Bunching in Residential Electricity Consumption

Laura Grant*
Becka Brolinson†
ⓡ‡

September 26, 2022

Abstract

In this paper, we study how households living on US military bases responded to the introduction, and subsequent removal, of residential electricity prices. From 2006 through March 2019, tenants on military bases only paid for electricity use beyond a monthly allocation that varied by their housing type. For any use below that allocation, they received a rebate. The pricing schedule also included a "donut" around the allocation where electricity was free, generating large incentives for the households to change their electricity use in response to the nonlinear electricity price. After March 2019 all electricity for these households unexpectedly became free, providing us with the unique opportunity to evaluate how households responded to their pricing schedule. Under three alternative empirical approaches, we find that 39.64 percent more households to 3.41 percent more households located near that nonlinearity in their pricing schedule. These estimates correspond to price elasticities of electricity demand that range from -0.12 to -1.44, suggesting that households respond to more than just the nonlinear incentives in their electricity pricing schedule.

JEL Codes: H31, Q41, Q53, Q54

*lgrant@cmc.edu
†bmb105@georgetown.edu
‡Author order was generated using the AEA's random author order tool. Verification is available at: https://www.aeaweb.org/journals/policies/random-author-order/search. The authors appreciate useful comments from and discussions with Arik Levinson, Blake Shaffer, participants at the Georgetown economics seminar, the AERE Summer Conference, and OSWEET. Appendix available on request. Feedback welcome.

1 Introduction

LG and BB co-writing Intro

In this paper, we investigate the trade-off between progressive electricity pricing and simplification of electricity pricing. In theory, simplifying electricity prices can generate inefficiencies based on how a rational household should responding to those changes in the pricing schedule. But, it is possible that in practice, households do not respond to those prices in theory predicts that they should. In this paper, we ask

We measure the distortion in household electricity use due to the non linearity in the pricing schedule and intermittent billing.

We measure the distortions in household electricity consumption due to two billing practices introduced to simplify a utility billing program: (1) a nonlinearity in the pricing schedule and (2) intermittent billing that is the result of payments being collected only after a household's cumulative balance surpasses a certain dollar threshold.

Two simplifications of billing practices that can undermine the goal of conservation.

Literature: These billing practices are widely used in other

Progressivity is a widely targeted goal and a multitude of other programs, such as taxes, use this type of schedule to encourage redistribution. Discuss literature where households don't respond to nonlinearities in taxes. This literature leaves a gap: (1) in the tax literature people actually don't have much flexibility. The frictions are too high for them to be able to respond. This research is in a place where households aren't actually able to respond. Households in electricity are able to respond. (2) Literature in electricity finds that households respond to average prices or mis-respond to nonlinear rates for residential electricity bills.

NOTE: BOE calculation for households of what they should be doing if we took the off the shelf price elasticity of electricity where would we see them coming from.

Maybe the question is: Are they responding to the marginal price or are they responding to the target? Or both?

Contribution: Can we say something more about the mechanism?

We evaluate the trade off between energy conservation, progressivity, and the dimp

progressive pricing policies, simplification Pricing leads to We observe a progressive pricing schedule that, in theory, could generate inefficienct responses rela We evaluate the tradeoff between simplifications of residential electricity bills and RQ: testing household response to simplifications in billing

New Intro

1. Nonlinear prices are used in a lot of settings 2. Frequently it's been shown that people don't respond 3. But, if people have a lot of information and actually have control over the outcome variable, do they respond? 4. Yes, we find that they do respond. 5. Information component– when they get hit with the $15 subsidy they do respond, but these are probably the people that weren't bunching anyways because they are far enough away from the bunching range.?

In this paper, we investigate whether households respond to marginal prices in the way standard economic theory predicts. The setting is residential electricity consumption for households living on US military bases. This setting is particularly helpful for understanding whether households respond to changes in their marginal prices because before 2006, the US military did not charge service members for their electricity use in on-base housing. Starting in 2006, in order to encourage conservation and save taxpayer dollars, the military introduced a form of nonlinear pricing, known as a feebate. The households received a free allocation of electricity they could use each month, based on the average usage of similar types of homes. If the household used more electricity than that free allocation, they had to pay a flat marginal price per kWh. If the household used less than that allocation, then they receive a subsidy equal to the same price per kWh for the kWhs that they conserved.

The military and partner utilities understood that exactly targeting the free allocation might be difficult for households, so each base introduced a "free range" of electricity around the allocation. At both of the bases we study, that free range is equal to plus or minus 10 percent of the monthly allocation. This pricing schedule generates large, nonlinear incentives, to encourage households to cut back on their electricity consumption. This nonlinear pricing schedule leads us to ask—did these households respond to this nonlinear pricing schedule in a way that standard economic theory would predict?

In this paper, we develop a simple theoretical model based on Saez (2010) and Kleven and Waseem (2013) to generate predictions of how a rational household *should* respond to this pricing schedule for their residential electricity consumption. Theory suggests that household should "bunch" at the second nonlinearity in the pricing schedule at the upper end of the free range.

To empirically determine whether households bunched in respond to the nonlinearity, we leverage another unique aspect of the experience of households living on US military bases. In March 2019 suddenly and without warning reversed the pricing policy, reverting back to the old system of free electricity for all tenants in on-base housing. That reversal gives us a useful, exogenously set price of zero against which to contrast usage when prices were nonlinear.

We estimate counterfactual electricity consumption to understand what households would have done had the price been linear using data from 31,381 living at two US military bases from October 2014 through February 2020. We use a simple differences approach to estimate what the households *would have done* if there had been no "free range" generating the large changes in marginal electricity prices. This approach is similar to that of Kucko, Rinz, and Solow (2018), where we predict household electricity use using data from the free period as the counterfactual, controlling for household fixed effects, location fixed effects, and weather.

We use that counterfactual electricity use, and compare it to observed electricity use to show that 27.61percent more households consumed electricity at the upper end of the free-range than otherwise would have if there was no nonlinearity in the pricing schedule. This finding is quite surprising as previous research investigating household responses to nonlinearities in their tax schedules finds that only self-employed income earners responded to nonlinearies in their tax schedules (Saez 2010, Chetty et al. 2011, Kucko, Rinz, and Solow 2018, Mortenson and Whitten 2018). Previous work in electricity finds that households respond to average prices or that households mis-interpret their pricing schedule (Ito (2014), Shaffer (2020)).

So, in light of finding that households do respond to the nonlinearity, we investigate further how they could be responding to this pricing schedule. First, we calculate the average amount of electricity that the typical "bunching" houeshold would have had to use to take advantage of the full free range of electricity consumption. We find that amount to be XX kWh per day, which corresponds to using the typical air conditioner for X more hours per day, running your dryer for X number of loads more a month, or leaving your lights on for X number of hours per day. **To LG: version of possible paragraph for the BOE calculations.**

Other notes: Core of the paper is focused on the bunching. We also have these asymmetric RD results which are quite interesting. In the tax space– what are the inefficiencies caused by nonlinearities in the tax code? So, these papers ask how distortionary are these nonlinearities. I think that's a bit less relevant in our context, but we could characterize by how much they undermine the program by introducing this free range. It's probably not going to be huge so that might not be what we want to focus on. To tie the bunching with the RD:

So, what can we say about people's behavior? They bunch and they seem to respond to the subsidy component but not the charge. We see with income taxes that people don't bunch. This is a context where households do bunch. But people don't

Second, we use a secondary component of the pricing schedule where households only pay a bill, or receive a rebate, if their cumulative balance exceeds \$15. We find that households increased their electricity consumption by 4.2 percent if they *just missed* the cutoff for receiving the rebate. That 4.2 percent is similar to the average amount of electricity that the household would have to use to take advantage of the full "free range" of electricity consumption. This finding suggests that if households learn that they aren't conserving enough electricity to be eligible for receiving a rebate, that they then move in the opposite direction and increase their electricity use to consume the full amount of the free-electricity range. **NOTE TO LG: is that 4.2 close to what you find for the back of the envelope calculations?? If so, I think this is how we maybe square the RD results with the other bunching results**.

Third, we estimate the change in probability that a household bunches based on the season. We find that households, on average, are 1.6 percent more likely to use electricity at the upper end of the free-range when the nonlinear pricing schedule is in effect relative to the counterfactual. We also find that households are 3.5 percent more likely to bunch in winter months and 2.8 percent more likely to bunch in the summer. We also find that households that bunched in the previous month are more likely to bunch in the current month, suggesting that there is a share of the population that are "persistent bunchers."

Last, we calculate the implied price elasticities of electricity demand that result from the estimated levels of bunching. We follow Kleven and Waseem (2013) to calculate an upper bound of the price elasticity of demand. We find price elasticises around -0.73, which are high relative to previously estimated price elasticities of electricity demand, suggesting that households may be responding to more than just the nonlinear pricing schedule.

Former Intro

Before 2006, the US military did not charge service members in on-base housing for their residential electricity use. Starting in 2006, that benefit changed and households were charged market rates for their electricity use. But that price only applied to electricity exceeding a specific monthly allocation that varied by household size and location. Households that used less than the allocation received a refundable credit, a cash rebate, of that same price for every kWh they used less than the monthly allocation. Under the new pricing policy, electricity cost some positive marginal price per kWh, based on location and local electricity prices, either in the form of a foregone credit or an out-of-pocket payment.

In addition, the military did two other things that make this whole program especially interesting. First, they created a buffer zone around the allocation—typically 10 percent—in which electricity was free. So the typical household forgoes a credit for each kWh used up to 95% of the allocation, pays nothing for each kWh between 95 and 105% of the allocation, and pays out-of-pocket for each kWh used beyond that point. This nonlinear price schedule provides us with a new and valuable opportunity to test how households respond to nonlinear prices in an important setting.

Second, in March 2019 the military suddenly and without warning reversed the policy, reverting back to the old system of free electricity for all tenants in on-base housing. That reversal gives us a useful, exogenously set price of zero against which we can contrast usage when prices were nonlinear.

In this paper, we use these stark price nonlinearities and sudden changes to ask the following questions. Do households respond to this pricing schedule? And if so, are those responses consistent with what standard economic theory would predict? And if not, what other factors could be driving household electricity consumption decisions? We address these questions using monthly electricity billing data from 31,381 households at two US military bases from October 2014 through February 2020.

The unexpected change in electricity prices allows us to estimate how many more households located near the nonlinearity in the pricing schedule than otherwise would have if the nonlinearity did not exist. We use three alternative empirical approaches to predict what households would have done had there been no nonlinearity in the price. Two of the three approaches combine panel data on monthly household electricity use with the change in prices to estimate counterfactual electricity consumption under a flat electricity price. The third approach follows the standard estimation procedure in the bunching research from Saez (2010) and Kleven and Waseem (2013) and uses only the data prior to free electricity to fit a predicted counterfactual.

Under all three approaches we find that households do respond to the pricing schedule.

The estimates of the extra households in the distribution nearby the nonlinearity in the pricing schedule range from 39.64 percent more households to 3.41 percent more households under the most conservative approach.

Next, to evaluate whether these responses are consistent with the households responding to the sharp change in the marginal price we estimate an upper bound on the price elasticities of electricity demand following Kleven and Waseem (2013). We use the change in the electricity consumption of the marginal household caused by the pricing schedule and estimate elasticities that range from -0.12 to -1.44. These elasticities are large relative to previously estimated price elasticities of electricity demand (Ito 2014; Brolinson 2020).

There are several possible explanations for the household's large responses to their electricity pricing schedule. The first is the way the allocation system was designed. Households received information about their electricity use relative to their peer group of like-type homes (Lincoln RECP 2019). Allcott and Rogers (2014) shows that when households receive information about their electricity use relative to their peers, that they respond by changing their electricity consumption. However, in our sample, households received this information before and after the change in electricity prices. If they responded to that information in the same way in both time periods our estimates control for that response.

Second, standard economic theory predicts that households should respond to the rebate and the charge in the same way because the marginal cost is the same on either side of the allocation. However, Kahneman and Tversky (1979) show that losses may outweigh gains and so households may respond more to the marginal charge than to the rebate. We evaluate how individuals change their electricity use after receiving a rebate versus making a payment.

Third, electricity demand is derived from demand for electricity services such as lighting, heating, and cooling. Households faced uncertainty both in what their monthly allocation was an in what changes they would need to make to directly target that allocation. Households may better be able to target the free range when they have larger electricity allocations.

We find that in months when the allocations are larger, households are more likely to locate near the nonlinearity, suggesting that the optimization frictions present in this context may be smaller than those present in the taxation literature for wage and salary earners (Hausman 1982; Chetty et al. (2011)). We also find that households who have recently bunched are more likely to bunch in the following month. This finding aligns with previous research on household response to nonlinear tax rates, which shows that households who have previously located nearby refund-maximizing nonlinearities are more likely to do so in the future (Mortenson and Whitten 2018).

Our paper makes several contributions by providing: (1) The first study that uses individual panel data to compare a period with the incentive to bunch to its absence, an ideal counterfactual,[1] Previous techniques to estimate how households respond to a nonlinearity in their budget constraint fit polynomials across a discontinuity make strong parametric assumptions that can ultimately impact the estimates of elasticities yielded by a bunching estimator Blomquist et al. 2019. (2) The first evidence that nonlinear electricity prices can in fact induce bunching of energy consumption (3) Evidence of a real response to a change in marginal electricity prices whereas to date, most of the bunching literature demonstrates that *some households*, i.e. households with self-reported income, do respond to changes in tax incentives by changing their income (Saez 2010; Chetty et al. 2011; Mortenson and Whitten 2018; Kucko, Rinz, and Solow 2018). (4) Comparisons of estimates of bunching using a natural experiment and panel data to estimates under the standard approach of predicting a polynomial. And (5) estimates of the resultant elasticities, which demonstrate that elasticities estimated using bunching are sensitive to assumptions.

On the consumer side, literature in resource economics finds that households are generally insensitive to block-rates, a type of nonlinear marginal pricing that utilities implement (Borenstein 2009; Ito 2014; Shaffer 2020; Wichman 2014). Instead, the households typically

1. The most closely related paper is Kucko, Rinz, and Solow (2018) which measures bunching in response to the Medicare coverage gap. Theirs was the first paper to use a panel data approach to compare bunching across state borders. But, they do not use individual-level data.

respond to average price, which is the equivalent of responding to the total cost on their electricity bill. It is natural to ask if the household knows how much they are consuming. The answer is probably not. Electricity is an invisible force and indirectly consumed; a consumer derives all benefits from powering lights, refrigerators, computers, and mobile phones. The power meter is often outside the home. And, even if households read the power meter, it must translate changes in behavior into usage then into dollars. The effort to figure out where they fall in the pricing schedule and how to change behavior likely outweighs the possible savings on their bills, known as rational inattention (Sallee 2014; Grubb 2015). Utility bills provide some information, but these arrive in the mail or inbox after the electricity has been consumed. In our case, households needed to adjust their electricity use *relative* to their peer group, suggesting that optimization frictions in our context were less of a barrier.

This paper is closely aligned with the research in labor economics and public finance which demonstrates that income for wage earners is unresponsive to the discontinuities in marginal tax rates (Hausman 1982; Chetty et al. 2011). The exceptions are the self-employed, who report their own income, indicating that these taxpayers may lower their stated earnings to maximize their tax refunds (Saez 2010; Kucko, Rinz, and Solow 2018; Mortenson and Whitten 2018; Chetty et al. 2011). On reconsideration, the above authors attribute an employee's lack of responsiveness to frictions in their labor supply; most workers cannot choose the hours they work for an employer.

2 Household Pricing Schedule

Service members often rent, rather than buy, a home where they are stationed because of frequent moves due to reassignments. Bissell, Crosslin, and Hathaway (2010) estimate that 30 percent of military households live in on-base housing. The rent for these homes came with 'utilities included' until 2006. Then, the armed forces introduced a payment program for on-base households for electricity. If the price incentives saved the military money, as intended, the savings could be used to improve local amenities such as parks, playgrounds, and community centers (Lincoln RECP 2019). At the outset of the program, military families resisted paying for electricity. So, the military made concessions to engender community support. Utilities serving on-base households apply a flat rate per kilowatt hour, r in \$/kWh, but the pricing schedule has two unique aspects.

First, all households receive an allocation of electricity, which we refer to as $\bar{A}_{gt}$. Each household's allocation, $\bar{A}_{gt}$, is the average electricity use of the homes in their billing group, g, in month t. Billing groups are defined by similar home types in the same location. For example, at one base in the US one billing group included four bedroom houses in a subdivision built in the 1990s, whereas, a different billing group included two bedroom houses in another part of the base that were built in the 2010s. Households with electricity use above the free allocation, $\bar{A}_{gt}$, pay a price per kWh for those excess kWh; and households with use below the allocation receive a rebate for the difference between the allocation and electricity use in that billing cycle.[2] The allocation is a subsidy that offsets the *average* household's cost of electricity, aligning with residents' desire for 'utilities included,' while the marginal price provides the incentive to conserve electricity.

Second, all households received free electricity for an additional $\pm$%-range around each month's allocation. Each residential military location established the $\pm$%-range, which we will refer to as K, for the additional free electricity prior to implementing its program.[3] As a result, the pricing schedule has three parts, the rebate range , the free-electricity range, and the payment range. The rebate range covers use from 0 kWh to $(1-K)\bar{A}_{gt}$, the free range covers from $(1-K)\bar{A}_{gt}$ to $(1+K)\bar{A}_{gt}$, and the payment range covers all use above $(1+K)\bar{A}_{gt}$. The free-electricity range benefits high-use households, which made them more accepting of introduction of the program.

If c_{it} was household i's electricity use in month t, their bill in month $t+1$ would be

2. The pricing schedule is a version of a two-part tariff with a marginal price and a fixed monthly subsidy, i.e., a negative fixed fee.

