



Federal Reserve
Bank of Dallas

Do Bill Shocks Induce Energy Efficiency Investments?

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

Working Paper 2405

Research Department

<https://doi.org/10.24149/wp2405>

September 2024

Working papers from the Federal Reserve Bank of Dallas are preliminary drafts circulated for professional comment. The views in this paper are those of the authors and do not necessarily reflect the views of the Federal Reserve Bank of Dallas or the Federal Reserve System. Any errors or omissions are the responsibility of the authors.

Do Bill Shocks Induce Energy Efficiency Investments?*

Corey Lang[†], Kevin Nakolan[‡], David S. Rapson[§] and Reid Taylor[‡]

July 16, 2024

Abstract

Inattention can lead to suboptimal investment in energy efficiency. We study whether electricity bill shocks 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 average sized bill shocks are 22 percent more likely to invest in energy efficiency than households with normal bills. This result suggests that inattention is indeed a factor in residential energy decisions and utilities may be able to leverage bill shocks to promote efficiency investments.

JEL Classification: Q40, Q50, D12

Keywords: inattention, bill shock, energy efficiency investments

*We gratefully acknowledge research funding from the Alfred P. Sloan Foundation (Grant #G-2017-10005) 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.

[†]Corey Lang, University of Rhode Island, clang@uri.edu.

[‡]Kevin Nakolan, Energy Information Administration, Department of Energy, kevin.nakolan@eia.gov.

[§]David S. Rapson, University of California, Davis and Federal Reserve Bank of Dallas, dsrapson@ucdavis.edu.

[‡]Reid Taylor, Federal Reserve Bank of Dallas, ReidBTaylor@gmail.com.

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 typically leaving the rest to heuristics or habits. This has implications for societal welfare. When ignored decisions have external benefits or costs, understanding and remedying the source of inattention can improve social welfare (Allcott et al. (2023)). This paper seeks evidence of inattention, and an opportunity to overcome it, 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)). However, 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); Chuang et al. (2022)). 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 residential electricity customers exposed to exogenous and large bill shocks subsequently invest more in home energy efficiency upgrades. The intuition is that customers are inattentive to bills as long as they are within a normal range, and independently or as a result, they are also inattentive to advantageous energy efficiency investments. However, an abnormally large bill may “shock” a customer into shifting attention towards their electricity consumption and energy efficiency investment opportunities.¹

¹A similar channel has been documented in the rooftop solar industry, in which short-run weather fluctuations influence the homeowner decision to invest in solar panels (Liao (2020) and Lamp (2023)). Note,

Our empirical setting is a utility district in Connecticut, United Illuminating (UI), which serves about 300,000 residential electricity customers. We observe household-level monthly electricity bills with information on the usage and total billed amount for all customers for years 2008-2017. During this period, like many utilities, UI had active energy efficiency programs that offered UI customers reduced prices on home audits and energy efficiency investments. Data on investments, including audits, are recorded at the household-day level and we match these records to the billing data. We collect daily temperature data in the service territory to be used as part of our empirical strategy. Our sample focuses on customers living in single-family homes with prolonged (more than 3 years of) continuous service from UI. Our final sample comprises of 120,000 customers with 11.5 million monthly observations. At some point during our sample, 19% of households make some energy efficiency investment through our observed investment channel.

We develop a novel identification strategy to estimate the causal effect of bill shocks on energy efficiency investments. The endogeneity concern in our setting is that electricity consumption and the choice of durable good attributes are jointly determined by heterogeneous consumer preferences. To address this, we implement an instrumental variables model based on a combination of temperature extremes (heat waves and cold snaps) and idiosyncrasies of billing patterns across customers. UI's customers are split across 17 billing cycles that have associated recurring billing months that start and end at different times for different groups, spread relatively evenly throughout a given month. For example, the February bill for billing cycle 1 may go from January 2 to February 1 while billing cycle 2 may go from January 5 to February 4, and so on. 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. We define treatment households as those who experienced a single bill encompassing the seasonal peak temperature event, whereas

however, that the bill shock channel that we study is related to, but distinct from, weather.

households that experienced the same peak temperature event split nearly evenly across two bills are control units. Treatment is our instrument for the endogenous bill shock variable, which we define as percent deviation from the average bill over the prior 12 months. The instrument is powerful, with F-statistics over 1,000 in all specifications, and we show supporting evidence that the required identifying assumptions for an instrumental variables strategy are satisfied in our setting. Through our identification strategy, we are able to isolate the mechanism to the size of the bill itself, as opposed to changes in total outlay for energy costs, as treatment and control households experience the same underlying demand shock.

