis underlying our IV approach is that treatment has a positive relationship with investments, and that treatment operates exclusively through its impact on bill amounts. In Table 2 we presented the fixed effects regression estimate of the reduced form relationship between treatment status and post-period investment rate. We can further test the relationship using a standard event study design. The benefit of using an event study design is twofold. We are able to test the first part of the hypothesis that treatment and post-period investments are positively correlated, as well as characterize the dynamics of the relationship over the duration of the post-period. Secondly, we are able to provide further evidence that treatment is randomly assigned and orthogonal to pre-period investment decisions, characterized by parallel trends in the event study for the months leading up to the peak weather event.

We estimate the following equation:


$$Invest_{it} = \sum_{k=-12}^{12} \beta_k \cdot D_{ikt} + \alpha_i + \delta_t + \epsilon_{it} \quad (4)$$

where $Invest_{it}$ is defined differently than before, and is instead an indicator variable for

whether the household made an energy efficiency investment in month t . D_{ikt} is a series of event time dummy variables for the time, in months, before and after the month during which the weather shock occurred. As before, we include household and time fixed effects to account for unobserved heterogeneity.

We construct a customer-month sample by selecting observations for the 12 months before and after the peak weather event for the same treatment and control households as our main IV seasonal panel. As in the IV sample, we omit observations associated with seasons where a household does not qualify as either treatment or control. In light of the recent advancements in the difference-in-differences literature concerning differential timing of treatment, we omit from the analysis observations for a household if that household was treated in the previous 2 seasons. For example, if a household was treated in Winter of 2014, we would not include observations associated with Summer or Winter 2015, if that household was classified as treatment or control in those seasons. This creates a control group of households who are either never-treated, not-yet treated, or sufficiently distanced from their prior treatment.

Figure 4: Effect of Treatment on Investments



Notes: Estimated coefficients and the 95% confidence intervals from an event study specification for the reduced-form effect of treatment on subsequent household energy efficiency investments are shown.

A plot of the estimated coefficients for the event time dummies is shown in Figure 4, with event time $t = -1$ being the omitted category and representing the timing of the peak weather

event. Treatment is associated with a statistically significant increase in investments for each of the five months that follow the weather event, before dissipating. This supports our choice of a six month post-period investment window in the IV specification. For the nine months immediately prior to the event, we estimate parallel pre-trends between the treatment and control households. While the 9 months preceding exhibit parallel pre-trends, there are negative coefficients estimated 10 and 11 months before the weather event. This could be a result of those households which are most likely to have their attention drawn by the event are those which underinvested in previous seasons.

The results from the event study model are consistent with our narrative that the large bill shock event attracts attention as both treatment and control households experienced the same weather event. In keeping with our IV approach, we believe that the impact of treatment operates exclusively through the shock's impact on the electricity bill.

4.2 Instrumental Variables

In order to establish a baseline, we first estimate the naive OLS specification shown in Equation (1). We report the coefficient estimates in Columns 1 and 2 of Table 3, controlling for prior household energy efficiency investments and season-by-year fixed effects in both Columns 1 and 2, and adding household fixed effects in Column 2. The estimated coefficient on $\Delta Bill_{it}$ is -0.001 in both specifications and statistically significant at the 10% level at least. These coefficients suggest that bill shocks are associated with a decrease in energy efficiency investments, albeit a very small decrease, which is opposite of our hypothesis.

Previously discussed concerns regarding the endogeneity of the main independent variable, $\Delta Bill_{it}$, lead us to believe there may be substantial bias in the results from the OLS estimation. *A priori*, the direction of bias was ambiguous. Given these results, one likely explanation is that anticipated bill shocks, such as those from home renovations, birth of a child, or switching to work from home, do not lead to energy efficiency investments. To address the concerns of omitted variables bias, we turn next to the results from our instrumental variables approach.

