Something for Nothing: The Who and the How Much of Bill Salience Effects on Energy Consumption

Brett Close, TRC, Seattle, WA
Derek Dinsmoor, TRC, Seattle, WA
Galib Rustamov, SCE, Rosemead, CA

ABSTRACT

Utilities across the country have relied on behavioral programs as a source of residential energy savings for over a decade. These programs provide information and behavioral "nudges" to help customers reduce their energy consumption. Academic research has shown that receiving a utility bill can have a similar effect, leading to average energy savings on the order of 0.6-1.0% in the week following the receipt of a bill (Gilbert and Zivin 2014). The hypothesized mechanism is that receiving the bill and seeing the financial impact of energy consumption raises the salience of energy costs for the customer, who then allocates more attention to their energy consumption for a period of time and decreases low-valued energy consumption (such as turning off lights or appliances when not in use, turning off HVAC when not needed, etc.). Previous research was unable to disaggregate impacts between customers who receive and pay their bills in different modes, such as email vs. mailed paper bills, and autopay vs. mailing in a check. However, based on the differences in customer engagement involved with these different methods, impacts are likely to vary between these different customer groups. That is, customers who receive email bills (which they may not even open) and have autopay set up (and may not even actively notice how much they paid for their utility bill) are likely to have less of a response than customers who receive a paper bill that they open to see the amount in order to write a check or actively submit a payment. This project uses utility meter data as well as information from customer databases to better understand salience for customer groups and their corresponding energy impacts. Based on estimates of the bill at a customer level, we will segment the results to understand which customer groups experience these bill salience impacts, and how the magnitude of those impacts varies between groups.

Introduction

This paper discusses the results from a study of bill salience impacts among Southern California Edison (SCE) customers. The purpose of this research was to help SCE understand bill salience for its residential energy customers, specifically if there is a decline in energy usage following the receipt of a bill, and to what extent. In other words, does receiving a bill make the cost of energy more salient, thus decreasing usage relative to the period before the bill in order to decrease the cost? If so, how long does the decrease last and for whom? Under what conditions does the decline happen? By utilizing this research, utilities can better-understand which groups of customers experience different effects, shedding light on areas of opportunity to increase energy savings. Evaluators can leverage this research to better-understand how to incorporate metadata and machine learning techniques to estimate savings for various customer segments, going beyond an average treatment effect for all customers. Implementors may see this research as an area of opportunity for improvement or refinement of the utility bill process to help increase salience.

Overview of Relevant Theory and Evidence

The notion of bill salience impacts is based on the idea that attention and cognitive resources are scarce. People have countless demands on their attention and they must allocate that attention among competing demands. One of those demands is their energy use: did they remember to turn the lights out in the room they left? Is the thermostat set for the right balance of cost and comfort? Receiving an electric bill and seeing the price for the electrical consumption one has enjoyed during the previous month may raise the salience of energy consumption, leading people to pay slightly more attention to it, at least for a period of time, before other demands for attention supplant it.

The most relevant research on this topic comes from Gilbert and Zivin (2014). The authors found that there was an average decrease in customer's usage after they received a utility bill. That study, however, is limited by the inability to analyze the data by whether the bills were sent by mail or email, or if customers paid through an automatic bill pay channel or another method. Arguably, when studying bill salience, these details must be known to go beyond the average impacts to understand how the responses differ. This is highlighted by the finding in a study by Sexton (2015) that automatic bill payment (ABP) increased residential electricity consumption by 4% for a publicly owned electric utility. The importance of payment channel is also highlighted in DNV GL's research for the California Public Utility Commission on Pacific Gas & Electric residential customers where they found that customers who enrolled in ABP or Budget Billing (BB) programs have higher energy usage than other customers (DNV GL 2017; Getachew, Agnew, and Sadhasiva 2018). That study, however, noted this increase in energy usage was mitigated for customers enrolled in Home Energy Report (HER) programs. Home Energy Report programs provide customers with information about energy usage (price and quantity); therefore, the provision of this information seems to be important.

The importance of energy usage information is supported by a randomized control trial studied by Jessoe and Rapson (2014) that mixed increases in prices during an event period with information provided by an in-home display (IHD).[1] The authors of that study were interested in understanding the price elasticity of electricity consumers and asked if the lack of complete information made consumers appear inelastic to price or if consumers actually were price inelastic. They found that relative to the control, residential electricity customers who were exposed to information feedback via their IHD's during an increased pricing event exhibited an 8-22% decrease in energy usage in that event period and households who only experienced the price increase reduced demand by 0-7% during pricing events (households in both treatment groups were informed of upcoming price increases via email, voicemail, or text message). Thus, with real-time information, households became more elastic to price. Furthermore, they found that the treatment effects spilled over into non-event hours and days as customers formed conservation habits.

Timing of the bill has also been investigated as an important factor in energy demand. Households in Germany receive their energy bills just once a year, in contrast to monthly billing common in the United States. A recent study by Singhal (2020) found that German households used less energy for heating annually if their bill preceded the winter heating season as compared to households that were billed during off-winter months.

