Policy Research Working Paper 9207

Optimal Pricing of a New Utility Service

The Case of Piped Water in Vietnam

Quy-Toan Do
Hanan Jacoby

WORLD BANK GROUP
Development Economics
Development Research Group
April 2020

Policy Research Working Paper 9207

Abstract

As utility services expand throughout the developing world, providers must grapple with how to set prices to recover average costs. Data from a multi-year randomized pricing experiment among nearly 1500 recently-connected piped water customers in Vietnam reveal month-to-month demand persistence. Based on structural demand estimation, the authors document how endogenous preferences, if unaccounted for, can lead to low take-up and thereby threaten the financial viability of the new water utility. They also show that such demand persistence, while distinct from credit constraints, calls for pricing schemes that similarly defer payment, effectively allowing future consumers to subsidize their present selves.

This paper is a product of the Development Research Group, Development Economics. 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://www.worldbank.org/prwp. The authors may be contacted at abonfield@worldbank.org.

Optimal Pricing of a New Utility Service : The Case of Piped Water in Vietnam*

Quy-Toan Do†   Hanan G. Jacoby‡

Originally published in the Policy Research Working Paper Series on *April, 2020.* Previous versions were made on *May, 2021 and December 2021.* To obtain the originally published version, please email prwp@worldbank.org.

JEL codes: L95, D91, O18

Keywords: Endogenous preferences, two-part tariffs, take-up, willingness-to-pay.

*This study is registered in the AEA RCT Registry and the digital object identifier (DOI) is: "10.1257/rct.6589-1.0". The authors gratefully acknowledge financial support from the World Bank's Strategic Impact Evaluation Fund (SIEF), the Knowledge for Change Program (KCP), the Impact Evaluation to Development Impact (i2i) Fund, and the Research Support Budget. Tung Duc Phung and the Mekong Development Research Institute (MDRI) managed the price experiment and collected the baseline data. Trung Dang Le and RTAnalytics (RTA) collected quarterly monitoring data. Irène Hu provided excellent research assistance. We gratefully acknowledge comments from Jishnu Das, Alejandro Molnar, Claudia Ruiz, Gil Shapira and seminar participants at the World Bank applied-micro seminar. The findings, interpretations, and conclusions expressed in this work do not necessarily reflect the views of the World Bank, its Board of Executive Directors, or the governments they represent. The World Bank does not guarantee the accuracy of the data included in this work.

†World Bank: qdo@worldbank.org

‡World Bank: hjacoby@worldbank.org

1 Introduction

The provision and pricing of utilities remains a major policy issue in developing countries, where the expansion of piped water, electricity, and other large service networks is ongoing. Economists have long been concerned with how such natural monopolies should be priced, subsidized and/or regulated to ensure access to socially valuable services while covering their large capital costs (see, e.g., Laffont and Tirole 1993; Joskow 2007). A common premise in this literature is that consumer preferences are exogenous and unchanging. Once connected, however, it may take time for consumers to habituate to the new service or to learn about their preferences. As a result, current consumption may shift out future demand. If such a dynamic is sufficiently strong, and consumers are unaware of it prior to connecting, then they may substantially under-value a connection ex-ante, leading to sub-optimal take-up and even non-viability of an ex-post socially desirable utility service. In this setting, an informed policy-maker could improve welfare by pricing a new utility so as to account for endogenous preferences and consumers' lack of awareness thereof.

To quantify the endogenous formation of preferences for a new utility service, we ran a three-year price experiment in a rural commune of Vietnam's Red River Delta Region served by a single piped water provider. Our setting is fitting because households were in the process of transitioning from traditional rainwater harvesting to modern piped water, both of which systems were installed in their homes concurrently. Rainwater needs to be stored in and pumped from dedicated tanks, as well as filtered and boiled for cooking purposes, whereas piped water provides greater convenience on these dimensions as well as higher pressure for showers, house cleaning and the like. Reflecting the transition process, in the three years from the start of the water utility's operations to the start of our experiment, piped water use per customer grew at an average rate of 0.9 percent per month by volume, pointing toward a behavioral model that delivers at least an initial phase of non-stationary consumption. While our focus on piped water in Vietnam is, of course, particular, we believe that learning through experience and/or habit formation are generic features of newly introduced utility services in a wide range of settings, both contemporary and historical. For instance, Taylor and Trentmann (2011) vividly recount how bathing habits and routines in Victorian England evolved with both changes in the availability and in the pricing of piped water.

For the experiment, about 1,500 piped water customers were randomly divided into three equal groups. We then rolled out a subsidy scheme sequentially, each group receiving a 50 percent discount on the price of piped water for six months, followed by a 75 percent discount for another six months after that. At any point in time, therefore, at least one group did not receive a subsidy and could thus serve as a control. To make the changes in price throughout the experiment immediately salient, we undertook a personalized communication campaign, both at the home and at the time of monthly bill payment, which was always done in-person at a central location. Meanwhile, we collected data from the utility on monthly piped water usage and, at a quarterly frequency, from households on, among other things, the number and duration of stay of residents (including visitors) and on water-sharing practices.

In addition to a modest response to the current price, the experimental data reveal persistence: households that faced a lower water price during months in the recent past (and thus had higher usage in those months) also consumed more in the current month, *conditional* on the current month's price. We take this as evidence that households were building consumption capital, either in the form of habits or learning through experience.[1] Alternative explanations can be plausibly ruled out: lags in adjustment to new prices were obviated by our proactive information campaign; loss aversion or reference dependence, which could lead to apparent demand persistence, is inconsistent with the symmetric demand responses to upward and downward price changes observed in the data and cannot account for demand growth prior to the experiment; finally, household investments in piped water infrastructure (such as plumbing and washing machines) cannot be the source of the experimental findings given the size and duration of the price subsidies.

To quantify the welfare gains from utility pricing strategies that leverage such persistence, we next use our experimental data to structurally estimate the dynamics of piped water demand. Under our maintained assumption that, in deciding on their current piped water usage, consumers do not internalize its effect on future demand, consumption depends only on the current price and on past usage.[2] Our experiment, by generating random variation in price, not only allows us to identify the (short-run) price elasticity of piped water demand, but also delivers valid instruments, namely lagged prices, for causally estimating the effect of past con-

[1]Dupas (2014) interprets experimental evidence on demand persistence for bednets in Kenya through the lens of a structural model to demonstrate a large experience good effect. Habit formation is a less plausible source of persistence in her setting.

[2]We justify this assumption of consumer lack of awareness in detail below as, among other things, consistent with projection bias (Loewenstein et al., 2003).

sumption on current consumption. Before turning to our policy simulations, we show that the structural model successfully predicts the dynamic pattern of piped water consumption out-of-sample. Specifically, we start each household off with zero past consumption at the time they connect and use our estimated model to forecast their consumption trajectories from the month of connection up until the start of the price experiment in January 2016. Our forecasts fit the actual out-of-sample consumption trajectories well.

We next extend the paradigmatic optimal utility pricing framework to allow endogenous preferences. In the seminal treatment of two-part tariffs (Auerbach and Pellechio, 1978), a utility chooses a fixed (i.e., connection) fee and a marginal price to maximize a static welfare function defined over heterogeneous consumers who are free not to connect.[3] The utility must also earn enough revenue per customer to cover average costs. When the demand for the utility service grows endogenously and consumers are unaware of this, their ex-ante valuation of the service will be lower than their ex-post valuation. Taking this wedge into account in its pricing strategy, the optimizing utility should defer charges until the "long-run", the point at which consumption reaches its steady-state.[4] In the context of a two-part tariff, demand persistence thus militates in favor of a higher markup and a lower connection fee. Further, by replacing the connection fee with a recurring subscription fee (as exist in most cell-phone plans, for example), one that only comes into effect in the long-run, the utility can reduce the distortionary markup and potentially eliminate it altogether, thereby improving welfare.[5] Such back-loading of payment for a new utility service is a novel consequence of endogenous preference formation. Deferment of the connection fee could also be a way of providing credit to customers limited in their ability to pay upfront. Even with credit, however, a consumer may still be unwilling to take up the service, whereas her future-self would. Deferred payment thus effectively taxes the high-willingness-to-pay future consumer to subsidize her low-willingness-to-pay present-self.

The final step is to numerically solve for alternative optimal water pricing schemes in our setting, given estimated (endogenous) preferences. As a benchmark, we first consider a scenario in which the utility incorrectly assumes static or exogenous preferences, as might be obtained

[3]Increasing block tariffs, whereby the marginal price increases with usage, may further improve efficiency if the price elasticity of demand is decreasing in usage (see the discussion of "Ramsey pricing" in Wilson, 1993).

[4]A related literature on the monopoly pricing of experience goods also shows that low initial prices followed by higher prices later may be optimal in some circumstances (Shapiro, 1983; Bergemann and Välimäki, 2006).

[5]Only in a dynamic setting, such as ours, is the distinction between connection and subscription fees meaningful.

from an ex-ante willingness-to-pay elicitation (e.g., Lee et al. 2020). In this case, the optimal connection fee and marginal price combination from the utility's point of view is just the Auerbach and Pellechio (1978) two-part tariff. Next, we recalculate the two-part tariff when the utility correctly accounts for the endogenous formation of preferences as well as their customers' lack of awareness thereof. We find that substantial welfare gains can be achieved by the "sophisticated" utility, depending on its costs. Given average cost in our setting, pricing that ignores demand persistence would fail to generate enough revenue to finance the utility in the first place, whereas a sophisticated two-part tariff *would* cover the cost of this ex-post socially desirable project. Lastly, we compute the optimal deferred subscription fee, which, by allowing for zero distortionary markup, leads to a 12 percent welfare gain over two-part pricing with connection fee and markup.

This study contributes to a growing body of evidence, from both developed and developing countries, of persistence and adjustment lags in the demand for utility services, whether in response to price changes (Ito et al., 2018; Ito and Zhang, 2020; Deryugina et al., 2020) or to conservation programs (Allcott and Rogers, 2014; Costa and Gerard, 2018). Two features of our work set it apart, however. First, we focus on the case of a new utility service, in which the policy-maker faces a tradeoff between expanding the customer base and extracting more from established customers. Second, while these existing studies also rationalize their empirical findings by appealing to some form of habit formation and/or learning, they are not designed to structurally estimate endogenous preferences. In doing precisely this, we are uniquely able to make quantitative welfare statements about counterfactual pricing structures that incorporate such preferences.[6]

More broadly, in randomizing a relevant policy instrument, here the marginal price of the service, across a real population of utility customers, this research advances the "lab-in-the-field" approach to questions of utility provision and optimal pricing. In pioneering work, Lee et al. (2020) randomize electricity service connection fees across rural households in Kenya to trace out the extensive margin demand curve and compare it to the average cost curve for connections to decide the social desirability of the service. Our paper cautions that such an ex-ante willingness-to-pay elicitation alone may be insufficient for two-part pricing and, ultimately, for provision decisions when preferences for the utility service form endogenously. In this setting,

[6]In this respect, our paper extends the static demand literature focusing on how utility price policy affects consumer welfare; see, e.g., Reiss and White (2005), Ito (2014), McRae (2015), and Szabó (2015).

an ex-ante elicitation of willingness-to-pay measures consumer surplus at some implicitly assumed and hence unknown marginal price, rather than at *any* marginal price. However, the determination of optimal provision and two-part pricing requires, firstly, recovering the short-run price elasticity of demand and, secondly, requires estimating the long-run price elasticity of demand when it differs from its short-run analog. Our approach does both of these things while highlighting a source of low elicited willingness-to-pay, complementary to but distinct from credit constraints (Devoto et al., 2012b; Berkouwer and Dean, 2021).

The rest of the paper is organized as follows: We describe the experiment, data and demand persistence results in section 2. In section 3, we set up a generic model of consumer demand with intertemporally dependent preferences and, in section 4, present the structural estimation of this model. With these results in hand, in Section 5, we assess optimal pricing schemes for new utility services, and specifically for piped water in our setting. We conclude the paper in Section 6.

2 Experiment and data

Working with the Government of Vietnam, the overall objective of the experiment was to assess willingness-to-pay for piped water in a rural area where the service had recently been introduced and had begun supplanting the traditional mode of water supply, rainwater collection.

2.1 Experimental design

The price experiment took place between January 2016-May 2018 in three villages of My Huong commune in the Red River Delta region of northern Vietnam. My Huong is relatively prosperous, with average per capita expenditures at about the 80th percentile for rural Vietnam. Since the arrival of piped water in 2012, My Huong has been served by a single provider, An Thinh utility, with whom we partnered over the course of the 3-year project. Among other things, An Thinh shared their customer billing records, which include monthly water usage (in m^3) and payment, from the month of connection onward, including throughout the price experiment.

Household listing An exhaustive door-to-door listing in the three villages was completed in mid-2015, yielding a total of 1660 households, which includes multi-family living arrangements with shared kitchen and electrical meter. After matching with An Thinh's customer records, we verified that 267 households or 16 percent of the population were unconnected to piped water as of July 2015, at which time we offered them free connections (i.e., waived the connection fee). While almost all of these 267 households eventually did connect, only about half did so in time for the baseline survey (see below).

Pre-experiment price linearization Prior to the price randomization, starting in July 2015, we "linearized" the water tariff schedule once and for all to avoid having to deal with self-selection of households onto price blocks (as in, e.g., Reiss and White 2005, McRae 2015, and Szabó 2015). At the time, the official schedule had three blocks: for consumption below 10 m^3, the price was VND5,300 per m^3, increased to VND7,200 per m^3 for any consumption above 10 m^3 and below 20 m^3, and again to VND8,000 per m^3 for any higher consumption.[7] With the linearization, every household in the three villages was charged a uniform VND5,300 per m^3 price and the water utility was reimbursed for the revenue loss relative to the block-price schedule. This linearized price schedule prevailed throughout the experiment, with randomized discounts applied at various points as described below. In June 2017, however, provincial authorities raised the water price to a uniform VND7,950 per m^3, a 50 percent increase, which we passed on to all consumers and which was broadly announced by An Thinh.

Baseline data collection We collected baseline data in October 2015. Out of the universe of 1,526 connected households in that month, including some of those that we had offered free connections to in July, the survey covered 1,488. Reasons given for non-participation were refusal (11 households), family events (10 households), commercial water users mis-registered as residential users (8 cases), physical incapacity (6 households), and the remaining 3 cases were households who no longer used piped water. The survey questionnaire collected data on household demographics, assets, income, and consumption, as well as a module on domestic water use.

Quarterly monitoring In November 2015, we began a series of "monitoring" surveys, implemented at the time of bill payment, in which we collected quarterly information on household demographics (whether households had family or friends returning/staying with them and for

[7]The exchange rate at the time was approximately US$ 1 = VND21,000.

how many days) as well as water sharing practices (whether households were giving/receiving piped water to/from other households). All but 23 households covered by the baseline survey provided quarterly information throughout the experiment.

Experimental sample Our final experimental sample consists of 1,462 water customers.[8] Of these, 114 or 8 percent had been unconnected as of July 2015, but subsequently received a (free) piped water connection through the project in time to be included in the baseline survey. As noted above, prior to our interventions, a larger fraction of commune households (16 percent) did not have connections, presumably because they were not willing to pay the fee. When we turn to optimal pricing in section 4.4, and specifically to the computation of social welfare for the commune, we augment the estimation sample by randomly drawing from the 114 households so that previously unconnected households (inclusive of the original 114) make up precisely 16 percent of the sample.

Price discounts The actual price experiment started in January 2016, 6 months into the price linearization; see Figure 1 for a timeline. All the households surveyed at baseline agreed to participate in the price discount scheme and were randomly divided into three groups. Each group would, according to staggered schedules, receive price discounts of 50 or 75 percent and thus be charged 50 or 25 percent of the July 2015 flat price, respectively (Figure 1).[9] Each price subsidy regime ran for 6 months with the water utility reimbursed for the difference between official and subsidized price. However, households were not informed of the overall duration of the price experiment or of any particular discount more than a month in advance.

The experimental design incorporated a six-month pause in all discounts beginning January 2017, one year into the price experiment. The purpose of the pause was to allow a reduced-form test of demand persistence, exploiting the fact the each of the three groups had been exposed to different price discounts over the previous year yet were currently facing the same price. We discuss this test in section 2.4.

[8] While the baseline survey collected data on 1,488 households, 3 were dropped due to consumption exceeding 50 m^3 for 5 or more months during the experiment; they were most likely operating a business out of their homes. We dropped a further 23 households due to non-participation in the quarterly monitoring surveys (so that no time-varying demographic information could be collected).

[9] As mentioned earlier, a universal 50 percent increase in the base price was implemented in the whole province of Bac Ninh in June 2017. This price increase was also passed on to the experimental sample; total subsidies are reflected accordingly in the lower panel of Figure 1, which indicates that 50 (resp. 75) percent rebates translate into a price equal to 75 (resp. 37.5) percent of the July 2015 price.

Information campaign Households were informed of the discounts the month before they took effect in two ways: (1) through a note left on the "Comments" section of the bill, which was handed to each customer upon their monthly in-person payment at a central location in the village, shortly followed by (2) a letter sent by the water utility to each customer's home.[10] To make the discounts particularly salient just before they were to go into effect, any rebate as a result of a discount was applied retroactively to the prior month's bill and consumers were informed that henceforth this would be the prevailing rate.