Results are consistent with our hypothesis. Households exposed to an average bill shock are 22 percent more likely to invest in energy efficiency upgrades in the following six months than those that are not, despite both groups being exposed to the same conditions. While the relative change is large, the absolute change is small due to the low baseline levels of energy efficiency program participation among the population in a given period. This result is robust to varying definitions of the peak weather event used to assign treatment, alternative definitions of what qualifies as an investment, and alternative post-event investment window lengths used to construct the main outcome variable. However, two features of our results suggest that the seasonality of energy investments in our setting were surprising. Event study estimates reveal a decline in treated household energy investments roughly a year before the bill shock. Moreover, in an extension of our main findings we estimate heterogeneous effects by season and find the effect is concentrated in investment responses to bill shocks as a result of a cold snap. We propose that this could be due to both the increased salience of bill shocks that occur in the winter in the region of study as well as home renovation seasonality (which is more active in the warmer spring and summer months that follow cold snap season relative to fall and winter that follow heat wave season). Future research is needed to help clarify these potential mechanisms.

This paper makes three main contributions to the literature. First, our results contribute to the literature on customer inattention and price salience by showing that

large bills can draw the attention of consumers and lead to future investment expenditures. Intermittent billing has become a feature of many products in recent years, however bills which vary from period to period as a function of usage are less common. Prior literature that has studied the effect of “bill shocks” on customer behavior in the cell phone industry (Grubb and Osborne (2015)) and health care (Hoagland et al. (2023)). The related literature on residential electricity, however, is mixed with respect to whether customers are inattentive to their usage and prices. Sallee (2014) presents a model in which rational inattention arises from costly information acquisition.² In our setting, customers are attentive to energy usage and prices, even if intermittently. This result can be unified with rational inattention to the extent bill shocks overcome the costs of acquiring information about usage and expenditures. Related to our work, Gilbert and Graff Zivin (2014) shows that the arrival of electricity bills increases the salience of prices in the short-run, with homeowner’s reducing usage in the weeks that follow. Sexton (2015) finds an increase in usage when accounts move to automatic intermittent billing. Jessoe et al. (2014) find that unexpected (and quasi-random) changes in electricity price tariffs appear to affect consumer behavior in a manner that is consistent with intermittent attentiveness to electricity prices. Our work also relates to the literature on price salience in demand for durable goods; Myers (2019) and Houde and Myers (2021) show that customers are attentive to local electricity prices when making high-value durable household investment decisions.

Second, our results suggest that an opportunity exists for electric utilities to target information about energy efficiency programs to customers who have recently experienced bill shocks. A large literature has documented the effect feedback and nudging can have on subsequent energy consumption (Faruqui et al. (2010); Buckley (2020); Allcott and Kessler (2019); Houde et al. (2013); Jessoe and Rapson (2014); Gilbert and Graff Zivin (2014)). The information is readily available and the cost for firms to notify customers is small. Demand reduction and demand response are key features of

²This is supported by a body of empirical work in the energy setting (Houde (2018); Allcott (2013); Allcott and Rogers (2014); Davis and Metcalf (2016); Allcott and Taubinsky (2015); Allcott and Knittel (2019)).

reaching the energy efficiency goals of decarbonization policy. Our work contributes to the understanding of how existing business practices can play a role in reaching these goals.

Lastly, we make a methodological contribution to the literature through our novel identification strategy based on weather and billing cycles. Our research design could relatively easily be applied elsewhere to estimate causal effects of bill shocks. We have identified a source of exogenous variation that is common to settings where billing is intermittent, but costs are a function of past demand and not fixed. These types of billing structures are common to utility companies, residential natural gas, heating oil, and water. Similar billing structures are present in cell phones, internet service, and health setting and could benefit from our methodology using related first stage variation from sources other than weather.


