Environmental externalities and free-riding in the household*

B. Kelsey Jack[†]

Seema Jayachandran[‡]

Flavio Malagutti[§]

Sarojini Rao[¶]

May 10, 2023

Abstract

In addition to generating a negative environmental externality, a household's water consumption entails another "market failure": household members free-ride off each other and overconsume. The problem stems from consumption being billed at the household level and the difficulty of monitoring one another's consumption. 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 increases. Another prediction is that this overconsumption problem is more acute when the financial benefit of a lower utility bill is shared unevenly among household members. We show that households are indeed more responsive to a change in the household-level price of water when their financial arrangements are more equal. Our results offer a novel explanation for the low price sensitivity of residential water (and electricity) consumption.

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

[‡]Princeton University, jayachandran@princeton.edu

[§]University of California, Santa Barbara, fmalagutti@ucsb.edu

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

1 Introduction

Because of negative environmental externalities, the level of water or energy consumption that is privately optimal typically exceeds what is socially optimal. This paper highlights a second reason that water and energy are over-consumed: Household members have an opportunity to free-ride off each other, due to the fact that their usage is pooled into one bill. This intrahousehold inefficiency is a potential contributor to the inelastic demand for water and electricity that is observed among residential customers in many settings.

This 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. 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 combines all household members' usage. Second, it is difficult for household members to back out individual-level use.[1]

We develop a conceptual framework for understanding free-riding incentives within the home, based on the observation that household members 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 "primary residual claimant": the person who claims most of the residual household income if the utility bill is lower. Our main prediction from the model is that a person-specific price change will have a larger effect on consumption for someone who is not the primary residual claimant. 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. Another prediction of the model is that the effect of a household-level (i.e., standard) price change depends on how evenly residual claimancy is shared. All else equal, a household whose members have more similar conservation incentives will respond more to the price change. Both predictions stem from a convex cost of individual-level water conservation.

---

[1]We conducted a small survey (a) in Lusaka, Zambia and (b) among mTurk users in the US. When asked for which consumption categories is tracking their spouse's consumption most difficult, water and electricity were the most common responses. They were also the most common responses for tracking own consumption. See Appendix Figure A.1.

We test these predictions 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.

To generate a person-specific price change, we inform either the man or the woman about the rewards program (*individual incentive* treatment). 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 primary residual claimant. We refer to the other spouse as the "non-residual claimant." We also have a treatment arm in which the prospect of the reward is communicated to both spouses (*couple incentive*), which acts like an increase in the household-level price. Our main outcome is the household's water usage for the two to nine months the incentives are in place.

As predicted, the response to the incentive treatments varies with status quo incentives for conservation. First, water use declines considerably more if the individual incentive is given to the non-residual claimant rather than the residual claimant. Men are most often the residual claimants, 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-residual claimant. Second, the effect of the couple incentive on water use is larger when the couple reports sharing residual claimant status more equally, though this result is underpowered.

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 misper-

---

[2]Livingstone'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.

ceptions 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 2020; 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.[3]

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; Duflo and Udry 2004; Mazzocco 2007; Robinson 2012; Angelucci and Garlick 2016). We highlight hidden action (specifically, limited information about consumption) as the source of inefficiency, unlike most previous work which explores limited commitment or hidden income.

2 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 that are characterized by greater observability of actions or individual-level purchases.

---

[3]The structure of the problem is similar in settings where multiple households share a water or electricity meter, which is common in low and middle income countries.

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 or meter sharing arrangements.

2.1 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 $\mu > 2$.[4] Individuals consume a maximum quantity of water given by $\bar{w}$ if they exert no effort at all towards conserving water.[5] 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 cost, utility is concave in water consumption and linear in other consumption.

We model a household as comprising two individuals, a husband and a wife. Assuming equal welfare weights on each person's utility, the household's optimal choice of conservation effort is symmetric across individuals and is given by:

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

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

2.2 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

---

[4]Footnote 8 gives the intuition for why our prediction requires that the effort cost function is steeper than quadratic.

[5]The 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.

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 determines the ex post division of income that remains after the household pays the water bill; each spouse controls a share $\lambda_i \in [0, 1]$, with $\lambda_i + \lambda_{-i} = 1$. (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.) 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:

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

In addition, individuals internalize some share $0 \leq \alpha_i \leq 1$ of their spouse's utility. Thus, i's utility function is given by $u_i = v_i + \alpha_i v_{-i}$. The prediction that we derive and test does not require intrahousehold altruism, but we include it in the model because it is realistic, and to emphasize that a non-cooperative framework still allows for altruism among household members.[6]

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,

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

If $\lambda_i = 1$ (full control over savings from water conservation) or $\alpha_i = 1$ (perfect altruism toward one's spouse), then person i fully internalizes the household's cost of water consumption. Her

---

[6]A person might also internalize how her water use affects her spouse's income because of enforcement of household agreements around water use (if individual water use is partly observable). The parameter α_i can also be thought of as a reduced-form representation of this enforcement.

conservation effort is at her first-best level: $e_i^* = e_i^{FB} = \left(\frac{1}{\mu}\frac{p\bar{w}}{c}\right)^{\frac{1}{\mu-1}}$. However, if $\lambda_i = 1$, then $\lambda_{-i} = 0$, and individual $-i$ only exerts effort insofar as he is altruistic toward his spouse.

More generally, equation (2) 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 one's spouse all lead to lower water consumption.[7]

Our empirical focus is on how λ_i, the individual claim of residual income from the water bill, 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 then becomes $v_i' = \lambda_i(Y - pW) - ce_i^\mu - P_iW$. Importantly, the new term that depends on P_i enters into v_i without being diluted by λ_i. The individual's optimal effort is as follows:

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

2.3   Effect of an individual-level price change

Our experimental treatments make water use more costly, effectively increasing the price of water. In our data, we observe household-level water use, which we take to be the sum of individual optimization and denote $W^* = w_i^* + w_{-i}^*$. We thus derive a testable prediction about how the sensitivity of household water use to the individual-level price, $\frac{\partial W^*}{\partial P_i}$, depends on the household's status quo financial arrangements. We present the result here and the proof in Appendix A.1. Note that the first derivative of water consumption with respect to price is negative, so a positive cross-derivative maps to a decrease 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 primary residual claimant (lower λ in the household) is more responsive

---

[7]The 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.

to changes in the individual-level price.

The intuition for the result is that the spouse with lower λ_i has weaker incentives to conserve based on the household-level price, so she is exerting less effort toward conservation to start with. Thus, given the convexity of the cost of effort, she faces a lower marginal cost of effort and changes effort more in response to the price change.[8] 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 primary residual claimant (lower λ_i) will have a larger effect on aggregate consumption than if the individual price is directed to the residual claimant.

2.4   Effect of a household-level price change

In addition to our main prediction about individual-level incentives, we derive an additional prediction about how the effect of a household-level price change (p) depends on how evenly residual claimancy is shared between the spouses.

Prediction 2: $\frac{\partial^2 W^*}{\partial |\lambda_i - \lambda_{-i}| \partial p} > 0$, or equivalently, $\left| \frac{\partial W^*}{\partial p} \right|$ is decreasing in $|\lambda_i - \lambda_{-i}|$. In words, households with a smaller difference in λ's are more responsive to changes in the household-level price.

Individuals with a higher λ_i are more sensitive to a change in the household level price, but at a rate that is decreasing in λ_i given the convexity of conservation effort costs. Given the diminishing returns to λ_i at the individual level, household-level price sensitivity is maximized when $\lambda_i = \lambda_{-i}$, and, more generally, households with smaller differences between their λ's are more price responsive than those with large differences. In a household with a large gap in λ's, the spouse with high λ_i has already internalized the household's water price and faces a high marginal cost of effort, which is not fully offset by the lower effort cost of his or her spouse, given the curvature in the cost of conservation.

---

[8] The curvature of the cost of effort function is important for this prediction. A higher λ_i individual starts at 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 conservation effort (i.e., she internalizes p more). The marginal cost of effort must be increasing steeply enough to offset the marginal savings for the high λ_i individual. For our set up, this requires the cost function to be greater than quadratic.

In the appendix, we present an discuss an additional prediction about the ordering of these predictions: the effect of a household-level price change on water use is smaller than the effect of an individual-level price change directed to the spouse with lower residual claim but larger than the effect of an individual price change for the spouse with higher residual claim.

2.5 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. 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. 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.[9]

Adding to these observability challenges, knowing one's own consumption is often diffi-

---

[9]This 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).

cult.[10] 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: $\hat{w}_{-i} = W - (w_i + \epsilon)$. 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).

3   Data and experimental design

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

3.1   Study sample

Our full sample comprises 1,282 married couples who are SWSC customers. Since the setting is urban and everyone in our sample has piped water, they are mostly middle class. 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.