Price salience was confirmed by households inasmuch as 97 percent of them responded within rounding error of the actual price of VND5,300 when asked in the quarterly monitoring survey of December 2015 (at the end of the price linearization phase) what price per cubic meter they expected to be charged the following month. Evidently, households were not only well aware of the price of piped water but of the existence of the on-going linearization scheme as well. Broad awareness of the price flattening is suggested by the fact that even among the 16 percent of households who consumed 10 m^3 or more in December 2015, and who would therefore have paid a higher marginal price absent the linearization scheme, 97 percent still accurately reported within rounding error of the exact price. During this initial phase of the experiment, households were also informed of the imminent subsidy program, although its timing and magnitude were not revealed until the last minute.

2.2 Preliminary analysis

Most of the experimental sample households (70 percent) were connected to piped water in the initial roll-out of An Thinh's network from September 2012 to February 2013, with the remainder connecting throughout the rest of 2013 (10 percent), 2014 (9 percent), and 2015 (11 percent). Prior to the arrival of piped water, rainwater stored in concrete tanks was the principal source of domestic water supply; for 83 percent of households, internal piping and plumbing fixtures were already in place for conveying rainwater. Appendix Table B.1 shows household water usage patterns at baseline; all but three households report using piped water. More than half of the sample (806 households) report still using some rainwater for cooking and drinking (rainwater is filtered for household use and typically boiled before drinking), and half also report using either rainwater or groundwater for shower and bathroom needs. In short,

[10]Nonpayment or delinquency on water bills is virtually non-existent in our setting.

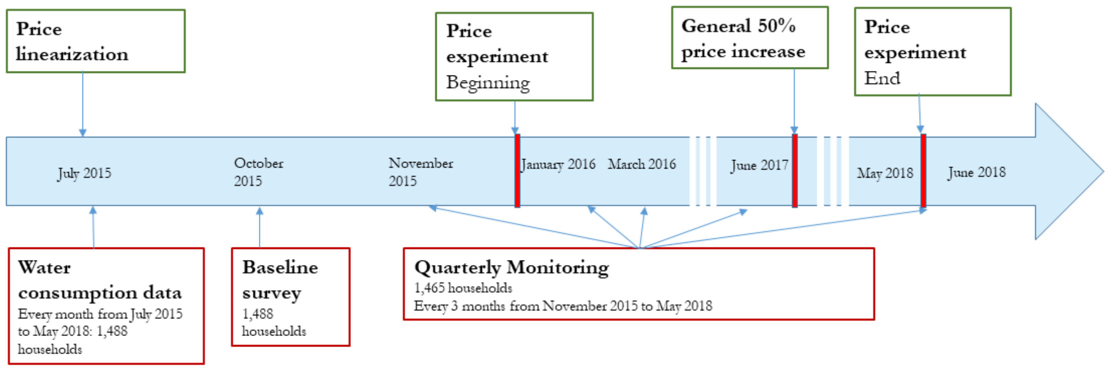

Subsidized price rollout

	Jun-Dec 15	Jan-Jun 16	Jul-Dec 16	Jan-Jun 17	Jul-Dec 17	Jan-May 18
Group 1	100%	50%	25%	100%	150%	150%
Group 2	100%	100%	50%	100%	37.5%	150%
Group 3	100%	100%	100%	100%	75%	37.5%

Figure 1: Experiment timeline and price subsidy rollout

the transition to piped water was far from complete.

Sample balance Appendix Table B.2 presents descriptive statistics for the entire sample and for each experimental group as well as balance tests. In the first row, we see that average monthly water consumption in the pre-subsidy phase (second half of 2015), when all three experimental groups faced the same linearized price, was statistically indistinguishable across groups. We also find balance for quarterly person-days present in the household, constructed from information provided in the quarterly monitoring surveys. Despite being collected during the experiment, it is plausible to assume that this variable, reflecting the number of potential water users, is not influenced by the experiment. Finally, the data show adequate balance on characteristics of the household and the household head, all measured at baseline.

Intra- and inter-household arbitrage By generating price differences between nearby households and over time within households, the experiment risked creating arbitrage opportunities that households could exploit by "trading" water among themselves or storing water. However, there is a strong *prima facie* case against such intra- or inter-household arbitrage in our context. The main benefits of piped water where there is an abundant substitute in the form of rainwater is its higher pressure (for uses such as house cleaning, individual showers and so forth) as well as convenience; i.e., not having to collect, filter, and store rainwater or maintain a water pump, in the case of groundwater. The amenity value of high pressure and convenience would not carry over to water stored or shared between households.[11]

That said, our data allow a test for the potential effects of arbitrage, both inter and intra-household. With respect to the former, data on water sharing obtained in the quarterly monitoring surveys rules out significant inter-household trade. Just under 2 percent of the experimental sample (about 28 households per quarter on average) shared in piped water from their neighbors over the course of the price experiment, whereas less than 2 percent (about 25 households per quarter) shared out their piped water to a neighboring household that itself had piped water. Moreover, we find no evidence that household water consumption or its response to price discounts were affected by water sharing (in or out) during a quarter in which experimental price differentials prevailed (see Appendix C).[12] With respect to intra-household arbitrage, we use information on the volume of the concrete rain water tank, to which households could potentially divert piped water to store for future use. A larger tank implies greater scope for storage; median tank volume is 5 cubic meters (interquartile range 2-9; 4 percent of households have no tank) against median monthly piped water consumption of 6 cubic meters. If such arbitrage were important, we would expect households with larger tanks to have a greater response to discounts. However, as shown in Appendix Table D.1, row (6), we fail to reject equality of price responses across households with varying water tank sizes (p-value=0.920).

2.3 Results of price experiment

Overview The upper panel of Figure 2 summarizes the experimental prices faced by each group of households. Recall that during the six months prior to January 2016 all households regardless of experimental group faced the same linear price, which we normalize to 1 in the figure. As also noted, once the experiment was underway, price discounts were rolled out successively to the different groups and then withdrawn, with each discount prevailing for six months. A break in this pattern, which we shall return to shortly, was a six-month pause in all

[11]While the option value of not running out of rainwater was also cited as a benefit of piped water, and would be amenable to water sharing, this is an unlikely occurrence for most households.

[12]Since we collected identifiers of all households involved in piped water "trade," we also checked, using dyadic regressions, whether the likelihood of a transaction is associated with the water price differential across trading households. In results available upon request, we find no significant relationship.

discounts beginning January 2017. Starting July 2017 (month 25), the water utility increased the price by 50 percent, which was fully passed on to the households enrolled in the experiment.

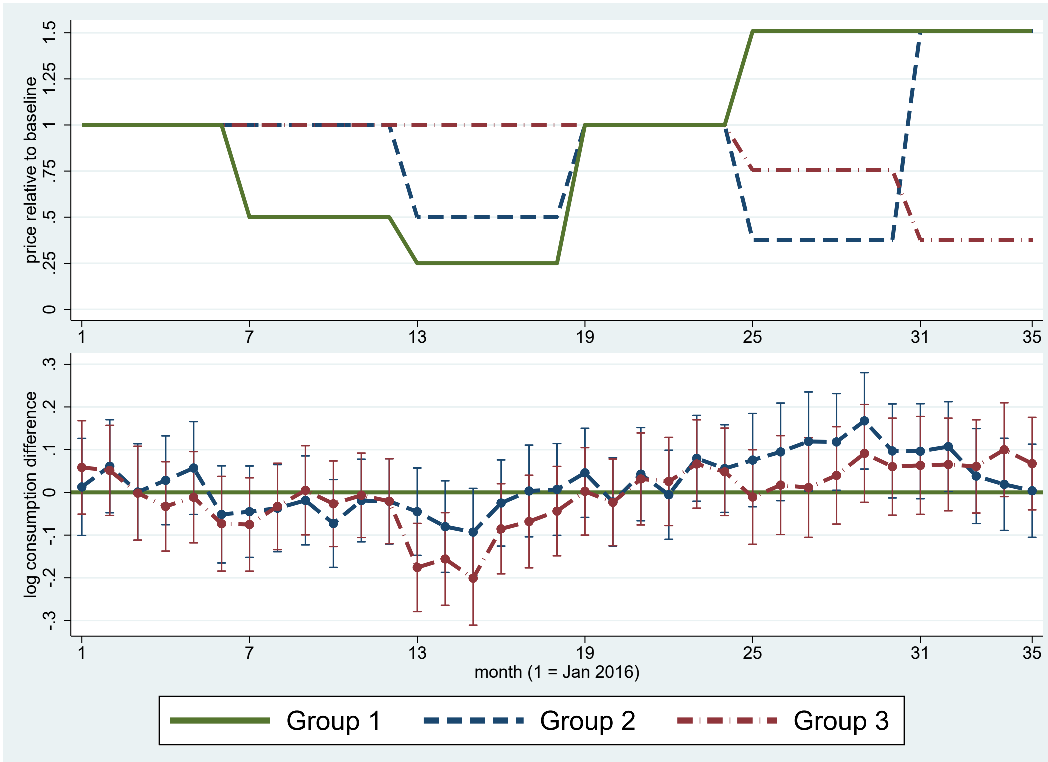

Notes: Top panel: Experimental prices by month and treatment group (base price normalized to 1). Month 1 = July 2015; month 7 = January 2016 (start of experiment); month 19 = January 2017 (start of no discount semester); month 25 = July 2017 (start of general price increase by utility). Bottom panel: Treatment effects for group 2 and 3 (group 1 normalized to zero) by month estimated from a regression of log monthly water consumption on person days, month dummies, and month dummies interacted with treatment group.

Figure 2: Experimental price and group-wise treatment effects by month

The lower panel of Figure 2 summarizes the experimental results using the base sample of 1,462 households, displaying monthly coefficients b_t and associated 95 percent confidence intervals from the regression

$$\log c_{it} = a_0 + a_1 N_{it} + \sum_k b_t \mathbb{1}_{group=k} + d_t + u_{it}, \tag{1}$$

where $\log c_{it}$ is log monthly water consumption and N_{it} is quarterly person-days present in the household. These monthly treatment effects capture log consumption differences between groups 2 or 3 and the base group 1 (normalized to zero), netting out seasonality and common time shocks and/or trends in water demand using month dummies d_t. For most months of the experiment, average log consumption of the three experimental groups are statistically indistinguishable, the exception being months 13-15 (group 1 gets the 75 percent discount and thus has higher consumption than control group 3) and months 28-29 (group 2 gets the 75 percent discount and thus has higher consumption than control group 1)

Contemporaneous response to price discounts Next we investigate water customers' static demand response to price discounts. In Table 1, we report regressions of (log) monthly household water consumption on two dummy variables indicating the discount received (columns 1 and 2), the omitted category being no discount. In columns 3 and 4, we pool the two discount categories into one single "any discount" dummy variable and columns 5 and 6 report regressions of log consumption on price.[13] Each specification controls for person-days and includes dummies for every month of the experiment. In addition to OLS specifications (columns 1, 3, and 5), we also report household fixed effects estimates (columns 2, 4, and 6). There is a small, but precisely estimated, increase in water consumption in response to the discounts, significantly larger with respect to the 75 percent discount than to the 50 percent discount. The coefficient on the price is of the opposite sign; demand for piped water is downward sloping.[14] As we would also expect, quarterly person-days have a statistically strong (and positive) impact on piped water demand. A comparison of the coefficients between columns 1, 3, and 5, with those in columns 2, 4, and 6, indicates that the inclusion of household fixed-effects slightly lowers standard errors without affecting point estimates much. This result confirms that our sample is balanced on unobserved preferences for piped water (as well as on pre-intervention consumption and observed covariates).

[13]In addition to the discounts, price captures the general rate increase of July 2017, the effect of which interacts with experimental group because each group faced a different discount at the time.

[14]We have good reasons to attribute this negative price response to a substitution effect rather than an income effect induced by the price discounts. First, the total cost of piped water accounts for a tiny share (less than 1 percent) of the typical household budget in our sample. Second, in Appendix Table D.1 (rows 2-4), we cannot reject homogeneity of price responses across households with respect to income and wealth measures. Thus, one implication of income effects, that poorer households have higher gross price elasticities than richer households, finds no support in the data.

Table 1: Contemporaneous response to water price discount

	Dependent variable: log monthly piped water consumption					
	(1)	(2)	(3)	(4)	(5)	(6)
50 percent discount	0.0406	0.0336				
	(0.0147)	(0.0146)				
75 percent discount	0.0893	0.0847				
	(0.0155)	(0.0151)				
Any discount			0.0639	0.0580		
			(0.0135)	(0.0133)		
Water price					-0.0801	-0.0760
					(0.0167)	(0.0154)
Person-days/1000	2.945	1.559	2.945	1.561	2.945	1.560
	(0.106)	(0.102)	(0.106)	(0.102)	(0.106)	(0.102)
R^2	0.222	0.659	0.221	0.658	0.221	0.659
Month dummies	YES	YES	YES	YES	YES	YES
HH fixed effects	NO	YES	NO	YES	NO	YES

Notes: Robust standard errors in parentheses clustered at household level. Sample size is $32 \times 1,462 = 46,784$ for all regressions. Estimation is by OLS (cols 1,3, and 5) or household fixed effects (cols 2, 4, and 6). Person-days is measured at the quarterly frequency.

2.4 Persistence in piped water demand

A simple test of demand persistence

Our experimental design provides a simple test of demand persistence that can be read directly from the bottom panel of Figure 2. First, as seen in the top panel, during months 7-18 each of the groups received a different sequence of price discounts to induce different average consumption levels by month 18. If there is persistence, then these exogenous differences in past consumption should, in turn, induce corresponding differences in month 19 consumption, at which point all groups entered the no discount semester and thus faced the same price. Our test was powered given the plausible value of 0.8 for the price elasticity of demand for piped water (see Diakité et al. 2009 and Appendix E). As we will see, however, the price elasticity in our setting turns out to be much lower, only about 0.05, which means that our reduced-form test for persistence is substantially under-powered ex-post.[15] As it happens, the bottom panel of Figure 2 shows no detectable difference in consumption in month 19 between group 1, the most subsidized of the three during months 7-18, and groups 3, the least subsidized. Next, we investigate demand persistence using a more restrictive specification than that of equation (1), one that also exploits

[15]Because of the dissipation of demand persistence over time, power is substantially lower for the same test performed using month 20 consumption (see Appendix E), and even lower using month 21, and so on.

Table 2: Persistence in piped water consumption

	Dependent variable: log monthly piped water consumption								
	previous 2 months			previous 3 months			previous 4 months		
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
Current month variables:									
50 percent discount	0.0239			0.0223			0.0182		
	(0.0186)			(0.0182)			(0.0184)		
75 percent discount	0.0582			0.0601			0.0603		
	(0.0189)			(0.0187)			(0.0190)		
Any discount		0.0298			0.0289			0.0272	
		(0.0171)			(0.0168)			(0.0170)	
Water price			-0.0400			-0.0420			-0.0441
			(0.0190)			(0.0187)			(0.0186)
Previous month variables:									
#months with 50 percent discount	0.00665			0.00585			0.00543		
	(0.00920)			(0.00646)			(0.00527)		
#months with 75 percent discount	0.0177			0.0119			0.00899		
	(0.00931)			(0.00656)			(0.00537)		
#months with any discount		0.0175			0.0131			0.0108	
		(0.00836)			(0.00586)			(0.00478)	
Mean water price			-0.0463			-0.0480			-0.0481
			(0.0189)			(0.0202)			(0.0221)
Significance tests (p-values):									
All prev. month variables	0.137	0.036	0.014	0.186	0.025	0.018	0.246	0.024	0.029
75 percent discount only	0.057			0.071			0.094		
Observations	43,860	43,860	43,860	42,398	42,398	42,398	40,936	40,936	40,936
R-squared	0.663	0.663	0.663	0.666	0.666	0.666	0.668	0.668	0.668
Month dummies	YES	YES	YES	YES	YES	YES	YES	YES	YES
HH fixed effects	YES	YES	YES	YES	YES	YES	YES	YES	YES

Notes: Robust standard errors in parentheses clustered at household level. All regressions include person-days (measured at quarterly frequency). Columns 1, 4, and 7 include dummies for whether household is currently receiving the 50 percent or 75 percent discount as well as number of previous months with 50 percent or 75 percent discount. Columns 2, 5, and 8 include dummies for whether household is currently receiving any discount as well as number of previous months with any discount. Columns 3, 6, and 9 include current water price inclusive of discount as well as mean water price over the previous months.

data from the entire experiment.

Reduced-form evidence on demand persistence

In Table 2, we extend the specifications reported in Table 1 by including information on past discounts. Corresponding to the current discount dummies, we construct variables representing the number of months that the household faced the same type of discount in the previous 2, 3, or 4 months. In columns 1, 4, and 7 of Table 2, for instance, along with dummies for whether the household is currently receiving the 50 or 75 percent discount, we have the number of months out of the last 2, 3, and 4, respectively, that they received either of these discounts. In the water price specifications (columns 3, 6, and 9), we simply include the mean price that the household faced over the previous 2, 3, or 4 months, which incorporates any discounts. Note that, while specifications that include separate dummies for each type of discount are more flexible in terms of functional form, they are also likely to be the least precisely estimated given the rather high monthly serial correlation across discounts induced by the experiment. At any rate, the upshot of Table 2 is that greater discounts (or a lower average price) in the recent past leads to higher consumption in the current month conditional on the current discount. The evidence is strongest in the specifications with the effective water price (p-value=0.014 in column 3) and weakest, as expected, when we separate out the 50 and 75 percent discounts, in which case only the effect of past exposure to the 75 percent discount approaches significance.