The paper proceeds as follows. 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 data set contains electricity usage and bill amounts at the customer billing-month level for 302,046 unique customers in the UI service territory for the years 2008 to 2017. Due to our focus on durable investments, it is desirable to observe households with a long occupancy duration. We restrict the data set 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. All customers in a given billing cycle have 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 data set 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.³ 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 tangible 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 an impact 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. We use the resulting data 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, a relationship referred to as the principal-agent problem (Gillingham et al. (2012), Davis (2012)). We merge to our billing and investment data an identifier for single-family dwelling from the assessor data. 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.⁴ 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

³Appendix Table A.1 reports the number of investments made during the sample period by investment category as reported by UI. Appendix Table A.2 reports the sub-category for the site-visit category.

⁴While 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).

our sample still contains some renters, our estimates of investment responses to bill shocks may be slightly attenuated. Lastly, we include in our data a binary indicator variable whether the assessor data indicates a housing unit uses electric heat.

Lastly, we collect daily temperature data from the National Oceanic and Atmospheric Administration (NOAA) for our entire sample time frame. 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 (CDD) are defined as $(65 - \text{average daily temperature})$ 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 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 monthly bill of \$182. Over our entire sample period, 19.4% of households make at least one energy efficiency investment that we observe in our investment data set. On average, there are .26 energy efficiency investments per household in the sample. Lastly, only 1% of customers have electric heat (the bulk of customers 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

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
Investments per HH	0.26 (0.62)	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

Notes: Utility data is sourced from customer-level data provided by United Illuminating for the 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.

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 causally 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 heat wave (or cold snap) occurs, this will increase electricity use and thus the amount paid 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 heat wave, 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 and the identical increase in energy consumption. 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.

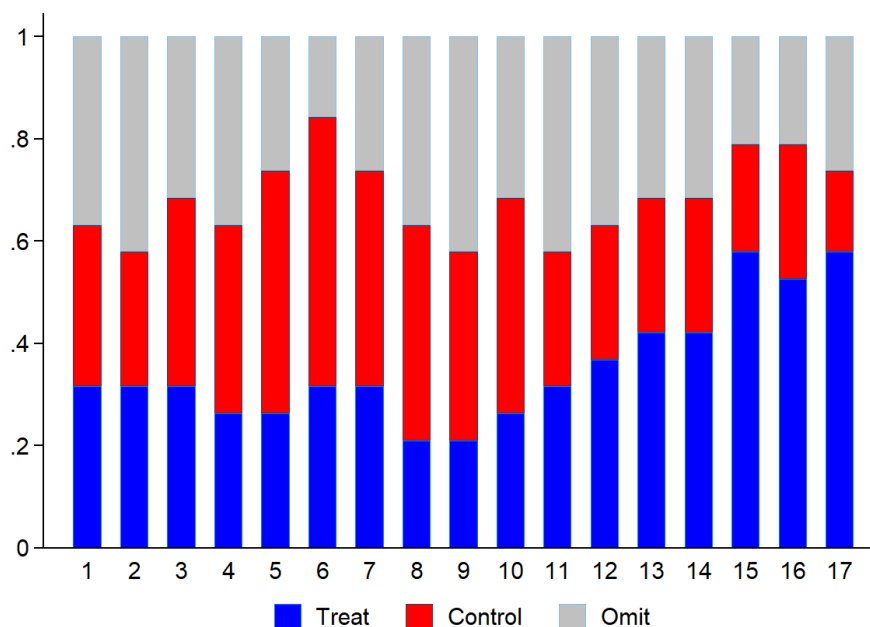
We operationalize this intuition by first defining "winter" and "summer" seasons during which the cold snaps and heat waves can occur. "Winter" is December 1 through March 31 and "Summer" is June 1 through September 30. A "heat wave" of window length " W " is defined as the W consecutive days during 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, defined below.

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 heat wave or cold snap 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. Appendix Figure A.1 reports the dates for the 20-day peak weather event and the start date for the investment window for each season.

