

Rational, Lazy or Confused? Evidence of misperception in consumer responsiveness to nonlinear prices

By BLAKE SHAFFER*

This paper examines how consumers respond to nonlinear prices. Using a rich dataset that exploits a natural experiment on electricity consumers in British Columbia, I find evidence of households misunderstanding nonlinear prices—specifically that marginal price applies to all consumption, not simply incremental. While small in number, these households have a large effect in aggregate, masking an otherwise predominant response to average price. Previously largely unexplored in the literature, misperception has important economic and policy implications beyond electricity markets. I estimate the welfare loss of these misperceiving households to be the equivalent of 10% of annual electricity expenditure.

JEL classifications: D12, Q41, C23, C24, C26

HOW DO CONSUMERS RESPOND to nonlinear prices? Standard economic theory predicts consumers will optimize *at the margin*, to the point where marginal benefit equals marginal cost. However, information costs and rational inattention may result in consumers behaving differently (Sallee, 2014). For example, in a study of Californian households, Ito (2014) finds electricity consumers respond to average, not marginal, cost. If this is correct, policies based on theoretical assumptions may lead to unintended outcomes. Going beyond whether consumers respond to marginal or average price, this paper explores a novel and more fundamental question: To what extent do consumers misunderstand nonlinear prices? In other words, are consumers *rational*, *lazy* or *confused*?¹

To answer this question, I take advantage of a quirk in the structure of British Columbia’s electricity market that creates a natural experiment. In October 2008, BC Hydro, the

* Department of Economics, University of Calgary (email: blake.shaffer@ucalgary.ca). I wish to thank Werner Antweiler, Pamela Campa, Koichiro Ito, Andreas Lange, Ken McKenzie, Lucija Muehlenbachs, Rob Oxoby, Stefan Staubli, Atsuko Tanaka, Scott Taylor, Trevor Tombe, and Frank Wolak for advice and suggestions. I would also like to thank participants at the 2016 CREE conference and 2017 Toulouse conference on energy and climate change for helpful comments. Finally, I am indebted to John Duffy and Michael Li at BC Hydro, and Rod Carle at the City of New Westminster, for their willingness and enthusiasm to contribute data for this project.

¹Throughout this paper, I affectionately refer to marginal, average, and misperceived price responders as *rational*, *lazy* and *confused* types. No judgement is implied.

provincial electric utility serving 95% of the province, changed from a flat rate to a two-tier increasing block tariff whereby the price of electricity increases for incremental consumption beyond a quantity threshold each billing cycle. Meanwhile, the City of New Westminster—one of the few places in the province, which for historical reasons sets its own electricity rates—serves as a near-ideal control group in that it chose to remain on a flat rate. The data for this paper consists of a rich set of monthly billing records from 2005–2013, covering the universe of households in New Westminster and the neighbouring regions served by BC Hydro.

Using a mix of reduced form and structural methods, I uncover behavior that on the surface indicates marginal cost responsiveness: there are large changes in electricity consumption for households near the tariff threshold. However, upon closer inspection I find this to be explained by heterogeneity in price perception among households. I find that most households respond to average cost (85%), a small share respond to marginal cost (7%), and a small but important share of households misperceive the nonlinear price schedule (8%). Specifically, these *confused* households appear to mistakenly perceive jumps in marginal cost to apply to all consumption, not just incremental. While small in number, these *confused* households have a significant effect on aggregate results, leading to sizeable welfare losses and important policy implications.

This finding contributes in three ways. First, from a policy perspective, the presence of misperception has implications for optimal rate setting, potentially misleading policymakers from achieving their goals. Initially, misperception is helping achieve the policy goal of conservation: relative to a population consisting solely of average cost responders, this small share of *confused* households reduces aggregate electricity consumption by roughly 1%. However, I find that the number of households bunching at the threshold diminishes over time, suggesting as consumers educate themselves on how the tariff works it is leading to more average cost responsiveness and thus, ultimately, less conservation. I estimate a simple flat rate would deliver 1% more conservation, or roughly the equivalent of a 10% price increase, versus the two-tier pricing structure.

Second, in terms of welfare, I find that misperception causes *confused* households to over-respond near the threshold, resulting in significant welfare losses of up to 10% of their annual electricity expenditure.

Third, from a methodological perspective, this paper serves as a caution against the sole use of average treatment effects when examining consumer response in the presence of price perception heterogeneity. If one were to assume homogeneous households, one would conclude from bunching and panel data analysis that the population consists solely of marginal

cost optimizers. However, allowing for heterogeneity in price perception, I use conditional difference-in-differences and indirect inference to uncover a population consisting largely of average cost responders, and a small share of *confused* and marginal cost responders. Nevertheless, this mix of heterogeneous households is capable of mimicking marginal cost responsiveness in the eyes of certain reduced form estimators, confounding inference from average treatment effects alone.²

Early attempts to investigate nonlinear tariff responsiveness in the tax literature relied on difference-in-differences estimation using different groups experiencing different changes in rates after a tax schedule change (Eissa, 1996). Those that are subject to little change serve as control, whereas those experiencing larger changes serve as treatment. The problem with this approach is that these groups are likely to be compositionally different (e.g. high versus low income) and likely to respond to nonlinear tariffs in different ways. In the context of electricity consumption, Borenstein (2009) points out that while it may be tempting to compare changes between high-consuming and low-consuming households (i.e. groups that face different prices as a result of the nonlinear tariff), the presence of natural mean reversion biases the results. Borenstein (2009) states, “separating the household mean reversion from the effect of rate changes is possible in theory, but fairly challenging in practice” [p.23]. In this paper, using New Westminster as a control group, I can observe changes for households in the same place in the distribution of consumption, but facing different price schedules. Intuitively, observed mean reversion by New Westminster households can be subtracted from the effect observed by BC Hydro customers, leaving the residual change the result of the nonlinear tariff.

More recently, modern econometric techniques exploit quasi-experimental variation to identify consumer responsiveness to nonlinear tariffs. Nataraj and Hanemann (2011) employ a regression discontinuity (RD) design to estimate household responsiveness to nonlinear tariffs for water consumption, comparing households just below and above a newly introduced threshold in a water tariff, both before and after implementation. Ito (2014) adds a second dimension, employing a spatial discontinuity design that compares usage over time across two differently-affected regions. My identification strategy adds a third dimension, considering changes across time, space and deciles of household electricity consumption. Since changes in marginal and average prices differ between low and high consumers facing nonlinear tariffs, this third dimension allows for identification of heterogeneous responsiveness, and ultimately the presence of misperception.

²This methodological caution is consistent with recent arguments made by Blomquist and Newey (2017) regarding inference from bunching estimators when faced with heterogeneity in consumer preferences.

The notion that consumers may not optimize according to marginal cost when faced with complex nonlinear tariffs is not new. Bartolome (1995) investigates individuals’ responses to income tax and finds that roughly half use average tax rate ‘as if’ it is the marginal tax rate. In other words, they view jumps in marginal tax rates applying to all their income, not simply incremental income, creating the same type of misperception as the *confused* households discussed here. Interestingly, he finds this misperception to be socially beneficial in that it reduces the distortionary behavioural labour response to income taxes. In the context of electricity, Schneider and Sunstein (2017) find that “when transaction costs and decision biases are taken into account, the most cost-reflective electricity policies are not necessarily the most efficient”.

This paper is structured as follows. Section 1 provides the necessary background on increasing-block tariffs, as well as historical context for the British Columbia electricity market, and Section 2 describes the data. The empirical analysis is divided into two parts. In the first part (Section 3), I perform reduced form empirical analysis using three methods (bunching, instrumental variables, and difference-in-differences) to estimate the causal effect of the nonlinear tariff on consumption. In the second part (Section 4), I present a simplified model of heterogeneous consumer behaviour based on household types that respond to marginal, average and misperceived prices, and invoke indirect inference to solve for the mix of types that best fits the reduced form estimates. Section 5 discusses the policy and welfare implications of misperception and Section 6 concludes.

I. Background

Electricity provides a suitable setting to examine consumer responsiveness to nonlinear pricing for several reasons. First, despite its everyday usage, consumers are generally unaware of both their actual electricity consumption and its cost. This *lack of salience* suggests consumers are unlikely to respond according to the predictions of standard theory (Chetty, Looney and Kroft, 2009). Even in the case where attention is paid, complex rate tariffs can often lead consumers to misperceive their marginal price (McRae and Meeks, 2015). Second, a widely-used residential electricity tariff provides the necessary nonlinear structure in which to empirically examine the research question. An *increasing-block tariff* involves a low rate for all household consumption up to a defined quantity in each billing cycle, followed by a higher rate for all incremental consumption above this threshold. This tariff provides the needed variation between marginal and average prices to separately estimate responsiveness. Lastly, the availability of *large administrative data* provides the necessary power to empirically analyse the question.

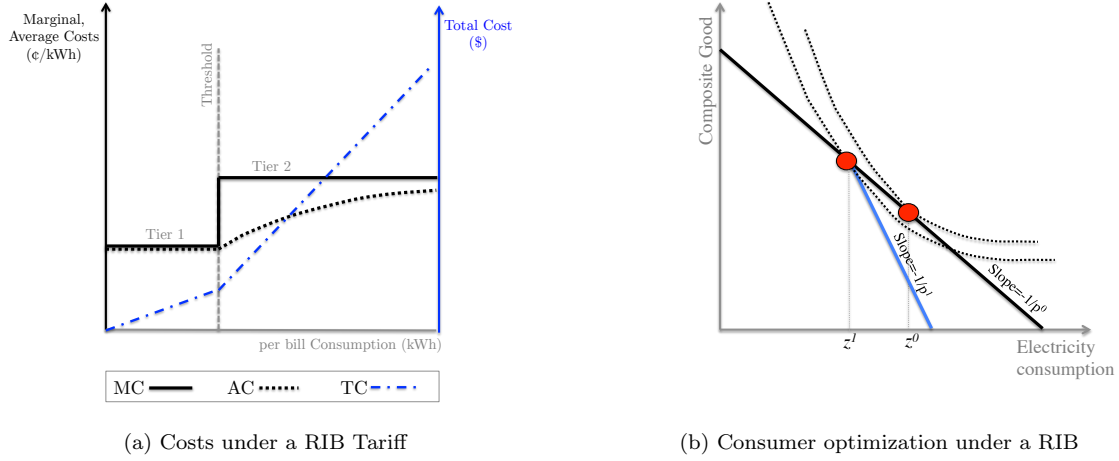


Figure 1. : THE ECONOMICS OF A RIB TARIFF

Notes: In figure 1a, the solid black line illustrates the marginal price of electricity under a 2-step RIB tariff. The marginal price jumps higher at the threshold. The dotted black line illustrates the effect on average price. It matches marginal price below the threshold, but increases asymptotically beyond the threshold. Total cost is shown in blue (right axis), with its slope matching the respective marginal price before and after the threshold. Figure 1b presents a stylized representation of consumer optimization under a RIB. The solid black line illustrates the budget constraint of a representative consumer prior to the introduction of a RIB. The optimal bundle occurs at the point z^0 , where their indifference curve is tangent to the budget line. The RIB introduces a kink in the budget line at the threshold. The new optimal allocation, shown here, shifts left to z^1 .