Habit formation/learning through experience vs. alternative mechanisms

Both provincial water authorities and the director of An Thinh utility told us they expected households to increase piped water consumption over time. Indeed, in its first 3 years of operation (preceding our experiment), piped water use among An Thinh's customers was growing by an average of 0.9 percent per month.[16] While households had previously collected rainwater that needed to be stored in and pumped from dedicated tanks, piped water is more convenient and provides higher pressure; it also does not need to be filtered and boiled for cooking purposes like rainwater. As these benefits become manifest through continued use, An Thinh's customers were expected to increasingly shift out of rainwater and into piped water, especially

[16] This is a within household estimate and thus accounts accounts for selectivity into An Thinh's customer-base by connection month.

for intensive uses like showers, cleaning, and so forth.[17] Whether we call this phenomenon habit formation or learning through experience, or some combination thereof, it implies that exogenously inducing greater piped water consumption in the past should raise today's consumption conditional on the current price, which is precisely what we find in Table 2.

Alternative explanations for this same evidence include: (i) lagged awareness of price discounts; (ii) loss aversion; and (iii) complementary investment in piped water infrastructure, none of which are likely to be quantitatively important in our setting, as we now discuss. Price awareness or salience was already mentioned in the context of our experimental design. Water customers were informed of the discounts just in advance in two attention-getting ways, on their water invoice handed to them during in-person bill-paying and in an official letter sent to their homes. Through this publicity campaign, we minimized lack of awareness or inattention to the discounts as they came into force.

If loss averse households form a mental account and/or a reference point for piped water expenditures, then the impact of the current price would, in general, depend on the past price; in particular, on whether the former is higher or lower than the latter (see, e.g., Thaler 1985; Kőszegi and Rabin 2006). While such a phenomenon could account for the evidence of demand persistence in Table 2, loss aversion also has the directly testable prediction that demand responses are more elastic to price increases than to price decreases (see, e.g., Ahrens et al. 2017). In Appendix F, however, we present evidence to the contrary, finding that price increases and decreases have symmetric effects on piped water use. Thus, loss aversion is not relevant in our setting and, hence, cannot be an important source of demand persistence. More broadly, unlike the habit formation or experience good mechanism, loss aversion cannot explain rising water consumption in a static price environment, as prevailed prior to our experiment.

Investment in piped water infrastructure is also unlikely to be driving the results in Table 2. First, we find persistence over as little as 2 months, not enough time for households benefiting from a sudden price discount to install plumbing inside their homes or build new showers. Second, as noted previously, most households already had such infrastructure installed prior to their connection to the piped water service. Third, and perhaps most telling, the absolute magnitude of the price discounts and their uncertain duration, while sufficient to marginally incentivize monthly piped water consumption in the treated group as compared to the con-

[17] Monthly water consumption during the hot season from May to September averages about 12 percent higher than during the rest of the year in our sample, largely reflecting the greater demand for showers.

trol, are unlikely to induce differential big-ticket investments or appliance purchases (e.g., a washing machine) across these groups, especially when the average monthly water bill, at the undiscounted price, is less than US$2.

3 Endogenous Preferences for Piped Water

In light of our experimental findings, we next specify a model of consumer demand for piped water that allows for preferences to change with past consumption and yet is empirically tractable as well as convenient for our investigation of optimal pricing.

3.1 Setting the stage

We begin with quasi-linear instantaneous utility with time-varying preference parameter ρ_t

$$u_t(c, z) = (\rho_t c - c \log c) + \alpha z, \tag{2}$$

where c and z are current piped water and numeraire consumption, respectively. The parameter ρ_t has a time-invariant component θ, an exogenous time-varying component ξ_t, and an endogenous time-varying component $\bar{C}_{t-1}$ as follows:

$$\rho_t = 1 + \theta + \beta \bar{C}_{t-1} + \xi_t. \tag{3}$$

For the sake of tractability, we assume that $\bar{C}_t$ follow law of motion

$$\bar{C}_{t-1} = \sum_{k=1}^{\tau} \gamma^{k-1} \log c_{t-k}, \tag{4}$$

where τ is a finite number of periods. In other words, consumption in period t is assumed to directly affect future preferences for τ periods onward. A geometric lag, moreover, implies that the sensitivity of future preferences to past consumption decays at a constant rate over time; when $\tau = 1$, we simply have $\bar{C}_{t-1} = \log c_{t-1}$.

The generic preference specification given by equations (2)-(4) can be rationalized by a model of habit formation (Pollak, 1970; Becker and Murphy, 1988) or of experience goods (Nelson, 1970; Shapiro, 1983). In the first interpretation, $\bar{C}_{t-1}$ is the stock of habits and $\beta > 0$

implies habit formation. In the second interpretation, Shapiro (1983) distinguishes two cases: The optimistic case, in which the consumer anticipates high service quality, but is ultimately disappointed, and the pessimistic case, in which the consumer anticipates low service quality but is pleasantly surprised. In the optimistic case, β, the coefficient on consumption experience $\bar{C}_{t-1}$, is negative, whereas in the pessimistic case it is positive.

3.2 Water consumption choices

A "rational addict" (Becker and Murphy, 1988) or a "strategic experimenter" (Bergemann and Välimäki, 2006) internalizes the effect of current consumption on both current and future preference. Optimal consumption choice, in particular, takes prices $\{p_t\}$ as given and solves dynamic program

$$V_t(\theta, \bar{C}_{t-1}, p_t, y_t|\xi_t) = \max_{\{c_{t+k}, z_{t+k}\}_{k\geq 1}} \mathbb{E}_t \sum_{k\geq 0} \delta^k u(c_{t+k}, z_{t+k}|\theta, \bar{C}_{t-1+k}, \xi_{t+k}) \tag{5}$$

subject to the law of motion given by (4) and the budget constraint, $\forall k \geq 1$,

$$z_{t+k} + p_{t+k} \cdot c_{t+k} = y,$$

where y is consumer income, which we assume to be time invariant. Recognizing, however, that consumers might not be fully sophisticated about how their future preferences are shaped by current choices, we posit instead that consumers are naive or myopic, i.e., unaware of process (4), and thus behave so as to solve the static program

$$V_t(\theta, \bar{C}_{t-1}, p_t, y|\xi_t) = \max_{c,z} u(c, z|\theta, \bar{C}_{t-1}, \xi_t) \tag{6}$$

subject to budget constraint

$$z + p_t \cdot c = y.$$

Our assumption of static optimization is an application of projection bias (Loewenstein et al., 2003) whereby agents perceive that their future preferences will be identical to their current preferences. Whether one interprets our preference structure in terms of habit formation or experience good consumption, projection bias equally predicts that consumers underestimate their future taste for piped water; either they are oblivious to their future habits or they are pessimistic about how much they will value the service after having experienced it. Loewenstein et al. (2003) cite a range of research providing suggestive evidence of projection bias. More recent experimental studies also provide powerful evidence against substantial sophistication (Augenblick and Rabin, 2019; Acland and Levy, 2015).[18] In addition, and more generally, since our bounded-rationality assumption amounts to a functional form restriction on the intertemporal Euler equation, we can provide a formal statistical test of naive/myopic/static versus fully sophisticated/forward-looking/dynamic optimization in our setting. The result of this test, reported in Appendix G, indicates that piped water consumers do not appreciably internalize the future effects of their current consumption, whether due to habit formation or to learning.

With a binding budget constraint, optimal consumption is given by first-order condition:

$$\log c_t(\theta, \bar{C}_{t-1}, p_t|\xi_t) = \theta - \alpha p_t + \beta \bar{C}_{t-1} + \xi_t. \tag{7}$$

Equation (7) defines the short-run demand function as in Pollak (1970). The short-run price elasticity (actually semi-elasticity, or elasticity at $p_t = 1$) is given by α, whereas β captures the extent to which habits formed in the past influence today's consumption.

3.3 Short- vs. long-run demand

While consumer short-run demand is given by equation (7), we define long-run consumption by first adding time $t = 0$, a pre-connection period, and then write

$$\log c_\infty(\theta, p) = \lim_{t\to\infty} \mathbb{E}_0 \log c_t(\theta, \bar{C}_{t-1}, p|\xi_t), \tag{8}$$

i.e. the expectation taken before connection at $t = 0$ of the limit value of (log) consumption, given that price p is time-invariant. Further, assume that $\beta\frac{1-\gamma^T}{1-\gamma} < 1$, or

$$\lambda \equiv \left[1 - \beta\frac{1-\gamma^T}{1-\gamma}\right]^{-1} > 1. \tag{9}$$

[18] In a carefully designed experiment, Hussam et al. (2017) find a degree of rational addiction to handwashing in India. As they note, however, "our design sets up the optimal scenario to facilitate rational habit formation: households are fully aware that we want to help them develop a habit of handwashing, and we reiterate the future dates at which the value of the behavior will change." It is unclear, therefore, how this sophistication result generalizes to a setting like ours in which households are not primed for habit change.

We can state our first intermediate result (proofs in the Appendix):

Lemma 1 If inequality (9) holds, then long-run demand given by (8) is well-defined and

$$\log c_{\infty}(\theta, p) = \lambda(\theta - \alpha p). \blacksquare \tag{10}$$

As the stock of consumption builds up over the long-run, demand for piped water increases as does its price elasticity. Intuitively, an increase in price not only decreases contemporaneous consumption but also reduces past consumption, thus amplifying the static price response.

4 Structural estimation

In this section, we take up the estimation of λ and the other preference parameters using the experimental data, and then validate the results using out-of-sample data.

4.1 Preference Estimation

We first detail the steps in our procedure to estimate consumer preferences.

Empirical specification

Equation (7) describes consumer behavior; its empirical counterpart is

$$\log c_{it} = -\alpha p_{it} + \beta \sum_{k=1}^{\tau} \gamma^{k-1} \log c_{it-k} + \omega N_{it} + \theta_i + \zeta_t + \varepsilon_{it}, \tag{11}$$

where i indexes consumers. Here, we also decompose the exogenous preference shock ξ_{it} into an observed component N_{it}, reflecting time-varying household demographics, and an unobserved component ε_{it} including aggregate shock ζ_t, i.e. $\xi_{it} = \omega N_{it} + \zeta_t + \varepsilon_{it}$. Note that price p_{it} varies over time and (randomly) across consumers by experimental group.

To remove the consumer fixed effect θ_i, we difference (11) between period t and $t - L$, for some to-be-determined time lag L, such that

$$\Delta^L \log c_{it} = -\alpha \Delta^L p_{it} + \beta \sum_{k=1}^{\tau} \gamma^{k-1} \Delta^L \log c_{it-k} + \omega \Delta^L N_{it} + \Delta^L \zeta_t + \Delta^L \varepsilon_{it}, \tag{12}$$

where, e.g., $\Delta^L p_{it} \equiv p_{it} - p_{it-L}$ and $\Delta^L \zeta_t$ is a month fixed effect (as included in the reduced form models above). Any choice of lag used for differencing is equally valid econometrically. Below, we use a data-driven approach to select L.

Identification

Experimentally induced variation in the price of piped water can be used to identify, not only the short-run demand curve, but also the causal impact of past water consumption on current demand. Consider identification of the key preference parameters α and β in the case $\tau = 1$. The identification argument is the same for $\tau > 1$, with the additional complication of estimating γ, the discussion of which we leave until later.

With $\tau = 1$, equation (12) becomes

$$\Delta^L \log c_{it} = -\alpha \Delta^L p_{it} + \beta \Delta^L \log c_{it-1} + \omega \Delta^L N_{it} + \Delta^L \zeta_t + \Delta^L \varepsilon_{it}, \tag{13}$$

from which we see that, *conditional* on the first lagged consumption change, the change in current consumption depends only on the current price change $\Delta^L p_{it}$ (which identifies α) but not on the change in lagged prices $\Delta^L p_{it-k}$ or any function of lagged prices for $k \geq 1$. This is the potential set of theoretically valid exclusion restrictions. Theory, moreover, implies that the lagged consumption change $\Delta^L \log c_{it-1}$ (unconditionally on $\Delta^L \log c_{it-2}$) *does* depend on the change in lagged prices or functions thereof. Hence, the intertemporal dependence parameter β is identified in principle.

Our identification argument relies exclusively on experimental variation in price induced by the sequence of randomized discounts. Typically, however, estimation of dynamic error components models in differenced form, such as equation (13), involves using various lags of the dependent variable as instruments.[19] This approach, originally proposed by Anderson and Hsiao (1981) and refined by Arellano and Bond (1991), *requires* that the ε_{it} be serially uncorrelated or, at minimum, follow a restricted moving-average process. By contrast, we do not restrict the time-series properties of the preference shock and, in particular, we do not assume that ε_{it} is serially uncorrelated; indeed, we find evidence to the contrary below.

[19] Along these lines, see, e.g., Meghir and Weber (1996); Dynan (2000); Carrasco et al. (2004) for estimation of Euler equations allowing for intertemporal nonseparability.

Estimation in the general case $\tau \geq 1$

In the general case, we want to recover both γ and τ, along with the other parameters of equation (12). Conditional on values of γ and τ, the parameters α, β, and ω are estimable as in the previous subsection. Given such estimates, we may compute residuals $\hat{e}_{it} = \Delta^L \hat{\varepsilon}_{it}$ and form the concentrated sum of squared errors

$$S(\gamma, \tau) = \sum_{i=1}^{N} \sum_{t=1}^{T} \hat{e}_{it}^2(\gamma, \tau). \tag{14}$$

Following Hansen (1999) for a similar type of problem, the least squares estimate of (γ, τ) minimizes $S(\gamma, \tau)$ and can be found by grid search. Once we have $(\hat{\gamma}, \hat{\tau})$, the other coefficients can be obtained as $\alpha = \alpha(\hat{\gamma}, \hat{\tau}), \beta = \beta(\hat{\gamma}, \hat{\tau})$, and so on. In practice however, we restrict $\tau < 5$ for the purpose of our estimation. While Hansen (1999) notes that the dependence of (α, β, ω) on $(\hat{\gamma}, \hat{\tau})$ is of second-order asymptotic importance, which means that "naive" standard errors of these parameter estimates are asymptotically correct, we do not rely on this argument. Rather, we cluster-bootstrap the entire procedure to get our standard errors.

We present our results as the following sequence of estimation steps, beginning with data-driven algorithms for choosing both lag-length L for differencing and the lagged price instruments:

1. With $\tau = 1$ and using a 'basic' instrument set $\left\{\Delta^L p_{it-1}, \Delta^L \log p_{it-1}\right\}$ for $\Delta^L \log c_{it-1}$, select the L that yields the best-fitting first-stage for equation (12). Call this L^*.

2. Given L^*, use machine-learning (post-double selection methodology of Belloni et al., 2012) to select a sparse instrument set from a successively expanded basic instrument set, i.e., $\left\{\Delta^{L^*} p_{it-1}, ..., \Delta^{L^*} p_{it-k}, \Delta^{L^*} \log p_{it-1}, ..., \Delta^{L^*} \log p_{it-k}\right\}$.

3. Repeat step 1 with sparse instrument set until convergence.

4. Given L^* and the sparse instrument set, estimate γ and τ.

We now turn to the implementation of the estimation procedure.

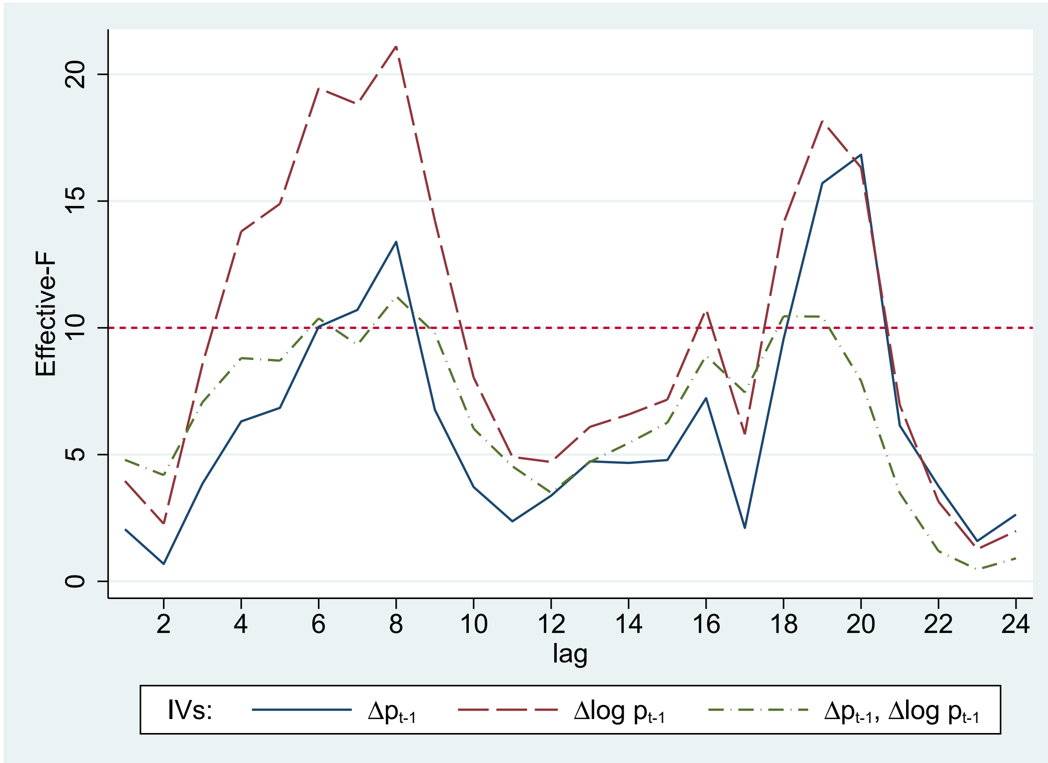

Notes: Effective F-statistic Olea and Pflueger (2013) at different values of differencing lag-length L under alternative instrument sets. Horizontal line indicates rule-of-thumb threshold value for weak instrument inference suggested by Andrews et al. (2019).

