Journal of Economic Behavior & Organization 107 (2014) 308–323

Contents lists available at ScienceDirect

Journal of Economic Behavior & Organization

journal homepage: www.elsevier.com/locate/jebo

Perceived price in residential water demand: Evidence from a natural experiment☆

Casey J. Wichman*

Department of Agricultural and Resource Economics, University of Maryland, 2200 Symons Hall, College Park, MD 20742, United States

ARTICLE INFO

Article history:
Received 5 October 2013
Received in revised form 15 February 2014
Accepted 21 February 2014
Available online 13 March 2014

JEL classification:
Q25
D12
L95
H42

Keywords:
Water demand
Price elasticity
Difference-in-difference-in-difference
Regression discontinuity
Perceived price

ABSTRACT

Under complicated billing structures, the price to which consumers respond remains inconclusive. In this paper, I exploit a quasi-experiment to estimate a causal effect of price for residential water customers during the introduction of increasing block rates for a North Carolina utility. Perceived price is identified through a billing anomaly in which changes in marginal and average prices move in opposite directions. Empirical results contribute evidence that residential water customers respond to average price. Average price elasticity estimates vary from −0.43 to −1.14 across the distribution of consumption in triple-difference models, with an estimate of −0.31 in the tightest bandwidth of regression discontinuity specifications.

1. Introduction

Increasing block rate structures are typically adopted by water utility managers to promote conservation among high-users and affordability for low-to-moderate-users. With growing concern for water scarcity, as well as the need for revenue stability at the utility-level, the introduction of block rates for residential water customers is becoming increasingly common. However, upward pressure on the costs of providing water to households necessitates a better understanding of how consumers respond to unclear price signals. This knowledge will help utility managers craft rate structures commensurate with their goals and provide researchers with a framework for studying water demand that better conforms to observed consumer behavior. Specifically, price elasticity of residential water demand is the key parameter of interest because price

☆ I thank conference participants at NC State, the 2nd Annual AERE Summer Conference, the 2012 AAEA Annual Meetings, and the Identification of Causal Effects in Energy and Environmental Economics workshop at the University of Tennessee. Klaus Moeltner, Christian Vossler, Anna Alberini, Roger von Haefen, Laura Taylor, and Kerry Smith provided helpful comments, criticisms, and suggestions that substantially improved this manuscript. In addition, Pat Davis and Kevin Ray at the Orange Water and Sewer Authority, as well as Jeff Hughes, Mary Tiger, and Shadi Eskaf at the UNC Environmental Finance Center, deserve special thanks for their help with data requests and constructive feedback. Travel funding and honorarium from the Howard H. Baker Jr. Center for Public Policy at the University of Tennessee, Knoxville is gratefully acknowledged.

* Tel.: +1 301 405 1293; fax: +1 301 314 9091.
E-mail address: wichman@umd.edu

http://dx.doi.org/10.1016/j.jebo.2014.02.017

is used as an instrument of conservation during periods of acute scarcity and utilities are often constrained by zero-profit mandates such that the impact of changes in prices is important for revenue planning (Olmstead et al., 2007).

In this paper, I motivate a conceptual framework for residential water demand that relies on a customer's consumption patterns last month as a heuristic for this month's prices. This identification strategy expands upon previous research by exploiting the assignment of billing cycles for residential customers of Orange Water and Sewer Authority (OWASA) in Chapel Hill, North Carolina during the introduction of increasing block rates. Conceptually, this treatment allows for identification of perceived price for one subset of households relative to nearly identical households on a different billing cycle who face similar weather patterns and utility-specific shocks but may be responding to different price information.

Most of the literature on the price elasticity of water demand suffers from the non-experimental nature of utility pricing, relying on cross-sectional variation across households or municipalities, or variation over time. The former provides a potential avenue for omitted variables to bias results and the latter fails to exploit exogenous, unanticipated changes to the pricing structure for a given household (Olmstead et al., 2007; Klaiber et al., 2014). Further, the co-movement of marginal and average price tend to confound results that examine price perception. In this analysis, three complementary quasi-experimental methods are applied to explore the behavioral response to a change in the rate structure in an attempt to isolate a causal estimate of price elasticity for water customers. Difference-in-difference (DD) methods are applied to capture an average treatment effect for all water customers based on the notion that consumers respond to last month's bill as a proxy for this month's price. Triple-difference (DDD) estimates allow for the identification of heterogeneous price effects by segregating the sample into decile groups that correspond to a household's average historical consumption. Finally, the transition to a block rate structure allows for the exploitation of a discontinuity in price in which changes in marginal and average prices move in opposite directions for a portion of the water use distribution. Fuzzy regression discontinuity (FRD) models that focus on this anomaly reinforce the DDD results and allow for testing whether the identification strategy is plausible in light of potential confounding factors. The complementarity of this suite of quasi-experimental techniques provides strong empirical evidence that residential water consumers respond to average price.

The results of empirical models indicate that the short-run response to the adoption of increasing block rates is a net increase in consumption due to lower prices in the first and second consumption blocks. Further, DDD estimation results lend evidence that customers across most deciles of consumption respond to average price, while customers in both tails of the distribution exhibit no significant response to price. Average price elasticity estimates range from −0.43, for consumption around 3000 gal per month, to −1.14, for consumption around 7000 gal per month, which are within the range of previous studies (Espey et al., 1997; Dalhuisen et al., 2003). Elasticity estimates calculated with changes in marginal price display either implausibly large magnitudes, wrong signs, or statistical insignificance. Lastly, by exploiting the divergence in average and marginal prices at 7000 gal per month after the rate change, FRD results provide further evidence that residential water customers respond to average price.

2. Perceived price

Despite an extensive literature estimating the effect of price changes on customer demand under complicated rate structures, the price signal that residential water customers use to make consumption decisions remains unclear to researchers (Nataraj and Hanemann, 2011). Standard economic theory stipulates that consumers should respond to the marginal price for the next unit of water consumed, as well as the marginal price in each block below the final block of consumption (Hewitt and Hanemann, 1995; Olmstead et al., 2007; Strong and Smith, 2010). However, it is plausible that consumers are not aware of the marginal price they face in each of the blocks (Nataraj and Hanemann, 2011) and consumers might not fully understand how their bill is calculated within the tiered rate structures (Nieswiadomy and Molina, 1989). Recent work in electricity demand suggests that consumption patterns better reflect a response commensurate with changes in either expected marginal price (Borenstein, 2009) or average price (Ito, 2014) under tiered rate structures, while results from a quasi-experiment in water demand imply that high-volume residential water customers seem to respond to marginal price (Nataraj and Hanemann, 2011).

2.1. Marginal, expected marginal, and average price

In a simple model of demand under block pricing without uncertainty, consider a consumer with quasi-linear utility:

$$u(\mathbf{w}, x) = V(\mathbf{w}) + x \tag{1}$$

where $\mathbf{w}$ is a vector of water consumption in each price block for that period and x is a linearly separable numeraire good with its price normalized to unity. Let $I = x + \mathbf{p}'\mathbf{w}$ represent the consumer's budget constraint with $\mathbf{p}$ being the price schedule for water consumption and wealth, I, is determined exogenously. The consumer maximizes her utility subject to the prices she faces in two blocks of water consumption:

$$\max_{\mathbf{w}} u(\mathbf{w}, x) = I + V(\mathbf{w}) - p_1 w_1 - p_2 w_2 \tag{2}$$

where w_1 and w_2 are the quantities demanded in each consumption block corresponding to marginal prices p_1 and p_2. The solution to the consumer's problem results in the following piece-wise demand function for water:

$$\mathbf{w}^{\star} = \begin{cases} w^{\star}(p_1) & \text{if } w^{\star}(p_1) \leq k \\ k & \text{if } w^{\star}(p_1) = k = w^{\star}(p_2) \\ w^{\star}(p_2) & \text{if } w^{\star}(p_2) \geq k \end{cases} \tag{3}$$

where k is the "kink point" in the consumer's budget constraint (Ito, 2014).[1] This model of water demand with non-linear budget constraints, formalized by Hewitt and Hanemann (1995), is typically estimated using discrete-continuous choice methods in which the consumer chooses her consumption block and the optimal amount to consume within that block simultaneously. In the discrete-continuous choice framework, marginal prices in each block are necessary pieces of information for the consumer to make her consumption decision. While this model conforms to utility theory, it makes the assumption that a consumer performs a complex utility-maximizing decision based on perfect information of the water provider's rate schedule and her precise level of consumption throughout the billing period (Borenstein, 2009).

