

NBER WORKING PAPER SERIES

ENVIRONMENTAL EXTERNALITIES AND FREE-RIDING IN THE HOUSEHOLD

Kelsey Jack
Seema Jayachandran
Sarojini Rao

Working Paper 24192
<http://www.nber.org/papers/w24192>

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
January 2018

We thank the Southern Water and Sewerage Company for their collaboration on the project. We have benefited from comments made by audience members at numerous seminars and conferences, and from conversations with Nava Ashraf, Rebecca Dizon-Ross, Gilbert Metcalf, James Sallee, Dmitry Taubinsky, and Alessandra Voena. Flavio Malagutti, Lorenzo Uribe, Lydia Kim, and Alejandro Favela provided outstanding research assistance. We also thank Amanda Kohn for designing an information flyer used in the study. The International Growth Centre and J-PAL's Urban Services Initiative and Governance Initiative provided financial support. This RCT was registered in the American Economic Association Registry for randomized control trials under number AEARCTR-0000660. All errors are our own. The views expressed herein are those of the authors and do not necessarily reflect the views of the National Bureau of Economic Research.

NBER working papers are circulated for discussion and comment purposes. They have not been peer-reviewed or been subject to the review by the NBER Board of Directors that accompanies official NBER publications.

© 2018 by Kelsey Jack, Seema Jayachandran, and Sarojini Rao. All rights reserved. Short sections of text, not to exceed two paragraphs, may be quoted without explicit permission provided that full credit, including © notice, is given to the source.

Environmental Externalities and Free-riding in the Household
Kelsey Jack, Seema Jayachandran, and Sarojini Rao
NBER Working Paper No. 24192
January 2018
JEL No. D10,H21,H23,O10,Q56

ABSTRACT

Water use and electricity use, which generate negative environmental externalities, are susceptible to a second externality problem: with household-level billing, each person enjoys private benefits of consumption but shares the cost with other household members. If individual usage is imperfectly observed (as is typical for water and electricity) and family members are imperfectly altruistic toward one another, households overconsume even from their own perspective. We develop this argument and test its prediction that intrahousehold free-riding dampens price sensitivity. We do so in the context of water use in urban Zambia by combining billing records, randomized price variation, and a lab-experimental measure of intrahousehold altruism. We find that more altruistic households are considerably more price sensitive than are less altruistic households. Our results imply that the socially optimal price needs to be set to correct both the environmental externality and also the intrahousehold externality.

Kelsey Jack
Department of Economics
Tufts University
314 Braker Hall
Medford, MA 02155
and NBER
kelsey.jack@tufts.edu

Sarojini Rao
Department of Economics
University of Chicago
5757 S. University Avenue
Chicago, IL 60637
sarojini25@gmail.com

Seema Jayachandran
Department of Economics
Northwestern University
2211 Campus Dr
Evanston, IL 60208
and NBER
seema@northwestern.edu

1 Introduction

Water use and energy use generate environmental externalities. They are also subject to a second externality problem if the billing unit is a household: each individual enjoys the private benefits of consumption but shares the costs with other household members. (Or, conversely, each individual bears the private costs of conservation, but shares the savings with others.) In the absence of perfect altruism within the household or a way to ascertain each person’s usage, households will over-consume and be less price-sensitive, relative to their first best. The socially optimal price must be set to correct the environmental externality and this second externality within the home.

The degree to which household members internalize each others’ welfare likely varies across households, either because of variation in altruism or in the ability to observe and enforce individual levels of usage. When the this misalignment of incentives within the household is more severe, the household will be less price sensitive, all else equal. We test this prediction in the context of water usage in Livingstone, Zambia.¹ We survey customers of the regional water utility and combine monthly administrative billing data, a lab-experimental measure of intrahousehold altruism, and a randomized intervention that generates exogenous variation in the effective price of piped water. In our context, water is a significant expense for households: on average, nearly 5 percent of monthly income goes to the water bill.

We start by laying out the intrahousehold problem with a simple model that adapts the moral hazard in teams framework.² Each individual decides how much costly effort to put into water conservation, taking others’ decisions as given. Because the individual bears the full private cost of conservation while sharing the benefits (via the water bill) with the rest of the household, conservation will tend to be below the household’s Pareto optimal level. We allow for heterogeneity across individuals and households in how much they internalize other household members’ utility. We model this as altruism, but it could also be thought of as a reduced-form representation of monitoring and enforcement. As shorthand, we refer to households that internalize others’ utility more as “more efficient” because their intrahousehold free-riding inefficiency is smaller.

¹Surface water from the Zambezi River is the water source that supplies the local water utility and, in turn, households in Livingstone, and it is a scarce resource. Seasonal scarcity results in periods of water shortage in Livingstone, as in many developing country cities (NWASCO 2015). Use in Livingstone also imposes negative externalities on other who depend on the river, such as farmers downstream. Poor infrastructure resulting in distribution loss is one reason that water scarcity might be a larger problem in developing countries than developed countries, and increased weather volatility due to climate change is exacerbating the problem (Jacobsen et al. 2013; Van der Bruggen et al. 2010).

²As we discuss later in the paper, our framework does not assume that households are non-cooperative in general. The fact that individual water use is not observable and cannot be purchased individually precludes a cooperative equilibrium for the specific domain of piped water.

The model generates two main predictions that we test with our data. First, more efficient households will be more price sensitive. Second, individual-level prices lead to larger effects, relative to the status quo of household-level prices, for someone who is not usually responsible for the water bill and thus typically has weak incentives to conserve.

To test the predictions, we sample 1,282 married couples who are customers of the water utility, and survey each spouse separately.³ Each plays a modified dictator game with his or her spouse. Spouses will share more with each other either if they are more altruistic toward one another or are able to recoup more of the money they transfer to their spouse in the game. Both of these traits would also likely make individuals internalize each other’s utility when making water consumption choices, so we use this measure as a proxy for the degree of intrahousehold inefficiency in water use. We also ask other survey questions such as who bears the financial upside or downside if the water bill decreases or increases. (Hereafter, we refer to this person as the “effective bill payer.”)

To generate price variation, randomly selected households are given a financial incentive to conserve water. If household water use, as measured by the monthly bill, falls below a household-specific target, they are entered into a lottery with a 1 in 20 (or better) probability of winning a substantial payout. This reward for reducing consumption is akin to a discrete increase in the price of water over some range of consumption. In one third of the treated households, the prospect of this reward is conveyed to the couple together, and in the other two thirds it is conveyed to either the man or the woman only. This design effectively allows us to turn a household-level price into an individual-level price. We follow households for three to nine months after treatment and observe both water usage and bill payment behavior.⁴

Overall, we find an average decrease in monthly water use of 6.2 percent in response to the

³All of the households in our sample have piped water to their homes, and around 60 percent have indoor plumbing. Half own their homes, 16 percent employ a maid, and 90 percent have regular incomes from either salaried work or self-employment. In other words, these households are decidedly middle class by Zambian standards.

⁴We also investigate two other features of the billing environment that might affect households’ response to our price incentive treatment, described in detail in Section 3.3. First, water is priced on an increasing block tariff, which results in a poor understanding of the marginal price. We elicit price perceptions using a carefully framed procedure that elicits beliefs over quantities that are salient to water use decisions. While the median response is close to the truth, we observe a wide variance in beliefs across respondents and wide confidence intervals around individual beliefs. All households that receive the price incentive also receive information about the actual price of water. In addition, a sub-sample of household receive information about prices but no incentive to conserve. Second, a lack of trust in the water provider or a misunderstanding of the billing process might undermine customers belief that their water use translates directly into their bill. For example, when asked whose water use is to blame for a high monthly bill, many respondents placed blame on the water provider. We introduce a cross-cutting “provider credibility” treatment that explains how bills are generated. Neither the price information nor the credibility treatment result in measurable impacts on water use, even when prior beliefs about prices or the provider are taken into consideration.

financial incentive to reduce consumption. Consistent with the predictions from our model, the average response is driven almost entirely by more efficient households: couples that share more with each other in the modified dictator game — our proxy for intrahousehold efficiency — reduce their water use more in response to the price incentive treatment. In addition, it matters which member of the couple is given the incentive to conserve: incentives targeted at women lead to reductions in water use that are twice as large as either incentives targeted toward men or to couples, although the differences across incentive sub-treatments are not statistically significant. This greater sensitivity to a change in price targeted to the woman is consistent with our household model given the context: in the majority of households, men are the effective bill payers.⁵ We test this interpretation directly and find that, controlling for the gender of the person given the price incentive, a higher individual-level price leads to a larger reduction in household water use when given to the non-bill payer.