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 heat wave or cold snap 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 heat wave or cold snap 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 cycle 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. Appendix Figure A.2 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 the 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 18 and August 6, with this time span indicated by the solid vertical lines. Billing cycles depicted in dashed lines are the treated group as their respective billing dates contain the entire heat wave on one bill. Billing cycles depicted in dash-dot lines are the control group because the heat wave is split 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 27 and would end 180 days later for all bill cycles.

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.

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 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 a binary variable, *PastInvest_{it}*, which is not displayed in equations (2) and (3) for simplicity, to account for whether the customer has invested in energy efficiency in the preceding 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 results are not driven by a few bill cycles which may be different in unobservable ways.

Table 2: Summary Statistics

	Treatment	Control	Regression Adjusted 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)
Monthly Bill Amount	226.19 (150.48)	213.45 (134.63)	11.98 (68.28)
CDD	225.93 (37.42)	242.47 (36.45)	-18.61 (-155.76)
HDD	821.32 (86.27)	823.35 (98.77)	-3.88 (-13.42)

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 and include household and season fixed effects as in our main regression specifications except CDD and HDD which only include household 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 regression 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.

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 four 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. However, there are negative coefficients estimated 10 and 11 months before the weather event. Given the quasi-randomness conferred by the identification strategy, a pre-bill shock treatment effect was unexpected. Extending the causal interpretation, the pre-event effects may

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 household energy efficiency investments are shown.

reveal an aspect of selection – those households which are most likely to have their attention drawn by the event are those which underinvested in previous seasons.

Overall, the results from the event study model are consistent with the narrative that the large bill shock events attract 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 signifi-

cant 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.

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: Estimates for the effect of electricity bill shocks on subsequent household energy efficiency investments are shown. Columns 1 and 2 present OLS estimates of the effect of bill shocks on investments. In Columns 3 and 4, the percent deviation in the electricity bill amount the month of the weather shock from the prior year's average bill (delta bill) is instrumented for by treatment status. An indicator variable capturing whether the household made past investments is included. Columns 1 and 3 control for season-year fixed effects while Columns 2 and 4 include 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.

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 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 implies 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 heat waves and Columns 3 and 4 show results for cold snaps. 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 the set of first stage instruments.⁵ The reported first stage coefficients continue to satisfy the relevance assumption necessary for identification of the

⁵CDD 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.

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 heat waves are associated with higher bill shocks than those from cold snaps for treatment households compared to that season's control group.

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

	Heat Wave		Cold Snap	
<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 electricity bill amount the month of the weather shock from the prior year's average bill (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-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 entirely concentrated with those customers that received a bill shock from a cold snap in the winter. The coefficients on $\Delta Bill_{it}$ for cold snaps are 0.194 and 0.167 and are highly statistically significant. In contrast, the coefficients for heat waves 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 from a cold snap

is associated with a 16.7 percentage point increase in the probability of investing.

These contrasting results are somewhat puzzling for two reasons. First, we estimate larger bill shocks occurring with heat waves rather than cold snaps, so intuitively we might expect larger investment responses from heat waves. Second, the assessor data imply that only 1% of our sample uses electric heat as their primary heating source. We hypothesize several reasons why we could expect to see different responses across seasons even given these puzzles. First, the heightened cold snap response may be a function of our setting. Connecticut has a harsher winter climate that may make these shocks more salient. Second and relatedly, bills may simply be more likely to be read or scrutinized during those harsh winter months when daylight is short and fewer vacations are taken. Third, we capture investments made during the six-month window following the peak temperature event, which for cold snaps typically occurs in February as shown in Table A.1. This leads to an investment window running from March through September. In contrast, for heat waves, the peak events all occur in late July and August, yielding investment windows spanning from August to February. Households are traditionally more likely to make house renovations and improvements during the summer months. Private residential construction spending is highly cyclical in nature with summer months having 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 can be made. Lastly, in terms of heat source, it is possible that households are using space heaters to selectively heat rooms instead of relying on expensive delivery of heating oil or propane, or it is possible that the assessor data is inaccurate or out-of-date.


