

Stanford | King Center on
Global Development

Environmental Externalities and Free-riding in the Household

Seema Jayachandran

B. Kelsey Jack

Sarojini Rao

November, 2019

Working Paper No. 1057

Environmental externalities and free-riding in the household*

B. Kelsey Jack[†]

Seema Jayachandran[‡]

Sarojini Rao[§]

November 9, 2019

Abstract

Besides generating negative environmental externalities, a household's water consumption entails another "market failure": household members free-ride off each other and overconsume. This problem stems from the difficulty of attributing usage to specific individuals. We document the importance of this phenomenon in urban Zambia by combining utility billing records and randomized person-specific price variation. We derive and empirically confirm the following prediction: Individuals with weaker incentives to conserve under the household's financial arrangements reduce water use more when their person-specific price rises. Our results offer a novel explanation for the low price sensitivity of residential water (and electricity) consumption.

*We thank the Southern Water and Sewerage Company for their collaboration on the project. We have benefited from comments made by audience members at numerous seminars and conferences, and from conversations with Nava Ashraf, Rebecca Dizon-Ross, Gilbert Metcalf, James Sallee, Dmitry Taubinsky, and Alessandra Voena. Flavio Malagutti, Lorenzo Uribe, Lydia Kim, and Alejandro Favela provided outstanding research assistance. We also thank Amanda Kohn for designing an information flyer used in the study. The International Growth Centre and J-PAL's Urban Services Initiative and Governance Initiative provided financial support. This RCT was registered in the American Economic Association Registry for randomized control trials under number AEARCTR-0000660. All errors are our own.

[†]Department of Economics, University of California, Santa Barbara, kelseyjack@ucsb.edu

[‡]Department of Economics, Northwestern University, seema@northwestern.edu

[§]Virginia Department of Planning and Budget, sarojini.rao@dpb.virginia.gov

Introduction

Because of negative environmental externalities, water or energy consumption that is privately optimal exceeds the socially optimal level. This paper highlights a second reason that water and energy are over-consumed: Household members have an opportunity to free-ride off each other. This arises because the billing unit is the household, and consumption is difficult to attribute to specific individuals. The misalignment of incentives within the household is a potential contributor to the inelastic demand for water and electricity observed among residential consumers in many settings.

This intrahousehold problem is analogous to moral hazard in teams. Because an individual bears the full cost of her conservation effort but shares the benefits (savings on the utility bill) with the rest of the household, conservation is below the household’s Pareto optimal level. While intrahousehold decision-making is often modeled in a cooperative framework, two features of household utilities lead to non-cooperative decision-making in this domain. First, piped water and electricity are not purchased individually; the utility bill pools all household members’ usage. Second, it is difficult for household members to back out individual-level use.¹

We derive and implement a test of free-riding based on the observation that members of a household often differ in their financial stake in lowering the utility bill. We refer to the person who bears most of the financial cost of high utility bills as the “effective bill payer.” The testable prediction is that a person-specific price change will have a larger effect on consumption for someone who is not the effective bill payer. The intuition is that this person has weak status-quo incentives to conserve, so the personal price change represents a larger proportional change in her financial incentive to conserve.

We test this prediction by collaborating with the water utility in Livingstone, Zambia.² We combine surveys of 1,282 married couples who are customers of the water utility with monthly billing data. We overlay a randomized intervention that varies the effective price of water at either the individual or household level. Specifically, the intervention offers a financial incentive to reduce water consumption, which is akin to a price increase over a certain range of consumption.

¹We conducted a small survey (a) in Lusaka, Zambia and (b) among mTurk users in the US. When asked for which consumption categories is tracking quantities most difficult, water and electricity were the most common responses. This was true both for own and spouse’s consumption. See Appendix Figure A.1.

²Livingstone’s water source is the Zambezi River. The city faces periodic water shortages when the river level is low (NWASCO 2015). Externalities from Livingstone’s water use also include water shortages for farmers downstream and for wildlife.

In some households, we inform only the man about the rewards program, and in other households, only the woman. These individual-specific incentives, in essence, generate a price change that is fully borne by the individual. While the randomization is based on gender, our prediction pertains to who has the strongest status-quo stake in keeping the water bill low. Thus, we asked survey questions that allow us to ascertain which spouse is the “effective bill payer.” We also have a treatment arm in which the prospect of the reward is communicated to the couple, which acts like an increase in the household-level price. We then examine households’ water usage for the two to nine months the incentives are in place.

As predicted, the response to the individual incentive depends on who is targeted: water use declines considerably more if the incentive is given to the non bill payer rather than the bill payer. Men are most often the effective bill payers, but our finding does not simply reflect heterogeneity by gender. When we simultaneously control for the gender of the incentive recipient, we continue to find a larger effect when the recipient is the non bill payer.

Our study links two previously unconnected strands of literature, on environmental externalities and on intrahousehold decision-making. Our contribution to the literature on corrective pricing in environmental economics is to highlight a previously undiscussed reason that consumers under-respond to utility prices. We complement previous work on misperceptions and lack of information about prices (Kahn and Wolak 2013; Ito 2014; Jessoe and Rapson 2014; McRae and Meeks 2016) and lack of salience (Allcott 2011; Allcott et al. 2014) as factors that dampen the price elasticity of demand. Our set up is closely related to the misalignment between landlords and tenants (Levinson and Niemann 2004; Gillingham et al. 2012; Myers 2015; Elinder et al. 2017). A key difference is that aggregation of usage across many people is at the root of the intrahousehold incentive problem we study.

We contribute to the household economics literature by studying implications of intrahousehold decision-making for a novel domain of consumption, namely environmental-externality-generating utilities. We contribute to a small set of papers showing Pareto inefficiency in consumption as opposed to production (Dercon and Krishnan 2000; Dufflo and Udry 2004; Mazzocco 2007; Robinson 2012; Angelucci and Garlick 2016). We highlight hidden action, or limited information about consumption, as the source of inefficiency, unlike most previous work which explores limited commitment or hidden income.

Test of intrahousehold free-riding

We present a formal model of intrahousehold free-riding in water consumption in Appendix A.1, and briefly sketch the key prediction here.

Consider two spouses who must exert costly effort to conserve water. They split the household’s income net of the water bill according to some sharing rule. The spouse with more claim on the residual income (the “effective bill payer”) has more financial incentive to conserve, so puts in more conservation effort.

Now consider adding a price component that one spouse is fully responsible for paying. The model predicts that the spouse who is not the effective bill payer (the “non bill payer”) will be more responsive to this price. The intuition is that it represents a larger change in her incentives to conserve because she is not internalizing the standard household-level price as much. Another way to see this is that she is exerting less effort toward conservation under the status quo, so with a convex cost of conservation, she will increase effort more in response to a given individualized price increase.³

We thus test the following prediction: An individual-specific financial incentive for household conservation (which is akin to a price increase the individual fully bears) reduces water use more when directed at the spouse with a low residual claim on savings on the household water bill. If the prediction does not hold, it suggests that certain assumptions of the model, such as conservation effort being unobservable or households having non-unitary preferences, may be invalid.

Data and experimental design

Implementing our test of intrahousehold free-riding requires (1) data on household water use, (2) a measure of who in the household is the effective bill payer for water, and (3) person-level variation in the price of water. We partnered with the private regulated utility, Southern Water and Sewerage Company (SWSC), that provides piped water in Livingstone, Zambia to survey their customers and implement a randomized experiment that offered financial rewards for reducing water use.

³The curvature of the effort cost function is important for this prediction because the effective bill payer also benefits more from savings on the regular household bill that result from his effort. The prediction holds if the marginal cost of effort increases steeply enough in effort, e.g., with effort costs that are steeper than quadratic if utility is linear in income.

Study sample

Our full sample comprises 1,282 married couples who are SWSC customers. We selected the sample from the universe of SWSC’s metered residential accounts by imposing restrictions based on billing data and an in-person screening visit. We summarize the procedure here, with details provided in Appendix A.3.

Using billing data as of February 2015 (N=9,868), we eliminated households with a broken or unreliable meter; zero water consumption in more than half of the preceding four months; very low month-to-month variation in usage (indicative of meter tampering); low usage;⁴ extremely high usage (likely misclassified firms); a high outstanding balance with SWSC; or a high amount owed to them by SWSC. Applying these filters yielded 7,425 households that we targeted for in-person screening. We conducted screening visits and then full surveys on a rolling basis across neighborhoods between May and December 2015.⁵

A surveyor visited the short-listed households to screen on other study inclusion criteria: the water meter was not shared with other households; the household was headed by a married (or cohabiting) couple; both spouses lived at that address; the household resided at that address for at least the four months prior to April 2015; and they did not plan to move in the following six months. We screened 6,594 households, of which 2,051 met our inclusion criteria.

We scheduled a follow-up visit with 1,817 of the screened-in couples, explaining that both spouses needed to be present for the survey and they would be compensated 40 Kwacha (4 USD) for participation.⁶ We completed surveys with 1,282 of these households. The main reason for not surveying the remainder is that we ended fieldwork in December 2015 once we reached our target sample size.

Treatments were delivered (i.e., individuals or couples were told about the incentives to conserve) in conjunction with the surveys. The incentives then remained in place through February 2016.

Billing data

SWSC bills households monthly based on monthly in-person meter readings. We measure water consumption using this billing data.

⁴This ensures that we were not encouraging unhealthily low usage. The sample (urban, has piped water) is mostly middle class.

⁵We conducted the sampling in two waves because the budget-feasible sample size depended on our success rate and pace for completing surveying. The first wave used more stringent inclusion criteria.

⁶We report 2015 USD values using an exchange rate of 10 Kwacha per USD.

We create a panel of billing data from March 2014 to February 2016. The end date aligns with when we ended the incentive program, and the start date ensures at least a year of pre-period billing data and two calendar years of outcome data for all households. With a balanced panel in calendar time, the estimated treatment effect averages across households with different treatment duration; treated households have at least 2 and up to 9 months of treatment. We show that the results are similar if we instead use a panel balanced in event time, restricting to the first two months that incentives are in place, which is the minimum treatment duration. We also show robustness to alternative panel lengths.

Our main outcome is the log of household water use. The log transformation drops a small number of months with a reading of zero (which are likely billing errors or months the entire household was away, in any case). We drop months in which meter readings were estimated or the meter was reported as broken or disconnected. We control for an indicator for the month following a missing observation to account for the fact that the first reading after an estimated or missing reading might not map to the current month's consumption.

The average water price for our sample households is 5.1 Kwacha (0.51 USD) per cubic meter (m^3). Average household consumption is 19 m^3 per month, about half of typical US household consumption, resulting in monthly consumption charges of around 95 Kwacha (9.50 USD), or about 4 percent of median income.⁷

Survey data

A pair of surveyors (always a woman and a man) visited each screened-in household for the household survey. After a few preliminary demographic questions, husbands and wives were separated and surveyed in different rooms. After finishing their individual questionnaires, both surveyors and respondents reconvened in a common room for final questions.