Borenstein (2009) relaxes this assumption by introducing a model that allows for consumption decisions to be made in response to local marginal prices in the context of electricity demand. This conceptualization of consumer behavior provides a more intuitive model of demand under block rates while avoiding the restrictive assumptions in Hewitt and Hanemann's framework. In this model, consumers maximize expected utility,

$$\max_{\mathbf{w}} E[u(\mathbf{w}, x)] = I + E[V(\mathbf{w})] - E[p_1 w_1 + p_2 w_2], \tag{4}$$

where the first-order condition states that a consumer will choose consumption $\mathbf{w}^*$ that sets her marginal utility equal to the expected marginal price, which is a probability-weighted average of the marginal prices in each consumption block. This formulation allows for smooth demand functions, though it still requires complete knowledge and understanding of the utility's rate structure. Borenstein tests an empirical model in which customers set behavioral consumption rules at the start of the period, such as setting the thermostat to a fixed temperature, based on the marginal price they expect to face while allowing for exogenous demand shocks within the consumption period. He finds evidence that electricity customers are more likely to respond to expected marginal price than marginal price in Southern California, but it is possible that they are responding to even less precise information.

Further, several empirical papers have suggested that even expected marginal price places too large of a computational burden on customers who face block rates and that the total bill, or average price, is a more accurate representation of a customer's perceived price (Foster and Beattie, 1981; Liebman and Zeckhauser, 2004; Ito, 2014). In this framework, consumers optimize their consumption in an ad-hoc fashion such that the marginal utility of consumption is equal to the ex-post average price. Additionally, Ito (2014) and Shin (1985) formalize models in which the response to average price is motivated by the costs of obtaining the necessary price information to maximize welfare in the standard framework. Despite these conceptual advances, empirical evidence has not provided a conclusive answer to which price is perceived in residential water demand.

2.2. A heuristic approach

In this paper, the restrictive assumptions inherent in Hewitt and Hanemann's (1995) model, and subsequent structural representations of water demand, are relaxed by assuming that the average household does not actively seek out information from the water utility about rate changes nor do they habitually monitor their water use throughout the billing period. Conceptually, this model is plausible because the primary means of communication between the water customer and the utility is the periodic water bill. Thus, this framework implies that the customer does not respond to the price she is charged in the current time period, rather she updates this month's consumption based on her utility bill for water use in the previous month. In effect, this model is nested within Eq. (4) with the following assumption on the customer's conditional expectation of price:

$$E[p_t|\mathbf{w}_{t-1}] = p_{t-1} + S_t + \mu_t \tag{5}$$

where p_t is the price a customer uses to make consumption decisions in period t, $\mathbf{w}_{t-1}$ is the customer's consumption last billing period, S_t allows for known differences in seasonal water usage, and μ_t allows for error in the customer's prediction of her price due to unexpected shocks to water use.

This model of price perception has several attractive features for residential water demand: (1) in the absence of a well-publicized rate change, customers are generally informed about changes to their rate structure through their water bill after

[1] In a model of quasi-linear utility, there are no income effects on water consumption. The income effect, likely to be small for residential water consumption, is typically important in estimating price elasticity for tiered rate structures through infra-marginal rate changes affecting virtual income. See Olmstead et al. (2007) for sufficient treatment of the effect of virtual income on price elasticity in water demand under block rates. Following Borenstein (2009) and Ito (2014), this analysis assumes the effect of income is negligible and proceeds with quasi-linear utility as a plausible model of residential water demand.

the changes have taken place; (2) utility bills rarely include more information than the charges incurred by consumption in each block, so this model assumes imperfect information about the full price schedule and block endpoints; and (3) it is plausible that customers learn about their consumption habits ex-post and alter their behavior for the next billing period. Further, this framework is flexible enough to accommodate consumption behavior that responds to either marginal, expected marginal, or average price. In fact, using last month's as a heuristic for this month's consumption requires minimal information costs. Empirically, many researchers have modeled water demand with lagged price variables in an informal fashion to minimize the effect of endogeneity and avoid contemporaneous correlation with consumption (Arbués et al., 2003; Ito, 2014; Renwick and Archibald, 1998; Wichman et al., 2014).

3. Empirical strategy

To avoid common problems of simultaneously determined price variables, I approach the question of perceived price with a quasi-experimental model that produces partial price effects as a function of treatment assignment, rather than explicit inclusion of a price variable in the regression. Nataraj and Hanemann (2011) and Klaiber et al. (2014) are the only researchers to develop a quasi-experimental strategy to estimate a causal effect of price on water demand. Nataraj and Hanemann estimate a regression discontinuity model to exploit the introduction of an additional price block among customers just above and just below the new block cut-off. The treatment effect is interpreted as an elasticity estimate of −0.12 for a price increase of nearly 100% (Nataraj and Hanemann, 2011). But, this interpretation is confounded by the co-movement of average and marginal price. By performing back-of-the-envelope calculations with the price schedule and statistics reported in Nataraj and Hanemann's analysis, the increase in average price is roughly 10% for the difference between treatment and control groups. This price increase translates to an average price elasticity estimate of −1.16 which is a plausible response for high volume water customers with a large proportion of extraneous water use.[2] Klaiber et al. (2014) use seasonal variation in marginal prices to assess heterogeneous responses to price. However, they do not consider alternative measures of price perception. These studies make an important advancement in the literature on perceived price by employing quasi-experimental methods to elicit a causal effect of price while avoiding potential sources of bias.

3.1. Treatment assignment

The notion that water customers respond to last month's bill as a proxy for this month's prices allows for the analysis of a unique natural experiment. Since Chapel Hill water customers are segmented into one of three billing cycles, residential customers in different cycles receive their utility bill at different points each month. If customers respond only to the information provided in their monthly bill, it is possible that customers in adjacent billing cycles would respond to different price information if a change in the rate structure occurred between the bill dates for each cycle.[3]

To illustrate the assignment of treatment status, consider two identical households who do not actively seek out rate information from the utility—households A and B. Household A is billed at the end of the month and household B is billed in the beginning of the month. If a rate change occurs on the first day of the month, household A will not recognize that rates have changed until they receive their next month's bill. Household B, however, will receive information about the new rate structure from this month's bill and update consumption in the current period based on the new rates. Thus, two otherwise identical households will make decisions on water use during similar time periods based on different price information. The differential consumption behavior between these two households can be interpreted as the effect of the new price structure since both households faced similar weather conditions and the same exogenous demand shocks in the overlapping billing period.

This anecdote describes the assignment of treatment and control status in this analysis. I focus on two periods—the month before and the month after tiered block rates were introduced for Chapel Hill water customers. I restrict my analysis to two sequential billing periods for two billing cycle groups. Within the time frame of the study, the first cycle group received monthly water bills on August 29th, September 28th, and October 31st in 2007 which results in the number of days between bills at 30 and 33, respectively. The second cycle group received bills on September 9th, October 10th, and November 7th in 2007 resulting in the number of days between bills at 33 and 28 respectively.[4] The possible divergence in billing lengths

[2] Using the difference between Nataraj and Hanemann's estimate of 43 ccf for pre-treatment water-use and 35 ccf for a control estimate within the tightest band of RD models (note: 1 ccf = 100 cubic feet = 748 gal), average price is calculated at \$1.56 for the treatment group and \$1.42 for the control group. The percent change in consumption after the treatment effect is −5.1ccf/43ccf = −0.119 %. The average price elasticity was estimated by dividing the percent change of the treatment effect by the percent change in average price between treatment and control groups.

[3] There are roughly 10 days between the bill dates for each of the billing cycles and subsequently anywhere from 1 to 10 days between the date on which the meter was read and the date on which the bill was generated, depending on the account location within the meter route. On average, sequential cycles have overlapping consumption of 20 days. The average bill length is designed to be between 27 and 30 days and remains consistent throughout the year, however the exact length varies due to the inability to read meters on weekends, among other constraints.

[4] While there is variation in the number of days between bills among different billing cycles, it is unlikely that this variation significantly affects the amount of billed consumption within a household's billing period. Bill dates represent the point at which bills are mailed to customers and thus subject to the utility's administrative work schedule. In fact, it is likely that this variation is driven by the inability to mail water bills on the weekend. Discussions with utility officials suggest that billing periods remain relatively uniform over the year in order to maintain consistency in billed amounts from period to period within customers.

Table 1
Summary statistics.

	Obs	Mean	Std. Dev.	Min	Max
Volume (1000 gal/month)	96,702	5.28	5.73	0	191.00
Total bill ($/month)	96,702	61.73	60.47	9.84	2425.38
Average price ($/1000 gal)	96,702	13.77	4.85	3.08	48.45
Marginal price ($/1000 gal)	96,702	8.11	2.28	1.98	17.21
Maximum temperature (°F)	96,702	82.06	6.04	71.10	94.08
Evapotranspiration (in)	96,702	4.30	0.93	3.00	6.24

Notes: Summary statistics are for September and October household consumption from 2006, 2007, and 2008. Consumption and bill data was obtained from customer billing records. Price variables were calculated from billing records and utility rate sheets. Weather variables were obtained from the NC State Climate Office.

between the two groups indicate that changes consumption could simply reflect the relatively shorter billing period for the latter group in the second period. However, this effect works against the main findings of the paper, thus making the results presented here conservative estimates of the true effect. Since each bill reflects water use in the preceding month, the first group's October 31st bill will be the first instance in which rates have been calculated under the new rate structure; hence, this group's October consumption will reflect the old rates despite the fact that they are being charged on the new rate structure. Alternatively, customers in the second cycle group will have received their October 10th bill and updated their consumption to reflect the new rates, thus consumption observed on the November 7th bill will be commensurate with the rate change. Households in the first group are designated as control households and households in the second group are designated as treatment households.