Our result on heterogeneity by the household’s level of intrahousehold efficiency uses randomized price variation but also relies on existing, non-randomized variation. This raises the concern that unobservables correlated with both intrahousehold efficiency (specifically, sharing in the dictator game) and other determinants of household-level price sensitivity might drive our findings. We address this concern in two ways. First, and most importantly, we develop multiple predictions about the types of households that should respond most to our price incentive intervention, including differential predictions based on which member of the couple (randomly) receives the incentive, a test that relies only on the design-based variation in the price incentive. Support for multiple predictions makes it less likely (though of course not impossible) that the patterns simply reflect omitted variables. Second, we investigate observable possible sources of heterogeneity in price sensitivity, such as household size or wealth, and show that our results are robust to controlling for them in parallel to our intrahousehold efficiency measure. While the line between controlling for spurious correlation and eliminating relevant variation is somewhat arbitrary, our main results largely stand up to a series of robustness checks.

After presenting our regression estimates, we discuss the normative and policy implications for corrective pricing. The welfare implications of our findings are analogous to those of Allcott et al. (2014), where correcting the environmental externality also corrects the intrahousehold “internality”, and therefore leads to an additional welfare gain and a higher optimal corrective tax. However, like in Taubinsky and Rees-Jones (2016), heterogeneity

⁵This finding raises the question of why households do not make the woman the bill payer in the absence of our intervention, given that she is reported to use more water than her husband in the majority of households in our sample. We discuss qualitative evidence on this puzzle in the conclusion.

in the degree of the intrahousehold free-riding lowers the welfare gain relative to the case with homogenous households. We calibrate the optimal corrective price in light of both the average distortion and the heterogeneity that we measure in our experimental setting. Only considering the average level of intrahousehold inefficiency, the optimal price would be over 60 percent higher than the optimal price set to internalize the marginal damages from consumption. This, however, would overshoot by roughly 40 percentage points the welfare-maximizing price when one takes into account the variation in the degree of intrahousehold efficiency across households. Given the welfare cost of a homogenous tax to correct a heterogenous intrahousehold distortion, it may be better to separately address the intrahousehold free-riding problem and the environmental externality. We show suggestive results that better information about consumption increases price sensitivity: households in which both spouses can read their water bills or are roughly aware of one another’s consumption are more price sensitive. Information interventions or individual-level incentives for conservation that address intrahousehold inefficiencies in water and energy use offer a promising direction for future research.

The study links two previously unconnected strands of literature, one on environmental externalities and one on intrahousehold decision-making. Our contribution to the literature on corrective pricing in environmental economics is to highlight a previously undiscussed reason that consumers might under-respond to utility prices.⁶ We thus add to literature on misperceptions of price (Ito 2014; McRae and Meeks 2016), lack of information about the price (Jesoe and Rapson 2014; Kahn and Wolak 2013), and lack of salience (Allcott 2011; Allcott et al. 2014) as factors that dampen the price elasticity of demand. The incentive mis-alignment within the household that we study resembles the incentive problem between landlords and tenants, which has been shown to lead to over-consumption of electricity (Levinson and Niemann 2004; Elinder et al. 2017) and underinvestment in efficiency (Gillingham et al. 2012; Myers 2015).⁷ Intrahousehold inefficiencies might be especially important in poor countries, where gender roles are particularly imbalanced. However, while our empirical applications focuses on husbands and wives, an analogous intrahousehold inefficiency arises from children, which is likely equally applicable in richer and poorer countries.

Our main contribution to the household economics literature is to study implications of intrahousehold decision-making for a novel domain of consumption, namely goods whose con-

⁶Rungie et al. (2014) point out that intrahousehold heterogeneity in preferences over water quality affect stated preference measurement of household-level preferences. They propose an alternative approach to discrete choice modeling of preferences that takes into consideration the influence of the individual on the collective household choice.

⁷A related problem exists for bill sharing among groups in other situations, such as dining at a restaurant (Gneezy et al. 2004).

sumption imposes environmental externalities beyond the household. Much of the previous literature either focuses on the implications of intrahousehold decision-making on investments in children (which might also be a source of societal externalities) or tests between different models of the household: unitary, cooperative non-unitary, and non-cooperative. A non-cooperative framework is applicable when, despite altruism, shared information and long-term interactions of households, there is limited information or limited commitment (see Lundberg and Pollak 1994). The imperfection in our context is that individual consumption of piped water is not easily observed (which also applies to home energy), combined with the fact that the good is delivered and billed to the household, not the individual.⁸ We contribute to a small set of papers showing Pareto inefficiencies in consumption or expenditure outcomes for households (examples include Dercon and Krishnan 2000; Duflo and Udry 2004; Mazzocco 2007; Robinson 2012; Angelucci and Garlick 2016).⁹ We also add to a literature that uses lab-experimental methods to measure differences across households in intrahousehold decision-making (e.g., Ashraf 2009; Mani 2011; Kebede et al. 2013; Castilla and Walker 2013). Specifically, we join a growing literature that tests for heterogeneity in the impact of interventions or in outcomes outside of the lab based on lab-experimental measures of a household’s decision-making (Schaner 2015; Hoel 2015; Ashour et al. 2017; Fiala 2017).

The paper proceeds as follows. The next section presents a simple model of water use in the household. Section 3 describes the experimental design and implementation. Section 4 presents the results, and Section 5 discusses the implications for optimal pricing. Section 6 concludes.

2 Model of water use within the household

In this section, we model a household’s water consumption, which is a function of effort spent on conservation. We start by benchmarking the household’s choice 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

⁸Respondents in a qualitative survey conducted in markets around Lusaka (N=96) were most likely to list water as one of the three most difficult consumption categories to keep track of. This was true both for own and spouse’s consumption. See Appendix Figure A.1.

⁹Many papers fail to reject efficiency in consumption (see Donni and Chiappori (2011) for a review). Efficiency in saving behavior is more often rejected than is efficiency in consumption, perhaps because of the greater commitment challenge posed by intertemporal savings decisions (e.g., Schaner 2015; Ashraf 2009). The timing of water and energy bills introduces an intertemporal dimension to the consumption decision, which may make them more prone to consumption inefficiencies than other types of consumption. We discuss the particulars around utility bills and intrahousehold decisions in Section 2.4. A much larger literature has examined, and often rejected, Pareto efficiency in investment and production (e.g., Udry (1996)).

effort. Second, water is not purchased at the individual level; a utility bill for piped water pools all household members' usage. We discuss these features of water in more detail at the end of this section. Because of these features, we model water use as a non-cooperative game. In the literature, households are more often modeled in a cooperative framework, befitting the altruism and long-term relationship among family members. Our model setup should not be interpreted as implying households are not cooperative over other domains characterized by greater observability of actions or individual-level purchases.

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

2.1 Optimal water conservation

Household aggregate water use, W , is the sum of water use by each individual i within the home, which is given by $w_i = \bar{w}(1 - e_i)$, where conservation effort $e_i \in (0, 1)$ lowers water use but at a convex cost, ce_i^2 . Individuals consume a maximum quantity of water given by \bar{w} if they exert no effort at all towards conserving water.¹⁰ The water utility charges the household pW , where $W \equiv \sum_i w_i$. The household has total income Y and we assume $pW < Y$. We assume that utility is linear in income remaining after the water bill is paid, i.e., 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:

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

Solving the first order condition, the household achieves its first best outcome if each member exerts effort, $e_i^{FB} = \frac{p}{2c}\bar{w}$.

2.2 Individual best response

The first best equilibrium might not obtain, however, if the conservation effort of the other member of the household, $-i$, is difficult to observe. We assume that each individual i takes

¹⁰This 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.

her spouse's conservation effort e_{-i} as given, assuming that e_{-i} is difficult to observe and therefore to contract over. We discuss this assumption in greater detail below.

Individual bargaining weights $\lambda_i > 0$ determine the ex post division of income Y that remains after the household pays the water bill. (In practice, households might have different sharing rules for different expenses. What is specifically relevant is the identity of the "effective bill payer," or residual claimant on the water bill, and the sharing rule he/she applies to the savings that accrue from water conservation.) Bargaining weights sum to 1 ($\lambda_i + \lambda_{-i} = 1$), and aggregate water use is given by $W = w_i + w_{-i} = 2\bar{w}(1 - \frac{e_i + e_{-i}}{2})$.