The survey elicited the respondent's beliefs about the price of water, understanding of the water bill, view on which spouse uses more water, and demographic characteristics, among other information.

A key variable for us is who is the household's effective water bill payer (BP), that is, who is the residual claimant over changes in the water bill. We use two survey questions asked of each spouse to construct this measure: whose income is used to pay the bill and who physically pays the bill (payment is in person). For each question, the possible answers are oneself, one's spouse, both jointly, or someone else. We code the BP variable as follows

⁷We do not have income data for our sample, so calculate median income (220 USD/month) for households with piped water in Livingstone in the 2010 Living Conditions Monitoring Survey.

(illustrated for when the wife is respondent). If her response is herself for both questions, we code her as the BP. We also code her as BP if she either says she pays the bill but the income used is from both of them or someone else, or the income used is hers but both/someone else pays the bill. Less clear-cut is when she says she pays the bill but her husband’s income is used. We code the payer as the BP in these cases because a follow-up question asked of a subsample suggests that the payer usually has control over money left over after paying the bill. Finally, we code as 0.5 cases where for both questions, she says both of them or someone else play the role. Thus, at the respondent level, BP equals 0, .5, or 1 for each spouse, and the sum of BP across the two spouses equals 1 by construction. As there is subjectivity in how we code this variable, we show extensive robustness checks using alternative coding.

We similarly code BP according to the husband and then average the husband’s and wife’s BP variables to obtain the household measure.⁸ This yields four cases. (1) BP equals 1 for the husband and 0 for the wife, or vice versa. The spouses agree on who the BP is. (2) BP equals 0.75 for one spouse, and 0.25 for the other; this arises when one respondent identifies a specific person as the BP, while the other views the role as evenly split. (3) BP equals 0.5 because one member of the couple says the husband is BP and the other says the wife; here they strongly disagree. (4) BP equals 0.5 because both think the BP role is shared.

For the first two cases, there is within-couple variation in the BP variable. Thus, when we randomize the incentive to conserve at the individual level, we induce randomized variation in the recipient’s BP status. The variable’s values differ between the first two cases (0 and 1 versus 0.25 and 0.75), but note that there is no between-couple variation in the expected value of BP, which is, by construction, always 0.5. We exclude the last two cases from the analysis (222 and 8 households, respectively) because the incentive recipient has $BP = 0.5$ regardless of the randomization outcome. The results are similar when we add in these households.

Randomized “price” variation

Varying the regulated price charged by the utility was infeasible in our setting. Moreover, our key prediction is based on individual-level price variation. Thus, we manipulate the experienced water price through an intervention that increases the financial returns to water conservation.

Half of households were randomized into an incentive treatment. They were informed

⁸Appendix Table A.1 summarizes the wife’s and husband’s BP variables.

during their survey visit that they were being offered a monetary incentive to reduce water use.⁹ In the months following the visit, if the household reduced its consumption by a specified amount, they were entered into a monthly lottery that paid out 300 Kwacha (30 USD) prizes. Specifically, the household had to reduce its consumption by at least 30 percent relative to its average usage during a two-month reference window. The mean (median) reduction required to qualify was 5.8 (5.0) cubic meters.¹⁰ Those who qualified for the lottery in a given month had a 1 in 20 (or better) chance of winning the prize, so the expected prize was around 1.5 USD, which represents a roughly 40 percent increase in the price of water.¹¹

A fixed reward for reaching a consumption threshold differs from a standard price increase in that the effective price of water increases over a particular range of consumption and only if a usage threshold is not exceeded. This format of rewards was easy for participants to understand. In addition, the reward is an expected reward; we randomly select some households for payment to simplify the logistics and reduce the field costs of paying the prizes. These design decisions were guided by pragmatic considerations but still allow us to test our predictions.

The incentive treatment consists of three sub-treatment arms, each comprising one sixth of the total sample.¹² Appendix Figure A.2 summarizes the experimental design. In one sub-treatment, both spouses learn about the incentive, and know that the information is provided to both. In this case, the intervention generates incentives similar to an increase in the household's price. The other two sub-treatments inform only the wife or only the husband about the incentive. These treatments are similar to an increase in an individual-specific price; that is, the price of household consumption increases, but only one individual is aware of this and he or she fully bears the price increase. This raises the recipient's unilateral payoff from water conservation. The treated individual could, of course, share the information with his or her spouse, or the spouse might find out about it, but the individual-specific treatment comes closer to an individual price change than does the couple treatment. Individuals who won a prize were informed and paid privately in the latter two sub-treatments.

We test our main prediction by comparing the two individual incentive sub-treatments.

⁹The script is provided in Appendix A.4.

¹⁰The reference period was March-April 2015 for households surveyed in May-early August; June-July for households surveyed early August-September, and July-August for households surveyed in October-December.

¹¹Given a target reduction of 5.8 m³ and a price of 5.1 Kwacha/m³ over that range of consumption, the expected value of the prize was 2.06 Kwacha/m³. This calculation accounts for the increasing block tariff.

¹²The randomization was within four strata based on the household's pre-period monthly water usage and outstanding balance due to SWSC.

The control group is helpful for being able to gauge the absolute magnitude of the effect and calculate a price elasticity within our sample. We include the couple incentive arm as a benchmark because it is the most similar to standard household-level pricing used by utilities.¹³

Other interventions

We also varied two other factors that might affect water use. First, water is priced on an increasing block tariff (i.e., the marginal price increases discretely at certain thresholds of usage), which results in a poor understanding of the marginal price. All households that receive the incentive treatment also receive information about the actual price of water. In addition, a subsample of the households that received no incentive to conserve were given the information about the price. This intervention was intended to serve as an additional source of (perceived) price variation and to homogenize price beliefs among those receiving incentives. Second, distrust of the water provider or a misunderstanding of the billing process might undermine customers’ belief that their water use directly maps into their bill. We implemented a cross-cutting “provider credibility” treatment that explains how bills are generated. Neither the price information nor the credibility treatment had measurable impacts on water use, even when prior beliefs about the price or provider are taken into consideration. Details of these interventions and analysis of their impacts are presented in the appendix.

Summary statistics

Table 1 summarizes characteristics of the sample and tests for balance between treatment arms. The first column reports the mean and standard deviation of several variables for the subsample that received no incentive. Average water consumption prior to the start of the intervention is 19 cubic meters per month. Household size is 6 people, living in 3.5 rooms. (To illustrate the importance of intrahousehold frictions, our study focuses on husband-wife dynamics, but as the household size underscores, intrahousehold decision-making is often more complex.)

In 49 percent of households, both husband and wife agree that the husband is the effective bill payer. In 6 percent of households, they “weakly agree” that the husband is the effective

¹³A previous, longer version of the paper included a second test of intrahousehold free-riding derived from the model that also relied on the couple incentive: Households in which spouses are more altruistic toward one another will be more responsive to household-level pricing.

bill payer, meaning that one spouse says that the husband is the bill payer and the other says the responsibility is shared. In less than 1 percent of households, they agree that the responsibility is shared. In 20 percent of households they “strongly disagree,” meaning that one spouse says that the husband is the bill payer and the other says that the wife is. In the remaining 24 percent of households, the spouses either agree or weakly agree that the woman is the effective bill payer. The table also shows that in 80 percent of households, spouses agree that the woman is the bigger water user.

Subsequent columns of Table 1 report regression coefficients and standard errors that assess the difference between a subsample and its comparison group. Column 2 compares the subsample in which the couple received an incentive to the no-incentive group. Column 3 does the same for households that received an individual-level incentive. Our main test zeros in on the individual incentive arm, so columns 4 and 5 break down this group. Column 4 shows the subsample where the woman received the incentive (gender was the basis of the randomization); the relevant comparison is to households where men received the incentive. Finally, column 5 shows the subsample where the non bill payer received the incentive, with the comparison group being households where the bill payer received it.¹⁴ F-tests indicate that, in all cases, we cannot reject balance between a subsample and its comparison group.¹⁵

Estimation strategy

We use the randomized individual incentive as a person-specific price increase to test the following prediction: The individual incentive is more effective in reducing household water use if it is offered to the spouse who is not the effective bill payer.

Our main analysis focuses on the 870 households in the individual incentive arm or control group for whom there is within-couple variation in bill payer status. In robustness checks, we add in the 182 households that received the couple incentive and the 230 households with no within-couple variation in BP status.

We use monthly household-level outcome data from before and after the intervention and estimate a difference-in-differences regression. The regressor $IndivTreat_{it}$ equals 1 for households assigned to the individual treatment in months after the survey; treatments were

¹⁴The bill payer variable is not binary. For this table, we pool individuals with bill payer status of 1 or 0.75 as bill payers, and those with status of 0 or 0.25 as non bill payers. We omit households where bill payer status has no within-couple variation.

¹⁵Appendix Figure A.3 shows that there are parallel pre-trends between subsamples in average water use in the months leading up the survey.

delivered at the end of the survey visit.¹⁶

The main hypothesis is that the individual treatment will be more effective if targeted to the non-BP (where $NonBP \equiv 1 - BP$.) Thus, the key regressor is $IndivTreat_{it} \times NonBP_{it}$. This is not a standard interaction in that $NonBP$ (the incentive being given to the non-BP) is only defined within the incentive group.

In the estimating equation below, β_1 identifies the effect of the bill payer receiving the individual-level incentive relative to the control group, and β_2 measures the additional effect of the incentive going to the non-BP. The main hypothesis is, thus, $\beta_2 < 0$.

$$y_{it} = \alpha + \beta_1 IndivTreat_{it} + \beta_2 IndivTreat_{it} \times NonBP_{it} + \gamma_1 Post_{it} + \gamma_2 Wave_i \times Month_t + \tau_t + \eta_i + \gamma_3 MissingFlag_{it} + \epsilon_{it} \quad (1)$$

The other regressors are control variables. $Post_{it}$ is a post-survey indicator, as the survey itself may have had effects; it varies across households within month because the survey was rolled out over time. We drew our sample in two waves, and the subsamples exhibit different time trends, so we conclude the interaction of $Month_t$, a continuous month-year variable, and $Wave_i$, a sampling-wave indicator, to absorb additional residual variation. We also include month-year fixed effects, τ_t , and household fixed effects, η_i . $MissingFlag$ is an indicator for months immediately following a missing observation. We cluster standard errors at the household level.

Results

The individual incentive reduced water consumption across most of the distribution, as shown in Figure 1. The figure plots post-treatment water consumption, normalized by the household’s average consumption in the two pre-survey reference months, separately for the control group and the pooled individual incentive group.