3.2. Data

For this analysis, a panel of monthly billing data for all residential customers from January 2006 to December 2008 was obtained from OWASA. Each residential billing record contains quantities and charges incurred by water and sewer use. Meters are read monthly and consumption is rounded to the nearest thousand gallons for billing purposes.[5] OWASA assumes the amount of water and sewer usage is identical each month. The billing data are supplemented with temperature and evapotranspiration information from the North Carolina State Climate Office. Lastly, information about rates and outreach to customers regarding rate changes was obtained from OWASA. Summary statistics for consumption, prices, and weather are reported in Table 1. Only customers residing in single-family dwelling units who did not change premises within the time frame of the study, as indicated by the billing records, are analyzed. Customers with consumption billed through irrigation meters are not included in the data set.[6] University accounts, which are billed on a separate cycle, are also removed. In the final sample of customers, there are 10,249 household consumption observations in September 2007 and 10,435 in October 2007.

3.3. Overview of rate structure change

Increasing block rates were introduced for OWASA customers on October 1st, 2007. Prior to the rate change, residential customers paid a uniform price for both water and sewer usage. After the rate change, the same customers faced an increasing five-block rate structure. In April 2007, OWASA mailed a brochure about the introduction of block rates to customers. In addition, customers were encouraged to attend a public hearing on the new rate schedule with the Board of Directors in May 2007. These rates were introduced to help meet revenue needs as well as to encourage conservation among high-volume users while allowing water bills to remain affordable (Orange Water and Sewer Authority, 2007). Incidentally, below-average rainfall in the summer of 2007 resulted in severe drought conditions throughout the fall of 2007 and into the spring of 2008. To encourage conservation, OWASA implemented voluntary watering restrictions on September 27, 2007 and mandatory watering restrictions on October 18, 2007 which remained in effect until the spring of 2008. Due to the delayed nature of billing cycles, which is central to the identification of this analysis, the treatment group faced a longer exposure to drought restrictions than did the control group. This effect would pose significant concerns for the accuracy of this research if treatment households reduced consumption by more than the control households. However, results show that the treatment effect moves in the opposite direction of the drought restrictions. This finding alleviates concerns about confoundedness and allows for the treatment effects to be interpreted as lower bounds. In addition, FRD models are estimated with data from the treatment group only to examine the robustness of the results to this potential threat to identification.

The rates customers faced before and after the rate change are presented in Table 2 and Fig. 1. The marginal price for water and sewer rates was $9.17 for all units of consumption prior to October 1st. After the rate change, customers paid

[5] Rounding to the nearest thousand gallons of water consumption results in classical measurement error. Thus, it is likely that there is attenuation present in parameter estimates.

[6] Water billed through an irrigation meter is charged on a separate rate schedule without sewer rates and comprises less than one percent of the billing observations in this sample, thus removing these observations does not exclude households who have in-ground irrigation systems or strong preferences for outdoor water use.

Table 2
Changes in marginal and average prices before and after rate change.

Volume (gal)	Before rate change		After rate change		%Δ MP	%Δ AP
	Marginal (price)	Average (price)	Marginal (price)	Average (price)		
1000	$9.17	$9.17	$6.14	$6.14	−33.04%	−33.04%
2000	$9.17	$9.17	$6.14	$6.14	−33.04%	−33.04%
3000	$9.17	$9.17	$6.14	$6.14	−33.04%	−33.04%
4000	$9.17	$9.17	$8.86	$6.82	−3.38%	−25.63%
5000	$9.17	$9.17	$8.86	$7.23	−3.38%	−21.18%
6000	$9.17	$9.17	$8.86	$7.50	−3.38%	−18.21%
7000	$9.17	$9.17	$9.69	$7.81	5.67%	−14.83%
8000	$9.17	$9.17	$9.69	$8.05	5.67%	−12.21%
9000	$9.17	$9.17	$9.69	$8.23	5.67%	−10.25%
10,000	$9.17	$9.17	$9.69	$8.38	5.67%	−8.62%
11,000	$9.17	$9.17	$9.69	$8.50	5.67%	−7.31%
12,000	$9.17	$9.17	$11.62	$8.76	26.72%	−4.47%
13,000	$9.17	$9.17	$11.62	$8.98	26.72%	−2.07%
14,000	$9.17	$9.17	$11.62	$9.17	26.72%	0.00%
15,000	$9.17	$9.17	$11.62	$9.33	26.72%	1.74%
16,000	$9.17	$9.17	$11.62	$9.47	26.72%	3.27%

Notes: Marginal price is defined as the dollar amount paid for the next 1000 gal of water consumed for both water and sewer use. Average price is defined as the total bill for water and sewer use without inclusion of base service fees divided by the amount of water consumed.

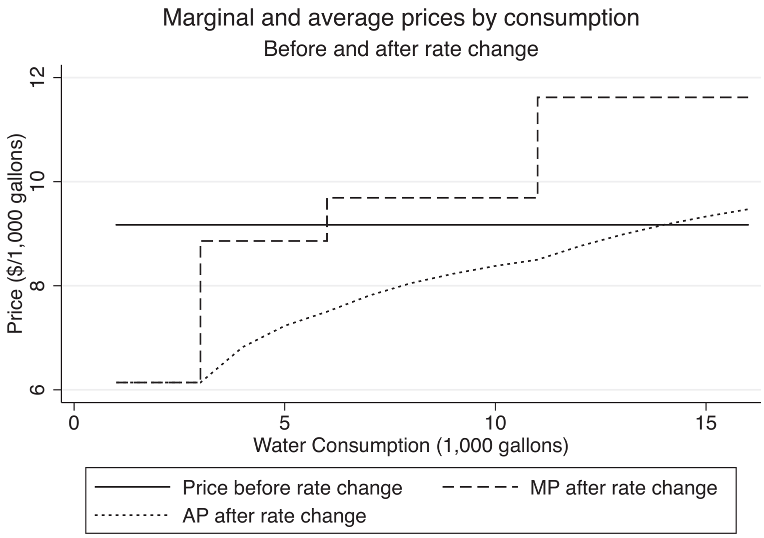

Fig. 1. Marginal and average prices before and after price change.

$6.14 for the first 3000 gal, $8.86 for the next 3000 gal, $9.69 for consumption between 7000 and 11,000 gal inclusive, $11.62 for consumption between 11,000 and 16,000 gal inclusive, and $17.21 for consumption beyond 16,000 gal.[7] Average price is defined as the total volumetric charge for water and sewer use divided by the ex post level of consumption and marginal price is defined as its "local" counterpart in that it is assumed consumers only respond to the highest marginal price they face within the block rate structure. This formulation of average price represents a proxy for the total bill that a customer would observe after their consumption period. Descriptive statistics for both variables are also presented in Table 2.[8] A customer who consumed less than 7000 gal before and after the rate change saw a decrease in both the marginal and average price she faced in each time period. A customer who used between 7000 and 14,000 gal each period, however, faced an increase in marginal price but a decrease in average price. This anomaly occurred because the price paid for the first units of water consumed in the month after the rate change were charged at a lower rate than the uniform price prior to the rate change.

[7] Due to rounding of consumption to the nearest thousand gallons, the rate schedule is defined in regards to discrete blocks of 1000 gal per month. If a customer displayed 3499 gal of monthly water use, for example, she would be billed at the rate for 3000 gal.

[8] Volumetric wastewater charges are included in the price variables in the analysis. While the rate structure for water changed from uniform rates to increasing block rates, the charges for wastewater remained uniform throughout the time frame of this study.

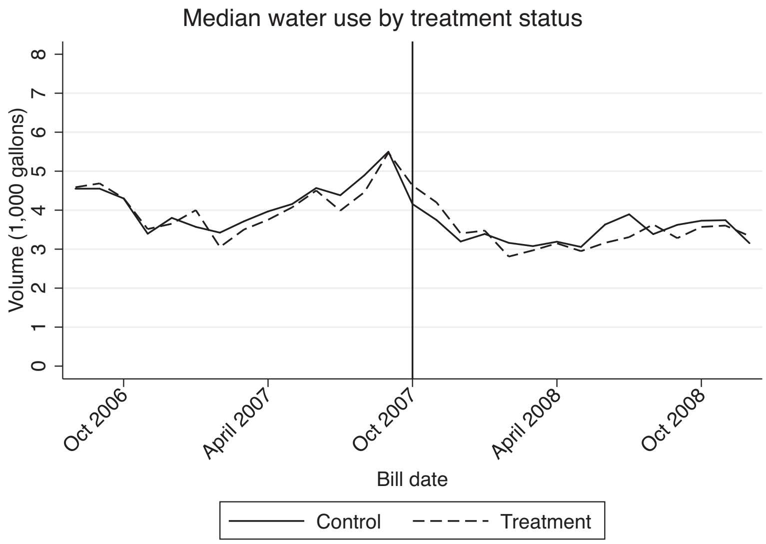

Fig. 2. Median monthly water consumption for treatment and control households.