A. What is an Increasing-Block Tariff?

A residential increasing block tariff (henceforth ‘RIB’) involves an increasing marginal price for electricity. In a 2-tier RIB, consumers pay a low per-unit rate for all consumption up to a defined threshold within each billing cycle and a higher per-unit rate for all consumption above the threshold.³ Figure 1a illustrates marginal, average, and total costs under a RIB. The change in marginal cost is abrupt; there is a step change at the threshold. Whereas, average cost does not have a step, but rather a gradual and asymptotic increase towards the higher level. The total cost curve does not contain any step, but rather a kink at the threshold. The slope of the total cost curve steepens beyond the threshold in accordance with the higher Tier 2 marginal price.

The basic idea behind a RIB is that by raising the price on consumption beyond a specific threshold, large consumers responding to marginal cost will conserve. A simple model helps us develop the intuition for this result. Suppose we have a representative consumer allocating their income, m , across electricity, z , and a composite good, x , by

³A progressive income tax schedule is another form of increasing block tariff. The marginal tax rate increases with income, with the higher rate only applying to incremental income.

optimizing their utility in the standard manner:

$$(1) \quad \max_{x,z} U(x, z) \text{ subject to: } x + pz \leq m$$

where the price of the composite good is normalized to 1. The first order condition would lead to an optimal amount of electricity consumption, z^* , such that the marginal rate of substitution (MRS) between z and x equals p . Now, if we introduce a kink in the budget constraint by way of a nonlinear tariff, we change the effective “ p ” to which the MRS must equal. In Figure 1b this reduces the optimal level of consumption from z^0 to z^1 . Of course, underpinning this theory is the assumption that consumers respond at the margin, to the point where marginal benefits equal marginal cost—an assumption we examine in detail in this paper.

The increasing-block tariff structure is widely-used around the world. In a survey by BC Hydro (2014), they find that 35% of utilities surveyed used an increasing-block tariff (31 out of 88). Of those, over half use the simplest two-step tariff. Despite their widespread use, there is little empirical evidence as to how consumers respond to such tariffs.

B. The British Columbia context

British Columbia is a province in Canada with over 4.5 million residents. Over 95% of the province’s electricity demand is served by the provincially-owned electric utility, BC Hydro (BC Hydro, 2015). Regions not covered by BC Hydro include Fortis BC in the interior of the province (formerly West Kootenay Power) and various cities that for historical reasons retain local distribution and price-setting ability, of which the city of New Westminster, located in the populated Greater Vancouver regional district, is one.

In October 2008, BC Hydro switched its residential rate to an increasing-block tariff. The motivation was to promote conservation by large-users, while maintaining revenue neutrality by lowering the first tier rate (BC Hydro, 2014). BC Hydro describes the introduction of the RIB in its *Evaluation of the Residential Inclining Block Tariff Report*:⁴

“In August 2008 the British Columbia Utilities Commission determined that it was in the public interest for BC Hydro to implement the new RIB rate and required the new RIB rate structure go into effect October 1, 2008 for approximately 1.6 million residential customers [accounts]. The Step 1 to Step 2 threshold was set at 1,350 kWh per billing period, which was

⁴BC Hydro uses the term *inclining* block rate whereas I use the more common term *increasing* block rate. The intent of the terms is synonymous.




Figure 2. : MAP OF BC LOWER MAINLAND

Notes: This map shows the Greater Vancouver region in the southwest portion of British Columbia. New Westminster is shown in yellow, with the surrounding BC Hydro service territory used for the analysis shown in Orange. The remaining region (green) is also served by BC Hydro.

approximately 90 per cent of the median consumption of BC Hydro’s residential customers. The Step 2 rate was established at BC Hydro’s current estimate of the cost of new energy supply, grossed up for losses and the Step 1 rate was calculated to achieve revenue neutrality for the residential class.” (BC Hydro (2014), p.ii)

The City of New Westminster, did *not* match the switch to a RIB, creating a near-ideal natural experiment. Consumers in New Westminster remained on a flat rate tariff, while BC Hydro customers in the neighbouring cities of Burnaby, Coquitlam, Richmond and Surrey switched to a RIB, in many cases across the street from one another. Figure 2 depicts the Greater Vancouver region of BC, with the City of New Westminster shown in yellow. All other regions in Figure 2 are served by BC Hydro. The orange areas are the six forward sortation areas (FSAs) bordering New Westminster used here.⁵

Figure 3 shows the evolution of BC Hydro and New Westminster residential electricity prices over time. Prior to October 2008, New Westminster and BC Hydro shared near identical rates.⁶ After October 2008, the BC Hydro rate splits into two: Tier 1 and Tier 2. New Westminster remains on a single rate, with annual changes intended to track BC Hydro’s average rate change. Of note, the BC Hydro Tier 1 rate falls upon RIB implemen-

⁵The 6 BC Hydro FSAs used for this analysis are V3K, V3N, V3V, V4C, V5E and V5J.

⁶The City of New Westminster generally matched any changes that BC Hydro made to their rates prior to October 2008. The small deviations between the two rates prior to October 2008 were unintentional, and instead were temporary delays in getting municipal council approval.




Figure 3. : BC HYDRO AND NEW WESTMINSTER ELECTRICITY RATES; 2005–2013

Notes: This figure presents time series of electricity rates. Prior to implementation of the RIB, rates in NW and BCH were roughly the same. After October 2008, BCH splits into two lines, representing the two rate tiers. Rate increases can be seen occurring roughly annually in both regions.

tation and remains consistently below the New Westminster single rate afterwards. The Tier 2 rate is consistently above the New Westminster rate.⁷

II. Data

The data consist of billing information for over 190,000 households from 2005–2013, containing over 6 million observations. This covers the universe of households in the City of New Westminster and the six neighbouring forward sortation areas (FSAs) in BC Hydro’s service territory.

The raw data contain anonymized premise ID, bill start and end dates, and consumption in kilowatt-hours (kWh). These data were then merged with publicly available price data for both regions to complete the dataset. For the empirical analysis, I use a balanced dataset containing 34,592 households and 3.7 million monthly observations with accounts spanning the entirety of the study period.⁸

Table 1 presents summary statistics from the balanced dataset. The demographic infor-

⁷In conversations with New Westminster officials, the reasoning behind remaining on the flat rate were twofold. First, there was a sense that the public preferred the simplicity of a flat rate. And second, as we will see in our review of the data, New Westminster customers are on average smaller users as compared to the average BC Hydro customer. This would have led to a revenue shortfall if New Westminster adopted the same threshold and rate tiers as BC Hydro.

⁸Details of the balanced dataset creation are described in the Appendix. The reduction in households and observations in the balanced versus raw datasets reflects the large number of accounts that do not span the entirety of the 9 year period.

Table 1—: SUMMARY STATISTICS

	New West		BCH (6 FSAs)	
	Pre	Post	Pre	Post
Daily consumption (kWh)				
Mean	21.1	21.1	27.8	27.5
Median	16.6	16.5	24.1	23.5
Standard deviation	16.7	16.9	18.6	18.4
Skewness	2.3	2.2	2.0	1.9
Average bill prices (cents/kWh)				
Marginal Price	6.17	7.75	6.18	7.88
Average Price	6.17	7.75	6.18	6.94
Number of observations	458,055	641,277	1,098,585	1,538,019
Number of households	10,179		24,413	
Median household income (\$2010)	54,932		63,949	
Share of renters	44%		32%	
Mean number of rooms per dwelling	5.0		5.8	

Note: All statistics relate to the balanced panel dataset.

mation is obtained by matching FSA information to 2011 Census data. Several features are worth noting. First, both the mean and median daily consumption in New Westminster is significantly lower than the neighbouring BC Hydro region. This is a function of New Westminster's higher density arising from a larger share of apartment units and lower number of rooms per dwelling. Second, New Westminster has a greater share of renters (44% versus 32%), raising a potential problem for comparison of consumer responsiveness: If renters do not pay utilities, or do not have control over them, they may be less responsive (Levinson and Niemann, 2004).⁹ I overcome this problem in various ways. First, my initial empirical strategy is a bunching estimator that relies largely on BC Hydro data alone. Second, the difference-in-differences strategy controls for level differences between the regions; it focusses only on changes in trends. Valid identification requires that consumption trends are parallel prior to the reform, not that levels are the same. I improve the validity of this assumption by performing the difference-in-differences estimation conditional on decile of pre-reform consumption, in a manner similar to a triple difference estimation. This delivers conditional average treatment effects (CATE) for similarly sized households. This is discussed further in Section 3.C.

⁹Of course, a counterargument could be made that control by landlords of a large number of units would raise the attentiveness to electricity bills and serve to *increase* responsiveness.

Table 2 presents the same BCH and NW data broken out by decile of pre-reform consumption. The deciles are determined across *all* households, not separately for BCH and NW, allowing for better comparisons of like-for-like households across the regions.