For both the treatment and control groups, there is not much mass below the target level to be eligible for the prize: Most of the treatment effect is due to reductions not large enough to qualify for the prize. If households could perfectly choose their consumption level, we would expect bunching just below the target among treated households. However, the

¹⁶We set $IndivTreat_{it}$ equal to 1 as of the survey date because the intervention could have immediate effects; the household was only eligible for a prize based on the next full bill cycle, however. We drop the month in which the survey occurred since it is partially treated. Our definition of month corresponds to the billing cycle, which starts on the 20th of each month.

difficulty of knowing one’s own and others’ water use makes the pattern less surprising. The continuity in the reductions suggests that households responded to the lumpy financial incentive similarly to how we would expect them to respond to a standard price increase.

The regression version of the comparison in Figure 1 finds that the individual treatment significantly reduced water use, by 0.059 log points (see Appendix Table A.2).

Recall that all households that received the incentive also received information on the price of water, as did a subset of the no-incentive group. Appendix Table A.2 also estimates the effects of these treatments, and shows that the individual incentive treatment effect is largely unaffected. In the rest of the paper we pool all of the no-incentive households, both pure controls and those that received price information. This increases statistical power when we estimate the overall effect of the individual incentive treatment; it does not affect the identification our main coefficient of interest ($IndivTreat \times NonBP$), which does not rely on the control group. We also ignore the cross-cutting credibility treatment. In other words, we impose the restriction, which we cannot empirically reject, that these other interventions have zero effect.¹⁷

The main result of the paper is shown in Table 2: The individual incentive causes a larger reduction in household water consumption when the recipient has lower bill payer status. Column 1, which estimates equation (1), shows that when the incentive goes to the BP, there is an estimated 0.03 log point reduction in water use, an effect statistically indistinguishable from zero. When the incentive goes to the non-BP, there is a significantly larger effect. The point estimates of -0.109 log points implies a total effect of -0.14 log points, equivalent to a short-run price elasticity of about -0.26, when the non-BP receives the incentive. (Appendix A.3 provides the details of this elasticity calculation, which has many caveats.)

When the bill payer received the incentive, in the absence of intrahousehold frictions, he should have been able to reproduce the effects of the non bill payer incentive arm: He could tell his spouse about the rewards program and promise her almost all of the prize. Our results are suggestive that bill payers may not have thought to do this, or a commitment problem prevented it from being effective. We return to discussing this puzzle in the conclusion.

One might expect the effects to be strongest in the first few months that the incentives are in place due to greater salience. The bottom panel of Appendix Figure A.3, which plots the month-by-month coefficients corresponding to the specification in Table 2, column 1, shows that the effect is negative in the first three months, and then bounces around in subsequent months. (The coefficients beyond the first two months are less precise because

¹⁷Appendix Table A.3 shows that the price information and provider credibility treatments had no detectable impacts, even after allowing for heterogeneity based on priors about the price or about SWSC.

fewer households contribute to their estimation.)

Columns 2 and 3 of Table 2 disentangle whether the effect is due to bill payer status or gender. These two variables are correlated; in most cases, the husband has higher bill payer status. Column 2 examines heterogeneity by gender instead of BP status. There is no significant differential effect of the individual incentive when the wife receives it, though the point estimate is negative. Column 3 simultaneously estimates the effects of the incentive being given to the wife and to the non-BP. In other words, we estimate the effect of targeting the non-BP, controlling for gender, to determine if the former effect is driven entirely by gender. It is not: the interaction with effective bill payer status remains significant and similar in magnitude to column 1, while recipient gender per se does not seem to affect responsiveness to the incentive.

To summarize, the existing household arrangement regarding who has claim on savings from water conservation is an important determinant of the effectiveness of the incentive treatment.

Robustness checks

One potential concern is that BP status is correlated with other individual characteristics, so we are not measuring heterogeneity by BP status per se. We can partially address this by controlling for individual-level observables such as employment status in parallel to BP status. In Table 3, we include interactions of *IndivTreat* with other characteristics, first one at a time (column 1) and then all at once (column 2). While we do see some heterogeneity in price sensitivity by these other measures, our main coefficient of interest (*Incentive* \times *Non-BP*) is stable in magnitude and significance in column 2.

Table 4 tests for the sensitivity of our main result (Table 2, column 1) to how we define bill payer status. In column 1, we switch to putting precedence on the “whose income” variable instead of “who pays” when those two variables disagree. In column 2, we drop the cases where the those two underlying variables disagree. With both of these variations, we continue to find that the incentive has a larger effect when given to the non-BP. Column 3 restricts the sample to households where the couple strongly agrees on who the BP is and we find similar results. Finally, in columns 4 and 5, instead of using the average of the husband’s and wife’s BP assessments, we use just one respondent’s answers. Results are similar to our main results.

Appendix Table A.4 brings in the households where BP equals 0.5 for both spouses, which occurs if they each specify different people as BP or if they both think the role is shared.

When we include these households (and this non-randomized variation), the key coefficient changes only beyond the third decimal place.

Appendix Table A.4 also shows the results adding in the couple incentive arm. The point estimates suggest that the individual incentive to the non-BP reduces water use by over twice as much as the couple incentive, but this difference is statistically insignificant.

Finally, we test for sensitivity of our results to how we construct the panel (Appendix Table A.5). The first three columns show different pre-treatment panel lengths. The next three columns include only two post-survey observations per household, to ensure that treated households contribute equally to the estimated treated effect; the somewhat larger point estimate is consistent with the slight decay in the effect size after the third month seen in Appendix Figure A.3. Finally, the last column shows the results using a panel balanced in event time rather than calendar time, with 14 pre-treatment months and 2 post-treatment months. The point estimate is similar to our main specification.

Conclusion

This paper highlights how intrahousehold free-riding exacerbates households' overconsumption of piped water and electricity. These utilities have the features that usage is billed to the household, and household members cannot easily observe each individual's consumption. Thus, they cannot apportion the bill based on each individual's consumption. In the face of this free-riding problem, targeting an individual-level price increase to the household member who normally has the least incentive to conserve water should — and does — lead to a larger response.

This moral hazard problem between spouses would exist even if men and women were perfect equals, but the problem is exacerbated by traditional gender roles, with women doing most of the chores and men controlling the money. This husband-wife power imbalance might be more common in developing countries (Jayachandran 2015). However, other forms of intrahousehold free-riding — for example, children wasting water and energy — are likely equally applicable in rich and poor countries.

Limited information on individual-level consumption is a fundamental constraint for households, but why do they seem to compound the problem by assigning bill responsibility to the man, or more precisely, the smaller water user? Why is the bigger water user the effective bill payer for only one third of households? Traditional gender roles is not a fully satisfactory answer because many husbands give their wives an allowance for groceries.

In follow-up discussions with 40 households, most stated that using a similar arrangement for water had never occurred to them. One conjecture is that “optimal” intrahousehold contracting norms emerge slowly, while piped water is a new phenomenon. When women fetched water from rivers or springs, they were the “effective bill payers”; wasting water meant more time spent fetching water.

Even if households improve how they share the bill, limited information about usage will still lead to over-consumption. One policy lever to reduce water use is corrective pricing, i.e., a tax, which would now need to correct both the environmental externality and the intrahousehold “internality” (Allcott et al. 2014). In an extension to this paper (Jack et al. 2018), we calculated the optimal tax on water to correct for both intrahousehold free-riding and the environmental externality, adapting the framework of Taubinsky and Rees-Jones (2018). The key take-away is that if the internality problem varies substantially across households, as it does in our context, then corrective pricing is a highly imperfect instrument to fix it.

The more promising solution is to design policies based on the specific intrahousehold constraints. Individual-level pricing was a useful way to test our predictions, but may not be viable to scale up. That said, a potentially scalable analog to our experimental variation is a rewards program for conservation that uses demographically-targeted in-kind rewards (e.g., gift cards especially valued by women). Another tack is to reduce information frictions. For example, giving households better information about household-level usage through smartphone apps with real-time data would enable better monitoring; detailed information about household use is a first step toward backing out each person’s use. In addition, technologies that lower the effort cost of conservation (e.g., automatic shut-offs for faucets or lights), might be especially valuable in the face of intrahousehold moral hazard.

References

- Allcott, H. (2011). Consumers' perceptions and misperceptions of energy costs. *American Economic Review* 101(3), 98–104.
- Allcott, H., S. Mullainathan, and D. Taubinsky (2014). Energy policy with externalities and internalities. *Journal of Public Economics* 112, 72–88.
- Angelucci, M. and R. Garlick (2016). Heterogeneity in the efficiency of intrahousehold resource allocation: Empirical evidence and implications for investment in children. *Working Paper*.
- Dalhuisen, J. M., R. J. Florax, H. L. De Groot, and P. Nijkamp (2003). Price and income elasticities of residential water demand: A meta-analysis. *Land Economics* 79(2), 292–308.
- Dercon, S. and P. Krishnan (2000). In sickness and in health: Risk sharing within households in rural Ethiopia. *Journal of Political Economy* 108(4), 688–727.
- Duflo, E. and C. Udry (2004). Intrahousehold resource allocation in Cote d'Ivoire: Social norms, separate accounts and consumption choices. *National Bureau of Economic Research Working Paper*.
- Elinder, M., S. Escobar, and I. Petré (2017). Consequences of a price incentive on free riding and electric energy consumption. *Proceedings of the National Academy of Sciences* 114(12), 3091–3096.
- Gillingham, K., M. Harding, and D. Rapson (2012). Split incentives in residential energy consumption. *The Energy Journal* 33(2), 37–62.
- Ito, K. (2014). Do consumers respond to marginal or average price? Evidence from nonlinear electricity pricing. *American Economic Review* 104(2), 537–563.
- Jack, K., S. Jayachandran, and S. Rao (2018). Environmental externalities and free-riding in the household. *National Bureau of Economic Research Working Paper*.
- Jayachandran, S. (2015). The roots of gender inequality in developing countries. *Annual Review of Economics* 7(1), 63–88.
- Jessoe, K. and D. Rapson (2014). Knowledge is (less) power: Experimental evidence from residential energy use. *American Economic Review* 104(4), 1417–1438.