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

$$v_i = \lambda_i(Y - pW) - ce_i^2.$$

Individuals may also internalize some share $0 \leq \alpha_i \leq 1$ of their spouse's utility, with i 's utility function given by $u_i = v_i + \alpha_i v_{-i}$. We model and refer to α_i as a measure of i 's altruism toward his or her spouse, but it might also reflect enforcement of agreements around water use if individual water use were partially observable. Person i chooses e_i to satisfy the first order condition:

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

or, equivalently,

$$w_i^* = \bar{w}[1 - \frac{p}{2c}\bar{w}(\lambda_i + \alpha_i(1 - \lambda_i))] \quad (2)$$

For $\lambda_i = 1$ or $\alpha_i = 1$, person i fully internalizes the household's cost of water consumption, and the individual conservation decision is equal to the decision in the first best: $e_i^* = e_i^{FB} = \frac{p}{2c}\bar{w}$. However, if $\lambda_i = 1$, then $\lambda_{-i} = 0$, and individual $-i$ only exerts effort insofar as she is altruistic toward her spouse.

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

Finally, we consider a hypothetical individual price on water, P_i , such that person i effectively becomes the bill payer by bearing the full cost of water consumed by the household. i 's indirect utility function becomes $v_i = \lambda_i Y - P_i W - ce_i^2$, and her optimal $e_i^* = \frac{1}{2c}P_i\bar{w}$. Her incentive in this case is equivalent to her incentive under the household-level price if her

$\lambda_i = 1$.

2.3 Effect of a price change

Our experimental treatments are designed to make water use more costly to the household, effectively increasing the price. We are interested in how price sensitivity $\frac{\partial w_i^*}{\partial p}$ differs based on the inner workings of the household. Note that because $\frac{\partial w_i^*}{\partial p} < 0$, a negative cross-derivative represents an increase in price sensitivity.

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

This result states that individuals that are more altruistic are more price sensitive. We observe household level price responses, so the empirical prediction is that households that have higher average levels of altruism are more price sensitive. Intuitively, i will reduce her water use more when the price increases if she internalizes the savings that will accrue to her spouse more. Note also that the individual effect will be stronger as λ_i decreases: $\frac{\partial^2 w_i^*}{\partial p \partial \alpha_i} |_{\lambda_i \rightarrow 0} < \frac{\partial^2 w_i^*}{\partial p \partial \alpha_i} |_{\lambda_i \rightarrow 1} < 0$. This result that the marginal effect of α_i is greatest when λ_i is low means that altruism matters more for water use when the other person pays the bill.

Result: $\frac{\partial^2 w_i^*}{\partial p \partial \lambda_i} < 0$

In words, the greater the individual's stake in the water bill, the more sensitive she is to the price. Since we assume that $\lambda_i + \lambda_{-i} = 1$, there is no cross-household variation in the average value of λ to test this prediction. Instead, to identify how being the effective bill payer affects individual (and in turn household) water use, we add a person-specific component to the price, which we denote P_i . The individual utility function then becomes $v_i = \lambda_i(Y - pW) - ce_i^2 - P_iW$. The effect of such a manipulation depends on i 's existing incentive to conserve water, λ_i .

Result: $\frac{\partial}{\partial \lambda_i} \left(\frac{\partial w_i^*}{\partial p} - \frac{\partial w_i^*}{\partial P_i} \right) < 0$

The difference between the individual response to the household water price and to a person-specific price is smaller for someone with a larger residual claim over any savings on the water bill (higher λ_i). This is because this person already internalizes the household water price p , so a change in P_i represents a smaller proportional change in her effective price, and thus leads to a larger change in her water consumption. Conversely, as $\lambda_i \rightarrow 0$, the difference between offering an individual a price incentive through the household price p and an individual price P_i increases. When observing household-level water consumption, the prediction is that directing the individual price P_i to the individual with less stake in the bill will have a larger effect on aggregate consumption.

2.4 Discussion of assumptions

What makes water special A key feature of water consumption, or the choice of conservation effort, 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. Note that 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 from the cookie jar. If water meters were more accessible and easier to interpret, an individual could check the meter before and after a spouse’s shower to observe consumption.¹¹ Adding to these observability challenges, knowing one’s own consumption is often difficult.¹² Even ex post, if i can only observe own consumption with some error ϵ , then she can only infer w_{-i} from the total bill with error: $w_{-i} = W - (w_i + \epsilon)$. Moreover, the fact that some part of water consumption is a public good at the household level (e.g., washing the family’s dinner dishes) further complicates the problem of quantifying others’ effort toward conservation. (Note that even when water is used to produce public goods, there is still some “private” consumption if, conditional on how clean you get the dishes, washing them in a manner that wastes less water requires more effort and hence higher private costs.)

¹¹This 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).

¹²The 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 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.

Altruism versus enforcement Our modeling set up abstracts from enforcement and takes as given that the spouse’s water use is not observable. In practice, water use might be partly observable, in which case monitoring and enforcement of intrahousehold agreements becomes relevant. Even if water use were observable, difficulty enforcing intrahousehold agreements is sufficient to result in inefficient levels of aggregate consumption, and might lead to variation across individual-level α_i and household-level average $\bar{\alpha}$. Going forward, we refer to α_i and $\bar{\alpha}$ as measures of intrahousehold efficiency, to accommodate the possibility that either higher levels of altruism or better monitoring and enforcement might drive individuals within a household to consume closer to the household optimum. While we do not explicitly model the nature of the intrahousehold friction (i.e., what allows intrahousehold free-riding to persist), we conjecture that households in which water use is more observable will behave like households with higher $\bar{\alpha}$ and be more price sensitive.

3 Experimental design and data

An empirical test of our predictions requires three inputs: (1) a measure of aggregate household water consumption W , (2) variation in water prices p and P_i , and (3) measures of λ_i and α_i . We describe how we operationalize each of these in turn.

3.1 Water use

We partnered with the private regulated utility, the Southern Water and Sewerage Company (SWSC), that provides piped water to residents in Livingstone, Zambia. Households are billed based on monthly meter readings, and charged according to an increasing block tariff (i.e., the unit price is higher for usage beyond a threshold, and continues to increase in steps).¹³ Our main outcome measure is household water use per month. We obtain monthly billing records for January 2012 through September 2016. For each household in our sample, we create a panel that extends four months after treatment and 20 months before treatment, ensuring that observations for all households cover a two year window, regardless of when they were surveyed.¹⁴ Water use is measured in cubic meters based on in-person water meter readings collected monthly between the 20th and 25th of each calendar month. We keep only successful meter readings (i.e. drop the months in which meter readings were estimated or a

¹³Tariffs are regulated by the National Water Supply and Sanitation Council, and are intended to recover operating and maintenance costs, with cross subsidization from high to low tariff blocks and across customer types (NWASCO 2014).

¹⁴Households received up to 8 months of treatment, so we discard some treated months in favor of allowing all study households to contribute equally to the estimated treatment effect. As a robustness check, we include all treatment months in the analysis.

meter was reported as broken or disconnected). We log transform the consumption outcome variable, which drops a small number of zero reading months (which are likely billing errors or months the entire household was away, in any case). We generate an indicator for the month following a zero or missing observation to account for the fact that the first actual reading after an estimated reading or month with a broken meter may only partly reflect that month’s consumption.

The tariff schedule for 2015, when our study took place, is shown in Appendix Figure A.2. The average price in the pre-intervention period (2013-15) among households that we survey is 4.36 Kwacha, or around 0.44 USD, per cubic meter.¹⁵ Average household consumption is around 19 cubic meters per month, a little under half of typical US household consumption, resulting in monthly consumption charges of around 85 Kwacha or 8.50 USD per month.¹⁶ While we do not have household level monthly income or expenditure measures for our sample, we use the 2010 wave of the Living Conditions Monitoring Survey (LCMS), restricted to households with piped water in urban Livingstone, to calculate a median monthly expenditure of 192 USD (CPI adjusted to 2015 USD) and a median monthly income of 220 USD. Thus, the average water bill is around 4 percent of median income.

In addition to the measure of water consumption, we estimate the impacts of the intervention on other customer outcomes including payment behavior and missing meter readings. Our outcome measures and other relevant statistics related to the monthly bill are shown in the top panel of Table 1, which also tests for balance across the treatments, as discussed below.