We present both first and second stage estimates from our main IV specification in Columns 3 and 4 of Table 3. Column 3 includes season-by-year fixed effects to account for temporal shocks common to all households, such as exceptionally hot/cold seasons, transitory shocks to fuel and energy costs, or changes in investment incentives. Column 4 adds the additional household fixed effect, restricting estimation to variation within a household, over time, to account for inherent differences between households which are correlated with both electricity

Table 3: Effect of Electricity Bill Shocks on Energy Efficiency Investments

| | OLS Estimates | | IV Estimates | |
|--------------------------|---------------------|----------------------|----------------------|----------------------|
| | (1) | (2) | (3) | (4) |
| <i>Estimated Effects</i> | | | | |
| Delta Bill | -0.001* (0.000) | -0.001*** (0.000) | 0.041*** (0.004) | 0.043*** (0.004) |
| Past Investment | 0.021*** (0.001) | -0.135*** (0.001) | 0.022*** (0.001) | -0.134*** (0.001) |
| <i>First Stage</i> | | | | |
| Treatment | | | 0.064*** (0.001) | 0.059*** (0.001) |
| Past Investment | | | -0.026*** (0.002) | -0.026*** (0.002) |
| R ² | 0.003 | 0.156 | 0.229 | 0.334 |
| F-stat | | | 8,165 | 6,556 |
| N | 1,025,572 | 1,025,151 | 1,025,572 | 1,025,151 |
| Season-Year FE | Yes | Yes | Yes | Yes |
| Household FE | No | Yes | No | Yes |

Notes: IV estimates for the effect of electricity bill shocks on subsequent household energy efficiency investments are shown. The percent deviation in the seasonal electricity bill amount from the prior season (delta bill) is instrumented for by treatment status. An indicator variable capturing whether the household made past investments is included. Column 1 controls for season-year fixed effects while column 2 includes both season-year fixed effects and a household fixed effect. The Cragg-Donald F-statistic from the first stage result is reported in the bottom panel. *** = significant at 1 percent level, ** = significant at 5 percent level, * = significant at 10 percent level.

consumption and the decision to invest. Both specifications control for whether the household has made a previous energy efficiency investment the prior two years, which is plausibly correlated with both contemporaneous energy consumption and future investment decisions.

For causal inference to be valid using an instrumental variables approach, the instrument must be sufficiently correlated with the endogenous regressor to satisfy the relevance assumption. We report the Cragg-Donald F-statistics for our first stage estimates, which both exceed 6,500, indicating a very strong statistical relationship between treatment status and $\Delta Bill_{it}$. Interpretation of the first stage coefficients imply that treatment is associated with bill increases that are 5.9 to 6.4 percentage points higher on average than the control group. Given that the average bill in our sample is \$182, this relative bill increase is equivalent to an increase of \$10.71 to \$11.65. Thus, bill cycle timing alone has a causal impact on bill increases.

Turning towards our second stage estimates in the top panel of Table 3, the estimated coefficients on $\Delta Bill_{it}$ are now positive, ranging from 0.042 to 0.045, and are statistically significant at the 1% level. The coefficient changes very little with the inclusion of household fixed effects, which we attribute to the random nature of the shocks and the balance of shocks across bill cycles. We treat Column 4 as our preferred specification. We can interpret the coefficient on $\Delta Bill_{it}$ as the percentage point increase in the probability of making an energy investment resulting from a 100% increase in the customer's electricity bill. In particular, in our preferred specification, a 100% increase in the electricity bill results in a 4.5 percentage point increase in the probability of making a green energy investment. Putting these numbers into perspective, the first stage indicates treatment increases $\Delta Bill_{it}$ 5.9 percentage points on average, which then would yield a 0.27 percentage point increase in investment. The baseline investment rate from Table 2 is 1.2 percent, meaning that treatment increases investment 22.1%. These findings clearly support the idea that heightened attention through bill shocks leads to meaningful increases in energy efficiency investments.

While not the focus of our research, it is worth discussing how past investments in energy efficiency influence future investments. We see the same patterns between Columns 1 and 2 and Columns 3 and 4. When household fixed effects are not included in the model, the coefficient on $PastInvest_{it}$ is positive, but switches to a negative sign when household fixed effects are included. We interpret this pattern as follows. There is a selection process into which type of households invest in energy efficiency, and thus compared to other households, those who have invested in the past are more likely to do so again. However, when household fixed effects are included, only within-household variation is used to estimate coefficients and

that selection process is accounted for. In this case, the coefficient is negative because past investments reduce opportunity or benefit of additional investments.