- Kahn, M. E. and F. A. Wolak (2013). Using information to improve the effectiveness of nonlinear pricing: Evidence from a field experiment. *Working Paper*.
- Levinson, A. and S. Niemann (2004). Energy use by apartment tenants when landlords pay for utilities. *Resource and Energy Economics* 26(1), 51–75.
- Mazzocco, M. (2007). Household intertemporal behaviour: A collective characterization and a test of commitment. *Review of Economic Studies* 74(3), 857–895.
- McRae, S. and R. Meeks (2016). Price perception and electricity demand with nonlinear tariffs. *Working Paper*.
- Myers, E. (2015). Asymmetric information in residential rental markets: Implications for the energy efficiency gap. *Working Paper* (246R).
- National Water Supply and Sanitation Council, Zambia (2015). Strategic plan: 2016-2020. Technical report, NWASCO.
- Robinson, J. (2012). Limited insurance within the household: Evidence from a field experiment in Kenya. *American Economic Journal: Applied Economics* 4(4), 140–164.
- Taubinsky, D. and A. Rees-Jones (2018). Attention variation and welfare: Theory and evidence from a tax salience experiment. *Review of Economic Studies* 85, 2462–2496.
- Worthington, A. C. and M. Hoffman (2008). An empirical survey of residential water demand modelling. *Journal of Economic Surveys* 22(5), 842–871.

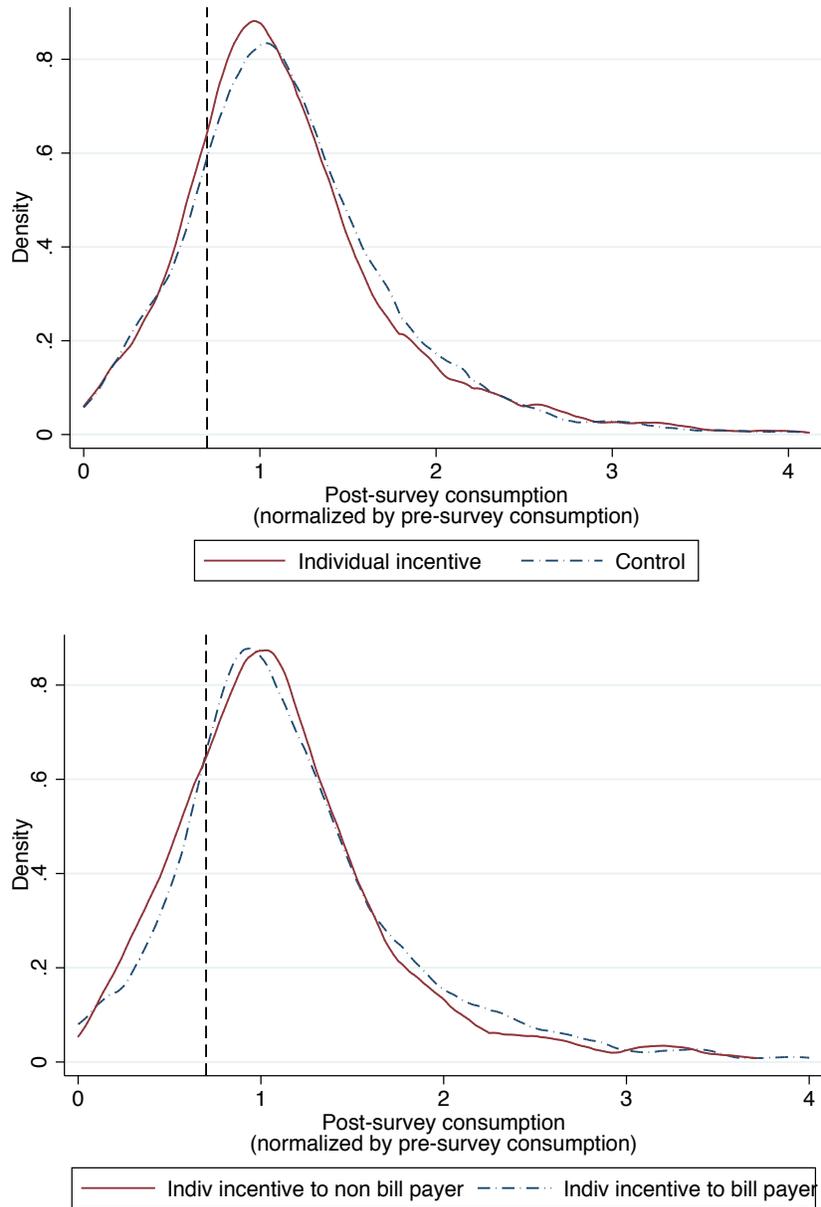


Figure 1: Post-intervention water consumption, relative to pre-intervention

Notes: Density plots of post-intervention monthly consumption relative to average monthly consumption in the reference months (pre-survey) used to determine incentive treatment eligibility. The dashed vertical line shows the 70 percent threshold for lottery eligibility. The control group includes all households not assigned to an incentive arm.

Table 1: Balance: Incentive treatment arms

Treatment arm:	No incentive (1)	Couple incentive (2)	Individual incentive (3)	Wife incentive (4)	Non BP incentive (5)
Quantity of water consumed	18.844 (11.922)	-0.001 (0.002)	-0.002 (0.002)	0.005 (0.003)	0.004 (0.003)
Household size	5.860 (2.286)	0.005 (0.006)	-0.000 (0.007)	0.003 (0.011)	-0.005 (0.012)
HH has maid	0.169 (0.375)	-0.015 (0.041)	-0.032 (0.043)	0.076 (0.072)	-0.096 (0.082)
HH owns home	0.512 (0.500)	-0.012 (0.030)	-0.015 (0.031)	-0.030 (0.050)	-0.027 (0.057)
Rooms in home	3.529 (1.264)	0.002 (0.013)	0.016 (0.013)	-0.031 (0.019)	0.008 (0.024)
Both know bill quantity	0.104 (0.305)	0.134 (0.045)	0.065 (0.048)	-0.022 (0.076)	0.085 (0.086)
Both know bill charge	0.678 (0.468)	0.002 (0.031)	0.024 (0.032)	-0.068 (0.053)	-0.073 (0.059)
Agree W is bigger water user	0.795 (0.404)	0.022 (0.036)	0.095 (0.039)	-0.044 (0.068)	0.050 (0.073)
Agree H is bill payer	0.530 (0.499)	0.021 (0.038)	0.029 (0.039)	-0.064 (0.063)	0.074 (0.064)
Weakly agree H is bill payer	0.051 (0.221)	0.099 (0.068)	0.025 (0.073)	0.002 (0.118)	0.150 (0.120)
Agree bill payer role is shared	0.009 (0.095)	-0.162 (0.171)	-0.117 (0.177)	-0.599 (0.359)	0.000 (0.000)
Strongly disagree on bill payer	0.206 (0.405)	0.061 (0.044)	-0.002 (0.047)	-0.212 (0.077)	0.000 (0.000)
F-statistic		1.328	1.070	1.749	0.783
Comparison group		No incentive	No incentive	Husband incentive	BP incentive
Households in treatment	664	182	436	213	173
Households in comparison		664	664	223	176

Notes: Column 1 reports means and standard deviations of time-invariant household characteristics preceding the intervention in the no incentive arm. Columns 2-5 show output from regressing an indicator for treatment status (column headers) on covariates. The F-statistic associated with the regression is reported at the bottom of the table. *Non BP* varies from 0 to 1; it equals 1 if the individual incentive goes to the person whom both spouses agree is not the effective bill payer. Households with no within-couple variation in bill payer status are excluded from column 5.

Table 2: Individual price incentive effects, by recipient payer status

	log(Quantity)		
	(1)	(2)	(3)
Individual incentive	-0.029 [0.037]	-0.056 [0.037]	-0.031 [0.040]
Indiv incentive x Non BP	-0.109** [0.047]		-0.113** [0.053]
Incentive x Wife		-0.055 [0.046]	0.006 [0.052]
Observations (HH)	870	870	870
Observations (HH-months)	17,253	17,253	17,253

Notes: Individual incentive treatment arms interacted with heterogeneity variables: *Non BP* varies from 0 to 1; it equals 1 if the individual incentive goes to the person whom both spouses agree is not the effective bill payer. *Wife* equals 1 if the individual incentive goes to the wife. The couple incentive treatment is excluded, as are households with no within-couple variation in bill payer status. The omitted category is the no incentive control group. The panel begins in March 2014 and ends in February 2016. Standard errors are clustered at the household level, and all columns control for household and month-year fixed effects, an indicator for months following a missing quantity observation, and a continuous month-year variable interacted with sampling wave.

Table 3: Robustness to controlling for individual characteristics

	log (Quantity) (1)	log (Quantity) (2)
Indiv incentive x Non BP	-0.109** (0.047)	-0.110** (0.056)
Indiv incentive x Over 50	0.099** (0.044)	0.085 (0.058)
Indiv incentive x Has regular employment	0.022 (0.053)	0.065 (0.056)
Indiv incentive x Fluent in English	0.020 (0.095)	0.013 (0.102)
Indiv incentive x Low education	-0.013 (0.050)	-0.017 (0.058)
Indiv incentive x Uses more water	-0.014 (0.046)	0.053 (0.053)
Indiv incentive x Knows bill quantity	0.022 (0.048)	-0.009 (0.060)
Indiv incentive x Knows bill price	0.001 (0.050)	-0.005 (0.064)
Observations (HH)	870	866
Observations (HH-months)	17,253	17,174

Notes: Robustness check on the results reported in column 1 of Table 2. *Indiv incentive* refers to the individual incentive arm. Each coefficient is an interaction between *Indiv incentive* and a characteristic of the recipient. Column 1 shows separate regressions in each cell. Column 2 reports results of a single regression. Regressions include the post-survey indicator interacted with the heterogeneity variables. The couple incentive treatment is excluded, as are households with no within-couple variation in bill payer status. The panel begins in March 2014 and ends in February 2016. Standard errors are clustered at the household level, and all columns control for household and month-year fixed effects, an indicator for months following a missing quantity observation, and a continuous month-year variable interacted with sampling wave.

Table 4: Robustness to different ways of defining non bill payer variable

	log(Quantity)				
	(1)	(2)	(3)	(4)	(5)
Individual incentive	-0.028 [0.036]	-0.006 [0.039]	-0.020 [0.038]	-0.027 [0.037]	-0.034 [0.038]
Indiv incentive x Non BP	-0.114** [0.048]	-0.136*** [0.051]	-0.126** [0.052]	-0.111** [0.046]	-0.101** [0.047]
Specification notes	Income var	Drop income \neq payer	Drop disagree	Husband defn	Wife defn
Observations (HH)	870	816	725	870	870
Observations (HH-months)	17,253	16,237	14,422	17,253	17,253

Notes: *Incentive* \times *Non BP* is the product of the someone in the household having received the individual lottery and (1 minus) the BP status of the individual. Columns vary how the bill payer variable is constructed relative to our main specification. Column 1 uses the income variable if income and payer disagree. Column 2 drops cases where income and payer disagree. Column 3 drops cases where either income or payer are both/other for at least one of the individual. Column 4 uses the husband's definition of bill payer. Column 5 uses the wife's definition of bill payer. The couple incentive treatment is excluded, as are households with no within-couple variation in bill payer status. The panel begins in March 2014 and ends in February 2016. Standard errors are clustered at the household level, and all columns control for household and month-year fixed effects, an indicator for months following a missing quantity observation, and a continuous month-year variable interacted with sampling wave.