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.4, we report estimates for the specifications that allow for heterogeneous responses for heat waves and cold snaps, and results suggest qualitatively similar conclusions as those seen in Table 4 across the spectrum of window length.

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

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.

Columns 3 and 4, respectively.⁶ 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 are shown. The percent deviation in the electricity bill amount the month of the weather shock from the prior year's average bill (delta bill) is instrumented for by treatment status in all specifications. Column 1 replicates our main results as shown in Column 2 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 change the definition for the post-shock investment period to 3 months and 9 months, respectively, from the baseline level of 6 months. The dependent variable in Columns 3 and 4 follow our main results and do not include administrative-only investments. All models include season-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.

⁶We 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. This paper makes use of exogenous variation in the impacts of extreme temperature events to document one such circumstance, bill shocks. This work adds to the growing literature which empirically documents price salience and customer attention in residential household energy settings.

In this paper, we make use of the random timing of extreme high and low temperature events during the year with respect to the timing of electricity billing periods across customers to identify households that received anomalously large electricity bills. Given our design, we can isolate the effect to the increase in the amount on the bill, as opposed to an effect of the increase in household energy use, as both treatment and control households experienced the same energy-use shock from the weather event, only differing in the amount of the high-use period contained in a single billing month. This design, coupled with our rich data on household level energy efficiency investments, allows us to present novel causal estimates of impact of billing shocks on subsequent investment activity in the months that follow.

Customers in our setting exhibit a willingness to invest in home energy efficiency upgrades in the months after being exposed to the weather-induced bill shock. Households exposed to the average bill shock amount are 22 percent more likely to invest than households that did not receive a bill shock. This effect is largely concentrated to the peak cold temperature events that occur in the winter, likely a function of both the study's geographic location in Connecticut and the cyclical nature of the home construction industry. During this window of time, an opportunity may exist for targeted outreach that makes use of the increased attention to encourage these households to consider energy efficiency investments and inform them of the potential benefits.

This paper offers a strong, novel identification strategy applied to an important

research question. We acknowledge that a couple elements of our findings are not intuitive, though plausible. Future work is needed applying these methods to other utility data in different geographies to assess if bill shock responses are dependent on location, climate, etc., and investigate mechanisms.

References

- Allcott, Hunt**, “The welfare effects of misperceived product costs: Data and calibrations from the automobile market,” *American Economic Journal: Economic Policy*, 2013, 5 (3), 30–66.
- **and Christopher Knittel**, “Are consumers poorly informed about fuel economy? Evidence from two experiments,” *American Economic Journal: Economic Policy*, feb 2019, 11 (1), 1–37.
- **and Dmitry Taubinsky**, “Evaluating behaviorally motivated policy: Experimental evidence from the lightbulb market,” *American Economic Review*, aug 2015, 105 (8), 2501–2538.
- **and Judd B. Kessler**, “The welfare effects of nudges: A case study of energy use social comparisons,” *American Economic Journal: Applied Economics*, 2019, 11 (1), 236–276.
- **and Michael Greenstone**, “Measuring the Welfare Effects of Residential Energy Efficiency Programs,” *SSRN Electronic Journal*, may 2017.
- **and Todd Rogers**, “The short-run and long-run effects of behavioral interventions: Experimental evidence from energy conservation,” oct 2014.
- **, 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.
- Buckley, Penelope**, “Prices, information and nudges for residential electricity conservation: A meta-analysis,” jun 2020.
- 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.

Chuang, Yating, Magali A. Delmas, and Stephanie Pincetl, "Are Residential Energy Efficiency Upgrades Effective? An Empirical Analysis in Southern California," *Journal of the Association of Environmental and Resource Economists*, jul 2022, 9 (4), 641–679.

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.

— **and Gilbert E Metcalf**, "Does better information lead to better choices? Evidence from energy-efficiency labels," *Journal of the Association of Environmental and Resource Economists*, 2016, 3 (3), 589–625.

Faruqui, Ahmad, Sanem Sergici, and Ahmed Sharif, "The impact of informational feedback on energy consumption-A survey of the experimental evidence," *Energy*, 2010, 35 (4), 1598–1608.