Table 4 presents results that incorporate the two main variations on our main IV specification: heterogeneity in treatment by CDD/HDD and heterogeneity in the effect of bill shocks by season. Columns 1 and 2 show results for summer seasons and Columns 3 and 4 show results for winter seasons. Columns 1 and 3 use the preferred specification from Table 2, and Columns 2 and 4 add the additional interaction term $treat_{it}*(C|H)DD_t$ to first stage instruments.³ The reported first stage coefficients continue to satisfy the relevance assumption necessary for identification of the IV estimator with the Cragg-Donald F-statistic indicating a very strong relationship between our instruments and the endogenous variable in all specifications. We also see the first-stage relationship between the instruments and $\Delta Bill_{it}$ is in the expected direction, with both the binary treatment variable and the interaction term between treatment and CDD/HDD being associated with higher bill shocks, on average. Further, we find that peak weather events in the summer are associated with higher bill shocks than those from winter for treatment households compared to that season's control group.

Table 4: Effect of electricity Bill Shocks on Energy Efficiency Investments

| | Summer | | Winter | |
|--------------------------|---------------------|---------------------|---------------------|---------------------|
| <i>Estimated Effects</i> | | | | |
| Delta Bill | -0.010** (0.004) | -0.002 (0.003) | 0.194*** (0.013) | 0.167*** (0.011) |
| <i>First Stage</i> | | | | |
| Treatment | 0.092*** (0.001) | 0.100*** (0.001) | 0.032*** (0.001) | 0.032*** (0.001) |
| Treatment x CDD or HDD | | 0.138*** (0.002) | | 0.027*** (0.001) |
| R ² | 0.634 | 0.638 | 0.566 | 0.567 |
| F-stat | 10,913 | 7,286 | 2,000 | 1,275 |
| N | 509,741 | 509,741 | 500,236 | 500,236 |
| Sargan Statistic | | 18.48 | | 27.13 |

Notes: IV estimates for the effect of electricity bill shocks on future household energy efficiency investments broken out by season are shown. The percent deviation in the seasonal electricity bill amount from the prior season (delta bill) is instrumented for by treatment status in columns 1 and 3, and by treatment status and its interaction with Cooling Degree Days or Heating Degree Days in columns 2 and 4. CDD and HDD are demeaned. An indicator variable capturing whether the household made past investments is included as a control variable. All models include season by year fixed effects and a household fixed effect. *** = significant at 1 percent level, ** = significant at 5 percent level, * = significant at 10 percent level.

Examining the second stage coefficients, we see that the energy investment response is

³CDD and HDD are demeaned so that the coefficient on the interaction term with *treat* is the average effect of treatment at the average CDD or HDD level.

entirely concentrated with those customers that received a bill shock from a winter peak event. The coefficients on $\Delta Bill_{it}$ for winter are 0.194 and 0.167 and are highly statistically significant. In contrast, the coefficients for summer are actually negative, but very small in magnitude. Focusing on Column 4, the coefficient on $\Delta Bill_{it}$ implies that a 100% increase in a customer's electricity bill during the winter season is associated with a 17.6 percentage point increase in the probability of investing.

We hypothesize two reasons why we could expect to see different responses across seasons in our setting. First, it is possible that consumers are more sensitive to shocks during winter months and more likely to react to shocks during this time period. Given that our data is from households in Connecticut, the more notable harsh winter climate may make these shocks more salient, leading to this phenomenon of a difference in investment rates by season. If customers are more likely to assume that peak winter events will recur, while peak summer events may in fact just be anomalies, then we would observe the results seen in our estimates. Secondly, we capture investments made during the six-month window following the peak temperature event, which for winter, typically occurs in late February and early March. This leads to an investment window running from March through September. For the summer window, the peak events all occur in August, except for 2009, yielding investment windows spanning from August to February. Households are traditionally much more likely to make house renovations and improvements during the summer months. Appendix Figure XXXX shows the highly cyclical nature of total private residential construction spending over time. Summer months, have on average 30% higher spending compared to winter months. With construction at its highest when the peak weather event occurs in summer, when contractors are at their highest demand, it is reasonable to think that consumers are unable to immediately react to the bill shock, and to the extent that their attention to the bill shock declines over time, any inertia created dissipates before an energy investment is able to be made.


4.3 Robustness

In addition to testing the robustness of our results to the inclusion of various controls, we next explore the robustness of our key findings to the major sample selection criteria made: the peak weather event 20-day window length, the 6-month investment period length, and the exclusion of administrative category only investments.