The above evidence suggests that timing, program enrollment, payment channel, and information feedback are important considerations for bill salience research. Through the use of analytical approaches using advanced metering infrastructure (AMI)[2] data and SCE's customer data, we were able to identify whether a decline in energy use is occurring, for how long and for whom, and under what conditions.

[1] The study looked at effects of a real-time pricing on energy use. Participants in the study were randomized into a control group, a group with pricing information conveyed by day-ahead notification, and a group that received day-ahead notification and an IHD. The participants who received an IHD were able to view real-time information about their energy use, such as how much usage went up when they turned on an appliance.

[2] AMI meters provide frequent reads such as daily, hourly, or sub-hourly, as opposed to monthly.

Overall Approach

Our approach uses statistical techniques to detect whether customers' energy consumption after receiving a bill is lower than it was before receiving the bill, controlling for other key drivers of consumption, such as temperature. Using AMI data, we estimated these energy usage changes using a within-subject modeling approach to provide estimates for each unique account and bill combination. After joining a series of metadata for each account, we identified key variables correlated with these estimates across all customers using machine learning techniques. Then, we performed an impact segmentation analysis to understand the differences between these groups of customers and quantify bill salience estimates within various subgroups. Our approach is described in greater detail in the methodology section of this paper.

Key Findings

We found a number of interesting results, all of which are described in detail in the main body of the paper.

Positive Salience Effects

- On average, our models indicate that there is a positive salience effect of approximately 0.18% during the 10 days after a billing statement is created.
- These savings primarily occur on the 6th and 7th days after a statement is created, with savings of approximately 0.8% on those days.

Engagement with Delivery and Payment are Key

- Our models indicate that accounts which receive an electronic bill and use a passive form of payment (e.g., auto-pay) have approximately no salience effect.
- Accounts which receive a paper bill and/or use an active form of payment, have a positive salience effect.

Income-Related Differences

- On average, accounts with lower income have higher salience effects.
- Several programs, payment information, and demographic variables show higher salience effects for lower income groups.

Methodology

The modeling portion of the analysis involved two main components, identifying key variables and estimating salience impacts.

Identifying Key Variables

We estimated bill salience impacts for each unique account and bill combination, controlling for weather using a within-subject modeling approach. That is, for each account, we estimated a bill-level salience impact controlling for outdoor air temperature. Then, we used machine learning techniques, such as generalized linear models, gradient boosting methods, and others including ensemble models that combine multiple component models, to identify which variables in the metadata showed a high correlation with these estimates. The process involved using auto machine learning techniques to construct several models to fit the data and selecting the top-performing model based on root mean square error (RMSE). We experimented with automated machine learning on the regular dataset as well as one-hot encoded[3] variables to see how the results varied by data preparation. The top performing model was a stacked ensemble model composed of several underlying models. These models suggested that salience impacts are correlated with month, year, income, bill delivery type, and various payment information associated with the account. We investigated these variables, along with several other variables of interest to identify the bill salience impacts within each segment.

Estimating Impacts

We estimated overall percent savings, percent savings by "post day," and percent savings by hour of the day during the 10 days following the statement. Looking at these three different categorizations of savings allowed us to understand better when and if there is a bill salience effect. We modeled the hyperbolic arcsine of the hourly energy use as a function of a post-period indicator and hyperbolic arcsines of cooling- and heating-degree hours. We include account-level fixed effects to control for unobserved factors that do not vary over time (such as demographic information), and month-year fixed effects to control for seasonal variation not captured by outdoor air temperature. The hyperbolic arcsine transformation, much like the logarithmic transformation, allows us to estimate percentage changes of the marginal impacts, but unlike a logarithmic transformation, hyperbolic arcsine is defined at zero, allowing us to include observations with zero cooling degree days or heating degree days. We clustered the standard errors at the account level to allow for arbitrary correlation between observations for the same account. The following equation presents a representative model formula:

$$\text{asinh}\,(kWh)_{it} = \beta_1 Post_{it} + \beta_2 \text{asinh}\,(CDH)_{it} + \beta_3 \text{asinh}\,(HDH)_{it} + \beta_4\,\text{asinh}(CDH^2)_{it} + \beta_5 \text{asinh}\,(HDH^2)_{it} + \beta_6 \text{asinh}\,(CDH * HDH)_{it} + \alpha_t + \gamma_i + \epsilon_{it}$$

Where:

kWh_{it}	is the average kWh consumption by account *i* on day *t*;
$Post_{it}$	is a binary variable taking the value of 0 during the 7 days prior to a bill statement being created and 1 during the 10 days following creation of the statement for account *i* at time *t*;
CDH_{it}	is the number of cooling degree hours for account *i* at time *t*;
HDH_{it}	is the number of heating degree hours for account *i* at time *t*;
β_k	are regression slope coefficients;
α_j, *and* γ_i	are year-month and account fixed effects; and
ϵ_{it}	is an idiosyncratic error term.

Results

Overview

This section describes the results of the impact segmentation analysis. Table 1 presents summary statistics for key groups based on bill delivery and payment type. The number of accounts reflects the total number of accounts that exist in that group for at least one bill, so the sum will be greater than the total number of accounts as many accounts exist in more than category in our sample.

[3] One-hot encoding is a way of representing mutually exclusive categorical variables in numeric form.

