

Does Energy Consumption Respond to Price Shocks?

Evidence from a Regression-Discontinuity Design

Paulo Bastos

Lucio Castro

Julian Cristia

Carlos Scartascini

The World Bank
Development Research Group
Trade and International Integration Team
February 2014



Abstract

This paper exploits unique features of a recently introduced tariff schedule for natural gas in Buenos Aires to estimate the short-run impact of price shocks on residential energy utilization. The schedule induces a nonlinear and non-monotonic relationship between households' accumulated consumption and unit prices, thus generating exogenous price variation, which is

exploited in a regression-discontinuity design. The results reveal that a price increase causes a prompt and significant decline in gas consumption. They also indicate that consumers respond more to recent past bills than to expected prices, which argues against the assumption that consumers have perfect awareness of complex price schedules.

This paper is a product of the Trade and International Integration Team, Development Research Group. It is part of a larger effort by the World Bank to provide open access to its research and make a contribution to development policy discussions around the world. Policy Research Working Papers are also posted on the Web at <http://econ.worldbank.org>. The authors may be contacted at pbastos@worldbank.org.

The Policy Research Working Paper Series disseminates the findings of work in progress to encourage the exchange of ideas about development issues. An objective of the series is to get the findings out quickly, even if the presentations are less than fully polished. The papers carry the names of the authors and should be cited accordingly. The findings, interpretations, and conclusions expressed in this paper are entirely those of the authors. They do not necessarily represent the views of the International Bank for Reconstruction and Development/World Bank and its affiliated organizations, or those of the Executive Directors of the World Bank or the governments they represent.

DOES ENERGY CONSUMPTION RESPOND TO PRICE SHOCKS? EVIDENCE FROM A REGRESSION-DISCONTINUITY DESIGN^{*}

PAULO BASTOS[†]

LUCIO CASTRO[‡]

JULIAN CRISTIA[§]

CARLOS SCARTASCINI^{**}

JEL classification: L95, D12, L51.

Keywords: Energy consumption, elasticity of demand, regulation of public utilities, regression discontinuity design.

Sector Board: EPOL

^{*} We would like to thank three anonymous referees for very helpful comments and suggestions. We are also grateful to seminar participants at the Inter-American Development Bank and the Northeast Workshop on Energy Policy and Environmental Economics, and in particular to Matías Busso, Omar Chisari, Sebastian Galiani and Howard Chong for thoughtful discussions. Mauricio Cordiviola (tariff manager), Hernan Maurette and Jorge Montanari from the Public Affairs Department, and their teams at MetroGAS S.A., provided valuable cooperation and guidance in the extraction and processing of the company's customer data, as well as in distilling the information contained in the data and in official documents regarding the tariff changes. Gastón Astesiano and Ramón Espinasa provided insightful comments on the overall project. María Antonella Mancino provided superb assistance to this project, and Margherita Calderone and Melisa Iorianni collaborated at different stages of the process. Results have been screened to ensure that no confidential customer data are revealed or could be retrieved. The opinions expressed herein are those of the authors and do not necessarily reflect those of the institutions they are affiliated with.

[†] Development Research Group, World Bank, 1818 H Street NW, Washington DC, 20433, United States.

e-mail: pbastos@worldbank.org

[‡] Centro de Implementación de Políticas Públicas para la Equidad y el Crecimiento (CIPPEC), Av. Callao 25, 1B - C1022AAA Buenos Aires, Argentina.

e-mail: lcastro@cippec.org

[§] Corresponding author. Research Department, Inter-American Development Bank, 1300 New York Avenue NW, Stop E1007, Washington DC, 20577, United States.

e-mail: jcristia@iadb.org

^{**} Research Department, Inter-American Development Bank, 1300 New York Avenue NW, Washington DC, 20577, United States.

e-mail: carlossc@iadb.org

I. INTRODUCTION

Suppose that energy prices experience a shock. Does energy consumption respond? How much and how promptly? These are key questions in the study of a wide range of macroeconomic, regulatory and environmental issues, such as the transmission channels of energy price shocks, optimal taxation and pricing policies in energy markets, and interventions to address climate change. Naturally, economists have a long-standing interest in estimating consumption responses to price changes in energy markets.¹ Progress towards this aim has been complicated by an important identification challenge, however. Since consumers typically experience the same events at essentially the same time, it has been difficult to construct the equivalent of randomly assigned treatment and control groups and thereby ground the estimated price-elasticities on a well-defined counterfactual (Reiss and White [2008]).

In this paper, we exploit unique features of a recently introduced tariff schedule for residential natural gas in Greater Buenos Aires to estimate the short-run impact of price shocks on residential gas utilization.² The revised schedule induces a non-linear and non-monotonic relationship between annual previously accumulated consumption and unit prices, thus giving rise to exogenous price variation. In particular, the introduction of a threshold for defining unit prices based on previously accumulated consumption approximates a randomly assigned price differential for a large number of consumers on each side of the threshold, allowing us to build treatment and control groups to estimate the impact of interest. We estimate the demand effect of a price shock using a regression discontinuity (RD) design whereby the gas consumption levels of households situated barely above the tariff discontinuity are compared with those of households located barely below.

We implement the RD design on administrative records of the natural gas distribution company, gathering longitudinal data on the unit prices and consumption levels of each consumer located near the tariff discontinuity – the same information that these households received in their utility bills. We find that an increase in the average price of natural gas in the utility bill received by consumers induces a statistically significant and prompt decline in residential gas utilization: a 25% price increase reduced residential gas consumption by 3.8% in the subsequent two-month period. This result suggests that policy interventions via the price mechanism may constitute a powerful instrument to influence the patterns of residential energy utilization, even in a relatively short time span.³

An important feature of our research design is that it exploits the specific information set available to consumers to estimate the effect of interest. In particular, because households are reclassified every billing period on the basis of their annual accumulated consumption, fully informed consumers in the treatment and control group face the same expected prices moving forward. However, we provide survey evidence suggesting that consumers have highly imperfect knowledge about the price determination mechanism, and appear to infer prices from recent past bills. In addition, we offer strong statistical evidence that consumers do not manipulate strategically their annual accumulated consumption. For these reasons, the resulting estimates are especially relevant for residential energy markets characterized by ex-post billing where

households infer changes in prices from the utility bill. Importantly, while it has long been emphasized that this feature of residential energy markets plays an important role in shaping consumption responses to price changes (Shin [1985]), there is still little direct evidence on whether and how promptly energy consumption responds to price variations inferred from utility bills.

We complement and extend several strands of literature. A number of studies on the price-elasticity of energy demand employ time series methods using data on energy prices and aggregate energy consumption (Liu and Lin [1991]; Krichene [2002]; Bushnell and Mansur [2005]; Fezzi and Bunn [2010]). A related body of work draws on cross-sectional survey data, including influential papers by Parti and Parti [1980], Dubin and McFadden [1984], Dubin [1985] and Reiss and White [2005]. While these methods allow for the estimation of long-term impacts, the aggregated or cross-sectional nature of the data imposes relatively strong identifying assumptions. Furthermore, estimates yielded by cross-sectional data are, by construction, silent on the speed with which energy consumption adjusts to price shocks, an issue that is of key interest in a variety of policy contexts.

Another strand of research estimates price-elasticities in the context of tariff field experiments, including early work by Hausman et al. [1979], Acton and Mitchell [1980], Caves and Christensen [1980] and Parks and Weitzel [1984]. Whereas this approach addresses some limitations of the time-series and cross-sectional evidence, it has been criticized on the ground that the (most often voluntarily-selected) set of participants are thoroughly informed about price changes at the outset, generating an informational context that differs significantly from real-world situations in which households learn about price changes from utility bills or the press (Acton [1982]; Reiss and White [2008]).

Two recent papers using disaggregate billing data on electricity consumption from California are perhaps the closest to our own. Reiss and White [2008] examine how price shocks and conservation appeals impact residential electricity consumption. Their estimates point to sizable short-run impacts on energy utilization. In independent work, Ito [2013] exploits a spatial discontinuity in electric utility service areas in southern California, which leads to nearly identical households experiencing different nonlinear price schedules. His contribution is highly complementary to ours. Consistent with our results for the natural gas market, he finds that residential consumers in the electricity market respond to (lagged) average price, rather than marginal or expected marginal price.⁴

The remainder of the paper is structured as follows. Section II provides background information on the market for natural gas in Greater Buenos Aires and describes the data employed. Section III describes the research design and provides important complementary evidence from a survey of consumers located near the tariff discontinuity of interest. Section IV presents the econometric results. Section V provides a discussion of the results in the context of the literature. Section VI offers some concluding remarks.

II. BACKGROUND AND DATA

We focus on the residential market for natural gas in Greater Buenos Aires. In late 2008, the national energy regulatory agency (ENARGAS) introduced the first increase in residential gas tariff since 2002. The revised schedule was composed of eight new tariff groups, each facing different variable fees per cubic meter; see Table I. This schedule introduced significantly higher unit prices for those consumers with higher levels of annual accumulated consumption. We exploit the resulting discontinuity in variable prices at the annual consumption level that divides categories R32 and R33 to examine consumption responses to price shocks.

"Place Table I approximately here."

Our empirical analysis draws on administrative records from MetroGAS S.A., one of the largest residential gas distributors in Argentina. The company has a client base in Greater Buenos Aires of about 2.5 million residential households, who receive their gas bills every two months. The administration of MetroGAS agreed to provide us with data on a small share of its customer base around the discontinuity R32-R33.⁵ Given this constraint, we have defined the sample with the view to optimize the implementation of the RD design for this discontinuity. Specifically, consumers were selected into the sample provided by the company if: (1) they had a residential bill issued in May 2009 with an annual accumulated consumption between 1480 and 1520 cubic meters; and (2) they had been customers for at least six bimonthly cycles by that month. The resulting estimation sample contains 7190 households.