Online Appendices

A.1 Model of intrahousehold free-riding in water use

We model a household’s water consumption as a function of effort spent on conservation. We start by benchmarking the household’s water use in the absence of any intrahousehold frictions. We then allow for individual-level water conservation choices that diverge from the household’s first best. Two features of water use guide our modeling decisions. First, there is limited observability of others’, and to an extent one’s own, conservation effort. Second, water is not purchased at the individual level; a utility bill for piped water pools all household members’ usage. We discuss these features of water in more detail at the end of this section. Because of these features, we model water use as a non-cooperative game. In the literature, households are more often modeled in a cooperative framework, befitting the altruism and long-term relationship among family members. Our model setup should not be interpreted as implying households are not cooperative over other domains characterized by greater observability of actions or individual-level purchases.

Our model is, in essence, a moral hazard in teams model, and similarly generates a free-riding problem, with each individual exerting inefficiently low effort to conserve water. Within this model set-up, we generate predictions about price sensitivity. We model a household as consisting of two individuals, whom we describe as husband and wife, but the intuition extends to other household structures.

A.1.1 Model setup and household optimum

Household aggregate water use, W , is the sum of water use by each individual i within the home, given by $w_i = \bar{w}(1 - e_i)$. Conservation effort $e_i \in [0, 1]$ lowers water use but at a convex cost, ce_i^μ , where $\mu > 2$. Individuals consume a maximum quantity of water given by \bar{w} if they exert no effort at all towards conserving water.¹ The water utility charges the household pW , where $W \equiv \sum_i w_i$. The household has total income Y , and we assume $p \sum_i \bar{w} < Y$. We model utility as quasi-linear in the income remaining after the water bill is paid. Given the convex conservation costs, utility is, thus, concave in water consumption and linear in other consumption.

We model a household as comprising two individuals, a husband and a wife. Assuming

¹The maximum level can be thought of either as the level of consumption where marginal benefits are equal to zero (i.e., a satiation point) or some physical constraint on water use associated with, for example, running all of the household’s taps for 24 hours a day.

equal welfare weights on each person's utility, the household's optimal choice of conservation effort is symmetric across individuals and is given by:

$$\max_{e_i} Y - 2p\bar{w}(1 - e_i) - 2ce_i^\mu. \quad (2)$$

Solving the first order condition, the household achieves its first best outcome if each member exerts effort, $e_i^{FB} = \left(\frac{1}{\mu} \frac{p\bar{w}}{c}\right)^{\frac{1}{\mu-1}}$.

A.1.2 Individual best response

The first best equilibrium might not be obtained, however, if the conservation effort of the other member of the household, $-i$, is difficult to observe. We assume that each individual i takes her spouse's conservation effort e_{-i} as given, assuming that e_{-i} is difficult to observe and therefore to contract over.

A sharing rule, $\lambda_i \geq 0$, determines the ex post division of income that remains after the household pays the water bill. (In practice, households might have different sharing rules for different expenses. What specifically is relevant is residual claim on the water bill, or the sharing rule that applies to the savings that accrue from water conservation.) Across spouses, the sharing parameters sum to 1 ($\lambda_i + \lambda_{-i} = 1$), and aggregate water use is given by $W = w_i + w_{-i} = 2\bar{w}(1 - \frac{e_i + e_{-i}}{2})$.

Individual i receives utility from income available for non-water consumption and disutility from water conservation effort:

$$v_i = \lambda_i(Y - pW) - ce_i^\mu.$$

Individuals may also internalize some share $0 \leq \alpha_i \leq 1$ of their spouse's utility, with i 's utility function given by $u_i = v_i + \alpha_i v_{-i}$. We model α_i as a measure of i 's altruism toward his or her spouse, but a spouse might also internalize how her water use affects her spouse's income due to enforcement of household agreements around water use (if individual water use is observable). Thus, α_i can be thought of in broader terms as a (downward) shifter of i 's propensity to free-ride.

Person i chooses e_i to satisfy the first order condition:

$$e_i^* = \left(\frac{1}{\mu} \frac{p\bar{w}}{c} (\lambda_i + \alpha_i(1 - \lambda_i))\right)^{\frac{1}{\mu-1}}$$

or, equivalently,

$$w_i^* = \bar{w} \left[1 - \left(\frac{1}{\mu} \frac{p\bar{w}}{c} (\lambda_i + \alpha_i(1 - \lambda_i)) \right)^{\frac{1}{\mu-1}} \right]. \quad (3)$$

For $\lambda_i = 1$ or $\alpha_i = 1$, person i fully internalizes the household's cost of water consumption, and the individual conservation decision is equal to the decision in the first best: $e_i^* = e_i^{FB} = \left(\frac{1}{\mu} \frac{p\bar{w}}{c} \right)^{\frac{1}{\mu-1}}$. However, if $\lambda_i = 1$, then $\lambda_{-i} = 0$, and individual $-i$ only exerts effort insofar as she is altruistic toward her spouse.

More generally, equation (3) shows that w_i^* is decreasing in p , λ_i , and α_i . A higher price, enjoying the monetary upside of lower water bills, and more altruism toward's one spouse all lead to lower water consumption.²

Our empirical focus is on how λ_i affects price sensitivity. Because $\lambda_i + \lambda_{-i} = 1$, there is no cross-household variation in the average value of λ to identify how existing incentives within the household affect individual (and in turn household) water use.

To measure the effect of λ_i on price sensitivity, we add an individual-specific component to the price, denoted P_i . The individual utility function, with the introduction of individual-specific pricing, then becomes $v_i' = \lambda_i(Y - pW) - ce_i^\mu - P_iW$ and her optimal effort is now $e_i'^* = \left[\frac{1}{\mu} \left(\frac{p\bar{w}}{c} (\lambda_i + \alpha_i(1 - \lambda_i)) + \frac{P_i\bar{w}}{c} \right) \right]^{\frac{1}{\mu-1}}$.

A.1.3 Effect of a price change

Our experimental treatments make water use more costly, effectively increasing the price of water, and, in our data, we observe household-level water use, W . Here we derive comparative statics for how the price sensitivity of household water use, $\frac{\partial W}{\partial p}$ and $\frac{\partial W}{\partial P_i}$, depend on the inner workings of the household. We summarize the results in this section; the proofs are presented in Section A.1.5. Note that the first derivative of water consumption with respect to price is negative, so a negative cross-derivative represents an increase in price sensitivity.

Prediction 1: $\frac{\partial^2 W}{\partial \lambda_i \partial P_i} > 0$, or equivalently, $\left| \frac{\partial W}{\partial P_i} \right|$ is decreasing in λ_i . In words, the individual who is not the effective bill payer (lower λ in the household) is more responsive to changes in the individual-level price.

²The theoretical predictions characterize the marginal change in water use with respect to a marginal price change, but they also hold for a discrete price change associated with a threshold quantity change. Similarly, here we derive predictions for water use in levels, while our empirical results test for effects on log water use; rewriting the model in logs generates the same predictions.

The intuition for the result is that the lower λ_i individual has weak incentives to conserve based on the household-level price, so she is exerting less effort toward conservation, and thus faces a lower marginal cost of effort given the convexity of the cost of effort.³ Since $W = w_i + w_{-i}$ and w_{-i} is unaffected by a change in P_i (we assume changes in P_i are not observed by $-i$), the change in household water use W is identical to the change in w_i . Directing the individual price P_i to the individual who is not the effective bill payer (lower λ_i) will have a larger effect on aggregate consumption.

Prediction 2: $\frac{\partial^2 w_i^*}{\partial p \partial \alpha_i} < 0$, or equivalently, $\left| \frac{\partial w_i}{\partial p} \right|$ is increasing in α_i . In words, individuals who are more altruistic toward their spouse are more sensitive to the household-level price.

The intuition is that i will reduce her water use more in response to a price increase the more she internalizes the savings that will accrue to her spouse. We observe household-level, not individual-level water usage, so the testable prediction is that households in which members have higher average levels of altruism toward one another are more sensitive to household-level prices.

Note that the individual-level effect described above will be stronger as λ_i decreases: $\left| \frac{\partial^2 w_i^*}{\partial p \partial \alpha_i} \right|$ is decreasing in λ_i . The marginal effect of α_i being larger when λ_i is low means that altruism matters more for water use when the other person pays the bill.

A.1.4 Discussion of assumptions

What makes water (and electricity) special A key feature of water consumption implicit in our setup is that the household — not the individual — pays for water. Household utilities such as water or electricity tend to have this feature in contrast with, for example, clothing, where a couple could divide up income and make individual purchases. This point is distinct from saying water is a public good; (some) water consumption is rival and excludable (e.g., drinking a glass of water) but purchases are not made individually.

There are also goods such as food for which households could choose to make individual purchases but do not typically do so; this seems natural for ingredients used to prepare shared meals, but some food consumption, such as snack food, is more often individual consumption.

³The curvature of the cost of effort function is important for this prediction. A higher λ_i individual starts at a higher effort level, so faces a higher marginal cost of effort, but also benefits more from the savings on the household-level water bill that result from his conservation effort. The marginal cost of additional effort must be increasing in the level of effort steeply enough for the first effect to dominate ($\mu > 2$). In contrast, Prediction 2 does not require $\mu > 2$.

The fact that households could but do not purchase snack food separately raises the other key feature of water assumed in this setup: lack of observability of individual consumption. A spouse’s water use is difficult to observe. First, it is hard to match water quantities to activities (e.g., how many gallons used in a 5 minute shower, how many gallons used to wash dishes). Second, feedback on consumption is infrequent since it typically arrives once a month with the water bill. This compounds the observability problem. Contrast this with snack food, where the household has more information to assign consumption to each individual: if you notice that the number of cookies in the cookie jar has decreased since the last time you were in the kitchen, you know one of your family members stole a cookie from the cookie jar. If water meters were more accessible and easier to interpret, an individual could check the meter before and after a spouse’s shower to observe consumption.⁴

Adding to these observability challenges, knowing one’s own consumption is often difficult.⁵ Even ex post, if i can only observe her own consumption with some error ϵ , then she can only infer w_{-i} from the total bill with error: $w_{-i} = W - (w_i + \epsilon)$. Moreover, the fact that some part of water consumption is a public good at the household level (e.g., washing the family’s dinner dishes) further complicates the problem of quantifying others’ effort toward conservation. (Note that even when water is used to produce public goods, there is still some “private” consumption if, conditional on how clean you get the dishes, washing them in a manner that wastes less water requires more effort and hence higher private costs.) Of course, other consumption goods within the home may be susceptible to one or more of these challenges, though qualitative survey data is consistent with worse observability for water and electricity than other common categories of consumption (see Appendix Figure A.1).