---

[10]The 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.

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 (to ensure that we were not encouraging unhealthily low usage); 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.[11]

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.[12] 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.

3.2 Billing data

SWSC conducts in-person meter readings each month and bills households monthly based on the meter reading. We measure water consumption using this billing data.

We create a panel of billing data from March 2014 to February 2016. The end date aligns

---

[11]We 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.

[12]We report 2015 USD values using an exchange rate of 10 Kwacha per USD.

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. The panel is balanced in calendar time, so the estimated treatment effect averages across households with different treatment duration; treated households have at least 2 and up to 9 months of treatment, with an average of 5.3 months. 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 (i.e., no meter reading took place) 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.[13]

3.3 Survey data and construction of residual claimant variable

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.

---

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

A key variable for us is who the household's primary residual claimant (RC) on the water bill is. This concept is difficult to measure directly; in piloting, many respondents were unable to understand direct questions about which spouse enjoyed the financial benefit if the bill was low and bore the cost if the bill was high. Instead, we construct the measure using two survey questions that, based on our piloting, typically map to residual claimancy: whose income is used to pay the bill and who physically pays the bill (payment is in person). Spouses seem to have claim over the income they earn, and the person who pays the bill usually does so out of a larger pool of money and has control over the balance, so both of these factors are proxies for residual claim.

We asked each spouse these questions about whose income is used and about who pays the bill. The possible answers are oneself, one's spouse, both jointly, or someone else. We code the RC variable as follows (illustrated for when the wife is the respondent). If her response is herself for both questions, we code her RC variable equal to 1. It is equal to zero if her response is her spouse for both questions. If her response to both question is both spouses or someone else, we code as 0.5. Less clear-cut is when she says she pays the bill but her husband's income is used, or vice versa. We prioritize the payer variable 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. Specifically, we code her RC variable as equal to 1 if she says she pays the bill but the income used is from both of them or someone else, and 0.5 if her income is used but both or someone else pays the bill. Thus, at the respondent level, RC equals 0, .5, or 1 for each spouse, and the sum of RC 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 RC according to the husband and then average the husband's and wife's RC variables. These averages are our preferred measures of the husband's RC status and wife's RC status.[14] There are four cases. (1) RC equals 1 for the husband and 0 for the wife, or vice versa. The spouses agree on who the RC is. (2) RC equals 0.75 for one spouse, and 0.25 for the other; this arises when one respondent identifies a specific person as the RC, while the other views the role as evenly split. (3) RC equals 0.5 because one member of the

---

[14]Appendix Table A.1 summarizes the wife's and husband's RC variables.

couple says the husband is RC and the other says the wife; here they strongly disagree. (4) RC equals 0.5 because both think the RC role is shared.

For the first two cases, there is within-couple variation in the RC variable. Thus, when we randomize the incentive to conserve at the individual level, we induce randomized variation in the recipient's RC 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 RC, which is, by construction, always 0.5. We exclude the last two cases from the analysis (227 and 31 households, respectively) because the incentive recipient has $RC = 0.5$ regardless of the randomization outcome. We show that the results are similar when we include these households.

We also use the payer and income variables to construct a measure of shared residual claimant status that we use to test Prediction 2. The *Equal RC status* variable equals 1 if both spouses reply "both" to at least one of the two questions about bill payer status and income toward the bill and equals 0.5 if one member of the couple replies "both" to at least one of the questions. For all other households, i.e., where all responses refer to one member of the couple or to someone else in the household, the variable is set to zero. Note that we need to use a somewhat expansive definition of equality for there to be enough variation in *Equal RC status* to test Prediction 2. We show robustness to coding the main residual claimant variable to match the *Equal RC status* variable.

3.4   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. One or both spouses were informed during their survey visit that they were being offered a monetary incentive to reduce water use.[15] In the months following the visit, if the household reduced its consumption by a specified amount, the individual or household was entered into a monthly lottery

---

[15]The script is provided in Appendix A.4.

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.[16] 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.[17]

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.[18] 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, p. 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, P_i; 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. Individuals who won a prize were informed and paid privately in the latter two sub-treatments.

In the notation of the household model, the incentive treatment when both spouses are told about the prize adds a term to the indirect utility function, $v_i = \lambda_i(Y - pW + L \times \mathbf{1}(W \leq \bar{W})) - e_i^\mu$, where L is the expected value of the lottery payout. The individual sub-treatments,

---

[16]The 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.

[17]Given a target reduction of 5.8 m^3 and a price of 5.1 Kwacha/m^3 over that range of consumption, the expected value of the prize was around 2 Kwacha/m^3. This calculation accounts for the increasing block tariff.

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

in which surveyors informed only one spouse about the prize, move the lottery payoff outside of the λ_i term: $v_i = \lambda_i(Y - pW) + L \times \mathbf{1}(W \leq \bar{W}) - e_i^\mu$, similar to an increase in P_i in the model. This increases i's unilateral payoff from water conservation, which has a larger effect on overall household consumption if $\lambda_i < \lambda_{-i}$. Of course, individuals could share the information with their spouse or the spouse might inadvertently find out about it in some cases, but the individual-specific treatment comes closer to an individual price than does the joint treatment.

We test our main prediction by comparing the two individual incentive sub-treatments. Our secondary treatment relies on heterogeneity within the couple incentive treatment. The couple incentive arm also provides a benchmark similar to standard household-level pricing used by utilities.[19] The control group is helpful for being able to gauge the absolute magnitude of the effect and calculate a price elasticity within our sample.

3.5   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

---

[19]A previous, longer version of the paper included a different 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.

the appendix.

3.6 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 50 percent of households, both husband and wife agree that the husband is the primary residual claimant for the water bill. In 17 percent of households, they agree that the wife is the primary residual claimant for the water bill. In 32 percent of households, residual claimant status is shared, at least to some degree. In 17 percent of households they "strongly disagree," meaning that one spouse says that the husband is the residual claimant and the other says that the wife is. Table 1 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-residual claimant received the incentive, with the comparison group being households where the residual claimant received it.[20] F-tests indicate that, in all cases, we cannot reject balance between a subsample and its

---

[20]The residual claimant variable is not binary. For this table, we pool individuals with residual claimant status of 1 or 0.75 as residual claimants, and those with status of 0 or 0.25 as non-residual claimants. We omit households where residual claimant status is 0.5 for each spouse, so in which the variable has no within-couple variation.

comparison group. In addition, Figure 2 shows that there are parallel pre-trends between subsamples in average water use in the months leading up the survey.

4 Estimation strategy

We use the randomized variation in the individual and couple incentive to conserve water as person-specific and household-level price increases. These allow us to test the two predictions from our model: (1) The individual incentive is more effective in reducing household water use if it is offered to the spouse who is not the primary residual claimant. (2) The effect of the couple incentive is larger in households where residual claimant status is more evenly shared.

We use monthly household-level outcome data from before and after the intervention and estimate a difference-in-differences regression. The regressors $IndivTreat_{it}$ and $CoupleTreat_{it}$ equal 1 for households assigned to the individual treatment and couple treatment, respectively, in months after the survey; recall that treatments were delivered at the end of the survey visit.[21]

Our main prediction is that the individual treatment will be more effective if targeted to the non-RC (where $NonRC \equiv 1 - RC$.) Thus, the key regressor is $IndivTreat_{it} \times NonRC_{it}$. This is not a standard interaction in that $NonRC$ (the incentive being given to the non-RC) is only defined within the individual incentive group. Our main analysis omits 258 households with no within-couple variation in residual claimant status, because the random treatment assignment cannot generate random variation in the regressor of interest. (When testing Prediction 2, we add back the 31 households where the lack of variation comes from both spouses indicating shared residual claimant responsibilities.)

In the estimating equation below, β_1 identifies the effect on water consumption of the residual claimant receiving the individual incentive, relative to the control group. β_2 measures the additional effect of the incentive going to the non-RC. The main hypothesis is,

---

[21]We set these indicators 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.

thus, $\beta_2 < 0$.

$$
\begin{aligned}
y_{it} = {} & \alpha + \delta CoupleTreat_{it} + \beta_1 IndivTreat_{it} + \beta_2 IndivTreat_{it} \times NonRC_{it} \\
& + \gamma_1 Post_{it} + \gamma_2 Wave_i \times Month_t + \tau_t + \eta_i + \gamma_3 MissingFlag_{it} + \epsilon_{it}
\end{aligned}
\tag{4}
$$

The other regressors are control variables. $Post_{it}$ is a post-survey indicator, as the survey itself may have had effects; the variable varies across households within a 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 include 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.