Figure 3: Effective-F statistic by lag length

4.2 Lag-length and instrumental variable selection (steps 1 to 3)

The choice of the lag length in Step 1 is guided by the maximization of the relevant price variation. With experimental price regimes of 6-month duration, using an L much below 6 would not fully exploit transitions across regimes. Beyond that, however, there is little we can say a priori. As our first-stage fit criterion, we use the effective-F statistic of Olea and Pflueger (2013), which is preferred for weak instruments diagnostics (Andrews et al., 2019). Note, however, that we do *not* use the effective-F as a criterion for selecting instruments, as this could lead to biased inference. Rather, for a *fixed* instrument set, Step 1 finds the L that maximizes effective-F for the first-stage corresponding to equation (13).

Aside from the basic instrument set consisting of $\Delta^L p_{it-1}$ and $\Delta^L \log p_{it-1}$, Figure 3 also shows effective-F statistics using these two instruments individually and, for reference, the rule-of-thumb threshold value of 10; below this threshold, Andrews et al. (2019) suggest deploying weak-instruments-robust inference. As expected, first-stage fit is relatively poor for $L \ll 6$, but also for $12 \leq L < 18$, in the latter case because we are, among other things, differencing across the no discount semester and the first (pre-discount) semester between which prices did not change for any of the three groups (see Figure 2). First-stage fit improves again at $L = 18$ months and beyond, although estimation sample sizes drop steadily. Based on the effective-F statistic for the basic instrument set, we choose $L^* = 8$. We get an identical result if we instead base lag-length selection on the average effective-F across all three instrument sets in Figure 3.

Table 3: Alternative specifications ($\tau = 1$)

	$L = 8$		$L = 7$	$L = 19$
	(1)	(2)	(3)	(4)
β	0.616	0.654	0.747	0.688
	(0.130)	(0.116)	(0.133)	(0.122)
α	-0.0420	-0.0400	-0.0309	-0.0338
	(0.0119)	(0.0113)	(0.0125)	(0.0116)
ω	0.601	0.546	0.414	0.570
	(0.193)	(0.172)	(0.189)	(0.210)
Effective-F	11.25	9.170	9.330	10.44
Over-id	[0.563]	[0.679]	[0.682]	[0.286]
AR(1)	[0.031]	[0.006]	[0.002]	[0.015]
Month dummies	Yes	Yes	Yes	Yes
Observations	35,087	35,087	36,549	19,005

Notes: Robust standard errors in parentheses clustered at household level (1,462 households); p-values for tests of overidentifying restrictions and of zero first-order serial correlation in the residuals (AR(1)) in square brackets. Dependent variable in all regressions is the change in log monthly consumption from period t to $t - L$. Estimation is by two-step GMM. Instruments selected by IV-LASSO (machine-learning) starting from $\{\Delta^8 p_{t-1}, \Delta^8 \log p_{t-1}\}$ in cols 1 (selected $\Delta^8 p_{t-1}$ and $\Delta^8 \log p_{t-1}$) and from $\{\Delta^8 p_{t-1}, \Delta^8 p_{t-2}, \Delta^8 \log p_{t-1}, \Delta^8 \log p_{t-2}\}$ in cols 2 (selected $\Delta^8 p_{t-1}, \Delta^8 \log p_{t-1}$, and $\Delta^8 p_{t-2}$). Instruments in columns 3 and 4 are $\{\Delta^L p_{t-1}, \Delta^L \log p_{t-1}\}$.

We report the results of Step 2 in Table 3, beginning in column 1 with LASSO-based instrument selection (Ahrens et al., 2019). Both instruments, $\Delta^8 p_{it-1}$ and $\Delta^8 \log p_{it-1}$, are selected, so our basic instrument set coincides with the sparse set. Next, starting from $\{\Delta^8 p_{it-1}, \Delta^8 p_{it-2}, \Delta^8 \log p_{it-1}, \Delta^8 \log p_{it-2}\}$, the basic instrument set is again selected, albeit augmented by $\Delta^8 p_{it-2}$ (column 2). Nevertheless, we cannot reject the hypothesis that $\Delta^8 p_{it-2}$ is a redundant instrument using an LM test (p-value$= 0.13$).[20] Moreover, adding a third lag to the starting instrument set leads to the same selected instruments as the two-lag case. In sum, expanding the

[20]By contrast, starting from the basic instrument set, we *can* reject the hypothesis that either $\Delta^8 p_{it-1}$ or $\Delta^8 \log p_{it-1}$ is redundant with p-values of, respectively, 0.043 and 0.002.

Table 4: Estimation of general preferences

	$\tau = 1$	$\tau = 2$	$\hat{\tau} = 3$	$\tau = 4$
γ	—	0.47	0.41	0.38
			[0.080]	
β	0.624	0.432	0.402	0.403
	(0.131)	(0.0966)	(0.0939)	(0.0949)
			[0.0789]	
α	-0.0416	-0.0442	-0.0458	-0.0463
	(0.0120)	(0.0124)	(0.0126)	(0.0127)
			[0.0148]	
ω	0.584	0.612	0.633	0.640
	(0.193)	(0.199)	(0.203)	(0.203)
			[0.216]	
Effective-F	11.1	11.3	11.6	11.8
Over-id p-value	0.551	0.591	0.532	0.522
Month dummies	YES	YES	YES	YES
Observations	33,522	33,522	33,522	33,522
$S(\gamma, \tau)$	20000	18929	18813	18823

Notes: Naive standard errors in parentheses clustered at household level (1,462 households); cluster bootstrapped standard errors in square brackets (for best fitting model). Dependent variable in all regressions is log monthly consumption. Estimation conditional on (γ, τ) is by two-step GMM with additional instruments Δp_{t-1} and $\Delta \log p_{t-1}$.

number of price lags beyond one does not appreciably add to the power of our first-stage.

While our procedure thus converges on $L^* = 8$, for robustness, we also report alternative differencing in columns 3 and 4 of Table 3, with little change in the results. Finally, to validate the importance of having experimental identification in this setting, we also report p-values from an Arellano and Bond (1991) type AR(1) test in Table 3. Our rejection of zero first-order serial correlation in the residuals suggests that dynamic panel methods, which use lags of consumption as instruments, may be invalid.[21]

4.3 Estimating γ and τ (step 4)

Table 4 presents the results of our grid search over values of $\gamma \in (0, 1)$ and $\tau < 5$ following the procedure outlined in section 4.1. Column 1 in Table 4 replicates the result in column 1 of Table

[21]With $L = 1$, only the tests for 2nd and higher-order autocorrelation would be relevant, since first-order autocorrelation results mechanically by first-differencing the data. Given, however, that $L > 1$ here, autocorrelation of first-order and higher is relevant to the validity of dynamic panel data methods.

3.[22] The remaining columns of Table 4 correspond to the successive values of τ and, in each case, we report the best fitting γ (based on a grid search in increments of 0.01). Comparing the value of $S(\gamma, \tau)$ across these models leads us to select $\hat{\tau} = 3$ with $\hat{\gamma} = 0.41$, the latter estimate being quite large relative to its bootstrapped standard error of 0.08.

Thus, we find that greater piped water consumption over the last three months encourages current consumption, although the influence of consumption three months ago is only about ($0.41^2 =$) one-sixth of that undertaken in the previous month. At our preferred parameter values, the short-run price elasticity of piped water demand is 0.046 with a standard error of 0.015, whereas the implied long-run elasticity is 0.124 with a standard error of 0.061. For comparison, if we estimate a "reduced-form" price response by simply regressing $\Delta^8 \log c_{it}$ on $\Delta^8 p_{it}$, $\Delta^8 N_{it}$, and month dummies, we obtain an elasticity of 0.085 (0.016), which is somewhere in between our estimated short and long-run price elasticities. By conflating the contemporaneous price response with the past price responses via lagged consumption (recall that prices are highly correlated across months), the reduced form elasticity overestimates the short-run elasticity. The reduced form elasticity also underestimates the long-run elasticity, as only the latter takes into account the steady-state implications of the first-order condition.

To sum up, we find evidence of intertemporal dependence in piped water consumption. Households that were experimentally induced to consume more water in a given month through price discounts continued to consume more water in subsequent months conditional on the prevailing price. This behavior implies a long-run price elasticity that is nearly three times higher than the short-run price elasticity or, equivalently, that $\hat{\lambda} = \left[1 - \hat{\beta}\frac{1-\hat{\gamma}^{\hat{\tau}}}{1-\hat{\gamma}}\right]^{-1} = 2.7$.

4.4 Model validation: out-of-sample consumption dynamics

Before turning to policy implications, it is worth asking how well our dynamic structural model captures piped water consumption patterns over time. Connections to the new water utility in My Huong commune began in September 2012 and continued up until our pricing experiment started in January 2016 (see Appendix Figure B.1). We use our estimates $\hat{\alpha}, \hat{\beta}, \hat{\gamma}$, and $\hat{\theta}_i$ to predict water consumption for each household from the first month of their connection ($t = 1$) up until December 2015, assuming that $\bar{C}_0 = 0$ and computing piped water demand under the

[22] The small discrepancy between column 1 in Table 4 and column 1 in Table 3 is due to the requirement that, for the estimation of the coefficients in Table 3, $\Delta \log c_{t-4}$ needs to be non-missing, which leads to a slightly smaller sample.

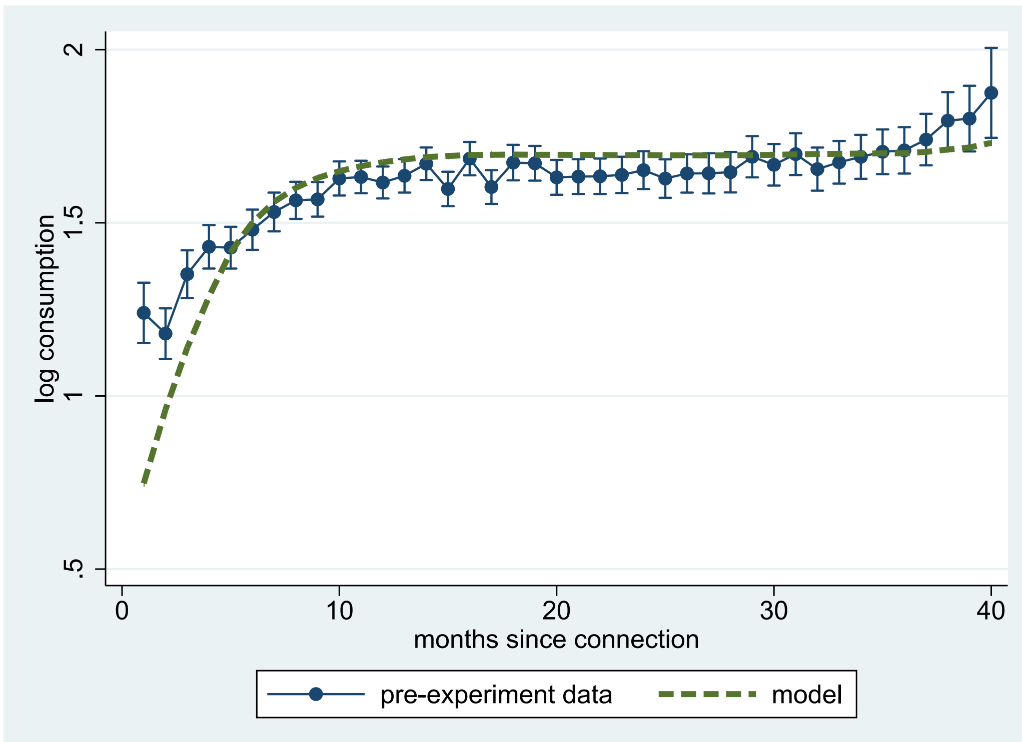

Notes: The solid line plots average log piped water consumption for each month since connection using data from September 2012 to December 2015. The dashed line does the same for model-generated consumption. In both cases, we display predicted margins from a regression of log consumption on month of connection dummies, estimated preferences, and month $\times$ year fixed effects. Bars denote 95% confidence intervals for the predicted margins using the actual data.

Figure 4: Actual vs. model-generated pre-experiment log consumption

prevailing official block tariff schedule.[23]

The solid line in Figure 4 shows monthly predicted margins (and confidence intervals) from a regression of *actual* log monthly consumption on month since connection dummies, which also controls for preferences $\hat{\theta}_i$ and for calendar month $\times$ year fixed effects to absorb aggregate shocks.[24] Water usage data cover the first 40 months of the utility's operation in the commune prior to our experiment and includes all (and only) the 1462 households in the experimental sample. The dashed curve in Figure 4 is derived from exactly the same regression procedure but with the dependent variable now being model-generated log monthly consumption; i.e., an unbalanced panel of household specific time-series beginning at initial connection. With the exception of the first and last few months of connection, our model's out-of-sample fit is quite good. Most encouraging is that model and data seem to agree on how long it takes, on average, for consumption to reach its steady state level starting from the initial month of connection; a static model of demand, by contrast, would trivially, but counter-factually, predict that steady state consumption is achieved at $t = 1$.

5 Optimal pricing of a new utility service

We are now ready to investigate the optimal pricing of a new utility service. Returning to the model of Section 3, the *ex-ante* expected indirect utility, i.e., prior to consumption of the service, may be defined as

$$\mathbb{V}_0(\theta, p, y) = \frac{1-\delta}{\delta}\left[\alpha y + \mathbb{E}_0 \sum_{t\geq 1} \delta^t V_t(\theta, \bar{C}_0, p, y|\xi_t)\right]. \tag{15}$$

Ex-post utility, i.e., once consumption has reached its long-run steady state level, is similarly given by

$$\mathbb{V}_\infty(\theta, p, y) = \frac{1-\delta}{\delta}\left[\alpha y + \mathbb{E}_0 \sum_{t\geq 1} \delta^t V_t(\theta, \bar{C}_\infty, p, y|\xi_t)\right]. \tag{16}$$

In defining these utilities, we abstract from transition dynamics. Thus, equation (15) assumes that $\bar{C}_0 = 0$ and remains so over time, whereas equation (16) assumes that $\bar{C}_{t-1}$ reaches $\bar{C}_\infty$ in

[23]While time-varying demographics, N_{it}, were not collected prior to the experiment, the estimated household fixed effect $\hat{\theta}_i$ incorporates the household average demographics $\bar{N}_i$ over the course of the experiment.

[24]Note that, in an alternative household fixed specification, we could not distinguish cohort (month of connection) and time (aggregate shock) effects.

period $t = 1$. Assuming further that the preference shock ξ_t is "small", we obtain:[25]

$$\mathbb{V}_0(\theta, p, y) \approx e^{\theta - \alpha p} + \frac{1}{\delta}\alpha y \tag{17}$$

and

$$\mathbb{V}_\infty(\theta, p, y) \approx e^{\lambda(\theta - \alpha p)} + \frac{1}{\delta}\alpha y. \tag{18}$$

Before turning to optimal pricing decisions, we specify the utility's cost structure. Let K denote per customer cost, which includes plant and equipment, operating costs and home connection costs. To avoid dependence on discount rate δ, we interpret K as a perpetual flow starting once the utility begins operating in period 1. Finally, consistent with our setting, we assume that the marginal cost of supplying piped water is zero so that price and markup are the same. The optimization program for the public utility is thus to maximize aggregate welfare under the constraint that fees and water sales revenues per customer cover average cost K.

5.1 Optimal pricing: homogeneous preferences

To isolate the implications of demand persistence on optimal pricing, we first consider the homogeneous preferences case, in which the entire population shares the same value of θ. Prior to setting up the water utility to serve this population, we suppose that the social planner undertakes a preference elicitation, as in Lee et al. (2020), wherein consumers truthfully reveal their willingness-to-pay for piped water. While thus knowing θ, this unsophisticated planner does not take into account endogenous preferences. As is well-known, the utility's mark-up in this scenario will be zero – hence, $p = 0$ given zero marginal cost – and it will cover its costs, insofar as it can, entirely through a fixed usage or connection fee. Ex-ante willingness-to-pay of consumers is, therefore, given by $\mathbb{V}_0(\theta, 0, 0) = e^\theta$, whereas the ex-post willingness-to-pay is given by $\mathbb{V}_\infty(\theta, 0, 0) = e^{\lambda\theta}$. It follows that a decision to undertake the project based on ex-ante willingness-to-pay leads to under-investment when costs are in the range $e^\theta < \alpha K \leq e^{\lambda\theta}$. We illustrate this range of potentially inefficient investment by the unshaded region in Figure 5.

A novel rationale for two-part tariffs A more sophisticated social planner can alleviate the inefficiency by offering a two-part tariff consisting of an upfront fee combined with a markup

[25] Note that for every $t \geq 1$, $V_t(\theta, 0, p, y|\xi_t) = e^{\theta - \alpha p + \xi_t} \approx (1 + \xi_t)e^{\theta - \alpha p}$. Thus, $\mathbb{E}_0 V_t(\theta, 0, p, y|\xi_t) \approx e^{\theta - \alpha p}$, which is time-invariant and hence yields expression (17). A similar argument holds for (18).

p.[26] To make our analysis of optimal pricing independent of discount rate δ and in a way consistent with our definition of average cost K, we express this fee as a flow of payments F starting in period 1, which is the same as a flow of payments δF starting in period 0.[27] In equating the upfront payment of the fee to perpetual installment payments of δF on a loan of size F taken at (monthly) interest rate $\frac{1}{\delta}$, we abstract from credit constraints. While credit constraints may limit what utilities can charge upfront for a connection (Devoto et al., 2012a; Lee et al., 2020; Berkouwer and Dean, 2021), the relevant policy instrument to deal with this issue is not the utility's price structure but rather the provision of credit to households.