This sample is composed of longitudinal records corresponding to the bills issued in May 2009, in the five previous bimonthly periods, and in the three subsequent ones. For each period, the data set comprises information on the amount billed, quantity of gas consumed (in cubic meters), type of reading (measured or imputed), category assigned to the consumer, and the exact dates of reading and issuance. It further contains information on the region and neighborhood of residence of each consumer.

III. RESEARCH DESIGN

An ideal experiment designed to estimate the impact of a price shock on residential energy consumption would randomly assign some consumers to a treatment group, facing price P_H , and other consumers to a control group, facing price P_L . Unfortunately, a large-scale experiment of this kind has yet to be implemented, making the task of estimating this behavioral response rather difficult. To approximate such an ideal experiment, we exploit unique features of the price schedule described in Section II, along with the specific information set available to consumers.

In May 2009, households with annual accumulated consumption above 1500 cubic meters were assigned a unit price roughly 25% higher than those with an annual accumulated consumption barely below this level. This discontinuity of the unit price schedule makes it possible to apply a RD design in which the outcome variable corresponds to the consumption

level in the subsequent two-month period and the running variable to the annual accumulated consumption.

However, as we explain in detail below, the interpretation of the RD design in this application is made difficult by two important features of the price determination mechanism: (1) the category to which consumers are assigned, and hence the unit price they are charged, is determined by the accumulated consumption of the previous 12 months; and (2) the categorization of consumers is revised every two months, in line with the variation of the 12-month accumulated consumption over that period.

III(i). *The price determination mechanism*

Let us define the key variables underlying the determination of the amount billed in a given bimonthly period 0. The annual accumulated consumption AAC_0 corresponds to the sum of the consumption level C in period 0 and in the five previous bimonthly periods:

$$(1) \quad AAC_0 = \sum_{j=-5}^0 C_j$$

The unit price in period 0 is a function of whether accumulated consumption is above or below a given threshold:⁶

$$(2) \quad P_0 = \begin{cases} P_L & \text{if } AAC_0 \leq 1500 \\ P_H & \text{if } AAC_0 > 1500 \end{cases}$$

The total bill B in period 0 can be expressed as:

$$(3) \quad B_0 = FC + P_0 C_0 + \mu_0$$

where FC is a fixed cost, P_0 is the unit price in that period, and μ_0 is an idiosyncratic shock capturing the fact that the bill received by consumers sometimes contains idiosyncratic adjustments and retroactive charges (e.g., taxes and other charges defined by the regulator in a rather ad hoc manner).

Finally, while consumers may target consumption levels in the following period, they are unable to perfectly control their gas consumption patterns. Hence, actual consumption will differ from targeted consumption (CT) by a random shock. That is:

$$(4) \quad C_1 = CT_1 + e_1$$

In our setting, period 0 corresponds to that billed in May 2009. Consumers in the treatment group are those with annual accumulated consumption barely above 1500 cubic meters in period 0, while consumers in the control group are those with annual accumulated consumption barely below this level. Whether or not we would expect this price shock to have a differential impact

on gas consumption in period 1 crucially depends on the specific information set held by consumers.

Since households are reclassified every period on the basis of their annual accumulated consumption, fully informed consumers in the treatment and control groups face the same expected price moving forward. Hence, if perfectly informed, both groups would have essentially the same incentive to restrain consumption so as not to surpass the 1500 cubic meters threshold in period 1, despite the fact that the bill received in period 0 contained sharp differences in unit prices. However, in the light of the well-documented prominence of information imperfections in residential energy markets with ex-post billing (Shin [1985]), and considering the complexity and novelty of the price determination mechanism in the residential market for natural gas of Greater Buenos Aires, we would expect consumers to be imperfectly informed. To provide further evidence on the information set held by consumers, we have surveyed a subset of households in the estimation sample.

III(ii). *Survey evidence on consumers located near the discontinuity*

We have administered a telephone survey to 353 households from the estimation sample. The sub-sample surveyed was stratified by district to ensure an adequate geographical representation. The survey was conducted in September 2010 and targeted the member of the household that was responsible for paying the gas bill.⁷

The survey questionnaire consists of two sections. The first section collected information on basic socioeconomic characteristics of the household head (age, education, occupation) and housing conditions. The second section explored knowledge about the amount billed and payment method, and assessed the extent to which consumers read their utility bills. In addition, this section collected extensive information on perceived and objective knowledge on how the amount billed is computed; specifically, on: (1) how frequently tariffs are determined; (2) the past consumption periods used for their determination; and (3) the consumption threshold that leads to a higher unit price.

"Place Table II approximately here."

Table II reports summary statistics on surveyed households. For comparison, it also provides summary statistics from the national Household Survey (Encuesta Permanente de Hogares) on households living in Greater Buenos Aires. Relative to the average resident, surveyed household members from our estimation sample are more likely to be female and married. They are also more likely to have tertiary education, and tend to be considerably older than the average resident. These differences are likely to reflect the fact that the survey targeted the household member that was responsible for paying the gas bill. The fact that surveyed households are more likely to have tertiary education might also suggest that they tend to have higher-than-average income. Table II further suggests that surveyed households are more likely to be homeowners, and have a slightly larger number of rooms and different families. This evidence is consistent

with the fact that our estimation sample focuses on households of relatively high gas consumption. Further data from the questionnaire suggests that natural gas is the dominant source of energy used by these households: 76.1% of households use it for space heating, 95.2% use it for water heating, and 99.4% use it for cooking.

"Place Table III approximately here."

Table III reports the key results from the survey.⁸ Nearly 92% of households reported that they were able to remember the amount charged in the last gas bill. The proportion of households who paid their bill by direct debit is relatively small (14%), which alleviates the concern that consumers might not be aware of how much they are charged every period. About 75% of consumers stated that they regularly read their gas bills.

However, knowledge about the price determination mechanism proved to be almost non-existent. Among surveyed households, 31% stated that they know how the price is determined. However, the questions aimed at assessing precise knowledge of the price determination mechanism suggest that the proportion of well-informed consumers is considerably lower. First, only 14% of households knew that consumers are re-categorized (and unit prices are determined) in each billing cycle. Second, 39% of consumers knew that their billing category is determined on the basis of the accumulated consumption over the previous year. Third, only 4% of consumers knew that the threshold that divides categories R32 and R33 is 1,500 cubic meters. Overall, less than 1% of households provided correct answers to the three objective questions posed.

In summary, the survey reveals that consumers tend to know how much they are paying for their gas consumption, but have scant information about the actual pricing scheme. Consequently, in the remainder of this paper, we will assume that the vast majority of consumers have imperfect information about the prevailing price determination mechanism and infer prices from the utility bill.

III(iii). *Econometric model*

Under the assumption that most if not all consumers have imperfect information about the price determination mechanism, we can estimate the short-term effects of price shocks by applying a sharp RD design. That is, we can compare gas utilization in period 1 for consumers that in period 0 had annual accumulated consumption barely above and below the 1500 cubic meters threshold, as both these sets of consumers are expected to be very similar along observed and unobserved characteristics but experienced very different unit prices. Since we can reasonably assume that households infer prices from recent past bills, differences in consumption in period 1 across both groups of consumers can be interpreted as the short-run behavioral response to the price shock.

To implement the RD design we estimate the following regression model:

$$(5) \quad C_{i,1} = \beta_0 + \beta_1 Treatment_{i,0} + f(\overline{AAC}_{i,0}) + \omega_{i,0}$$

where $C_{i,1}$ corresponds to consumption in period 1 for consumer i . $\overline{AAC}_{i,0}$ corresponds to the running variable, the normalized annual accumulated consumption in period 0. That is:

$$(6) \quad \overline{AAC}_{i,0} = AAC_{i,0} - 1500$$

The treatment variable is a binary indicator of whether individual i in period 0 was assigned the higher unit price. It is determined as:

$$(7) \quad Treatment_{i,0} = \begin{cases} 0 & \text{if } \overline{AAC}_0 \leq 0 \\ 1 & \text{if } \overline{AAC}_0 > 0 \end{cases}$$

Parameter β_1 in (5) captures the average effect of barely surpassing the threshold, and hence having received a substantially larger utility bill, once we flexibly control for the running variable, \overline{AAC}_0 . The intuition behind this approach is that all observable and unobservable variables should evolve smoothly around this threshold, and hence any jump in consumption in period 1 can be attributed to the discontinuous increase in the amount billed. In other words, in a valid RD design, a key assumption is that observations are randomly distributed between the treatment and control group in a local neighborhood of the threshold (Lee and Lemieux [2010]). In our setting, this assumption seems plausible for three reasons. First, as documented above, consumers have highly imperfect knowledge about the price determination mechanism. Second, they tend to be unaware of the location of the relevant threshold. Third, even if they had perfect information about the pricing scheme, it would be difficult and time-consuming – if at all possible – to precisely manipulate gas utilization so as to avoid passing this cutoff. Indeed, this would require obtaining access to gas readings, forecasting factors affecting future demand, and knowing exactly when the billing period ends (i.e. the exact timing of the meter reading).

We focus on behavioral reactions during the period of mid-May to mid-July 2009, as these two winter months account for the bulk of annual gas consumption (about 30% of the total). It would have been possible to examine consumption patterns of these households in later billing cycles. However, beginning in June 2009, public reactions motivated by the increase in gas prices generated a disruption in the normal billing process of the company. This generated differences in the timing of bill issuance for individuals on both sides of the cutoff, making it difficult to disentangle between timing effects and those stemming from the price shock.

Though we would like to compare observations just above and below the threshold, in practice a larger window around the cutoff has to be used to obtain precise estimates. In choosing the width of this window, researchers typically face a trade-off between bias and precision: a wider window provides greater precision but at the expense of expected higher bias. In our setting, because the utility company had records on almost two million consumers, we were able to select a narrow window around the threshold and still have a sizable sample size of 7200 consumers.