Table 1.Summary Statistics for Key Groups

Delivery	Payment	Average Usage (kWh/day)	Usage Standard Deviation	Number of Accounts	Number of Bills
Electronic Delivery	All	22.2	18.1	13,055	249,059
Paper Delivery	All	20.3	17.5	10,616	194,386
Unknown	All	20.6	19.5	9,409	60,496
All	Active Payment	21.9	18.1	18,205	312,203
All	Passive Payment	19.6	15.7	4,633	88,776
All	Unknown	21.0	19.7	17,332	102,962
Electronic Delivery	Active Payment	22.7	18.4	10,135	152,313
Electronic Delivery	Passive Payment	20.2	15.9	3,557	66,095
Electronic Delivery	Unknown	24.4	20.4	7,787	30,651
Paper Delivery	Active Payment	20.6	17.3	9,495	144,460
Paper Delivery	Passive Payment	17.9	14.9	1,268	22,681
Paper Delivery	Unknown	21.3	20.0	6,621	27,245
Unknown	Active Payment	26.3	20.6	2,490	15,430
Unknown	Unknown	18.6	18.7	9,154	45,066
All	All	21.3	18.1	21,871	503,941

The figures in this section display percent savings Average Treatment Effect (ATE) point estimates along with their 90% confidence interval using cluster-robust standard errors at the account level. A positive ATE indicates our models show a decrease in energy usage following the creation of a bill statement. The confidence intervals show the uncertainty level around the point estimate, attributable to uncontrolled variation in the data or small sample sizes. The n value shown next to each segment displays the total number of accounts which represent the particular subgroup at least once in the data.

Overall, we found an average bill salience effect of 0.18% across the 10 days of the post-period we analyzed. This effect was statistically significantly different from 0 at the 90% level. Because of variation in prices across customers and across time, there is not direct way to convert this energy reduction into a bill reduction. The largest impacts occurred on days six and seven after the statement was created, as shown in Figure 1. This general pattern of savings (no savings on the first few days, savings on the following days peaking around day six or seven and then declining to zero) was repeated across all the sub-group analyses we conducted.

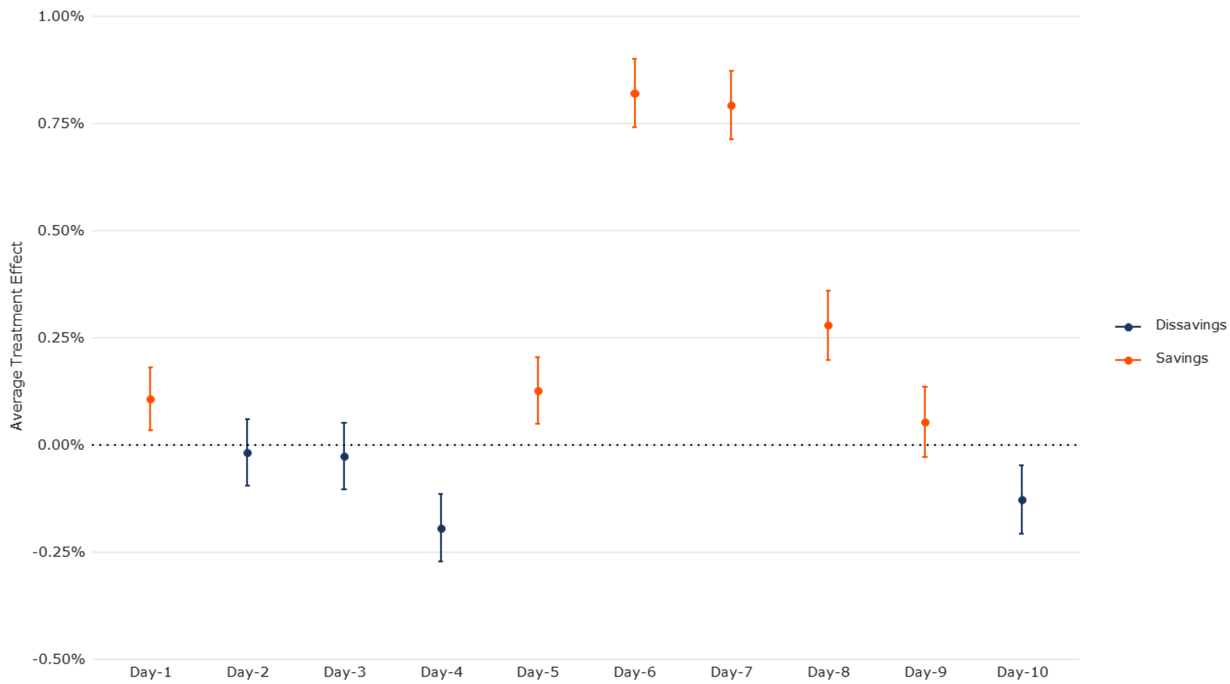

Figure 1. Average Percent Savings by Day

An hourly analysis of the results shows that the impacts are greatest in the morning between 6 am and 8 am and in the afternoon and evening between 4 pm and 11 pm, as shown in Figure 2.