A.1.5 Derivations and proofs

We show the derivation of the optimal e_i^* and w_i^* and then prove our two results. Recall that p is the household level price, P_i is an individual price, α_i is the weight i puts on $-i$ ’s utility, and λ_i is i ’s share of the household income net of the water bill.

⁴This improvement in intrahousehold observability may explain part of the decline in electricity use associated with the introduction of smart metering (e.g., Jessoe and Rapson 2014).

⁵The fact that even one’s own consumption is difficult to gauge means that, even leaving aside the free-riding problem within a group, an individual might not consume the amount of water she is targeting. For example, if there were a prize for reducing water, a person living alone might unintentionally miss the target. This problem of only being able to choose consumption with error is a distinct one from the free-riding problem we are focused on, and could lead to over- or under-consumption of water.

Individual utility for person i is given by

$$u_i = \lambda_i(Y - pW) - ce_i^\mu + \alpha_i(1 - \lambda_i)(Y - pW) - \alpha_i ce_{-i}^\mu - P_i W$$

or, substituting in $W = 2\bar{w}(1 - \frac{e_i + e_{-i}}{2})$,

$$u_i = \lambda_i(Y - 2p\bar{w}(1 - \frac{e_i + e_{-i}}{2})) - ce_i^\mu + \alpha_i(1 - \lambda_i)(Y - 2p\bar{w}(1 - \frac{e_i + e_{-i}}{2})) - \alpha_i ce_{-i}^\mu - 2P_i\bar{w}(1 - \frac{e_i + e_{-i}}{2}).$$

Solving the first-order condition with respect to e_i gives:

$$e_i^* = \left[\frac{1}{\mu} \left(\frac{p\bar{w}}{c} (\lambda_i + \alpha_i(1 - \lambda_i)) + \frac{P_i\bar{w}}{c} \right) \right]^{\frac{1}{\mu-1}}, \text{ and}$$

$$w_i^* = \bar{w} \left(1 - \left[\frac{1}{\mu} \left(\frac{p\bar{w}}{c} (\lambda_i + \alpha_i(1 - \lambda_i)) + \frac{P_i\bar{w}}{c} \right) \right]^{\frac{1}{\mu-1}} \right).$$

Our first prediction is that as λ increases, responsiveness to changes in the individual-level price decreases.

Prediction 1: $\frac{\partial^2 W^*}{\partial P_i \partial \lambda_i} > 0$.

Proof:

$$\frac{\partial w_i^*}{\partial P_i} = -\frac{\bar{w}^2}{c\mu(\mu-1)} \left[\frac{1}{\mu} \left(\frac{p\bar{w}}{c} (\lambda_i + \alpha_i(1 - \lambda_i)) + \frac{P_i\bar{w}}{c} \right) \right]^{\frac{2-\mu}{\mu-1}}.$$

Next, taking the derivative of the expression above with respect to λ_i gives,

$$\frac{\partial^2 w_i^*}{\partial P_i \partial \lambda_i} = -\left(\frac{2-\mu}{\mu-1} \right) \frac{p\bar{w}^3(1-\alpha_i)}{c^2\mu^2(\mu-1)} \left[\frac{1}{\mu} \left(\frac{p\bar{w}}{c} (\lambda_i + \alpha_i(1 - \lambda_i)) + \frac{P_i\bar{w}}{c} \right) \right]^{\frac{3-2\mu}{\mu-1}}.$$

The expression above is positive for $\mu > 2$. Finally, w_{-i} is not affected by a change in P_i that is unobserved by individual $-i$, so $\frac{\partial^2 W^*}{\partial P_i \partial \lambda_i} = \frac{\partial^2 w_i^*}{\partial P_i \partial \lambda_i}$. ■

The second prediction is that an individual's responsiveness to the household-level price is increasing in her altruism toward her spouse (her α_i).

Prediction 2: $\frac{\partial^2 w_i^*}{\partial p \partial \alpha_i} < 0$.

Proof:

Here we set $P_i = 0$.

$$\frac{\partial w_i^*}{\partial p_i} = -\frac{\bar{w}^{\frac{\mu}{\mu-1}} (\lambda_i + \alpha_i(1 - \lambda_i))^{\frac{1}{\mu-1}}}{(c\mu)^{\frac{1}{\mu-1}} (\mu-1)} p^{\frac{2-\mu}{\mu-1}}.$$

Thus, the second derivative is given by:

$$\frac{\partial^2 w_i^*}{\partial p \partial \alpha_i} = - \frac{\bar{w}^{\frac{\mu}{\mu-1}} (1 - \lambda_i) ((\lambda_i + \alpha_i (1 - \lambda_i))^{\frac{2-\mu}{\mu-1}})}{(c\mu)^{\frac{1}{\mu-1}} (\mu - 1)^2} p^{\frac{2-\mu}{\mu-1}}.$$

$$< 0 \quad \blacksquare$$

Note that the result above holds even for $\mu \leq 2$.

Finally, we show, as a corollary, that the derivative of $\frac{\partial^2 w_i^*}{\partial p \partial \alpha_i}$ with respect to λ_i is positive, and hence that the cross-partial with regard to α_i is strongest when $\lambda_i \rightarrow 0$ and vice-versa; i.e. i 's altruism towards her spouse matters most when i is not the primary bill-payer.

$$\begin{aligned} \frac{\partial^3 w}{\partial \lambda_i \partial \alpha_i \partial p} &= \frac{\partial}{\partial \lambda_i} \left[- \frac{\bar{w}^{\frac{\mu}{\mu-1}} (1 - \lambda_i) ((\lambda_i + \alpha_i (1 - \lambda_i))^{\frac{2-\mu}{\mu-1}})}{(c\mu)^{\frac{1}{\mu-1}} (\mu - 1)^2} p^{\frac{2-\mu}{\mu-1}} \right] \\ &= \left(-p^{\frac{2-\mu}{\mu-1}} \right) \left(\frac{\bar{w}^{\frac{\mu}{\mu-1}}}{(c\mu)^{\frac{1}{\mu-1}} (\mu - 1)^2} \right) \frac{\partial}{\partial \lambda_i} \left[(1 - \lambda_i) ((\lambda_i + \alpha_i (1 - \lambda_i))^{\frac{2-\mu}{\mu-1}}) \right] \\ &= \left(-p^{\frac{2-\mu}{\mu-1}} \right) \left(\frac{\bar{w}^{\frac{\mu}{\mu-1}}}{(c\mu)^{\frac{1}{\mu-1}} (\mu - 1)^2} \right) \left[-(\lambda_i + \alpha_i (1 - \lambda_i))^{\frac{2-\mu}{\mu-1}} + \left(\frac{2-\mu}{\mu-1} \right) (1 - \lambda_i) (1 - \alpha_i) (\lambda_i + \alpha_i (1 - \lambda_i))^{\frac{3-2\mu}{\mu-1}} \right] \\ &> 0 \end{aligned}$$

since both of the summed terms in square brackets are < 0 for $\mu > 2$.

A.2 Appendix figures and tables

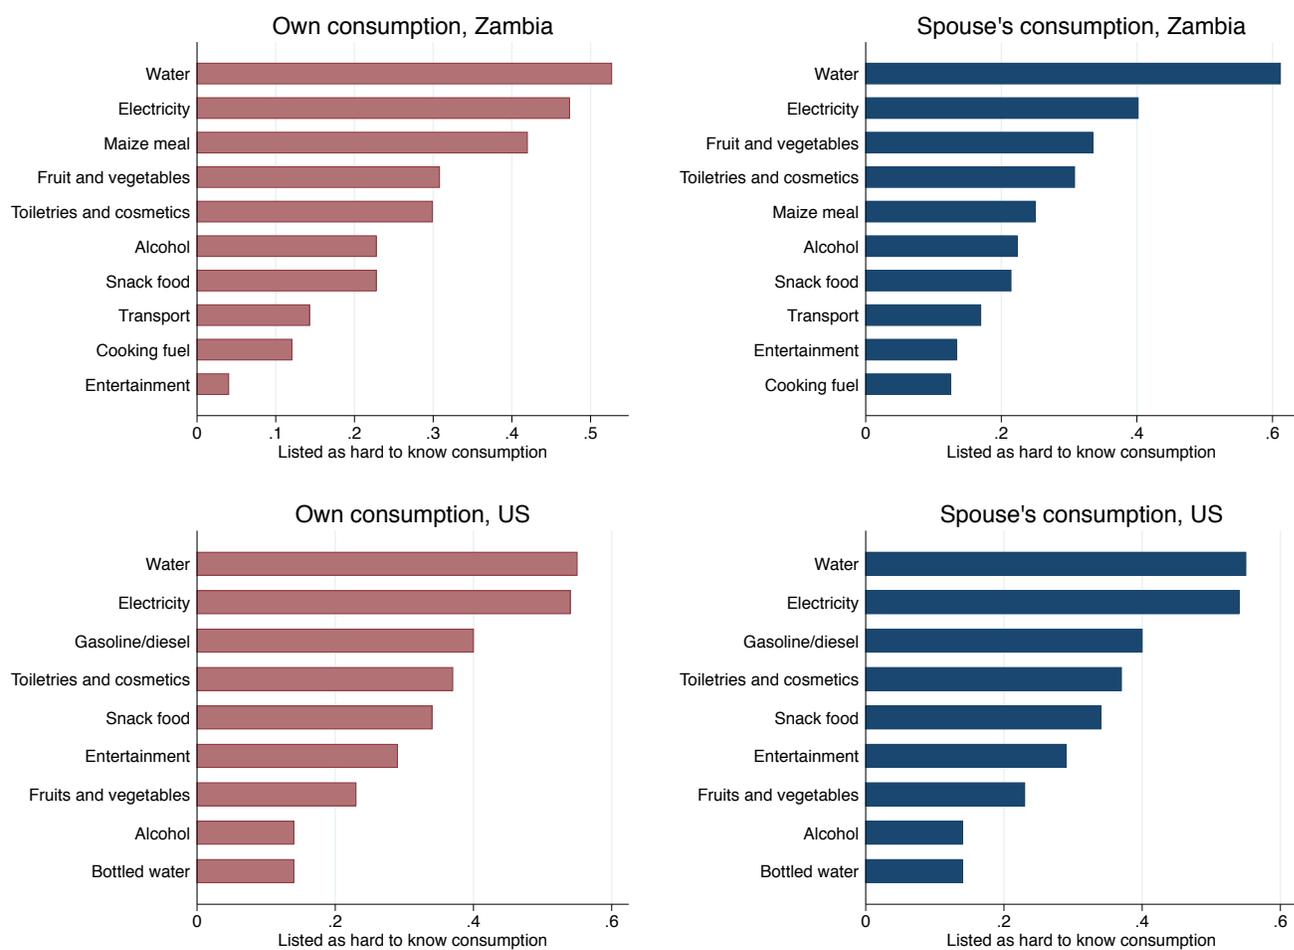


Figure A.1: Observability of consumption

Notes: Share of respondents reporting that a consumption category was among the top three most difficult to observe own (left) and spouse's (right) consumption. Respondents in the top panel are a convenience sample of market-goers in Lusaka (N=96). Respondents in the bottom panel are a sample of Mechanical Turk users in the United States (N=116).