Now consider the welfare improvement that could be achieved by a social planner who correctly anticipates future preferences. Compared to our earlier planner, who only uses knowledge of ex-ante willingness-to-pay, this sophisticated planner's program is

$$\max_{\{F,p\}} \mathbb{V}_{\infty}(\theta, p, -\delta F), \tag{19}$$

subject to the same participation constraint (since consumers are unaware of future preferences),

$$\mathbb{V}_0(\theta, p, -\delta F) \geq 0 \tag{20}$$

and budget constraint

$$pe^{\lambda(\theta-\alpha p)} + F \geq K. \tag{21}$$

Revenues, on the left-hand side of equation (21), are based on long-run consumption, which the social planner correctly anticipates. In sum, both the maximand (19) and budget constraint differ from those of the unsophisticated planner.

Since a higher price p distorts consumption and reduces welfare, budget balance implies that the optimum will be reached at the lowest possible price and highest corresponding connection fee. If $e^{\theta} \geq \alpha K$, then the optimal pricing scheme has $F = K$ and $p = 0$; at the other extreme, $e^{\lambda\theta} < \alpha K$, full cost-recovery is infeasible and piped water is thus not socially desirable

[26] An optimal tariff schedule does not necessarily have only two parts. However, if the demand elasticity is constant across consumers, which is the case here as we will establish empirically below, then the optimal nonlinear price is a two-part tariff (provided that the social planner does not have a redistribution motive).

[27] To see why, note that a one-time connection fee F assessed at time $t = 0$ is equivalent to a flow of payments equal to $(1 - \delta)F$ paid in every period or to a flow of payments equal to $\frac{1-\delta}{\delta}F$ paid from $t = 1$ onward. By the same token, a flow of payments equal to F paid in every period starting at $t = 1$ is equivalent to payments equal to δF made from $t = 0$ onward.

(rightmost shaded region in Figure 5). In the nontrivial case of $e^{\theta} \leq \alpha K \leq e^{\lambda\theta}$, so that fixed costs are low enough to be fully recoverable but not so low as to be recoverable through the connection fee alone, the social planner simply chooses the lowest price that satisfies the budget constraint.[28] Summarizing, we have

Proposition 1 *When average costs are in the range of potentially inefficient investment, the optimal two-part tariff consists of a positive markup and a connection fee below average cost.* ■

Continuing with the graphical illustration in Figure 5, for αK below e^{θ}, in the leftmost shaded region, a connection fee alone achieves full cost-recovery. For higher K, however, this fee is not sufficient and a positive markup is required to allow both participation and budget balance (segment AC). Such a markup entails a deadweight loss that, at sufficiently high K, eventually leads consumers to no longer connect (DB) despite long-run surplus still exceeding average cost. A two-part tariff allows the water utility to be viable for average costs along the segment AC, where the unsophisticated social planner would have abstained from undertaking the project in the first place. The welfare gain attributable to sophisticated pricing thus corresponds to the area between the orange dashed line and the 45-degree line. This gain is achieved by, in effect, taxing future selves through a markup in order to subsidize current selves in the form of a connection fee reduction. In the theory of utility pricing, this argument constitutes a novel rationale for a two-part tariff with price in excess of marginal cost. Next, we argue that welfare gains can be even larger by expanding the set of pricing contracts.

A subscription plan Instead of a connection fee F, which is paid by the present consumer, i.e., before endogenous preference formation, the social planner can use a monthly subscription fee f, which is paid *only* by the future consumer. Note that the recurrent subscription fee is conceptually distinct even from our amortized representation F of the upfront connection fee

[28] That is, we may rewrite the social planner's program as

$$\max_{p}(1 + \alpha p)e^{\lambda(\theta-\alpha p)} - \alpha K$$

subject to modified budget constraint (incorporating the binding participation constraint)

$$\alpha p e^{\lambda(\theta-\alpha p)} + e^{\theta-\alpha p} = \alpha K.$$

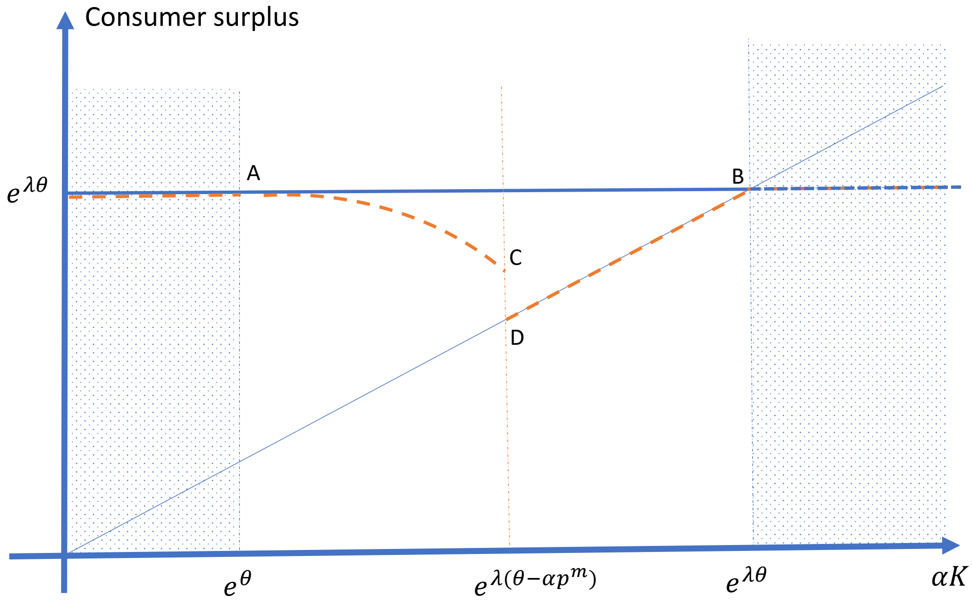

Notes: The horizontal axis represents average costs for the utility including capital costs; the vertical axis is consumer surplus. Projects on the 45-degree line generate zero net social surplus, while projects above (resp. below) generate positive (resp. negative) surplus. The dashed orange line plots consumer surplus under a two-part tariff with a connection fee; the plain blue line plots consumer surplus under a two-part tariff with a deferred subscription fee. The vertical lines at e^{θ} and $e^{\lambda\theta}$ are the ex-ante and ex-post utilities derived from piped water, respectively. The shaded area on the left of e^{θ} represents cases where the project is both ex-ante and ex-post desirable: connection or subscription two-part tariffs with zero markup are equivalent. A project in the shaded area on the right of $e^{\lambda\theta}$ is both ex-ante and ex-post undesirable and no investment takes place. The area in-between is where ex-ante and ex-post assessments differ; p^m is the optimal markup in this region.

Figure 5: Two-part tariffs: connection fee vs. subscription plan

(which is also a payment flow starting at $t = 1$) inasmuch as the latter is an *unconditional* payment, whereas f is paid only if consumption takes place. A consumer who declines to pay amortized connection fee K would also not be willing to pay the same amount as a monthly subscription if required to commit to it at time 0. However, the social planner, aware of the time inconsistency, knows that, despite the time 0 consumer's unwillingness, the time 1 consumer will be willing to pay a flow of K beginning at time 1. Thus, since connection fee and monthly subscription are equivalent for the utility's budget, it always (weakly) prefers to charge zero connection fee and recover its costs through the monthly subscription.[29] So, we have

Proposition 2 *The optimal subscription plan consists of zero markup and a subscription fee equal to average cost; it leads to efficient investment decisions.* ■

Deferred subscription pricing, by shifting lump-sum payment to the future (long-run) consumer without a distortionary markup, fully restores efficiency (line AB); it dominates the two-part tariff of Proposition 1 (line AC) and allows projects to be financed when the latter pricing strategy fails to balance the budget (line DB). The welfare gain in moving from a connection fee to a subscription fee model is thus represented by the distance between AB and the dashed line $ACDB$.

5.2 Optimal pricing: heterogeneous preferences

The case of homogeneous preferences highlights how a two-part tariff effectuates *within* consumer (future to present) cross-subsidization. With a non-degenerate distribution $G(.)$ over $[0, +\infty[$ of preference parameters θ, two-part tariffs have the additional feature of allowing *between* consumer cross-subsidization as established by Ng and Weisser (1974) and Auerbach and Pellechio (1978), the seminal neoclassical treatments with preference heterogeneity.

In particular, the sophisticated planner's program becomes

$$\max_{F,p} \mathbb{V}^*(\bar{\theta}, p) = \int_{\bar{\theta}}^{\infty} e^{\lambda(\theta - \alpha p)} dG(\theta) - \alpha F[1 - G(\bar{\theta})], \tag{22}$$

[29]Implementation issues with the subscription plan could arise along the transition path to the long-run steady state, which we have abstracted from in our discussion. For example, in practice, there may have to be an "introductory" period of zero or low subscription fees, which could induce households to connect that would not want to stay connected at the higher long-run subscription fee (i.e., if preferences are heterogeneous). Nevertheless, over a long enough horizon, the welfare costs of such potential "leakages" would be small.

subject to a participation constraint that defines marginal consumer $\bar{\theta}$

$$e^{\bar{\theta}-\alpha p} = \alpha F, \tag{23}$$

and budget constraint

$$p \int_{\bar{\theta}}^{\infty} e^{\lambda(\theta-\alpha p)} dG(\theta) + F[1 - G(\bar{\theta})] \geq K. \tag{24}$$

When $\lambda = 1$, this problem reduces to Auerbach and Pellechio (1978), in which case a markup $p > 0$ is optimal for K sufficiently high. This markup acts as a tax, the greater burden of which is borne by high-θ consumers, to finance a subsidy on the connection fee, the main beneficiaries of which are low-θ consumers (who would not have otherwise connected). For $\lambda > 1$, we also have the wedge between ex-ante and ex-post willingness-to-pay discussed above. Cutting against what would seem like this second incentive for a markup, however, is the fact that, when $\lambda > 1$, long-run demand is also high, thus accentuating the deadweight loss from the first (Auerbach and Pellechio, 1978) markup. So, unlike the case of homogeneous preferences, here there is no clear-cut range of average costs over which a markup is called for (as in Proposition 1), nor is it clear when a deferred subscription fee plan entails zero markup (as in Proposition 2). We thus turn to numerical simulations to gain further insight.

Using estimates of $G(\theta)$ and α derived from our structural estimation (see Section 4 and next subsection), Figure 6 shows optimal markup regimes for alternative combinations of average costs K and intertemporal dependence parameter λ. When $\lambda = 1$, the optimal markup under the connection fee contract (F^*, p^*) and the subscription fee contract (f^{**}, p^{**}) are identical; deferring payment is useless when demand is time invariant. As noted, the markup in this case serves only as a way for high-θ consumers to cross-subsidize those with low θ.

While a markup may continue to be warranted to redistribute between households (green domain) for $\lambda > 1$, higher values of λ also imply greater deadweight loss from this taxation. The optimal markup falls to zero faster for the subscription contract as the fee in this case is incurred only once consumption has reached its long-run level and the markup does not serve as a within-household intertemporal transfer (orange domain).[30] Reading Figure 6 along the vertical dimension, for a given λ, lumpsum transfers dominate distortionary markups (gold domain) for small K. As K increases, cost recovery requires between-consumer cross-subsidization under

[30]For high enough λ, the optimal markup becomes zero even in the connection fee setting (gold domain); e.g., at $K = 30$, p^* is negligible for $\lambda > 6$.

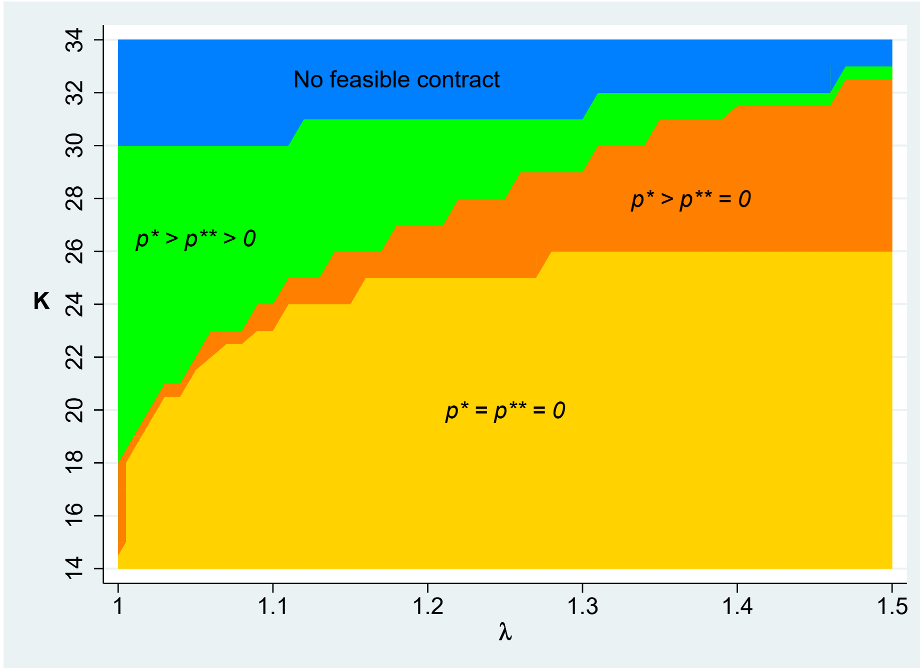

Notes: The graph plots solutions to the optimal connection fee with markup p^* and subscription fee with markup p^{**} at different average costs (K) and degree of intertemporal dependence (λ). In the blue domain, no contract achieves cost-recovery; in the green domain, both contracts have a markup; in the orange domain, only the connection fee contract has a markup; in the gold domain, neither contract has a markup.

Figure 6: Markup regimes for two-part tariffs with preference heterogeneity

a connection fee (orange domain) and, eventually, under both a connection and subscription fee (green domain). For high enough K, the utility, facing downward-sloping demand, cannot recover average costs (blue domain).

5.3 Utility pricing scenarios

To simulate the welfare implications of alternative water price structures in our setting, we first extract the distribution of preferences for piped water in My Huong commune. Using (11), we recover $\hat{\theta}_i$ for our estimation sample. Since, as noted above, previously unconnected households are under-represented in our estimation sample, we reweight the data to ensure that the empirical distribution of $\hat{\theta}_i$ corresponds to that of the commune population.[31] Next, we parametrically estimate the distribution of preferences G using a truncated normal density fit to the empirical distribution of $\hat{\theta}_i$, where the right truncation point is the sample maximum.

[31]Specifically, we create an augmented representative sample by drawing 120 $\hat{\theta}_i$ (with replacement) from those of the 114 previously unconnected households included in our estimation sample; by construction, this artificial sample has 16 percent previously (as of July 2015) unconnected households as in our commune household listing.

Note that $\hat{\lambda} = 2.7$ implies a substantial wedge between ex-ante and ex-post willingness-to-pay; the median value of this wedge, expressed as the ex-post to ex-ante surplus ratio (at marginal cost), $e^{(\lambda-1)\hat{\theta}_i}$, is 3.25 with an interquartile range of 2.35-4.40.

Exercise 1: Gains from a sophisticated two-part tariff

Consider the unsophisticated social planner, who errs in assuming that preferences for piped water are exogenous and unchanging ($\lambda = 1$), but who ascertains the distribution of preferences $G(\theta)$ along with the short-run price elasticity of demand α.[32] As noted, the optimal two-part tariff (F, p) in this scenario was originally discussed by Auerbach and Pellechio (1978). Given alternative values of average cost K, we compute ex-post social welfare at the contracts offered by the unsophisticated planner, where ex-post social welfare is the long-run ($\lambda = 2.7$) average consumer surplus net of average costs. We repeat this exercise for a sophisticated social planner, one who (correctly) assumes endogenous preferences for piped water and who knows that consumers are ex-ante unaware of having such preferences.

Figure 7 illustrates the benefits of taking endogenous preferences into account in pricing a new utility service. The percentage welfare gain attributable to sophisticated pricing, the difference between the solid and dashed curves, can be extremely large depending on the value of K. To pin down the relevant scenario in our setting, we obtained actual cost figures from An Thinh water utility, consisting of their initial investment in plant and equipment (including the underground pipe network for the three villages of the commune), annual operating expenses, as well as the cost of a home connection and water meter. The water itself is pumped from a nearby river and hence is effectively free. Since the utility is privately owned and run as a business (albeit subject to government rate regulation), these figures probably represent the social opportunity cost of piped water provision in rural Vietnam quite accurately. Be that as it may, at the plant's reported capacity of 10,000 customers and a plausible real interest rate of 5 percent, the implied K works out to about 33 (in the appropriate units).[33]

[32] For concreteness, suppose that this planner commissions a study in which alternative (F, p) contracts are randomly offered to potential customers and the resulting data are used to estimate take-up as a function of (F, p). From the participation constraint, $e^{\theta - \alpha p} \geq F$, and a parametric assumption on G, it is easy to see that both G and α are identified. Note that this experiment is different from the one actually conducted by Lee et al. (2020) in rural Kenya inasmuch as the latter only randomizes F.

[33] Specifically, initial investment was reported to us as US$9.7 million in current dollars, annual operating costs at US$30,000, household connection cost at US$84. We convert all of these figures to monthly per customer flows. In the case of plant and equipment, we do so using a real interest rate equal to the lending rate minus the rate of inflation (both for 2019) as reported for Vietnam in World Bank statistics.