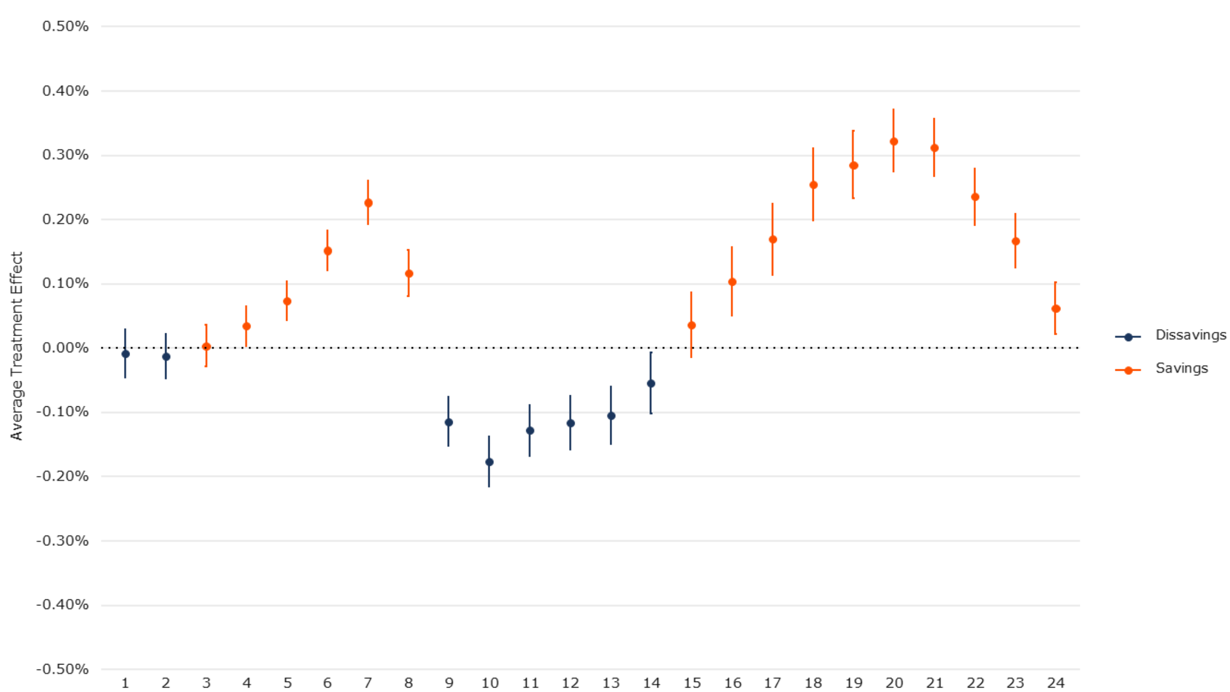

Figure 2. Average Percent Savings by Hour

Bill Delivery Type

As discussed above, knowing the type of bill a customer receives (such as a paper bill in the mail or an email from the utility website) is important for understanding the mechanism by which bill salience operates. The data in our analysis contain three primary bill delivery types: CheckFree, paper bill, and SCE.com. CheckFree bills and SCE.com bills are electronic forms of delivery whereas paper bills are mailed to customers. We explored how salience impacts varied between these three delivery types.