Our second prediction is that the couple treatment will be more effective in households where residual claimant status is more equally shared between spouses. We define $Equal_i$ to equal one when both spouses indicate shared bill payment responsibilities (see Section 3.3). We modify equation (4) to test the prediction that $\delta_2 < 0$: an increase in the couple price decreases water use more in households where residual claimant status is more equal.

$$
\begin{aligned}
y_{it} = {} & \alpha + \delta_1 CoupleTreat_{it} + \delta_2 CoupleTreat_{it} \times Equal_i + \beta IndivTreat \\
& + \gamma_1 Post_{it} + \gamma_2 Wave_i \times Month_t + \tau_t + \eta_i + \gamma_3 MissingFlag_{it} + \epsilon_{it}
\end{aligned}
\tag{5}
$$

Other estimation details follow equation (4).

5   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, most of the mass is above the target level to be eligible for the prize: The treatment effect is due in large part to reductions not large enough to qualify for the prize.[22] If households could perfectly choose their consumption level, we would expect bunching just below the target among treated households. However, the difficulty of knowing one's own and other household members' 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 is shown in Table 2. The couple incentive treatment reduced water use by 0.50 log points, which is not significantly different from zero at conventional levels (column 1). The individual treatment significantly reduced water use, by 0.059 log points ($p < 0.05$). The individual incentive results in a larger magnitude reduction when delivered to the wife than to the husband, though the difference between the coefficients is not precisely estimated (column 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. Table 2 also estimates the effects of these treatments, and shows that the individual incentive treatment effect is largely unaffected (column 3). In the rest of the paper we pool all of the no-incentive households, both pure controls and those that received only price information. This increases statistical power when we estimate the overall effect of the individual incentive treatment; importantly, it does not affect the identification of our main coefficient of interest (β_2), 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.[23]

The main result of the paper is associated with Prediction 1 and shown in Table 3: The individual incentive causes a larger reduction in household water consumption when the recipient has lower residual claimant status. Column 1, which estimates equation (4), shows that when the incentive goes to the RC, there is an estimated 0.02 log point reduction in

---

[22]Out of 2,335 treated household-months, 431 had reductions large enough to qualify for the lottery.

[23]Appendix Table A.2 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.

water use, an effect statistically indistinguishable from zero. When the incentive goes to the non-RC, there is a significantly larger effect. The point estimates of -0.11 log points implies a total effect of -0.13 log points when the non-RC receives the incentive, equivalent to a short-run price elasticity of about -0.25. (Appendix A.3 provides the details of this elasticity calculation, which has many caveats.) As a benchmark, the price elasticity for the couple incentive treatment, which most closely resembles a standard household-level price change, is -0.13.

When the residual claimant received the incentive, in the absence of intrahousehold frictions, she should have been able to reproduce the effects of the non-residual claimant incentive arm: She could tell her spouse about the rewards program and promise her almost all of the prize. Our results are suggestive that residual claimants may not have thought to do this, or a commitment problem prevented it from being effective. We discuss this puzzle further 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 Figure 2, which plots the month-by-month coefficients corresponding to the specification in Table 3, 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 fewer households contribute to their estimation.)

Column 2 of Table 3 disentangles whether the effect is due to residual claimant status or gender. These two variables are correlated; in most cases, the husband has higher residual claimant status. Column 2 simultaneously estimates the effects of the incentive being given to the wife and to the non-RC. In other words, we estimate the effect of targeting the non-RC, controlling for gender, to determine if the former effect is driven entirely by gender. It is not: the interaction with residual claimant status remains significant and similar in magnitude to column 1, while gender of the recipient per se does not seem to affect responsiveness to the incentive. Column 3 shows robustness to including households where both spouses say that residual claimant status is shared.

We can also compare the effect of giving the individual incentive to the non-RC with the effect of the couple incentive treatment. The effect of providing the incentive to both

spouses lies in between providing it to the RC and providing it to the non-RC, consistent with the model (see Corollary in Appendix A.1), but we lack the statistical power to distinguish coefficients. The p-value on a test for equality of the coefficients is 0.18.

Next, Table 4 tests how the effect of the couple treatment depends on how residual claimant status is shared between spouses, following equation (5). Column 1 includes the full sample and shows that delivering the couple treatment to a household where spouses say that residual claimant status is shared results in a decrease in water use that is nearly 10 times larger than the effect of the couple treatment in households where residual claimant status is not shared. Eliminating households where spouses strongly disagree about residual claimant status (column 2) or households with no variation in residual claimant status (column 3) only increases the magnitude of the interaction term.[24] Across all sample restrictions, the interaction term is statistically insignificant at conventional levels, since it is estimated off of a very small sample of households. Nonetheless, the sign of the interaction term is consistent with Prediction 2: greater intrahousehold equality increases price sensitivity.

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 treatments.

5.1   Robustness checks

We only have statistically significant results for our test of Prediction 1, so we focus our robustness checks on this result.

The household's financial arrangements are, of course, not randomly assigned, so one potential concern is that our effect is not due to heterogeneity by RC status, but instead heterogeneity by some characteristic of the recipient that is correlated with RC status. To address this concern, we control for individual-level observables, measured in the baseline survey, in parallel to RC status.

For example, the non-RC might have a higher marginal utility from income because she has limited control of the household budget. This would make the incentive more valuable to her, explaining her greater responsiveness to it. We therefore control for whether the

---

[24]These sample restrictions match those in Table 3.

incentive recipient is employed, her education level, and her age, as proxies for income and generalized bargaining power.[25] Another concern is that the non-RC might usually be inattentive to water use, so the incentive has a greater salience effect for her. Her larger reduction in water use in response to the incentive could be due to salience rather than the incentive representing a larger effective price change for her. This alternative is in some ways similar in spirit to what we are highlighting — the household's financial arrangements make one spouse put in inefficiently low conservation effort. Nonetheless, we address this concern by controlling for the recipient's knowledge about the price of water and the household's water use on its most recent bill. Finally, the non-RC might have an easier time reducing water use because her baseline level of water use is higher. To address this, we control for whether the recipient is the biggest water user in the household.

Table A.3 reports the results when we include interactions of $IndivTreat$ with these other characteristics of the incentive recipient, first one at a time (column 1) and then all at once (column 2). While we do see some heterogeneity in price sensitivity based on some of these characteristics, our coefficient of interest ($IndivTreat \times Non\text{-}RC$) is stable in both magnitude and significance when we control in parallel for them in column 2.

We next test the sensitivity of our results to how we construct the residual claimant variable. In Table A.4, 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 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-RC. Column 3 restricts the sample to households where the couple strongly agrees on who the RC is and we find similar results. In columns 4 and 5, instead of using the average of the husband's and wife's RC assessments, we use just one respondent's answers. Finally, in column 6, we code the residual claimant variable to match the coding in the equal RC status variable. The results are similar to our main results.

Finally, we test for sensitivity of our results to how we construct the panel (Appendix Table A.5). The first three columns use different pre-treatment panel lengths, which does

---

[25]Education and age are both measured in coarse categories. We split both as close to the median as the data allow.

not affect the coefficients very much. The next three columns include only two post-survey months per household, to ensure that treated households contribute equally to the estimated treatment effect; the somewhat larger point estimate is consistent with the slight decay in the effect size after the third month seen in Appendix Figure 2. 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.

6   Discussion and 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 how much each person consumed. 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 reduction in the household's water use.

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. Women have the most scope to reduce household water use, but also the least incentive to do so. 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 usually assigning bill responsibility to the man, or more precisely, the smaller water user? Why is the bigger water user the primary residual claimant 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 our setting. In follow-up discussions with 40 households, most stated that using

a similar allowance-like 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 "primary residual claimants"; wasting water meant they had to spend more time fetching water.

Even if households improve how they split the bill, limited information about individual-level 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 smart-phone apps with real-time data would enable better monitoring of family members; 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. (2020). Asymmetric information in residential rental markets: Implications for the energy efficiency gap. *Journal of Public Economics 190*, 104251.

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.