Our baseline specification consists of a local linear regression that controls for normalized annual accumulated consumption as expressed in (6). Although in general it is recommended to control for the running variable to minimize the potential bias (Imbens and Lemieux [2008]), in this particular application the close relationship between the outcome and the running variable justifies following this approach. For robustness, we additionally report estimates: (1) reducing the window in the running variable used to select consumers; (2) controlling for higher order polynomials of the running variable; and (3) allowing for a differential slope between the outcome and the running variable on both sides of the cutoff.

We cluster standard errors by the running variable, as suggested by Lee and Card [2008] for cases in which this variable is discrete.⁹ It has been pointed out that when the number of clusters is small, standard errors may be downward biased (Angrist and Lavy [2009]). In our baseline specification, there are 41 clusters which may be a large enough number (Cameron, Gelbach and Miller [2008]). Nonetheless, this problem becomes more serious when the number of clusters is reduced, as a narrower window in the running variable is used. However, the results are robust to two proposed solutions to this problem: (1) computing averages of the relevant variables by clusters and reproducing the analysis at this level; and (2) selecting the highest between robust-clustered and conventional standard errors (Angrist and Pischke [2008]).

IV. RESULTS

IV(i). *Testing the validity of the research design*

The basic identifying assumption of the RD design is that the outcome variable would have been continuous at the assignment threshold in the absence of the treatment (Lee [2008]). Albeit this assumption cannot be tested directly, we provide evidence on this issue by examining whether a number of covariates are continuous at the threshold. We define a treatment group composed of consumers in the [1,20] cubic meters of normalized annual accumulated consumption by May 2009 and a corresponding control group for those in the [-20,0] cubic meters range. Given that a small bandwidth is used to select these two groups of consumers, differences in the running variable between them are minimal, as the average difference in this variable is only 20 cubic meters – which represent only 1.3% of the mean (20/1500). Hence, as a first approximation it is possible to compare average values between the treatment and control groups to inspect for evidence in favor of the identifying assumption. We additionally run local linear regressions to test for the existence of jumps in covariates at the threshold once we control for differences in normalized annual accumulated consumption.

"Place Table IV approximately here."

Table IV presents results for key dates and period lengths. Panel A shows that the timing of events is very similar between the treatment and control groups, while Panels B and C document that the length of critical periods is almost identical across groups. Similar patterns emerge when

we test for jumps in these variables, by regressing them on the treatment dummy while controlling for the normalized annual accumulated consumption in period 0 (Column 4). Results not reported (but available on request) show that there are significant differences between the analyzed dates and period lengths across regions, suggesting the relevance of exploring balance in timing patterns. Inspecting balancing in period lengths for Period 1 is critical to attribute differences in total consumption during that period to changes in consumer behavior. The results provide evidence that actual gas consumption is recorded every two months. Hence consumption reported in the administrative records corresponds to actual consumption and not to imputations by the firm. Moreover, gas bills are issued approximately one week after the final measurement for the period and should be received by consumers approximately 10 days after a period ended, according to the firm. This leaves about 50 days for consumers to react.

"Place Table V approximately here."

In Table V we examine differences in the geographic distribution across groups. Though in general the results point to adequate balancing along this dimension, in 4 out of 10 regions we find statistically significant differences at the 10% level between the treatment and control groups. However, the results show that only in one case there are statistically significant differences at the 10% level once we control linearly for normalized annual accumulated consumption (adjusted difference column).

"Place Table VI approximately here."

In Table VI we analyze differences in historical consumption patterns. Panel A presents raw and adjusted differences between consumption in periods -5 to 0. By construction, the annual accumulated consumption of the treatment group will be higher than that of the control group. Hence the existence of some statistically significant differences when analyzing raw differences is not surprising. Though expected, these findings highlight the importance of adjusting for the running variable. When doing this, we tend not to observe statistically significant differences across groups. Panel B presents a cleaner test of pre-treatment differences between groups by comparing consumption in each period as a share of the total annual consumption. In this dimension, the treatment and control groups present strikingly similar patterns, suggesting once more that the RD design is able to yield unbiased estimates in this context.

"Place Figure 1 approximately here."

Figures 1 to 3 depict the results presented in Tables IV to VI. The same patterns highlighted in the tables clearly stand out from these figures: the covariates considered are smooth around the discontinuity.

"Place Figures 2 and 3 approximately here."

Importantly, Table VII and Figure 4 clearly show that the average amount billed in period 0 slightly increases as normalized annual accumulated consumption rises but jumps drastically when the latter crosses the cutoff.¹⁰

"Place Table VII and Figure 4 approximately here."

It has been stressed in the RD design literature that this approach will not be suitable if agents can manipulate the running variable, implying that the condition that individuals on both sides of the discontinuity are similar is not fulfilled (McCrary [2008]; Lee and Lemieux [2010]). For the reasons outlined above, we would expect the scope for this manipulation to be limited in our setting. Nevertheless, to explore this issue further we follow McCrary [2008] and examine the density distribution of the running variable, in particular whether there is a jump in this density around the threshold. Figure 5 shows that the density is quite flat and does not point to the existence of any discontinuity around the threshold.

"Place Figure 5 approximately here."

IV(ii). *Estimating the short-run impacts of the price shock*

Main results. Given the evidence confirming the validity of the research design, we now turn to the primary focus of the paper: the impact of the price shock on gas consumption in the subsequent billing period. Table VIII presents the results. In specification (1), we regress gas consumption in period 1 on the treatment dummy, while controlling linearly for normalized annual accumulated consumption. The results reveal that experiencing a price shock induces a statistically significant drop in gas consumption of 15.9 cubic meters in the subsequent period (or roughly 3.8% of the average gas consumption).

"Place Table VIII approximately here."

Figure 6 depicts these results. There is a clear positive relationship between consumption in period 1 and normalized annual accumulated consumption in period 0, as would be expected given that consumers with higher consumption in the past also tend to consume more in the future. But, most importantly, gas consumption seems to fall discontinuously at the threshold, suggesting that households react to the price shock by significantly reducing consumption in the subsequent two-month period. The estimated effect is sizable if considering that, given the short time span, it is unlikely that consumers will adjust to the new price via investments in more efficient appliances or improvements in insulation. Moreover, since consumers typically learn about the new price approximately 10 days after the beginning of the period, they have only about 50 out of approximately 60 days to adjust to the price shock.

"Place Figure 6 approximately here."

Results from other specifications show that these estimates are quite robust. In specifications (3), (5) and (7) the control function becomes progressively more flexible as we add second, third and fourth order terms. The estimated coefficient remains remarkably stable, hovering between 15.2 and 16.5. In all cases, the results are statistically significant at least at the 10% level, although as expected the standard errors are larger in more flexible specifications. The remaining columns present the corresponding estimates when allowing for differential shapes of the control function at both sides of the cutoff. Once again, the estimated coefficient remains robust, although its precision falls markedly in more flexible specifications.

"Place Table IX approximately here."

Table IX presents further evidence on the robustness of the results as we use an increasingly narrow bandwidth, thus restricting our attention to observations progressively closer to the threshold. Although the coefficients become less precisely estimated when restricting to observations located closer to the cutoff, the estimated impact remains virtually unchanged. Table A.I in the Appendix further shows that the estimated impacts are robust to the inclusion of neighborhood fixed-effects.

Alternative hypotheses. In addition to the issues addressed in the previous section, a potential threat to the validity of our estimates is that they might be partially driven by mean reversion. If gas consumption does not follow a random walk, households who experience a positive demand shock in period 0 might be expected to reduce consumption in period 1 relative to identical households with a negative or zero consumption variation in period 0.¹¹

To account for this hypothesis, in Table X we examine differential consumption patterns in period 1 for consumers located just above and just below two placebo AAC thresholds (where consumers did not experience a price shock). If negative serial correlation were to explain our estimates, we would expect to observe significant differentials in consumption patterns around other AAC thresholds. But the estimates in Table X show that this is not the case. We find no differential consumption patterns across consumers located just above and just below such fake thresholds, while we do confirm the results relative to the threshold of interest.

"Place Table X approximately here."

As mentioned above, an important concern about our analysis is that the estimated treatment effects might partially reflect strategic behavior of households in the treatment group. In particular, rather than just reacting to the price shock, these households might have an incentive to reduce consumption in order to fall below the threshold of interest in the subsequent billing

period. Under this alternative scenario, those consumers in the treatment group that are located just above the threshold should be expected to reduce consumption by less than those consumers in the treatment group located farther away from the threshold. To examine this hypothesis, we divide the treatment group into two different sets of consumers: a "Close" group, including those consumers with AAC in period 0 between 1501 and 1510 cubic meters; and a "Far" group encompassing those consumers with AAC in the 1511 to 1520 interval. The control group continues to be composed of consumers with AAC in Period 0 between 1480 and 1500 cubic meters. To fall below the thresholds of interest, individuals in the "Close" group need to reduce their consumption in Period 1 (compared to Period 0) by 5 cubic meters, whereas those in the "Far" group need to reduce it by 15 cubic meters, on average. If consumers in the treatment group were to behave strategically, both the "Close" and "Far" groups would be expected to reduce consumption in Period 1, but the latter group would be expected to reduce it by 10 cubic meters more on average. Alternatively, if consumption responses were mainly driven by the fact that all consumers in the treatment group experienced a price shock, both the "Close" and "Far" groups would be expected to decrease consumption by a similar amount. To distinguish between these alternative hypotheses, we estimate an equation of the form:

$$(8) \quad C_{i,1} = \beta_0 + \beta_1 Close_{i,0} + \beta_2 Far_{i,0} + f(\overline{AAC}_{i,0}) + \omega_{i,0}$$

where $Close_{i,0}$ and $Far_{i,0}$ have the meanings described above. Table XI documents that the reduction in consumption for the "Close" and "Far" groups, as compared to the control group, is very similar (-16.3 versus -18.5) providing little support to the hypothesis that consumption responses are motivated by strategic behavior.