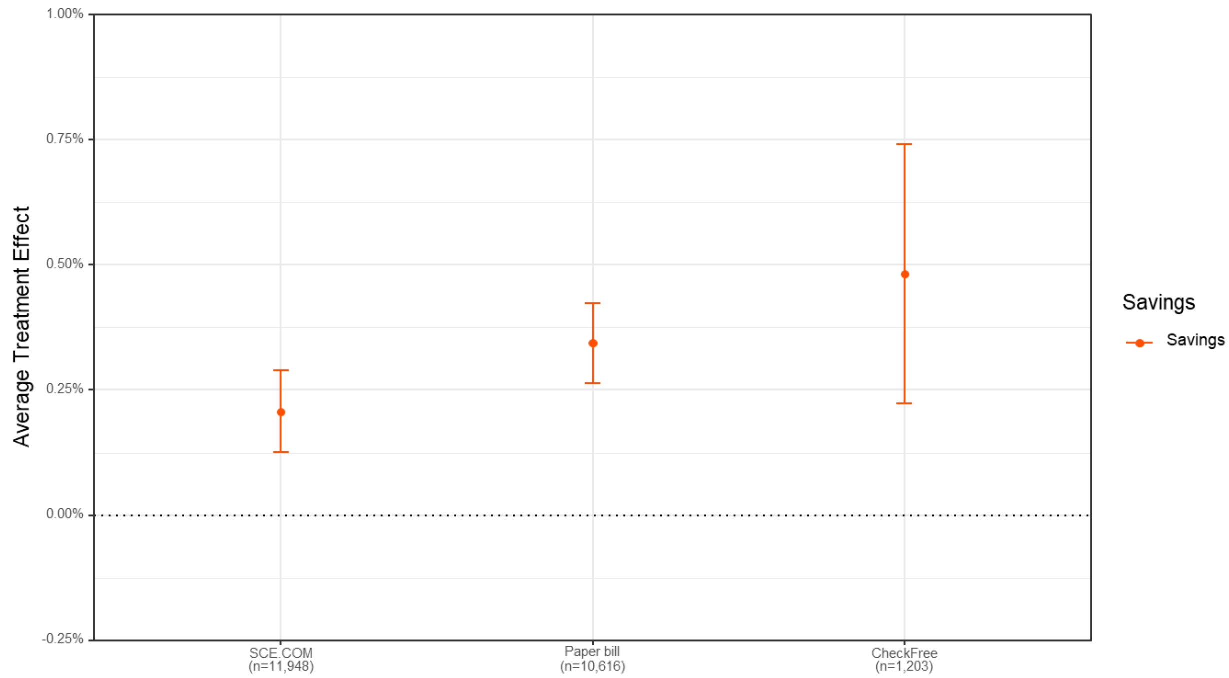

Figure 3. Average Percent Savings by Bill Delivery Type

All three delivery types have positive point estimates, meaning on average, there is a reduction in energy usage after the statement is created for each of these channels. None of the point estimates appear to be statistically significantly different from one another. When broken out by post-day, the positive salience effect for paper bills appears on 9 of the 10 days after the statement is created whereas bills delivered through SCE.com have roughly half of the days.

Overall, the results from the impact segmentation on bill delivery type suggest each channel has a positive salience effect. The savings for paper bills appear to be slightly higher than bills through SCE.com, although this difference is not statistically significant. Intuitively this seems plausible – a paper bill shows up in the mailbox, providing a signal to the customer that a payment is due. Further, the paper bill appears to have several days of savings. An illustrative example may be a customer opens the bill and leaves it on the counter, serving as a signal for several days. Alternatively, a digital bill appears to have a slightly lower salience effect – the bill may not be accessed or simply may be lost in the mix of other digital information, thus lowering the overall point estimate.

Bill Payment Channel

SCE offers a variety of payment channels through which customers may pay the balance on their bill. We estimated impacts by payment channel for channels with enough accounts in the data to explore these differences. The models identified small, positive salience effects for authorized payment agencies, credit/debit cards, and mailed payment channels. Further, payments through authorized payment agencies and credit/debit cards appear to have positive impacts for several days after the statement is created. The other channels appear to have approximately no salience effect. We did not identify any substantial differences within each payment channel broken out by income. Results are shown in Figure 4.

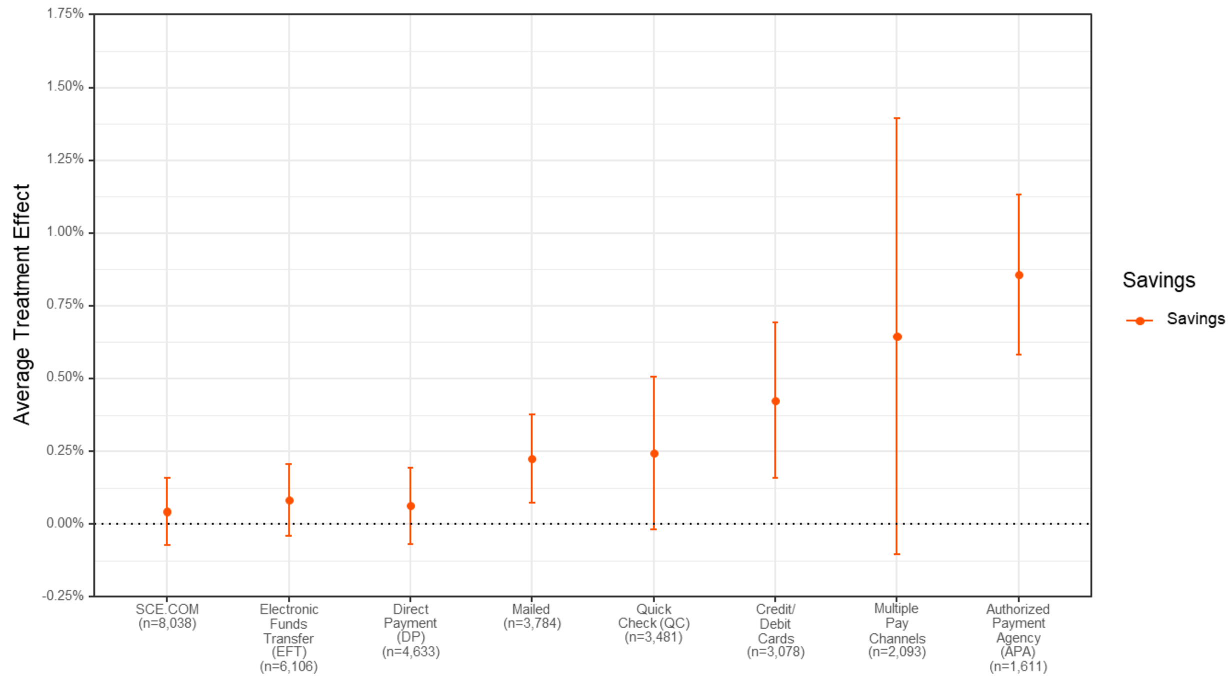

Figure 4. Average Percent Savings by Payment Channel

Some of these payment channels require an active form of payment, meaning the customer must take action each billing period to pay their bill (e.g., mailing a check). Other payment channels allow a passive form of payment, allowing the customer to utilize automatic payment options online. Active channels such as “Mailed” shows positive impacts in our model. This seems plausible – the customer actively has to take time to pay their bill, thus increasing the salience of the bill. Passive channels such as “Direct Payment” (a common autopay method), show approximately zero salience impact in our model, suggesting that lower levels of engagement with the bill may lead to less salience.