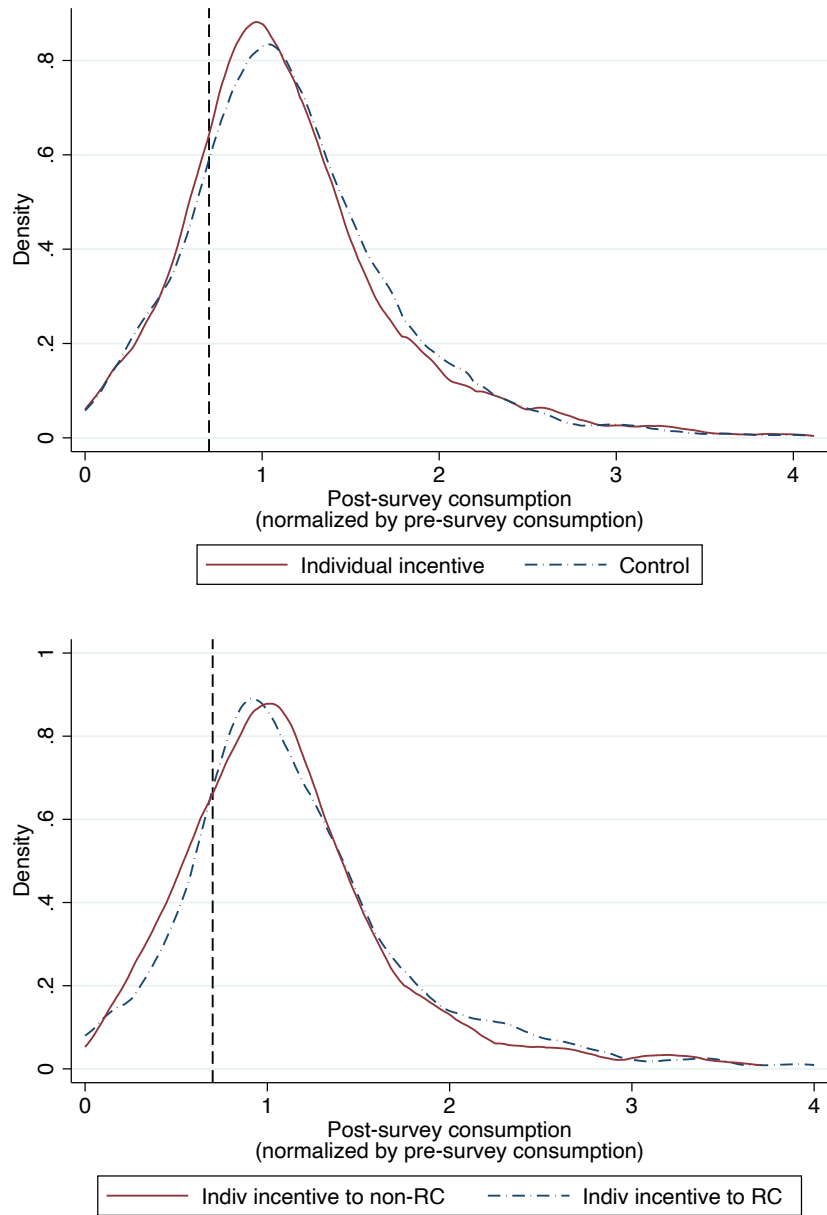Figure 1: Post-intervention water consumption, relative to pre-intervention

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.
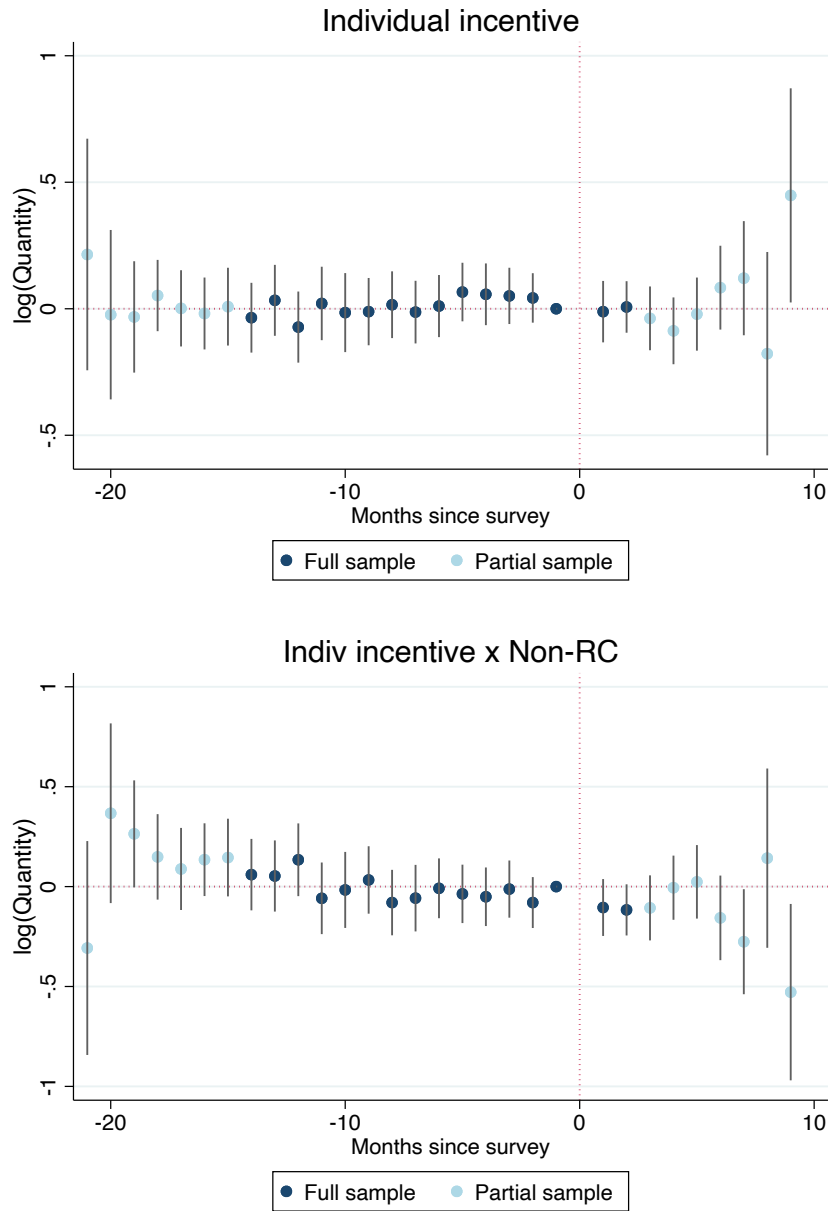
Figure 2: Water outcomes by month, pre- and post-treatment

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. Households with no variation in residual claimant status are excluded.

Table 1: Balance: Incentive treatment arms

Treatment arm:	No incentive (1)	Couple incentive (2)	Individual incentive (3)	Wife incentive (4)	Non-RC incentive (5)
Quantity of water consumed	18.844	-0.002	-0.002	0.005	0.005
	(11.922)	(0.002)	(0.002)	(0.003)	(0.003)
Household size	5.860	0.005	-0.000	0.003	-0.007
	(2.286)	(0.006)	(0.007)	(0.011)	(0.012)
HH has maid	0.169	-0.016	-0.032	0.073	-0.080
	(0.375)	(0.041)	(0.043)	(0.072)	(0.082)
HH owns home	0.512	-0.010	-0.015	-0.035	-0.033
	(0.500)	(0.030)	(0.031)	(0.050)	(0.057)
Rooms in home	3.529	0.002	0.017	-0.033	-0.000
	(1.264)	(0.013)	(0.013)	(0.019)	(0.024)
Both know bill quantity	0.104	0.140	0.066	-0.023	0.081
	(0.305)	(0.045)	(0.048)	(0.076)	(0.086)
Both know bill charge	0.678	0.002	0.024	-0.058	-0.075
	(0.468)	(0.031)	(0.032)	(0.053)	(0.059)
Agree W is bigger water user	0.795	0.019	0.094	-0.056	0.078
	(0.404)	(0.036)	(0.039)	(0.068)	(0.075)
Agree H is RC	0.489	-0.042	0.022	-0.118	-0.003
	(0.500)	(0.040)	(0.043)	(0.069)	(0.076)
Agree W is RC	0.167	-0.062	-0.005	-0.064	-0.045
	(0.373)	(0.049)	(0.052)	(0.084)	(0.090)
RC role is shared	0.320	-0.004	-0.014	-0.000	0.015
	(0.372)	(0.040)	(0.042)	(0.067)	(0.076)
Strongly disagree on RC	0.172	0.025	0.004	-0.272	0.000
	(0.377)	(0.048)	(0.051)	(0.084)	(0.000)
F-statistic		1.305	1.010	1.816	0.757
Comparison group		No incentive	No incentive	Husband incentive	RC incentive
Households in treatment	664	182	436	213	171
Households in comparison		664	664	223	180

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-RC incentive* varies from 0 to 1; it equals 1 if the individual incentive goes to the person whom both spouses agree is not the primary residual claimant. Households with no within-couple variation in residual claimant status are excluded from column 5.