3. In the two locations we study, the range is ± 5 percent of the group average, $\bar{A}_{gt}$, spanning consumption from 95% to 105% of $\bar{A}_{gt}$ for each billing group.

defined as:

$$Bill_{i,t+1}(c_{it}) = \begin{cases} r[(1-K) * \bar{A}_{gt} - c_{it}] & \text{if } c_{it} < (1-K)\bar{A}_{gt} \\ r[c_{it} - (1+K)\bar{A}_{gt}] & \text{if } c_{it} > (1+K)\bar{A}_{gt} \\ 0 & \text{Otherwise} \end{cases}$$

The first line represents a rebate: net out the free-electricity range and receive a check for the remainder of the amount under the group average. The second line represents a charge: net out the total amount up to the upper end of the free-electricity range and the household is charged for the overage.

The free-electricity range generates incentives for some households to increase their electricity use. In the next section, we describe a theoretical framework to generate predictions about household response to the discontinuities in the electricity pricing schedule.

3 Conceptual Framework

In this section, we present a conceptual framework to generate predictions and comparative statics about consumer behavior in response to their nonlinear pricing schedule. Our framework for a representative household builds on the model of taxpayers presented in Kleven and Waseem (2013). A household's utility depends on electricity consumption, c, and consumption of a composite good, x. The household faces prices of x, normalized to \$1, and electricity, r in \$/kWh, and has income \$$I$.

Figure 1 shows a stylized version of the household's optimal choice under three alternative budget constraints. First, suppose that the household faces a flat electricity price, without an the free allocation. The household chooses bundle c_0 to maximize utility. Second, suppose the household receives an electricity allocation A. The effect is theoretically identical to additional income—the budget constraint shifts out by the \$-value of the allocation, \$$r * A$ to the dotted line in Figure 1. The allocation is a lump-sum transfer, rather than a discontinuity, and should not distort incentives.

Figure 1: Optimal Consumption

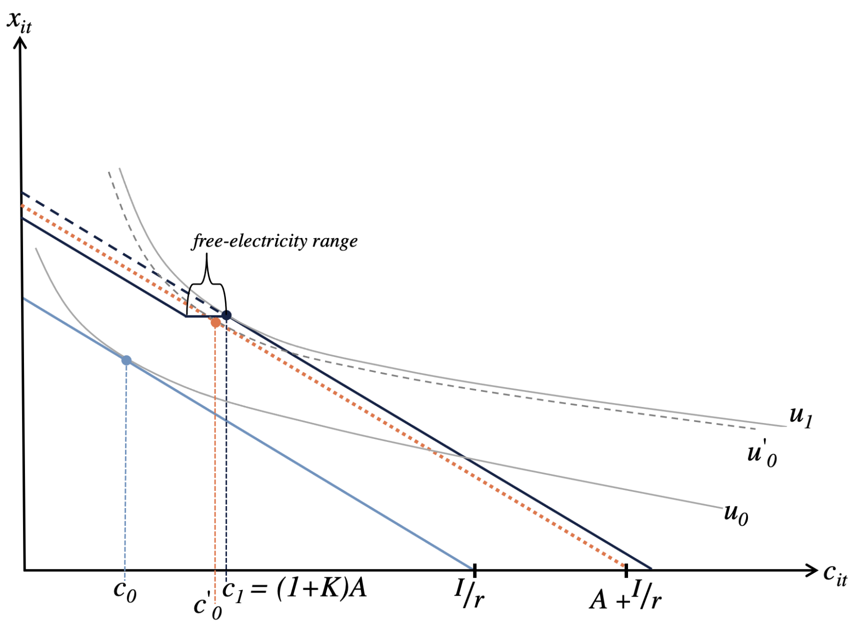

Note: Utility maximization, conditional on three alternative budget constraints. The solid blue budget is a constant electricity pricing schedule at rate r, with optimum of c_0. The orange dashed line indicates that the allocation A of electricity acts like income, increasing consumption to $c_0^{'}$. The solid black line shows the same allocation with a free-electricity range causing a nonlinearity and moving optimal consumption to c_1, the top of the free-electricity range. *Source*: Generated by the authors.

Third, suppose that the military now introduces the free electricity allocation. The solid nonlinear budget constraint portrays the free-electricity range as ± 5 percent of A, spanning consumption $[0.95A, 1.05A]$. The flat region of the budget dark, solid, nonlinear budget constraint in Figure 1 represents the free-electricity range. In the terms laid out by Kleven and Waseem (2013), the free-electricity range generates a "pure notch" in the budget constraint—a discrete change in the budget constraint at a cutoff but no change in the marginal rate on either side.[4]

We use this conceptual framework provide testable hypotheses, which compare the households' responses to the free-electricity range to a counterfactual without the free-electricity

4. The pure notch represents a discrete change in bill liability on either side of the threshold, but no change in marginal rates. A pure notch also generates nonlinearities in a household's average price schedule. This stands in contrast to a "proportional notch," which combines both a discrete change in the bill liability with a change in marginal rates in either side of the cut off (a kink). For more information on notches versus kinks see Kleven and Waseem (2013) and Kleven (2016)).

range. First, the free-electricity range should induce any household with electricity use within the free-electricity range, in the counterfactual, to consume the entire free amount, causing this household to bunch at the maximum of free electricity, $1.05A$. Any household with electricity use from $0.95A$ to $1.05A$ could increase their utility by increasing their electricity use to the top of the free-electricity range at $1.05A$. These households could consume more electricity free of charge, increasing their utility.

Moreover, there exist households with preferences such that their optimal counterfactual consumption is less than the minimum of the free-electricity range, $0.95 * A$; that with the free-electricity range, these households' optimal electricity use induces them to bunch at $1.05A$. Because electricity is a normal good, these households obtain a higher utility and consume more electricity. One example for these households is shown in Figure 1 where utility increases from $u_0^{'}$ to u_1 and consumption from $c_0^{'}$ to $1.05A$.

The free-electricity range transfers income from low-usage households to higher-usage households by cleaving the budget set into two segments. There are two types of low-use households that are worse-off under a pricing schedule with the free-electricity range. The first are households with such low electricity use, $c < A$, such that with the free-electricity range they decrease their electricity use. This is because the total subsidy they could receive is $0.95A$ which is less than A. The second are households for whom their optimal usage was originally less than A, however, with the free-electricity range their best option is to consume at $1.05A$, but they obtain a lower utility. In general, a household with counterfactual consumption of $c > A$ is strictly better off with the free-electricity range, whereas, a household with counterfactual consumption of $c < A$ may be better or worse off. Note, these responses assume strong monotonicity of preferences and perfect information.

Our model deviates from Saez (2010) because rather than a small change in marginal prices, at the first discontinuity in the budget constraint households experience a 100 percent decrease in prices. Then, at the second discontinuity point, prices revert back to their previous levels. Because the price in the buffer range is zero, we cannot compute arc elasticities using the methods in Saez (2010). And, the marginal price of electricity on either side of the free-electricity range is the same.

Instead, to calculate elasticities we follow Kleven and Waseem (2013) and consider an approximation of the change in marginal price for the lowest-use buncher to calculate an upper bound on the elasticity of electricity demand. In this framework, there is a household that is perfectly indifferent between consumption $c_l < 0.95A$ and the upper end of the free-electricity range at $1.05A$. Figure 2 shows a hypothetical indifference curve for a household that is perfectly indifferent between c_l and $1.05A$. All households with electricity use between c_l and $1.05A$ should bunch at $1.05A$ because their utility will be higher. Thus, we expect to see "missing mass" in the distribution from c_l to $1.05A$ with these households bunching at $1.05A$, where prices switch from 0 back to r.

Figure 2: Marginal Buncher Determines Elasticity Upper Bound

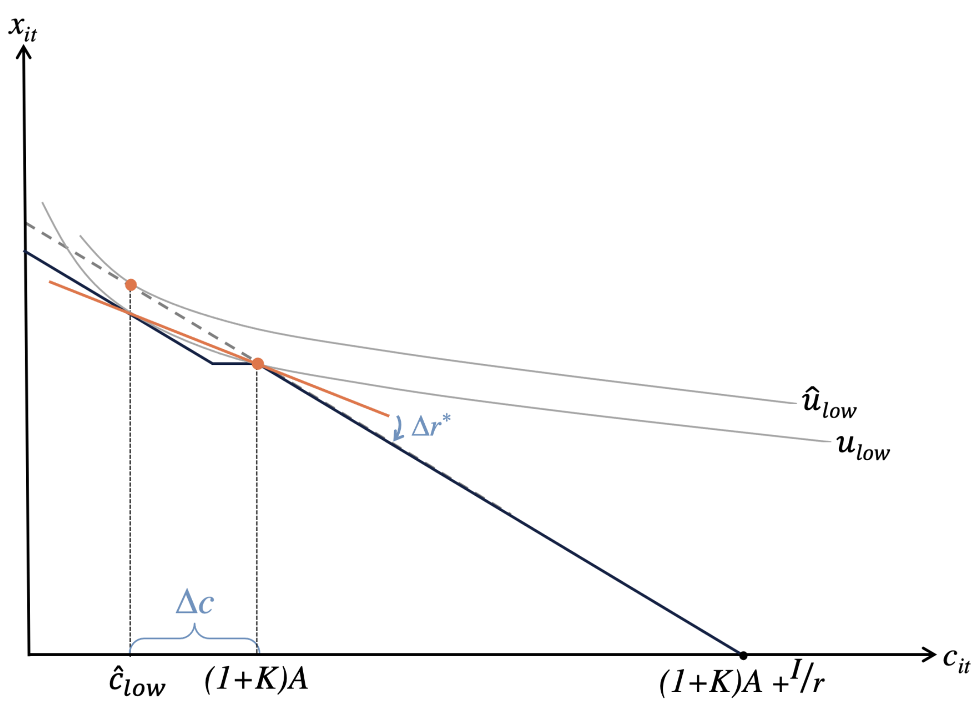

Note: We estimate the change in consumption for the household that is just indifferent to bunching, i.e., marginal. Following Kleven and Waseem 2013, an upper bound for the price elasticity of demand is percent Δc due to the percent Δr. *Source*: Generated by the authors.

The change in marginal price that the household would experience for a move from c_l

to $1.05A$ is the *implicit marginal price*, represented by the dashed orange link in Figure 2. Using this implicit marginal price overestimates the price elasticity of demand as the figure shows there are points that are strictly preferred to u_l on the dashed line (the implicit marginal price). We can use the change in consumption from c_l to $(1 + K) * A$ and the implicit marginal price r^* to estimate the price elasticity of demand:

$$e = \frac{\frac{\Delta c}{c}}{\frac{\Delta r^*}{r}} \tag{1}$$

where we treat the discontinuity in the budget constraint as the hypothetical change in prices from the implicit marginal price r^* to the original marginal price r. Because this representation does not take into the full change in price that the household experiences, it estimates an upper bound on the household's true price elasticity of electricity demand.

In section 5.2 we estimate the range with missing mass and the range with bunching mass. These estimates directly correspond to the cutoff point c_l of the marginal buncher, and allow us to calculate elasticities in Section 7 using the implicit change in marginal price at that consumption level and the elasticity estimated by Equation (1).

4 Data

We use monthly electricity billing data for households living on two Army bases in the continental US. For confidentiality, we cannot reveal which bases we use. Both locations have the same pricing structure. More data provides additional power, particularly because we only observe one year with completely free electricity. We clean the data, dropping outliers and observations with missing information; details are provided in Appendix A.1. We normalize all data to average daily values to account for differences in days per billing cycle. The data for analysis contain 455,351 observations for 31,381 households, tracking 9,307 homes from October 2014 through February 2020.

Figure 3 summarizes the average electricity allocation, in kWh/day, over time for the two locations. We weight these averages by billing group size, which vary. This figure shows several important characteristics of our data: First, we have more than four years of data with households facing the nonlinear pricing schedule and one year of data with households receiving free electricity. Second, electricity use is highly seasonal due to weather and the characteristics of the homes: the billing group averages, which are equivalent to the allocations, vary month to month. Location 1 has a cooler climate, using a large amount of electricity in the winter, despite half the households having gas heating. Summer air conditions usage also leads to high usage at Location 1. Location 2 is in a warmer climate with mild winters, leading to larger summer usage and resultant allocations. Lastly, Location 1 has higher overall usage, and allocations, than Location 2, on average; represented by the horizontal lines on the figure. Location 1 has more need for climate control year-round, despite the larger swings in usage at Location 2.

Figure 3: Allocation in kWh/Day for all Months of Observation, by Location

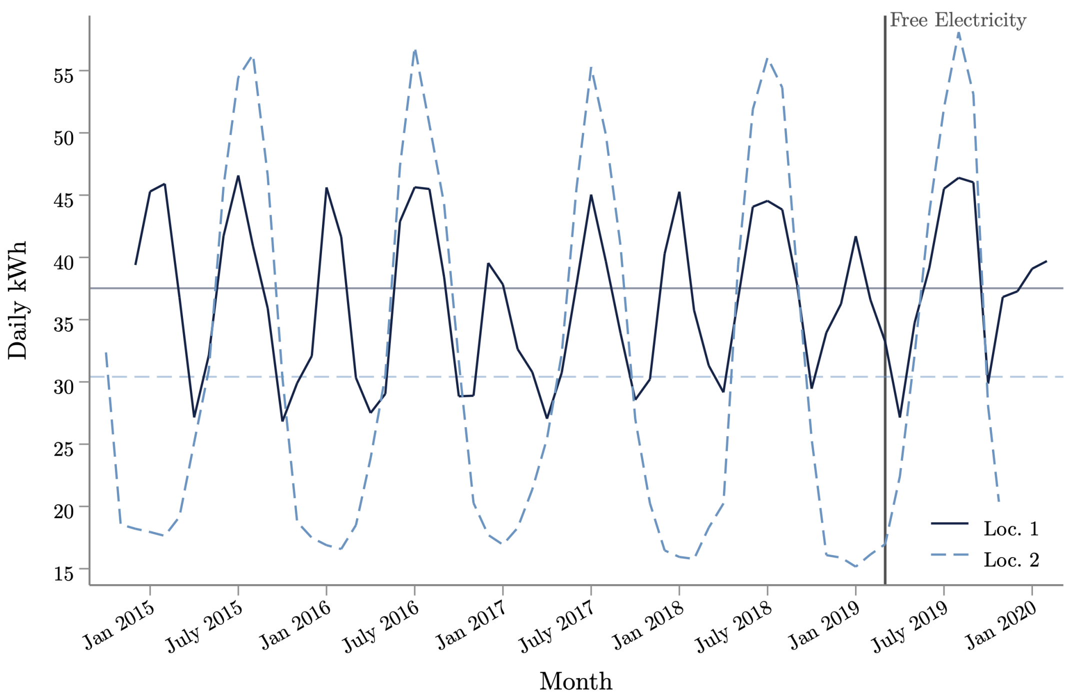

Note: The vertical line indicates when electricity became free. The solid horizontal line represents the free amount of electricity a household living in location one receives each day, on average, and the dashed horizontal line marks the same for the second location.

Table 1 reports basic summary statistics for the households in our sample, pooled and for each location. Mean electricity consumption is 33.36 kWh/day at a rate of $0.07/kWh, a $2.42 value of electricity per day, or approximately $74 per month. As noted above in Figure 3, the two locations have different mean consumption, 37.5 and 30.4 kWh/day, respectively.[5] We also characterize the bottom and top of the free-electricity range; the average range is 3.28 kWh/day. On average, 24.8 percent of the households have a gas connection, 11.7 percent of them had electronic billing, and 3.9 percent of them had autopay.[6]

We collect weather data from NOAA's Climate Data Online service (NOAA 2020). We combine these with monthly billing data by matching each location with monthly heating degree days (HDD), cooling degree days (CDD), and precipitation from its closest weather station. The weather stations are 4.2 and 5.1 miles from their respective bases. HDD and CDD measure how cool or warm a location is by representing how much the average daily temperature deviates from a pre-specified temperature (US Energy Information Administration 2020). We use 65°F as the set point. For example, if the average daily temperature is 40°F, the HDD for that day would be 25°F (or 65-40). The measures reported in Table 1 represent the average daily "degree days" the household experienced. On average, households in our sample experienced more cooling degree days than heating degree days. The correlation between the group average and CDD is 0.63, demonstrating that warmer weather is correlated with higher electricity use.

Military personnel frequently move between bases both within the US and around the world. The average household lived in one location for 17 months in our sample. There were 95 billing groups across both locations with 43 at location one and 52 at location two. The average number of households in a billing group in a month was 150 households and

5. The averages mask heterogeneity between billing groups. We pick four billing groups out of the 95, two from each location, to illustrate this; Figure D.5. The groups' distributions of electricity use differ in both the mean and the variance, reflecting that homes of different sizes and vintages have different consumption patterns.

6. Sexton (2015) demonstrates that automatic bill pay decreases the salience of electricity prices and causes electricity consumption to increase. However, in the context in that paper, almost 80 percent of households enrolled in an automated payment program for their electricity bills, which is much higher than the 3.9 percent in our sample.