Bill Delivery and Payment

Understanding bill delivery type and bill payment type differences are only one piece to the puzzle. It’s plausible that the salience effect is correlated with both of these channels. To investigate this hypothesis, we classified bill delivery types and payment types into broader groups. We classified bill delivery types into either paper or electronic delivery and classified payment channels into either active or passive payments. Our model results from this two-by-two classification are shown in Figure 5 below.

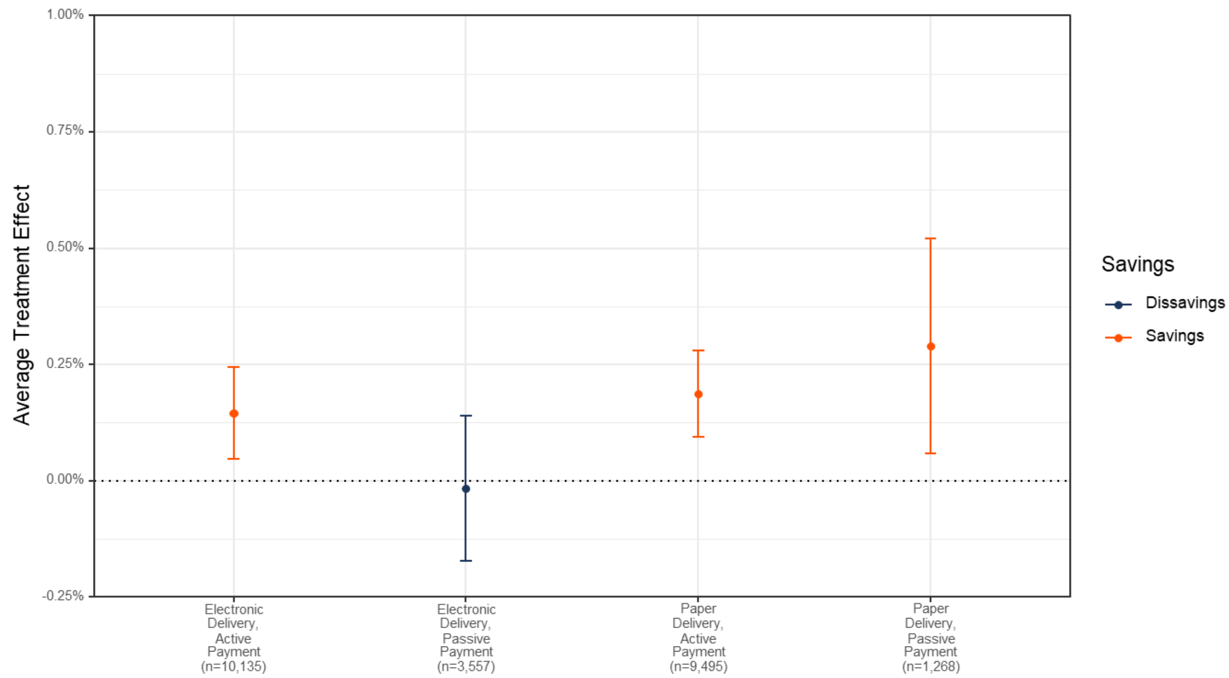

Figure 5. Average Percent Savings by Bill Delivery Type and Payment Type

Our two-by-two classification and modeling approach showed that bills with an electronic delivery and a passive form of payment have approximately no salience effect whereas all other classifications show positive effects. This suggests that customers who either receive a physical (paper) bill or actively have to make a payment for their bill, have higher salience effects. We did not identify any substantial differences by post day or income level.

Income-Related Findings

We assigned accounts into income quantiles based on each customer's discrete income level and the number of persons living in their household. We investigated the relationship between income and bill salience across various programs, demographics, and payment information. The results of the impact segmentation on income and our income-related findings are described below.

Overall, lower-income accounts had greater bill salience impacts than higher-income accounts, as shown in Figure 6. In general, we found that indicators of less disposable income aligned with findings of greater bill salience impacts, as would make sense given the underlying theory: if a household's disposable income is more constrained, allocating scarce attention to energy use could lead to greater increases in utility in the presence of decreasing marginal utility of income. Given the data available to us, most of the evidence comes from billing system data. There are many reasons a bill could go unpaid or be paid late, not all of which are related to disposable income, but we assume that conditional on other factors (including an account fixed effect), bills are more likely to be paid late or not be paid if the account holder is experiencing a greater income constraint.