Table 2: Average effects of all treatments

| | log (Quantity) | | |
	(1)	(2)	(3)
Couple incentive	-0.050	-0.050	-0.046
	[0.039]	[0.039]	[0.041]
Individual incentive	-0.059**		-0.055*
	[0.026]		[0.030]
Post-survey	-0.026	-0.026	-0.036
	[0.025]	[0.025]	[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)	0.829		0.828
Couple = Wife (p-val)		0.477	
Husband = Wife (p-val)		0.247	
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 3: Individual price incentive effects, by recipient payer status

| | log(Quantity) | | |
	(1)	(2)	(3)
Couple incentive	-0.071	-0.071	-0.071
	[0.044]	[0.044]	[0.044]
Individual incentive	-0.020	-0.026	-0.025
	[0.038]	[0.040]	[0.039]
Indiv incentive x Non-RC	-0.111**	-0.125**	-0.116**
	[0.047]	[0.052]	[0.052]
Incentive x Wife		0.024	0.009
		[0.049]	[0.049]
Total effect, incentive to non-RC	-0.131***	-0.150***	-0.141***
	[0.037]	[0.052]	[0.050]
Sample	Drop no variation	Drop no variation	Drop disagree
Observations (HH)	1,024	1,024	1,055
Observations (HH-months)	20,365	20,365	20,965

Notes: Individual incentive treatment arms interacted with heterogeneity variables: *Non-RC* varies from 0 to 1; it equals 1 if the individual incentive goes to the person whom both spouses agree is not the primary residual claimant. *Wife* equals 1 if the individual incentive goes to the wife. Households with no within-couple variation in residual claimant status are excluded from columns 1 and 2. Column 3 adds back in households with no variation in residual claimant status because both spouses say they share billing responsibilities. 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 4: Treatment effects by intrahousehold bill-sharing

| | log(Quantity) | | |
	(1)	(2)	(3)
Couple incentive	-0.019	-0.033	-0.034
	[0.051]	[0.060]	[0.060]
Couple incentive x Equal	-0.103	-0.122	-0.123
	[0.106]	[0.122]	[0.122]
Individual incentive	-0.059**	-0.078***	-0.075***
	[0.026]	[0.028]	[0.029]
Total effect, incentive to equal hh	-0.121	-0.155*	-0.156*
	[0.082]	[0.092]	[0.093]
Equal mean	0.315	0.316	0.313
Sample		Drop disagree	Drop no variation
Observations (HH)	1,282	1,055	1,024
Observations (HH-months)	25,506	20,965	20,365

Notes: Incentive treatment arms interacted with heterogeneity variables. *Equal* equals 1 if both the husband or wife indicates that they share responsibility on bill payments and equals 0.5 if either the husband or the wife indicates that they share responsibility on bill payment. Column 2 excludes households with strong disagreement on residual claimant status and column 3 excludes households with no variation in residual claimant status (to match sample restrictions in Table 3). 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.

Online Appendices

A.1   Model proofs

Set up

We follow the set up from the main text, which gives the following.

1. Individual water consumption: $w_i = \bar{w}(1 - e_i)$, where $\bar{w}$ is the individual satiation level of water use and $e_i \in [0, 1]$ is the individual's water conservation effort.

2. Household water consumption: $W = \bar{w}(1 - e_i) + \bar{w}(1 - e_{-i}) = 2\bar{w}(1 - \frac{(e_i + e_{-i})}{2})$.

3. Individual water conservation cost: $C(e_i, \mu) = ce_i^\mu$, with $\mu > 2$.

4. Household income: Y. Household price of water: pW, with $pW < Y$. Individual-specific price of water: $P_i W$.

5. Sharing rule, ex-post division of residual income from water bill: $\lambda_i \in [0, 1]$ and $\lambda_i + \lambda_{-i} = 1$.

6. Agents derive utility from residual income, after the water bill is paid, minus the cost of having conserved water: $U = (Y - pW) - ce^\mu$. In the presence of an individual price, this becomes $U = (Y - pW) - ce^\mu - P_i W$.

Prediction 1

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

Agent i's best response function takes e_{-i} as given. Each spouse chooses water consumption to maximize individual utility.

$$\max_{e_i} \lambda_i(Y - pW) - ce_i^\mu - P_i W$$

$$= \max_{e_i} \lambda_i(Y - p[2\bar{w}(1 - \frac{(e_i + e_{-i})}{2})] - ce_i^\mu - P_i[2\bar{w}(1 - \frac{(e_i + e_{-i})}{2})]$$

$$\text{FOC } e_i: \quad \lambda_i p 2\bar{w}\frac{1}{2} - c\mu e_i^{\mu - 1} + P_i 2\bar{w}\frac{1}{2} = 0 \tag{6}$$

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

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

Since w_{-i} is unaffected by a change in P_i, the household aggregate response to a change in P_i is the same as the change in w_i. In other words:

$$\frac{\partial W^*}{\partial P_i \partial \lambda_i} = \frac{\partial w_i^*}{\partial P_i \partial \lambda_i} + \frac{\partial w_{-i}^*}{\partial P_i \partial \lambda_i} = \frac{\partial w_i^*}{\partial P_i \partial \lambda_i}. \tag{7}$$

Individual i's response to a change in P_i is given by:

$$\frac{\partial w_i^*}{\partial P_i} = \frac{-\bar{w}}{\mu - 1} \left(\frac{\bar{w}}{c\mu}\right)^{\frac{1}{\mu-1}} \left(\lambda_i p + P_i\right)^{\frac{2-\mu}{\mu-1}} < 0 \tag{8}$$

The effect of λ_i is then obtained by differentiating $\frac{\partial w_i^*}{\partial P_i}$ with respect to λ_i:

$$\frac{\partial^2 w_i^*}{\partial P_i \partial \lambda_i} = \frac{-\bar{w}(2-\mu)p}{(\mu-1)^2} \left(\frac{\bar{w}}{c\mu}\right)^{\frac{1}{\mu-1}} \left(\lambda_i p + P_i\right)^{\frac{3-2\mu}{\mu-1}} > 0, \quad \forall \, \mu > 2 \tag{9}$$

When $\mu > 2$, the first term is positive, making the expression positive. Individual-price sensitivity decreases (change in consumption is less negative) as residual claim increases. ∎

Prediction 2

Prediction: $\frac{\partial^2 W^*}{\partial |\lambda_i - \lambda_{-i}| \partial p} > 0$, or equivalently, $\left|\frac{\partial W^*}{\partial p}\right|$ is decreasing in $|\lambda_i - \lambda_{-i}|$. In words, households with a smaller difference in λs are more responsive to changes in the household-level price.

Here, we are interested in how the effect of a change in the price on the water bill (p) depends on the difference in residual claim within the household. From the optimal consumption result in (6) and setting $P_i = 0$, individual i's response to a change in p is given by

$$\frac{\partial w_i^*}{\partial p} = \lambda_i^{\frac{1}{\mu-1}} \frac{-\bar{w}^2}{c\mu(\mu-1)} \left(\frac{p\bar{w}}{c\mu}\right)^{\frac{2-\mu}{\mu-1}} < 0. \tag{10}$$

By definition, for any $\lambda_i \in [0, 1]$, we have $\lambda_{-i} = 1 - \lambda_i$. So in each household, individual $-i$'s response to a change in p can be expressed as a function of her spouse's residual claim, λ_i:

$$\frac{\partial w_{-i}^*}{\partial p} = \lambda_{-i}^{\frac{1}{\mu-1}} \frac{-\bar{w}^2}{c\mu(\mu-1)} \left(\frac{p\bar{w}}{c\mu}\right)^{\frac{2-\mu}{\mu-1}}$$

$$= (1-\lambda_i)^{\frac{1}{\mu-1}} \frac{-\bar{w}^2}{c\mu(\mu-1)} \left(\frac{p\bar{w}}{c\mu}\right)^{\frac{2-\mu}{\mu-1}} < 0 \tag{11}$$

Equations (10) and (11) lead to the following cross-partial derivatives, with respect to λ_i:

$$\frac{\partial^2 w_i^*}{\partial p \partial \lambda_i} = \lambda_i^{\frac{2-\mu}{\mu-1}} \frac{-\bar{w}^2}{c\mu(\mu-1)^2} \left(\frac{p\bar{w}}{c\mu}\right)^{\frac{2-\mu}{\mu-1}} < 0, \quad \forall \, \lambda_i \in [0, 1], \, \mu > 2. \tag{12}$$

$$\frac{\partial^2 w^*_{-i}}{\partial p \partial \lambda_i} = (1 - \lambda_i)^{\frac{2-\mu}{\mu-1}} \frac{\bar{w}^2}{c\mu(\mu - 1)^2} \left(\frac{p\bar{w}}{c\mu}\right)^{\frac{2-\mu}{\mu-1}} > 0, \quad \forall \ \lambda_i \in [0, 1], \ \mu > 2. \qquad (13)$$

Notice that (12) and (13) have opposite signs. This is because when the residual claim of one spouse increases, the other's must decrease; as individual i increases her residual claim (marginal increase in λ_i), individual $-i$'s residual claim falls by the same amount.