3.2 Change in the effective price of water

The ideal variation to test the price sensitivity implications of our model would be a change in the marginal price of water. However, randomly varying the (regulated) water price charged by the utility, SWSC, was infeasible in our setting. Instead, we manipulate the household’s experienced water price by increasing the returns to water conservation through a randomized intervention. Figure 1 summarizes the experimental design.¹⁷

The treatment was implemented in conjunction with a household survey, run between May and December 2015. Households in the *incentive treatment* are provided with a monetary incentive to reduce water use. In the months following onset of the treatment, households are entered into a lottery for 300 Kwacha (30 USD) for reductions in billed consumption,

¹⁵We use an exchange rate of 10 Kwacha / USD and adjust for inflation to 2015 USD values throughout.

¹⁶Customers are charged for meter rental at a rate of 5 Kwacha per month and for sanitation and sewerage as a fixed proportion of monthly water use.

¹⁷The randomization was within four strata defined by whether the household’s pre-period average monthly water usage and outstanding balance due to SWSC were above or below our sample median.

relative to an account-specific reference level. Conditional on qualifying for the lottery, households had a 1 in 20 (or better) chance of winning. To qualify, households had to reduce their consumption by at least 30 percent relative to their average water usage in a two-month reference window, which resulted in a mean reduction target of 5.8 (median 4.95) cubic meters. The reference window was updated twice over the course of roughly eight months of fieldwork.¹⁸

In the notation of the household model, the incentive treatment adds a term to the indirect utility function: $v_i = \lambda_i(Y - pW + R \times \mathbf{1}(W \leq \bar{W})) - e_i^2$, where R is the expected value of the lottery payout. With 1 in 20 probability of winning, the expected value of the lottery incentive is around 1.5 USD or 15 Kwacha.¹⁹ The average price associated with the reference window is 5.1 Kwacha / cubic meter. With a target reduction of 5.8 cubic meters, the expected value of the (net) price shock is 2.06 Kwacha per cubic meter or a roughly 40 percent increase in the average price.²⁰ The treatment – in which a household receives a fixed reward for reaching a threshold that is proportional to past usage – differs from an increase in the marginal price of water in that there is a discrete change in the effective price of water over a particular range of consumption. In addition, the reward is an expected reward; we randomly select some households for payment to reduce implementation costs and simplify the logistics of paying the prizes. These design decisions were based on feasibility and transparency considerations, and we note that extrapolating from the treatment effects that we measure to price elasticities requires a number of assumptions. Rewriting our model in terms of a discrete change in the price associated with a quantity threshold does not change the predictions: households with higher $\bar{\alpha}$ are more likely to meet the threshold.

Our incentive treatment consists of three sub-treatment arms. In the first, both spouses learn about the lottery, and know that the information is provided to both. In this case, the intervention is analogous to an increase in p . The second and third sub-treatments provide only the wife or only the husband with information about the lottery. (Prize winners were also informed and paid privately in these arms). These individual sub-treatments move the

¹⁸The first reference period (March-April 2015) was used for households surveyed in May-early August; this was updated using a June-July reference window for households surveyed through late September. In September, we expanded the sample; new households had July-August 2015 as their reference period. This sample was used from late September through December, when fieldwork was completed.

¹⁹The probability of winning could be higher than 1 in 20 since we drew one winner for every 20 eligible households. Thus, if 21 households were eligible, we drew 2 winners. This was explained to households. At the same time, eligibility was made somewhat more difficult by the fact that the utility bills in round numbers; a household with a reduction target in fractions of a cubic meter would have to cut back to the nearest whole cubic meter to qualify.

²⁰Note that this calculation accounts for the increasing block tariff which causes the average price to fall from 5.1 to 4.6 Kwacha per cubic meter. We therefore calculate the “net” price shock after accounting for this mechanical reduction in the average price associated with lower consumption.

payoff from the lottery to outside of the λ_i term: $v_i = \lambda_i(Y - pW) + R \times \mathbf{1}(W \leq \bar{W}) - e_i^2$, so are analogous to an increase in P_i in the model. This increases i 's unilateral payoff from water savings, which has the greatest effect on overall household consumption if $\lambda_i < \lambda_{-i}$. Of course, individuals could share the information with their spouse or the spouse could find out about it, but the individual-specific treatment comes closer to an individual price than does the joint treatment.

3.3 Other experimental manipulations

We introduce two additional sources of variation through our experimental design. The first, which provides price information, acts both as a potential additional source of price variation and as a way to homogenize price beliefs. The second, which provides information about the credibility of the water provider, was intended to address misperceptions about the billing process.

3.3.1 Price information

We leverage the fact that most households are unaware of their marginal price to generate additional variation in prices. As part of the survey, we elicit price beliefs for both spouses, and then, in a *price information treatment*, provide accurate information about water prices. All incentive treatment households are also in the price information treatment. A challenge in communicating price information to households is that the units of consumption are difficult to map on to consumption. Based on extensive pre-testing, we both elicit price beliefs and provide price information in units of time spent using water rather than in cubic meters. Specifically, the survey asked “*suppose you wanted to save 20 Kwacha from your monthly bill; then, by how many minutes would your household as a whole have to reduce the use of the tap each day?*”²¹ Treated households received information that cutting back by 20

²¹The question text included clarifications that we meant running the tap at a normal rate, as they would for daily activities like washing their hands, and also that we were asking them to think about the minutes that the tap was running during the various activities they did, and not the overall time spent doing chores. If the respondent said they did not know and could not provide an estimate, the questions was repeated once and they were given a second chance to respond. If they were still unable to answer the question, we asked them about a series of narrowing intervals, e.g. less than 20, 20-40, 40-60, more than 60, and then given their chosen 20-minute interval, we asked about 5 minute intervals within it, and then re-asked the main price belief elicitation question. 81% of men and 83% of women answered the question the first time it was asked, and an additional 10% of respondents answered the question in the second or third attempt. We then asked them about the highest and the lowest that they thought the number of minutes could be, and then asked them again for a best guess, giving them a chance to revise their previous answer if they wanted. The price belief elicitation question was asked after a series of questions on their own and their spouse’s water use, as we found during piloting that thinking about water-intensive chores beforehand made it easier for respondents to understand the question. The price elicitation module was piloted with almost 300 households, who were then excluded from our sample.

minutes per day would save the average household 20 Kwacha on their monthly bill.

The effect of the price information depends on prior beliefs about water prices. Specifically, if households fully update, then the “price treatment” associated with the information is just the difference between the true price and the prior. Thus, individuals with a prior below the true price receive a positive price shock and vice versa. We categorize individuals into beliefs above and below the price information, and construct a household level measure that equals one if either spouse underestimated the price. The price information intervention also serves to remove variation in price beliefs, which allows us to calculate price elasticities associated with the incentive treatment.²² Thus, all households enrolled in the incentive treatment also receive the price information treatment.

3.3.2 Provider credibility

In a cross-cutting *provider credibility treatment*, households were offered assurance that water bills are based on actual water use, to address a worry that both the lottery and price information treatments might be ineffective if households do not believe that bills reflect consumption. Treated households were given information about the timing of the billing cycle and how their bill is calculated in the event that a meter reader is unable to read the meter, either because it cannot be accessed (e.g., the gate is locked) or because it is too unclear to be read (e.g., due to condensation). The information was paired with a reassurance that the provider is committed to honest billing practices and tries to ensure that households are only charged based on their actual water usage.²³

²²Note that we report price elasticities because they are familiar units, but our tests of our hypotheses do not rely on this conversion. There are several caveats to converting our treatment into a change in price, such as the fact that we assume risk neutrality, the price change only applies to a certain range of consumption, and individuals might differ in whether they believe that the required reduction is feasible.

²³The script for the provider credibility treatment is as follows: “We have collected this information purely for research and will not share any details with SWSC. However, we want to provide you with a little bit of extra information about how SWSC calculates your bill. SWSC tries to ensure that bills are accurate by reading your meter monthly and using the amount of water consumption shown on your meter to calculate your bill. That is, the amount that you are charged is based on the amount of water you use. The meter readings taken this month measure your usage since the time when last month’s reading was taken. Once SWSC has collected all the readings for this month, this is used to calculate the bill that will be given to you next month. For example, when you received your water bill in March you were charged for the water your household used between the 21st of January and the 20th of February, roughly speaking. When you received your water bill in April, you were charged for the water your household used between the 21st of February and the 20th of March, and so on. If there are some months that they cannot get a meter reading, then you are charged an estimate based on your previous consumption, and they try to get meter readings again as soon as possible. Then the next time they read your meter, they adjust your bill for any over- or under-charges from the months when they were not able to do the reading. SWSC is taking measures to make sure that bills are fair and based on actual water usage. They are committed to honest billing practices.”