Table 1: Summary Statistics

	Pooled Sample	Location 1	Location 2
Electricity Use, kWh/Day	33.36	37.52	30.41
	(18.80)	(19.29)	(17.87)
Price, $/kWh	0.07	0.07	0.08
	(0.01)	(0.00)	(0.00)
Value of Electricity Use, $/Day	2.42	2.55	2.33
	(1.35)	(1.31)	(1.38)
Cooling Degree Days, °F/day	6.27	4.90	7.24
	(6.37)	(5.24)	(6.90)
Heating Degree Days, °F/day	8.17	10.87	6.26
	(8.69)	(10.43)	(6.56)
Charge, $/Day if Charged	0.59	0.58	0.59
	(0.55)	(0.56)	(0.55)
Rebate, $/Day if Rebated	0.39	0.40	0.38
	(0.46)	(0.44)	(0.48)
Bottom, Free-Electricity Range, kWh/Day	31.12	34.96	28.39
	(15.25)	(15.37)	(14.56)
Top, Free-Electricity Range, kWh/Day	34.39	38.64	31.38
	(16.85)	(16.99)	(16.10)
Value of Free-Electricity Range, $/Day	0.24	0.25	0.23
	(0.12)	(0.11)	(0.12)
Share with Gas, %	24.79	56.52	2.26
	(26.78)	(2.25)	(0.46)
Gas Use, Therms/Day	1.35	1.36	1.18
	(1.38)	(1.39)	(0.98)
Share with E-bill, %	11.73	12.41	11.26
	(1.98)	(2.22)	(1.62)
Share with Autopay, %	3.92	4.73	3.35
	(1.12)	(1.10)	(0.68)
Time as Resident, Months	16.53	16.83	16.32
	(11.22)	(11.37)	(11.11)
Billing Group Size	150.32	148.23	151.80
	(119.22)	(106.63)	(127.39)
Billing Groups	95	43	52
Homes	9,307	3,981	5,326
Residents	31,381	12,716	18,665
Observations, Nonlinear	370,029	151,994	218,035
Observations, Free	85,322	37,080	48,242
Observations, Total	455,351	189,074	266,277

Note: Each column reports the average for the pooled sample, and by location. Standard deviations in parentheses.

the median was 105 households. For privacy reasons, the addresses of the other households in a billing group were not disclosed. Because the households were not aware of the other members of their billing group, we have little concern that households coordinated amongst themselves to increase their monthly electricity allocations.

5 Do Households Respond to Nonlinearities in their Pricing Schedule?

5.1 A Natural Experiment with Free Electricity

To determine whether households respond to the military's pricing schedule, we must understand what they would have done in the absence of the nonlinearities in the budget constraint. We use three different approaches to predict that counterfactual electricity use and understand households' responses to their electricity pricing schedule. Two of the three approaches rely on the fact that all electricity for families in on-base housing became free in March 2019. After the change, households continued receiving information about their monthly allocation but neither paid nor received rebates for any electricity consumed thereafter.

As a result of this abrupt and unexpected change in prices, we observe household electricity use under two pricing schedules: nonlinear pricing due to a free-electricity range and completely free electricity. The best possible counterfactual would be observed electricity use under a constant price to estimate the counterfactual distribution of electricity consumption. But, we do observe households with a marginal price of zero. We use observed electricity consumption under that flat price of zero to predict what households would have done if the did not have the free-electricity range up to March 2019. Observing electricity use under both pricing schedules allows us to predict counterfactual electricity use when there is no nonlinearity and compare those outcomes to the observed distribution where there is one.

In Section 5.2, we quantify bunching using household-level panel data on electricity use. We estimate changes in electricity consumption using a simple differences framework. Then, we use those estimates to predict counterfactual electricity consumption without the free-electricity range and calculate the amount of bunching using that approach.

5.2 Quantifying Bunching Using Household-Level Data

We estimate how households with varying electricity use responded to the incentives created by the free-electricity range. To determine whether households increased (or decreased) their electricity use as a result of the nonlinearity, we use observed consumption when electricity was free to split households into 10 percent bins relative to the group average (10-20 percent of the average, 20-30 percent of the average, etc.). We use this characterization because the free electricity period is our control period without a nonlinear pricing schedule. We characterize households into bins to estimate the average response of the typical household in one of those bins. For example, a household with average electricity use of 75 percent of the group average is assigned to the 70–80 percent bin. These estimates allow us to predict each household's counterfactual electricity consumption.

The bins are defined by the households' *average* electricity use when electricity was free, so our specification keeps each household's bin constant. But, households infrequently moved between 10 percent bins.The standard deviation in household electricity use was 0.11 for households that used between 70 and 130 percent of the group average.[7]

After we calculate each household's bin, we estimate the average change in electricity use as a result of the nonlinear pricing schedule using:

$$ln(c_{it}) = \beta_0 + \beta_1 X_{lt} + \sum_{j=1}^{j=20} \beta_{2j}[\mathrm{Bin}_j * \mathrm{N}_t] + \mathrm{N}_t + \gamma_i + \tau_t + \theta_t + \epsilon_{it} \tag{2}$$

where $ln(c_{it})$ is the natural log of household i's average daily electricity use in month-of-

7. The standard deviation for households across the whole pricing schedule is 0.19, suggesting that households usually did not deviate more than two bins from their average post period bin.

sample t, X_{lt} includes HDD and CDD, and Bin_j is the average electricity use bin for household i. We define the treatment variable, N_t, as moving from the post-March-2019 sample to the pre-March-2019 sample because we want the coefficients to reflect the change caused by the free-electricity range, not what happens when it is removed. This approach facilitates interpretation but has no impact on the estimation.[8] γ_i is a household-level fixed effect, τ_t is a month-of-year fixed effect, and θ_t is a year fixed effect. Errors are clustered at the billing group level because it is possible that errors are correlated between billing groups (Cameron, Gelbach, and Miller 2011).

The parameters of interest, the β_{2j}s, estimate the average percent change in electricity use for the Bin_j households. Controlling for weather and the included fixed effects, this parameter captures how households changed their electricity consumption in response to the nonlinear pricing schedule. Note that because all electricity became free at the same time for all households, the month-of-year and year fixed effects control for the fact that average electricity use was higher when all electricity was free. Thus, the parameter estimate of β_{2j}s captures the average change in electricity use for households whose incentives changed at different parts in the pricing schedule. In other words, if a household typically used between 110-120 percent of their allocation was electricity was free, under the nonlinear pricing schedule they would have to pay for the kWh in excess of the allocation. The estimate for those households differs from households who typically used less than the group average, even when electricity was free.

Causal identification requires that conditional on the fixed effects and weather controls, there are no time-varying trends in household electricity consumption that change simultaneously with the change in prices. Household fixed effects absorb time-invariant characteristics at the household level, such as consumption type or other household-specific behavioral tendencies. In support of that identification assumption, Figure 4 shows electricity use at the two locations under the two pricing regimes: the nonlinear regime, the dark blue densities,

8. If the treatment indicator were changed to reflect the removal of prices, the estimates would be the same but with the opposite sign.

and the free regime, the orange densities. Strikingly, the distributions at both locations shifted in the same general way when all electricity became free. Note that these figures display the distribution of daily kWh, rather than electricity use relative to the free allocation. Thus, we would not expect to be able to see bunching in this figure because the allocation changes each month.

Figure 4: Distribution of Monthly Electricity Use by Location

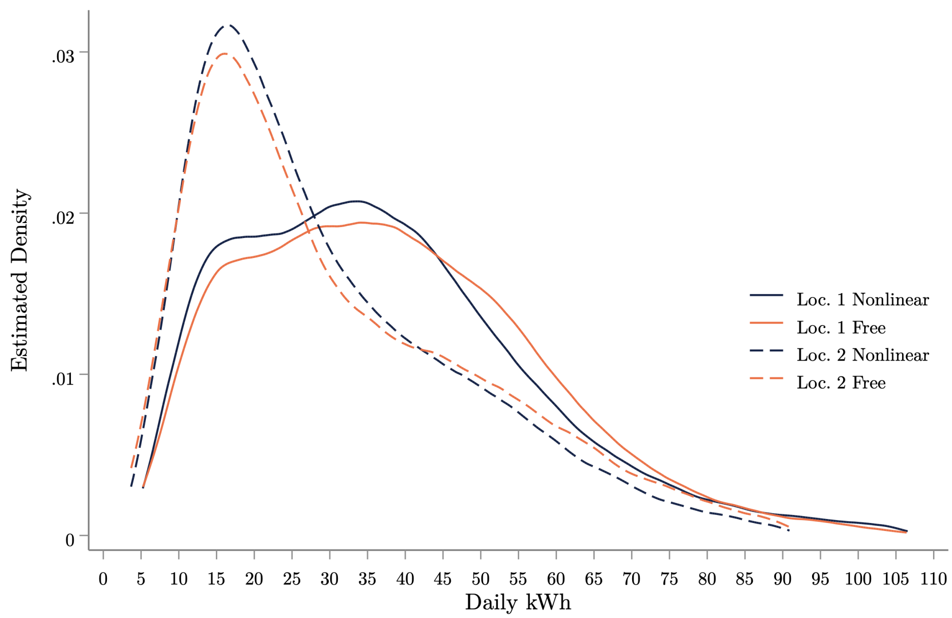

Note: This figure shows the distribution of daily electricity use at both locations under the two alternative pricing schedules. The distributions represent the data across months.

Figure 5 show the β_{2j}s from Equation (2).[9] All estimates are relative to the 7th bin (households with average use of 70 to 80 percent of the group average). Bins one through four show statistically significant increases in electricity use. Whereas, bins eight through twenty show statistically significant decreases in electricity use, ranging from 1.8 to 8.3 percent. The figure shows that pattern: high users decreased their electricity use under nonlinear prices relative to free electricity.[10]

9. Appendix Table D.8 shows the point estimates displayed in Figure 5.

10. We also estimate versions of Equation (2) separately for each season and an alternative definition of the bin sizes. The outcomes for the point estimates can be seen in Appendix Figure D.6. The pattern in each Figure is similar to the pattern in Figure 5.

Figure 5: Changes in Electricity Use

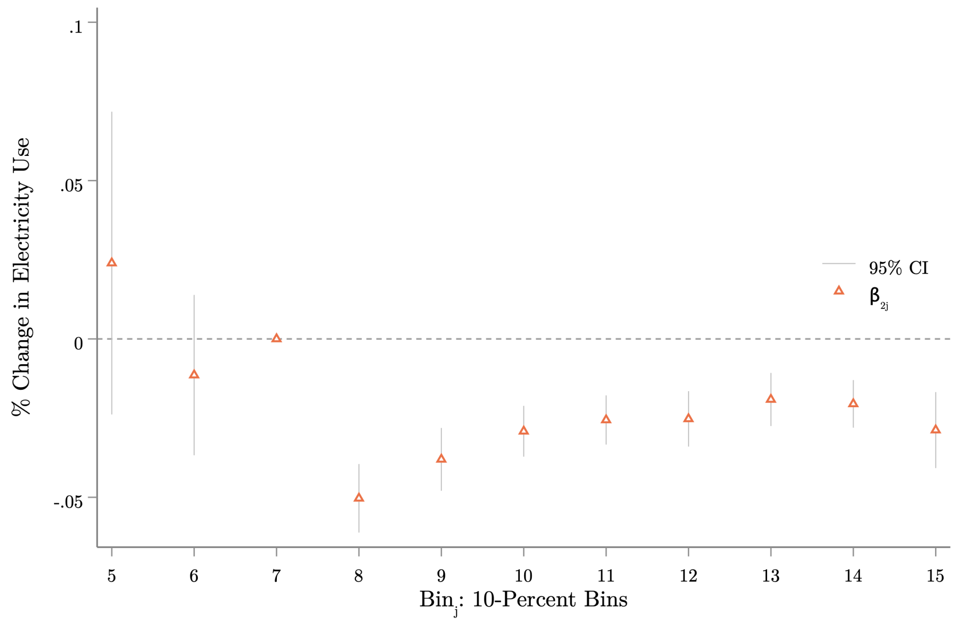

Note: Each Bin_j represents average electricity use in 10-percent bins when electricity is free. Note that Bin_7 is the omitted category. The point estimates β_{2j} are shown by the triangles and represent the log change in electricity use for each bin-type.

Next, we predict each household's counterfactual electricity consumption as though the free-electricity range did not exist. In practice, we predict c_{it} using Equation (2) while omitting the β_{2j}s. Chetty et al. (2011) and Saez (2010) use this approach to predict the counterfactual distribution of electricity use by omitting the estimates from the pre-specified bunching range. This predicted counterfactual electricity consumption represents what households *would have done* if there was no nonlinearity in their pricing schedule.

Figure 6 shows the outcome of predicting the counterfactual. The figure shows versions of the distribution in one percent bins of the group average. Each point on the figure shows the average share of households that fell within that bin. The circles show the shares for the observed data and the triangles show the shares for the predicted data resulting from Equation (2). The figure shows missing mass and excess mass in the observed distribution versus the predicted distribution. The figure provides suggestive evidence that there is more mass in the distribution in the free-electricity range.

Figure 6: Individual Approach, Observed and Predicted Distributions

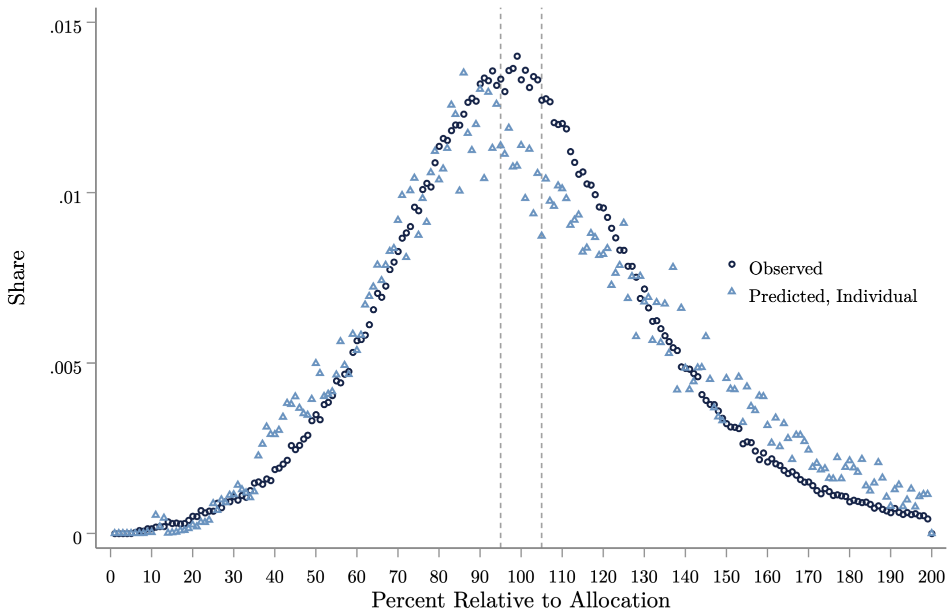

Note: This figure shows the average share of households that fall into one percent bins relative to the group allocation. The circles show the observed data and the triangles show the predicted data resulting from Equation (2).

While Figure 6 appears to demonstrate bunching, are the shares in each bin statistically different from one another? We complete t-tests to determine whether the observed distribution shows missing and excess mass relative to the predicted distribution. The t-tests allow us to determine the range containing missing mass and the range containing excess mass based on where the differences are negative or positive and statistically significant.

Figure 7 shows the results for each t-test for each percent bin. The orange triangles show the percentage point difference between the share of households in that bin in the observed data and the share of households in that bin in the counterfactual data. The grey lines represent the 95 percent confidence interval on for each t-test. If the difference estimate is positive, the observed data has a higher share of households in that bin than the counterfactual data. A positive estimate indicates that there is excess mass, or bunching, in that bin. Similarly, if the difference estimate is negative, the observed distribution has a lower share of households in that bin than the counterfactual distribution. A negative estimate indicates that there is missing mass in that bin.

Figure 7: Individual Approach, T-Test Estimates

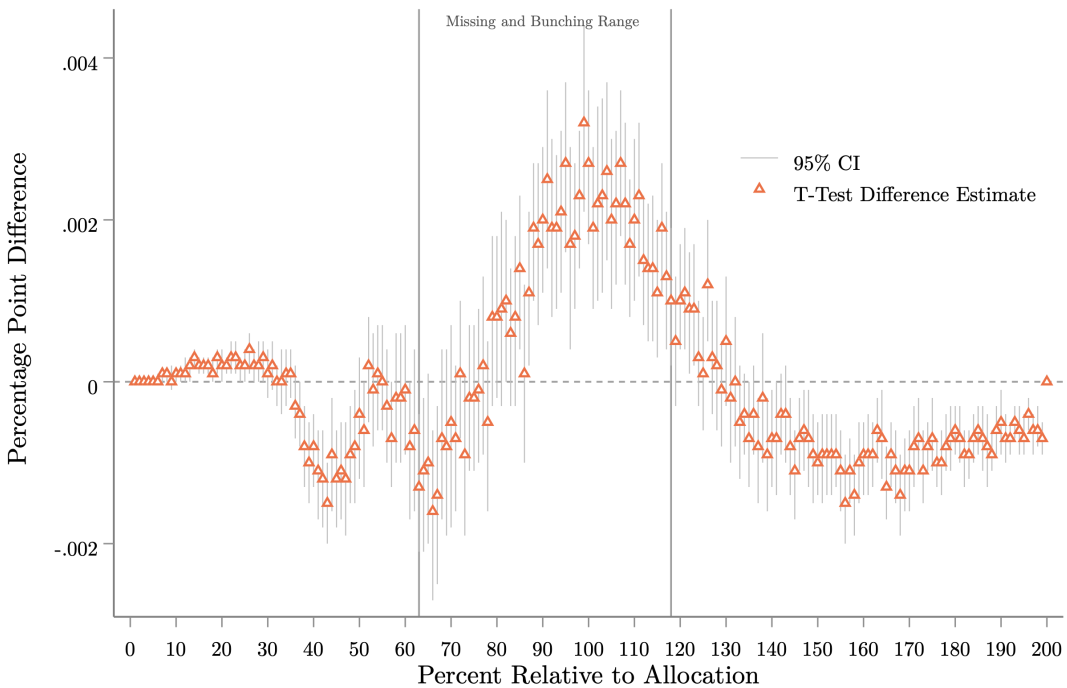

Note: This figure shows the result of 200 t-tests for the difference in the share of households falling into a percent bin in the observed data versus the predicted data. A negative and statistically significant estimate suggests that bin is in the missing range and a positive and statistically significant estimate suggests that bin is in the bunching range.

Figure 7 shows that the percentage point difference between the observed and the counterfactual is negative and statistically different for bins 63 to 67 percent of the group average. This is the "missing" range predicted by the individual-level specification. These missing households jump to the bunching range, which occurs from bin 87 to 118 of the group average, or the bunching range predicted by this specification.[11] This figure and the corresponding t-tests provide further evidence that the mass of missing households shifts from the left of the free-electricity range into the free-electricity range as our theory predicted.[12]

11. We calculate this missing (bunching) range based on where the estimates are negative (positive) and statistically different from 0 at the 5 percent level.