Table 2—: DESCRIPTIVE STATISTICS BY DECILE

BC HYDRO (ALL 6 NEIGHBOURING FSAs)

Decile	Demand (kWh)						Prices (cents/kWh)					
	Mean			Median			MC			AC		
	Pre	Post	Δ	Pre	Post	Δ	Pre	Post	Δ	Pre	Post	Δ
1	5.9	6.6	11.2%	5.9	6.1	3.1%	6.18	6.47	4.7%	6.18	6.44	4.1%
2	10.5	11.1	5.5%	10.2	10.3	1.3%	6.18	6.52	5.4%	6.18	6.45	4.3%
3	14.4	15.0	4.5%	14.0	14.2	1.1%	6.18	6.67	7.9%	6.18	6.48	4.8%
4	17.9	18.3	2.4%	17.4	17.4	(0.0%)	6.18	6.95	12.4%	6.18	6.52	5.6%
5	21.4	21.8	2.1%	20.9	20.7	(0.6%)	6.18	7.49	21.1%	6.18	6.64	7.4%
6	25.1	25.2	0.6%	24.5	24.0	(2.0%)	6.18	8.05	30.3%	6.18	6.79	9.8%
7	29.1	29.0	(0.4%)	28.4	27.6	(2.7%)	6.18	8.49	37.3%	6.18	6.98	13.0%
8	34.0	33.5	(1.5%)	33.1	31.9	(3.7%)	6.18	8.77	41.9%	6.18	7.20	16.5%
9	40.9	39.7	(2.9%)	39.8	38.0	(4.7%)	6.18	8.95	44.8%	6.18	7.46	20.7%
10	59.3	55.6	(6.2%)	54.4	51.0	(6.3%)	6.18	9.01	46.8%	6.18	7.85	27.0%

NEW WESTMINSTER

Decile	Demand (kWh)						Prices (cents/kWh)					
	Mean			Median			MC			AC		
	Pre	Post	Δ	Pre	Post	Δ	Pre	Post	Δ	Pre	Post	Δ
1	5.9	6.5	8.5%	5.9	6.0	2.7%	6.17	7.75	25.6%	6.17	7.75	25.6%
2	10.4	10.7	3.1%	10.1	10.1	0.1%	6.17	7.75	25.6%	6.17	7.75	25.6%
3	14.3	14.7	2.8%	13.8	13.8	(0.6%)	6.17	7.75	25.6%	6.17	7.75	25.6%
4	17.8	18.3	2.7%	17.1	17.0	(0.8%)	6.17	7.75	25.6%	6.17	7.75	25.6%
5	21.4	21.9	2.6%	20.6	20.6	(0.1%)	6.17	7.75	25.6%	6.17	7.75	25.6%
6	25.1	25.6	2.2%	24.2	24.1	(0.3%)	6.17	7.75	25.6%	6.17	7.75	25.6%
7	29.1	29.3	0.7%	28.3	27.8	(1.7%)	6.17	7.75	25.6%	6.17	7.75	25.6%
8	34.1	33.4	(2.0%)	33.1	32.0	(3.4%)	6.17	7.75	25.6%	6.17	7.75	25.6%
9	40.7	39.8	(2.3%)	39.6	38.2	(3.7%)	6.17	7.75	25.6%	6.17	7.75	25.6%
10	60.5	57.2	(5.5%)	54.6	51.5	(5.6%)	6.17	7.75	25.6%	6.17	7.75	25.6%

Note: Deciles are determined across all households in the dataset, not separately by region. All statistics relate to the balanced panel dataset.

The data-by-decile in Table 2 provide a first glimpse of the pre- versus post-reform consumer responsiveness. The first thing to note is the appearance of mean reversion in demand. The lowest deciles are increasing their consumption in the post-reform period, while the higher deciles are decreasing consumption. Importantly, this is occurring in both regions and cannot be directly attributed to the RIB. The empirical analysis in Section 3 strips out the responsiveness attributable to the RIB from the change observed in the control group. Second, looking at prices, we see the uniform change in New Westminster’s tariff across the deciles, which is the same for average and marginal costs since the flat rate creates no difference between the two. For BC Hydro, the top deciles see their prices increase more than the lower deciles, with the change being larger for marginal cost, as compared to average cost. I use the different changes in average and marginal cost before and after RIB implementation, between the regions and across usage deciles to provide three dimensions of variation for the empirical analysis.

III. Empirics

In the first part of the empirical analysis, I employ three reduced form quasi-experimental strategies, each exploiting different facets of the data, to investigate consumer responsiveness to the nonlinear tariff. In the subsections that follow, I detail the empirical methodologies and their respective results separately.

A. Bunching Design

We begin with the first of these methods—*bunching*. This method uses less of the available data, however, it offers simple and clear visual evidence of consumer responsiveness.

If consumers respond to marginal prices, we should expect to see bunching at the threshold of the higher Tier 2 rate. The intuition is illustrated in Figure 4a.¹⁰ Originally, Consumer A sets her optimal amount of electricity consumption at $z + dz^*$. With the introduction of the nonlinear tariff, the budget constraint changes, with the slope steepening to the right of the threshold at z^* . This causes consumer A to shift to a lower indifference curve until it is tangent to the new budget constraint precisely at the kink, z^* . Whereas, Consumer B, whose original indifference curve was tangent at z^* , is left unaffected. This creates a bunching of households originally in the region $[z^*, z^* + dz]$ at z^* .

This bunched mass can be used to estimate elasticity. Saez (2010) shows that, by defini-

¹⁰I follow Saez (2010) for this derivation and in the illustrative graphs that follow.




Figure 4. : BUNCHING

Notes: These three figures present stylized illustrations of bunching at the threshold. Similar to Figure 1b, Figure 4a illustrates the effect of the introduction of a RIB on consumer optimization. Prior to the RIB, consumer A is at location $z + dz$. Upon implementation, consumer A shifts to z^* . Whereas, consumer B remains at z^* both before and after the introduction of the RIB. Figure 4b shows the mass of consumers that will attempt to shift towards the threshold. Figure 4c shows an example of what bunching mass may look like in the ex-post distribution.

tion, for a small price change, dp , the price elasticity of consumption is given by:

$$(2) \quad \frac{dz^*}{z^*} = e \frac{dp}{p}$$

We know dp, p and z^* from the tariff; all that is left to calculate e is to estimate dz^* . To do so I estimate the mass of bunching at the threshold relative to a counterfactual distribution with no nonlinear tariff. Figure 4b illustrates a stylized example of the shift in mass to the threshold. If consumers optimize perfectly to the change in marginal price, the area under the pre-reform curve between z^* and $z^* + dz^*$ would shift entirely to the vertical line at z^* . Realistically, it is impractical to suggest the bunching would all occur precisely at the point z^* . Thus, in practice we expect a modest area surrounding z^* to be “attempts” at perfect optimization. Figure 4c illustrates the area of attempted bunching, B , above a counterfactual distribution. To determine dz , I calculate the ratio of B over the counterfactual amount of mass at the threshold, h_0 . Specifically, $dz = \frac{B}{h_0} * binsize$. In order to do so, a credible counterfactual must be identified. In this paper, I identify three such counterfactuals (for robustness). But before describing the counterfactuals in detail, let us take the first step in visually checking for the presence of any bunching.

EVIDENCE OF BUNCHING

In both Ito (2014) and Borenstein (2009), there is little evidence of bunching at the kink points of the nonlinear tariff used for Southern California Edison electricity customers. In




Figure 5. : DISTRIBUTION OF BC HYDRO CONSUMPTION BY HOUSEHOLD

Notes: This figure presents two overlaid histograms of BC Hydro consumption by household. The solid (green) bars represent the distribution of households in 2007, prior to the RIB. The open (white) bars show the post-RIB distribution in 2009. An increase in density is observed near threshold (shown by the red line).

this case, looking at BC Hydro customers with a single newly-introduced threshold, the picture is different. Figure 5 plots two overlapping histograms of household consumption for the years immediately before and after the RIB implementation (2007 and 2009).¹¹ The 2007 distribution appears smooth across the threshold (shown as the red vertical line). The 2009 distribution, however, displays evidence of bunching, with a clearly visible “bump” at the threshold. There appears to be a decrease in mass in the region to the right of the threshold (specifically the 30–50 daily kWh range), while the region close to the threshold (22.2 kWh) has markedly increased mass. This superimposed figure provides a clear picture of bunching at the threshold, and also serves as the basis of one of the counterfactual methods discussed below.

CONSTRUCTING THE COUNTERFACTUAL

Quantifying this evidence requires the construction of a counterfactual distribution; in other words, what would the distribution of household consumption be in the absence of the nonlinear kink? For robustness, I identify three such counterfactuals, with the use of available data increasing with each one.

¹¹Plotting other “before and after” years displays similar visual evidence of bunching. Although, as noted in the Appendix, the amount of bunching appears to dissipate over time—another hint that a learning process may be going on.

Method 1 - Polynomial counterfactual

The first method closely follows Chetty et al. (2011), constructing the counterfactual distribution by fitting a flexible polynomial to the actual data, excluding data in the region of observed bunching, by estimating the following regression equation:

$$(3) \quad C_j = \sum_{n=0}^p \beta_n \cdot (Z_j)^n + \sum_{i=z^L}^{z^U} \gamma_i \cdot \mathbf{1}[Z_j = i] + \epsilon_j$$

where C_j is the density of household bills in bin j ; Z_j is the consumption observed in bin j ; and $\mathbf{1}[Z_j = i]$ is a dummy variable indicating whether bin j is in the excluded region. The order of polynomials, p , and the excluded range, $[z^L, z^U]$ are subjective. I use a seven order polynomial and set the excluded range to be three bins on either side of the threshold based on visual inspection.¹²

The size of bunching, B , is calculated as the area between the actual data and the counterfactual, within the excluded region. By taking the ratio of this size to the predicted density of the counterfactual bin at the threshold, h , I get an estimate of the change in consumption induced by the nonlinear tariff, dz .

However, by omitting the large mass in the area of observed bunching, the counterfactual has a cumulative density less than that of the actual data. This has the effect of overstating the amount of excess mass in the bunching region and consequently overstating elasticity. Thus the counterfactual distribution must be “corrected” in order for its cumulative distribution to match that of the actual distribution. Chetty et al. (2011) allocate the missing mass to the right side of the distribution on the basis that this is where individuals would have shifted away from to remain under an income tax threshold. In this case, there is reason to believe shifting is occurring *both* from the right and left sides of the distribution on account of the marginal price dropping for small consumers relative to the pre-reform prices. I correct the polynomial by uniformly scaling all bins such that the sum of the corrected mass matches that of the actual data.¹³

Figure 6a plots the actual distribution of BC Hydro household bills in 2009 and the counterfactual distribution constructed by this method.¹⁴ The shaded region represents

¹²The results are robust to the choice of polynomial order, p . Increasing the range of z^L and z^U increases the estimated elasticity from approximately -0.05 to -0.10, however, the standard error increases as well.