The household's aggregate price sensitivity is the sum of both individuals' price sensitivities. We are interested in how aggregate price sensitivities is affected by a change in λ_i; for $\lambda_i > 1/2$, an increase in λ_i reflects an increase in the absolute value of the difference between λ_i and λ_{-i}, i.e., an increase in inequality within the household. The household's price sensitivity's response to a change in λ_i is:

$$\frac{\partial^2 W^*}{\partial p \partial \lambda_i} = \frac{\partial^2 w^*_i}{\partial p \partial \lambda_i} + \frac{\partial^2 w^*_{-i}}{\partial p \partial \lambda_i}$$

$$= (12) + (13) = \frac{\bar{w}^2}{c\mu(\mu - 1)^2} \left(\frac{p\bar{w}}{c\mu}\right)^{\frac{2-\mu}{\mu-1}} \left((1 - \lambda_i)^{\frac{2-\mu}{\mu-1}} - \lambda_i^{\frac{2-\mu}{\mu-1}}\right) \qquad (14)$$

$$> 0, \quad \forall \ \lambda_i \in (1/2, 1], \ \mu > 2$$

Intuition. When the spouse with a higher residual claim increases her λ_i, the household's residual claim levels become less equal, and aggregate price sensitivity declines (becomes less negative). ∎

Comparison of individual-level to household-level price changes

Corollary: $\left|\frac{\partial W^*}{\partial P_i}\right|_{\lambda_i \in (\frac{1}{2}, 1)} < \left|\frac{\partial W^*}{\partial p}\right| < \left|\frac{\partial W^*}{\partial P_i}\right|_{\lambda_i \in (0, \frac{1}{2})}$. In words, the effect of a change in the household-level price falls between the effect of the individual-level price directed to the individual with smaller claim and the individual with the larger claim on savings on the household water bill.

To compare the effects of the couple and individual price treatments, we will calculate the difference (in levels) in total household water consumption after (i) a marginal increase in the couple-price p experienced by both spouses and (ii) a marginal increase in P_i experienced by either the high residual claimant (High-RC, $\lambda_i > 1/2$) or low residual claimant (Low-RC, $\lambda_i < 1/2$) in the household. For a marginal change in the couple or individual price, a bigger change in levels implies greater price sensitivity. Since the effect of a price increase on aggregate water use is negative, a higher level in one treatment implies a *smaller* effect of price on water use.

The household's aggregate response to a change in the price on the water bill (couple

price) is given by:

$$\frac{\partial W^*}{\partial p} = \frac{\partial w_i^*}{\partial p} + \frac{\partial w_{-i}^*}{\partial p} = \frac{-\bar{w}}{(\mu-1)}\left(\frac{\bar{w}}{c\mu}\right)^{\frac{1}{\mu-1}} p^{\frac{2-\mu}{\mu-1}}\left(\lambda_i^{\frac{1}{\mu-1}} + (1-\lambda_i)^{\frac{1}{\mu-1}}\right) \qquad (15)$$

The household's aggregate response to a change in the individual price is given by (8).

We compare consumption levels as the difference in optimal use after a marginal change in the couple price, W^{*p}, and optimal aggregate use after a marginal change in the individual price, W^{*P_i}.

$$W^{*p} - W^{*P_i} = \left(w_i^* + w_{-i}^* + \frac{\partial w_i^*}{\partial p} + \frac{\partial w_{-i}^*}{\partial p}\right) - \left(w_i^* + w_{-i}^* + \frac{\partial w_i^*}{\partial P_i}\right)$$

$$= \left(\frac{\partial w_i^*}{\partial p} + \frac{\partial w_{-i}^*}{\partial p}\right) - \frac{\partial w_i'^*}{\partial P_i} = (15) \text{ - } (8) \qquad (16)$$

$$= \frac{-\bar{w}}{\mu-1}\left(\frac{\bar{w}}{c\mu}\right)^{\frac{1}{\mu-1}}\left(\underbrace{p^{\frac{2-\mu}{\mu-1}}\left(\lambda_i^{\frac{1}{\mu-1}} + (1-\lambda_i)^{\frac{1}{\mu-1}}\right)}_{Couple} - \underbrace{(\lambda_i p + P_i)^{\frac{2-\mu}{\mu-1}}}_{Individual}\right)$$

We evaluate this expression both at $P_i = 0$ and $P_i > 0$ for λ_i above and below 1/2.

Couple vs. high residual claimant We start by comparing the effect of a marginal change in the couple price to a marginal change in the individual price delivered to the high residual claimant ($\lambda_i \in (1/2, 1)$).

Evaluate at $P_i = 0$.
Factor out common terms to arrive at[1]

$$W^{*p} - W^{*P_i} = \frac{-\bar{w}}{\mu-1}\left(\frac{\bar{w}}{c\mu}\right)^{\frac{1}{\mu-1}} p^{\frac{2-\mu}{\mu-1}}\left(\underbrace{\lambda_i^{\frac{1}{\mu-1}} + (1-\lambda_i)^{\frac{1}{\mu-1}}}_{Couple} - \underbrace{\lambda_i^{\frac{2-\mu}{\mu-1}}}_{HighRC}\right) < 0, \; \forall \; \lambda_i \in (1/2, 1), \; \mu > 2$$

$$\iff \left.\frac{\partial W}{\partial P_i}\right|_{\lambda_i \in (1/2,1)} < \frac{\partial W}{\partial p} \quad \blacksquare$$

(17)

Evaluate at $P_i > 0$.
Focus on terms inside the parentheses. When the two terms are equal, the treatments are

---

[1]To highlight the relevant case, notice that our prediction implies $W^{*p} - W^{*P_i} < 0$ (in levels, the optimal consumption after an increase in the couple-price is lower than after an increase in the high RC price). This requires the term inside the parentheses to be positive, which we achieve with $\lambda_i^{\frac{1}{\mu-1}} + (1-\lambda_i)^{\frac{1}{\mu-1}} - \lambda_i^{\frac{1}{\mu-1}} > 0 \iff \lambda_i^{\frac{1}{\mu-1}}(1 - \lambda_i^{-1}) + (1-\lambda_i)^{\frac{1}{\mu-1}} > 0, \; \forall \; \lambda_i \in [0.5, 1].$

equally effective. Find which values of (p, P_i) yield the equality result. Denote as $\widehat{P}_i$ the optimal value of

$$\underbrace{p^{\frac{2-\mu}{\mu-1}}\left(\lambda_i^{\frac{1}{\mu-1}} + \left(1 - \lambda_i\right)^{\frac{1}{\mu-1}}\right)}_{Couple} = \underbrace{\left(\lambda_i p + \widehat{P}_i\right)^{\frac{2-\mu}{\mu-1}}}_{HighRC}$$

(18)

$$\iff \frac{\widehat{P}_i}{p} = \left(\lambda_i^{\frac{1}{\mu-1}} + \left(1 - \lambda_i\right)^{\frac{1}{\mu-1}}\right)^{\frac{\mu-1}{2-\mu}} - \lambda_i \in [-1, 0], \ \forall \ \lambda_i \in (0.5, 1).$$

Intuition. The ratio $\widehat{P}_i/p$ that satisfies this equation is the one that makes the performance of the couple and High-RC treatments exactly the same. The individual-price that generates this equality is $\widehat{P}_i(p, \lambda_i, \mu)$, a function of three parameters: couple-price, residual claim distribution, and cost of conservation. From the equality level, increasing $\widehat{P}_i$ makes the High-RC treatment *less effective* than the couple treatment (more conservation in the couple arm). Equality in treatment outcomes happens with $P_i^*/p \in [-1, 0]$, implying $P_i^* < 0$. This means that there are no possible combinations of household characteristics (λ_i, μ) and *positive* individual-price P_i where the household as a whole would be more price-responsive to the High-RC than to the couple treatment. The couple treatment is always more effective. The High-RC would need to be in a condition where she is being individually paid to use water $(P_i < 0)$ for a shock in prices to result in the High-RC treatment delivering more response than the couple treatment.