We first look at the robustness of the main results to the choice of window length. Too short of a window length means our designation of treatment is unlikely to lead to a meaningful

shock to electricity bills in addition to limiting the number of good control billing waves by construction. Too long of a window length can smooth over peak weather events that cause reasonably large shocks to electricity bills, as well as limits the number of billing waves that can be designated treatment (since by definition a treatment billing wave must encompass all the window on one bill). In Figure 5, we vary the window length from 10 to 25 days along the x-axis and report the coefficient on $\Delta Bill_{it}$ from the second stage estimation of our preferred IV specification along with the coefficient's 95% confidence interval. Excluding the extreme ends of the window length distribution, our estimates are robust to changes in the window length. Specifically, we see near identical results for windows of length 15-20 days. In Appendix Figure A.3, we report estimates for the specifications that allow for heterogeneous responses by winter and summer seasons, and results suggest qualitatively similar conclusions as those seen in Table 4 across the spectrum of window length.

Figure 5: Coefficient Plot by Window Length



Notes: This figure plots the estimated second stage coefficient from the IV regression of investments on delta bill and the 95% confidence interval, varying the peak weather event window definition used to designate treatment and control groups from 10 to 25 days.

Next, in Table 5, we test the robustness of our results to the other two main sample selection criteria, investment period and investment type. Our main results from Column 4 of Table 3 are replicated in Column 1 of Table 5 for ease of reference. First, we add back in investments which were classified as administrative only to our main specification in Column 2. These

results are nearly identical. As such, in the next two columns we revert to excluding administrative only investments. Next, we examine how results change when the investment period length, which is used in the construction of our outcome variable, is changed to 3 months or 9 months, which appear in Columns 3 and 4, respectively.⁴ For an investment period of three months, the estimated coefficient is 0.028, and for an investment period of nine months, the estimated coefficient is 0.053. These results combined with the estimated coefficient of 0.045 for an investment period of six months lead to three conclusions. First, we see that as the window length increases, more investments are made in total, consistent with results from our earlier event study results. Secondly, the main results cannot be explained by the treatment group simply making investments sooner than the control group, with control group catching up as treatment group demand is satisfied. Thirdly, we see diminishing additional effects the longer we extend the window. The diminishing incremental change in the effect makes sense because as time goes on the attention focused on electricity dissipates.

Table 5: Effect of Electricity Bill Shocks on Energy Efficiency Investments

| | (1) | (2) | (3) | (4) |
|-------------------------------|----------------------------------|----------------------------------|----------------------------------|----------------------------------|
| <i>Second Stage Estimates</i> | | | | |
| Delta Bill | 0.043 ^{***} (0.004) | 0.043 ^{***} (0.004) | 0.028 ^{***} (0.003) | 0.048 ^{***} (0.005) |
| Past Investment | -0.134 ^{***} (0.001) | -0.131 ^{***} (0.001) | -0.061 ^{***} (0.001) | -0.164 ^{***} (0.001) |
| N | 1,025,151 | 1,025,151 | 1,025,151 | 1,025,151 |
| Investment Period | 6 Mo. | 6 Mo. | 3 Mo. | 9 Mo. |
| Administrative Investments | No | Yes | No | No |

Notes: IV estimates for the effect of electricity bill shocks on future household energy efficiency investments broken out by season are shown. The percent deviation in the seasonal electricity bill amount from the prior season (delta bill) is instrumented for by treatment status in all specifications. Column 1 replicates our main results as shown in column two of table 3. Column 2 reports estimates including administrative only investments in the dependent variable, which are excluded from our main results. Results in Columns 3 and 4 report results from changing the definition for the investment period to 3 months and 9 months, respectively, from the baseline level of 6 months. The dependent variable in Columns 3 and 4 do not include administrative only investments. All models include season by year fixed effects and a household fixed effect and control for past investments made by the household. *** = significant at 1 percent level, ** = significant at 5 percent level, * = significant at 10 percent level.

⁴We cap the investment length at 9 months in order to not contaminate the outcome variable with effects from the following year's peak weather event of the same season.

5 Conclusion

Behavioral obstacles appear to be a major factor inhibiting investment in residential energy efficiency upgrades. The limits of human capacity for attention point to the potential benefits of capitalizing on circumstances that draw peoples' focus to the costs of habits and inaction. In this paper, document one such circumstance.

Energy customers in our setting exhibit a willingness to invest in energy efficiency after being exposed to a weather-induced bill shock. This period begins a shortly after the bill-shocked household receives their utility bill and lasts for 4 months, on average. During this window of time, an opportunity may exist for targeted outreach that encourages these households to consider energy efficiency investments and informs them of the potential benefits.