3.4. Difference-in-differences

First, I specify a DD model to estimate an overall effect of the rate structure change on consumption. This initial specification, similar to that of Nataraj and Hanemann (2011), takes the form:

$$w_{it} = \beta_1 post_t + \beta_2(treat_i \times post_t) + Z_{it}\theta + \alpha_i + \epsilon_{it} \tag{6}$$

where w_{it} is the quantity of water consumed by household i in month t in thousands of gallons, $post_t$ is a dummy variable equal to 1 if the period of consumption is after the rate change and 0 otherwise, $treat_i$ is a dummy variable equal to 1 for treatment households, Z_{it} is a vector of control variables, α_i are the household fixed effects, and ϵ_{it} is the residual error term. If the common trend assumption between treatment and control groups is satisfied, β_2 will represent the average causal effect of the change in rate structure on the consumption of treatment households.

Fig. 2 illustrates consumption patterns by treatment status for monthly consumption at the 50th percentile. As shown, there is strong co-movement prior to October 2007. At the time of the rate change, represented by the vertical line, the difference in trends between treatment and control households diverges in both magnitude and slope. This graphical analysis evidences common trends prior to the rate change and suggests that the change in the rate structure has a positive effect on consumption for treatment households. While there appears to be a slight deviation in the trends directly before the rate change, that this effect is sustained for the succeeding two billing periods suggests that there is a positive response among treated households.

3.5. Difference-in-difference-in-differences

Next, I exploit the assignment of billing cycles by applying a difference-in-difference-in-difference (DDD) strategy similar to Gruber (1994) and Davidoff et al. (2005) to assess the relative response to the introduction of tiered rates among different levels of consumption. The baseline model is,

$$w_{it} = \beta_1 post_t + \beta_2(treat_i \times post_t) + \sum_{j=2}^{10} \gamma_j(post_t \times D_i^j) + \sum_{j=2}^{10} \delta_j(treat_i \times post_t \times D_i^j) + Z_{it}\theta + \alpha_i + \epsilon_{it} \tag{7}$$

where D_i^j is a dummy variable equal to 1 if average fall consumption in the year prior to treatment is in the jth decile and 0 otherwise; the other variables are the same as in Eq. (6).[9]

[9] The D_i^j term is determined by the decile of mean consumption for September, October, and November consumption in 2006 across all households. This allows the assignment of a decile group that reflects typical, pre-treatment, fall consumption for each household. All models were estimated with decile groups assigned by mean fall consumption in 2006 and 2008 to (1) allow for a more robust assessment of a consumer's typical fall consumption, (2) mitigate the effect of mean reversion for a single month with unusually high consumption, and (3) control for decreasing medium-run trends in water consumption observed at the utility-level. These models produced qualitatively similar results and are available from the author upon request. Pre-treatment decile groups are preferred since all households eventually received the treatment when using 2008 consumption in the assignment of decile groups.

Table 3
Difference-in-difference-in-difference average treatment effects by decile.

Decile	Before rate change			After rate change			Diff.	Std. Err.
	Mean	Std. Err.	Obs	Mean	Std. Err.	Obs		
A. Water consumption for treatment households by decile								
1	2.578	(0.173)	526	2.297	(0.130)	535	−0.281	(0.214)
2	2.679	(0.120)	405	2.414	(0.101)	408	−0.265	(0.156)
3	3.532	(0.156)	496	3.004	(0.096)	500	−0.528	(0.178)
4	4.198	(0.201)	516	3.662	(0.117)	521	−0.535	(0.224)
5	4.631	(0.143)	583	4.198	(0.112)	586	−0.433	(0.180)
6	5.352	(0.190)	457	4.862	(0.168)	458	−0.490	(0.253)
7	6.947	(0.240)	624	5.978	(0.179)	630	−0.969	(0.296)
8	7.426	(0.227)	476	6.608	(0.161)	479	−0.819	(0.274)
9	9.755	(0.309)	584	7.64	(0.225)	591	−2.116	(0.377)
10	17.535	(0.483)	757	12.37	(0.300)	759	−5.166	(0.553)
B. Water consumption for control households by decile								
1	2.677	(0.186)	575	2.133	(0.141)	562	−0.543	(0.231)
2	3.389	(0.192)	453	2.785	(0.156)	442	−0.603	(0.246)
3	4.403	(0.266)	506	3.175	(0.126)	503	−1.228	(0.277)
4	4.953	(0.300)	551	3.666	(0.116)	548	−1.287	(0.294)
5	5.804	(0.241)	556	4.431	(0.156)	559	−1.373	(0.280)
6	6.525	(0.256)	520	5.039	(0.172)	512	−1.486	(0.303)
7	7.652	(0.319)	592	5.676	(0.186)	586	−1.976	(0.357)
8	8.296	(0.298)	426	6.156	(0.179)	423	−2.140	(0.337)
9	10.747	(0.360)	451	8.336	(0.336)	446	−2.411	(0.492)
10	17.005	(0.841)	395	12.03	(0.524)	387	−4.979	(0.966)
C. DDD Average treatment effect								
1	0.262	(0.315)						
2	0.339	(0.286)						
3	0.700	(0.322)						
4	0.751	(0.368)						
5	0.940	(0.325)						
6	0.996	(0.394)						
7	1.007	(0.461)						
8	1.321	(0.431)						
9	0.295	(0.608)						
10	−0.187	(1.013)						

Notes: Treatment status is determined by billing cycle. Households are assigned to decile groups based on mean consumption in fall of 2006. DDD mean effects are calculated by subtracting the difference estimates in panel B from those in panel A for each decile group.

The model illustrated in Eq. (7) differs from traditional triple-difference models in that it allows for household-specific fixed effects which prevent the inclusion of time-invariant regressors.[10] The intuition of this equation, however, remains the same: β_1 absorbs any utility-specific changes over time, β_2 is the treatment effect for the omitted decile dummy, γ_j controls for decile-specific changes over time, and the set of third-level interaction terms, δ_j, captures the DDD treatment effect of water consumption, conditioned on the decile, relative to control households after the rate structure change. The assumption necessary for clean identification of a treatment effect across deciles is that there be no contemporaneous shock that affects the relative response of treatment households in the same time period as the rate change (Gruber, 1994). This assumption is plausible since treatment status is based on billing cycles that have overlapping consumption periods in a small geographic area serviced by the same utility. Thus, any variation in weather, utility-specific effects such as drought conservation programs, or other exogenous shocks to demand are likely to affect all households uniformly within the decile. As mentioned previously, the main threats to this assumption are (1) the increased exposure of treatment households to drought restrictions and (2) the relative decrease in billing period length among treatment households. These confounding effects are examined through several robustness checks and since positive treatment effects are found, the main findings of the DDD models are interpreted as lower bounds of the true treatment effect.

To quantify heterogeneous treatment effects, means for the treatment and control groups are presented in Table 3, which displays average consumption in each decile before and after the rate change for treatment and control households. As shown, mean consumption decreased for each decile after the rate change. For the first nine deciles, mean consumption decreased by a smaller amount for treatment households than it did for control households. Only in the tenth decile did consumption decrease by an amount larger than that of control households. These dynamics are captured in Panel C of Table 3, which are changes in consumption conditional on consumption decile for treatment households relative to control households in response to the rate change. The DDD effects indicate that for average consumption less than 11,000 gal, the

[10] The triple-difference moniker is typically reserved for specifications with an additional group of controls, whereas I simply adopt the terminology to explore heterogeneous treatment effects.

change in the rate structure increased relative consumption between treatment and control groups. For households in the tenth decile, the rate structure change reduced relative consumption. The standard errors, however, indicate the treatment and control households in the top and bottom 20% of the distribution are statistically similar, though regression will improve upon the precision of these effects.

3.6. Regression discontinuity framework

Lastly, I employ a fuzzy regression discontinuity (FRD) framework in the spirit of Nataraj and Hanemann (2011) to examine consumption behavior at a discontinuous block endpoint due to the introduction of tiered pricing in Chapel Hill. Nataraj and Hanemann assess the treatment effect for water customers above and below the price kink point for the addition of a third price block at bi-monthly consumption levels of 40 ccf, while I assess one point along the distribution of consumption (7000 gal per month) at which the change in marginal and average prices move in opposite directions. The FRD model has the same general form as the DD model developed in Eq. (6), however, treatment status is predicted by whether average fall consumption in 2006 was observed at or above 7000 gal. Customers in this consumption range are expected to face an increase of 5.7% in marginal price and a 14.8% decrease in average price due to the rate change. Thus, the treatment effect, β_2, from local linear regression of Eq. (6) around 7000 gal would be negative if customers respond to marginal price and positive if customers respond to average price, relative to control households below the discontinuity.