12. The four panels of Appendix Figure D.7 show the results of the t-test for the seasonal specifications shown in Appendix Figure D.6. These Figures show more imprecise estimates than the main results but do show a similar pattern with missing households coming from the left and bunching around the free-electricity range.

Using these predicted missing and bunching ranges, we can aggregate the share of households falling into each range in each billing group in each month to characterize the proportion of missing and bunching households. We find the percent difference for the share of households in the bunching bins in the observed distribution and the counterfactual distribution. The bunching and the missing range are set based on the bins with statistically significant differences between the observed and counterfactual distributions. The average estimated amount of bunching is 27.61 percent more households in the bunching region in the observed than the counterfactual distribution. The average estimated amount of missing mass is -27.10 percent fewer households in the missing region in the observed versus the counterfactual.

This estimated bunching is striking because previous literature using administrative tax data snows almost no evidence of bunching at kink points in the marginal tax schedule. Rather, these papers typically find that only households with self-reported income are able to locate at the kink in the pricing schedule (Saez 2010; Chetty et al. 2011; Kucko, Rinz, and Solow 2018; Mortenson and Whitten 2018).

These results demonstrate that households bunch in response to discontinuities in the pricing schedule and that households with relatively low amounts of electricity use, 63 to 67 percent of their group's average, may be incentivized to increase their electricity use to take advantage of the free-electricity range.

In addition to leveraging the individual panel data, we explore two alternative empirical approaches. The first uses aggregated data to explore how the shape of the distribution changes overall. This approach leverages the unexpected change in the electricity pricing schedule to estimate the share of the bunching and the missing mass in the distribution overall, rather than at the individual level. This approach and its corresponding results are described in Appendix Section B. The second uses a combination of the standard bunching approaches in Chetty et al. (2011) and Kleven and Waseem (2013). This approach counts the number of households that ever fall in a particular bin relative to the kink or notch point.

Then, the approach fits a predicted polynomial through the data surrounding the bunching range, but not including the bunching range. This approach and the corresponding results can be seen in Appendix Section C.

5.3 Confirming Bunching Using a Regression Discontinuity Design at Billing Thresholds

Households at both locations were sent monthly bills but only received a check or had to make a payment if their cumulative balance exceeded \$15. This \$15 threshold for payments and rebates allows us to implement a sharp regression discontinuity approach. Miller and Alberini (2016) find that income elasticities of electricity demand vary from 0.18 to 0.06 and therefore if households behave rationally, a \$15 payment or rebate should have little to no effect on their electricity use.

A regression discontinuity approach allows us to estimate differences in daily average electricity consumption for households just above versus just below the \$15 payment and rebate threshold. Figure 8 shows average daily electricity consumption binned in 50 cent intervals by last month's balance. The left panel shows households with cumulative balances just above and just below the \$15 threshold for receiving a rebate check. The figure shows a sharp discontinuity at \$15 where households that received the rebate use, on average, around 2.2 kWh less electricity per day than households that were just above that \$15 threshold. The right panel shows the average daily electricity use for households nearby the \$15 payment threshold. There is no visually discernible discontinuity in electricity use for households that were just above that \$15 threshold versus just below.

To estimate whether there is a discontinuity in daily electricity consumption for households just above versus just below the rebate or payment threshold, we estimate:

$$ln(c_{it}) = \beta_0 + \beta_1 \tilde{x_{it}} + \beta_2 D_{it} + \beta_3 \left[\tilde{x_{it}} * D_{it}\right] + \beta_4 X_{it} + \epsilon_{it} \tag{3}$$

where $ln(c_{it})$ is the natural log of household i's average daily electricity consumption in month t. $\tilde{x_{it}}$ is the difference between the household's billing statement and the \$15 rebate or payment, for the rebate it is Act. Statement$+15$ and for the payment it is Act. Statement-15. In both cases, $\tilde{x_{it}} = 0$ at exactly the threshold where households either received a rebate or had to make a payment. D_{it} is a dummy variable equal to one if the household's account statement was less than \$15, $\mathbb{1}\left[Act.Statement \leq -15\right]$, and they were owed a rebate, or greater than \$15 and they had to make a payment, $\mathbb{1}\left[Act.Statement \geq 15\right]$. $[\tilde{x_{it}} * D_{it}]$ represents the interaction between the difference from the threshold multiplied by receiving the rebate (making a payment). This term allows us to flexibly control for differing trends in electricity use on either side of the threshold. We also include a polynomial for this term to allow for nonlinearity in some specifications. X_{it} represents other controls such as heating degree days, cooling degree days, and lags of the household's account balance. These lags control for the fact that the rebate or payment is determined as a function of the household's previous consumption if their previous month's statement was not beyond the \$15 threshold. And ϵ_{it} represents an error term with mean zero.

The parameter of interest is β_2, which represents the average percent difference in daily electricity use for households that just received a rebate (had to make a payment) versus those that did not. To estimate Equation (3) for receiving the rebate, we limit the sample to households with account statements within \$10 of the rebate threshold, from \$-25 to \$-5. By limiting the sample in this way, we are limiting our analysis to households nearby the rebate threshold. To interpret estimates from Equation (3) causally, it must be the case that households just above the threshold versus households just below the threshold are similar and that receiving the treatment is as good as random (Cattaneo, Idrobo, and Titiunik 2019).

Table 2 shows the results for households that received a rebate. These estimates demonstrate that households who received a rebate used 3.7 percent to 4.3 percent less average

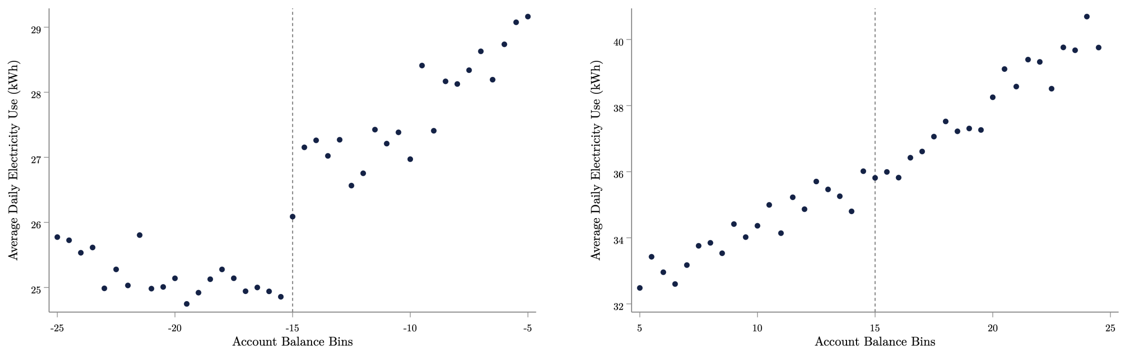

Figure 8: Average Daily Electricity Use by Last Month's Statement

daily electricity than otherwise similar households that did not receive a rebate.[13] Miller and Alberini (2016) find that income elasticities of electricity demand vary from 0.18 to 0.06 using a national sample of households. Our finding here of a 4.2 percent decrease in electricity use as a result of receiving a rebate check is the opposite and larger in magnitude than what we would find if this estimate represented only an income effect.

For example, suppose that the household earned around $60,000 per year or $5,000 per month. A $15 rebate check represents only a 0.3 percent increase in income. If the household's income elasticity of electricity demand was 0.018, then the percent change in electricity use would have been 0.035 percent.[14] When the household receives a rebate, the income effect predicts that they should increase their electricity consumption, indication that our finding that households decrease their electricity consumption after receiving a rebate is a behavioral response to the rebate check.

To BB: make a figure in event time where we plot average electricity use leading up to getting the subsidy. Does their electricity use go up or down on average after getting the subsidy.

13. We estimate several robustness tests for these results. The first, shown in Appendix Table D.9 shows the same results by season– winter versus summer. Those results show a stronger effect in winter than in summer, suggesting that households are better able to target the $15 rebate in the winter than in the summer. Appendix Table D.10 shows a robustness test controlling for previous account balances in months t-2 and t-3 because the rebate is a function of the household's cumulative balance. This table shows qualitatively similar findings to Table 2.

14. The income elasticity of electricity demand is -.118 which should equal the percent change in demand over the percent change in income: $0.118 = \frac{Pct.\Delta D}{0.003}$.

Table 2: RDD: Rebate Feebate Electricity

	(1) ln(daily kWh)	(2) ln(daily kWh)	(3) ln(daily kWh)
Act. Bal. + 15 ($)	0.011*	0.010*	0.012
	(0.001)	(0.004)	(0.010)
1[Act. Bal. $\leq$ -15]	−0.037*	−0.042*	−0.043*
	(0.007)	(0.010)	(0.014)
1[Act. Bal. $\leq$ -15] * (Act. Bal + 15)	−0.002*	−0.004	−0.008
	(0.001)	(0.004)	(0.013)
HDD	0.002*	0.002*	0.002*
	(0.000)	(0.000)	(0.000)
CDD	0.002*	0.002*	0.002*
	(0.000)	(0.000)	(0.000)
(Act. Bal. + 15)^2 ($)		0.000	0.000
		(0.000)	(0.002)
1[Act. Bal. $\leq$ -15] * (Act. Bal + 15)^2		0.000	−0.001
		(0.000)	(0.003)
(Act. Bal. + 15)^3 ($)			0.000
			(0.000)
1[Act. Bal. $\leq$ -15] * (Act. Bal + 15)^3			0.000
			(0.000)
Constant	2.187*	2.189*	2.187*
	(0.079)	(0.077)	(0.076)
Month-Year FE	Y	Y	Y
N	100,521	100,521	100,521

* p<0.05. The standard errors reported in parenthesis have been clustered at the billing group level.

Columns 1, 2, 3, and 4 show the results from three alternative specifications.

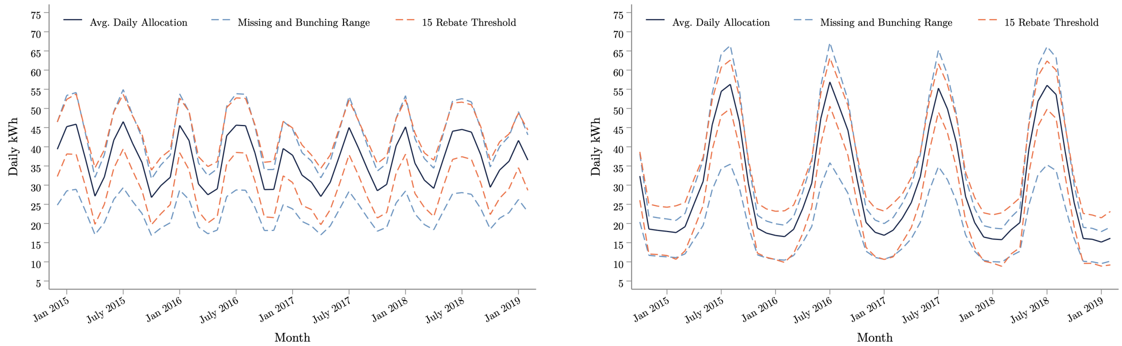

Figure 9: Average Daily Electricity Use, Bunching Range, and $15 Threshold

Table 3 shows the same estimates from Equation 3 for households nearby the payment threshold. The table demonstrates that there is no statistically significant effect of being just above the $15 threshold to make a payment. The stark difference between the findings for the rebate versus the payment suggest that households are more responsive to the rebate than they are to the payment.

Figure 9 shows the average daily allocation, the average upper and lower bounds of the bunching range, and the number of kWh a household would have to use (or conserve) to pay or recieve the rebate check. The left panel shows those ranges for location 1 and the right panel shows those ranges for location 2.

At location 1, the upper end of the bunching range almost perfectly overlaps with the $15 threshold for a payment. In other words, if a household was bunching, they would have to make the $15 payment each month if they were a buncher. At location 2, the range almost perfectly overlaps with the $15 rebate range in the winter months.

6 Seasonality and Persistence

We find that households bunch in response to nonlinearities in their residential electricity pricing schedule. But, what is causing them to bunch? One possible mechanism is seasonality: the free-electricity range of is larger in summer than it is in winter because group averages are higher, or in winter there is less variance in the group average, so households

Table 3: RDD: Charge Feebate Electricity

	(1) ln(daily kWh)	(2) ln(daily kWh)	(3) ln(daily kWh)
Act. Bal. - 15 ($)	0.009*	0.007*	0.008
	(0.001)	(0.002)	(0.006)
1[Act. Bal. ≥ 15]	−0.002	−0.003	−0.007
	(0.006)	(0.009)	(0.011)
1[Act. Bal. ≥ 15] * (Act. Bal - 15)	0.002	0.007	0.009
	(0.001)	(0.004)	(0.010)
HDD	0.002*	0.002*	0.002*
	(0.000)	(0.000)	(0.000)
CDD	0.001*	0.001*	0.001*
	(0.000)	(0.000)	(0.000)
(Act. Bal. - 15)^2 ($)		0.000	0.000
		(0.000)	(0.002)
1[Act. Bal. ≥ 15] * (Act. Bal - 15)^2		0.000	−0.001
		(0.000)	(0.003)
(Act. Bal. - 15)^3 ($)			0.000
			(0.000)
1[Act. Bal. ≥ 15] * (Act. Bal - 15)^3			0.000
			(0.000)
Constant	2.809*	2.805*	2.806*
	(0.063)	(0.063)	(0.065)
Month-Year FE	Y	Y	Y
Year FE	Y	Y	Y
N	69,559	69,559	69,559

* p<0.05. The standard errors reported in parenthesis have been clustered at the billing group level.

Columns 1, 2, 3, and 4 show the results from three alternative specifications.

may better able to target the bunching range.

To test whether seasonality plays an important role in household bunching in this context, we estimate the probability that a household bunches splitting the sample into the four seasons. How much more likely are households to bunch when they face a nonlinear pricing schedule? We measure the increase in probability that a household bunches using the predicted counterfactual electricity use from Section 5.2. The observed data act as the "treatment" sample and the predicted data act as the "control" sample for the analysis in this section. The predicted data represent each household's counterfactual electricity use under a flat marginal price. By comparing bunching under the observed data versus bunching under the counterfactual data, we are able to predict how much more likely a household was to bunch when they faced nonlinear incentives in their pricing schedule. To estimate that probability, we use a simple linear probability model where we compare bunching in the observed distribution to bunching in the counterfactual using:

$$\mathbb{1}[\mathrm{B}_{it}] = \beta_0 + \beta_2[\mathrm{O}_i * \mathrm{N}_t] + \beta_3\mathrm{O}_i + \beta_4\mathrm{N}_t + \epsilon_{it} \tag{4}$$

where $\mathbb{1}[\mathrm{B}_{it}]$ is an indicator variable equal to one if household i fell into the bunching region (defined as 88 percent to 118 percent of the group average as found in Section 5.2), O_i is an indicator equal to one if the data represent an observed data point for that household, N_t is an indicator variable equal to one if the month occurred when the nonlinear pricing schedule was in place, and ϵ_{it} is an error term with mean zero.

The estimate of β_2 is our outcome of interest and measures how much more (or less) likely an individual was to bunch under the nonlinear pricing schedule.[15] Table 4 shows the estimates from Equation 4. Column 1 shows that the average household is 1.6 percent more likely to bunch in the observed versus the predicted distribution. Columns 2 through 5 show

15. Recently, researchers studying energy efficiency use machine learning to predict counterfactual electricity use (Burlig et al. 2020 and Christensen et al. 2020). Our approach uses a simple differencing approach rather than machine learning, but we are comparing outcomes between the observed data and a predicted counterfactual in the spirit of that recent literature.

seasonal estimates for winter, spring, summer, and fall respectively. Households are 3.5 percent more likely to bunch in winter months and 2.8 percent more in summer months. These seasonal estimates demonstrate that when group averages are higher, households are better able to target landing in the bunching region and benefiting from the free-electricity range. These findings are consistent with the analysis in Section 5.3 that finds that households are more likely to respond to the \$15 rebate incentive in the winter months.

Table 4: Change in Probability of Bunching

	(1) $\mathbb{1}\,[B_{it}]$	(2) $\mathbb{1}\,[B_{it}]$	(3) $\mathbb{1}\,[B_{it}]$	(4) $\mathbb{1}\,[B_{it}]$	(5) $\mathbb{1}\,[B_{it}]$
$\mathbb{1}\,[O_{it} * N_t]$	0.016*	0.035*	0.004	0.028*	−0.005
	(0.004)	(0.007)	(0.005)	(0.006)	(0.006)
$\mathbb{1}\,[O_{it}]$	0.048*	0.021	0.010*	0.104*	0.057*
	(0.007)	(0.017)	(0.005)	(0.005)	(0.007)
$\mathbb{1}\,[N_t]$	−0.009*	−0.013*	0.010	−0.018*	−0.009
	(0.004)	(0.006)	(0.005)	(0.006)	(0.006)
Constant	0.339*	0.337*	0.344*	0.329*	0.345*
	(0.005)	(0.012)	(0.007)	(0.006)	(0.006)
N	385,758	111,356	82,348	89,334	102,720

* $p<0.05$. The standard errors reported in parenthesis have been clustered at the billing group level. Column 1 shows estimates for the full sample. Column 2 shows estimates for winter. Column 3 shows estimates for spring. Column 4 shows estimates for summer. And Column 5 shows estimates for fall.

We can also estimate the additional effect of having been a buncher in previous months. Figure 10 shows the share of household-months that did not bunch in the previous month, did bunch in the previous month, did bunch in the previous two months, and so on up to six months. We refer to this as the length of a household's "bunching streak."