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 and Joshua Graff Zivin, "Dynamic salience with intermittent billing: Evidence from smart electricity meters," *Journal of Economic Behavior and Organization*, nov 2014, 107 (PA), 176–190.

—, **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.

Hoagland, Alex, David Anderson, and Ed Zhu, "Medical Bill Shock and Imperfect Moral Hazard," 2023.

Houde, Sébastien, "How consumers respond to product certification and the value of energy information," *RAND Journal of Economics*, jun 2018, 49 (2), 453–477.

— **and Erica Myers**, "Are consumers attentive to local energy costs? Evidence from the appliance market," *Journal of Public Economics*, sep 2021, 201, 104480.

- , **Annika Todd, Anant Sudarshan, June A. Flora, and K. Carrie Armel**, “Real-time feedback and electricity consumption: A field experiment assessing the potential for savings and persistence,” *Energy Journal*, 2013, 34 (1), 87–102.
- Jessoe, Katrina and David Rapson**, “Knowledge is (Less) power: Experimental evidence from residential energy use,” *American Economic Review*, 2014, 104 (4), 1417–1438.
- , —, **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.
- Lamp, Stefan**, “Sunspots that matter: the effect of weather on solar technology adoption,” *Environmental and Resource Economics*, 2023, 84 (4), 1179–1219.
- Liao, Yanjun**, “Weather and the decision to go solar: evidence on costly cancellations,” *Journal of the Association of Environmental and Resource Economists*, 2020, 7 (1), 1–33.
- 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.
- Myers, Erica**, “Are home buyers inattentive? Evidence from capitalization of energy costs,” *American Economic Journal: Microeconomics*, may 2019, 11 (2), 165–188.
- Sallee, James M.**, “Rational Inattention and Energy Efficiency,” *Journal of Law and Economics*, 2014, 57.
- Sexton, Steven**, “Automatic bill payment and salience effects: Evidence from electricity consumption,” *Review of Economics and Statistics*, 2015, 97 (2), 229–241.
- Zivin, Joshua Graff and Kevin Novan**, “Upgrading efficiency and behavior: Electricity savings from residential weatherization programs,” *Energy Journal*, 2016, 37 (4), 1–23.