Typically, regression discontinuity models require the forcing variable to lie either above or below a particular threshold (Imbens and Lemieux, 2008). In this paper, the forcing variable is whether a water customer's average fall consumption in 2006 (the year before the rate change) exceeds 7000 gal. The 7000 gal indicator is chosen as the cutoff since that is the point at which marginal and average prices diverge under the new rate schedule. Since the forcing variable is not identified with contemporaneous consumption, it has the benefit of being plausibly exogenous within the study period. But, the imprecise nature of this assignment results in a fuzzy assignment to treatment (i.e., as a household's historical fall consumption approaches the cut-off, the probability of treatment assignment does not jump cleanly from zero to one at the cutoff (Imbens and Lemieux, 2008)).

The appropriateness of an FRD methodology in this context is highlighted graphically in Fig. 3 which depicts mean water consumption in 50 gallon bins relative to the distance from the 7000 treatment cutoff under varying bandwidth levels.[11] Additionally, I examine the distribution of the forcing variable according to McCrary (2008) to test for discontinuities at the 7000 gal cutoff for the same set of bandwidths. This procedure is performed by fitting kernel densities to either side of the cut-off as well as the entire domain of the forcing variable. For all bandwidths considered, there is an observable jump in consumption at the cutoff which provides evidence that there is a discontinuous change in consumption behavior due to treatment. But, there is no observable change in the distribution of the assignment variable at the point of discontinuity for the density of the forcing variable. While the magnitude of the consumption discontinuity wanes as the bandwidth increases in Panels A, C, E, and G, the graphical analysis presented in Fig. 3 provides convincing evidence that the regression discontinuity specifications capture a positive and plausible response to the treatment.[12]

4. Results and discussion

In this section, I present empirical results that identify causal effects of price across the distribution of residential water consumption. DD models indicate a positive average treatment effect arising from the change in rate structure. Additionally, the DDD specification provides evidence that most residential water customers respond to average price rather than marginal price as determined by their previous bill. To examine this effect more closely, FRD methods are applied at a point of divergence in marginal and average price and reinforce the notion that customers are indeed responding to average price. Finally, I assess the robustness of these results by varying sets of controls and assigning false treatment status. The results of falsification tests imply that the DDD and FRD models capture a valid response to changes in average price.

4.1. Difference-in-difference estimation results

First, I estimate a traditional DD model to capture the overall treatment effect of the rate structure change on consumption. The results of this model are presented in Column 1 of Table 4. The treatment effect, estimated at 433 gal per month, indicates that the rate structure change induced an overall increase in consumption due to the lower rates for the first units of consumption. This effect is significant at the 1% level. While one of the reasons that increasing block rates are typically adopted is to encourage conservation among high users, this result indicates that lower prices among low-users exhibits a perverse effect of raising aggregate consumption. This positive effect also assuages concerns that the treatment effect is estimating the effect of something other than the change in the rate structure since the primary sources of concern—seasonal

[11] Bandwidth refers to the absolute distance (in units of the forcing variable) from the discontinuity within which observations are included in the regression (McCrary, 2008).

[12] For a more detailed discussion on the appropriateness of regression discontinuity (and the difference between RD and FRD), see Imbens and Lemieux (2008), Lee and Lemieux (2010), McCrary (2008), and Angrist and Pischke (2009).

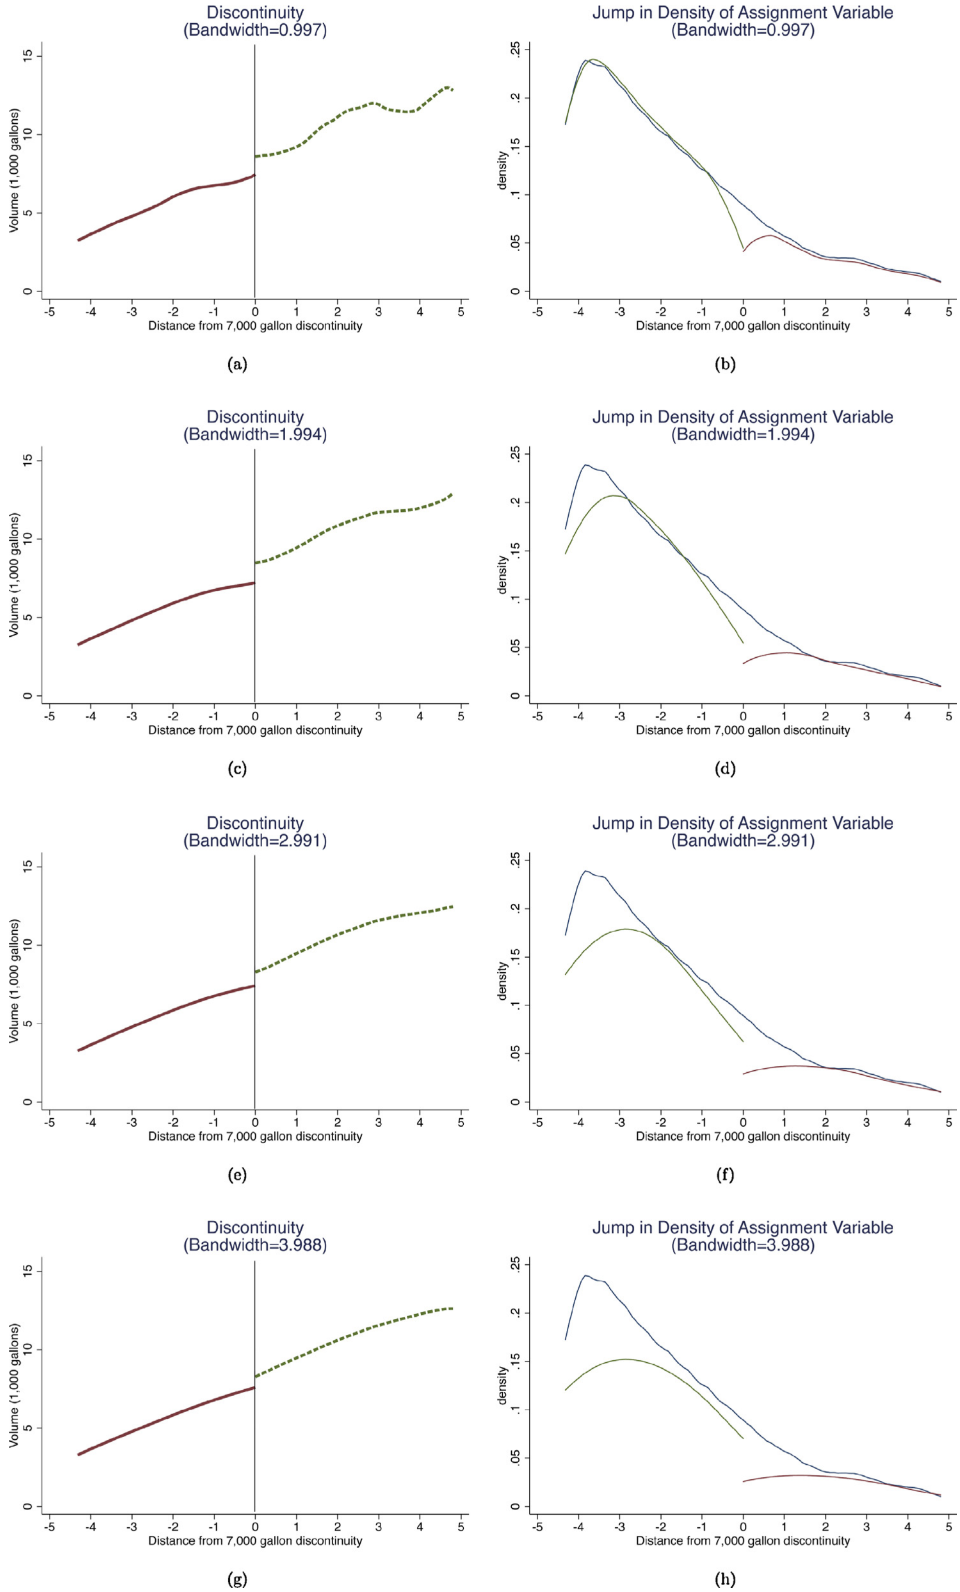

Fig. 3. Regression discontinuity consumption at 7000 gal cut-off.

Table 4
Difference-in-difference regression results.