"Place Table XI approximately here."

V. DISCUSSION

While the quasi-experimental setting we adopt provides unique features for examining short-run consumption responses to price shocks, some caution is needed in extrapolating to other situations. An important drawback of the RD design is that the results are local, in that they refer to the particular threshold studied, and may not generalize to the other points in the distribution of the running variable. This caveat applies to the present analysis. In the specific market we study, households with annual accumulated consumption around the 1500 cubic meters threshold have above average gas consumption. The extent to which the results generalize to other consumers is an open question. Households with lower levels of annualized consumption may present higher price sensitivity, if they are lower-income consumers and, therefore, more price-conscious. On the other hand, these consumers may be less sensitive to price changes, as lower baseline gas utilization may signal a low weight of this good in their overall consumption basket. Nevertheless, the study of price sensitivity for high-consumption households may be interesting in its own right, as they account for a sizable share of overall gas consumption.

A second limitation of our analysis is that the estimated impact refers to the consumption response to a positive price shock. It is difficult to assess the extent to which the results would generalize to policy interventions inducing a fall in the price faced by consumers – e.g. binding price caps and subsidies. Indeed, existing research suggests that consumers react differently to price increases and decreases (see, e.g. Dargay [1993]).

A third caveat is that, while the tariff schedule we exploit offers a unique ground for examining consumption responses to price shocks, caution is needed in extrapolating our estimates to other price schemes. Upon exceeding the 1500 cubic meter threshold, households face both a discontinuous price shock and an increase in the marginal cost of each additional unit consumed. Furthermore, as discussed above, perfectly informed consumers could reverse this discontinuous shock on future bills with a small decrease in consumption. It is important to emphasize that this setting differs clearly from graduated price schemes, where consumers face a higher marginal cost for all consumption above a threshold but a lower rate for consumption below the threshold. Our estimates are potentially more informative on how consumers might be expected to respond to an across-the-board rate increase, in which the price of all units of gas (marginal and inframarginal) increases. However, valid extrapolation to those situations hinges on households believing that their efforts to reduce consumption have only a marginal effect, as opposed to a discontinuous effect, on subsequent bills. Although the evidence we provide suggests that consumers are generally unaware of (or do not understand well) the relevant threshold, and also that they do not manipulate strategically their gas consumption around it, we cannot exclude the possibility that (some) households realize that a marginal effort to reduce consumption might reduce discontinuously their subsequent gas bill.

With these caveats in mind, our paper offers solid evidence that prices do matter for energy consumption in energy markets, even over short-run durations. As Reiss and White [2008, pp.654] emphasize, this fact is important as it has not been widely recognized by policy makers, in part due to the absence of clear and unambiguous evidence of such behavior. It is therefore interesting to compare our estimates to those obtained in recent research on consumption responses to price changes in energy markets. At about -0.15, the short-run estimated elasticity we report is in line with recent estimates for the natural gas market.¹² Drawing on state-level panel data for the United States, Bernstein and Griffin [2006] estimate a short-run elasticity estimate of -0.12. Also for the United States, but using disaggregated panel data at the customer level in an instrumental variables approach, Davis and Muehlegger [2010] report an estimated price elasticity of demand of -0.28. Our estimates are also similar to those of recent studies using disaggregated data for the residential electricity market in California. Focusing on San Diego households, Reiss and White [2008] find that a price increase of 130% induced a consumption decline of 13% in about 60 days. Ito [2013] exploits a spatial discontinuity in service areas in southern California, which leads to nearly identical households experiencing different nonlinear price schedules, and estimates a short-run elasticity with respect to lagged average prices of -0.11.¹³

Another important conclusion from our analysis is that the way in which customers process information about complex tariff structures is an important driver of their behavior. Indeed, we provide survey evidence that consumers have highly imperfect knowledge about the price determination mechanism, and appear to infer prices from past utility bills. In addition, we offer strong statistical evidence that households do not manipulate strategically their gas consumption around the relevant threshold, despite strong incentives for doing so under the prevailing price determination mechanism. The evidence we provide is therefore especially relevant for residential energy markets characterized by complex tariff structures and ex-post billing. While it has long been stressed that imperfect information in residential energy markets plays an important role in shaping consumption responses to price changes (Shin [1985]), there is still little direct evidence on consumer behavior in the presence of complex price schedules and ex-post billing. In this regard, our findings are consistent with (and complementary to) those of recent studies by Bushnell and Mansur [2005] and Ito [2013]. Examining the electricity market in San Diego, Bushnell and Mansur [2005] provide evidence that customers respond more to outdated prices from their last bill than to current market conditions. Ito [2013] finds that consumers respond to lagged average prices rather than marginal or expected marginal prices when faced with nonlinear electricity price schedules. Taken together this evidence suggests that, in order to precisely estimate the impacts of complex price schedules on energy utilization, policy makers and regulators need to move beyond the assumption of perfect awareness of such schedules by consumers.

VI. CONCLUDING REMARKS

We have exploited unique features of the tariff schedule for natural gas in Greater Buenos Aires, along with survey evidence on the specific information set possessed by consumers, to estimate the short-run effect of price shocks on residential gas consumption. The revised tariff schedule induced a non-linear and non-monotonic relationship between annual accumulated consumption and unit prices, thus generating exogenous price variation. Drawing on administrative records on the utility bills of residential consumers, we have estimated the short-run consumption response to a price shock using an RD design whereby two-month consumption levels of households situated barely above an important tariff discontinuity are compared with those of consumers located barely below – hence focusing on a large group of relatively homogeneous consumers facing sizable differences in prices.

We provide evidence that a price increase in the utility bill received by consumers causes a prompt and significant decline in gas consumption. We also show that customers appear to respond more to recent past bills than to expected prices moving forward. Our estimates therefore suggest that policy interventions via the price mechanism are powerful instruments to influence residential energy utilization patterns, but argue against an assumption of perfect awareness of complex price schedules by consumers.

REFERENCES

- Acton, J., 1982, 'An Evaluation of Economists' Influence on Electric Utility Rate Reforms,' *American Economic Review*, 72, pp. 114–199.
- Acton, J. and Mitchell, B., 1980, 'The Effect of Time-of-Use Rates in the Los Angeles Electricity Study,' RAND Corporation Report No. N-1533-DWP/HF. Santa Monica, United States: RAND Corporation.
- Angrist, J. and Lavy, V., 2009, 'The Effect of High Stakes High School Achievement Awards: Evidence from a Group-Randomized Trial,' *American Economic Review*, 99, pp. 1384–1414.
- Angrist, J. and Pischke, J., 2008, *Mostly Harmless Econometrics: An Empiricist's Companion*, (Princeton University Press, New Jersey, United States).
- Bernstein, M. and Griffin, J., 2006, 'Regional Differences in the Price-elasticity of Demand for Energy,' National Renewable Energy Laboratory, Midwest Research Institute.
- Bushnell, J. and Mansur, E., 2005, 'Consumption Under Noisy Price Signals: A Study of Electricity Retail Rate Deregulation in San Diego,' *Journal of Industrial Economics*, 53, pp. 493–513.
- Cameron, C.; Gelbach, J. and Miller, D., 2008, 'Bootstrap-Based Improvements for Inference with Clustered Errors,' *Review of Economics and Statistics*, 90, pp. 414–427.
- Caves, D. and Christensen, L., 1980, 'Econometric Analysis of Residential Time-of-Use Electricity Pricing Experiments,' *Journal of Econometrics*, 14, pp. 287–306.
- Davis, L. and Muehlegger, E., 2010, 'Do Americans Consume too Little Natural Gas? An Empirical Test of Marginal Cost Pricing,' *RAND Journal of Economics*, 41, pp. 791–810.
- Dargay, J., 1993, 'Demand Elasticities: A Comment,' *Journal of Transport Economics and Policy*, 27, pp. 87–90.
- Dubin, J., 1985, *Consumer Durable Choice and Demand for Electricity*, (North-Holland, Amsterdam, Netherlands).
- Dubin, J. and McFadden, D., 1984, 'An Econometric Analysis of Residential Appliance Holdings and Consumption,' *Econometrica*, 52, pp. 345–362.
- Fezzi, C. and Bunn, D., 2010, 'Structural Analysis of Electricity Demand and Supply Interactions,' *Oxford Bulletin of Economics and Statistics*, 72, pp. 827–856.
- Hand, M., 2002, 'The Economists: On the Future of Energy Markets,' *Public Utilities Fortnightly*, 140, pp. 12–18.
- Hausman, J.; Kinnucan, M. and McFadden, D., 1979, 'A Two-Level Electricity Demand Model: Evaluation of the Connecticut Time-of-Day Pricing Test,' *Journal of Econometrics*, 10, pp. 263–289.
- Hsiao, C. and Mountain, D., 1985, 'Estimating the Short-run Income Elasticity of Demand for Electricity,' *Journal of the American Statistical Association*, 80, pp. 259–265.
- Krichene, N., 2002, 'World Crude Oil and Natural Gas: A Demand and Supply Model,' *Energy Economics*, 24, pp. 557–576.
- Imbens, G. and Lemieux, T., 2008, 'Regression Discontinuity Designs: A Guide to Practice,' *Journal of Econometrics*, 142, pp. 615–635.