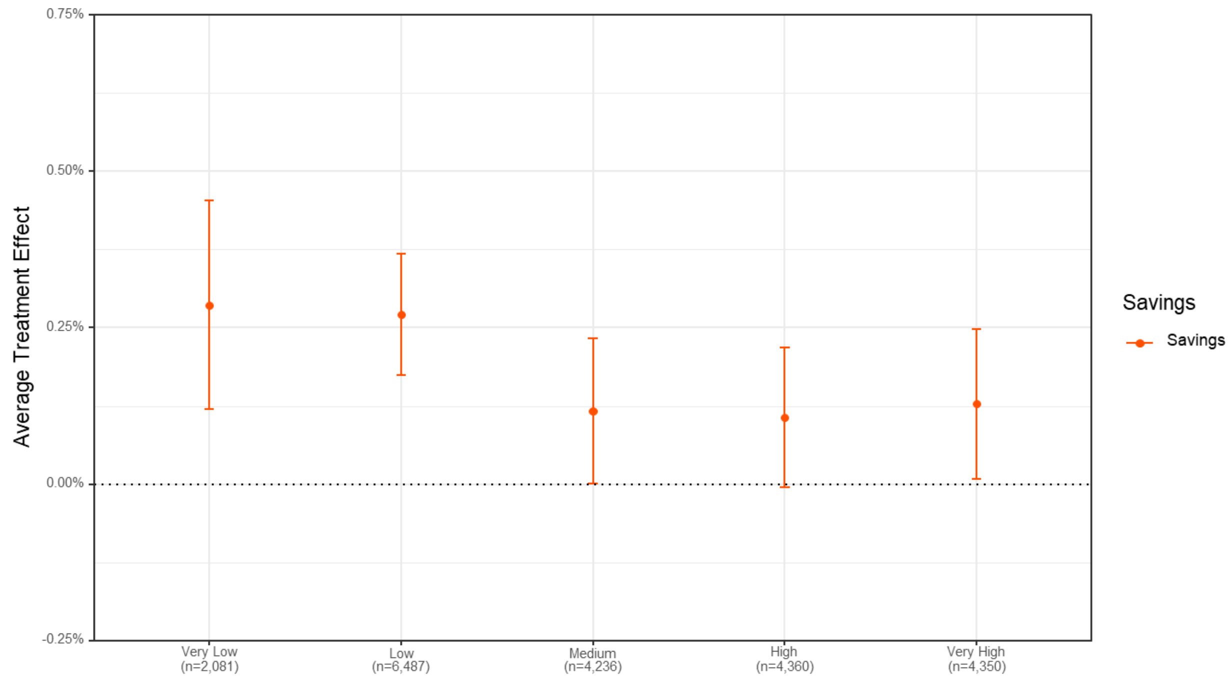

Figure 6. Average Percent Savings by Income Group

Due to the Covid-19 pandemic, SCE had a moratorium on power shut-offs for non-payment. We found that bills for which the customer was protected by the moratorium (i.e., would have been in danger of a shut-off if not for the moratorium policy) had a much higher salience impact than bills that were not.

Late payment and non-payment were two additional factors we looked at. We defined a bill as unpaid if no payment was submitted between the bill date and the following bill due date, and as late if there was a payment after the due date, but before the next bill due date. Accounts that had a positive number of unpaid bills had a larger bill salience impact than accounts that had no unpaid bills during the study period, and unpaid bills were associated with larger bill salience impacts than paid bills. We did not find a difference in bill salience impacts between accounts with a positive number of late payments and no late payments during the study period, but bills that were paid late had a larger bill salience impact than those paid on time. Accounts with fewer than ten late payments had a much larger average bill salience impact than those with ten or more late payments.

Other Notable Results

Impacts before and during the Covid-19 Pandemic

The AMI data received for the analysis spanned from mid-2018 through August 2020. In order to investigate how bill salience changed during the COVID-19 pandemic, we defined a “Post COVID” period as March 2020 through August 2020 and a “Pre COVID” period as March 2019 through August 2019. We modeled the average differences between these two periods to see how impacts change during these two partial-years. The results showed approximately equal impacts between the two periods, suggesting that the COVID pandemic did not have an overall effect on bill salience.

Although our data did not identify a clear salience difference pre- vs. post-COVID, we are limited by partial year data. Further research could be conducted to analyze if salience changes between the two periods by looking at a longer timeframe or various subsegments.

Net Energy Metering

The metadata allowed us to analyze differences between accounts with Net Energy Metering (NEM) vs other accounts. The results from our models indicate that NEM accounts exhibit higher impacts than other accounts. Further, these impacts are shown on each of the 10 days after the statement is created. When broken out by income, NEM accounts all exhibit roughly the same impacts; however, non-NEM accounts with lower income levels exhibit higher salience than other income levels.

Our Net Energy Metering results suggest that NEM accounts have higher levels of salience than other accounts; however, this could be correlation not causation. For example, a NEM customer is likely highly engaged with energy compared to the general population. It may be that this undefined characteristic of NEM accounts is the driver of the higher levels of bill salience, rather than simply being a NEM account. It is important to keep this in mind when interpreting results from these models.

Home Energy Reports

We looked at the impact of Home Energy Report (HER) program participation. Using HER enrollment start dates, we estimated impacts for bill-account combinations which received HERs vs those which did not. The results show that accounts which receive a HER appear to have slightly lower salience effects than those which did not receive a report.

We also looked at the type of report each HER participant received: print only, email only, print and email, or neither report. It is unclear what type of information is received for customers who did not receive either report, so we classified them as their own group. We classified non-HER participants as receiving no HER type. The results from our models show that customers receiving neither report or no HER type have higher impacts than the other three groups. The HER accounts receiving some form of report exhibit approximately no bill salience effect in our data.