3.4 Intrahousehold measures

We measure proxies for λ_i and α_i through a household survey conducted separately (and simultaneously) for the husband and wife. A series of questions documents water use, bill payment responsibilities, and intrahousehold cooperation, enforcement and altruism. In addition, we conduct modified dictator games between spouses as part of the survey visit. Both spouses played the game concurrently and in private, and the game was run by a trained surveyor who was of the same gender as the respondent. The game proceeded as follows.

The respondent is asked to pick one of two sealed envelopes and open it; the envelope contains either 20 or 30 Kwacha, and respondents only learn the value of their own draw, not the distribution. The surveyor explains that the money is theirs to keep and that they will be asked to make decisions using this money, but that they are under no expectation to share the amount. The respondent is then asked how much money they would send to their spouse versus keeping herself, and separately, how much money she would send to a water conservation NGO versus keeping. Before asking for their responses, the surveyor informs them that in each case, any money the respondent chooses to send will be doubled by the experimenter. The surveyor clarifies that the two decisions are mutually exclusive since the recipient of any money sent will be randomly selected after the respondent has made their decisions, with equal probability on each outcome, and that the respondent cannot influence which recipient is chosen at the end. The random endowments as well as the random selection of recipient ensures that the respondent can conceal her own earnings from her spouse and, thus, her actions are based more closely on her own preferences rather than concern about retribution from her spouse. If the respondent chooses to send some money, and the spouse is randomly chosen as the recipient, the spouse will know how much money was sent (since both spouses played the same game), but not how much money the respondent started with and hence how much the respondent kept for herself. Similarly, the respondent can also choose to send nothing to the spouse and claim that the NGO was chosen as the final recipient. The surveyor explains these aspects of the game to the respondent and asks questions during the explanation to check for respondent comprehension, so that respondents know what information can and cannot be hidden from their spouse.

The game provides a revealed preference measure of α_i . The more the respondent sends to her spouse, the more she cares about his financial resources.²⁴ Sharing with the spouse

²⁴For simplicity, our model assumed that total utility is a linear combination of own and spouse's utils, which predicts that spouses should share either nothing or everything in the dictator game. The fact that most sharing amounts are interior solutions is consistent with the aggregator function being non-linear, e.g., total utility is Cobb-Douglas in own and spouse's utils.

also reflects enforcement-based income sharing within the household, which may or may not carry over to water use; if a respondent expects that she can get back the amount she shared with her spouse, she would share more. The game-based measure does not allow us to separate altruism and enforcement.²⁵

As a proxy for λ_i , or who effectively pays the household's water bill, both spouses are asked whose income is used to pay the water bill and who physically pays the bill (which is done in person in this context). If the respondent's answers to those questions match, that person is labeled as the effective bill payer. If they do not, then follow up questions ask about how much discretion the person making the physical payment has over savings on the bill. We define respondent-specific indicators based on the respondent's perception of his or her claim on any savings on the bill (i.e. in some couples, each spouse may believe that he or she is the effective bill payer).

We ask respondents to compare their own water use directly to that of their spouse. We define the woman (man) as the larger water user if both members of the couple indicate that she (he) uses more water than her (his) spouse. Unlike the effective bill payer variable, where individual perceptions drive the incentive to reduce water use, we are interested in identifying which member of the couple actually uses more water, so we require that spouses' answers agree to assign one as the bigger water user. We also construct an indicator for the "stereotypical" intrahousehold arrangement, in which the husband is the effective bill payer and the wife is the larger water user.

We measure the couple's knowledge about household and each other's water use, which is necessary but not sufficient for monitoring and enforcement to reduce the free-riding problem. Specifically, we ask whether the respondent looks at the water meter and also test their knowledge of their household quantity of consumption and the total charge on the household bill. We also ask each spouse to name the top three water-using activities of his or her spouse, and construct a measure of whether their response matches the spouse's self-reported main water-using activities.

Finally, in addition to survey questions on water use and intrahousehold decision making, the survey collects information on demographics and socioeconomic status, as well as attitudes toward the water utility.

²⁵It is possible that sharing in the game is also predictive of bargaining power over income within the home, which may be correlated with λ_i . In the analysis, we use survey questions to measure which spouse has higher λ_i and is the effective water bill payer.

3.5 Sample construction and summary statistics

3.5.1 Sample construction

Our sampling takes the universe of metered household accounts as provided by SWSC and imposes some restrictions based first on billing data, and then based on a short screening exercise that was conducted in the field. Using the panel of billing data for metered residential customers as of February 2015 ($N=9,868$),²⁶ we eliminate households that did not have a working meter for at least 3 out of the 4 preceding months. We also excluded households that use no water (i.e. are billed for zero cu.m.) in more than half of the preceding 4 months. Households with very low variation in usage over the preceding four months were considered to have possibly tampered with the meter or have a delinquent meter reader. 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. Households with consistently low usage were also excluded since they would be least able to adjust their water consumption in response to a price shock, and, moreover, reducing water use could be harmful, e.g., in terms of hygiene, to households using very little to begin with; 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 were also dropped since they could also have had malfunctioning meters, or may not be as responsive to price, and may also have been significantly more difficult to survey (because they were presumably very wealthy households or firms mislabeled as residential customers by SWSC). Finally we drop households with an extremely high outstanding balance with the water utility, 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 also important for sampling. Specifically, we require that the water meter not be shared with other households, that the primary bill payer be married (or cohabiting) and that both spouses live at that address, and that the household was in residence for at least the 4-month period prior to April 2015. We also exclude households 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

²⁶This number excludes roughly 250 households involved in a pilot of the project, who were deemed ineligible for the full study.

our screening criteria.²⁷ 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. Appendix Table A.1 shows how sample characteristics evolve at each stage of sampling and randomization.²⁸

Households that met the screening criteria were informed about the survey. We scheduled a follow-up visit with the primary bill payer and his/her spouse, emphasizing that we needed both of them to be present for the full survey. We also informed respondents they would be compensated 40 Kwacha (4 USD) for participating in the survey. At the scheduled time and date, a pair of surveyors (always a woman and a man) visited the screened household for a full survey. After a few preliminary demographic questions, husbands and wives were separated and surveyed individually in different rooms of the house. Enumerators elicited water price beliefs, asked for perceptions of own and family members’ water usage, and conducted the modified dictator game. After finishing their individual questionnaires, both surveyors and respondents met back together in a common room for the last survey questions, and to receive the price information (if applicable). We brought the couple back together to avoid any awkwardness that might arise from ending the survey immediately following the game transfers with the couple separated.

3.5.2 Sample statistics

Table 1 summarizes characteristics of the sample, including elicited price beliefs and attitudes toward the water provider, and tests for balance between the incentive treatment and the control group. (Note that we use the term control group to mean the control group for the incentive treatment, which includes some households that received the price information and/or provider credibility treatment). The top panel shows statistics from the water bill, and so we report means and standard deviations for the households that were screened out of the survey for completeness. The middle panel shows household characteristics gathered through the survey. Around half of the sample owns their own home, and the average house-

²⁷Reasons 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.

²⁸As a robustness check, presented in Section 4.5, we also estimate our results adding in households that we sampled but were excluded during the screening stage. These households were sampled using the same criteria as the households that were ultimately surveyed, but were screened out after the surveyors’ initial visit. This adds 5,312 households to our survey sample, which are not systematically different from the surveyed or treated households in terms of pre-survey consumption patterns. Because we rely on the date of the survey to define the treatment timing in the panel billing data, we define a fake treatment date for the households that were screened but not surveyed. For households that were screened on a day that produced at least one completed survey, we use that survey date. When that strategy is not feasible, we use the average lag between screening and surveying (7 days).

hold size is close to six. Around 16 percent employ a maid. Our analysis of intrahousehold decision making around water use that focuses on the husband and wife clearly simplifies the household dynamics given that there are on average four additional members, plus, in some cases, a maid. Levels of English language proficiency are high. Price beliefs are reasonably accurate; in about 60 percent of households, at least one of the two respondents underestimated the price. Distrust of the service provider is high: in over 40 percent of households, both spouses say that high bills are the fault of the provider, i.e. not because of high consumption.