- Ito, K., 2013, 'Do Consumers Respond to Marginal or Average Price? Evidence from Nonlinear Electricity Pricing,' *American Economic Review*, forthcoming.
- Jessoe, K. and Rapson, D., 2011, 'Commercial and Industrial Demand Response Under Mandatory Time-of-Use Electricity Pricing,' UC-Berkeley, UCE³ WP-023.
- Lee, D., 2008, 'Randomized Experiments from Non-random Selection in U.S. House Elections,' *Journal of Econometrics*, 142, pp. 675–697.
- Lee, D. and Card, D., 2008, 'Regression Discontinuity Inference with Specification Error,' *Journal of Econometrics*, 142, pp. 655–674.
- Lee, D. and Lemieux, T., 2010, 'Regression Discontinuity Designs in Economics,' *Journal of Economic Literature*, 48, pp. 281–355.
- Liu, L. and Lin, M., 1991, 'Forecasting Residential Consumption of Natural Gas using Monthly and Quarterly Time Series,' *International Journal of Forecasting*, 7, pp. 3–16.
- Loewenstein, G. and Ubel, P., 2010, 'Economics Behaving Badly,' *The New York Times*, 14 July. Available at: <http://www.nytimes.com/2010/07/15/opinion/15loewenstein.html>
- McCrary, J., 2008, 'Manipulation of the Running Variable in the Regression Discontinuity Design: A Density Test,' *Journal of Econometrics*, 142, pp. 698–714.
- National Institute of Statistics and Censuses, 2003, '¿Qué es el Gran Buenos Aires?,' Press Release.
- Parks, R. and Weitzel, D., 1984, 'Measuring Consumer Welfare Effects of Time-Differentiated Prices,' *Journal of Econometrics*, 26, pp. 25–65.
- Parti, M. and Parti, C., 1980, 'The Total and Appliance-Specific Conditional Demand for Electricity in the Household Sector,' *Bell Journal of Economics*, 11, pp. 309–321.
- Pew Research Center, 2012, 'Assessing the Representativeness of Public Opinion Surveys,' Washington DC.
- Reiss, P. and White, M., 2005, 'Household Electricity Demand, Revisited,' *Review of Economic Studies*, 72, pp. 853–883.
- Reiss, P. and White, M., 2008, 'What Changes Energy Consumption? Prices and Public Pressures,' *RAND Journal of Economics*, 39, pp. 636–663.
- Shin, J., 1985, 'Perception of Price when Price Information is Costly: Evidence from Residential Electricity Demand,' *Review of Economics and Statistics*, 67, pp. 591–598.

NOTES

¹Research on this topic, discussed in more detail below, dates to Parti and Parti [1980], Dubin and McFadden [1984] and Hsiao and Mountain [1985]. Recent influential contributions include Reiss and White [2005/2008].

²Greater Buenos Aires (Gran Buenos Aires) is an urban metropolitan area that comprises the city of Buenos Aires and 24 adjacent municipalities (National Institute of Statistics and Censuses [2003]). According to the 2010 population census it has 12.8 million inhabitants, nearly a third of the total Argentinean population.

³The evidence we provide may therefore contribute to the discussion on the relative importance of prices versus nudges for steering consumers' behaviors (Loewenstein and Ubel [2010]).

⁴Focusing on the gas market in the US, Davis and Muehlegger [2010] exploit variation in wholesale prices to estimate the price elasticity of demand in the residential sector. We believe that our RD estimates add value to this literature in that they: (1) are less likely to be affected by unobserved heterogeneity across households; and (2) emphasize the role of imperfect information in shaping consumer behavior in this market. See also Jessoe and Rapson [2011] for a recent analysis of commercial and industrial demand response under mandatory time-of-use electricity pricing in Connecticut.

⁵To preserve confidentiality of its broad customer base, the company refused to provide us with data on consumers from other categories.

⁶For simplicity, in this section we focus on consumers with annual accumulated consumption between 1,251 and 1,800 cubic meters who can therefore face only two potential prices (see Table I).

⁷The survey used records from 2952 households in order to obtain 353 valid responses from the member of the household that is responsible for paying the gas bills. A failure to produce a valid interview occurred when the enumerator: (1) was unable to establish a phone conversation with any member of the household; (2) was unable to reach the person responsible for paying the gas bills; or (3) the person responsible for paying the gas bills refused to cooperate. The response rate of the survey (12%) is well in line with that of recent phone surveys in the US—a recent study by the Pew Research Center [2012] reports that a standard survey yields a response rate of 9%.

⁸The full set of results is available upon request.

⁹Measured consumption is always rounded to the nearest integer, and hence we cluster the standard errors at the unit levels of the running variable.

¹⁰To verify if the amount billed in period 0 actually followed the tariff structure prevailing at the time, we have predicted the bill by applying the prevailing unit prices, adding the fixed charge and applying the minimum billed amount. The correlation between predicted and actual bills is 0.98, suggesting that bills in period 0 followed closely the prevailing price schedule.

¹¹For example, if a household receives a guest in period 0 gas demand would likely increase in that period, returning to normal levels in period 1 once the guest leaves.

¹²This elasticity corresponds to the inverse of the ratio between the percentage increase in the unit price faced by consumers just above the threshold (25%) and the estimated change in gas consumption during the subsequent two-month period (-3.8% of average gas consumption).

¹³Reiss and White [2008] survey earlier literature using cross-sectional survey data for household electricity consumption, and conclude that the cross-sectional estimates vary widely, with typical values near -0.3.

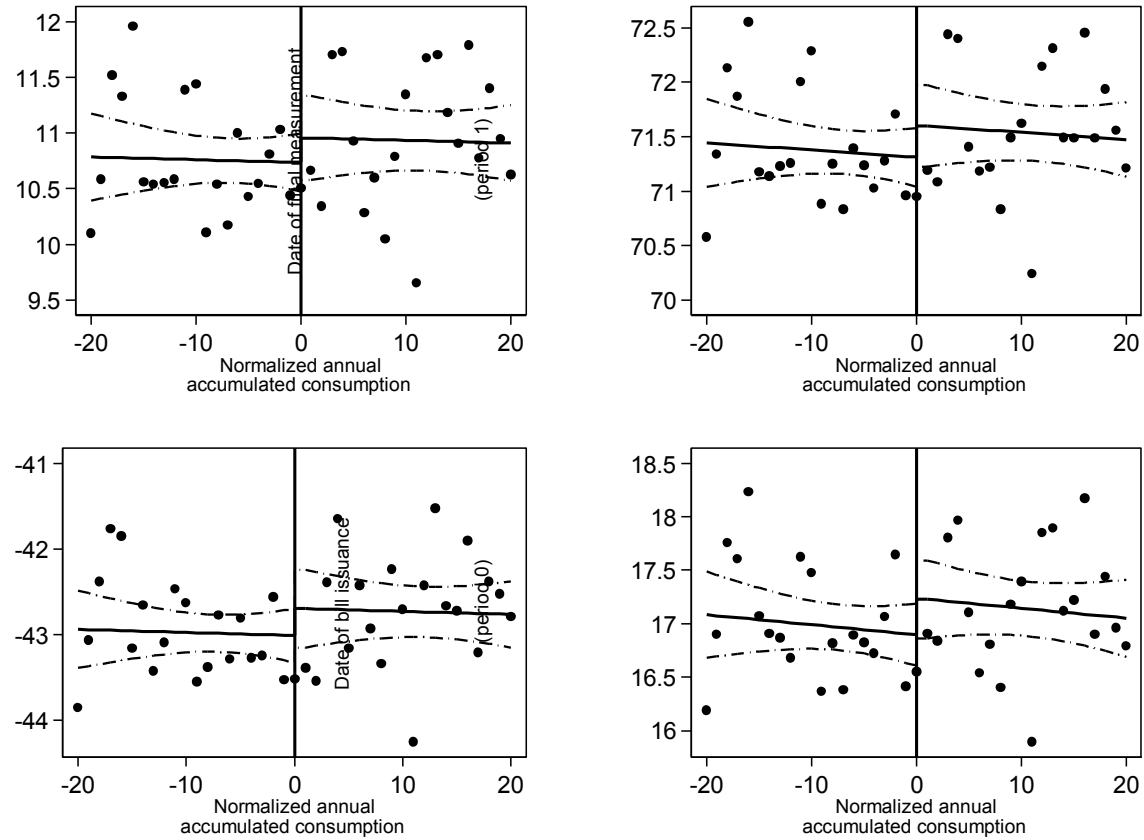


Figure 1
Dates by Annual Accumulated Consumption

Notes: Dates are normalized so that May 1st, 2009 corresponds to day 1. Scatter points correspond to local averages computed by bins with a bandwidth of 1 unit of normalized AAC. Solid lines correspond to parametric fits generated from regressions that include a constant, a treatment dummy and a linear term for normalized AAC. Dashed lines are the 95% confidence interval.