¹³As a robustness check, I also follow Chetty et al.’s correction method by adjusting only bins to the right of the threshold. I find no significant difference in elasticity estimates between the two correction methods.

¹⁴I use 2009 for all bunching estimates. This is largely chosen to be as close to possible to the policy




Figure 6. : COUNTERFACTUAL DISTRIBUTIONS

Notes: These three figures illustrate the 2009 distribution of BC Hydro consumption by household (in red), with three different counterfactual distributions. In all cases, the estimated area of bunching mass is shown in grey.

the excess mass due to bunching.

Method 2 - ‘Pre-reform treated group’ counterfactual

As an alternative to the polynomial method, I use the 2007 distribution of the treated group as the counterfactual (i.e. BC Hydro customers in the year before the policy change). This straightforward counterfactual avoids the parametric assumptions required by Method 1, and the subjective requirements of choosing the exclusion region for the polynomial regression and mass correction method. However, this method omits time-varying factors that could change the distribution of density by percentile between 2007 and 2009. Figure 6b plots the actual distribution of BC Hydro household bills in 2009 and the counterfactual distribution constructed by this method.

Method 3 - ‘Pre-reform treated group scaled by growth in control group’ counterfactual

To resolve the issue of omitted time-varying changes in Method 2, I make use of the observed changes over time in the control group, the City of New Westminster. It would be tempting to use the control group’s distribution as the counterfactual, however, in this case the pre-reform distribution of the control and treated groups are significantly different. Direct use of the control group’s distribution as counterfactual confounds time-varying changes with pre-existing compositional differences between treatment and control. Instead, I construct a third counterfactual by scaling each decile of the 2007 BC Hydro distribution with a growth factor specific to the change observed in each decile in New

change (Oct 2008) in an effort to limit error from time-varying factors in methods 2 and 3. Elasticity estimates for the subsequent years are listed in Table 5.

Westminster data between 2009 and 2007. The key assumption here is that changes to density-by-decile in New Westminster are similar to those observed in neighbouring BCH regions.

Figure 6c plots the actual distribution of BC Hydro household bills in 2009 and the counterfactual distribution constructed by this method. While there are slight differences between the counterfactuals, the presence of excess bunching remains clearly evident in all three.

RESULTS AND INTERPRETATION

Table 3 presents results using all three counterfactuals. The different counterfactual methods produce similar elasticity estimates in the range of -0.041 to -0.048 (s.e. 0.010–0.017). This is at the lower end of estimated price elasticity of electricity demand in the literature.¹⁵ To the extent consumers respond to average, not marginal, prices, a low elasticity by this method is not unexpected. The bunching estimator measures the elasticity purely with respect to changes in marginal price (i.e. a *marginal* price elasticity). In comparison, Ito (2014) found no significant marginal price elasticity when estimated by bunching methods.

Note that a bunching estimator is affected in two ways. The greater the degree of actual price elasticity, the greater will be the bunching and the estimated elasticity. Consumers do not, however, bunch perfectly at the threshold. Part of this sub-optimization is due to non-price demand shocks. The larger the variance of these shocks, the *less* bunching we will observe and thus lower elasticity estimates. As Borenstein (2009) notes, “if customers try to optimize, but have very large optimization error, then there would be little or no bunching, but there would also be less hope of identifying demand elasticity based on responses to the jumps in the *ex post* marginal price” [p.14]. Given these forces (elasticity and optimization error, or uncertainty) act in opposing ways, the calculated elasticity estimate from the bunching estimator—without incorporating uncertainty—should be viewed as a lower bound.

B. Instrumental Variables Design

The bunching estimator is compelling in its visual simplicity; however, inference is limited to responsiveness in the region near the threshold. It also makes little use of the rich

¹⁵In a survey of the literature by Jamil and Ahmad (2011), estimates of short-run price elasticities of electricity demand range from -0.06 to -0.33 (excluding the highest and lowest outliers of -1.06 and -0.02, respectively), with the median being -0.145.

Table 3—: BUNCHING ESTIMATES OF PRICE ELASTICITY

Polynomial	2007 BCH	2007 BCH scaled by NW
-0.048	-0.041	-0.045
(0.010)	(0.012)	(0.017)

Notes: Standard errors in parentheses. To calculate standard errors, I use a bootstrap method by sampling from 5% of the population and reconstructing counterfactuals and elasticity estimates 100 times. Method 3 demonstrates slightly larger standard errors due to higher variance in growth factors among smaller New Westminster samples.

household-level panel data and control group to which I have access. To exploit these data features, I regress annual percentage changes in consumption on percentage changes in marginal and average prices at the household level using monthly panel data. This *encompassing test* (Davidson and MacKinnon, 1993) tests whether the effect of one variable (eg. marginal price) is rendered insignificant with the inclusion of another (eg. average price). In other words, does one effect encompass the other?

A regression on prices, however, suffers from the problem of endogeneity. The structure of the RIB means higher consumption mechanically leads to higher marginal and average prices. As a result, OLS creates a spuriously positive correlation between price and consumption. To resolve this, I use an instrumental variable common to the public finance literature, a *simulated instrument* (Murray, 2005). Specifically, I take a prior period household consumption level and project it onto current tariffs *as if* their consumption level did not change. In doing so, the simulated instrument captures only the change in prices due to a change in tariff rates, not due to any change in behaviour.

To be a valid instrument, the simulated instrument must be correlated with price (non-weak instrument) and uncorrelated with consumption changes (exclusion restriction). The first stage regression shows strong correlation between the instrument and prices. Thus, the non-weak instrument requirement is satisfied. The exclusion restriction cannot be directly tested. Here I follow Ito (2014) in the logic of using a prior period consumption level that is halfway between the start and end periods used to calculate the change in relevant variables. Since I am calculating changes as 12 month differences, I select consumption in period $t - 6$. The argument is that initial period consumption levels would be affected by mean reversion—lower initial levels would be correlated with larger positive changes, and higher initial levels would be correlated with lower or negative changes. By using the midpoint of the difference period, the mean reversion concerns are reduced.

To further allay concerns over different trends in consumption between low and high

users, I include controls for prior period consumption into the regression. Following Ito (2014), I take a non-parametric approach for these controls by creating dummy variables of consumption percentiles, D_{qit} . Essentially this places households in bins of consumption at each period t based on their consumption in the same midway period consumption as used in the simulated instrument ($t - 6$). Lastly, I include region fixed effects (γ_{ci}) to control for trends in consumption that differ by region. The regression equation is:

$$(4) \quad \Delta \ln x_{it} = \beta_1 \Delta \ln MP_{it} + \beta_2 \Delta \ln AP_{it} + \sum_{q=1}^{100} D_{qit} + \gamma_{ci} + \epsilon_{it}$$

RESULTS

Table 4 presents the IV regression results. The regression is specified three different ways: first with only marginal price changes as the independent variable; second with only average price changes; and thirdly with both. Standard errors are clustered at the household level. Columns 1–3 summarize the results of these three regressions, but without the use of fixed effects. Column 1 shows a price elasticity of -0.136 with respect to marginal price only. Column 2 shows a price elasticity of -0.133 with respect to average price only. However, when the regression includes both marginal and average price, the effect from average price is rendered insignificant. In other words, consumers do not respond to average price changes once marginal price changes are accounted for. This is the opposite of Ito’s (2014) result and consistent with evidence of responsiveness to marginal prices found in Section 3A using a bunching estimator. Columns 4–6 run the same regressions as 1–3, but include percentile-by-year and region controls. The results are similar in that the effect of average price is rendered insignificant by the inclusion of marginal price. The estimated elasticity is slightly higher, at -0.155.

C. Conditional Difference-in-Differences Design

The bunching estimates and instrumental variables results suggest that BC Hydro consumers are responding to marginal price, unlike Ito’s (2014) finding of Californian consumers responding to average price. The question is why? To answer this, I focus on heterogeneity of behaviour across households by estimating conditional average treatment effects (CATE) in a manner similar to Wichman (2017) and Abrevaya, Hsu and Lieli (2015).

I estimate a difference-in-differences (DD) regression conditional on pre-reform usage deciles. This allows for comparison of like-for-like households between the two regions,

Table 4—: ELASTICITY ESTIMATES USING IV METHOD

	(1)	(2)	(3)	(4)	(5)	(6)
$\Delta \ln MP$	-0.136 (0.007)	· ·	-0.141 (0.010)	-0.154 (0.008)	· ·	-0.155 (0.011)
$\Delta \ln AP$	· ·	-0.133 (0.009)	0.010 (0.013)	· ·	-0.157 (0.010)	0.002 (0.014)
Percentile-time FE	NO	NO	NO	YES	YES	YES
Region FE	NO	NO	NO	YES	YES	YES

Notes: Standard errors clustered at the household level are shown in parentheses.

exploiting variation in marginal and average price that differs across both region, time and consumption levels. These DD estimates deliver the conditional average treatment effect, or CATE, for each decile.¹⁶

To verify the required parallel trends assumption, Figure 7 plots trends in consumption for each of the pre-reform usage deciles for both NW and BCH. While the plots demonstrate some noise around the trends, the general direction for both NW and BCH is consistent within each decile. In the low deciles, BCH households appear to be increasing their consumption post-reform at a greater rate than NW households. Decile 6 (just above the threshold) shows BCH consumption falling faster than NW, whereas the higher groups are indiscernible. Overall, the parallel trend assumption within each decile appears valid.