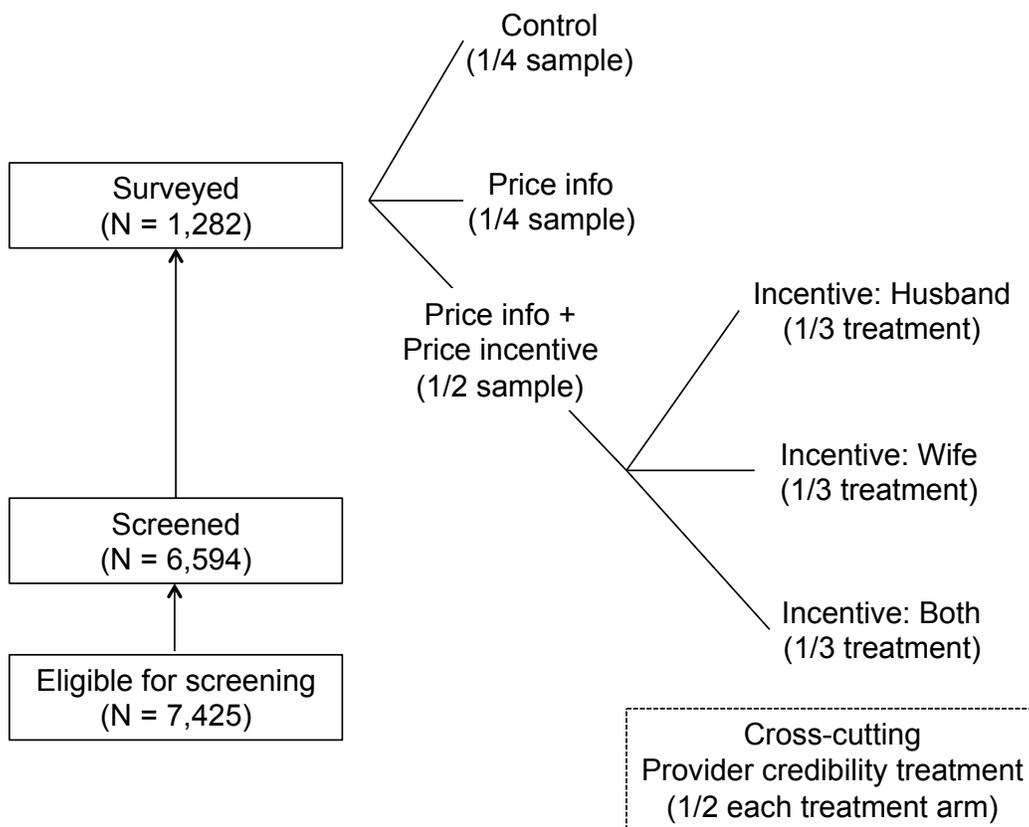


Figure A.2: Experimental design

Notes: Experimental design and sampling flow.

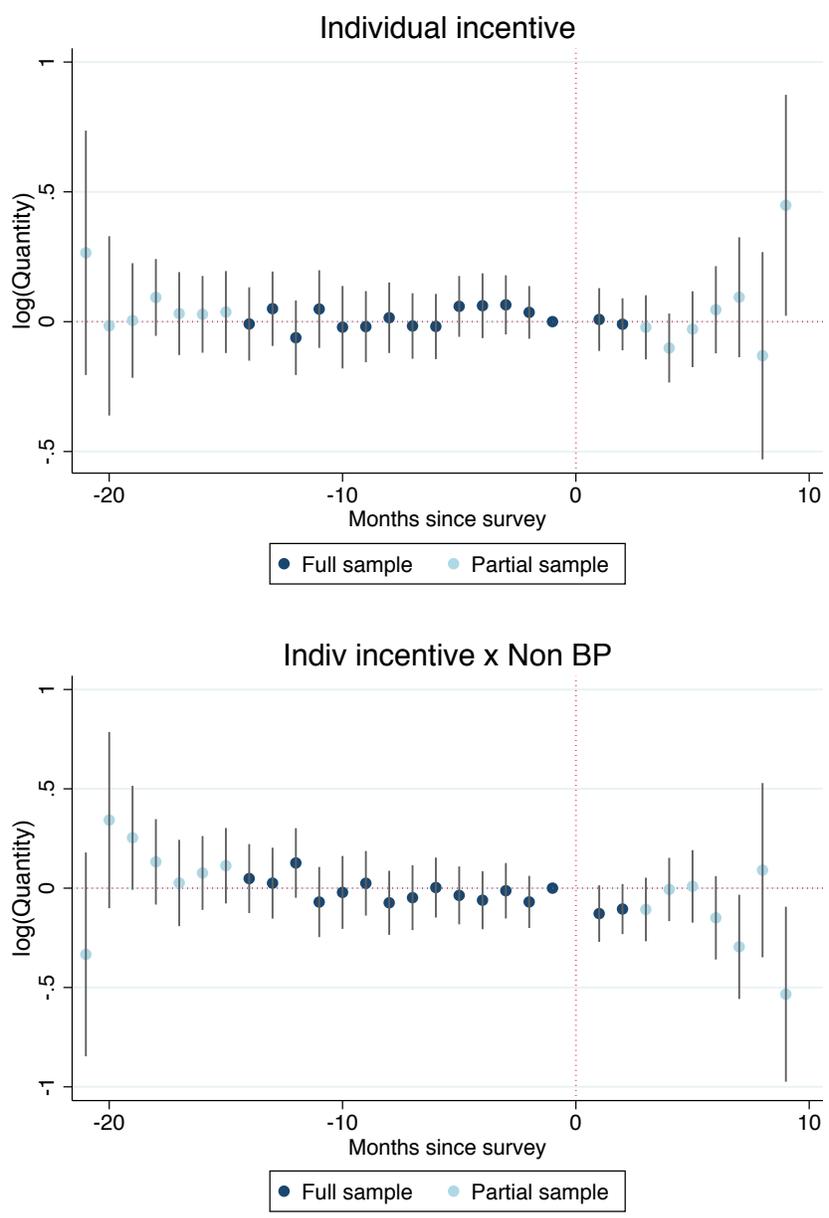


Figure A.3: Water outcomes by month, pre- and post-treatment

Notes: Regression coefficients from event-study specifications that interact treatment with an indicator for each month pre- and post-treatment. The top figure plots the coefficient for the individual incentive arm relative to the control group. The bottom figure plots the coefficient on the interaction term (mirroring our main specification). The darker colored markers indicate event-months that include the full sample; the light colored markers indicate event-months that are estimated off of only a sub-sample.

Table A.1: Bill payer definitions, by spouse

Using payer variable			
Husband's definition			
Wife's definition	Husband	Wife	Both/other
Husband	684	53	22
Wife	214	211	32
Both/other	49	9	8
Using income variable			
Husband's definition			
Wife's definition	Husband	Wife	Both/other
Husband	891	38	28
Wife	133	100	26
Both/other	52	6	8

Notes: Bill payer definitions, by spouse. The version shown in the top panel, which gives precedence to who physically pays the bill if that variable disagrees with whose income is used to pay the bill, is used in the main analysis.

Table A.2: Average effects of all treatments

	log (Quantity)		
	(1)	(2)	(3)
Couple incentive	-0.050 [0.039]	-0.050 [0.039]	-0.046 [0.041]
Individual incentive	-0.059** [0.026]		-0.055* [0.030]
Post-survey	-0.024 [0.025]	-0.024 [0.025]	-0.034 [0.033]
Husband incentive		-0.036 [0.032]	
Wife incentive		-0.083** [0.034]	
Provider credibility			0.029 [0.024]
Price information			-0.009 [0.032]
Couple = Indiv (p-val)	0.824		0.823
Couple = Wife (p-val)		0.473	
Husband = Wife (p-val)		0.246	
Observations (HH)	1,282	1,282	1,282
Observations (HH-months)	25,506	25,506	25,506

Notes: The panel begins in March 2014 and ends in February 2016. Standard errors are clustered at the household level, and all columns control for household and month-year fixed effects, an indicator for months following a missing quantity observation, and a continuous month-year variable interacted with sampling wave. *Provider credibility* and *Price information* are indicators that equal 1 post-survey for households assigned to receive the provider credibility treatment or price information treatment, respectively. All households in the incentive treatment also received the information treatment.

Table A.3: Heterogeneous effects of price information and provider credibility treatments

	log (Quantity) (1)	log (Quantity) (2)
Price information treatment	-0.003 [0.048]	
Info x Underestimated price	-0.016 [0.060]	
Provider credibility treatment		0.009 [0.033]
Provider credibility x Distrust billing		0.048 [0.048]
Observations (HH)	1,282	1,282
Observations (HH-months)	25,506	25,506

Notes: *Underestimated price* equals one if either spouse underestimated the marginal price of water. *Distrust billing* equals one if both spouses blame a high water bill on the provider. Regressions include the post-survey indicator interacted with the heterogeneity variables. The incentive treatment indicator is excluded. The panel begins in March 2014 and ends in February 2016. Standard errors are clustered at the household level, and all columns control for household and month-year fixed effects, an indicator for months following a missing quantity observation, and a continuous month-year variable interacted with sampling wave. Price beliefs are imputed for 257 households.

Table A.4: Individual incentive to non bill payer, robustness to different samples

	log(Quantity)		
	(1)	(2)	(3)
Individual incentive	-0.005 [0.035]	-0.027 [0.037]	-0.004 [0.035]
Indiv incentive x Non BP	-0.109** [0.047]	-0.109** [0.047]	-0.109** [0.047]
Couple incentive			-0.050 [0.039]
Included group	Strong disagree	Both say both/other	All
Observations (HH)	1,092	878	1,282
Observations (HH-months)	21,685	17,416	25,506

Notes: Columns add groups of households excluded in the main results to the specification in Column 1 of Table 2. Column 1 adds households where husband and wife strongly disagree on the definition of bill payer. Column 2 adds households where both spouses say “both/other”. Column 3 includes both of these groups and the couple incentive treatment. The panel begins in March 2014 and ends in February 2016. Standard errors are clustered at the household level, and all columns control for household and month-year fixed effects, an indicator for months following a missing quantity observation, and a continuous month-year variable interacted with sampling wave.