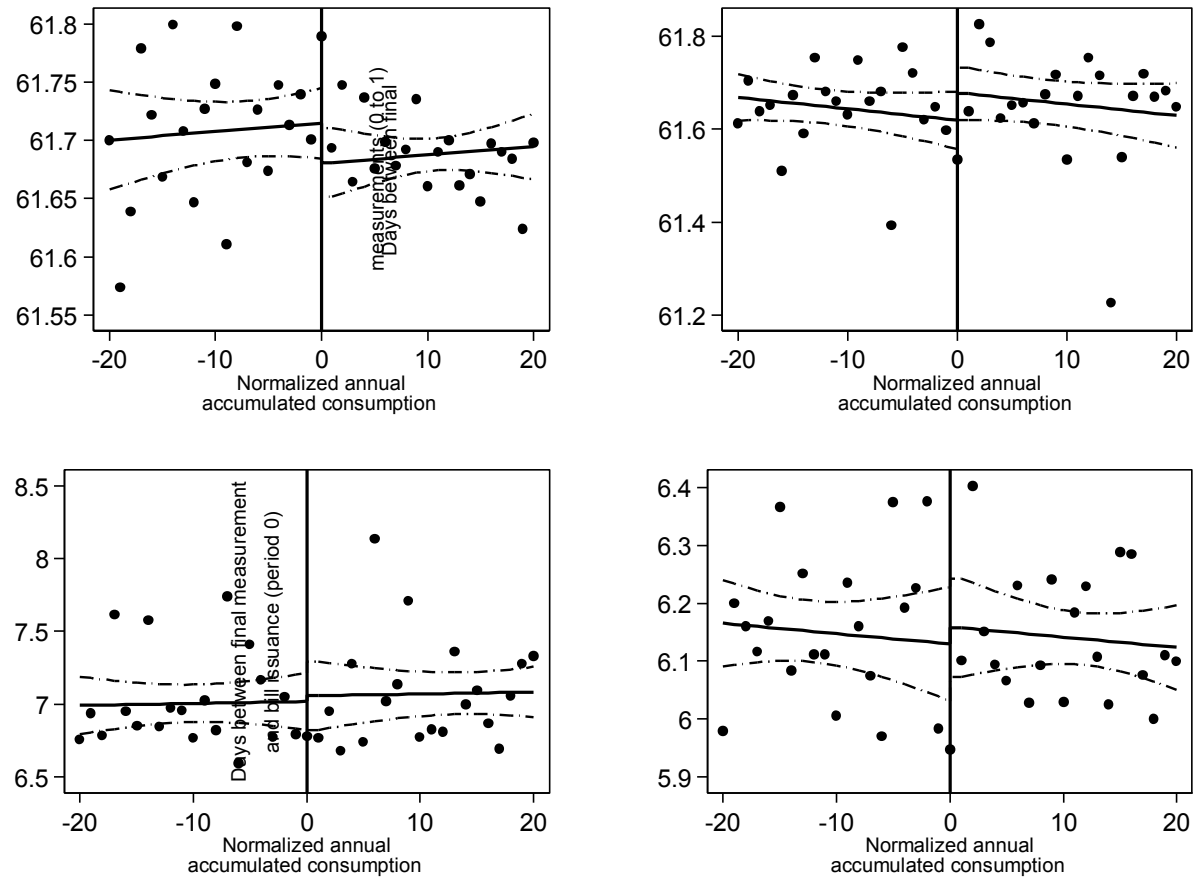


Figure 1 (cont.)
Dates by Annual Accumulated Consumption

Notes: Dates are normalized so that May 1st, 2009 corresponds to day 1. Scatter points correspond to local averages computed by bins with a bandwidth of 1 unit of normalized AAC. Solid lines correspond to parametric fits generated from regressions that include a constant, a treatment dummy and a linear term for normalized AAC. Dashed lines are the 95% confidence interval.

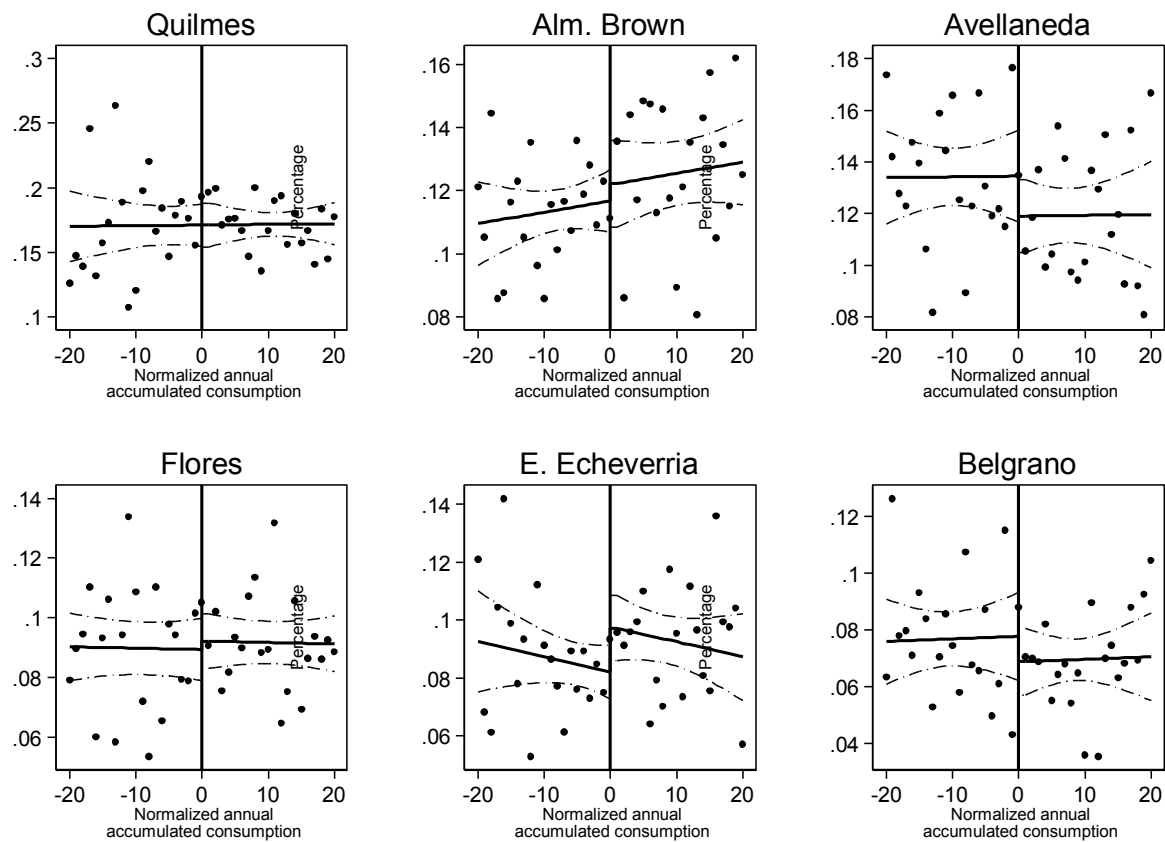


Figure 2
Region of Residence by Annual Accumulated Consumption (by 6 Largest Regions in Terms of Number of Users)

Notes: Scatter points correspond to local averages computed by bins with a bandwidth of 1 unit of normalized AAC. Solid lines correspond to parametric fits generated from regressions that include a constant, a treatment dummy and a linear term for normalized AAC. Dashed lines are the 95% confidence interval.

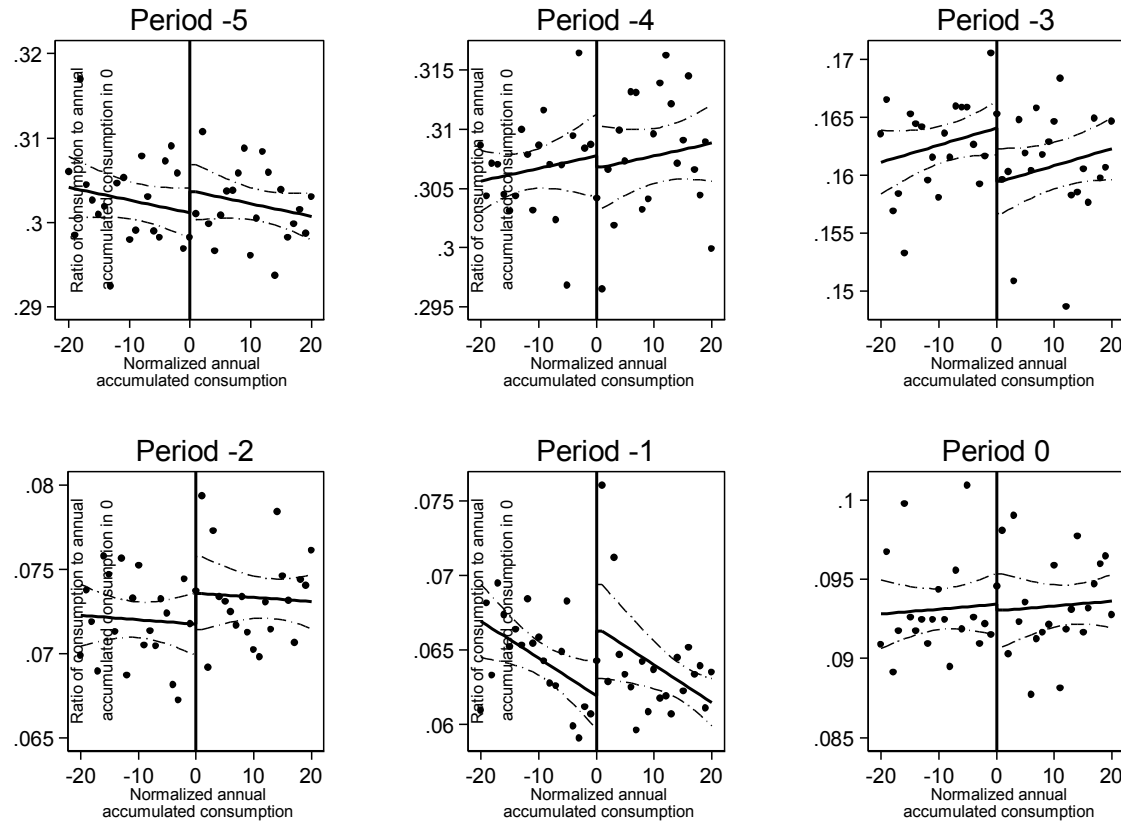


Figure 3
Ratio of Consumption to Annual Accumulated Consumption in Period 0

Notes: Scatter points correspond to local averages computed by bins with a bandwidth of 1 unit of normalized AAC. Solid lines correspond to parametric fits generated from regressions that include a constant, a treatment dummy and a linear term for normalized AAC. Dashed lines are the 95% confidence interval.