A Online Appendix

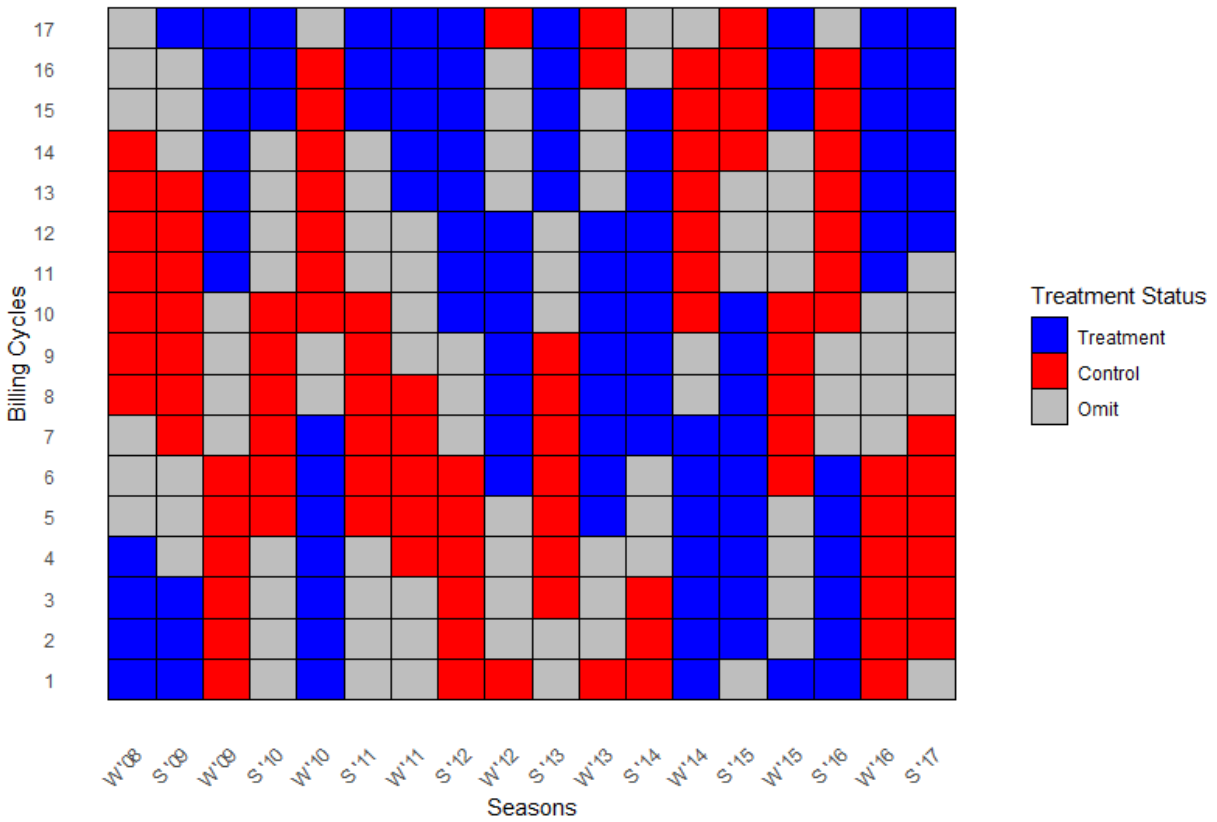
A.1 Figures

Figure A.1: Event Timing by Season

Season	Peak Weather Event	Invst. Window Start Date
Summer 2008	July 3 - July 22	Aug-15
Winter 2008	Jan 9 - Jan 28	Feb-20
Summer 2009	Aug 8 - Aug 27	Sep-22
Winter 2009	Dec 29 - Jan 17	Feb-9
Summer 2010	July 6 - July 25	Aug-18
Winter 2010	Jan 13 - Feb 1	Feb-25
Summer 2011	July 5 - July 24	Aug-17
Winter 2011	Jan 3 - Jan 22	Feb-14
Summer 2012	June 29 - July 18	Aug-13
Winter 2012	Jan 22 - Feb 10	Mar-6
Summer 2013	July 4 - July 23	Aug-16
Winter 2013	Jan 22 - Feb 10	Mar-6
Summer 2014	June 26 - July 15	Aug-8
Winter 2014	Feb 12 - Mar 3	Mar-26
Summer 2015	July 18 - Aug 6	Aug-27
Winter 2015	Jan 5 - Jan 24	Feb-16
Summer 2016	July 13 - Aug 1	Aug-26
Winter 2016	Jan 29 - Feb 17	Mar-14
Summer 2017	July 2 - July 21	Aug-14


Notes: This figure shows the dates for each season's 20-day peak heat wave or cold snap and the start date of the 6-month investment window.

Figure A.2: 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.3: 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 status 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.4: Coefficient Plot by Window Length



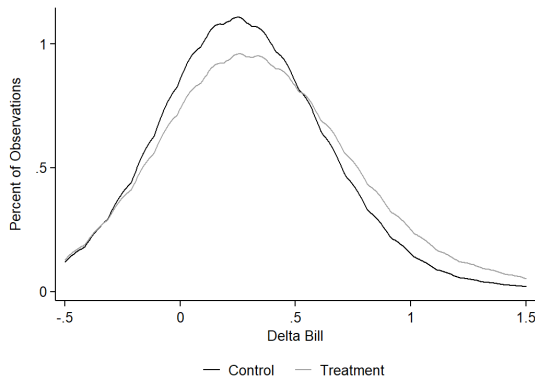
(a) Heat Wave




(b) Cold Snap

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

Figure A.5: K-Density Plot for Delta Bill




(a) Heat Wave



(b) Cold Snap

Notes: This figure reports the distribution of the delta bill variable for treatment and control observations separately for heat waves and cold snaps using a k-density plot.

Figure A.6: 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 the category for investments made by customers through United Illuminating's energy efficiency investment program

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 through United Illuminating's energy efficiency investment program in the Site visits: audits and inspections category.