	(1) Fixed Effects	(2) Random Effects	(3) Random Effects	(4) Random Effects
Post	-1.770^{***}	-1.758^{***}	3.410	2.376
	(0.079)	(0.081)	(3.887)	(3.109)
Treat		0.177	1.541	1.541
		(0.137)	(1.035)	(1.035)
Treat × post	0.433^{***}	0.410^{***}	0.904^{**}	1.747^{*}
	(0.111)	(0.111)	(0.388)	(1.012)
Evapotr			3.061	
			(2.301)	
Maxtemp				0.632
				(0.475)
Hausman test statistic:	143.845			
(p-value)	(<0.001)			
Observations	21,094	21,094	21,094	21,094
Number of households	10,654	10,654	10,654	10,654
Within R-squared	0.070	0.070	0.070	0.070

Notes: Dependent variable is monthly water consumption in thousands of gallons. Fixed effects are at the household level. Robust standard errors are clustered at the household level. The Hausman test is adjusted for heteroskedasticity and within-household error correlation.
* Significance at the 0.10 level.
** Significance at the 0.05 level.
*** Significance at the 0.01 level.

changes in consumption, conservation initiatives, and divergent billing durations between groups—would affect treatment consumption negatively.

Since weather varies only at the utility level, temperature and evapotranspiration effects are perfectly collinear with the *post* variable and the household fixed effects in a two-period framework. Thus, I also estimate the DD model in a random effects (RE) framework to assess whether the treatment effect is influenced by weather by exploiting cross-sectional variation.[13] The baseline random effects model in Column 2 of Table 4 is statistically similar to Column 1.[14] In Columns 3 and 4, I sequentially introduce evapotranspiration and maximum temperature to avoid further collinearity (since evapotranspiration is a function of temperature levels). Neither of the estimated parameters on weather regressors are statistically different than zero. Since RE models are inconsistent, these estimates should be interpreted with caution. But, the RE models provide justification for moving forward with the fixed effects specification without inclusion of weather controls.

4.2. *Difference-in-difference-in-difference estimation results*

The estimation results of the DDD model described in Section 3 are presented in Column 1 of Table 5. The coefficients listed are the third-level interaction effects diagrammed in Eq. (7). These effects represent the adjusted consumption response for treatment households relative to control households in each decile.[15] For the second through eighth decile interactions, the consumption response is positive and jointly significant at the 1% level, while the ninth decile is positive and significant at the 5% level. Due to aggregate seasonal decreases in water use, this result implies that the treatment households reduced consumption by an amount less than that of control households in the same decile. The insignificance in the first decile is not surprising since habitually low-use customers cannot cut back consumption in response to a rate increase and they likely have little preference to increase consumption significantly in response to lower rates. Conversely, the coefficient for the tenth decile of consumption is negative, though not significantly different from zero. Intuitively, this decile group faced increases in both marginal and average price so a negative effect is expected.

Mean volume for each decile prior to the rate change is presented in Column 2 of Table 5. In order to present results that compare this treatment effect with other studies, own-price elasticity of demand estimates are computed based on the percent change in consumption and the percent change in marginal and average price for treatment relative to control households. Elasticity estimates are presented in Columns 3 and 4 of Table 5. Standard errors for the elasticity estimates are simulated with 1000 draws using the Krinsky and Robb (1986) methodology assuming elasticities are distributed normally and centered at the calculated elasticity estimate with the standard deviation derived from the estimated

[13] To identify a comparable treatment effect in this framework an indicator for treatment households must be included to absorb unobserved variation across treatment and control groups since this variation was previously captured by the household-specific fixed effect.
[14] A Hausman test (adjusted for heteroskedasticity and within-household error correlation) rejects the hypothesis with p-value<0.001 that the strong exogeneity assumption holds, which is necessary for the RE model to be consistent. As such, the RE models only provide suggestive evidence of the implications of weather on treatment effects.
[15] Rather than reporting nominal regression parameters, I report the overall treatment interaction effect for the third-level DDD estimates relative to the omitted category. This allows for a more intuitive interpretation. A full set of nominal regression parameters is available from the author upon request.

Table 5
Difference-in-difference-in-difference regression results and elasticity estimates by decile.

	(1) Adj. DDD Estimate	(2) Volume (1000 gal)	(3) MP elasticity Estimate	(4) AP elasticity Estimate
Treat × post × D1	0.287	2.578	−0.336	−0.336
	(0.208)	(0.173)	[1.015]	[0.890]
Treat × post × D2	0.380***	2.679	−0.429	−0.429
	(0.032)	(0.120)	[0.290]	[0.185]
Treat × post × D3	0.664***	3.532	−0.569	−0.734
	(0.067)	(0.156)	[0.471]	[0.273]
Treat × post × D4	0.807***	4.198	−5.684	−0.750
	(0.088)	(0.201)	[0.592]	[0.305]
Treat × post × D5	0.968***	4.631	−6.182	−0.987
	(0.050)	(0.143)	[0.399]	[0.177]
Treat × post × D6	1.101***	5.352	−6.088	−0.972
	(0.062)	(0.190)	[0.513]	[0.189]
Treat × post × D7	1.134***	6.947	−4.828	−1.103
	(0.091)	(0.240)	[0.523]	[0.194]
Treat × post × D8	1.252***	7.426	2.974	−1.139
	(0.075)	(0.227)	[0.488]	[0.158]
Treat × post × D9	0.434**	9.755	0.784	−0.514
	(0.189)	(0.309)	[0.326]	[0.212]
Treat × post × D10	−0.083	17.535	−0.018	−0.048
	(0.395)	(0.483)	[0.201]	[0.221]
Observations	20,884			
Number of Households	10,542			
Within *R*-squared	0.128			

Notes: Dependent variable is monthly water consumption in thousands of gallons. Households are assigned to decile groups based on mean consumption in fall of 2006. Time dummy, second-level decile interactions, and constant term are omitted. Fixed effects are at the household level. Parameter estimates presented in Column 1 are the adjusted third-level DDD interaction terms. Elasticity estimates are calculated for the mean consumption at the decile group and price changes relative to prices prior to the rate change. Robust standard errors in Column 1 are clustered at the household level; standard errors of the mean are presented in parentheses in Column 2. Standard errors presented in brackets in Columns 3 and 4 represent are simulated with 1000 draws using the Krinsky and Robb (1986) methodology assuming elasticities are distributed normally and centered at the calculated elasticity estimate with standard deviation derived from the estimated variance-covariance matrix.
*Significance at the 0.10 level.
** Significance at the 0.05 level.
*** Significance at the 0.01 level.

variance-covariance matrix. For the first two deciles, the marginal and average price elasticity estimates are equivalent because prices are identical for this level of consumption. Marginal price elasticity estimates for the third through seventh decile range from −0.569 to −6.182, while average price elasticity estimates range from −0.734 to −1.103. Within this central region of the consumption distribution, marginal price estimates are implausibly large as all but one lie outside the range of previous elasticity estimates, though average price estimates are well within this range (Espey et al., 1997; Dalhuisen et al., 2003).[16]

The coefficients for the eighth and ninth decile are most illuminating with respect to identifying the price to which consumers respond. For this subset of customers, the marginal price increased roughly 5% for treatment households while average price decreased between 7% and 15% relative to control households. The estimated regression parameters for water customers in this group exhibit a response that corresponds to a positive marginal price elasticity, ranging from 0.784 to 2.974, and a negative average price elasticity estimate, ranging from −1.139 to −0.514. Lastly, the result for the tenth decile, though not statistically different from zero, displays plausible elasticity estimates (−0.018 and −0.048) for both marginal and average price, respectively. Collectively the DDD results indicate that consumers generally respond to average price.

4.3. Regression discontinuity estimation results

To further analyze the consumption response to changes in price, the FRD analysis focuses on a point in the rate schedule with a discontinuous divergence in marginal and average price for the same level of consumption after the rate change. For consumption levels at 7000 gal per month, marginal price increases by 5.67% more than the previous uniform rate, while average price decreases by 14.83% for treatment households relative to the control households. Consider a customer just above the discontinuity—if she responds to marginal price, she would reduce consumption commensurate with the increase in marginal price relative to customers below that level; however, if she responds to average price, then her consumption response will be positive.

[16] In a meta-analysis, Espey et al. (1997) find price-elasticity estimates of residential water demand to range from −0.02 to −3.33 with an average of −0.51.

Table 6
Fuzzy regression discontinuity results.