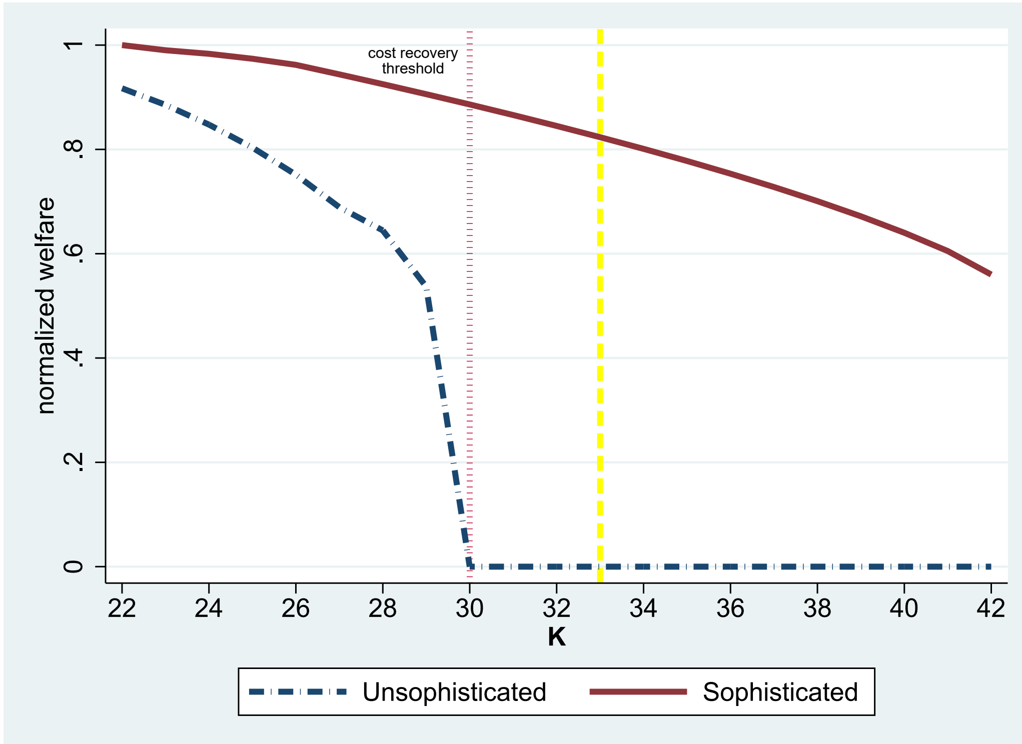

Notes: For different average cost K, the graph plots the ex-post welfare evaluated at the optimal two-part tariff. The vertical dotted pink line represents the threshold K beyond which the unsophisticated planner's pricing plan can no longer balance the utility's budget. Beyond this threshold, we set social welfare under the unsophisticated plan to zero. The vertical dashed yellow line represents the K reported by the An Thinh water utility.

Figure 7: Welfare with and without sophisticated pricing of new utility

At $K = 33$ in Figure 7, we see that the social value of the utility when decisions are taken by an unsophisticated social planner is zero. From this erroneous perspective, there is no combination of markup and connection fee that would allow the utility to cover its costs and, therefore, the project would not be undertaken in the first place. By contrast, a sophisticated social planner, who is attentive to endogenous preferences, would undertake the project, break even, and produce a positive social surplus. This finding suggests caution in relying *solely* on assessments of consumer demand or willingness-to-pay made before consumers have ever connected to judge the economic viability of a utility service.

Exercise 2: Sophisticated connection fee versus subscription plan

We next turn to the best way to price the utility service for a sophisticated planner. Figure 8 shows the welfare gain from switching to a deferred subscription plan from a two-part tariff with markup and connection fee as a function of average cost K. The pink dashed vertical

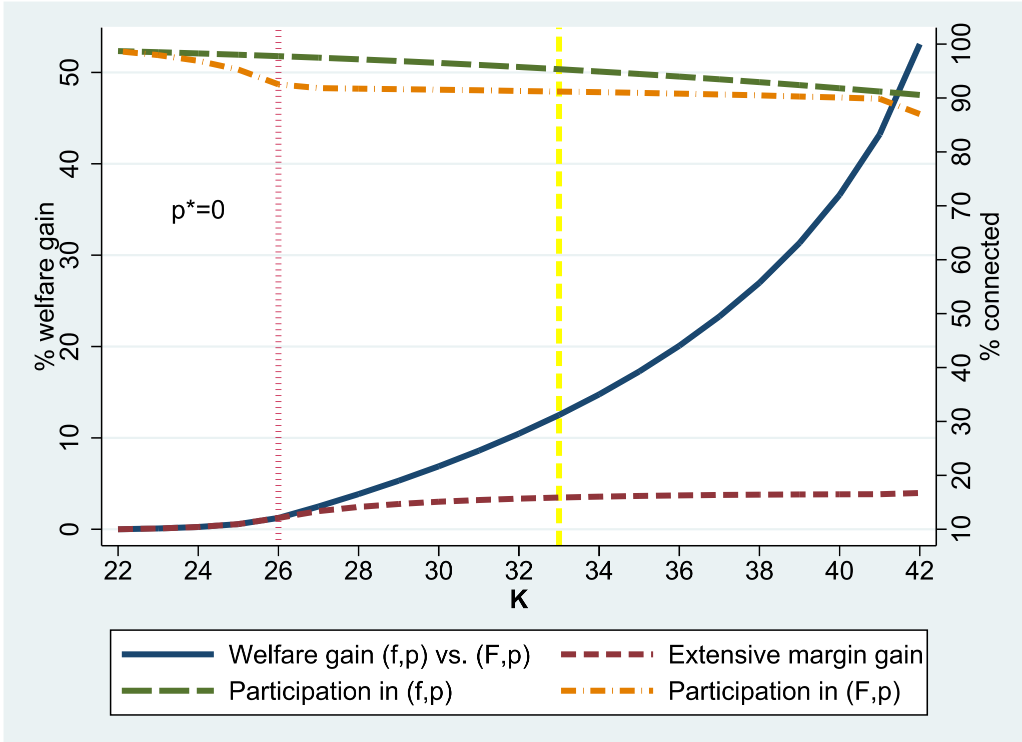

Notes: For different average cost K, the graph plots the percent welfare gain from switching from optimal connection fee contract to optimal subscription plan (solid blue); the welfare gain of same coming strictly at the extensive margin (short-dashed red); the participation rate under the subscription plan (long-dash green); the participation rate under the connection fee contract. The vertical dotted pink line represents the threshold K beyond which markup is positive; the vertical dashed yellow line represents the K reported by the An Thinh water utility.

Figure 8: Connection fee (F, p) vs. subscription fee (f, p) contract

line indicates the critical value of $\hat{K} = 26$ beyond which an optimal connection fee contract is characterized by a positive markup. Indeed, for any $K > \hat{K}$, ex-ante willingness-to-pay is not sufficient to fully recover cost so that a markup is welfare-improving. For lower values of K such that $K \leq \hat{K}$, an optimal connection fee contract does not require a markup but is still dominated by a recurring subscription (without a markup). The solid blue line shows the welfare gain from the contract switch, i.e., from deferring the fee. Depending on the level of costs that the utility needs to cover, the benefits from moving to a subscription plan can be substantial. As the long-dashed green and dash-dotted orange lines indicate, while a subscription plan allows higher participation, the difference in connection rates is never very dramatic, at most around 5 percentage points. Hence, the bulk of the welfare gains when K is large come at the intensive margin rather than at the connection margin, which is to say as a result of the lower deadweight loss of "taxation" that a zero-markup subscription plan allows. Accordingly, the red dashed line plots the welfare gains arising strictly from increased participation and shows that these gains fall in relative terms as K increases.

Finally, in the case of An Thinh water utility, reading off the dashed vertical (yellow) line in Figure 8, we assess that full cost recovery is possible for the utility given long-run consumer demand (as already noted). However, cost recovery cannot be achieved with a connection fee alone, i.e. without a distortionary markup. As a result, there would be a welfare gain of 12 percent in moving to a subscription plan with zero markup, about 70 percent of which gain is coming from the intensive margin.

6 Conclusion

New residential utilities are one of the mileposts of economic development, and yet their pricing and economics of provision have not heretofore been studied in the (plausible) case where preferences for the service evolve endogenously over time. To do so, we generated experimental variation in the price of residential piped water in a setting where households were still transitioning away from a traditional mode of water delivery. Consistent with the new piped water service being an experience good, being habit-forming, or some combination of the two, we find short-term persistence in demand: high consumption today increases consumption in subsequent months.

Such positive intertemporal dependence alters the calculus of two-part tariffs during the introductory phase of a new utility. If consumers do not internalize the impact of their present consumption on their future demand, two-part pricing allows future consumers to subsidize their present selves so as to encourage connection. We have shown in our setting that, whereas long-run demand for the new utility is such that it can fully cover its costs *in principle* through its connection fee and markup, it cannot do so *in practice* without accounting for endogenous preferences. We thus provided a rationale for low take-up distinct from credit constraints and a remedy in the form of deferred payment. We further proposed a subscription plan that fully defers payment of the connection fee until long-run preferences are formed, which achieves a further 12 percent welfare gain over the corresponding optimal connection fee.

References

Dan Acland and Matthew R Levy. Naiveté, projection bias, and habit formation in gym attendance. Management Science, 61(1):146–160, 2015.

Achim Ahrens, Christian Hansen, and Mark Edwin Schaffer. Pdslasso: Stata module for post-selection and post-regularization ols or iv estimation and inference, 1 2019. Achim Ahrens Christian B. Hansen Mark E Schaffer, 2018. "LASSOPACK: Stata module for lasso, square-root lasso, elastic net, ridge, adaptive lasso estimation and cross-validation," Statistical Software Components S458458, Boston College Department of Economics, revised 07 Apr 2018. Stata software for econometric estimation and testing: LASSOPACK and PDSLASSO.

Steffen Ahrens, Inske Pirschel, and Dennis J Snower. A theory of price adjustment under loss aversion. Journal of Economic Behavior & Organization, 134:78–95, 2017.

Hunt Allcott and Todd Rogers. The short-run and long-run effects of behavioral interventions: Experimental evidence from energy conservation. American Economic Review, 104(10):3003–37, October 2014.

T. W. Anderson and Cheng Hsiao. Estimation of dynamic models with error components. Journal of the American Statistical Association, 76(375):598–606, 1981. ISSN 01621459. URL http://www.jstor.org/stable/2287517.

Isaiah Andrews, James H. Stock, and Liyang Sun. Weak instruments in instrumental variables regression: Theory and practice. Annual Review of Economics, 11(1):727–753, 2019. doi: 10.1146/annurev-economics-080218-025643. URL https://doi.org/10.1146/annurev-economics-080218-025643.

Manuel Arellano and Stephen Bond. Some Tests of Specification for Panel Data: Monte Carlo Evidence and an Application to Employment Equations. The Review of Economic Studies, 58(2):277–297, 04 1991. ISSN 0034-6527. doi: 10.2307/2297968. URL https://doi.org/10.2307/2297968.

Alan J Auerbach and Anthony J Pellechio. The two-part tariff and voluntary market participation. The Quarterly Journal of Economics, 92(4):571–587, November 1978.

Ned Augenblick and Matthew Rabin. An experiment on time preference and misprediction in unpleasant tasks. Review of Economic Studies, 86(3):941–975, 2019.

Gary S. Becker and Kevin M. Murphy. A theory of rational addiction. Journal of Political Economy, 96(4):675–700, 1988. ISSN 00223808, 1537534X. URL http://www.jstor.org/stable/1830469.

Gary S. Becker, Michael Grossman, and Kevin M. Murphy. An empirical analysis of cigarette addiction. The American Economic Review, 84(3):396–418, 1994. ISSN 00028282. URL http://www.jstor.org/stable/2118059.

A. Belloni, D. Chen, V. Chernozhukov, and C. Hansen. Sparse models and methods for optimal instruments with an application to eminent domain. Econometrica, 80(6):2369–2429, 2012. doi: 10.3982/ECTA9626. URL https://onlinelibrary.wiley.com/doi/abs/10.3982/ECTA9626.

Dirk Bergemann and Juuso Välimäki. Dynamic pricing of new experience goods. Journal of Political Economy, 114(4):713–743, 2006.

Susanna Berkouwer and Joshua T. Dean. Credit, attention, and externalities in the adoption of energy efficient technologies by low-income households. Manuscript, UC Berkeley, 2021.

Raquel Carrasco, José M. Labeaga, and David J. López-Salido. Consumption and Habits: Evidence from Panel Data. The Economic Journal, 115(500):144–165, 12 2004. ISSN 0013-0133. doi: 10.1111/j.1468-0297.2004.00963.x. URL https://doi.org/10.1111/j.1468-0297.2004.00963.x.

Frank Chaloupka. Rational addictive behavior and cigarette smoking. Journal of Political Economy, 99(4):722–42, 1991. URL https://EconPapers.repec.org/RePEc:ucp:jpolec:v:99:y:1991:i:4:p:722-42.

Francisco Costa and Francois Gerard. Hysteresis and the welfare effect of corrective policies: Theory and evidence from an energy-saving program. Working Paper 24608, National Bureau of Economic Research, May 2018. URL http://www.nber.org/papers/w24608.

Tatyana Deryugina, Alexander MacKay, and Julian Reif. The long-run dynamics of electricity demand: Evidence from municipal aggregation. American Economic Journal: Applied Economics, 12(1):86–114, 2020.

Florencia Devoto, Esther Duflo, Pascaline Dupas, William Parienté, and Vincent Pons. Happiness on tap: Piped water adoption in urban morocco. American Economic Journal: Economic Policy, 4(4):68–99, 2012a.

Florencia Devoto, Esther Duflo, Pascaline Dupas, William Parient, and Vincent Pons. Happiness on tap: Piped water adoption in urban morocco. American Economic Journal: Economic Policy, 4(4):68–99, September 2012b. doi: 10.1257/pol.4.4.68. URL http://www.aeaweb.org/articles.php?doi=10.1257/pol.4.4.68.

Daouda Diakité, Aggey Semenov, and Alban Thomas. A proposal for social pricing of water supply in côte d'ivoire. Journal of Development Economics, 88(2):258–268, 2009.

Pascaline Dupas. Short-run subsidies and long-run adoption of new health products: Evidence from a field experiment. Econometrica, 82(1):197–228, 2014.

Karen E Dynan. Habit formation in consumer preferences: Evidence from panel data. American Economic Review, 90(3):391–406, 2000.

Jonathan Gruber and Botond Köszegi. Is Addiction Rational? Theory and Evidence*. The Quarterly Journal of Economics, 116(4):1261–1303, 11 2001. ISSN 0033-5533. doi: 10.1162/003355301753265570. URL https://doi.org/10.1162/003355301753265570.

Bruce E. Hansen. Threshold effects in non-dynamic panels: Estimation, testing, and inference. Journal of Econometrics, 93(2):345–368, December 1999. URL https://ideas.repec.org/a/eee/econom/v93y1999i2p345-368.html.

Reshmaan Hussam, Atonu Rabbani, Giovanni Reggiani, and Natalia Rigol. Habit Formation and Rational Addiction: A Field Experiment in Handwashing. Harvard Business School Working Papers 18-030, Harvard Business School, September 2017. URL https://ideas.repec.org/p/hbs/wpaper/18-030.html.

Koichiro Ito. Do consumers respond to marginal or average price? evidence from nonlinear electricity pricing. The American Economic Review, 104(2):537–563, 2014.

Koichiro Ito and Shuang Zhang. Reforming inefficient energy pricing: Evidence from china. Working Paper 26853, National Bureau of Economic Research, March 2020. URL http://www.nber.org/papers/w26853.

Koichiro Ito, Takanori Ida, and Makoto Tanaka. Moral suasion and economic incentives: Field experimental evidence from energy demand. American Economic Journal: Economic Policy, 10(1):240–67, 2018.

Paul L Joskow. Regulation of natural monopoly. Handbook of law and economics, 2:1227–1348, 2007.

Botond Kőszegi and Matthew Rabin. A model of reference-dependent preferences. The Quarterly Journal of Economics, 121(4):1133–1165, 2006.

Jean-Jacques Laffont and Jean Tirole. A theory of incentives in procurement and regulation. MIT press, 1993.

Kenneth Lee, Edward Miguel, and Catherine Wolfram. Experimental evidence on the economics of rural electrification. Journal of Political Economy, 128(4):1523–1565, 2020.

George Loewenstein, Ted O'Donoghue, and Matthew Rabin. Projection bias in predicting future utility. The Quarterly Journal of economics, 118(4):1209–1248, 2003.

Shaun McRae. Infrastructure quality and the subsidy trap. The American Economic Review, 105(1):35–66, 2015.

Costas Meghir and Guglielmo Weber. Intertemporal nonseparability or borrowing restrictions? a disaggregate analysis using a u.s. consumption panel. Econometrica, 64(5):1151–81, 1996. URL https://EconPapers.repec.org/RePEc:ecm:emetrp:v:64:y:1996:i:5:p:1151-81.

Phillip Nelson. Information and consumer behavior. Journal of political economy, 78(2):311–329, 1970.

Yew-Kwang Ng and Mendel Weisser. Optimal pricing with a budget constraint–the case of the two-part tariff. The Review of Economic Studies, 41(3):337–345, 1974.

Jos Luis Montiel Olea and Carolin Pflueger. A robust test for weak instruments. Journal of Business & Economic Statistics, 31(3):358–369, 2013. doi: 10.1080/00401706.2013.806694. URL https://doi.org/10.1080/00401706.2013.806694.

Robert A. Pollak. Habit formation and dynamic demand functions. Journal of Political Economy, 78(4):745–763, 1970. URL http://www.jstor.org/stable/1829929.

Peter C Reiss and Matthew W White. Household electricity demand, revisited. The Review of Economic Studies, 72(3):853–883, July 2005.