To understand whether households that previously bunched are more likely to bunch in the current month, we estimate a triple differences model where we include interactions for bunching in previous months:

Figure 10: Bunching Streaks in Months

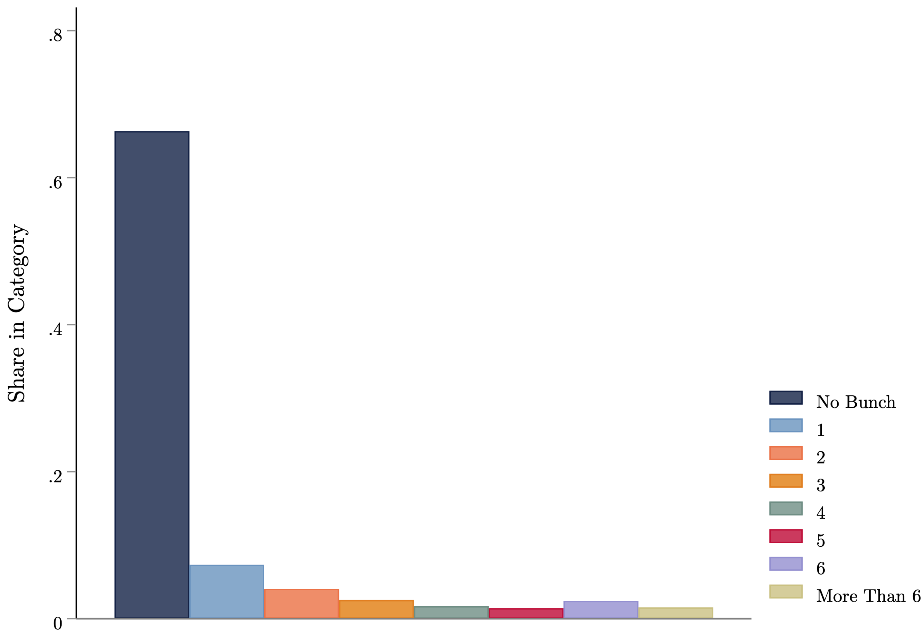

Note: This figure shows the average share of households that had a bunching streak of length n in months.

$$
\begin{aligned}
\mathbb{1}[\mathrm{B}_{it}] = &\beta_0 + \beta_2[\mathrm{O}_i * \mathrm{N}_t] + \beta_3[\mathrm{O}_i * \mathrm{B}_{it-n}] + \\
&\beta_4[\mathrm{N}_t * \mathrm{B}_{it-n}] + \beta_5[\mathrm{O}_i * \mathrm{N}_t * \mathrm{B}_{it-n}] + \beta_6\mathrm{O}_i + \beta_7\mathrm{N}_t + \beta_7[\mathrm{B}_{it-n}] + \epsilon_{it}
\end{aligned} \tag{5}
$$

Where all terms are the same as in Equation (4) apart from Bunch_{it-n}, which represents how many consecutive months the household has bunched. β_5 represents the average change in the probability of bunching for observed data (versus the counterfactual) for households with a bunching streak of n months.

Results from Equation (5) are reported in Appendix Table (D.11). The first column of the table shows that on average, households are 1.5 percent more likely to bunch, but if that household bunched last month, they are 2.8 percent more likely to bunch. Columns 2 and 3 show a similar pattern, households that had two or three month bunching streaks are more likely to bunch than households that did not. However, for individuals with bunching streaks of four to six months or more, there is no marginal increase in the probability of bunching relative to the full sample of households.

These findings suggest that households that were bunchers in the previous month, two months, or three months are more likely to bunch in the current month. This finding indicates that households that have previously bunched are more likely to bunch. Previous work by Mortenson and Whitten (2018) show that households that were "refund maximizers" (bunchers) in the previous tax year are more likely to be bunchers in this year. However, in our sample, we have the ability to test for bunching streaks and observe that the households that bunched in the one of or more of the previous three months are the most likely to bunch.

7 Optimization Frictions, Bunching, and Elasticities of Electricity Demand

Section 3 sets up a conceptual framework that allows us to estimate elasticities of electricity demand using our empirical context. Following Kleven and Waseem (2013) we characterize one measure of optimization frictions faced by the households by using the share of households that fall into the purely dominated region — the free-electricity range. If households are able to perfectly optimize, we should observe no households landing the free-electricity range: they could all be better off by increasing their consumption to exactly $1.05A$, without paying anything for that extra electricity use. Therefore, one measure of optimization frictions is the calculate the share of households that fall in the missing, or dominated, range even though they could be better off by increasing their electricity consumption.

Under our empirical approach, we can characterize optimization frictions by calculating the share of households that fall into the missing range predicted by that specification, defined as a. Then, we can calculate price elasticities of electricity demand using the predicted missing range to characterize the change in electricity use in response to the free-electricity range under the nonlinear pricing. We calculate elasticities using equation (1) presented in Section 3.

To estimate the elasticities, we must calculate the percent change in household electricity demand. We do so using the missing range predicted by the binned and individual specification. The percent change in electricity demand in response to the free-electricity range of the nonlinear pricing schedule, Δc, is the change from c_l to the bunching range. Next, we find the percent change in electricity price by finding the implicit marginal price, Δr, shown in Figure 2. That price equals the change in electricity expenditure for the move from c_l to the bunching range divided by the total change in electricity use. Elasticity estimates under equation (1) presented in Section 3 are overestimates for two reasons: first, the change from the flat marginal price to the implicit marginal price is an underestimate of the true change in price that the household experiences; second, these elasticities do not take into account optimization frictions, but rather they assume that all households in the missing range jump to the bunching range. As in Kleven and Waseem (2013), elasticity estimates adjusted for the optimization frictions can be obtained by multiplying the elasticities by $1 - a$.

Here, we use the missing and bunching range calculated using the t-tests for each bin between the observed and the counterfactual distribution shown in Figure 7 to calculate a range of elasticities. The predicted missing range is from 63 to 67 percent of the group average and the predicted bunching range is from bin 87 to 118. From bins 68 to 86 there is statistical uncertainty in whether those estimates are positive or negative. Therefore, we will calculate elasticities for a wide range of the missing and bunching bins.

Table 5 shows the range of elasticities estimated under different assumptions for where the bunching households jump to. The table shows price elasticities under the assumptions that the households respond to marginal pricies or to average prices. If households respond to marginal prices, the estimates range from XXXX. The estimates range from -1.38 if households jump from the top end of the missing range to the lowest end of the free-electricity range at $1.05A$ to -2.71 if households jump from the lowest end of the missing range to the top of the bunching range at $1.20A$.[16] If these estimates are scaled by the share of households that fall in the dominated region, they are around 1/2 as large, shown in column 6.

Table 5: Individual Specification: Elasticities

(1) Percent Bins Included	(2) Pct. Δ in Price	(3) Pct. Δ in Electricity	(4) Elasticity	(5) Share in Dom. Range [63-104]	(6) Attenuated Elasticity
Implicit Marginal Price					
63 to 87	0.00	27.59		0.47	
63 to 105	-23.81	40.00	-1.68	0.47	-0.89
63 to 118	-17.54	47.50	-2.71	0.47	-1.44
67 to 87	0.00	22.99		0.47	
67 to 105	-26.32	36.19	-1.38	0.47	-0.73
67 to 118	-18.87	44.17	-2.34	0.47	-1.25
Average Price					
63 to 87	0.00	27.59		0.47	
63 to 105	-9.52	40.00	-4.20	0.47	-2.23
63 to 118	-8.33	47.50	-5.70	0.47	-3.03
67 to 87	0.00	22.99		0.47	
67 to 105	-9.52	36.19	-3.80	0.47	-2.02
67 to 118	-8.33	44.17	-5.30	0.47	-2.82

Elasticities in this table are calculated by first finding the percent change in price from the nonlinear marginal price and either the implicit marginal price or the average price. Next, we calculate the percent change in electricity consumption for each bin group and find the elasticity, which is displayed in the final column.

We can characterize the share of households that fall into the dominated region from 63 to 104 percent of the group average in the sample of households we use for the individual level specification. On average, 46.79 percent of households use between 63 and 104 percent of the group average.[17] We use these estimates to calculate the attenuated elasticities shown in column 6 of Table 5. Taking into account optimization frictions, elasticities range from -0.73 to -1.44 if households respond to marginal prices or from -2.02 to -3.03 if households respond to average prices.

The elasticities for our altnerative empirical approach are reported in Appendix Section D

16. If instead households respond to average prices, then the estimates range from -3.80 to -5.70.
17. The standard deviation is 7.58 percent.

8 Discussion and Conclusion

This paper uses a theoretical model and three alternative empirical approaches to determine how households respond to a nonlinearity in their residential electricity pricing schedule. Furthermore, the model predicts that there will be missing mass and excess mass in the distribution and the three empirical approaches find just that. However, we find a range of of estimated ecxess mass depending on the approach. Our estimated bunching ranges from 3.41 percent more households near the top of the free-electricity range under the binned approach to 3.41 percent more households under the individual approach up to 39.64 percent more households under the Chetty et al. (2011) approach.

The binned approach represents our most conservative estimate of bunching, However, it provides us with little insight into exactly *who* is bunching. This approach relies heavily on the identification assumption that the shape distribution of electricity use remains relatively unchanged in aggregate when electricity becomes free.

The individual approach represents our middle estimate of bunching. However, this estimate may suffer from mean reversion in the distribution that we are unable to appropriately control for. Households with high (or low) electricity use one month are likely to drift back toward their average electricity use over time. Mechanically, the household's electricity use allocation is set exactly at the average electricity use of their assigned billing group. This mean reversion could bias us towards finding more bunching than actually exists.

Our largest estimate of bunching occurs under the Chetty et al. (2011) approach. This approach suffers from the fact that the researcher selects the bunching range. We selected a wide variety of ranges guided both by theory and our previous empirical estimates. This approach also does not allow us to account for the fact that the share of missing households should be equal to the share of bunching households.

It is also possible that the excess mass we find in the distribution is not only a result of the nonlinearity in the prices, but also the fact that the households are told how much electricity they use relative to other households in their billing group. We know from Allcott (2011) that when households are given information about how much electricity they use relative to their peers that they respond to that information. There also remains the possibility that households experience loss aversion and attempt to target the free electricity allocation for that behavioral reason. However, future work is needed to disentangle these separate effects from the marginal price effect.

References

Allcott, Hunt. 2011. "Social Norms and Energy Conservation." *Journal of Public Economics* 95 (9-10): 1082–1095.

Allcott, Hunt, and Todd Rogers. 2014. "The Short-Run and Long-Run Effects of Behavioral Interventions: Experimental Evidence from Energy Conservation." *American Economic Review* 104 (10): 3003–3037.

Auten, Gerald, and Robert Carroll. 1999. "The Effect of Income Taxes on Household Income." *Review of Economics and Statistics* 81 (4): 681–693.

Bissell, Kristie L, Robert L Crosslin, and James L Hathaway. 2010. "Military Families and Their Housing Choices." *Logistics Management Institute (LMI),* 108.

Blomquist, Soren, Whitney Newey, Anil Kumar, and Che-Yuan Liang. 2019. *On Bunching and Identification of the Taxable Income Elasticity.*

Borenstein, Severin. 2009. "To What Electricity Price Do Consumers Respond? Residential Demand Elasticity Under Increasing-Block Pricing," 37.

Brolinson, Becka. 2020. "Does Increasing Block Pricing Decrease Energy Use? Evidence from the Residential Electricity Market," Working Paper, gueconwpa~19-19-06.

Burlig, Fiona, Christopher Knittel, David Rapson, Mar Reguant, and Catherine Wolfram. 2020. "Machine Learning from Schools about Energy Efficiency," 37.

Cameron, A. Colin, Jonah B. Gelbach, and Douglas L. Miller. 2011. "Robust Inference With Multiway Clustering." *Journal of Business & Economic Statistics* 29 (2): 238–249.

Cattaneo, Matias D., Nicolas Idrobo, and Rocio Titiunik. 2019. "A Practical Introduction to Regression Discontinuity Designs: Foundations." *arXiv:1911.09511 [econ, stat].*

Chetty, Raj, John N. Friedman, Tore Olsen, and Luigi Pistaferri. 2011. “Adjustment Costs, Firm Responses, and Micro vs. Macro Labor Supply Elasticities: Evidence from Danish Tax Records.” *Quarterly Journal of Economics* 126 (2): 749–804.

Christensen, Peter, Paul Francisco, Erica Myers, and Mateus Souza. 2020. “Decomposing the Wedge Between Projected and Realized Returns in Energy Efficiency Programs,” 84.

Grubb, Michael. 2015. “Consumer Inattention and Bill-Shock Regulation.” *Review of Economic Studies* 82:40.

Hausman, Jerry A. 1982. *Taxes and Labor Supply.* Technical report Working Paper 245. National Bureau of Economic Research.

Ito, Koichiro. 2014. “Do Consumers Respond to Marginal or Average Price? Evidence from Nonlinear Electricity Pricing.” *American Economic Review* 104 (2): 537–563.

Kahneman, Daniel, and Amos Tversky. 1979. “Prospect Theory: An Analysis of Decision under Risk.” *Econometrica* 47 (2): 30.

Kleven, Henrik J., and Mazhar Waseem. 2013. “Using Notches to Uncover Optimization Frictions and Structural Elasticities: Theory and Evidence from Pakistan*.” *The Quarterly Journal of Economics* 128 (2): 669–723.

Kleven, Henrik Jacobsen. 2016. “Bunching.” *Annual Review of Economics* 8 (1): 435–464.

Kucko, Kavan, Kevin Rinz, and Benjamin Solow. 2018. “Labor Market Effects of the Affordable Care Act: Evidence from a Tax Notch.” *SSRN Electronic Journal.*

Lincoln RECP. 2019. *RECP Resident Energy Conservation Program.* https://www.lincolnrecp.com/how-it-works.php.

Miller, Mark, and Anna Alberini. 2016. "Sensitivity of Price Elasticity of Demand to Aggregation, Unobserved Heterogeneity, Price Trends, and Price Endogeneity: Evidence from U.S. Data." *Energy Policy* 97:235–249.

Mortenson, Jacob, and Andrew Whitten. 2018. "Bunching to Maximize Tax Credits: Evidence from Kinks in the U.S. Tax Schedule." *SSRN Electronic Journal.*

NOAA. 2020. *Climate Data Online (CDO) The National Climatic Data Center's (NCDC).* https://www.ncdc.noaa.gov/cdo-web/.

Saez, Emmanuel. 2010. "Do Taxpayers Bunch at Kink Points?" *American Economic Journal: Economic Policy* 2 (3): 180–212.

Sallee, James M. 2014. "Rational Inattention and Energy Efficiency." *The Journal of Law & Economics* 57 (781): 41.

Sexton, Steven. 2015. "Automatic Bill Payment and Salience Effects: Evidence from Electricity Consumption." *Review of Economics and Statistics* 97 (2): 229–241.

Shaffer, Blake. 2020. "Misunderstanding Nonlinear Prices: Evidence from a Natural Experiment on Residential Electricity Demand." *American Economic Journal: Economic Policy* 12 (3): 433–461.

Silicon Valley Power. *Appliance Energy Use Chart.* https://www.siliconvalleypower.com/residents/save-energy/appliance-energy-use-chart.

US Energy Information Administration. 2020. *Degree-Days - U.S. Energy Information Administration (EIA).* https://www.eia.gov/energyexplained/units-and-calculators/degree-days.php.

Wichman, Casey J. 2014. "Perceived Price in Residential Water Demand: Evidence from a Natural Experiment." *Journal of Economic Behavior & Organization* 107:308–323.

A Online Appendix

A.1 Data Preparation Further Descriptive Statistics

We drop observations with daily average electricity use in the bottom[18] and top[19] one percent of the distribution. We drop these outliers to prevent bias in our estimations. We also drop observations when the billing group includes 25 households or fewer.[20]

B Empirical Approach: Aggregated Data

Here rather than focusing on how electricity use changes at the household level, we focus on how the shape of the distribution of electricity use changes in response to the change in the pricing schedule. By using more aggregated data, we are both able to compare our results to other estimates in the existing literature and can address some potential concerns of mean reversion present in the individual approach. But rather than estimating the bunching mass using the cross-section as in the previous literature, we leverage the unexpected change in the electricity pricing schedule to estimate the share of bunching and missing mass.

After estimating the change in individual electricity use as a result of the free-electricity range, we use an alternative approach where we estimate the change in the shape of the distribution. These approaches use aggregated data to estimate overall changes in electricity use, rather than changes at the individual level.

B.1 Graphical Evidence of Bunching

To get a better idea of what the aggregate shifts in consumption look like when prices are set to zero, we transform our individual-level data to binned data. Each bin reports what

18. That value ranges from 3.7 kWh to 5.2 kWh per day depending on the location. The average Energy Star rated refrigerator uses between 35 and 45 kWh in a month (Silicon Valley Power). Low electricity use indicates that the home was likely unoccupied.

19. That value ranges from 90.9 kWh per day to 106.6 kWh per day depending on the location.

20. A handful of billing groups only have four houses. This presents a problem as the variation in the group average is much larger, and less predictable, for the smaller billing groups.

proportion of a billing group in a month falls into a given range of consumption relative to the group average A. The bins span from 0 to 300 percent of the group average in 5 percent intervals. This leaves 60 bins to characterize the distribution, providing 323,640 observations for the specifications in this section.

For example, if there are 100 people in a billing group and 5 of them fall between $0.70A$ and $0.75A$, then that bin would have a value of 0.05 for that billing group in that month.