Our estimates suggest that HER accounts exhibit less bill salience than non-HER accounts. It seems plausible that accounts which receive HERs use the report as their signal for energy usage rather than their utility bill. We do not know when the accounts receive HERs each month, but if the report comes at a different time than the billing statement is created, we may see salience effects around the time the HER is received rather than a salience effect when the statement is created.

Overview of Additional Results

Table 2 provides an overview of additional results by customer segment.

Table 2. Additional Results

Segment	Finding
Time of Use Rate (TOU)	Non-TOU accounts show slightly higher salience effects than TOU accounts. Both TOU and non-TOU accounts with low income have higher point estimates relative to other income levels.
Budget Assistance Program (BUDG)	SCE offers a tool called Budget Assistant that monitors projected bill amounts relative to budget goals and sends the customer alerts about usage. Accounts not enrolled in the Budget Assistance Program show slightly higher salience than accounts that are enrolled. Among both groups, accounts with low income have higher point estimates relative to other income levels.
Dwelling Type	There were no clear differences between single family and multifamily households except that single family accounts with low income have higher point estimates relative to other income levels.
Household Size	There were no clear differences between households of different sizes except that households with 1 person living in the home exhibited a linear relationship with income (as income increases, impact estimates decrease).

Owner vs Renter	There were no clear differences between homeowners and renters except that homeowners with low income have higher point estimates relative to other income levels.
Home Square Footage	There were no clear differences between homes of different square footage levels except that homes in the .60 to .80 quantile of the data show a linear relationship with income (as income increases, impact estimates decrease).
Marital Status	There were no clear differences between single vs married accounts except that single accounts with low income have higher point estimates relative to other income levels.
Children	There were no clear differences between accounts with children vs accounts without children. Both types of customers showed low-income customers with higher point estimates relative to other income levels in their respective group.
Age	We found no clear salience differences between customer age groups.
Home Year	We found no clear salience differences between home year (vintage) groups.
Climate Zone	Climate zones 14 and 15 (in the desert) may have slightly higher salience than other climate zones but large with uncertainty.
Year-Month	We found no clear seasonality differences apart from the weather-driven effects controlled for in the model.

Ongoing work

A separate part of this project that is finishing up as of December 2021 is follow-up work focused on understanding customer perceptions of their bills. The analysis presented here found that bill type and payment method were important determinants of bill salience impacts. The ongoing work conducts interviews with SCE customers about their perceptions of their bills and seeks to gain better insight into what elements of the bill drive the salience effect and how people respond to their bills.

Conclusions

This research adds to the literature on understanding the impact of receiving a bill on short-term household energy consumption. The hypothesized mechanism is through increased salience: receiving a bill increases the salience of energy consumption (and the price of energy), leading people to increase their scarce attention to it. This allocation of scarce attention wanes over time, but is re-activated by the receipt of another bill. We do find a statistically significant 0.18% average bill salience effect. This is somewhat lower than the impact found by Gilbert and Zivin (2014) and it is unclear if the differences are due to different populations (different times and places) or methodological approaches. We find that customers who receive a paper bill, use an active payment method, or both, have positive average salience effects, but that customers who receive electronic bills and use passive payment forms have average salience effects that are not statistically significantly different from zero. We also find that lower-income households have larger average salience impacts than higher-income households, and that, in general, indicators that are correlated with financial strain (being protected by a shut-off moratorium, paying a bill late or not paying it) are correlated with larger bill salience impacts.

The research presented here demonstrates both that a small bill salience effect exists, and that its magnitude varies significantly across customers. A second part of this study seeks to deepen this understanding through user experience interviews with a sample of customers to understand better how they engage with their bill. Extending this research to tie specific bill design modifications to increased or extended salience could lead to greater reductions in energy use. Identifying ways to achieve salience in low-engagement (electronic bill and autopay) customers could lead to broader-based salience impacts. While these bill salience impacts are a small percentage of an individual customer's bill, the impacts add up quickly across an entire customer base. And given that all customers receive bills, increasing the effectiveness of the billing process for eliciting a salience impact really could provide significant value to utilities and to society.

References

DNV GL (2017). Auto Bill Pay and Budget Billing Impact Evaluation - Residential, California Public Utilities Commission. www.calmac.org. CALMAC ID: CPU0163.01.

Getachew, L., Agnew, K., Sadhasivan, G. (2018). Auto Bill Pay and Budget Billing Impact Evaluation - Residential Gas.

Gilbert, B., & Zivin, J. G. (2014). Dynamic salience with intermittent billing: Evidence from smart electricity meters. Journal of Economic Behavior & Organization, 107, 176-190.

Jessoe, K., & Rapson, D. (2014). Knowledge is (less) power: Experimental evidence from residential energy use. American Economic Review, 104(4), 1417-38.

Sexton, S. (2015). Automatic bill payment and salience effects: Evidence from electricity consumption. Review of Economics and Statistics, 97(2), 229-241

Singhal, P. (2020). Inform Me When It Matters: Cost Salience, Energy Consumption, and Efficiency Investments.