Couple vs. low residual claimant The setup for this comparison is exactly the same as the one laid out above, except that now we evaluate the results with the individual price to the low residual claimant $(\lambda_i \in (0, 1/2))$.

$$W^{*p} - W^{*P_i} = \left(w_i^* + w_{-i}^* + \frac{\partial w_i^*}{\partial p} + \frac{\partial w_{-i}^*}{\partial p}\right) - \left(w_i^* + w_{-i}^* + \frac{\partial w_i^*}{\partial P_i}\right)$$

$$= \left(\frac{\partial w_i^*}{\partial p} + \frac{\partial w_{-i}^*}{\partial p}\right) - \frac{\partial w_i^*}{\partial P_i} = (15) \text{ - } (8)$$

(19)

$$= \frac{-\bar{w}}{\mu - 1}\left(\frac{\bar{w}}{c\mu}\right)^{\frac{1}{\mu-1}}\left(\underbrace{p^{\frac{2-\mu}{\mu-1}}\left(\lambda_i^{\frac{1}{\mu-1}} + \left(1 - \lambda_i\right)^{\frac{1}{\mu-1}}\right)}_{Couple} - \underbrace{\left(\lambda_i p + P_i\right)^{\frac{2-\mu}{\mu-1}}}_{LowRC}\right)$$

<u>**Evaluate at $P_i = 0$.**</u>
<u>Factor out common terms to arrive at[2]</u>

---

[2]Again, to highlight the relevant case, notice that the prediction implies $W^{*c} - W^{*h} > 0$ (in levels, the optimal consumption after an increase in the couple-price is higher than after an increase in the low-RC price). This requires the term inside the parentheses to be negative, which we achieve with $\lambda_i^{\frac{1}{\mu-1}} + \left(1 - \lambda_i\right)^{\frac{2-\mu}{\mu-1}} - \left(1 - \lambda_i\right)^{\frac{2-\mu}{\mu-1}} < 0 \iff \lambda_i \in [0, 0.5]$.

$$W^{*c} - W^{*h} = \frac{-\bar{w}}{\mu - 1}\left(\frac{\bar{w}}{c\mu}\right)^{\frac{1}{\mu-1}} p^{\frac{2-\mu}{\mu-1}}\left(\underbrace{\lambda_i^{\frac{1}{\mu-1}} + \left(1 - \lambda_i\right)^{\frac{1}{\mu-1}}}_{Couple} - \underbrace{\lambda_i^{\frac{2-\mu}{\mu-1}}}_{LowRC}\right) > 0, \ \forall \ \lambda_i \in (0, 1/2), \ \mu > 2$$

$$\iff \frac{\partial W}{\partial p} < \frac{\partial W}{\partial P_i}\bigg|_{\lambda_i \in (0,1/2)}$$

(20)

Evaluate at $P_i > 0$

Focus on terms inside the parentheses. When the two terms are equal, we achieve equality in treatment effectiveness. Find which values of p, P_i yield the equality result.

$$\underbrace{p^{\frac{2-\mu}{\mu-1}}\left(\lambda_i^{\frac{1}{\mu-1}} + \left(1 - \lambda_i\right)^{\frac{1}{\mu-1}}\right)}_{Couple} = \underbrace{\left(\lambda_i p + \widehat{P}_i\right)^{\frac{2-\mu}{\mu-1}}}_{HighRC}$$

(21)

$$\iff \frac{\widehat{P}_i}{p} = \left(\lambda_i^{\frac{1}{\mu-1}} + \left(1 - \lambda_i\right)^{\frac{1}{\mu-1}}\right)^{\frac{\mu-1}{2-\mu}} - \lambda_i \ \in \ (0,1], \ \forall \ \lambda_i \in (0, 0.5).$$

Intuition. Equality in treatment outcomes happens with $\widehat{P}_i/p \in (0,1]$, implying $0 < \widehat{P}_i \leq p$. This means that there are possible combinations of household characteristics (λ_i, μ) and *positive* individual-price P_i where the household would be more price-responsive to the couple rather than to the Low-RC treatment. The Low-RC treatment will only perform better than the couple if the effective individual-price is $P_i < \widehat{P}_i(p, \lambda_i, \mu) \leq p$. In other words, a marginal change in the individual price to the low residual claimant will be more effective than a marginal change in the couple price when the individual price lies below the couple price.

High residual claimant vs. low residual claimant. The comparison between a marginal change in the individual price targeted to the high versus low residual claimant is given in prediction 1: an increase in the individual price to the low RC leads to a bigger reduction in aggregate water use than does an increase in the individual price to the high RC.[3]

Result: Summary We show that $\left|\frac{\partial W^*}{\partial P_i}\right|_{\lambda_i \in (\frac{1}{2},1)} < \left|\frac{\partial W^*}{\partial p}\right| < \left|\frac{\partial W^*}{\partial P_i}\right|_{\lambda_i \in (0,\frac{1}{2})}$ holds as long as $P_i < p$. This will be the case in any household water billing situation, where $P_i = 0$. ∎

---

[3]Using a similar approach to evaluating the individual prices that lead to equality in the treatments when $P_i > 0$, we find that equality occurs at $P_l = P_h + p(2\lambda_i - 1) \geq 0$, where P_l denotes the price to the low RC and P_h denotes the price to the high RC. That is, the only condition in which a change to the high RC's individual price has a larger effect is one in which the low RC already faces a (much) higher individual price.
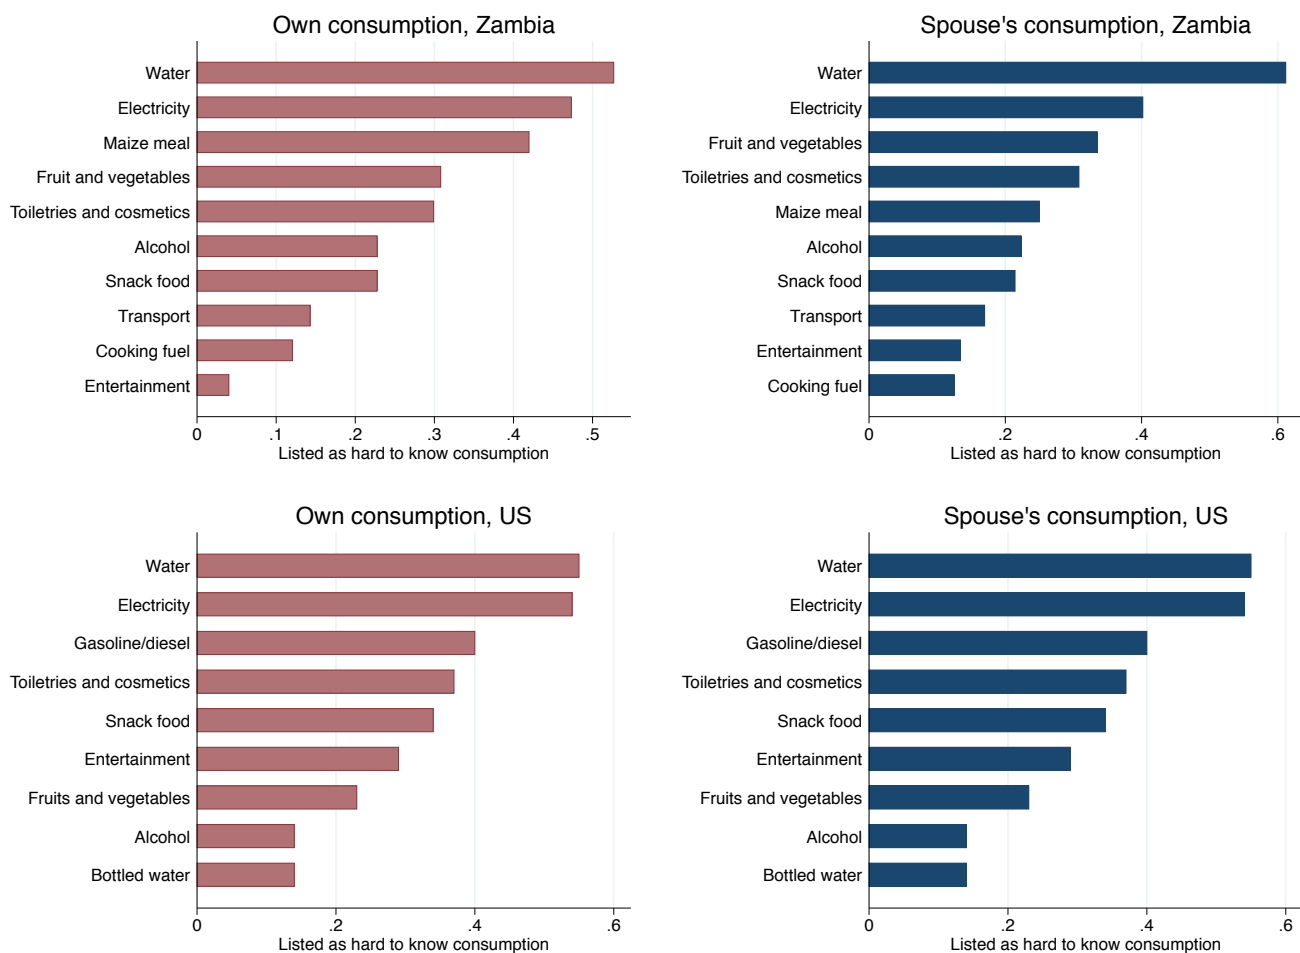
A.2   Appendix figures and tables