Using this binned data, we estimate how the shape of the distribution shifted when prices were removed in March 2019 using the following equation:

$$\begin{aligned} S_{bjt} = & \beta_0 + \sum_{b=2}^{b=60} \beta_{1b}(\text{Bin}_b * \text{Nonlinear}_t) + \sum_{b=2}^{b=60} \beta_{2b}(\text{Bin}_b) + \\ & \text{Nonlinear}_t + \beta_3 X_{lt} + \tau_t + \theta_t + \delta_j + \epsilon_{bjt} \end{aligned} \tag{6}$$

All terms are the same as in Equation (2) except S_{bjt}, which measures the share of households in percent bin b in billing group j in month of sample t. Bin_b is a dummy variable equal to one for each 5 percent bin between 0 and 300 percent of the group average. And δ_j is a billing-group fixed effect. And ϵ_{bjt} is the error term, which is clustered at the billing group level to account for correlation between groups (Cameron, Gelbach, and Miller 2011). All estimates are relative to the first bin which includes electricity use between 0 and 5 percent of the group average.[21]

The coefficients of interest are the β_{1b}s, which represent the average percentage point difference in the share of households in bin b when households were charged using the nonlinear schedule versus the free schedule.

The identification assumption for interpreting β_{1b} as the causal effect of the change in pricing regimes on the share of households in a particular bin b is that conditional on fixed effects and weather controls, there is no other time-varying trend that would affect the shape

21. We omit the first bin as there should be no effect of the change in prices on the share of households in the very first bin in the distribution.

of the distribution of electricity consumption. Said differently, conditional on fixed effects and other controls, there cannot be any other change in policy or behavior that might affect the distribution of electricity consumption except the price change.

Figure B.1 shows the estimates of β_{1b} from Equation (6) relative to the first bin (the omitted bin). Each orange triangle represents the percentage point change in the share of households in one bin. The vertical grey lines above and below each triangle show the 95 percent confidence interval. The grey dashed lines indicate the consumption range with free electricity.

The figure provides a few insights. First, there is a statistically significant increase in the share of households using electricity from 100 percent of the group average up to 105 percent of the group average. The estimate for electricity use from 105 up to 110 is positive, but not statistically significant and the estimate for electricity use from 110 up to 115 is positive and statistically significant. These estimates provide further evidence that there are more households locating near the nonlinearity in the pricing schedule when the free-electricity range is in effect. Second, while the estimates are not statistically significant, the estimates from bins 70 to up to 85 are negative, which closely aligns with the "missing range" predicted by the individual specification. Taken together, this graphical evidence suggests that the range of free electricity is causing bunching. In the next section, we take a simple differences approach to estimate the amount of bunching and missing mass and where in the distribution that missing mass comes from.

B.2 Quantifying Bunching Using Binned Data

In this section, we build on the approach in Section B.1 and quantify the bunching and missing households using a multi-step process. First, we estimate a modified version of equation (6) that accounts both for where in the distribution bunching occurs, the bunching range, and where in the distribution that bunching mass comes from, the missing range. We use an iterative procedure adapted from Kleven and Waseem (2013) to determine the

Figure B.1: Average Change in the Share of Households in Each Bin

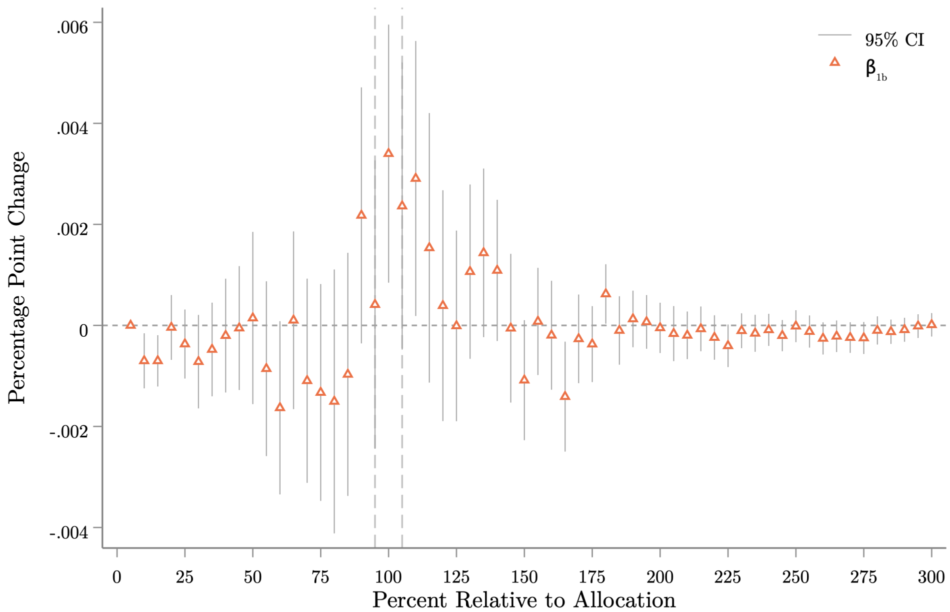

Note: This figure shows the estimates of β_{1b} from Equation (6). A positive and statistically significant estimate demonstrates that a higher share of households landed in that bin when the nonlinear pricing schedule was in effect versus the free pricing schedule.

bunching and missing ranges of the distribution. Second, we use those empirically determined bunching and missing ranges to predict the counterfactual distribution of electricity use in the absence of discontinuities in the pricing schedule similar to the approach in Saez (2010) and Kucko, Rinz, and Solow (2018).

We start by estimating a modified version of equation (6) with the following specification:

$$\begin{aligned} S_{bjt} = &\beta_0 + \beta_1(\text{Miss}_b * \text{Nonlinear}_t) + \beta_2(\text{Bunch}_b * \text{Nonlinear}_t) + \text{Miss}_b + \\ &\text{Bunch}_b + \text{Nonlinear}_t + \beta_5 X_{lt} + \mu_b + \tau_t + \theta_t + \delta_j + \epsilon_{bjt} \end{aligned} \tag{7}$$

where all terms are the same as in equation (6) apart from Miss_b, which is a dummy variable equal to one if bin b falls in the missing range of the distribution, Bunch_b which is a dummy variable equal to one if bin b falls into the bunching range, and μ_b, which represents a percent-bin fixed effect which controls for the differences in the distribution of electricity use common across all billing groups. For β_1 and β_2 to be interpreted as the causal impact of the change in the pricing schedule on the share of households in a missing bin, it must be that conditional on month-of-sample fixed effects, weather, billing-group fixed effects, and bin fixed effects, the distribution of electricity consumption under the free pricing regime is representative of electricity consumption under a flat pricing schedule without any discontinuities in the price.

The bins that fall into the Miss_b and Bunch_b range are not known ex-ante. To determine the best bin allocations, we use an iterative procedure. We estimate (7) to minimize the difference between the predicted number of missing households and bunching households (Kleven and Waseem 2013). We estimate (7) for every missing range and bunching range combination from bin 60 to 140 where at least two bins are included. This procedure leads us to 364 possible combinations of missing ranges and bunching ranges.

This iterative process allows us to determine the missing range and the bunching range empirically. The missing range and bunching range occur when the difference between the predicted missing mass and the predicted bunching mass is minimized.

The results from the top ten specifications that minimize the absolute value of the difference between the predicted and bunching mass are shown in Table B.1. The specifications are ordered from the smallest predicted difference to the largest predicted difference. The coefficients for the missing range, β_1, and bunching range β_2 are shown in columns 2 and 3 respectively. These estimates show the average percentage point change in the share of households falling into a bin in the missing or bunching range. To find the difference between the predicted missing mass and predicted bunching mass, we multiply the coefficients in columns 2 and 3 of Table B.1 by the number of bins that were used in the estimate (columns 4 and 5). Column 6 reports that difference in the predicted missing mass and the predicted bunching mass.

We use these specifications to determine which bins should be included in the missing range and the bunching range. While the estimates in the top three rows represent the three smallest estimated differences between the predicted missing mass and the predicted excess mass, the estimates for both β_1 and β_2 are not statistically distinguishable from zero. The estimates included in the 4^{th} row, the light grey highlighted row, are both statistically significant and correspond to the missing range from 65 percent below the average to (but not including) 90 percent. The bunching bins in the specification are from 90 percent of average to (but not including) 110 percent. The difference between the predicted missing mass and bunching mass is 0.000629.

Table B.1: Missing and Bunching Bins to Minimize Estimated Missing and Excess Mass

Miss and Bunch Bins Included	β_1	β_2	# Miss Bins	# Bunch Bins	Difference
Miss (60 to 75) Bunch (80 to 105)	-0.001189*	0.000825	4	6	0.000194
	(0.000525)	(0.000603)			
Miss (70 to 95) Bunch (100 to 105)	-0.000743	0.002335*	6	2	0.000212
	(0.000548)	(0.000909)			
Miss (65 to 95) Bunch (100 to 105)	-0.000611	0.002292*	7	2	0.000307
	(0.000507)	(0.000909)			
Miss (65 to 85) Bunch (90 to 105)	-0.001283*	0.001761*	5	4	0.000629
	(0.000610)	(0.000739)			
Miss (60 to 80) Bunch (85 to 105)	-0.001334*	0.001199	5	5	-0.000675
	(0.000553)	(0.000658)			
Miss (60 to 90) Bunch (95 to 105)	-0.000903	0.002354*	7	3	0.000741
	(0.000461)	(0.000830)			
Miss (70 to 85) Bunch (90 to 105)	-0.001652*	0.001838*	4	4	0.000744
	(0.000697)	(0.000746)			
Miss (60 to 70) Bunch (75 to 105)	-0.001055	0.000576	3	7	0.000867
	(0.000585)	(0.000548)			
Miss (65 to 80) Bunch (85 to 105)	-0.001232	0.001212	4	5	0.001132
	(0.000660)	(0.000660)			
Miss (60 to 85) Bunch (90 to 105)	-0.001365*	0.001747*	6	4	-0.001202
	(0.000522)	(0.000734)			

Note: * significant at 0.05 or better. Standard errors reported in parenthesis, clustered at billing group level.

Appendix Table D.7 shows the point estimates from Equation (7) for the specifications in Table B.1 where both β_1 and β_2 are statistically significant, which includes rows four, seven, and ten. This table shows that on average, there were 0.001 percentage points fewer households on average in bins in the missing range. It also shows that typically, that there were 0.002 percentage points more households on average in bins in the bunching range. These results demonstrate that the nonlinear pricing schedule incentivizes households using less electricity than the lower bound of the free-electricity range to increase their electricity use.

By iterating over alternative ranges for the missing and bunching bins, we empirically determine which bins should be included in the missing range and the bunching range. Then, we use the estimates from Equation (7) for the bins that minimize that estimate difference to predict a counterfactual distribution of electricity consumption. In practice, this is done by predicting S_{bjt} values for each bin, billing group, and month of sample while omitting the β_1 and β_2 coefficients. This counterfactual represents the distribution of electricity consumption under a pricing regime with a flat marginal price. Figure B.2 shows the observed and the predicted distributio resulting from Equation (7)

Figure B.2: Binned Approach, Observed and Predicted Distributions

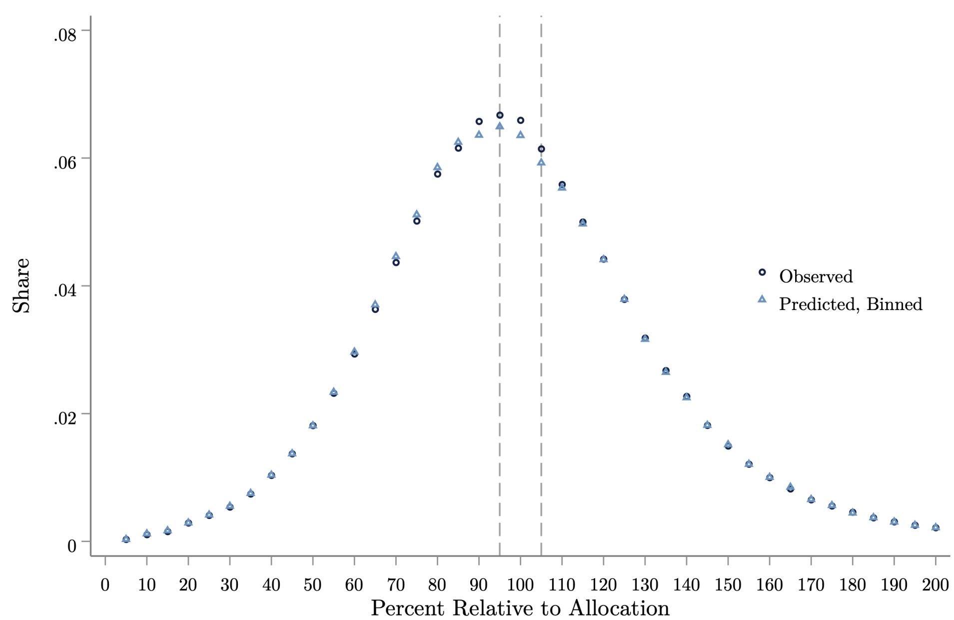

Note: This figure shows the average share of households that fall into each five percent bin relative to the group average and the predicted share of households.

To estimate the missing and excess mass, we can aggregate the share of bunching and missing households for each billing group for each month and calculate the percentage difference between the share of households in each range in each billing group. On average, there are 3.41 percent more households in the bunching range in the observed distribution than the counterfactual distribution. Similarly, there are -1.78 percent fewer households in the missing range in the observed distribution than the counterfactual distribution.

Taken together, these results show that households bunch close to the top of the free-electricity range of electricity at $1.05A$ in response to the nonlinear prices. The uncertainty around electricity consumption leads to some to undershoot the optimal discontinuity while others overshoot.

C Standard Bunching Methods

We also apply the standard approach to bunching analysis used in Chetty et al. (2011) and Kleven and Waseem (2013). The standard bunching approach counts the number of households that ever fall into a particular bin relative to the kink or notch point. Then, the approach fits a predicted polynomial through the data surrounding the bunching range but not including the bunching range.

We begin with the individual data. Rather than calculating the share of households that fall into bins relative to the group average, we count the total number of households across all months across all billing groups. Each bin then contains the total count of households that ever fell into that bin.

The Chetty et al. (2011) method uses a cross-section of the data to predict a counterfactual distribution of electricity use by fitting a polynomial through the data using:

$$C_b = \sum_{i=0}^{q} \beta_i^0 * (Z_b)^i + \sum_{i=-R}^{R} \gamma_i^0 * \mathbb{1}[Z_b = i] + \epsilon_b^0 \tag{8}$$

where C_b is the count of households that fall into percent bin b. Z_j is the percent of electricity use relative to the billing group average in one percent relative to the bunching bin ($Z_j = \{..., -10, -9, ..., 10, ...\}$); q is the order of the polynomial, and R indicates the width of the excluded region around the bunching bin at $1.05A$ (Chetty et al. 2011).

This strategy predicts a counterfactual distribution without the influence of the γ_is. That counterfactual is a smoothed polynomial that best fits the data outside of the bunching region. The share of bunching households is calculated by taking the difference between the observed count of households in the bin and the count of households in the bin under the predicted polynomial. We adjust the Chetty et al. (2011) estimation strategy to take into account the fact that households in the bunching region come from the left of the free-range of electricity.[22]

We use data from bins 35 to 175 to predict the polynomial and consider a range of excluded bins in R. Standard errors are calculated using a parametric bootstrap estimator. Table C.2 shows the resulting estimates for alternative ranges of R. The bunching estimates range from 39.64 percent extra households in the bins from 87 to 118 to 10.00 percent extra households in the excluded range from bin 100 to 109. The top line of the table corresponds to the bunching range predicted by the individual approach including bins 87 to 118. This shows that there are 39.64 percent extra households in this range.

Table C.2: Chetty Specification: Bunching Estimates

(1) Range	(2) Bunching	(3) S.E.	(4) Lower 95% CI	(5) Upper 95% CI
87 to 118	39.64	6.52	52.42	26.86
90 to 110	33.30	4.70	42.51	24.09
90 to 109	33.61	4.86	43.14	24.08
93 to 110	23.24	4.46	31.98	14.50
93 to 109	23.91	4.30	32.34	15.48
95 to 110	21.20	4.25	29.53	12.87
95 to 109	21.94	4.04	29.86	14.02
100 to 109	10.00	4.41	18.64	1.36

Note: Bunching estimates shown in column 2 are the resulting estimates under the Chetty et al. (2011) methodology. Standard errors in column 3 are the result of a bootstrap estimator.

22. Our modified code is available upon request.

Figure C.3 shows the resulting observed and counterfactual distributions under the assumption that the excluded range spans bins 87 to 118. The observed and predicted distributions are shown in Figure C.3. One reason this approach finds smaller estimates of the excess mass is that the Chetty et al. (2011) approach assumes that there are no income effects. However, the free-electricity range expands the household's budget constraint and may cause an income effect which our earlier approach would characterize as excess mass.

Figure C.3: Chetty Method, Observed and Predicted Distributions

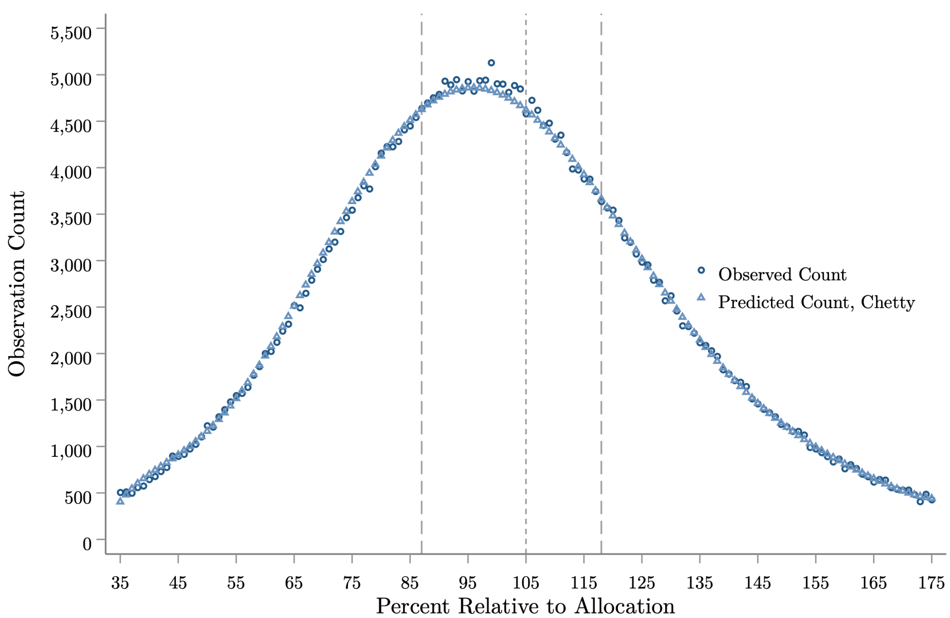

Note: This figure shows the total count of households that ever fell into that one percent bin relative to the group average as well as the counterfactual predicted using Equation (8).