Carl Shapiro. Optimal pricing of experience goods. The Bell Journal of Economics, pages 497–507, 1983.

Andrea Szabó. The value of free water: Analyzing south africa's free basic water policy. Econometrica, 83(5):1913–1961, 2015.

Vanessa Taylor and Frank Trentmann. Liquid politics: water and the politics of everyday life in the modern city. Past and Present, 211(1):199–241, 2011.

Richard Thaler. Mental accounting and consumer choice. Marketing science, 4(3):199–214, 1985.

Quang H. Vuong. Likelihood ratio tests for model selection and non-nested hypotheses. Econometrica, 57(2):307–333, 1989. ISSN 00129682, 14680262. URL http://www.jstor.org/stable/1912557.

Robert Wilson. Nonlinear pricing. Oxford University Press on Demand, 1993.

Appendix (Online)

Table of Contents

A Proofs

A.1 Proof of Lemma 1

We define the expected log consumption at the time of connection as $e_t(\theta, p) = \mathbb{E}_0 \log c_t(\theta, \bar{C}_{t-1}, p|\xi_t)$. Long-run consumption is thus $\lim_{t\to\infty} e_t(\theta, p)$.

From first-order condition (7) and equation (4), we write the law of motion for $\{e_t\}$ (and in doing so omit the reference to (θ, p) for notational simplicity):

$$e_t = \theta - \alpha p + \beta \sum_{k=1}^{T} \gamma^{k-1} e_{t-k}, \tag{A.1}$$

with $e_t = 0$ for all values of $t \leq 0$. First, we show by induction that $\{e_t\}$ is increasing. To do so, we take the difference

$$e_t - e_{t-1} = \beta \sum_{k=1}^{T} \gamma^{k-1} (e_{t-k} - e_{t-k-1}).$$

Given that the stock of past consumption is equal to zero at time $t = 0$, the difference is nonnegative for $t = 1$ and positive for $t = 2$. Suppose that for some $t > 2$, the difference $(e_t - e_{t-1})$ is positive for all periods prior to t. Then $e_{t+1} - e_t$, a sum of positive terms, is positive. By induction, therefore, the sequence $\{e_t\}$ is nondecreasing for all t. Next, we show that it is bounded above. To do so, we note that since it is increasing, we can bound e_t above by bounding each term under the sum sign in equation (A.1) by e_t and write

$$e_t \leq (\theta - \alpha p) + \beta \sum_{k=1}^{T} \gamma^{k-1} e_t. \tag{A.2}$$

If $\lambda > 1$, then we have for all t, $e_t \leq e_\infty$. Thus, since (9) is assumed to hold, sequence $\{e_t\}$ is increasing and bounded above by e_∞. It thus converges and since e_∞ is the unique fixed point defined by (A.1), sequence $\{e_t\}$ converges to e_∞.

A.2 Proof of Proposition 1

If $e^\theta \geq \alpha K$, the participation constraint is not binding for $F = K$; demand persistence has no efficiency implications. Likewise, when $e^{\lambda\theta} < \alpha K$, there is no price that both satisfies the participation and budget constraints so that the utility is not funded in the first place, which is also the socially optimal outcome. We now focus on $e^\theta < \alpha K \leq e^{\lambda\theta}$. With a contract consisting of a connection fee only, no fee will ensure participation while satisfying the budget constraint. If the utility can charge a markup, then it can raise $p \cdot e^{\lambda(\theta - \alpha p)}$.

A binding budget constraint thus implies that

$$F = K - pe^{\lambda(\theta - \alpha p)}$$

and the participation constraint becomes

$$e^{\theta - \alpha p} - \alpha F \geq 0,$$

so that we can rewrite the participation constraint as

$$(1 + \alpha p e^{\lambda}) \cdot e^{\theta - \alpha p} \geq \alpha K. \tag{A.3}$$

There exists a $p^m \in (0, \theta/\alpha)$ such that the left-hand side of inequality (A.3) is increasing over the interval $(0, p^m)$ and decreasing over $(p^m, 0)$. Thus, there exists $K^m \in (e^{\theta}, e^{\lambda\theta})$ such that the utility can be financed by a two-part tariff with positive markup ($p > 0$) and below-average-cost connection fee ($F < K$).

A.3 Proof of Proposition 2

Following the same logic as in the Proof of Proposition 1, a binding budget constraint implies that a monthly subscription is defined (as before) by

$$f = K - pe^{\lambda(\theta - \alpha p)}$$

so that the participation constraint for the future consumer can be written

$$(1 + \alpha p) \cdot e^{\lambda(\theta - \alpha p)} \geq \alpha K. \tag{A.4}$$

The left-hand side of inequality (A.4) is now decreasing in p so that the optimal price contract is characterized by zero markup ($p = 0$) and monthly subscription $f = K$.

B Additional Figures and Tables

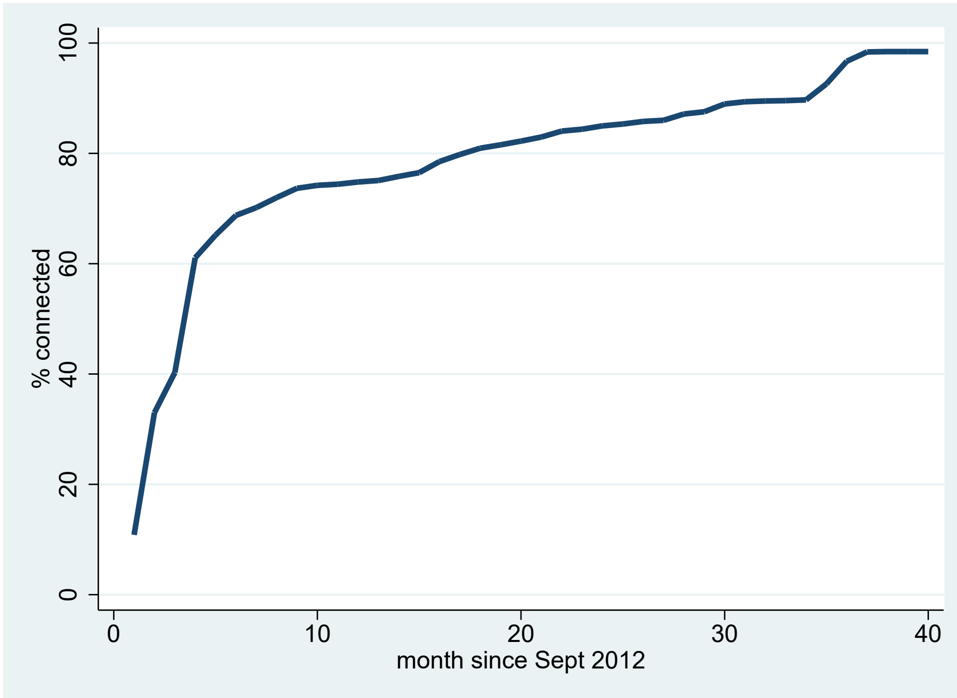

Notes: Percentage of experimental sample (1462 households) connected to piped water by month since utility began operations in September 2012.

Figure B.1: Evolution of piped water connections

Table B.1: Water sources and uses

		How used		
Source	use	cook/drink	shower/bathroom	general hygiene
Piped	1459	1155	1435	154
Rainwater	975	806	478	44
Groundwater	490	6	412	89
Bottled	102	101	1	0
River/lake	8	0	6	3

Notes: Figures are number of households (sample size = 1,462 households). Multiple uses possible for a given water source.

Table B.2: Descriptive Statistics and Balance Tests

	Experimental groups				*p*-values			
	1	2	3	Overall	1 vs 2	1 vs 3	2 vs 3	joint
Water use (pre-subsidy)	6.74	7.07	6.72	6.85	0.29	0.95	0.26	0.45
	(0.21)	(0.22)	(0.21)	(0.12)				
Person days/1000	0.279	0.289	0.284	0.284	0.24	0.55	0.59	0.51
	(0.006)	(0.006)	(0.006)	(0.003)				
Baseline survey characteristics:								
log PC expenditures	7.25	7.23	7.23	7.24	0.54	0.55	0.99	0.78
	(0.02)	(0.02)	(0.02)	(0.01)				
log dwelling area	4.40	4.41	4.41	4.41	0.74	0.79	0.95	0.94
	(0.03)	(0.03)	(0.03)	(0.02)				
log dwelling value	12.91	13.00	12.98	12.96	0.11	0.23	0.71	0.25
	(0.04)	(0.04)	(0.04)	(0.02)				
Rainwater use (0/1)	0.68	0.65	0.67	0.67	0.23	0.78	0.35	0.45
	(0.02)	(0.02)	(0.02)	(0.01)				
Rainwater tank volume (IHS)	2.18	2.15	2.24	2.19	0.64	0.30	0.13	0.30
	(0.04)	(0.04)	(0.04)	(0.02)				
Household size	3.44	3.56	3.43	3.47	0.20	0.95	0.18	0.32
	(0.07)	(0.07)	(0.07)	(0.04)				
Proportion male	0.46	0.48	0.46	0.47	0.41	0.91	0.36	0.61
	(0.01)	(0.01)	(0.01)	(0.01)				
Proportion over 65 yr	0.17	0.14	0.16	0.16	0.15	0.76	0.25	0.32
	(0.02)	(0.01)	(0.01)	(0.01)				
Proportion under 5 yr	0.06	0.07	0.07	0.07	0.13	0.39	0.52	0.32
	(0.01)	(0.01)	(0.01)	(0.00)				
Head is male	0.80	0.81	0.81	0.81	0.77	0.83	0.93	0.96
	(0.02)	(0.02)	(0.02)	(0.01)				
Head no educ	0.12	0.11	0.10	0.11	0.50	0.34	0.79	0.62
	(0.01)	(0.01)	(0.01)	(0.01)				
Head primary ed	0.17	0.20	0.21	0.20	0.21	0.08	0.64	0.20
	(0.02)	(0.02)	(0.02)	(0.01)				
Head lower sec ed	0.43	0.47	0.44	0.45	0.16	0.53	0.43	0.37
	(0.02)	(0.02)	(0.02)	(0.01)				
Head higher sec ed	0.21	0.15	0.16	0.17	0.02	0.03	0.91	0.03
	(0.02)	(0.02)	(0.02)	(0.01)				
N	487	485	490	1,462				

Notes: Figures in first four columns are means (standard deviations). Water use is measured in cu. meters/month. Pre-subsidy phase is July-Dec 2015 when all groups faced same linearized price. Person-days in the household is measured at the quarterly frequency and averaged over the entire experiment. For rainwater tank, inverse hyperbolic sine (IHS) of tank volume in m^3.

C Robustness to intrahousehold water arbitrage

If households took advantage of the discount offered to their experimental group by reselling piped water to neighboring households, then we should see a greater price response during these periods. Thus, we re-run the household fixed effects regressions in Table 1 by including dummies for whether the household shared in or out piped water during a quarter in which there were experimental price differentials, or, alternatively, by dropping observations in which either such dummy has a value of one. The resulting discount and price coefficients, reported in Table C.1, are extremely close to their counterparts in Table 1, indicating that intra-household arbitrage is not a threat to our experiment.

Table C.1: Robustness to potential water arbitrage

	(1)	(2)	(3)	(4)	(5)	(6)
50 percent discount	0.0338	0.0327				
	(0.0146)	(0.0144)				
75 percent discount	0.0846	0.0820				
	(0.0151)	(0.0148)				
Any discount			0.0581	0.0562		
			(0.0133)	(0.0130)		
Water price					-0.0761	-0.0734
					(0.0154)	(0.0151)
Share-in	-0.0581		-0.0579		-0.0578	
	(0.0485)		(0.0484)		(0.0484)	
Share-out	0.141		0.142		0.141	
	(0.0432)		(0.0432)		(0.0431)	
Person-days/1000	1.557	1.556	1.560	1.558	1.559	1.557
	(0.102)	(0.101)	(0.102)	(0.102)	(0.102)	(0.102)
R^2	0.659	0.660	0.659	0.660	0.659	0.660
Observations	46,784	46,658	46,784	46,658	46,784	46,658
Month dummies	YES	YES	YES	YES	YES	YES
HH fixed effects	YES	YES	YES	YES	YES	YES

Notes: Robust standard errors in parentheses clustered at household level. Dependent variable is log monthly piped water consumption. Shared-in is a quarterly-level household dummy that takes a value of one if the household received piped water from another household during a quarter in which there were price differences across experimental groups. Shared-out is a quarterly-level household dummy that takes a value of one if the household gave piped water to a household with its own piped water connection during a quarter in which there were price differences across experimental groups. Columns 2, 4, and 6 exclude any observation for which the share-in or the share-out dummy is equal to 1. Person-days is measured at the quarterly frequency.

D Heterogeneous consumption responses

Table D.1: Heterogeneity tests

Test (H_0 : No difference in partial elasticity)	p-value
1) Hot vs. cold season	0.127
2) Per-capita expenditure quintile	0.670
3) Log area of dwelling (m^2)	0.191
4) Log value of dwelling	0.152
5) Use of rainwater	0.975
6) Rainwater tank volume (IHS)	0.920

Notes: See notes to Table 1, specification (6), which includes both household and month fixed effects. Significance test p-values are for coefficient(s) of interaction(s) of respective variable(s) with price. In row (1), a dummy for whether month is in the hot season (May-Sept.); in row (2), a set dummies for quintile of per capita expenditures measured at baseline; in rows (3) and (4), respectively, the log area and value of the dwelling at baseline; in row (5), a dummy for whether household was using rainwater at baseline (yes = 67 percent); in row (6), the inverse hyperbolic sine (IHS) of rainwater tank volume in m^3 (4 percent have no tank).

Table D.1 reports test of heterogeneity in the responsiveness of monthly water usage to price based on specification (6) of Table 1 (effective water price with household fixed effects). First, piped water demand exhibits substantial seasonality. In our sample, monthly consumption during the hot season from May to September averages about 12 percent higher than during the rest of the year, presumably reflecting the greater demand for showers.[34] We test for seasonality in the price response by including an interaction between price and a hot season dummy. As shown in the first row of Table D.1, we find no evidence of this.

Next, we consider heterogeneity in price response by household income or wealth (using per-capita household expenditures quintiles), by the area and value of the dwelling, which may not only reflect wealth but also the number and/or quality of plumbing fixtures, and by the use of rainwater (for any domestic purpose) and by the capacity of the rainwater tank, all measured at baseline. These last two variables indicate ability to substitute out of piped water. In no case do we find significant demand interactions with the water price. These findings imply that household preferences are separable between piped water and other goods as well as with previous investments in piped water substitutes.

[34]Higher hot season demand for piped water is not due to a greater likelihood of rainwater tanks running dry. In Vietnam, rainfall is plentiful year-round; the hot season, in particular, is the wettest part of the year.

E Power calculations for the simple persistence test

We assume a data generating process described by

$$\log c_{it} = -\alpha p_{it} + \beta \log c_{it-1} + u_{it}, \tag{E.1}$$

where α is the price response (elasticity) and β is the degree of (first-order) persistence. Taking averages of (E.1) across households in group g yields

$$\overline{\log c}_{g,t} = -\alpha p_{g,t} + \beta \overline{\log c}_{g,t-1}. \tag{E.2}$$

Now, let ρ_t be the reduced-form treatment effect of being in experimental group 1 versus group 3 in month t, i.e., $\rho_t = \overline{\log c}_{1,t} - \overline{\log c}_{3,t}$. Recall from the upper panel of Figure 2 that until $t = 6$, groups 1 and 3 will face the same price and, hence, $\overline{\log c}_{1,6} = \overline{\log c}_{3,6}$. Subsequently, from month $t = 7$ till $t = 18$, group 1 will face first $p = 0.5$ for six months and then $p = 0.25$ for six months, whereas group 3 will face $p = 1$ throughout the relevant period of the experiment. From equation (E.2) and the definition of ρ_t, it follows that

$$\rho_{19} = \alpha(0.75 + 0.5\beta^6) \sum_{\tau=1}^{6} \beta^{\tau}. \tag{E.3}$$

Using this formula, we compute the effect size at $t = 19$ for our power calculations as a function of (α, β). Effect size at $t = 20$, two months into the no discount semester, is $\rho_{20} = \beta\rho_{19}$.

We assume a sample of 500 households per group. For the standard deviation of log monthly water consumption, we use a value of 1.07 based on all of An Thinh's customers in My Huong commune from July 2014 to June 2015. In terms of priors from the literature, there are few credible estimates of the demand elasticity for piped water in developing countries. Diakité et al. (2009) find an elasticity of about 0.8 for Cote D'Ivoire, which arguably would be an underestimate in our context given the high substitutability between piped water and rainwater. At any rate, in light of this uncertainty and the complete lack of priors on β, we compute power for a wide range of (α, β) values in panels (a) and (b) of Figure E.1 . In the red regions, power is adequate. Thus, for example, if the true $\alpha = 0.8$ (horizontal dashed line), we would have adequate power in month 19 for true values of $\beta > 0.15$. We have reasons to believe, therefore, that our simple reduced form test for demand persistence is well-powered. By contrast, under the same scenario in month 20, we would have adequate power for true values of $\beta > 0.35$. Intuitively, power goes down as we use data from later months because the effect of consumption differences induced earlier in the experiment dissipate over time. Since, even ex-ante, the month 20 test is well-powered only for moderately high demand persistence, we do not focus on it.