References

- Allcott, Hunt and Michael Greenstone**, "Measuring the Welfare Effects of Residential Energy Efficiency Programs," *SSRN Electronic Journal*, may 2017.
- , **Daniel Cohen, William Morrison, and Dmitry Taubinsky**, "When Do "Nudges" Increase Welfare?," 2023.
- Boomhower, Judson and Lucas Davis**, "Do energy efficiency investments deliver at the right time?," *American Economic Journal: Applied Economics*, jan 2020, 12 (1), 115–139.
- Burlig, Fiona, Christopher Knittel, David Rapson, Mar Reguant, and Catherine Wolfram**, "Machine learning from schools about energy efficiency," *Journal of the Association of Environmental and Resource Economists*, nov 2020, 7 (6), 1181–1217.
- Cabeza, L F, Q Bai, P Bertoldi, J M Kihila, A F P Lucena, É Mata, S Mirasgedis, A Novikova, and Y Saheb**, "2022: Buildings. In IPCC, 2022: Climate Change 2022: Mitigation of Climate Change. Contribution of Working Group III to the Sixth Assessment Report of the Intergovernmental Panel on Climate Change," Technical Report, Cambridge University Press, Cambridge, UK and New York, NY, USA 2022.
- Christensen, Peter, Paul Francisco, Erica Myers, and Mateus Souza**, "Decomposing the wedge between projected and realized returns in energy efficiency programs," *Review of Economics and Statistics*, jul 2023, 105 (4), 798–817.
- Davis, Lucas W**, "Evaluating the Slow Adoption of Energy Efficient Investments: Are Renters Less Likely to Have Energy Efficient Appliances?," in Don Fullerton and Catherine Wolfram, eds., *The Design and Implementation of U.S. Climate Policy*, University of Chicago Press, 2012, chapter 19, pp. 301–316.
- Fowlie, Meredith, Michael Greenstone, and Catherine Wolfram**, "Do energy efficiency investments deliver? Evidence from the Weatherization Assistance Program," *Quarterly Journal of Economics*, aug 2018, 133 (3), 1597–1644.
- Gilbert, Ben, Jacob LaRiviere, and Kevin Novan**, "Uncertainty and additionality in energy efficiency programs," *Journal of Environmental Economics and Management*, 2022, 115.

Gillingham, Kenneth, Matthew Harding, and David Rapson, "Split incentives in residential energy consumption," *Energy Journal*, 2012, 33 (2), 37–62.

Grubb, Michael D and Matthew Osborne, "Cellular service demand: Biased beliefs, learning, and bill shock," *American Economic Review*, jan 2015, 105 (1), 234–271.

Jessoe, Katrina, David Rapson, and Jeremy B Smith, "Towards understanding the role of price in residential electricity choices: Evidence from a natural experiment," *Journal of Economic Behavior and Organization*, 2014, 107 (PA), 191–208.

Kahneman, Daniel, "A Perspective on Judgment and Choice: Mapping Bounded Rationality," *American Psychologist*, sep 2003, 58 (9), 697–720.


Miller, George A., "The magical number seven, plus or minus two: some limits on our capacity for processing information," *Psychological Review*, mar 1956, 63 (2), 81–97.

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

A Appendix


A.1 Figures

Figure A.1: Treatment Status by Billing Cycle and Season

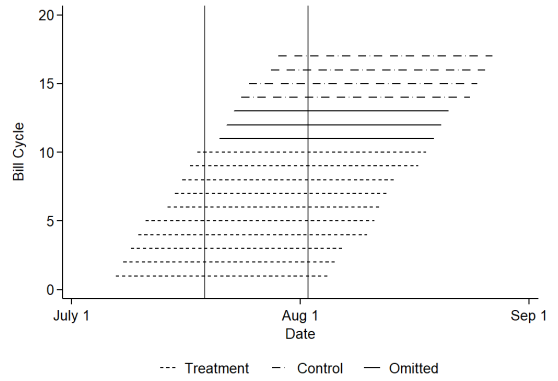


Notes: This figure shows the treatment status for each of the 17 billing cycles across the seasons of our sample period.