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).
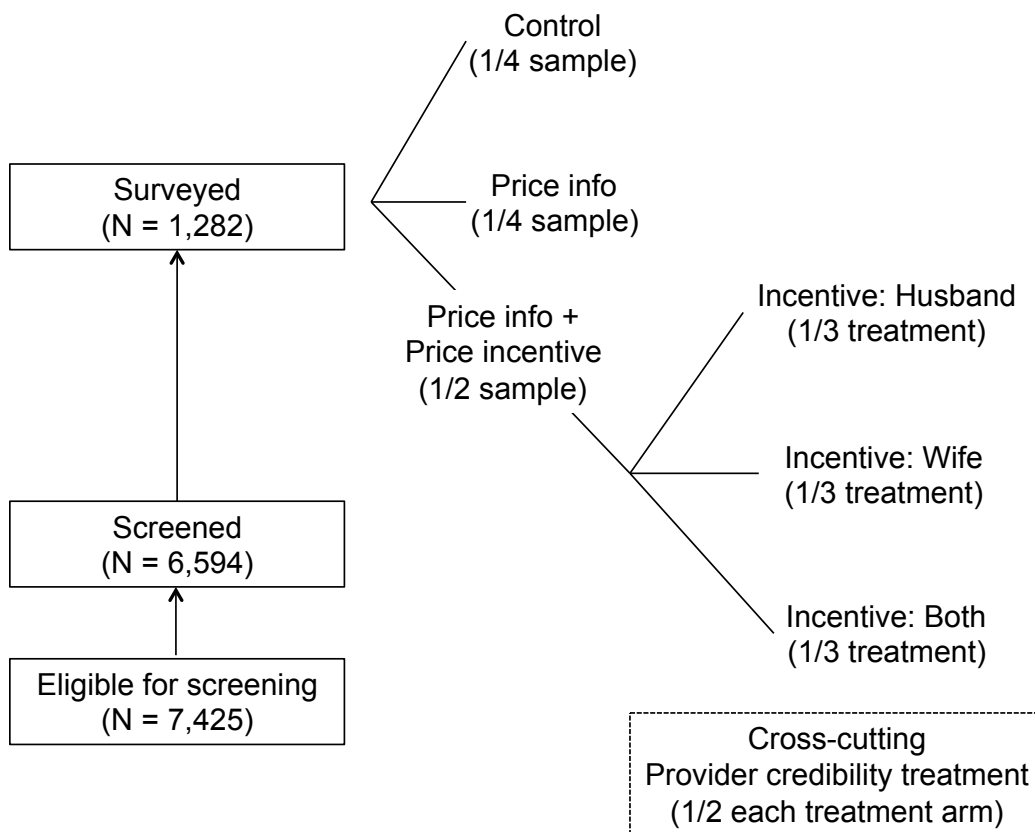
Figure A.2: Experimental design

Notes: Experimental design and sampling flow.

Table A.1: Residual claimant definitions, by spouse

Using payer variable			
	Husband's definition		
Wife's definition	Husband	Wife	Both/other
Husband	625	38	39
Wife	189	206	56
Both/other	73	25	31
Using income variable			
	Husband's definition		
Wife's definition	Husband	Wife	Both/other
Husband	621	1	111
Wife	17	30	46
Both/other	225	15	216

Notes: Residual claimant 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: Heterogeneous effects of price information and provider credibility treatments

	log (Quantity) (1)	log (Quantity) (2)
Price information treatment	-0.006 [0.048]	
Info x Underestimated price	-0.011 [0.060]	
Provider credibility treatment		0.018 [0.034]
Provider credibility x Distrust billing		0.024 [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.3: Robustness to controlling for individual characteristics

	log (Quantity) (1)	log (Quantity) (2)
Indiv incentive x Non-RC	-0.111**	-0.121**
	(0.047)	(0.056)
Indiv incentive x Over 50	0.086**	0.090
	(0.043)	(0.058)
Indiv incentive x Has regular employment	0.017	0.059
	(0.050)	(0.053)
Indiv incentive x Fluent in English	-0.015	-0.022
	(0.082)	(0.092)
Indiv incentive x Low education	-0.011	-0.031
	(0.048)	(0.059)
Indiv incentive x Uses more water	0.005	0.077
	(0.045)	(0.058)
Indiv incentive x Distrust billing	0.037	0.013
	(0.051)	(0.062)
Indiv incentive x Knows bill quantity	0.008	-0.024
	(0.047)	(0.058)
Indiv incentive x Knows bill price	-0.001	0.003
	(0.049)	(0.063)
Indiv incentive x High NGO sharing	-0.064	-0.066
	(0.047)	(0.046)
Observations (HH)	1,024	1,024
Observations (HH-months)	20,365	20,365

Notes: Robustness check on the results reported in column 1 of Table 3. *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. Households with no within-couple variation in residual claimant status are 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.

A.11

Table A.4: Robustness to different ways of defining non-residual claimant variable

			log(Quantity)			
	(1)	(2)	(3)	(4)	(5)	(6)
Individual incentive	-0.020 [0.038]	-0.025 [0.036]	-0.043 [0.039]	-0.019 [0.037]	-0.026 [0.037]	-0.027 [0.041]
Indiv incentive x Non-RC	-0.111** [0.047]	-0.080 [0.051]	-0.106** [0.048]	-0.111** [0.046]	-0.101** [0.047]	-0.098 [0.060]
Couple incentive						
Total effect, incentive to non-RC	-0.131*** [0.037]	-0.105*** [0.041]	-0.149*** [0.040]	-0.129*** [0.036]	-0.127*** [0.037]	-0.125*** [0.042]
RC definition	Main spec	Income variable	Drop intermediate RC	Husband definition	Wife definition	Equal definition
Observations (HH)	1,024	1,048	831	1,024	1,024	1,024
Observations (HH-months)	20,365	20,814	16,466	20,365	20,365	20,365

Notes: Incentive × Non-RC is the product of someone in the household having received the individual lottery and (1 minus) the RC status of the individual. Columns vary how the residual claimant variable is constructed relative to our main specification, which is shown in column 1. Column 2 uses the income variable if 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 residual claimant. Column 5 uses the wife's definition of residual claimant. Column 6 uses coding to match the intrahousehold billing equality used in Table 4. Households with no within-couple variation in residual claimant status are omitted in all columns. 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 Panel end	Jan 2014 Feb 2016 (1)	Mar 2014 Feb 2016 (2)	May 2014 Feb 2016 (3)	Jan 2014 2 mo post (4)	Mar 2014 2 mo post (5)	May 2014 2 mo post (6)	14 mo pre 2 mo post (7)
Individual incentive	-0.018 [0.038]	-0.020 [0.038]	-0.010 [0.037]	-0.017 [0.042]	-0.017 [0.042]	-0.008 [0.042]	-0.015 [0.041]
Indiv incentive x Non-RC	-0.120** [0.048]	-0.111** [0.047]	-0.103** [0.046]	-0.134*** [0.049]	-0.127*** [0.048]	-0.119** [0.048]	-0.100** [0.047]
Total effect, incentive to non-RC	-0.138*** [0.037]	-0.131*** [0.037]	-0.113*** [0.035]	-0.151*** [0.036]	-0.145*** [0.036]	-0.128*** [0.035]	-0.115*** [0.035]
Observations (HH) Observations (HH-months)	1,024 21,897	1,024 20,365	1,024 18,786	1,024 18,780	1,024 17,248	1,024 15,669	1,024 14,344

Notes: Robustness check on the results reported in column 1 of Table 3 (which is reproduced 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 2. 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),[4] 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.[5] 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 residual claimant 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.[6]

Households that met the screening criteria were informed about the survey. We scheduled a follow-up visit with the primary residual claimant 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 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

---

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

[5]They 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.

[6]Reasons 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.

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 (4), imposing the sample restriction in column 1 of Table 3, to calculate short-run price elasticities as follows.[7] First, with y_{it} equal to log of monthly water quantity, we can interpret the coefficient on $IndivTreat_{it}$ 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-residual claimant as a short run price elasticity. We observe a statistically significant 0.13 log point decrease in monthly consumption in response to treatment. For this sub-group, the average pre-intervention price is 4.88 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 5.85 cubic meters. The implied short run price elasticity is therefore -0.25.[8]

---

[7]We convert our treatment effects into elasticies 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.

[8]Our calculated short-run price elasticity of demand is within the ranges described by Dalhuisen et al. (2003) and 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

A.17

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.