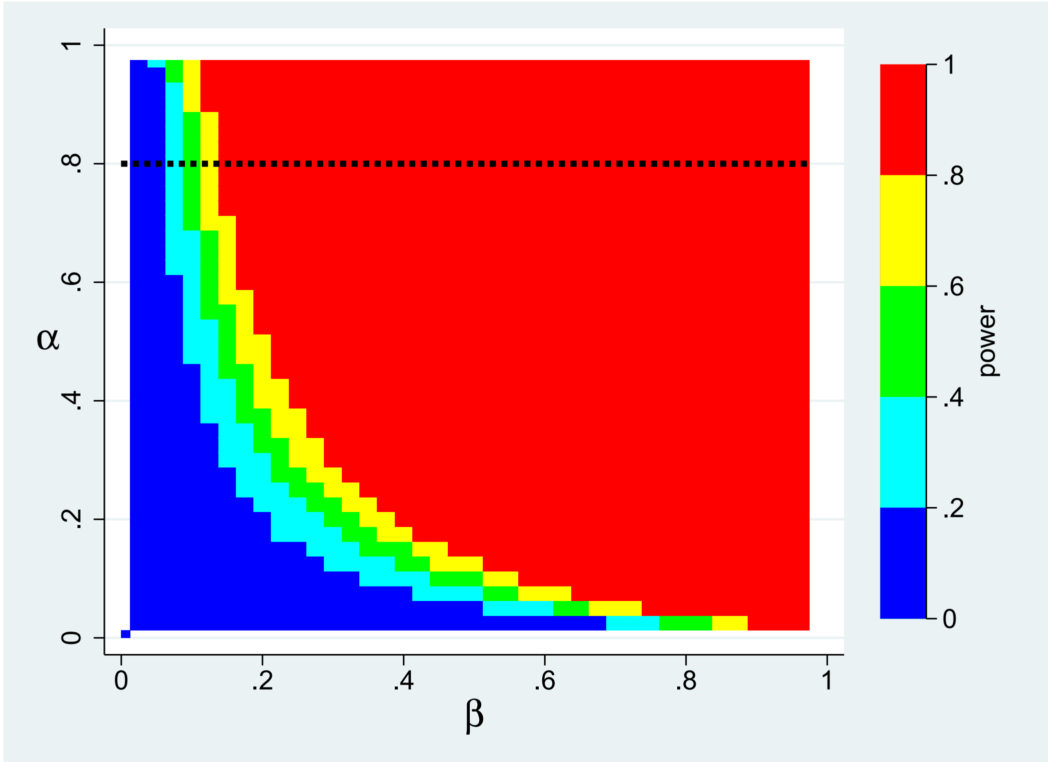

(a) Month $t = 19$

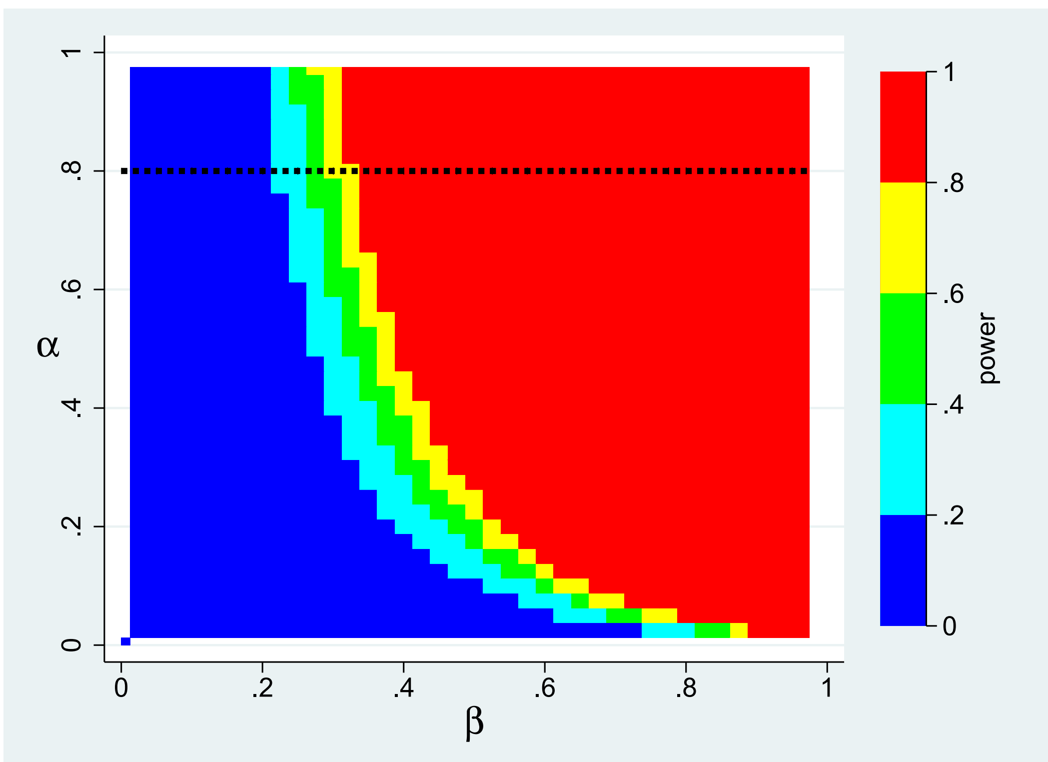

(b) Month $t = 20$

Notes: Heat map of power of test of null of no difference between groups 1 and 3 against true effect size ρ_t, where $t = 19$ in panel (a) and $t = 20$ in panel (b). On the vertical axis, are hypothetical values of the demand elasticity α with the horizontal dashed line indicating the Diakité et al. (2009) estimate of 0.8 and, on the horizontal axis, are hypothetical values of the persistence parameter β.

Figure E.1: Power calculations for comparison of group 1 vs. group 3

F Testing consumer loss aversion

As pointed out by, e.g., Ahrens et al. (2017), a corollary of loss aversion (as in Kőszegi and Rabin 2006) is that demand is more elastic with respect to a price increase than to a price decrease, as the former entails a utility loss quite apart from any income effect of the price change. To test the null hypothesis that piped water price increases and decreases induce symmetric demand responses, we move to the more appropriate differenced form

$$\Delta^L \log c_{it} = a^+ \mathbb{1}_{\Delta^L p_{it}>0} \cdot \Delta^L p_{it} + a^- \mathbb{1}_{\Delta^L p_{it}\leq 0} \cdot \Delta^L p_{it} + b\Delta^L N_{it} + d_t + u_{it}, \quad \text{(F.1)}$$

where, e.g., $\Delta^L p_{it} \equiv p_{it} - p_{it-L}$, a^+ is the elasticity with respect to price increases, a^- is the elasticity with respect to price decreases, and d_t is a month fixed effect. By using larger values of L, we exploit more of the variation across experimental price regimes, but at the efficiency cost of a smaller estimation sample. For instance, at the bottom of Table F.1, we see that at $L = 4$, a little over half of the month t to $t - L$ consumption changes are associated with no change in price at all, whereas this fraction falls to just one quarter at $L = 9$. At any rate, for no value of L can we reject the null hypothesis of symmetry ($a^+ = a^-$).

Table F.1: Testing symmetry of price responses

	$L = 4$	$L = 5$	$L = 6$	$L = 7$	$L = 8$	$L = 9$
a^+	-0.0577	-0.0726	-0.0819	-0.0795	-0.0821	-0.0824
	(0.0195)	(0.0190)	(0.0198)	(0.0212)	(0.0225)	(0.0231)
a^-	-0.0951	-0.101	-0.0936	-0.0935	-0.0902	-0.0767
	(0.0319)	(0.0307)	(0.0305)	(0.0318)	(0.0333)	(0.0348)
$a^+ - a^-$	-0.0374	-0.0280	-0.0117	-0.0140	-0.00810	0.00564
	(0.0374)	(0.0360)	(0.0366)	(0.0398)	(0.0443)	(0.0468)
$H_0 : a^+ = a^-$ (p-value)	0.317	0.436	0.749	0.725	0.855	0.904
Observations	40,936	39,474	38,012	36,550	35,088	33,626
R^2	0.072	0.082	0.094	0.093	0.089	0.098
% $\Delta^L p_{it} < 0$	27	35	41	41	42	42
% $\Delta^L p_{it} = 0$	54	41	30	28	26	25
% $\Delta^L p_{it} > 0$	19	25	29	31	32	33

Notes: Robust standard errors in parentheses clustered at household level. Dependent variable is change in log monthly piped water consumption and L is the lag (in months) for differencing the data. Coefficient on person-days (at the quarterly frequency) is not reported.

G Testing consumer sophistication vs. naivete

There is a considerable micro-econometric literature generalizing consumption Euler equations to allow for intertemporal nonseparability (see, e.g., Meghir and Weber, 1996; Dynan, 2000; Carrasco et al., 2004, among others). Identification involves using lagged endogenous variables (e.g., consumption) as instruments along the lines of Arellano and Bond (1991) under the assumption that preference shocks are serially uncorrelated. A second strand of this literature, emanating from Becker and Murphy (1988), considers whether consumers of addictive goods, principally tobacco, are "rational" in the sense that their consumption responds to future prices, under the assumption that future prices are anticipated.[35] Here we provide a test of sophisticated or forward-looking behavior in piped water consumption that relies neither on lack of serial correlation in preferences nor on consumers anticipating future price changes.

G.1 An Euler equation functional form test

In the full rational/sophisticated case, optimal consumption is given by Euler equation:

$$(\theta - \log c_t + \beta \bar{C}_{t-1} + \xi_t) + \delta\beta \mathbb{E}_t \sum_{k=1}^{\tau} (\delta\gamma)^{k-1} \frac{c_{t+k}}{c_t} = \alpha p_t. \tag{G.1}$$

The first term in the Euler equation is the marginal instantaneous utility of water consumption, while the second term captures the persisting effect of a marginal increase in today's consumption. Note that, owing to the quasi-linearity assumption, the second term does not directly include future prices. Nevertheless, future prices indirectly affect current choices via future consumption *insofar as they are anticipated by the consumer.*

Dropping the expectation operator in equation (G.1), replacing future values of consumption by realized values, differencing over time using lag-length L, adding month dummies and rearranging, yields

$$\Delta^L \log c_{it} \quad = \quad -\alpha \Delta^L p_{it} + \beta \sum_{k=1}^{\tau} \gamma^{k-1} \Delta^L \left(\log c_{it-k} + \delta^k \frac{c_{it+k}}{c_{it}} \right) + \omega \Delta^L N_{it} + \Delta^L \zeta_t + \nu_{it}, \tag{G.2}$$

where ν_{it} is an iid disturbance consisting of two components: $\Delta^L \varepsilon_{it}$, as in equation (12), and (a time difference of) the error in forecasting the future marginal utility of piped water consumption. Under rational expectations, ν_{it} is uncorrelated with anything in the period t information set; in particular, with the current price.

Since no parametric restriction leads from (G.2) to (12), the static/myopic demand function is not nested within the fully rational Euler equation. Nevertheless, we can still test the

[35] See Chaloupka (1991) and Becker et al. (1994), among others. Gruber and Köszegi (2001) call into question future price anticipation in the context of testing rational addiction models.

one against the other using the non-nested hypothesis test proposed in Vuong (1989), which compares the squared residuals from the two models observation-by-observation. In other words, we cast the test for sophistication as one for the best fitting functional form of the Euler equation.

Such a functional form test, however, would seem to require estimating equation (G.2). Yet, even after fixing δ, identification of β and γ from equation (G.2) is more tenuous than in the static case. To see this, suppose that $\tau = 1$ so that γ no longer appears in equation (G.2) and we simply have $\Delta^L \log c_{it-1} + \delta\Delta^L \frac{c_{it+1}}{c_{it}}$ on the RHS. As in the naive case discussed earlier, lagged prices and functions thereof are valid instruments but are unlikely to be strongly correlated with the second component of the endogenous RHS variable, $\Delta^L \frac{c_{it+1}}{c_{it}}$. In principle, *future* prices could also serve as instruments (as in Chaloupka, 1991; Becker et al., 1994), but only if they are in the agents' period t information set, i.e., insofar as they are anticipated. Unanticipated prices may have predictive power in the first stage for $\Delta^L \log c_{it-1} + \delta\Delta^L \frac{c_{it+1}}{c_{it}}$ but they would be correlated with the Euler equation forecast error and hence violate the exclusion restriction. Anticipation of future price changes was ruled out in our experiment by not announcing the discounts or their duration in advance. Hence, by construction, all future water price changes should be viewed as unanticipated.

To avoid these identification issues, or rather to test for sophistication when future prices are unanticipated, we consider an alternative strategy. Let β_0 and γ_0 be the true values of β and γ, respectively, from equation (G.2). First, we fix β and γ at some candidate values $\tilde{\beta}$ and $\tilde{\gamma}$ and construct

$$\tilde{Y}_{it} = \Delta^L \log c_{it} - \tilde{\beta} \sum_{k=1}^{\tau} \tilde{\gamma}^{k-1} \Delta^L \left(\log c_{it-k} + \delta^k \frac{c_{it+k}}{c_{it}} \right) \tag{G.3}$$

for given values of δ and τ. Next, we run the regression

$$\tilde{Y}_{it} = -\alpha\Delta^L p_{it} + \omega\Delta^L N_{it} + \Delta^L \zeta_t + \nu_{it} \tag{G.4}$$

and compare the squared residuals to those from equation (12) as per Vuong (1989). If we repeat this procedure on a grid of (β, γ) covering the entire feasible parameter space and always obtain a rejection of the sophisticated Euler equation in favor of the naive one, then we must also reject sophistication at the true values β_0 and γ_0, no matter what those are, provided that they are feasible. What makes this procedure work is that, even in the sophisticated case, the feasible parameter space is defined by (9), i.e. $\beta\frac{1-\gamma^\tau}{1-\gamma} \leq 1$, which implies that the set of β and γ combinations is bounded.

G.2 Results of the test

We run the horse-race between static and fully rational versions of the piped water Euler equation, fixing $\tau = 3$, the optimal value from the static estimation, and monthly discount factor $\delta = 0.98$ ($\delta = 0.95$ or 0.99 yield virtually identical results). We compute the rational Euler equation residuals, as described above, on a grid of points covering the feasible $(\beta, \gamma)-$space as defined by (9), i.e., the space bounded by $\beta = 0$, $\gamma = 0$, and $\beta(1+\gamma+\gamma^2) = 1$. We then compare these residuals to those from the static first-order condition estimated in the main text.

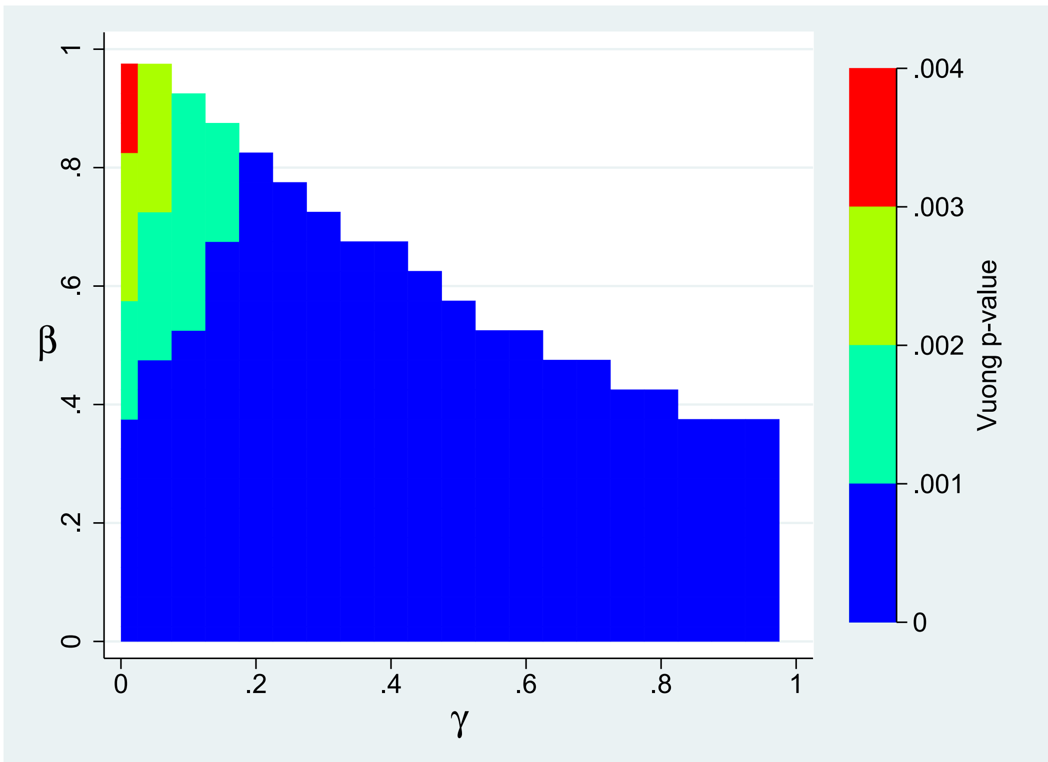

Notes: Each point on heat-map represents a p-value for a non-nested hypothesis test Vuong (1989) of the static first-order condition equation against the rational Euler equation (null) at a feasible (β, γ) combination. The feasible set is bounded by the curve $\beta(1 + \gamma + \gamma^2) = 1$.

Figure G.1: Test of sophistication vs. naivete

Figure G.1 shows the p-values of Vuong's (1989) non-nested hypothesis test at 270 equally spaced values of $(\tilde{\beta}, \tilde{\gamma})$ within the feasible set. In each case, the p-value is well above conventional significance levels, indicating that the static first-order condition fits the data better than the rational Euler equation. In particular, this must be the case at the true values (β_0, γ_0), assuming that they are feasible. Our results thus suggest that households are not fully sophisticated or forward-looking about their piped water use; they do not appear to "invest" in their piped water use today as though their future utility depended on it.