	(1) ±5000 (gal)	(2) ±4000 (gal)	(3) ±3000 (gal)	(4) ±2000 (gal)	(5) ±1000 (gal)
A. Regression discontinuity cut-off at 7000 gal (treatment and control)					
Treat × post	0.288***	0.431***	0.351***	0.320***	0.291***
	(0.086)	(0.082)	(0.080)	(0.080)	(0.070)
Observations	16,084	12,909	9398	6071	2944
Within R-squared	0.079	0.080	0.072	0.072	0.032
B. Regression discontinuity cut-off at 7000 gal (treatment only)					
Post	−0.270***	−0.105	−0.094	−0.048	0.204***
	(0.083)	(0.078)	(0.075)	(0.073)	(0.062)
Observations	1447	1290	1104	849	471
Within R-squared	0.017	0.003	0.004	0.002	0.105
C. False regression discontinuity cut-off at 7000 gal (treatment and control)					
Treat × post	0.093	0.012	−0.093	0.048	0.066
	(0.108)	(0.103)	(0.104)	(0.101)	(0.095)
Observations	15,174	12,330	8982	5807	2822
Within R-squared	0.006	0.009	0.013	0.003	0.015

Notes: Dependent variable is monthly water use in thousands of gallons. In Panel A, treatment status is assigned to households who satisfy two requirements: (1) were on the new price schedule as determined by their billing cycle and (2) had mean fall consumption in 2006 at or above 7000 gal per month. In Panel B, only treatment households are included such that the coefficient on post is simply the FRD treatment effect within treatment households. In Panel C, false treatment status is assigned to households who had mean fall consumption above 7000 gal per month for an artificial cut-off two months prior to the rate change. Fixed effects are at the household level. Robust standard errors in parentheses are clustered at the household level.
*Significance at the 0.10 level.
**Significance at the 0.05 level.
*** Significance at the 0.01 level.

In Panel A of Table 6, local linear regression estimates of Eq. (6) around 7000 gal are presented. At a bandwidth of 5000 gal around the FRD cut-off, the treatment effect is positive and significant at the 1% level. Within 4000 gal of the cut-off, the treatment effect tends monotonically towards zero, but remains significantly positive, at 291 gal per month, within 1000 gal of the cut-off. This result indicates that customers just above the discontinuous jump in price increased consumption by 4.54% in response to the change in price relative to control households.[17] Since this effect is positive, it lends credibility to the notion that consumers respond to average , not marginal, price. For the sake of comparison to previous research, this effect is interpreted as a local price elasticity estimate of −0.31, dividing the 4.54% increase in consumption by the 14.83% decrease in average price, for treatment households after the rate change within 1000 gal of the cut-off. This elasticity estimate is slightly different than the estimate for similar consumption levels in the DDD model because the control group in the RD model is composed of households on the new price schedule with average fall consumption below the 7000 gal cut-off as well as the control households who have not yet observed the new price schedule.

4.4. Robustness checks and falsification tests

Since the primary threats to identification in the DDD model are (1) the increased exposure to drought restrictions for the treatment group due to their delayed billing cycle and (2) the duration of billing lengths among treatment and control groups, the regression discontinuity model is estimated on a subsample of treatment households only. This strategy eliminates the potential confounding effect of drought restrictions or bill length since treated households are otherwise identical, though treatment is now predicted strictly by whether fall consumption lies above or below the 7000 gal discontinuity in price. Results are shown in Panel B of Table 6. Within 5000 gal of the cut-off, the treatment effect is negative and significant. As the bandwidth decreases, the treatment effect tends monotonically towards zero, becoming positive and significant within 1000 gal of the cut-off. The initial negative effect can be interpreted as habitually low users being poor controls for habitually high users in that they do not display qualitatively similar consumption patterns. As households become more similar in consumption patterns, the treatment effect becomes insignificant. The fact that a positive and significant treatment effect is identified within 1000 gal, and statistically similar to the treatment effect in Panel A, is reassuring given the lack of statistical power in the restricted sample. Regardless, this result provides strong evidence that the identification strategy is not confounded by the contemporaneous drought restrictions.

Additionally, to test whether the estimated treatment effects in the regression discontinuity models are valid, I estimate the same FRD model in Section 4.3 for two consecutive time periods two months prior to the rate change. Thus, in this falsification test, all households face the exact same uniform rate structure. The results of the false treatment FRD models

[17] The percent change in consumption is calculated using 6413 gal as the pre-treatment mean consumption for treatment households within 1000 gal of the cut-off.

Table 7
Difference-in-difference-in-difference regression results by decile with random effects and weather variables.

	(1) Fixed Effects	(2) Random Effects	(3) Random Effects	(4) Random Effects
Treat × post × D1	0.287	0.277	0.764	1.592*
	(0.208)	(0.207)	(0.489)	(0.937)
Treat × post × D2	0.380***	0.363***	0.852***	1.680*
	(0.032)	(0.032)	(0.230)	(0.911)
Treat × post × D3	0.664***	0.679***	1.164***	1.992**
	(0.067)	(0.067)	(0.262)	(0.885)
Treat × post × D4	0.807***	0.784***	1.268***	2.096**
	(0.088)	(0.087)	(0.283)	(0.877)
Treat × post × D5	0.968***	0.957***	1.439***	2.267***
	(0.050)	(0.049)	(0.241)	(0.867)
Treat × post × D6	1.101***	1.059***	1.542***	2.370***
	(0.062)	(0.062)	(0.254)	(0.888)
Treat × post × D7	1.134***	1.082***	1.567***	2.395***
	(0.091)	(0.090)	(0.287)	(0.864)
Treat × post × D8	1.252***	1.281***	1.767***	2.595***
	(0.075)	(0.076)	(0.249)	(0.900)
Treat × post × D9	0.434**	0.377**	0.861**	1.6890*
	(0.189)	(0.187)	(0.380)	(0.883)
Treat × post × D10	−0.083	−0.125	0.364	1.192
	(0.395)	(0.390)	(0.601)	(1.031)
Evapotr			3.005	
			(2.717)	
Maxtemp				0.620
				(0.414)
Hausman test statistic:	92540.8			
(p-value)	(<0.001)			
Observations	20,884	20,884	20,884	20,884
Number of households	10,542	10,542	10,542	10,542
Within R-squared	0.128	0.043	0.043	0.043

Notes: Dependent variable is monthly water consumption in thousands of gallons. Households are assigned to decile groups based on mean consumption in fall of 2006. Second-level interactions, baseline effects, and the constant term are omitted. Fixed effects are at the household level. Parameter estimates presented are the adjusted third-level DDD interaction terms. Robust standard errors are clustered at the household level. The Hausman test is adjusted for heteroskedasticity and within-household error correlation.

* Significance at the 0.10 level.
** Significance at the 0.05 level.
*** Significance at the 0.01 level.

are presented in Panel C of Table 6. Within all bandwidth specifications, this model produces no statistically significant treatment effects and the parameter estimates are small in magnitude relative to the FRD treatment effects found previously. With the lack of a non-zero treatment effect, this falsification test indicates that the true FRD model is estimating a plausible response to a discontinuous change in price.

Finally, the DDD model is estimated in a random effects framework including a set of weather controls. I use a random effects estimator to exploit between household variation which allows for the inclusion of weather parameters without being perfectly collinear with household fixed effects and the time indicators.[18] The first column of Table 7 presents the baseline fixed effects estimates discussed in Section 4.2 for comparison. Results from a baseline random effects model with no weather controls is presented in Column 2.[19] While the RE analysis is merely suggestive, the inclusion of weather controls in the empirical specification enlarge the coefficients of interest, though all significant treatment effects in Column 1 maintain conventional levels of significance in the RE framework. The weather covariates exhibit insignificant parameters. Overall, these results mimic the intuition of the estimates in Table 6 and provide justification that the fixed effects model adequately controls for changes in weather across deciles.

4.5. Lingering empirical concerns

This subsection outlines the implications of empirical concerns that could lead to potentially biased estimates and evaluates how these concerns might affect the external validity of this study. First, the most obvious shortcoming of this analysis is

[18] Similar to the DD model, the random effects estimates rely on inclusion of baseline treatment effects and additional secondary interactions such that Eq. (7) must include time-invariant regressors at the household-level. The results of interest are presented in Table 7.
[19] The Hausman test statistic (adjusted for heteroskedasticity and within-household error correlation) rejects the consistency of the RE model with p-value<0.001.

the inability to include demographic or household characteristic information to illustrate that demographics are qualitatively similar among treatment and control groups. Since the data used in this analysis were stripped of any unique geographic identifier for consumption records other than the billing cycle, demographic and spatial comparisons between treatment and control groups could not be performed. Thus, I rely on analyzing consumption patterns between billing cycles to argue that the difference between groups is negligible and these data provide a valid quasi-experimental setting. While this argument is admittedly weak, the econometric methods used mitigate the effect of this uncertainty. The DDD methods provide nonparametric control for decile-specific effects and relax the common trend assumption necessary for clean identification of a treatment effect (Gruber, 1994). Moreover, the FRD method further avoids this potential source of bias by comparing households just above and just below a discontinuous jump in prices for both treatment and control households, as well as treatment households only. Since there is reason to believe that demographics and water use are correlated (Nataraj and Hanemann, 2011), as we approach the discontinuity, the effect of heterogeneity in household composition is mitigated. Further, all econometric models control for time-invariant household characteristics. Thus, I contend that these concerns do not contaminate the results, though examining demographic and spatial heterogeneity in this context would be a fruitful area for further research.