Table A.5: Robustness check: Panel length

Panel start	Jan 2014	Mar 2014	May 2014	Jan 2014	Mar 2014	May 2014	14 mo pre
Panel end	Feb 2016	Feb 2016	Feb 2016	2 mo post	2 mo post	2 mo post	2 mo post
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Individual incentive	-0.029 [0.038]	-0.029 [0.037]	-0.019 [0.037]	-0.024 [0.042]	-0.024 [0.042]	-0.014 [0.041]	-0.018 [0.041]
Indiv incentive x Non BP	-0.117** [0.047]	-0.109** [0.047]	-0.103** [0.046]	-0.130*** [0.049]	-0.125*** [0.048]	-0.119** [0.048]	-0.103** [0.047]
Observations (HH)	870	870	870	870	870	870	870
Observations (HH-months)	18,552	17,253	15,922	15,891	14,592	13,261	12,151

Notes: Robustness check on the results reported in column 1 of Table 2 (which is reproduced in column 2 here). The specification shown in column 7 is balanced in event-time; the sample is restricted to the event-time months that are available for all households, those denoted with dark markers in Figure A.3. Standard errors are clustered at the household level, and all columns control for household and month-year fixed effects, an indicator for months following a missing quantity observation, and a continuous month-year variable interacted with sampling wave.

A.3 Data appendix

A.3.1 Sample selection

Using the panel of billing data for metered residential customers as of February 2015 ($N=9,868$),⁶ we eliminate households that did not have a working meter for at least 3 out of the 4 preceding months. We also exclude households that use no water (i.e., are billed for zero cubic meters) in more than half of the preceding 4 months. Households with very low variation in usage over the preceding four months are considered to have possibly tampered with the meter or have a delinquent meter reader.⁷ Households with consistently low usage are also excluded since they would be least able to adjust their water consumption in response to a price shock, and, moreover, reducing water use from a low base could be harmful, e.g., in terms of hygiene; we drop households if their usage was on the lowest price tier (less than 6 cubic meters) for more than 2 of the preceding 4 months. Households whose median water usage in the preceding four months was above the 99th percentile are also dropped. Finally we drop households with an extremely high outstanding balance with SWSC, or households that are owed a significant amount of money by SWSC, defined as 6 times or 4 times their median bill in the preceding four months, respectively. This yields a total of 7,425 households that we target for an in-person screening.

Households were visited by a surveyor to collect data on characteristics not observed in the billing data that were important for sampling. Specifically, we require that the water meter not be shared with other households, that the primary bill payer be married (or cohabiting) and that both spouses live at that address, and that the household was in residence for at least the 4-month period prior to April 2015. We also exclude households who say they are planning to move in the following 6 months.

Our surveyors made up to 3 attempts to screen each households; any adult member of the household could be given the screening questionnaire. In total, 6,594 households were screened, of which 31 percent (2,051) met all our screening criteria.⁸

Households that met the screening criteria were informed about the survey. We scheduled a follow-up visit with the primary bill payer and his/her spouse, emphasizing that we needed both of them to be present for the full survey. We also informed respondents they would be

⁶This number excludes roughly 300 households we included in a pilot, who were deemed ineligible for the full study.

⁷They were excluded based on the following criteria: if the coefficient of variation in this period was less than 0.05, or if the quantity reported was identical for 3 or more months.

⁸Reasons for not screening a household include that the home was vacant or under construction, that it was occupied by a business, or that no one was home for three consecutive attempts.

compensated 40 Kwacha (4 USD) for participating in the survey.

We scheduled survey appointments with 1,817 households from our eligible sample. Of these, we completed surveys with 1,282 households. This high “attrition” rate is due largely to stopping our attempt to survey households at the end of December 2015.

For the full survey, at the scheduled time and date, a pair of surveyors (always a woman and a man) visited the screened-in household. After a few preliminary demographic questions, husbands and wives were separated and surveyed individually in different rooms of the house. Enumerators elicited water price beliefs, asked for perceptions of own and family members’ water usage, and conducted the modified dictator game. After finishing their individual questionnaires, both surveyors and respondents met back together in a common room for the last survey questions, and to receive the price information treatment (if applicable).

A.3.2 Calculating price elasticities

To illustrate magnitudes, we use the estimates of β_1 associated with our incentive treatment in equation (1) to calculate short-run price elasticities as follows.⁹ First, with y_{it} equal to log of monthly water quantity, we can interpret the coefficient on $treat_{it}$ as $\partial \ln(q)/\partial treat$, which we divide by the impact of the treatment on price, $\partial p/\partial treat$. This results in $\partial q/q \times 1/\partial p$, which we multiply by the pre-intervention average price to deliver a short run elasticity. We calculate customer-specific average prices, accounting for the increasing block schedule and for inflation (Zambian consumer price index), prior to the intervention and use that as the basis for our subgroup-specific average marginal prices.

For example, in the main text, we interpret the impact of the effect of delivering the incentive to the non bill payer as a short run price elasticity. We observe a statistically significant 0.14 log point decrease in monthly consumption in response to treatment. For this sub-group, the average pre-intervention price is 4.89 Kwacha per cubic meter and the reduction in consumption required to qualify for the lottery (which pays 15 Kwacha in expectation based on a one in twenty chance of being drawn) is 5.74 cubic meters. The implied short run price elasticity is therefore -0.26.¹⁰

⁹We convert our treatment effects into elasticities to aid interpretation of the magnitudes. However, we note a number of caveats to this transformation. Specifically, the elasticity calculation requires a number of assumptions: (1) that households respond similarly to a discrete price change as to a continuous price change, (2) that households respond similarly to a quantity target as to a continuous price change, and (3) that households respond similarly to a probabilistic payout as to a certain payout from conservation with the same expected value.

¹⁰Our calculated short-run price elasticity of demand is slightly below the mean found in the literature reviewed by Dalhuisen et al. (2003) and in line with the short run elasticities summarized in Worthington and Hoffman (2008).

A.4 Scripts

A.4.1 Price incentive treatment

[Private – to be read to husband/wife before they are brought back together]

Thank you for answering these questions. Before I go to check with my colleague, I have good news: We are running a program that gives prizes to people who cut down their water bill.

We will run a raffle, which has a K. 300 cash prize, and you will be entered into the raffle if your household reduces your water use by 30% next month. Since we are now in [current month]'s billing cycle, we will not consider this month's water use, but use [next month's] water use instead. This shows up on the [next month + 1's] bill.

If your water use in [next month] is below X cubic meters, then you will be entered for the draw. You can check the actual [next month] usage on the bill in [next month + 1] to see if it is X or lower. *[Point out where to locate the water quantity on the bill.]*

The lottery winner will be picked on the 15th of [next month + 2].

If you make the required reduction, you will have a 1 in 20 chance of winning the prize. In other words, for every 20 people who qualify for the raffle based on their bills in [next month + 1], we will draw one winner.

If you are the winner, we will call you on the number you gave us previously to convey the good news.

You will be requested to come to our office in Mosi-oa-Tunya House to collect the prize money, and you will also be compensated K.20 for your transportation.

We will continue to run a raffle every month at least until the end of the year and maybe longer, so if you also reduce water use to X in the months after [next month], you will be entered into that month's raffle too, so if you don't win in one month, you could still win the next month as long as the usage on your bill for that month is less than X cubic meters. You could even be a winner in multiple months!

How do we figure out how much you have to cut back to qualify for the raffle? We look at how much your household used in this year's March and April bills. In these bills (March and April) your average use was for Y cubic meters. So you need to cut your household usage by Y-X cubic meters in order to achieve X cubic meters or lower and qualify to our draw. For every household in this program, the target water usage is based on their own past usage during those two months.

[If only the husband/only the wife is receiving the treatment]: Not all individuals or all

households are getting the opportunity to try for the raffle. In particular, you have been selected, so I am only informing you of this, and not your husband/wife.

My colleague is not informing your husband/wife about this either, because for your household, only you have been selected to participate. It is entirely up to you if you want to inform him/her or not.

If you would like to check whether your household cut back usage enough to qualify for the raffle, you may call 096-934-3167 after the 15th of [next month + 2]. You will not be charged any airtime to call this number.

When you call, the line will be cut immediately and you will automatically be called back from a different number. When you pick up the phone, you will hear a recorded message that tells you if you qualified for the raffle or not. The message is linked to the number you gave us, so please use the same sim card when you call.

You can also call that number after the 15th of each month following [next month + 2] to see if you qualified for that month's raffle.

You can also use that number to check if the raffle program is still going on.

If you win, we will ensure that we are speaking only with you when we call to inform you. Nobody else will know that you have won, unless you share the news.

[*If both are receiving the treatment*]: Not all individuals or all households are getting the opportunity to try for the raffle. In particular, your household has been selected. Just as I am informing you of this raffle, my colleague in the other room is informing your spouse about it as well.

If your household wins, we will inform both of you, and we would appreciate it if you both came to collect the prize. If you would like to check whether your household cut back usage enough to qualify for the raffle, you may call 096-934-3167 after the 15th of [next month + 2]. You will not be charged any airtime to call this number.

When you call, the line will be cut immediately and you will automatically be called back from a different number. When you pick up the phone, you will hear a recorded message that tells you if you qualified for the raffle or not. The message is linked to the number you gave us, so please use the same sim card when you call.

You can also call that number after the 15th of each month following [next month + 2] to see if you qualified for that month's raffle. You can also use that number to check if the raffle program is still going on.

If you win, we will ensure that we are speaking with you or your spouse when we call to inform you. Nobody else, other than your spouse, will know that you have won, unless you

share the news.

[*For everyone*]: Only people in some of the households we are surveying are eligible for this raffle, so others that you speak to may not have been given this opportunity. The raffle is sponsored by our research project, not SWSC – they will not be aware if you are eligible or not, or if you won or not.

A.4.2 Provider credibility treatment

We have collected this information purely for research and will not share any details with SWSC. However, we want to provide you with a little bit of extra information about how SWSC calculates your bill. SWSC tries to ensure that bills are accurate by reading your meter monthly and using the amount of water consumption shown on your meter to calculate your bill. That is, the amount that you are charged is based on the amount of water you use. The meter readings taken this month measure your usage since the time when last month's reading was taken. Once SWSC has collected all the readings for this month, this is used to calculate the bill that will be given to you next month. For example, when you received your water bill in March you were charged for the water your household used between the 21st of January and the 20th of February, roughly speaking. When you received your water bill in April, you were charged for the water your household used between the 21st of February and the 20th of March, and so on. If there are some months that they cannot get a meter reading, then you are charged an estimate based on your previous consumption, and they try to get meter readings again as soon as possible. Then the next time they read your meter, they adjust your bill for any over- or under- charges from the months when they were not able to do the reading. SWSC is taking measures to make sure that bills are fair and based on actual water usage. They are committed to honest billing practices.