To estimate the CATE, I interact indicators of pre-reform consumption deciles with treatment indicators in a manner similar to triple difference estimation:

$$(5) \quad \ln \text{dailykwh}_{it} = \alpha \text{Decile}_{id} + \beta (\text{Post2008}_t \times \text{Decile}_{id}) + \delta (\text{BCH}_i \times \text{Decile}_{id}) + \gamma_d (\text{BCH}_i \times \text{Post2008}_t \times \text{Decile}_{id}) + \eta_i + \phi_t + \epsilon_{it}$$

where $\ln \text{dailykwh}_{it}$ is the natural logarithm of consumption for household i in period t ; Decile_{id} is a dummy variable equal to 1 if household i 's pre-reform usage falls in decile $d \in [1, 10]$; BCH_i is a dummy variable equal to 1 if the household is served by BCH; Post2008_t is a dummy variable equal to 1 if the observation is in a time period after the RIB implementation in October 2008; and η_i and ϕ_t are household and time fixed effects,

¹⁶It would be tempting to perform a simple difference-in-difference (DD) approach, comparing BCH vs NW and pre vs post, but the parallel trend assumption would be violated given New Westminster households are on average smaller consumers than surrounding BC Hydro households (see Table 1). A simple DD would also only deliver the ATE, whereas the CATE by usage decile provides more information useful in the subsequent indirect inference stage of the analysis.

respectively.

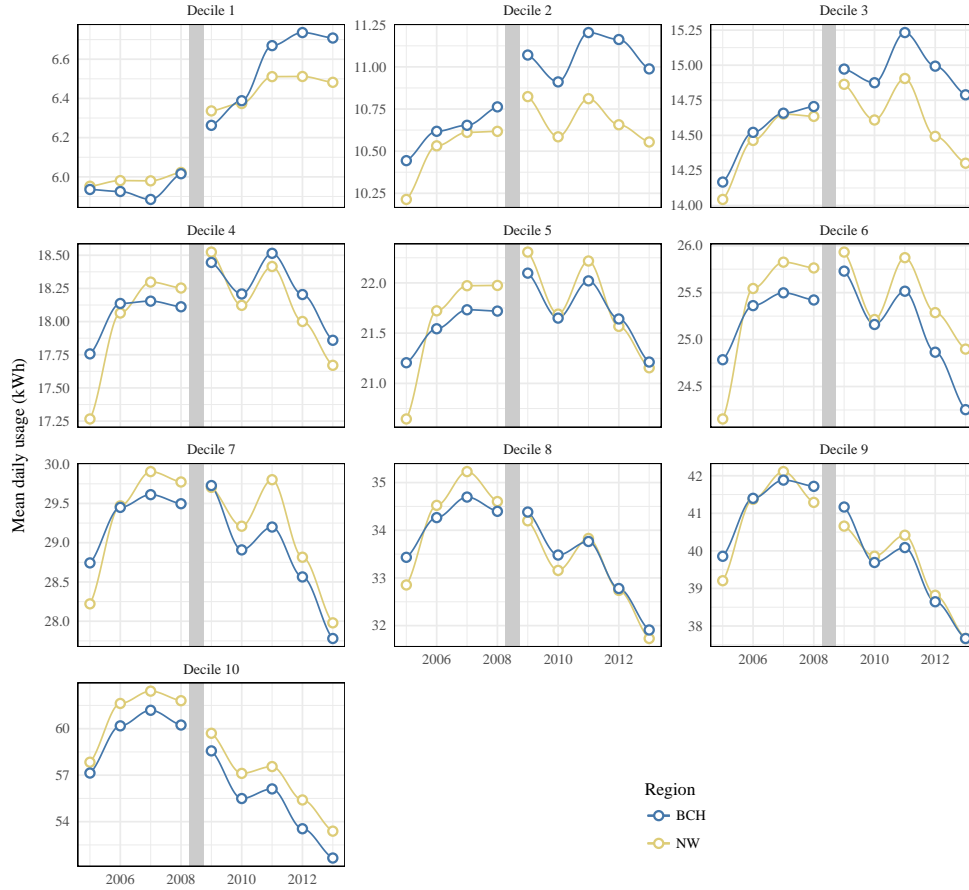


Figure 7. : CONSUMPTION TRENDS BY PRE-REFORM USAGE DECILES

Note: Each subplot represents a different decile of the consumption distribution. Time trends in mean consumption are shown for BCH (blue) and NW (yellow).

RESULTS

The estimated CATE are listed in Table A1 and plotted in Figure 8a. These estimates should be interpreted as the percentage change in average daily consumption for BCH households after the policy change relative to the change in NW households, for each decile separately. For the lower deciles, we see a 1–2% increase in BCH consumption relative to NW in the years following the RIB. This makes sense, as BCH households in these deciles face lower prices (both marginal and average), as seen in Figure 8b. Decile 6 is a different story. Decile 6 is situated just beyond the threshold—the point at which marginal cost jumps for BCH households relative to NW. Here we see a significant decrease

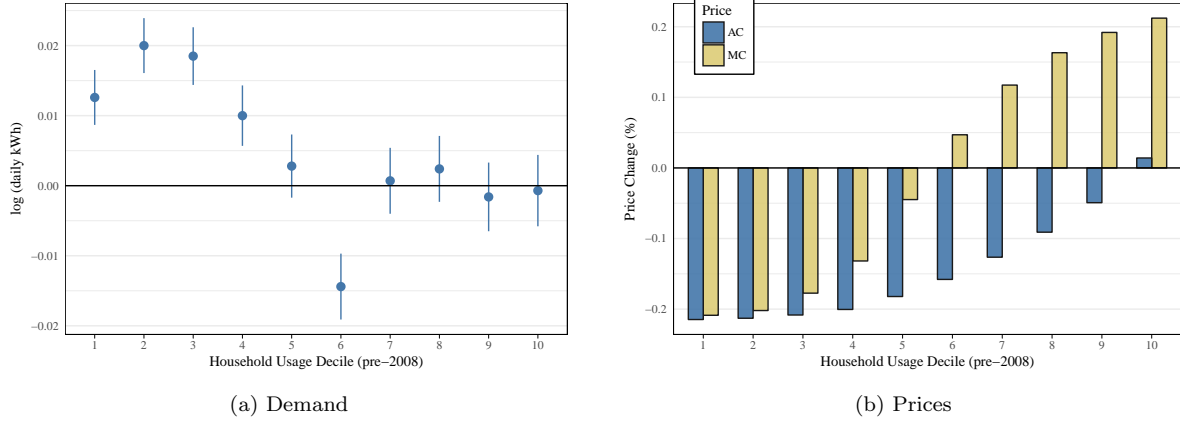


Figure 8. : RELATIVE CHANGE BY DECILE

Notes: Figure 8a plots the CATE for each decile with 95% error bars. The reduction in demand in BCH (relative to NW, after the RIB) can be seen clearly by the coefficient for the 6th decile, immediately above the RIB threshold. Figure 8b shows the change in prices between BCH and NW before vs after the RIB. Marginal price (yellow) increases for the higher deciles, on account of BCH's higher rate for larger consumption. The change in average prices (blue) is more gradual. If consumers were responding purely to either of these prices we would expect the shape of the coefficients to mirror the inverse of the price changes.

in consumption for BCH households relative to NW. This is consistent with marginal cost responsiveness identified by bunching analysis and encompassing tests. However, and critically so, deciles 7–10 show no significant difference in consumption changes between BCH and NW. This is inconsistent with the marginal cost responsiveness story identified by the earlier methods.

Figure 8b plots the relative price changes between BCH and NW after the introduction of the RIB, by decile. The changes in relative demand as expressed by the coefficient estimates in Figure 8a are not entirely consistent with either the marginal or average cost picture. Thus we have a puzzle.

IV. Indirect Inference

The puzzle presented by the reduced form empirical analysis is one of consumer behaviour that appears responsive to marginal cost, but may not be. One possibility is price elasticities differing across usage deciles due to income effects. The intuition being that the income effect diminishes at higher incomes, and higher incomes being correlated with larger residential users.

To investigate this possibility, Figure 9 follows Wichman (2014) by plotting price elasticities by decile of usage based on the CATE results with respect to both marginal and




Figure 9. : PRICE ELASTICITY OF DEMAND

Notes: Price elasticity of demand is plotted for each decile of household consumption. Elasticity with respect to marginal (blue) and average (yellow) price changes are similar for deciles 1-4 (-0.05 to -0.10), and again for high deciles (not significantly different than zero). At decile 6, the two elasticities diverge. Whisker lines represent standard errors.

average price changes. Comparing the lower deciles (1-4) to higher deciles (7-10) is consistent with decreasing elasticity (more inelastic) at higher usage levels. This is true regardless of whether the response is coming from changes in marginal or average cost. But the results at the threshold imply the story is not solely due to income effects. If one were to conclude that consumers are responding to marginal cost, it implies that a rather smooth decline in elasticity from -0.10 to zero is interrupted at decile 6—the decile adjacent to the threshold—where consumers are suddenly more elastic (-0.30). Similarly, to conclude average cost responsiveness requires accepting a positive elasticity for decile 6. Neither explanation is compelling.

To resolve the puzzle, I hypothesize the possibility of heterogeneous household behaviour based on three distinct household types responding to different perceived prices, which I affectionately refer to as *Rational*, *Lazy* and *Confused*.

- 1) *Rational types*: Well-informed and calculating households that optimize their consumption based on **marginal cost**. I represent preferences for this type of household with a quasi-linear utility function subject to the following budget constraint:

$$(6) \quad \max_{x,z} U(x,z) \text{ subject to } \begin{cases} x + p_1 z \leq m, & \text{if } z \leq \bar{z} \\ x + p_1 \bar{z} + p_2 (z - \bar{z}) \leq m, & \text{if } z > \bar{z} \end{cases}$$

- 2) *Lazy types*: Ito (2014) finds strong evidence of households whose cost to obtain rate and usage information is greater than the benefit of optimizing against marginal cost.

Lazy types respond to **average cost** according to the following preferences:

$$(7) \quad \begin{aligned} & \max_{x,z} U(x, z) \text{ subject to } x + p_a z \leq m \\ & \text{where } p_a = \frac{[p_1 \min(z, \bar{z}) + p_2 \max(z - \bar{z}, 0)]}{z} \end{aligned}$$

- 3) *Confused types*: This type of consumer **misperceives** the higher price to apply to *all* consumption, rather than only incremental consumption above the threshold. In other words, *confused* households interpret the threshold as a jump in average cost. This is represented by the following utility function:

$$(8) \quad \max_{x,z} U(x, z) \text{ subject to } \begin{cases} x + p_1 z \leq m, & \text{if } z \leq \bar{z} \\ x + p_2 z \leq m, & \text{if } z > \bar{z} \end{cases}$$

If consumers misperceive their rate tariff in this third manner, they have a strong incentive to reduce consumption from areas just to the right of the threshold to avoid the large inframarginal increase to their total costs.