Additionally, the identification strategy relies on the staggered nature of utility billing cycles such that drought regulations could potentially confound results due to the treatment group's relatively longer exposure to these policies. Further, the fact that the length of the billing cycle for treatment households becomes relatively shorter than that of control households potentially biases the findings of this paper. These effects, however, work in the opposite direction of the estimated treatment effects biasing the results toward zero. Since the main treatment effects are positive, it is likely that the estimated price elasticities are lower bounds of the true elasticity. The regression discontinuity model estimated only for treatment households allows for this effect to be isolated, and results are robust to this stratification.[20]

Further, the identification strategy relies on plausibly random assignment into treatment and control groups as well as an exogenous forcing variable that delineates the cutoff in FRD models. For the former, I assign treatment status as determined by a customer's billing cycle. Billing cycles are deterministic in that they are constructed for practical convenience such that a utility employee can conveniently drive along a meter route to read meters. While there is spatial correlation between households and demographic composition within billing cycles, there is little observed difference between household water consumption within the decile of consumption for the DDD models. Further, the main regression discontinuity specification includes households from the same billing cycle in both the control and treatment groups, thus mitigating any deterministic bias that might arise from non-random assignment into the treatment group. Additionally, the household itself has no ability to change its billing cycle short of moving across billing cycle boundaries, thus it is not likely that there is any manipulation of the treatment within the time frame of the study. For the latter, since the forcing variable in the RD models exploits the implementation of a five-block tiered price schedule adopted for the utility's revenue goals, conservation among high-users, and affordability among low-income customers, it is possible that the block endpoints were chosen to explicitly align with certain indicators in the distribution of customer consumption. If this is the case, then perhaps the 7000 gal discontinuity in price was chosen to penalize household water use above the mean of the total customer base. Since there is no observable discontinuity in the distribution of water consumption (as illustrated in Fig. 3) it is assumed that this assignment is as good as random.

Finally, though the results of this paper imply that residential water customers respond to average price, it should be noted that this study relies on the introduction of an entirely new price structure and not merely a marginal change in the price level. Thus, it is possible that the variation in consumption reflects uncertainty about the new rate structure. The treatment effect, however, is generally positive. If households were reacting to uncertainty about the new rate structure, the expectation would be a negative or null effect until households learn about how their bill is calculated. Because the positive treatment effect is robust across specifications, it is not likely that consumers are responding to this uncertainty. Lastly, while this analysis was motivated by the structural assumptions of consumer demand under tiered rate schedules, the result that most consumers respond to average price is conditional upon the alternative hypothesis that the price to which consumers respond is a local marginal price, rather than the entire rate schedule as discrete-continuous choice models predict. Thus, the main findings of this research provide evidence that reduced-form models of consumer water demand should incorporate behavior commensurate with changes in average price, though a comparison of elasticity estimates from (quasi-)experimental methods to structural models of water demand is a promising area for future research.

[20] In addition, I examine DDD mean effects for treatment and control groups in September and October of 2006 prior to the 2007 drought. During this time period, a small uniform rate increase occurred for all customers, which is identified exactly according the treatment-control strategy outlined previously given the staggered nature of utility billing. All of the deciles display similar trends in reductions due to seasonality and the effects and the triple-difference mean effects are primarily small and insignificant except for the upper and lower tails of the distribution. This serves to further alleviate concerns that a confounding effect of drought restrictions or billing period length is biasing results and ultimately strengthens the main findings of this paper. These results are available from the author upon request.

5. Concluding remarks

In this paper, a conceptual framework for residential water demand that relies on a customer's consumption patterns last month as a heuristic for this month's prices is developed. Then, by exploiting the introduction of block rates and the assignment of billing cycles for residential customers in Chapel Hill, North Carolina, I identify a causal effect of price across the entire distribution of water customers using DDD techniques. In addition, FRD models are estimated to focus on an anomaly in the rate change in which changes in average and marginal price move in opposite directions for customers with consumption at 7000 gal per month.

The results of this analysis imply that residential water customers exhibit consumption patterns commensurate with changes in average, rather than marginal, price. Across the distribution of consumption, price elasticity estimates from DDD models range from −0.43 to −1.14 and households consuming between 4000 and 7000 gal per month are found to be the most responsive to changes price. Additionally, empirical estimates from FRD models identify a positive treatment effect at a point where marginal prices increased and average prices decreased in response to the new rate structure. By exploiting the divergence in average and marginal prices, I obtain further support of the hypothesis that average prices are perceived by residential water customers. Within the tightest bandwidth of regression discontinuity models at 7000 gal per month, a local price elasticity is estimated at −0.31.

The results of this paper contribute in several ways to the literature on water demand and conservation, as well as the practitioner's guidebook on managing local water resources. First, this study provides empirical evidence that the price perceived by residential water customers is the average price from a customer's previous bill. This result adds to mounting evidence in the literature that consumers facing complicated pricing structures tend not to respond in a manner that standard utility theory predicts. Second, the introduction of increasing block rates can produce a perverse effect of increasing total demand due to price decreases in the lower blocks. Lastly, this paper provides evidence of heterogenous price elasticity estimates from quasi-experimental methods that support the well-accepted notion that residential water demand is generally price inelastic, but certainly not unresponsive to changes in price.

References

Angrist, J.D., Pischke, J.-S., 2009. Mostly Harmless Econometrics: An Empiricist's Companion. Princeton University Press.

Arbués, F., García-Vali nas, M.Á., Martínez-Espi neira, R., 2003. Estimation of residential water demand: a state-of-the-art review. J. Socio-Econ. 32 (1), 81–102.

Borenstein, S., 2009. To What Electricity Price do Consumers Respond? Residential Demand Elasticity Under Increasing-block Pricing , (University of California at Berkeley Working Paper).

Dalhuisen, J.M., Florax, R.J., de Groot, H.L., Nijkamp, P., 2003. Price and income elasticities of residential water demand: a meta-analysis. Land Econ. 79 (2), 292–308.

Davidoff, A., Blumberg, L., Nichols, L., 2005. State health insurance market reforms and access to insurance for high-risk employees. J. Health Econ. 24 (4), 725–750.

Espey, M., Espey, J., Shaw, W.D., 1997. Price elasticity of residential demand for water: a meta-analysis. Water Resour. Res. 33 (6), 1369–1374.

Foster Jr., H.S., Beattie, B.R., 1981. On the specification of price in studies of consumer demand under block price scheduling. Land Econ. 57 (4), 624–629.

Gruber, J., 1994. The incidence of mandated maternity benefits. Am. Econ. Rev., 622–641.

Hewitt, J.A., Hanemann, W.M., 1995. A discrete/continuous choice approach to residential water demand under block rate pricing. Land Econ. 71 (2), 173–192.

Imbens, G.W., Lemieux, T., 2008. Regression discontinuity designs: a guide to practice. J. Economet. 142 (2), 615–635.

Ito, K., 2014. Do consumers respond to marginal or average price? Evidence from nonlinear electricity pricing. Am. Econ. Rev. 104 (2), 537–563.

Klaiber, H.A., Smith, V.K., Kaminsky, M., Strong, A., 2014. Measuring price elasticities for residential water demand with limited information. Land Econ. 90 (1), 100–113.

Krinsky, I., Robb, A.L., 1986. On approximating the statistical properties of elasticities. Rev. Econ. Stat. 68 (4), 715–719.

Lee, D.S., Lemieux, T., 2010. Regression discontinuity designs in economics. J. Econ. Lit. 48 (2), 281–355.

Liebman, J.B., Zeckhauser, R., 2004. Schmeduling (Unpublished manuscript).

McCrary, J., 2008. Manipulation of the running variable in the regression discontinuity design: a density test. J. Economet. 142 (2), 698–714.

Nataraj, S., Hanemann, W.M., 2011. Does marginal price matter? A regression discontinuity approach to estimating water demand. J. Environ. Econ. Manage. 61 (2), 198–212.

Nieswiadomy, M.L., Molina, D.J., 1989. Comparing residential water demand estimates under decreasing and increasing block rates using household data. Land Econ. 65 (3), 280–289.

Olmstead, S.M., Michael Hanemann, W., Stavins, R.N., 2007. Water demand under alternative price structures. J. Environ. Econ. Manage. 54 (2), 181–198.

Orange Water and Sewer Authority, 2007. Our New Water and Sewer Rates.

Renwick, M.E., Archibald, S.O., 1998. Demand side management policies for residential water use: Who bears the conservation burden? Land Econ. 74 (3), 343–359.

Shin, J.-S., 1985. Perception of price when price information is costly: evidence from residential electricity demand. Rev. Econ. Stat. 67 (4), 591–598.

Strong, A., Smith, V.K., 2010. Reconsidering the economics of demand analysis with kinked budget constraints. Land Econ. 86 (1), 173–190.

Wichman, C.J., Taylor, L.O., von Haefen, R.H., 2014. Conservation Policies: Who Responds to Prices and who Responds to Prescription? Center for Environmental and Resource Economic Policy, North Carolina State University, Working Paper.