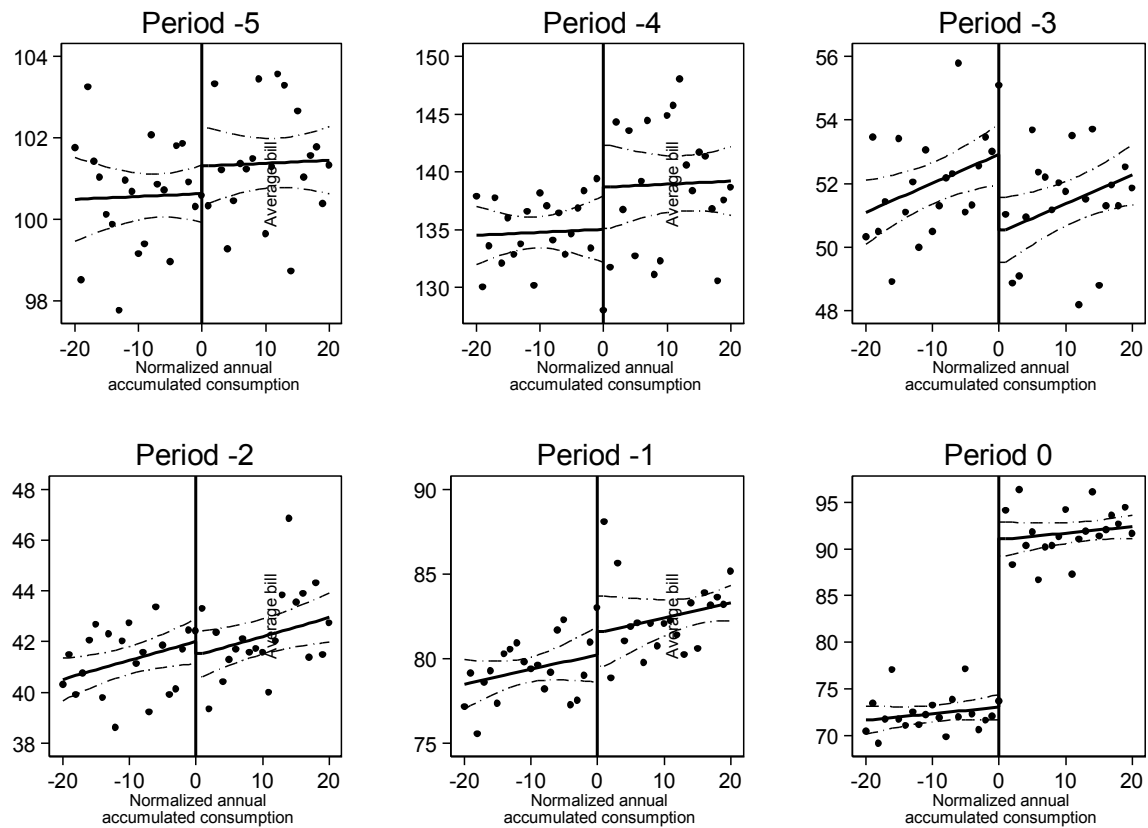


Figure 4
Average Bill by Annual Accumulated Consumption

Notes: Scatter points correspond to local averages computed by bins with a bandwidth of 1 unit of normalized AAC. Solid lines correspond to parametric fits generated from regressions that include a constant, a treatment dummy and a linear term for normalized AAC. Dashed lines are the 95% confidence interval.




Figure 5
Number of Observations by Annual Accumulated Consumption in Period 0




Figure 6
Average Consumption in July 2009 by Annual Accumulated Consumption (in m3)

Notes: Scatter points correspond to local averages computed by bins with a bandwidth of 1 unit of normalized AAC. Solid lines correspond to parametric fits generated from regressions that include a constant, a treatment dummy and a linear term for normalized AAC. Dashed lines are the 95% confidence interval.

TABLE I
TARIFF SCHEDULE VALID SINCE NOVEMBER 1ST, 2008

Category	Accumulated consumption (m ³ /year)		Unit price per m ³
	From	To	
R1	0	500	
R21	501	650	
R22	651	800	0.144
R23	801	1000	0.156
R31	1001	1250	0.247
R32	1251	1500	0.332
R33	1501	1800	0.437
R34	1801	+∞	0.517

Notes: The econometric analysis focuses on the discontinuity between the categories R32 and R33.

TABLE II
SOCIOECONOMIC CHARACTERISTICS OF HOUSEHOLDS IN THE ESTIMATION SAMPLE

Variable	Estimation sample survey	Household Survey
Female	75.4	52.0
Age	52.5	32.1
Married	60.8	38.6
Tertiary education	38.0	19.4
Home owner	88.9	71.2
Number of rooms	3.3	3.1
Number of families	1.1	1.0

Notes: Individual characteristics in the estimation sample survey are from the member of the household that is responsible for paying the gas bills. Statistics on households from Argentina's Household Survey (*Encuesta Permanente de Hogares*) refer to households living in Buenos Aires. "Number of families" refers to the number of different families living in the household. "Number of rooms" refers to the total number of rooms in the household.

TABLE III
KNOWLEDGE ON BILL AMOUNT AND PRICE DETERMINATION MECHANISM

Last amount billed		
Question		% Yes
Do you remember the amount of your last bill?		91.8
Price determination mechanism – Perceived knowledge		
Question		% Yes
Do you know how the total amount of the bill is computed?		30.7
Price determination mechanism – Objective knowledge		
Question	Correct Answer	% Correct Answer
How often does the company re-categorize consumers?	Every billing period	14.4
Re-categorization is calculated based on...	Last year's consumption	38.9
What is the consumption level that divides categories R32 and R33	1500 m ³	3.7

Notes: Results from a phone survey of 353 customers that had an annual accumulated consumption between 1480 and 1520 m³ in the bill issued in May 2009. For the questions about objective knowledge, four alternatives were presented. In the question “How often the company re-categorize consumers?” the options were: a) every billing period, b) every two billing periods, c) every six billing periods, d) other. In the question “Re-categorization is calculated based on ...” the options were: a) difference in consumption between the current bill and the previous; b) last year's consumption; c) last semester consumption; d) other. In the question “What is the consumption level that divides categories R32 and R33” the options were: a) 1000 m³; b) 2000 m³; c) 1,500 m³; d) Does not know.

TABLE IV: DATES AND PERIODS BY TREATMENT STATUS
(DAYS NORMALIZED: MAY 1ST, 2009 = DAY 1)

	Treatment	Control	Raw Difference	Adjusted Difference
Panel A: Dates				
Final Measurement Period 0	10.93 (0.14)	10.77 (0.11)	0.17 (0.17)	0.27 (0.29)
Bill Period 0 (<i>Treatment Date</i>)	17.07 (0.13)	16.91 (0.12)	0.16 (0.18)	0.30 (0.30)
Final Measurement Period 1	71.62 (0.13)	71.45 (0.12)	0.17 (0.17)	0.31 (0.29)
Panel B: Days Between Bill Period 0 and Final Measurement Period 1				
Days	54.55 (0.02)	54.54 (0.03)	0.01 (0.03)	0.01 (0.03)
Panel C: Days Between Final Measurement in Period 0 and Period 1				
Days	60.69 (0.02)	60.68 (0.02)	0.00 (0.03)	0.04 (0.06)

Notes: Means and standard errors in parentheses. Sample includes all customers with accumulated annual consumption between 1480 and 1520 m³ in the bill that was issued in May, 2009. Period 0 corresponds to the one whose bill was issued in May 2009. Dates in the table are normalized so May 1st, 2009 corresponds to day 1. Day 11 = May 11th, 2009. Day 17 = 17th May, 2009. Day 72 = July 12th, 2009. Standard errors clustered by accumulated consumption in Period 0.

TABLE V
REGION OF RESIDENCE BY TREATMENT STATUS

	Treatment	Control	Raw Difference	Adjusted Difference
Quilmes	0.172 (0.004)	0.171 (0.009)	0.001 (0.010)	0.000 (0.014)
Avellaneda	0.119 (0.006)	0.134 (0.006)	-0.015* (0.008)	-0.016 (0.014)
Ate. Brown	0.126 (0.005)	0.113 (0.004)	0.013* (0.006)	0.005 (0.011)
Flores	0.092 (0.004)	0.090 (0.004)	0.002 (0.006)	0.003 (0.008)
E. Echeverría	0.092 (0.004)	0.087 (0.005)	0.005 (0.006)	0.016* (0.009)
Belgrano	0.069 (0.004)	0.077 (0.005)	-0.007 (0.006)	-0.009 (0.013)
Floresta	0.060 (0.004)	0.066 (0.004)	-0.005 (0.006)	0.001 (0.010)
Devoto	0.064 (0.003)	0.054 (0.004)	0.010* (0.005)	-0.005 (0.011)
Norte	0.045 (0.004)	0.054 (0.003)	-0.009* (0.005)	-0.014 (0.011)
Other	0.160 (0.006)	0.154 (0.005)	0.006 (0.008)	0.019 (0.014)

Notes: Means and standard errors in parenthesis. The Raw Difference column reports mean difference between the Treatment and Control groups. The Adjusted Difference column presents the coefficient of regressing the respective variable on a dummy for treatment and a linear term for annual accumulated consumption by Period 0. Standard errors clustered by accumulated consumption in Period 0. *** p<0.01, ** p<0.05, * p<0.1. The Other category includes Almagro, Mataderos, Centro, Lomas de Zamora, Barracas, Lanús, San Vicente and Berazategui.

TABLE VI
CONSUMPTION LEVELS AND RATIOS TO ACCUMULATED CONSUMPTION BY TREATMENT STATUS

	Treatment	Control	Raw Difference	Adjusted Difference
Consumption Levels (m ³)				
Period -5	456.378 (1.469)	450.914 (1.714)	5.464** (2.230)	3.922 (4.423)
Period -4	464.935 (1.927)	456.953 (1.414)	7.982*** (2.360)	-1.592 (5.062)
Period -3	243.048 (1.576)	242.274 (1.391)	0.774 (2.076)	-7.062* (3.526)
Period -2	110.743 (0.970)	107.277 (0.809)	3.466*** (1.247)	2.760 (2.933)
Period -1	96.448 (1.293)	96.001 (0.940)	0.447 (1.579)	6.835 (4.075)
Period 0	140.970 (1.046)	138.733 (1.008)	2.237 (1.435)	-0.630 (3.080)
Ratio of Consumption in a Period and Annual Accumulated Consumption by Period 0				
Period -5	0.302 (0.001)	0.303 (0.001)	0.000 (0.002)	0.003 (0.003)
Period -4	0.308 (0.001)	0.307 (0.001)	0.001 (0.001)	-0.001 (0.003)
Period -3	0.161 (0.001)	0.163 (0.001)	-0.002 (0.001)	-0.004 (0.002)
Period -2	0.073 (0.001)	0.072 (0.001)	0.001 (0.001)	0.002 (.002)
Period -1	0.064 (0.001)	0.064 (0.001)	-0.001 (0.001)	0.005 (0.003)
Period 0	0.093 (0.001)	0.093 (0.001)	0.000 (0.001)	0.000 (0.002)

Notes: Means and standard errors in parenthesis. The Raw Difference column reports mean difference between the Treatment and Control groups. The Adjusted Difference column presents the coefficient of regressing the respective variable on a dummy for treatment and a linear term for annual accumulated consumption by Period 0. Standard errors clustered by accumulated consumption in Period 0. *** p<0.01, ** p<0.05, * p<0.1.