D Binned Approach: Characterizing Optimization Frictions and Elasticities

Using equation (7) we found that the predicted share of bunchers is equivalent to the predicted share of missing households if we include households from 65 up to 90 percent of the group average in the missing region and households from 90 up to 110 percent of the group average in the bunching region. Thus, c_l is 65 percent of the group average for our households.

We can calculate the implicit marginal price Δr, shown in Figure 2, for households that moved from the missing range to the bunching range. Column 1 of Table D.3 shows ranges of Δc that we consider to calculate a range of elasticity estimates. The percent changes in marginal price are shown in column 2 of Table D.3. Column 2 shows the percent change in electricity moving from the bottom and top of the missing range to the bottom and top of the bunching range. Column 4 shows the resulting elasticity estimate from dividing the percent change in electricity use in column 3 by the percent change in electricity prices in column 2.

Table D.3: Binned Specification: Elasticities

(1) Percent Bins Included	(2) Pct. Δ in Price	(3) Pct. Δ in Electricity	(4) Elasticity	(5) Share in Dom. Range [65-104]	(6) Attenuated Elasticity
Implicit Marginal Price					
65 to 90	0.00	27.78	-1.5e+15	0.45	-8.5e+14
65 to 105	-25.00	38.10	-1.52	0.45	-0.84
65 to 110	-22.22	40.91	-1.84	0.45	-1.01
90 to 90	.	0.00		0.45	
90 to 105	-66.67	14.29	-0.21	0.45	-0.12
90 to 110	-50.00	18.18	-0.36	0.45	-0.20
Average Price					
65 to 90	0.00	27.78		0.45	
65 to 105	-9.52	38.10	-4.00	0.45	-2.20
65 to 110	-9.09	40.91	-4.50	0.45	-2.47
90 to 90	0.00	0.00		0.45	
90 to 105	-9.52	14.29	-1.50	0.45	-0.82
90 to 110	-9.09	18.18	-2.00	0.45	-1.10

Elasticities in this table are calculated by first finding the percent change in price from the nonlinear marginal price and either the implicit marginal price or the average price. Next, we calculate the percent change in electricity consumption for each bin group and find the elasticity, which is displayed in the final column.

If households respond to marginal prices, the elasticity estimates range from -1.84 for households bunching from bin 65 up to bin 110 (not including 110) to -0.21 for households bunching from the top of the missing range to bin 105. If households respond to average prices, If households respond to marginal prices, the elasticity estimates range from -4.50 for households bunching from bin 65 up to bin 110 (not including 110) to -1.50 for households bunching from the top of the missing range to bin 105. Note that the elasticities for average price response are larger because the change in average price across the pricing schedule is smaller than the change in marginal price.

Next, we calculate the average share of households that fall into the dominated region from bins 65 to 104 weighted by billing group size. On average, 15.85 percent of households used between 93 and 104 percent of their group average.[23] Last, we compute the elasticity estimates taking into account the possibility of optimization frictions by multiplying the estimates in column 4 by $(1-a)$ shown in column 6. When we take into account optimization frictions, elasticities range from -1.01 to -0.12 if households respond to marginal price or from -2.47 to -0.82 if households respond to average price.

D.1 Simulated Instrument and Encompassing Test

Ito (2014) and Shaffer (2020) use an encompassing test combined with a simulated instrument to investigate whether households respond to marginal or average electricity prices. The encompassing test tests whether the effect of one variable "encompasses" another when they are both included in the same regression. In this case, the authors test whether average prices encompass the effect of marginal prices on electricity consumption:

$$\Delta ln(kWh_{it}) = \beta_0 + \beta_1 \Delta ln(MP_{it}) + \beta_2 \Delta ln(AP_{it}) + \sum_{j=1}^{10} D_{ijt} + \beta_3 \Delta X_{lt} + \eta_{it} \tag{9}$$

where $\Delta ln(kWh_{it}) \equiv ln(kWh_{it}) - ln(kWh_{it_{12}})$, $\Delta ln(MP_{it}) \equiv ln(MP_t(kWh_{it})) - ln(MP_{t_{12}}(kWh_{it_{12}}))$, $\Delta ln(AP_{it}) \equiv ln(AP_t(kWh_{it})) - ln(AP_{t_{12}}(kWh_{it_{12}}))$, D_{ijt} is a dummy variable indicating which decile of electricity use that household's consumption falls in, X_{lt} are heating degree days and cooling degree days for location l in month t.

However, the electricity prices in equation (9 are endogenous. Electricity prices are a

23. The standard deviation across billing groups is 6.11 percent.

function of quantity. One potential solution to this issue is using a simulated instrument Auten and Carroll (1999). The simulated instrument leverages variation in electricity prices that are induced by a change in policy. In Ito (2014) the policy change is temporal variation in electricity prices based on changes in both the tiered rates and where the nonlinearities occur. In Shaffer (2020) the policy change is the introduction of a nonlinear pricing schedule in one utility service territory.

The basic idea behind the simulated instrument is that given a constant level of electricity consumption, the household would still experience a change in their electricity prices over time because of the policy-induced price variation. This simulated instrument is $\Delta ln(\tilde{MP}_{it}) \equiv ln(MP_t(k\tilde{W}h_{it})) - ln(MP_{t_{12}}(k\tilde{W}h_{it}))$ based on constant electricity consumption $k\tilde{W}h_{it}$. We construct the simulated instrument based on two alternative selections for $k\tilde{W}h_{it}$: the six-month-lag of electricity consumption $kWh_{it_{-6}}$ and seasonal average electricity consumption $\overline{kWh_{it}}$.

One minor complication in calculating the simulated instrument is that the number of days in the month six months ago is frequently not the same as the number of days in the month in the time t (and $t-12$) month. The same is true for February in leap years. To adjust for these differences between the number of days in the month, we calculate the electricity prices based on daily average electricity consumption multiplied by the number of days in the month in time t (and $t-12$. This way, none of the variation in electricity prices is due to the different length of days in a given month.

We can estimate the linear model using a two stage least squares regression where the first stage is:

$$\begin{aligned} \Delta ln(MP_{it}) &= \pi_0 + \pi_1 \Delta ln(\tilde{MP}_{it}) + \pi_2 \Delta ln(\tilde{AP}_{it}) + \sum_{j=1}^{10} D_{ij\tilde{t}} + \pi_3 \Delta X_{lt} + \nu_{it} \\ \Delta ln(AP_{it}) &= \pi_0 + \delta_1 \Delta ln(\tilde{MP}_{it}) + \delta_2 \Delta ln(\tilde{AP}_{it}) + \sum_{j=1}^{10} D_{ij\tilde{t}} + \delta_3 \Delta X_{lt} + \nu_{it} \end{aligned} \tag{10}$$

The second stage estimates changes in electricity use as a function of the predicted changes in electricity price from the first-stage estimation.

Second stage:

$$\Delta ln(kWh_{it}) = \beta_0 + \beta_1 \Delta \widehat{ln(MP_{it})} + \beta_2 \Delta \widehat{ln(AP_{it})} + \sum_{j=1}^{10} D_{ijt_6} + \beta_3 \Delta X_{lt} + \eta_{it} \tag{11}$$

All terms are the same as in Equation (9) except the outcomes predicted using the price instruments: $\Delta \widehat{ln(MP_{it})}$ which is the predicted log-change in price from the first-stage regression based on the simulated instrument and $\Delta \widehat{ln(AP_{it})}$, which is the predicted same for average price.

Figure D.4 shows the average variation in the dependent and price variables in equation 9. The important thing to note is that there is very little variation in average prices. The variation in marginal prices arises from households moving in and out of the buffer over time. The structure of the electricity prices in our data lead us to be skeptical of using a simulated instrument in this case. Because electricity prices are a function of the household's group average, allocations are higher in summer months than in winter months. For example, if we are using the log difference in electricity prices between December, the instrument based on the six month lag would be using June electricity consumption. In that case, June electricity consumption will *almost always* fall above the buffer region simply because June electricity use is typically so much higher than December electricity use.

For illustrative purposes, we estimate the price elasticities via the ordinary least squares specification in Equation (11) and via 2SLS in Equation (9) using the six-month-lag instrument and the seasonal average instrument. Tables D.4 through D.6 via OLS and 2SLS with the two alternative price instruments. In particular, the results using the simulated instrument yield nonsensical results for the price elasticities for average price response. This is as a result of the minimal variation in average prices over time. The independent variable is the same in both specifications, but we are attempting to predict the same change in electricity consumption for almost no change in average electricity prices. For this reason, we do not rely on these estimates for any conclusions in the paper.

Figure D.4: Variation in Log-Difference in Electricity Use and Prices

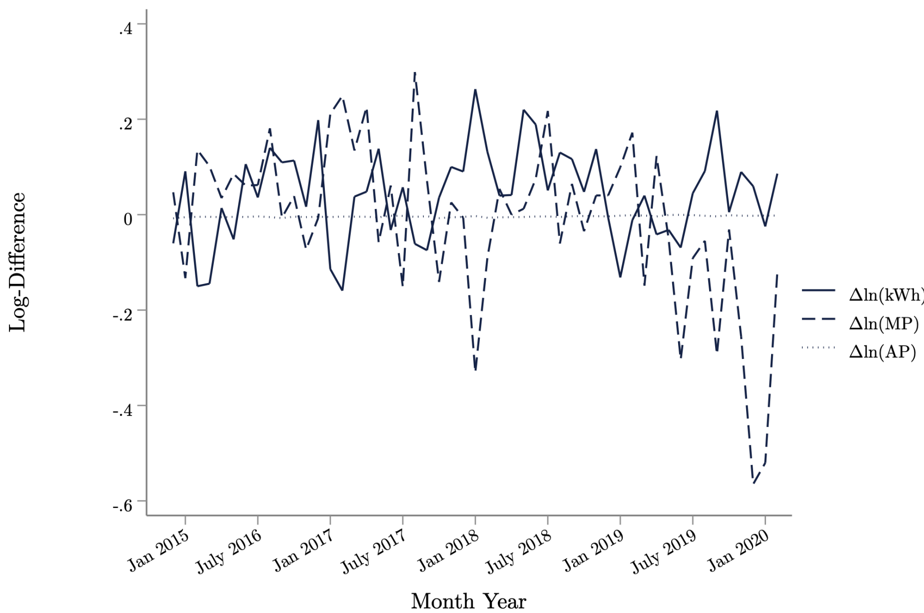

D.2 Tables and Figures

Table D.4: OLS: Encompassing Test

	(1) $\Delta ln(kWh_{it})$	(2) $\Delta ln(kWh_{it})$	(3) $\Delta ln(kWh_{it})$
$\Delta ln(MP_{it})$	-0.002^{*}		-0.001^{*}
	(0.000)		(0.000)
ΔHDD_{it}	0.000^{*}	0.000^{*}	0.000^{*}
	(0.000)	(0.000)	(0.000)
ΔCDD_{it}	0.002^{*}	0.002^{*}	0.002^{*}
	(0.000)	(0.000)	(0.000)
$\mathbb{1}\left[D_{it2}\right]$	-0.063^{*}	-0.010	-0.009
	(0.013)	(0.013)	(0.013)
$\mathbb{1}\left[D_{it3}\right]$	-0.016	0.033^{*}	0.033^{*}
	(0.011)	(0.012)	(0.012)
$\mathbb{1}\left[D_{it4}\right]$	0.015	0.056^{*}	0.055^{*}
	(0.009)	(0.009)	(0.009)
$\mathbb{1}\left[D_{it5}\right]$	0.033^{*}	0.060^{*}	0.059^{*}
	(0.009)	(0.009)	(0.009)
$\mathbb{1}\left[D_{it6}\right]$	0.050^{*}	0.065^{*}	0.065^{*}
	(0.008)	(0.009)	(0.009)
$\mathbb{1}\left[D_{it7}\right]$	0.063^{*}	0.077^{*}	0.076^{*}
	(0.011)	(0.012)	(0.012)
$\mathbb{1}\left[D_{it8}\right]$	0.079^{*}	0.091^{*}	0.092^{*}
	(0.011)	(0.013)	(0.013)
$\mathbb{1}\left[D_{it9}\right]$	0.103^{*}	0.114^{*}	0.115^{*}
	(0.010)	(0.015)	(0.015)
$\mathbb{1}\left[D_{it10}\right]$	0.160^{*}	0.186^{*}	0.187^{*}
	(0.009)	(0.013)	(0.013)
$\Delta ln(AP\_it)$		-1.260^{*}	-1.255^{*}
		(0.124)	(0.124)
N	166,383	165,923	165,923

* p<0.05. The standard errors reported in parenthesis have been clustered at the billing group level.

Table D.5: Simulated Instrument kWh t-6: Encompassing Test

	(1) $\Delta ln(kWh_{it})$	(2) $\Delta ln(kWh_{it})$	(3) $\Delta ln(kWh_{it})$
$\Delta ln(MP_{it})$	0.141		−0.058
	(0.139)		(0.060)
ΔHDD_{it}	0.000*	0.000*	0.000*
	(0.000)	(0.000)	(0.000)
ΔCDD_{it}	0.002*	0.002*	0.002*
	(0.000)	(0.000)	(0.000)
$\mathbb{1}[D_{it2}]$	0.059*	0.028*	0.016
	(0.025)	(0.005)	(0.015)
$\mathbb{1}[D_{it3}]$	0.055*	0.027*	0.017
	(0.024)	(0.004)	(0.013)
$\mathbb{1}[D_{it4}]$	0.042*	0.029*	0.025*
	(0.012)	(0.004)	(0.007)
$\mathbb{1}[D_{it5}]$	0.026*	0.024*	0.024*
	(0.007)	(0.004)	(0.005)
$\mathbb{1}[D_{it6}]$	0.011	0.025*	0.031*
	(0.017)	(0.003)	(0.008)
$\mathbb{1}[D_{it7}]$	0.009	0.019*	0.024*
	(0.016)	(0.004)	(0.006)
$\mathbb{1}[D_{it8}]$	0.005	0.023*	0.031*
	(0.022)	(0.004)	(0.008)
$\mathbb{1}[D_{it9}]$	−0.005	0.022*	0.034*
	(0.033)	(0.003)	(0.013)
$\mathbb{1}[D_{it10}]$	0.003	0.023*	0.033*
	(0.027)	(0.004)	(0.011)
$\Delta ln(AP_{it})$		0.177*	0.234*
		(0.041)	(0.081)
N	166,383	165,870	165,870

* p<0.05. The standard errors reported in parenthesis have been clustered at the billing group level.

Table D.6: Simulated Instrument Average Seasonal Electricity Use: Encompassing Test

	(1) $\Delta ln(kWh_{it})$	(2) $\Delta ln(kWh_{it})$	(3) $\Delta ln(kWh_{it})$
$\Delta ln(MP_{it})$	0.016		0.015
	(0.009)		(0.011)
ΔHDD_{it}	0.000*	0.000*	0.000*
	(0.000)	(0.000)	(0.000)
ΔCDD_{it}	0.002*	0.002*	0.002*
	(0.000)	(0.000)	(0.000)
$\mathbb{1}[D_{it2}]$	0.046*	0.036*	0.040*
	(0.005)	(0.006)	(0.006)
$\mathbb{1}[D_{it3}]$	0.041*	0.032*	0.035*
	(0.005)	(0.005)	(0.005)
$\mathbb{1}[D_{it4}]$	0.031*	0.022*	0.025*
	(0.005)	(0.004)	(0.005)
$\mathbb{1}[D_{it5}]$	0.027*	0.021*	0.022*
	(0.004)	(0.004)	(0.004)
$\mathbb{1}[D_{it6}]$	0.020*	0.017*	0.016*
	(0.004)	(0.004)	(0.004)
$\mathbb{1}[D_{it7}]$	0.019*	0.017*	0.014*
	(0.004)	(0.004)	(0.004)
$\mathbb{1}[D_{it8}]$	0.021*	0.018*	0.015*
	(0.005)	(0.004)	(0.004)
$\mathbb{1}[D_{it9}]$	0.015*	0.013*	0.010*
	(0.004)	(0.004)	(0.004)
$\mathbb{1}[D_{it10}]$	0.030*	0.026*	0.022*
	(0.005)	(0.003)	(0.004)
$\Delta ln(AP_{it})$		0.256*	0.240*
		(0.048)	(0.047)
N	166,383	165,870	165,870

* p<0.05. The standard errors reported in parenthesis have been clustered at the billing group level.

Table D.7: Simple Differences: Average Change in Bin Share After 0 Price

	(1) Bin Share	(2) Bin Share	(3) Bin Share
HDD	0.000*	0.000*	0.000*
	(0.000)	(0.000)	(0.000)
CDD	0.000*	0.000*	0.000*
	(0.000)	(0.000)	(0.000)
Nonlinear	0.000	0.000	0.000
	(0.000)	(0.000)	(0.000)
$\text{Miss}_b * N_t$ (65–85)	−0.001		
	(0.001)		
$\text{Bunch}_b * N_t$ (90–105)	0.002*		
	(0.001)		
$\text{Miss}_b * N_t$ (70–85)		−0.001	
		(0.001)	
$\text{Bunch}_b * N_t$ (90–105)		0.002*	
		(0.001)	
$\text{Miss}_b * N_t$ (60–85)			−0.001*
			(0.001)
$\text{Bunch}_b * N_t$ (90–105)			0.002*
			(0.001)
Constant	0.017*	0.017*	0.017*
	(0.000)	(0.000)	(0.000)
Loc. FE	Y	Y	Y
Month FE	Y	Y	Y
BG FE	Y	Y	Y
Year FE	Y	Y	Y
Bin FE	Y	Y	Y
N	323,640	323,640	323,640

* p<0.05. The standard errors reported in parenthesis have been clustered at the billing group level. Column 1 shows estimates for the full sample without individual fixed effects. Column 2 shows estimates for the full sample with individual fixed effects. And Column 3 shows estimates for the sample limited to households we observe under both pricing regimes.