3.5.3 Intrahousehold measures

Our main measure of intrahousehold efficiency comes from the respondents' incentive compatible decisions of how much to share with their spouse in the dictator game. We summarize the measure in the bottom panel of Table 1. Husbands send a larger fraction of their endowment to their wives than wives send to their husbands. Both spouses send a smaller share of their endowment to the water NGO than to their spouse. Table 1 also reports other water use measures associated with our theoretical predictions. In around 30 percent of households, the wife says she is the effective water bill payer. In about 80 percent of households, both spouses agree that the wife uses more water than her husband. Thus, in the typical household, the man is the effective bill payer (higher λ_i), and the woman is the bigger water user. In only around 36 percent of households are the incentives aligned for household water conservation, i.e. the effective bill payer is also the bigger water user.

The dictator game measure is correlated with a number of standard survey-based measures of intrahousehold cooperation (see Appendix Table A.2). Households in which respondents indicated that they decide on budgeting or extra spending together, and in which they make plans together and stick to their plans also share more in the dictator game. Respondents saying they can prevent their spouses from deviating from plans or that they do things that their spouse wants them to do predicts dictator game giving as well. These answers could be interpreted as measures of either altruism or enforcement with some more intuitively related to one or the other.

Of course, measures of intrahousehold efficiency may be correlated with other household characteristics that affect both water use and price sensitivity. Table 2 shows correlations between three dictator game outcomes – the share of the endowment sent by the husband, by the wife, and whether the average share sent was above median (our main measure of intrahousehold efficiency) – and individual and household characteristics. First, the share of endowment sent to the spouse is positively correlated with the share sent to the water NGO. While this may indicate that individuals who are more altruistic in general are also more

altruistic toward their spouses, it may also indicate some experimenter demand effect or confusion about the game, though we observe no correlation with a measure of social desirability bias (SDB score) constructed from responses to a shortened version of the Marlowe-Crowne questionnaire (Crowne and Marlowe 1960).²⁹ Second, neither of the other measures that generate predictions in our conceptual framework (effective bill payer status or who uses more water) are correlated with the dictator game measure, nor are variables describing knowledge of the bill. Third, our average dictator game measure is negatively correlated with household size, with age, and with home ownership, and positively correlated with employing a maid, household assets, number of rooms in the home and English language fluency. On the whole, wealthier respondents appear to share more in the dictator game, perhaps unsurprisingly. In our robustness checks, we revisit these variables to determine if they also are associated with differential responsiveness to our incentive treatment.

4 Results

4.1 Predictions

The experimental design and data collection described in the previous section allow us to test the following empirical predictions, associated with a financial incentive to conserve water:³⁰

Prediction 1: The incentive treatments decrease water consumption.

Prediction 2: The magnitude of the incentive treatment effect is increasing in intrahousehold efficiency, measured by $\bar{\alpha}$ (operationalized as the household’s average sharing in the dictator game).

Prediction 3: The individual-specific incentive treatment is more effective if it is offered to the individual who is not the effective bill payer, so who otherwise has weak incentives to conserve.

4.2 Estimation strategy

We use monthly outcome data before and after the intervention and estimate a difference-in-differences regression to quantify the treatment effects:

²⁹Social desirability bias might lead respondents to share more of their endowment if they think sharing is viewed favorably by the enumerator or researcher.

³⁰The theoretical predictions in Section 2 report the marginal change in water use with respect to a marginal price change. The predictions also hold for a discrete price change associated with a threshold quantity change. Note also that we derive predictions in Section 2 over water use levels while our empirical results test for effects on log water use. Rewriting the model in logs generates the same predictions.

$$y_{it} = \beta_1 \textit{treatedhh}_i + \beta_2 \textit{post}_{it} + \beta_3 \textit{treatedhh}_i \times \textit{post}_{it} + \epsilon_{it} \quad (3)$$

where $\textit{treatedhh}_i$ is a binary indicator for whether the household was assigned to the relevant treatment group and \textit{post}_{it} is a time-varying indicator that turns on for household i in the month after the survey. Note that \textit{post}_{it} varies across households and not just over time because the survey and treatment were rolled out over time. Treatments were delivered at the end of the survey visit, so \textit{post}_{it} also represents the post-intervention period. β_3 identifies the differential change in the outcome among treated households after the survey. Even though a household was only eligible for the lottery based on consumption in the first full billing cycle after the survey date, we set \textit{post}_{it} equal to 1 as of the survey date because it is possible the intervention had immediate effects. In our main specification, we drop the month in which the survey occurred since it is partially treated.³¹

To improve precision, we include neighborhood by time and household fixed effects in our preferred estimates. Defining $\textit{treat}_{it} \equiv \textit{treatedhh}_i \times \textit{post}_{it}$, we estimate:

$$y_{it} = \beta_1 \textit{treat}_{it} + \beta_2 \textit{post}_{it} + \tau_t + \eta_i + \epsilon_{it} \quad (4)$$

where τ are zone-month-year fixed effects (a zone is a neighborhood in Livingstone) and η_i are household fixed effects. In the presence of household fixed effects, β_1 identifies the treatment effect of interest, and β_2 captures any independent average difference in water use in the post period (which is possible if, for example, participating in the survey made even the control group more attentive to water conservation). We allow for arbitrary within-household correlation in water use over time by clustering standard errors at the household level. Because gaps in the panel are associated with meter disconnections and other meter reading issues, we add a time-varying indicator for months immediately following a missing observation to control for the fact that these months may record only a partial month of consumption. Our main predictions involve heterogeneity in the response to treatment by household type, so we interact \textit{treat}_{it} and \textit{post}_{it} with relevant household characteristics.

To illustrate magnitudes and as an input to our policy calibrations in Section 5, we use the estimates of β_1 associated with our incentive treatment in equation (4) to calculate short run price elasticities as follows.³² First, with y_{it} equal to log of monthly water quantity, we

³¹Note that because the billing cycle starts on the 20th of each month, our definition of month corresponds to the billing cycle, i.e. July runs from June 21 to July 20.

³²The 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 (smaller) certain payout from conservation.

can interpret the coefficient on $treat_{it}$ in the presence of household fixed effects as $\partial \ln(q)/\partial p$ or $\partial q/q \times 1/\partial p$, such that multiplying by the pre-intervention average price delivers a short run elasticity. We calculate customer specific average prices, accounting for the increasing block schedule and for inflation (Zambian consumer price index), in each pre-intervention month and use that as the basis for our subgroup-specific average marginal prices.³³

We show the exogeneity of treatment assignment to observable household characteristics in Table 1 for the incentive treatment and Appendix Tables A.3 and A.4 for the information and provider credibility treatments. We also plot average water use across treatment conditions and our dictator-game measure of intrahousehold efficiency in the months leading up to the survey (Appendix Figure A.3). Overall, we observe parallel trends by incentive treatment, by our binary measure of intrahousehold efficiency, and between study households and other customers in Livingstone.

The plot of average monthly consumption by our measure of intrahousehold efficiency (middle panel, Appendix Figure A.3) suggests that more efficient households actually consume slightly more than their less efficient counterparts. Our conceptual framework predicts the opposite, all else equal. As shown in Table 2, other covariates are correlated with our efficiency measure and may contribute to the higher average consumption among more efficient households. We regress household average pre-intervention consumption on our dictator game measure (column 1) and a vector of other household-level covariates (column 2) and show the correlations in Appendix Table A.5. Unconditional on other observables, we see significantly higher consumption among more efficient households. Conditional on observables, the coefficient shrinks and becomes insignificant (but remains positive). Our main predictions involve differences in how households respond to a price shock, but a parallel concern arises: if unobservables affect both our key sources of heterogeneity (dictator game sharing and who in the household is the effective bill payer) and price sensitivity, then we may attribute an effect of omitted variables to our measures of intrahousehold decision-making. We address this both through testing multiple theoretically motivated hypotheses and through robustness checks in Section 4.5.