There is some evidence of this type of consumer. McRae and Meeks (2015) perform a price-elicitation survey of electricity customers facing a newly-introduced nonlinear tariff in Krygzystan to test understanding. In their survey, only 24% of households correctly understood the price jumps of their increasing block tariff. Whereas, 42% misperceived the tariff as jumping at the threshold *for all consumption*, not simply the incremental amount above the threshold. The remaining households were either not aware of a nonlinear tariff (i.e. thought it was a flat rate) or did not display consistent patterns to enable classification.

Closer to this study, BC Hydro (2014) performed a survey of its customers to gauge awareness of the RIB. They found that 50% of their customers were aware of the nonlinear tariff. The remainder either thought they were still on a single rate tariff (31%), “didn’t know” (17%), or, interestingly, thought they were on a declining block rate (2%). Awareness of the RIB, however, does not ensure a correct understanding of the tariff. To test this, BC Hydro asked questions to those who identified as “aware of the RIB” to determine if they correctly identified their marginal price. Just over half (57%) answered in a manner consistent with the correct understanding of the RIB. While this is (in my opinion) an impressive number of well-informed households, it leaves a large segment of the population with the potential to misperceive the price.

METHODOLOGY

To test for the presence of misperception, I use the method of indirect inference. Indirect inference is a variant of the generalized method of moments (GMM), useful when nonlinear models make estimation by more efficient methods, such as maximum likelihood, intractable (Gourieroux, Monfort and Renault, 1993; Smith, 2008).¹⁷ Indirect inference involves using simulated data based on an economic model to estimate parameters in an auxiliary model. In this case, we want to find the mix of household types and price elasticity such that the CATE estimates from simulated and actual data are as close as possible.

Formally, let θ be a vector containing the structural parameters of interest (i.e. shares of types and elasticity), γ be a 10×1 vector of CATE estimates from the simulated data and γ_R be the reduced form CATE estimates calculated in Section 3C. Our optimization problem is thus to find the θ that minimizes the following criterion function:

$$(9) \quad \min_{\theta} [\gamma(\theta) - \gamma_R]' W [\gamma(\theta) - \gamma_R]$$

where W is a weighting matrix.¹⁸

In order to deliver estimates of $\gamma(\theta)$, I first simulate the data generating process for all three types of consumers in our economic model of heterogeneous household behaviour. I begin by creating a single representative household for each percentile of the pre-reform distribution, using their first year monthly means as a starting point. I then apply both a stochastic and deterministic process to simulate data for the entire 108 month time period.

The stochastic process involves both time-varying common shocks (ϕ_t in Eq.10) and idiosyncratic shocks (ϵ_{it} in Eq.10). Both are log-normal with variances calibrated from the pre-reform data.

The deterministic process is a rule-based response to price changes that varies depending on household type. For *rational* and *lazy* types, I use marginal and average price changes, respectively. I use percentage change in prices between period $t - 1$ and $t - 12$ as a way to deal with the endogeneity problem of using prices in period t that are themselves set by consumption in period t . Specifically, demand follows the equation:

¹⁷Indirect inference shares many similarities with *simulated method of moments* (SMM). The key difference is the use of an auxiliary model in indirect inference, whereas SMM is a direct matching of moments (Fackler and Tastan, 2008)

¹⁸ W should put more weight on more precisely estimated coefficients in γ_R . With similar standard errors, W converges to the identity matrix. The criterion function then simplifies to being the sum of squared errors between simulated coefficients and those estimated from the data.

$$(10) \quad \ln(x_{it}) = \ln(x_{i,t-12}) + \eta * \ln\left(\frac{p_{i,t-1}}{p_{i,t-12}}\right) + \phi_t + \epsilon_{it}$$

For *confused* types, I follow the same logic as the *lazy* types (responding to average prices), unless they fall within a specified distance from the threshold, in which case they reduce demand in an attempt to avoid crossing the threshold.¹⁹ Their attempt is imperfect, as the stochastic shocks occur after the deterministic rule is applied. This behavioural assumption offers a plausible heuristic of consumer behaviour that provides a buffer against the consumer finding themselves in a preference-dominated region.

New Westminster households are simulated to serve as the control group, responding to price changes in NW with the same constant elasticity (η in Eq.10). There is no difference between marginal and average prices in NW.

RESULTS

The simulated before and after distributions are shown in Figure 10. The *rational* households display the presence of bunching, as they are responding to the kink at the threshold. *Lazy* households do not. *Confused* households demonstrate a dramatic response at the threshold due to the fact that the notch creates a dominated range just to the right of the threshold. The pattern is far more extreme than observed in the actual data. Critically, none of the individual types, by themselves, visually matches the data.²⁰

Of note, the *rational* households display estimated elasticities, from both the bunching and IV methods, that are consistent with the actual data (Tables A2, A3). Whereas *lazy* households do not demonstrate any significant bunching, and their responsiveness to marginal cost disappears once average cost is accounted for. *Confused* households display overly strong responsiveness as compared to the actual data.

The CATE coefficients do not align with estimates from the actual data for any of the household types (Figure A1). The *rational* and *lazy* types do not reflect the response at the threshold, whereas the *confused* estimates are, again, overly strong.

None of the individual household types offers a good match for the actual data, suggesting a mix of households with heterogeneous behaviours is more plausible. Using the method of indirect inference, I solve for the mix of types that best fits the CATE between the simulated mix and the actual data using numerical optimization methods. I find a mix of

¹⁹I use 15% as my base case for distance from threshold to trigger the consumption shift. In robustness checks, I test a range of capture distances, from 10-20% and find the presence of *confused* types does not change significantly.

²⁰Using the simulated data, the same empirical methods are applied on each of the household types individually. Results are given in the Appendix.

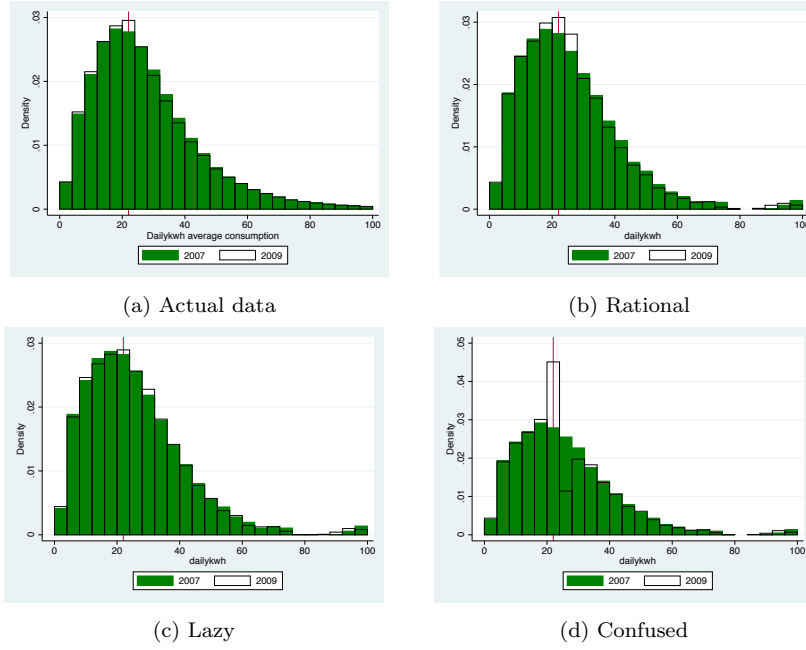


Figure 10. : SIMULATED DISTRIBUTIONS

Notes: Each figure shows two overlapping histograms for simulated distributions in 2007 (green) and 2009 (white). The vertical red line represents the Tier 2 threshold.

85% *lazy*, 7% *rational* and 8% *confused* produces simulated empirical estimates that most closely match the actual estimates. This is presented in Figure 11.

Strikingly, this simulated mix of average cost responders with a small fraction of misperceivers delivers bunching and IV estimates that would lead to the spurious conclusion of marginal cost responsiveness (Tables A2, A3). This speaks to the need to look beyond average treatment effects and consider heterogeneous behaviour by looking at the CATE using the conditional difference-in-differences method. Applying indirect inference reveals behaviour consistent with predominantly average cost responsiveness and a small presence of households misperceiving the nonlinear price schedule having a large aggregate effect.

PRECISION

To test how precisely these household type shares are estimated, I repeat the indirect inference procedure incorporating two sources of variation. The first are the stochastic shocks in the simulation process. The second is variation in the target estimates themselves (the CATE coefficients estimated in Section 3C). For this, I re-estimate the CATE coefficients using a bootstrap procedure. I then draw at random (with replacement) from the bootstrap results to select a new target set of coefficients for each iteration of the indirect inference




Figure 11. : SIMULATED “MIX” VERSUS ACTUAL CATE COEFFICIENTS

Note: Each point represents the difference-in-difference estimate (i.e. the CATE) for each decile of the pre-reform consumption distribution. The blue points represent the actual data, while the yellow points represent the simulated mix. Whisker lines represent standard errors (clustered at the premise level).

procedure.

The results are presented in Figure 12, which is a contour plot of the criterion function value from repeated indirect inference (200 iterations). The vertical and horizontal axes represent the share of *rational* and *lazy* types, respectively. A “valley of minima” is observed running from the top left to bottom right. This highlights that the specific shares of *rational* and *lazy* types is estimated less precisely, whereas their sum—and thus the estimated share of *confused* households—is estimated more precisely at 8%. This makes sense, as the variation between marginal and average price is less stark, and in fact, nearly identical below the threshold, as compared to the large perceived price changes for *confused* types.