Table D.8: Individual Regression: Average Change in Electricity Use by Bin Type

	(1) ln(daily kWh)	(2) ln(daily kWh)	(3) ln(daily kWh)
Nonlinear*Bin 1	−1.043*	0.329*	0.325*
	(0.014)	(0.006)	(0.006)
Nonlinear*Bin 2	−0.439*	0.459*	0.457*
	(0.161)	(0.175)	(0.175)
Nonlinear*Bin 3	−0.171*	0.282*	0.282*
	(0.056)	(0.071)	(0.070)
Nonlinear*Bin 4	−0.123*	0.156*	0.156*
	(0.025)	(0.044)	(0.044)
Nonlinear*Bin 5	−0.105*	0.024	0.025
	(0.018)	(0.024)	(0.024)
Nonlinear*Bin 6	−0.028*	−0.011	−0.011
	(0.008)	(0.013)	(0.013)
Nonlinear*Bin 8	−0.006	−0.050*	−0.049*
	(0.007)	(0.005)	(0.005)
Nonlinear*Bin 9	−0.011*	−0.038*	−0.037*
	(0.006)	(0.005)	(0.005)
Nonlinear*Bin 10	−0.012*	−0.029*	−0.028*
	(0.005)	(0.004)	(0.004)
Nonlinear*Bin 11	−0.021*	−0.026*	−0.025*
	(0.005)	(0.004)	(0.004)
Nonlinear*Bin 12	−0.035*	−0.025*	−0.024*
	(0.008)	(0.004)	(0.004)
Nonlinear*Bin 13	−0.030*	−0.019*	−0.018*
	(0.010)	(0.004)	(0.004)
Nonlinear*Bin 14	−0.046*	−0.021*	−0.020*
	(0.011)	(0.004)	(0.004)
Nonlinear*Bin 15	−0.044*	−0.029*	−0.028*
	(0.012)	(0.006)	(0.006)
Nonlinear*Bin 16	−0.039*	−0.030*	−0.029*
	(0.017)	(0.009)	(0.009)
Nonlinear*Bin 17	−0.055*	−0.042*	−0.041*
	(0.012)	(0.010)	(0.010)
Nonlinear*Bin 18	−0.033	−0.050*	−0.048*
	(0.026)	(0.013)	(0.012)
Nonlinear*Bin 19	−0.043	−0.070*	−0.069*
	(0.023)	(0.018)	(0.018)
Nonlinear*Bin 20	−0.022	−0.084*	−0.083*
	(0.024)	(0.014)	(0.014)
HDD	0.002*	0.001*	0.001*
	(0.000)	(0.000)	(0.000)
CDD	0.001*	0.003*	0.003*
	(0.000)	(0.000)	(0.000)
	(0.036)	(0.051)	(0.051)
Month-of-Year FE	Y	Y	Y
Year FE	Y	Y	Y
Individual FE		Y	Y
Decile Bin FE	Y	Y	Y
N	193,334	192,879	171,596

* p<0.05. The standard errors reported in parenthesis have been clustered at the billing group level. Column 1 shows estimates for the full sample without individual fixed effects. Column 2 shows estimates for the full sample with individual fixed effects. And Column 3 shows estimates for the sample limited to households we observe under both pricing regimes.

Figure D.5: Distribution of Monthly Electricity Use by Billing Group

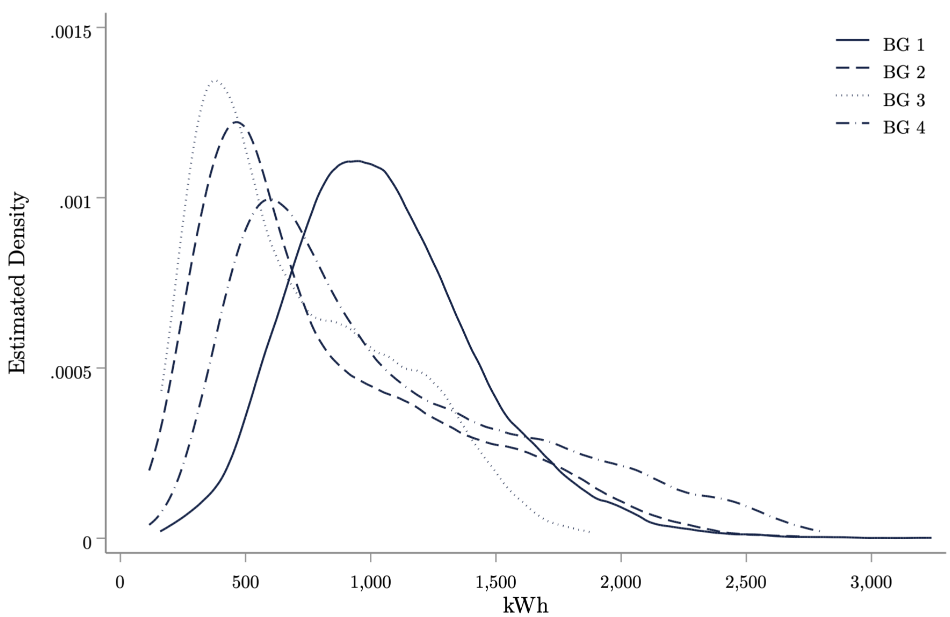

Table D.9: RDD: Rebate Feebate Electricity

	(1) ln(daily kWh)	(2) ln(daily kWh)	(3) ln(daily kWh)	(4) ln(daily kWh)	(5) ln(daily kWh)	(6) ln(daily kWh)
Act. Bal. + 15 ($)	0.012*	0.010*	0.004	0.020*	0.008	0.019*
	(0.002)	(0.001)	(0.005)	(0.003)	(0.013)	(0.008)
1[Act. Bal. $\leq$ -15]	−0.042*	−0.031*	−0.062*	−0.013	−0.059*	−0.018
	(0.011)	(0.006)	(0.016)	(0.008)	(0.019)	(0.012)
1[Act. Bal. $\leq$ -15] * (Act. Bal + 15)	−0.006*	0.002	−0.002	−0.007	−0.006	−0.011
	(0.002)	(0.001)	(0.007)	(0.005)	(0.018)	(0.011)
HDD	0.003*	0.001	0.003*	0.001	0.003*	0.001
	(0.000)	(0.001)	(0.000)	(0.001)	(0.000)	(0.001)
CDD	0.003*	0.001*	0.003*	0.001*	0.003*	0.001*
	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)
(Act. Bal. + 15)^2 ($)			0.001	−0.001*	0.000	−0.001
			(0.000)	(0.000)	(0.003)	(0.002)
1[Act. Bal. $\leq$ -15] * (Act. Bal + 15)^2			−0.001*	0.001*	0.000	0.000
			(0.001)	(0.000)	(0.004)	(0.003)
(Act. Bal. + 15)^3 ($)					0.000	0.000
					(0.000)	(0.000)
1[Act. Bal. $\leq$ -15] * (Act. Bal + 15)^3					0.000	0.000
					(0.000)	(0.000)
Constant	1.629*	2.988*	1.643*	2.971*	1.640*	2.972*
	(0.140)	(0.076)	(0.137)	(0.078)	(0.135)	(0.079)
Month-Year FE	Y	Y	Y	Y	Y	Y
N	58,581	41,940	58,581	41,940	58,581	41,940

* $p<0.05$. The standard errors reported in parenthesis have been clustered at the billing group level.

Columns 1, 2, 3, and 4 show the results from three alternative specifications.

Table D.10: RDD: Rebate Feebate Electricity

	(1) ln(daily kWh)	(2) ln(daily kWh)	(3) ln(daily kWh)
Act. Bal. + 15 ($)	0.009*	0.009*	0.004
	(0.001)	(0.003)	(0.010)
1[Act. Bal. $\leq$ -15]	−0.040*	−0.044*	−0.052*
	(0.007)	(0.010)	(0.015)
1[Act. Bal. $\leq$ -15] * (Act. Bal + 15)	−0.001	−0.003	−0.004
	(0.001)	(0.004)	(0.013)
HDD	0.002*	0.002*	0.002*
	(0.000)	(0.000)	(0.000)
CDD	0.002*	0.002*	0.002*
	(0.000)	(0.000)	(0.000)
Act. Bal t-2 ($)	0.001*	0.001*	0.001*
	(0.000)	(0.000)	(0.000)
Act. Bal t-3 ($)	0.001*	0.001*	0.001*
	(0.000)	(0.000)	(0.000)
(Act. Bal. + 15)^2 ($)		0.000	0.001
		(0.000)	(0.002)
1[Act. Bal. $\leq$ -15] * (Act. Bal + 15)^2		0.000	−0.003
		(0.000)	(0.003)
(Act. Bal. + 15)^3 ($)			0.000
			(0.000)
1[Act. Bal. $\leq$ -15] * (Act. Bal + 15)^3			0.000
			(0.000)
Constant	2.238*	2.238*	2.242*
	(0.080)	(0.078)	(0.076)
Month-Year FE	Y	Y	Y
N	83,043	83,043	83,043

* p<0.05. The standard errors reported in parenthesis have been clustered at the billing group level.

Columns 1, 2, 3, and 4 show the results from three alternative specifications.

Table D.11: Change in Probability of Bunching with Lags

	(1) $\mathbb{1}[B_{it}]$	(2) $\mathbb{1}[B_{it}]$	(3) $\mathbb{1}[B_{it}]$	(4) $\mathbb{1}[B_{it}]$	(5) $\mathbb{1}[B_{it}]$	(6) $\mathbb{1}[B_{it}]$	(7) $\mathbb{1}[B_{it}]$
$\mathbb{1}[O_{it}]$	0.047*	0.045*	0.045*	0.045*	0.044*	0.038*	0.038*
	(0.007)	(0.007)	(0.007)	(0.007)	(0.006)	(0.006)	(0.006)
$\mathbb{1}[B_{it_{-1}}]$	0.171*						
	(0.008)						
$\mathbb{1}[N_t]$	−0.011*	−0.006	−0.006	−0.010*	−0.011*	−0.009*	−0.009*
	(0.004)	(0.004)	(0.004)	(0.004)	(0.004)	(0.004)	(0.004)
$\mathbb{1}[O_{it} * N_t]$	0.015*	0.013*	0.017*	0.023*	0.026*	0.021*	0.021*
	(0.003)	(0.003)	(0.004)	(0.004)	(0.004)	(0.004)	(0.004)
$\mathbb{1}[O_{it} * B_{it_{-1}}]$	0.024*						
	(0.009)						
$\mathbb{1}[N_{it} * B_{it_{-1}}]$	−0.007						
	(0.009)						
$\mathbb{1}[O_{it} * N_t * B_{it_{-1}}]$	0.028*						
	(0.009)						
$\mathbb{1}[B_{it_{-2}}]$		0.199*					
		(0.009)					
$\mathbb{1}[O_{it} * B_{it_{-2}}]$		0.073*					
		(0.012)					
$\mathbb{1}[N_{it} * B_{it_{-2}}]$		−0.047*					
		(0.013)					
$\mathbb{1}[O_{it} * N_t * B_{it_{-2}}]$		0.039*					
		(0.012)					
$\mathbb{1}[B_{it_{-3}}]$			0.249*				
			(0.012)				
$\mathbb{1}[O_{it} * B_{it_{-3}}]$			0.048*				
			(0.014)				
$\mathbb{1}[N_{it} * B_{it_{-3}}]$			−0.082*				
			(0.017)				
$\mathbb{1}[O_{it} * N_t * B_{it_{-3}}]$			0.066*				
			(0.018)				
$\mathbb{1}[B_{it_{-4}}]$				0.223*			
				(0.014)			
$\mathbb{1}[O_{it} * B_{it_{-4}}]$				0.075*			
				(0.017)			
$\mathbb{1}[N_{it} * B_{it_{-4}}]$				−0.024			
				(0.019)			
$\mathbb{1}[O_{it} * N_t * B_{it_{-4}}]$				0.038			
				(0.020)			
$\mathbb{1}[B_{it_{-5}}]$					0.207*		
					(0.019)		
$\mathbb{1}[O_{it} * B_{it_{-5}}]$					0.089*		
					(0.023)		
$\mathbb{1}[N_{it} * B_{it_{-5}}]$					0.031		
					(0.036)		
$\mathbb{1}[O_{it} * N_t * B_{it_{-5}}]$					−0.025		
					(0.035)		
$\mathbb{1}[B_{it_{-6}}]$						0.238*	0.238*
						(0.020)	(0.020)
$\mathbb{1}[O_{it} * B_{it_{>-6}}]$						0.118*	0.118*
						(0.021)	(0.021)
$\mathbb{1}[N_{it} * B_{it_{>-6}}]$						0.006	0.006
						(0.041)	(0.041)
$\mathbb{1}[O_{it} * N_t * B_{it_{>-6}}]$						−0.017	−0.017
						(0.042)	(0.042)
N	371,498	357,912	344,604	331,802	319,330	296,784	296,784

* $p<0.05$. The standard errors reported in parenthesis have been clustered at the billing group level.

Column 1 shows estimates for a bunching streak of 1 month. Column 2 shows estimates for a bunching streak of 2 months. Column 3 shows estimates for a bunching streak of 3 months. Column 4 shows estimates for a bunching streak of 4 months. And Column 5 shows estimates for a bunching streak of 5 months. And Column 6 shows estimates for a bunching streak of 6 months.

Figure D.6: Seasonal and Alternative Bin Robustness Tests

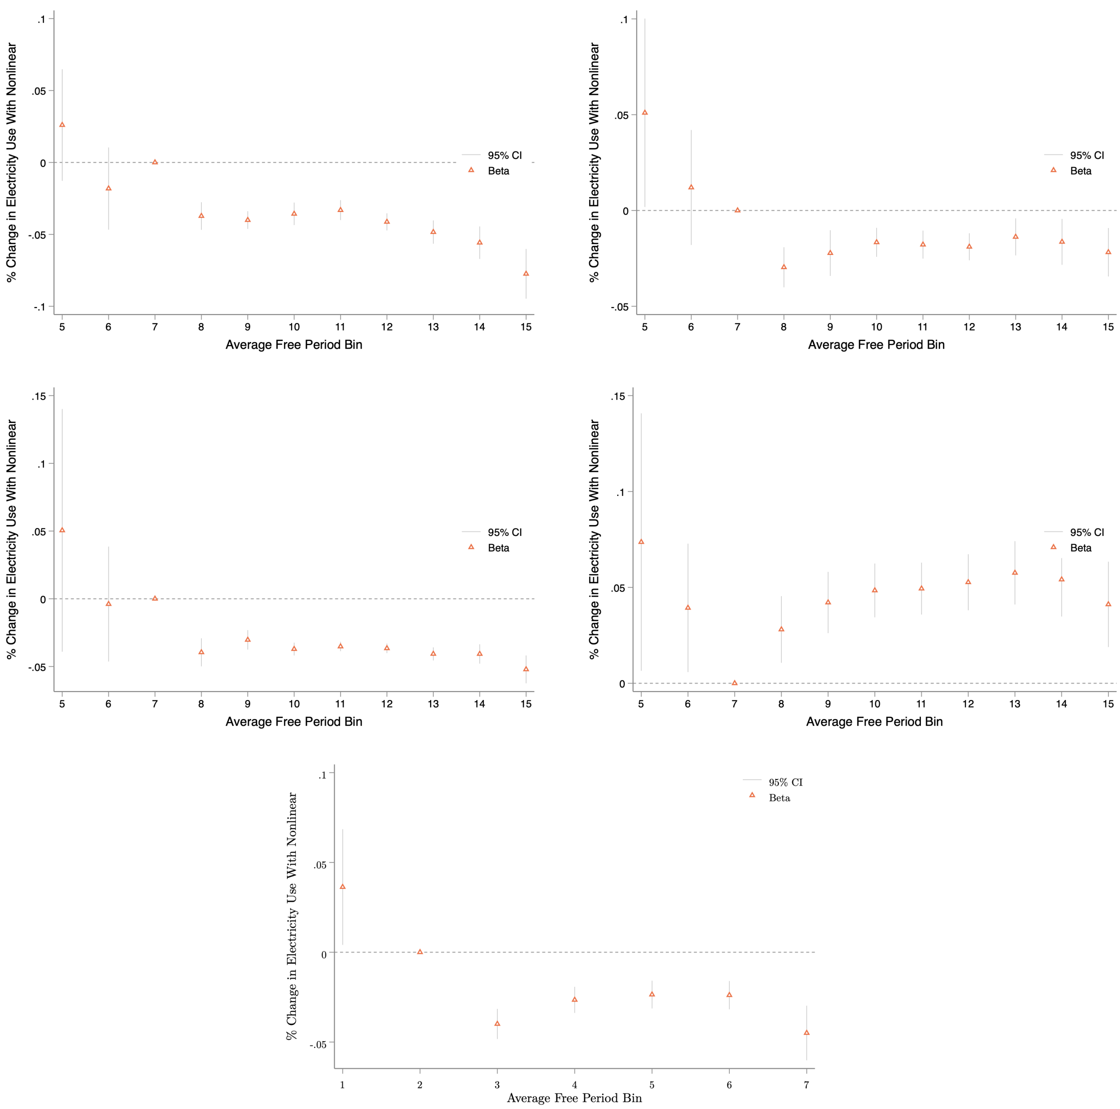

Figure D.7: Seasonal and Alternative Bin Robustness Tests