4.3 Average treatment effects

We begin by comparing the distribution of consumption in the incentive treatment group and control group, normalized by household-specific average consumption in the incentive

³³Note that this approach to calculating elasticities does not impose assumptions about how households perceive the price change, only that households knew their pre-treatment price. We increase the likelihood of this latter assumption by including all incentive treatment households in the price information treatment. However, given that these treatments were implemented concurrently, if the price information treatment affected price perceptions, then past usage – which we use to calculate elasticities – is unaffected.

treatment reference window. Figure 2 shows a decrease in consumption across most of the distribution.³⁴ It is worth noting that there is relatively little mass in either the treatment or control group below the target level to be eligible for the prize, which was 70 percent of reference window consumption, and most of the effect is due to reductions that were not large enough to make a household eligible for the prize. On the one hand, this is surprising because if households could perfectly choose their consumption level, there would be bunching just below the target level. On the other hand, the difficulty of knowing one’s own and others’ water use makes the pattern less surprising. In fact, the continuity in the reductions suggests that households responded to the lumpy financial incentive treatment in a similar way as we would expect them to respond to a standard price increase.

Table 3 reports the average treatment effect of the incentive treatment, as well as the other two treatments (price information and provider credibility). Column 1 reports the difference-in-differences results from estimating equation (3), without household or time fixed effects. Columns 2, 3 and 4 add household, month-year and zone-month-year fixed effects, respectively. The main effect of treatment group indicators are small and insignificant, indicating that the randomization was balanced and pre-intervention consumption is similar across arms (column 1).

Our main coefficient of interest is on $Incentive \times Post$. We observe a statistically significant 6.2 to 6.7 percent decrease in monthly consumption in response to the incentive treatment, consistent with prediction 1 (laid out in Section 4.1): the incentive treatments decrease average water use. The implied short run price elasticity is -0.27 (column 4, based on an average pre-intervention price 4.36).³⁵ We observe no significant average effect from the other treatments. Going forward, we focus on the specification shown in Column (4), which includes household and neighborhood-by-time fixed effects. Hereafter, we define the treatment variables as time-varying and report results from estimating equation (4).

While we observe no average effect of the other treatments, the effects of the price information and provider credibility treatments should depend on respondents’ prior beliefs about water prices and about the correspondence between water use and bills, respectively. Thus, we supplement the analysis in Table 3 with specifications that (a) interact the price information treatment with an indicator for whether the husband and wife underestimated the price, on average, and (b) interact the provider credibility treatment with trust in the

³⁴Appendix Figure A.4 shows the distribution of consumption pre- and post-treatment for the price incentive treatment and the provider credibility treatment.

³⁵Our calculated short-run price elasticity of demand is slightly below the mean found in the literature reviewed by Dalhuisen et al. (2003) and in line with the short run elasticities summarized in Worthington and Hoffman (2008). In the literature, the long run elasticity is generally shown to be larger than the short run elasticity.

water provider. Both should lead to reductions in water use because they imply that the treatments increased the effective perceived water price. Table 4 shows that these treatments had no detectable impacts, even after taking heterogeneous responses into consideration. For the remainder of the paper, we focus on the incentive treatment. To increase power, we pool the pure price information treatment with the control group and 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.³⁶

4.4 Intrahousehold heterogeneity

4.4.1 Household price incentives

Table 5 shows the main test of the prediction that households with more intrahousehold altruism/less inefficiency should be more price elastic (prediction 2). We pool the incentive treatment arms and estimate the effect of the incentive on household water use, allowing the effects to vary with the portion of the dictator-game endowment spouses shared with each other on average (column 1) or by each spouse separately (column 2). As predicted, column 1 shows a larger reduction for households that sent above the median on average; the effect in the above-median efficient households is roughly four times larger than the effect among less efficient households. These coefficients correspond to an average short-run price elasticity among households with below-median dictator game contributions of -0.10 (based on a pre-intervention average price of 4.3 Kwacha for this sub-sample), while the total effect for above-median households implies an elasticity of -0.44 (based on a pre-intervention average price of 4.4 Kwacha for this sub-sample). Appendix Table A.10 shows the robustness of these results to alternate approaches to aggregating the dictator game measure.³⁷

Columns 2 and 3 show the differential effect of the incentive treatments by each spouse's dictator game sharing. Column 2 includes the full sample, while column 3 restricts the sample to households that follow traditional gender roles, in which the woman is the bigger water user and the man is the effective bill payer. This restriction omits 551 households in which either the woman is not the bigger water user or the man is not the effective

³⁶Our main outcome, water use, is noisy, and even after conditioning on household and region-by-time fixed effects, we have relatively low power to detect impacts of the incentive intervention that is our main focus. Pooling the treatments improves power, particularly around the estimation of heterogeneous treatment effects. Results that separate out effects by each treatment arm are shown in the Appendix, and mirror the tables presented here. Note that the nested, rather than cross-randomized, design of the information treatment and the price incentive treatment (that is, all households that received the latter also received the former) implies that any interaction effect between the two treatments is not identified.

³⁷It may also matter whether spouses have similar levels of altruism. We observe similar decisions within couples: for over half of the households we study, the difference in the share of the endowment sent by the husband versus the wife is less than 0.25, and for only 15 percent of households is it more than 0.5.

bill payer. While we are under-powered to predict triple interactions between treatment, bill-paying (i.e., income sharing) arrangements and individual level efficiency measures, our model predicts that the marginal effect of individual-level altruism will be decreasing with the degree to which the person is the effective bill payer. We examine this qualitatively by comparing the results in column 3, where the woman’s altruism is predicted to matter more because she is not the effective bill payer, with those in column 2, where the prediction is more ambiguous. Though imprecisely estimated, we observe that the coefficient on the interaction between the incentive treatment and the woman’s dictator game measure indeed increases in magnitude by nearly five-fold, while the interaction with the man’s dictator game measure is largely unaffected by this sample restriction.

4.4.2 Individual price incentives

We now turn to looking at the incentive sub-treatments separately, in which the wife, husband, or couple are informed about the prize for reducing water consumption. If the household acts as a unitary agent, then these sub-treatments should have identical effects, but if interests are not fully aligned in the household, then their effects could differ and depend on the recipient’s existing incentives to conserve water. Specifically, if the individual incentive recipient is not the effective bill payer, and so typically has weak claim on savings from water conservation, then the effect should be larger. The results shown in Table 6 breaks the effect down by treatment arm. Recalling the patterns of water use and effective bill payer status shown in the bottom panel of Table 1, prediction 3 implies that the lottery directed to the wife should have a larger effect, given that in the typical household she is the larger water user and is not the effective bill payer. The pattern of coefficients in Table 6 confirms that the most effective lottery sub-treatment is the one directed toward the wife. Interestingly, even though one might have thought that the husband could just reproduce the effects of the wife incentive arm by promising her the prize (or most of it) in the other two arms, it seems that either husbands did not think of this or a commitment problem may have prevented it. Importantly, however, we lack the precision to reject that the effect of the wife lottery is statistically different from either of the other two sub-treatments.

We examine the individual incentives further in Table 7, aiming to disentangle whether the effect is due to gender per se or the recipients’s status-quo incentive and scope to conserve water. Column 1 interacts the individual-specific incentive treatment with an indicator for whether it was given to the non-bill-payer in the household. We find that the individual price incentive led to significant reductions in water use if and only if directed to the non-bill-payer. Another potentially relevant factor is whether the recipient has more ability to achieve the reduction in water use through changes to his or her water use; the bigger water

user in the household has more ability to do so. Thus, directing the incentive to the larger user should also be more effective. Column 2 shows weak support for this prediction. Given the considerable correlation between these two factors (the woman is likely to be both the non-bill-payer and the bigger water user), we first run both interactions together in column 3. The interaction with non-bill-payer remains significant, while the interaction with bigger water user becomes very small and positive. Next, we simultaneously estimate the effect of directing the individual incentive to the non-bill-payer and to the woman; that is, we estimate the effect targeting the non-bill-payer, controlling for gender to determine if the latter effect is driven entirely by gender. It is not: the interaction with effective bill payer status remains marginally significant and similar in magnitude to column 1, while recipient gender per se does not seem to affect response to the conservation incentive. Overall, the results presented in Table 7 suggest that the existing arrangements around who has a claim to any savings from water consumption is an important determinant of the effectiveness of the price incentive. This raises the question of why households are not able to resolve this conflict themselves, which we return to in Section 6.