V. Policy Implications

Ito (2014) found no such evidence of price misperception. Instead, he produced convincing evidence that Californian consumers respond to average not marginal prices. There are several reasons as to why we may find something different with BC Hydro customers. The first is salience. The BC Hydro RIB contains only one threshold, whereas the Californian tariffs contained up to 5 tiers (4 thresholds). The single BC Hydro threshold is likely much more salient for its customers and, in turn, more susceptible to being misperceived.

The second possible reason is experience. The BC Hydro RIB was introduced during the study period (2008), and as such consumers did not have significant experience in which




Figure 12. : PRECISION OF INDIRECT INFERENCE RESULTS

Notes: This contour plot presents the criterion function value for each iteration of repeated indirect inference for various shares of *rational* and *lazy* types. The dark blue region represents the lowest values of the criterion function. The lowest point (85% L, 6%R) is shown in the middle of the figure, while a “valley of minima” can be observed along the diagonal such that *confused* types are roughly 8%.

to fully understand the tariff. The Californian tariffs pre-existed Ito’s study period. As suggestive evidence that experience is a factor, I calculate the bunching elasticity estimates over several years (listed in Table 5). The price elasticity to *marginal* price obtained from this estimation method declines over time, suggesting the possibility of less *confused* households and a trend towards predominantly average cost responsiveness.

What then are the effects of price misperception? To answer this, I calculate the effect of the RIB policy using the estimate CATE-by-decile coefficients multiplied by the respective pre-reform consumption for each decile. The actual outcome shows a 0.2% increase in consumption in BCH relative to NW. This is driven by low deciles increasing consumption in response to a lower average price (the Tier 1 rate), largely offset by misperception causing a larger-than-expected decrease in the neighbourhood of the threshold.

Using the simulated household types and estimated CATEs, a counterfactual change can be calculated. If all households were responding to marginal price (i.e. *rational* types), the result would be an estimate 0.9% decrease in consumption. Given the decline in aggregate response shown in Table 5, it is reasonable to question whether this is the appropriate counterfactual. If, instead, the population consisted of entirely average cost responders (i.e. *lazy* types), the result would be an estimate 1.05% increase in consumption.

Table 5—: BUNCHING ESTIMATES OF PRICE ELASTICITY BY YEAR

2009	-0.048
2010	-0.035
2011	-0.033
2012	-0.032
2013	-0.020

Notes: Bunching estimates of price elasticity are calculate using the polynomial method for counterfactual construction. Use of the 2007 BCH distribution, either directly or scaled by New Westminster growth factors, exposes the counterfactual to concerns over omitted time-varying factors.

Thus, versus the latter scenario, misperception is actually helping deliver the conservation goal. However, a larger conservation result would be realised if consumers were better able to respond to marginal price, or if the average price simply reflected marginal price by returning to a flat rate structure.

If one expects that over time BC Hydro customers will better understand the tariff and reduce the number of confused households—a reasonable assumption given the observed decline in bunching over time—the result will trend towards weak average cost responsiveness. In other words: lazy, but not confused. Given this, the policy implication is that a simple flat rate is likely to achieve a greater amount of conservation than the two-tier RIB.

Finally, what are the welfare implications of misperception? We can answer this question in terms of the effect on consumer surplus, producer surplus and external social costs.

In terms of consumer surplus, the goal is to calculate the deadweight loss associated with households responding to a misperceived price, and thus consuming too much or too little relative to a “properly optimized” quantity. To calculate this, I simulate two sets of households, one responding based on the estimated mix of types, and the other 100% marginal cost responders. For both sets I use the same stochastic shocks; they differ only in their deterministic response to price. The deadweight loss can then be calculated based on the difference in monthly consumption between the two sets of households according to the following formula (a full derivation is shown in the Appendix):

$$(11) \quad DWL = \frac{1}{2} \frac{1}{\epsilon} PQ(\% \Delta Q)^2$$

The average deadweight loss per household is roughly \$5 per year, or slightly less than 1% of annual expenditure on electricity. However, reiterating the theme of this paper:

the average masks important heterogeneous effects. For *confused* households, the average deadweight loss is \$58 per household, or roughly 10% of expenditure. For average cost responders, the deadweight loss is small, roughly 50 cents per year.²¹

In terms of producer surplus, we return to the aggregate effects on consumption and ask how the change in aggregate consumption affects the marginal cost of supply. In the short run, for such a small change, there is likely to be very little material effect. In the long run, however, the goal of B.C.'s policy was to avoid investment in costlier new supply. In this case, relative to average cost responders, misperception is reducing consumption by close to 1%, or 200 gigawatt-hours per year. If the incremental marginal cost was \$20 per megawatt-hour higher than current supply, this is saving \$4 million per year, or slightly less than \$2 per household—a rather small savings in the context of electricity sales over \$1 billion per year. If the policy were changed to achieve the conservation results of *rational* types, a further 1%, or \$4 million per year could be achieved.

Lastly, conservation may deliver reduction in various external social costs such as avoided greenhouse gas emissions from additional generation. In the case of British Columbia, this may not be a material consideration since the generation mix is nearly 100% zero-GHGs. However, if B.C.'s generation were otherwise offsetting fossil fuel generation in neighbouring regions (or if this analysis were extended to a different region with a dirtier mix of electricity), the savings could be material. At a \$42 per tonne social cost of carbon, reducing coal generation by 1% via conservation works out to roughly \$8 million per year for the province as a whole.

VI. Conclusion

This paper emphasizes the important effect a small fraction of households misperceiving nonlinear prices can have on aggregate outcomes. By combining reduced form empirical methods, exploiting a natural experiment, with a simple structural model of heterogeneous household behaviour, I uncover underlying behavior consisting of a small but important share of households misperceiving the tariff.

The methodological implication is that caution must be heeded when claiming marginal cost responsiveness based on average treatment effects from bunching methods and encompassing tests with panel data. The strong response from *confused* households at the threshold—where marginal price changes are greatest—produces a spurious finding

²¹One could easily question whether average cost responders are indeed reducing their consumer surplus, or if their response reflects higher—but very much real to them—information costs of optimizing. If that were the case, they face no deadweight loss relative to marginal cost responders. It is more difficult, however, to justify price misperceivers in the same manner.

of marginal cost responsiveness.

From a policy perspective, this paper affirms Ito (2014)’s finding that a flat marginal price (rather than a nonlinear tariff) is the better policy choice to achieve greater conservation. In the case of BC Hydro, the presence of households misperceiving the tariff is likely masking an otherwise weak response to the RIB. Over time, as households gain more experience, the misperception behaviour is likely to dissipate, exposing a weaker aggregate response—one dominated by response to average prices.

How consumers respond to nonlinear tariffs has significant implications for policy and rate setting. In the presence of clear informational and attentiveness challenges, policies designed under the assumption of perfect information and optimization could fail in their goals. This study contributes to the literature using a unique dataset and clean methodologic approach combining reduced form and structural methods that allows for a deeper look at important heterogeneous responses. Critically, it demonstrates the important role misperception can play in determining outcomes.

REFERENCES

- Abrevaya, Jason, Yu-Chin Hsu, and Robert P Lieli.** 2015. “Estimating conditional average treatment effects.” *Journal of Business & Economic Statistics*, 33(4): 485–505.
- Bartolome, Charles.** 1995. “Which tax rate do people use: Average or Marginal.” *Journal of Public Economics*, 56: 79–96.
- BC Hydro.** 2014. “Evaluation of the Residential Inclining Block Rate F2009-F2012.”
- BC Hydro.** 2015. “BC Hydro Quick Facts 2015.”
- Blomquist, Soren, and Whitney Newey.** 2017. “The bunching estimator cannot identify the taxable income elasticity.” *NBER Working Paper No. 24136*.
- Borenstein, Severin.** 2009. “To what electricity price do consumers respond? Residential demand elasticity under increasing-block pricing.” Working Paper.
- Chetty, Raj, Adam Looney, and Kory Kroft.** 2009. “Salience and Taxation: Theory and Evidence.” *American Economic Review*, 99(4): 1145–77.
- 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*.
- Davidson, Russell, and James MacKinnon,** ed. 1993. *Estimation and Inference in Econometrics*. Oxford University Press.
- Eissa, Nada.** 1996. “Labor supply and the economic recovery Tax Act of 1981.” In *Empirical foundations of household taxation*. 5–38. University of Chicago Press.
- Fackler, P, and Huseyin Tastan.** 2008. “A framework for indirect inference.” Working paper.
- Gourieroux, C., A. Monfort, and E. Renault.** 1993. “Indirect Inference.” *Journal of Applied Econometrics*, 8: S85–S118.
- Ito, Koichiro.** 2014. “Do consumers respond to marginal or average price? Evidence from nonlinear electricity pricing.” *American Economic Review*, 104(2): 537–563.
- Jamil, Faisal, and Eatzaz Ahmad.** 2011. “Income and price elasticities of electricity demand: Aggregate and sector-wise analyses.” *Energy Policy*, 39(9): 5519 – 5527.

- Levinson, Arik, and Scott Niemann.** 2004. “Energy use by apartment tenants when landlords pay for utilities.” *Resource and Energy Economics*, 26: 51–75.
- McRae, Shaun, and Robyn Meeks.** 2015. “Price Perceptions and Electricity Demand with Nonlinear Prices.” Working Paper.
- Murray, Michael P.** 2005. *Econometrics: A modern introduction*. Pearson Higher Education.
- Nataraj, Shanthi, and W Michael Hanemann.** 2011. “Does marginal price matter? A regression discontinuity approach to estimating water demand.” *Journal of Environmental Economics and Management*, 61: 198–212.
- Saez, Emmanuel.** 2010. “Do Taxpayers Bunch at Kink Points?” *American Economic Journal: Economic Policy*, 2(3): 180–212.
- Sallee, James.** 2014. “Rational Inattention and Energy Efficiency.” *Journal of Law and Economics*, 57(3): 781 – 820.
- Schneider, Ian, and Cass Sunstein.** 2017. “Behavioral considerations for effective time-varying electricity prices.” *Behavioural Public Policy*, 1(2): 219–251.
- Smith, A.** 2008. “Indirect Inference.” *The New Palgrave Dictionary of Economics (2nd edition)*.
- Wichman, Casey J.** 2014. “Perceived price in residential water demand: Evidence from a natural experiment.” *Journal of Economic Behavior & Organization*, 107: 308–323.
- Wichman, Casey J.** 2017. “Information provision and consumer behavior: A natural experiment in billing frequency.” *Journal of Public Economics*, 152: 13–33.