Figure A.2: Billing Cycles During 2015 Summer Heat Wave by Treatment Status




(a) 10 Day



(b) 15 Day




(c) 20 Day




(d) 25 Day

Notes: This figure shows treatment starts across the 17 different billing cycles for the 2015 summer peak weather event. Each panel shows a different window length used to designate the hottest consecutive days for the peak event.

Figure A.3: Coefficient Plot by Window Length



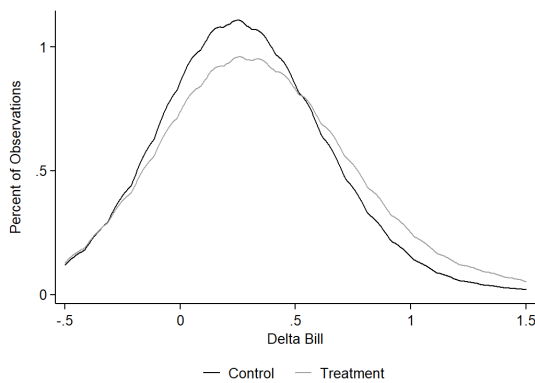
(a) Summer



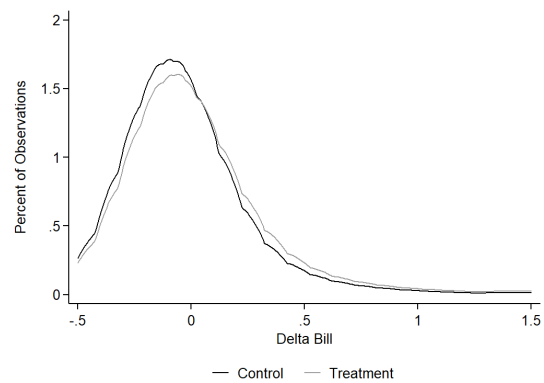
(b) Winter

Notes: This figure reports the estimated second stage coefficient on delta bill using various window lengths to designate the peak weather event. Results for summer and winter seasons are estimated and reported separately.

Figure A.4: K-Density Plot for Delta Bill




(a) Summer



(b) Winter

Notes: This figure reports the distribution of the delta bill variable for treatment and control observations separately by winter and summer seasons using a k-density plot.

Figure A.5: Treatment Status by Season for Event Study



Notes: This figure shows the count of households in the event study sample by treatment status for each season used in the event study analysis.

A.2 Tables

Table A.1: Energy Efficiency Investments by Categories

| | Freq. | Percent | Cum. |
|-------------------------------------|---------|---------|--------|
| Site visits: audits and inspections | 288,144 | 41.99 | 41.99 |
| HVAC | 40,002 | 5.83 | 47.82 |
| Custom Measures | 20,511 | 2.99 | 50.81 |
| Hot Water | 121,293 | 17.68 | 68.49 |
| Envelope | 93,117 | 13.57 | 82.06 |
| Incentive Bonus | 1,172 | 0.17 | 82.23 |
| Lights | 120,847 | 17.61 | 99.84 |
| Refrigeration | 1,066 | 0.16 | 100.00 |
| Total | 686,152 | 100.00 | |

Notes: The table reports investments made by customers through United Illuminating.

Table A.2: Site Visits: Detailed Subcategories

| | Freq. | Percent | Cum. |
|-------------------------------|---------|---------|--------|
| ADJUSTMENT, OIL, ARRA | 128 | 0.04 | 0.04 |
| ADMINISTRATIVE ADJUSTMENT | 378 | 0.13 | 0.18 |
| APPLIANCE EVALUATION | 24,172 | 8.39 | 8.56 |
| DATA ENTRY FEE, TEMPORARY | 3,043 | 1.06 | 9.62 |
| HEALTH AND SAFETY | 1,816 | 0.63 | 10.25 |
| HES SITE VISIT | 30,413 | 10.55 | 20.81 |
| HESCORE W/ CORE SERVICES | 9,035 | 3.14 | 23.94 |
| HOME AUDIT | 142,467 | 49.44 | 73.38 |
| HVAC TESTS | 36,342 | 12.61 | 86.00 |
| INSULATION VERIFICATION VISIT | 260 | 0.09 | 86.09 |
| KILL-A-WATT METER | 7,143 | 2.48 | 88.57 |
| SITE VISIT | 32,947 | 11.43 | 100.00 |
| Total | 288,144 | 100.00 | |

Notes: The table reports the subcategories for investments made in the Site visits: audits and inspections category.