TABLE VII
BILL AMOUNT BY TREATMENT STATUS

	Treatment	Control	Raw Difference	Adjusted Difference
Period -5	101.380 (0.305)	100.558 (0.288)	0.822* (0.414)	0.678 (0.731)
Period -4	138.959 (1.127)	134.761 (0.679)	4.198*** (1.384)	3.658 (3.034)
Period -3	51.406 (0.350)	52.004 (0.363)	-0.598 (0.498)	-2.473*** (0.891)
Period -2	42.233 (0.376)	41.255 (0.277)	0.978 (0.461)	-0.560 (0.795)
Period -1	82.449 (0.512)	79.348 (0.403)	3.101 (0.643)	1.299 (1.816)
Period 0	91.728 (0.565)	72.336 (0.436)	19.393*** (0.704)	17.965*** (1.474)

Notes: Means and standard errors in parenthesis. Bill amounts are in Argentine Pesos. The Raw Difference column reports mean difference between the Treatment and Control groups. The Adjusted Difference column presents the coefficient of regressing the respective variable on a dummy for treatment and a linear term for annual accumulated consumption by Period 0. Standard errors clustered by accumulated consumption in Period 0. *** p<0.01, ** p<0.05, * p<0.1.

TABLE VIII

IMPACTS OF PRICE INCREASE IN BILL 0 ON CONSUMPTION IN PERIOD 1, ALTERNATIVE FUNCTIONAL FORMS

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Treatment	-15.901** (6.491)	-15.266** (5.929)	-15.577** (6.153)	-15.597* (9.194)	-16.539* (9.201)	-14.595 (9.917)	-16.516* (9.174)	-17.126 (12.408)
AAC	0.720** (0.290)	1.004** (0.412)	0.707** (0.272)	2.201 (1.441)	0.812 (0.702)	7.515** (2.872)	0.811 (0.696)	2.352 (5.143)
AAC * Treatment		-0.619 (0.531)		-2.777 (1.968)		-12.721*** (3.936)		-2.244 (8.023)
AAC ²			-0.011 (0.012)	0.059 (0.069)	-0.011 (0.012)	0.738** (0.356)	-0.012 (0.044)	-0.487 (1.014)
AAC ² * Treatment				-0.014 (0.090)		-0.149 (0.465)		-0.009 (1.619)
AAC ³					-0.000 (0.002)	0.023* (0.012)	-0.000 (0.001)	-0.074 (0.079)
AAC ³ * Treatment						-0.040** (0.015)		0.136 (0.121)
AAC ⁴							0.000 (0.000)	-0.002 (0.002)
AAC ⁴ * Treatment								0.001 (0.003)
Constant	433.271*** (3.299)	436.134*** (3.992)	434.620*** (3.539)	439.962*** (6.161)	435.079*** (4.700)	447.792*** (6.506)	435.135*** (5.037)	443.718*** (8.874)
N	7190	7190	7190	7190	7190	7190	7190	7190
R-squared	0.001	0.001	0.001	0.001	0.001	0.002	0.001	0.002

Notes: The dependent variable is consumption in Period 1. Average of the dependent variable is 425.49. The estimation method is OLS. Standard errors clustered by accumulated consumption in Period 0. *** p<0.01, ** p<0.05, * p<0.1.

TABLE IX
IMPACTS OF PRICE INCREASE IN BILL 0 ON CONSUMPTION IN PERIOD 1, ALTERNATIVE BANDWIDTHS

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Treatment	-15.901** (6.491)	-15.266** (5.929)	-16.296* (8.105)	-14.780** (6.882)	-18.863* (10.386)	-15.872* (8.484)	-16.091 (13.525)	-13.176 (11.124)
AAC	0.720** (0.290)	1.004** (0.412)	0.759 (0.451)	1.443** (0.612)	1.060 (0.832)	2.403** (0.863)	0.703 (1.774)	1.898 (2.309)
AAC * Treatment		-0.619 (0.531)		-1.522* (0.782)		-3.134** (1.333)		-3.220 (2.821)
Constant	433.271*** (3.299)	436.134*** (3.992)	433.000*** (4.042)	438.133*** (4.647)	435.086*** (5.287)	441.855*** (5.386)	438.154*** (7.297)	441.205*** (8.397)
Bandwidth	20	20	15	15	10	10	5	5
N	7190	7190	5417	5417	3679	3679	1946	1946
R-squared	0.001	0.001	0.001	0.001	0.001	0.002	0.001	0.002

Notes: The dependent variable is consumption in Period 1. The bandwidths determine the range of normalized annual accumulated consumption used to select observations for a particular regression. For example, in Columns (1) and (2) only individuals with normalized annual accumulated consumption between [-20, 20] are included. Average of the dependent variable is 425.49. The estimation method is OLS. Standard errors clustered by accumulated consumption in Period 0. *** p<0.01, ** p<0.05, * p<0.1.

TABLE X
EXPLOITING PLACEBO AAC THRESHOLDS TO TEST FOR MEAN REVERSION

	(1)
AAC \geq -10	-1.646 (6.250)
AAC \geq 0	-16.969** (7.209)
AAC \geq 10	-3.129 (6.383)
AAC	0.890 (0.531)
Constant	435.767*** (7.941)
N	7190
R-squared	0.001

Notes: The dependent variable is consumption in Period 1. Explanatory variables include indicators for annual accumulated consumption (AAC) higher than -10, 0 and 10 and also AAC as a continuous variable. Individuals with normalized annual accumulated consumption between [-20, 20] are included. Average of the dependent variable is 425.49. The estimation method is OLS. Standard errors clustered by accumulated consumption in Period 0.*** p<0.01, ** p<0.05, * p<0.1.

TABLE XI
DIFFERENTIAL EFFECTS ACROSS CONSUMERS IN THE TREATMENT GROUP BY DISTANCE TO THE THRESHOLD

	(1)
Close	-16.339** (6.567)
Far	-18.500* (9.458)
AAC	0.794** (0.374)
Constant	434.015*** (3.835)
N	7190
R-squared	0.001

Notes: The dependent variable is consumption in Period 1. Explanatory variables include indicators for normalized annual accumulated between 1 and 10 (Close) and between 11 and 20 (Far) and 10 and also AAC as a continuous variable. Individuals with normalized annual accumulated consumption between [-20, 20] are included. Average of the dependent variable is 425.49. The estimation method is OLS. Standard errors clustered by accumulated consumption in Period 0.*** p<0.01, ** p<0.05, * p<0.1.

TABLE AI

ROBUSTNESS CHECK: IMPACTS OF PRICE INCREASE IN BILL 0 ON CONSUMPTION IN PERIOD 1

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)	(12)
Treatment	-15.901** (6.491)	-17.415*** (6.342)	-15.266** (5.929)	-16.637*** (5.656)	-15.577** (6.153)	-17.013*** (5.934)	-15.597* (9.194)	-16.670* (8.996)	-16.296* (8.105)	-18.000** (7.922)	-14.780** (6.882)	-16.371** (6.583)
AAC	0.720** (0.290)	0.822*** (0.272)	1.004** (0.412)	1.169*** (0.368)	0.707** (0.272)	0.806*** (0.251)	2.201 (1.441)	2.547* (1.282)	0.759 (0.451)	0.876* (0.450)	1.443** (0.612)	1.603** (0.604)
AAC * Treatment			-0.619 (0.531)	-0.753 (0.493)			-2.777 (1.968)	-3.335* (1.823)			-1.522* (0.782)	-1.620** (0.783)
AAC ²					-0.011 (0.012)	-0.013 (0.011)	0.059 (0.069)	0.068 (0.061)				
AAC ² * Treatment							-0.014 (0.090)	-0.011 (0.086)				
Constant	433.271*** (3.299)	461.555*** (33.802)	436.134*** (3.992)	464.847*** (33.767)	434.620*** (3.539)	463.133*** (33.715)	439.962*** (6.161)	468.975*** (35.078)	433.000*** (4.042)	465.119*** (43.858)	438.133*** (4.647)	470.180*** (44.095)
Bandwidth	20	20	20	20	20	20	20	20	15	15	15	15
Areas Effects	N	Y	N	Y	N	Y	N	Y	N	Y	N	Y
N	7190	7190	7190	7190	7190	7190	7190	7190	5417	5417	5417	5417
R-squared	0.001	0.056	0.001	0.056	0.001	0.056	0.001	0.057	0.001	0.057	0.001	0.058

Notes: The dependent variable is consumption in Period 1. The bandwidths determine the range of normalized annual accumulated consumption used to select observations for a particular regression. For example, in Columns (1) and (2) only individuals with normalized annual accumulated consumption between [-20, 20] are included. Average of the dependent variable is 425.49. The estimation method is OLS. Standard errors clustered by accumulated consumption in Period 0. Even columns present regressions with areas (neighborhoods) dummies. *** p<0.01, ** p<0.05, * p<0.1.