APPENDIX

A1. Balanced dataset preparation

From the raw data, I perform the following operations to prepare a suitable dataset for analysis. First, I consider only households with continuous bill data for the entire duration of the 2005–2013 period. This cuts the number of valid observations significantly, as many premise IDs include only sporadic readings. While this cuts the power of the empirical analysis, I have little reason to believe the remaining premises are systematically different than omitted accounts. Second, I exclude any households with negative consumption values or values that exceed ten times the median bill. Extreme values such as these occur on occasion due to meter misreads or corrections. Excluding households containing these extreme values eliminates less than 300 households. Finally, the individual household bills are not of uniform length, nor do they all follow the same cycle. To deal with the non-uniform length, I calculate a *dailykwh* variable as the total amount consumed divided by the (different) days in each billing cycle. To deal with different billing cycle start/end dates, I use the *dailykwh* variable to create a daily value for each day in the 2005–2013 period and subsequently collapse the daily data back to calendar months. This creates a balanced dataset with standard periods of length and aligned cycles. The number of households and total observations in the balanced dataset are reduced to 34,592 and 3.7 million, respectively.

A2. Derivation of consumer surplus

We start with the assumption of consumers optimizing based on a quasilinear utility function of the following form:

$$\max_{X,Q} U(X, Q) = X + \alpha Q^\beta \text{ subject to } I = X + \tilde{P}Q$$

where X is the composite good, Q is electricity, and \tilde{P} is the perceived price of electricity. The first order conditions deliver the following demand function for electricity:

$$Q = \left(\frac{\tilde{P}}{\alpha\beta} \right)^{\frac{1}{\beta-1}}.$$

Taking logs and letting $\epsilon = \frac{1}{\beta-1}$ gives:

$$\ln Q = -\alpha \ln(\alpha\beta) + \epsilon \ln \tilde{P}.$$

Differentiating with respect to $\ln \tilde{P}$ gives the price elasticity of demand, ϵ :

$$\epsilon = \frac{d \ln Q}{d \ln \tilde{P}} = \frac{\% \Delta Q}{\% \Delta \tilde{P}}.$$

We can now use this constant log elasticity in our estimate of the deadweight loss, for which we use the Harberger triangle approximation:

$$DWL = \frac{1}{2} \Delta Q \Delta \tilde{P}.$$

Noting that for small changes, ΔQ and $\Delta \tilde{P}$ can be approximated by:

$$\begin{aligned} \Delta Q &= \% \Delta Z \cdot Z \\ \Delta P &= \% \Delta \tilde{P} \cdot P \\ &= \frac{\% \Delta Q}{\epsilon} P. \end{aligned}$$

The deadweight loss is thus:

$$\begin{aligned} DWL &= \frac{1}{2} (\% \Delta Q \cdot Q) \left(\% \Delta Q \frac{1}{\epsilon} P \right) \\ &= \frac{1}{2} \frac{1}{\epsilon} PQ (\% \Delta Q)^2. \end{aligned}$$

A3. Tables and Figures

Table A1—: CONDITIONAL DIFFERENCE-IN-DIFFERENCES ESTIMATES

	(1)	(2)	(3)
Post#BCH#1.decile	0.0107*** (0.0024)	0.0119*** (0.0022)	0.0126*** (0.0020)
Post#BCH#2.decile	0.0200*** (0.0024)	0.0195*** (0.0023)	0.0200*** (0.0020)
Post#BCH#3.decile	0.0175*** (0.0025)	0.0180*** (0.0023)	0.0185*** (0.0021)
Post#BCH#4.decile	0.0099*** (0.0026)	0.0095*** (0.0024)	0.0100*** (0.0022)
Post#BCH#5.decile	0.0029 (0.0028)	0.0022 (0.0026)	0.0028 (0.0023)
Post#BCH#6.decile	-0.0142*** (0.0029)	-0.0151*** (0.0026)	-0.0144*** (0.0024)
Post#BCH#7.decile	0.0006 (0.0029)	0.0001 (0.0027)	0.0007 (0.0024)
Post#BCH#8.decile	0.0024 (0.0030)	0.0017 (0.0027)	0.0024 (0.0024)
Post#BCH#9.decile	-0.0016 (0.0030)	-0.0024 (0.0028)	-0.0016 (0.0025)
Post#BCH#10.decile	-0.0007 (0.0032)	-0.0015 (0.0030)	-0.0007 (0.0026)
Year-month FE	×	-	×
Household FE	-	×	×
Observations	3,721,963	3,721,963	3,721,963
R-squared	0.7606	0.7969	0.7780
Number of households	34,591	34,591	34,591

Note: Standard errors (clustered at premise level) shown in parentheses.

Table A2—: BUNCHING ESTIMATES USING SIMULATED DISTRIBUTIONS

Counterfactual	Actual	Simulated Mix	Rational	Lazy	Confused
Polynomial	-0.048 (0.010)	-0.098 (0.032)	-0.024 (0.032)	-0.009 (0.029)	-0.641 (0.104)
2007	-0.041 (0.012)	-0.078 (0.020)	-0.044 (0.016)	-0.007 (0.017)	-0.549 (0.053)
2007 Scaled	-0.045 (0.017)	-0.083 (0.021)	-0.039 (0.018)	-0.003 (0.018)	-0.544 (0.054)

Note: Bootstrapped standard errors shown in parentheses.

Table A3—: IV ESTIMATES USING SIMULATED DISTRIBUTIONS

	Actual			Simulated Mix					
$\Delta \ln MP$	-0.136 (0.007)	· (0.010)	-0.141 (0.010)	-0.133 (0.008)	· (0.011)	-0.137 (0.011)			
$\Delta \ln AP$	· (0.009)	-0.133 (0.009)	0.010 (0.013)	· (0.010)	-0.130 (0.010)	0.006 (0.014)			
	Rational			Lazy			Confused		
$\Delta \ln MP$	-0.131 (0.002)	· (0.004)	-0.125 (0.004)	-0.080 (0.003)	· (0.003)	0.008 (0.003)	-0.151 (0.006)	· (0.018)	-0.313 (0.018)
$\Delta \ln AP$	· (0.003)	-0.133 (0.003)	0.009 (0.004)	· (0.003)	-0.135 (0.003)	-0.142 (0.004)	· (0.005)	-0.067 (0.020)	0.240 (0.020)

Note: Standard errors (clustered at premise level) shown in parentheses.

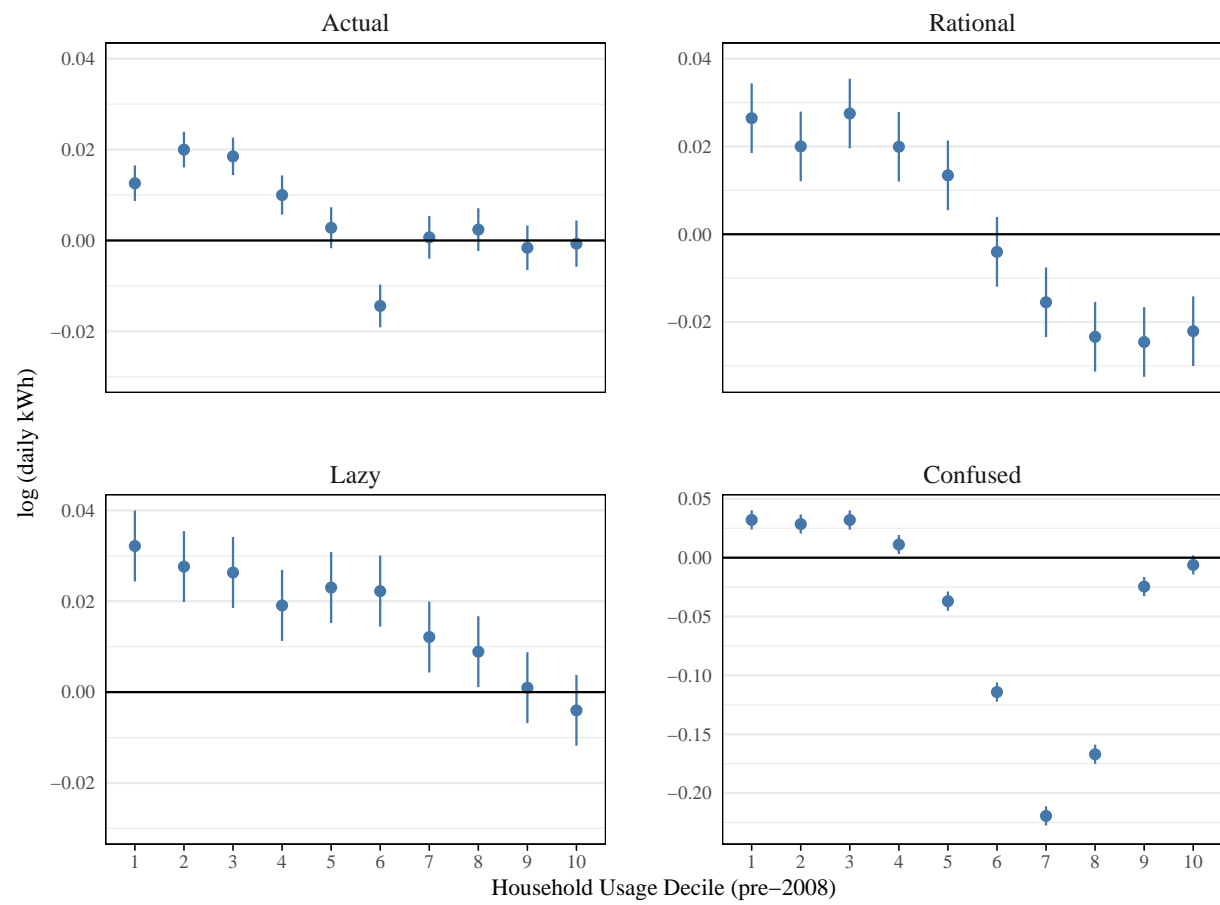


Figure A1. : SIMULATED DIFFERENCE-IN-DIFFERENCE COEFFICIENTS