4.4.3 Altruism versus enforcement

As discussed in Section 2.4, households may have a smaller household inefficiency in water use either because spouses are altruistic toward one another or because they are able to monitor and enforce water use. Our primary measure of intrahousehold efficiency – giving in the dictator game – may reflect either or both of these explanations. We collected survey measures of intrahousehold decision making that may offer a more nuanced look into which aspects of intrahousehold decision-making are associated with the incentive treatment effect (see Section 3.4 and Appendix Table A.2 for correlations of these measures with the dictator game outcomes). Table 8 shows the results from replacing the dictator game measure in equation (4) with each of the intrahousehold survey measures one by one (column 1) then with the intrahousehold survey measures included as parallel interactions (column 2) and, finally, with the first two principal components of the intrahousehold survey measures as parallel interactions (column 3).³⁸ Overall, the survey measures have relatively little explanatory power and the couple of significant interaction terms go in the wrong direction. Including them in parallel interactions (columns 2 and 3) has little effect on the significance or magnitude of the dictator game measure. It seems that the real-stakes game captures intrahousehold dynamics better than survey questions, limiting our ability to disentangle whether altruism or enforcement matters more.

³⁸Appendix Table A.6 shows how the variables contribute to the principal components.

A necessary (but not sufficient) condition for enforcement is that water use is observable, so we next examine the role of observability. We use measures of the accessibility of water use data to investigate whether the potential for intrahousehold monitoring predicts the household response to the incentive treatment. Table 9 shows heterogeneous treatment effects based on four measures of the observability of water use in the home. Columns 1 and 2 show binary measures (verified by the enumerator) of whether both spouses can identify the consumption quantity and total charge on the bill, respectively. Column 3 shows a binary measure of whether spouses show above-median awareness of each others' water use.³⁹ Column 4 sums these three measures into a knowledge sum. The three individual measures of information availability are associated with greater price sensitivity (all imprecisely estimated) with the largest differential effects based on knowing the consumption quantity on the bill and knowledge of spouse's water use. These results provide some suggestive evidence that efficiency may be driven, in part, by the information needed for monitoring and enforcing agreements around water use.

4.5 Robustness checks

4.5.1 Interpretation of the heterogeneity results

As shown in Tables 2 and A.5, both the dictator game measure and average pre-intervention water use are correlated with certain observable household characteristics. To address the concern that these other characteristics could confound our estimates of heterogeneous effects by intrahousehold efficiency, we add in interactions of these characteristics with the incentive treatment indicator in Table 10, first one at a time (column 1), then all at once (columns 2), and finally using the first two principal components of the household survey measures (column 3).⁴⁰ While we do see some heterogeneity in price sensitivity by these other measures, our main coefficient of interest (*Sent above average to spouse*) decreases only modestly and remains the largest magnitude of any of the interactions in column 3.

Note that some of these covariates should not necessarily be interpreted as sources of spurious correlation. For example, households that are more altruistic in general (as measured by *Sent above average to NGO*) may also be more altruistic within the home. Larger households may have a harder time enforcing income-sharing agreements of any kind. Controlling for these (and other) variables may therefore eliminate "good" variation in our intrahousehold efficiency measure. The robustness check shown in Table 10 is therefore a conservative test of our main result.

³⁹See Section 3.4 for further details on these variables.

⁴⁰Appendix Table A.7 shows how the survey variables contribute to the principal components.

These robustness checks focus on the dictator game heterogeneity. Support for our model comes equally from the predictions of differential responses to the individual price incentives based on the effective bill payer arrangement within the home. Given that the result is not just driven by the gender of the recipient of the individual price incentive (column 4, Table 7), the effect of unobservables has to be considerably more subtle to deliver our results. Namely, unobserved characteristics of the individual who is not the effective bill payer, conditional on gender, would have to be associated with greater price sensitivity (at the individual not the household level, since that result depends on the individual incentive treatment arms).

4.5.2 Specification and outcomes

We test for sensitivity to our specification by varying the outcome, sample, panel length and aggregation of the dictator game measure. First, Appendix Table A.8 shows similar results if we estimate the regression model in levels rather than logs (column 1) and if we use log of the total bill amount, which includes service charges and debts (column 2). Second, both the main effect of the incentive treatment and the interaction with the above-median dictator game measure are similar when we add in households that were screened but not surveyed (columns 1 and 2) or include all treated months in the analysis (columns 3 and 4). The bottom panel of Appendix Table A.9 shows the effect on quantity in levels and on the total bill (in logs). Results in levels are similar to the main results in logs, and the bill total changes in accordance with the observed consumption changes. Third, extending the panel to include more than four months post-treatment increases the magnitude of the treatment effects (Appendix Figure A.5), but only a subset of our sample (those surveyed first) contribute to these effects. Finally, Appendix Table A.10 shows results for alternative ways of aggregating dictator game decisions across spouses.

We also test whether our observed consumption responses reflect other margins of household adjustment, namely bill non-payment or meter reader evasion. Appendix Table A.12 shows little effect of any of the treatments on a measure of whether the household made a payment toward their bill (columns 1 and 2) or missing meter readings (columns 3 and 4).

Finally, as discussed in footnote 36, we pool treatments to improve power. Appendix Tables A.11 and A.13 repeat the analyses that involve heterogeneous treatment effects, and include the information and provider credibility treatments interacted with the heterogeneity measures of interest. We note a couple of results that appear to be influenced by the decision to pool treatments. First, Appendix Table A.11 shows that some of the heterogeneity with respect to average dictator game giving is associated with the information treatment. Consequently, the effects that we present in the main table might be better interpreted as the effect of the combined price incentive and price information treatment. Second, Appendix

Table A.13 shows that the summary measure of information about water use in the home appears to interact with the credibility treatment. Specifically, water use increases by around 5.5 percent in response to the credibility treatment for a one unit increase in the knowledge score (relative to those scoring zero). That said, the inclusion of this additional interaction term decreases the magnitude of the interaction with the incentive treatment only slightly.

5 Implications for optimal pricing

In our context, a household’s water use generates an environmental externality beyond the household due to competing needs for the scarce water drawn from the Zambezi River. In this section, we use our empirical results to calibrate the optimal corrective pricing in the presence of such an environmental externality. Specifically, we calculate the optimal adjustment to the price of water using an ad valorem tax to correct for both the inefficiencies associated with the intrahousehold externality and the environmental externality, using the framework developed by Taubinsky and Rees-Jones (2016).⁴¹ The intuition for the adjustment builds on Diamond’s (1973) tax adjustment using the elasticity-weighted marginal externality (but here reflecting heterogeneous distortions to household-level demand responses, as in Allcott et al. (2014)). We incorporate Taubinsky and Rees-Jones’s (2016) insight that heterogeneity in the distortion generates an additional source of welfare loss if the tax adjustment is based only on the average consumer’s elasticity. We explicitly account for the heterogeneity in the intrahousehold distortion that we documented in the preceding sections.

To adapt our model primitives to Taubinsky and Rees-Jones’s framework, we make the individual conservation effort choice binary, $e \in \{0, 1\}$ where each member of the household chooses whether to exert effort $e_i = 1$ at cost c . We further simplify our setup by setting the bargaining power parameter $\lambda_i = \frac{1}{2}$, which is always the average λ_i at the household level. We denote the amount of water use if the individual exerts effort as $\underline{w} < \bar{w}$. Individual i chooses to exert effort if and only if $\frac{1}{2}(1 + \alpha_i)(Y - p(\underline{w} + E(w_j))) - c \geq \frac{1}{2}(1 + \alpha_i)(Y - p(\bar{w} + E(w_j)))$. Individuals have some expectation of their spouse’s water use w_j . More importantly, the payoff from exerting effort depends not only on c , \bar{w} and \underline{w} , but also on α_i because individual i responds to a price p as if it were $\frac{1}{2}(1 + \alpha_i)p$.

Consuming water generates an environmental externality with marginal damages denoted by χ (measured in units of the numeraire), and a social planner wishes to introduce a tax τ to address the environmental externality. If individuals were perceiving the price with no distortion, the planner would just set τ equal to marginal damages χ , and the efficiency

⁴¹To simplify the utility’s water pricing problem, we assume marginal cost pricing and ignore any fixed cost recovery by the utility through, for example, a monthly service charge.

loss associated with the externality would be reversed (assuming no other taxes or frictions in the economy). However, in our model, the household responds to a price $p + \tau$ as if it were $\frac{1}{2}(1 + \bar{\alpha})(p + \tau)$, which means the tax will not sufficiently address the externality and will need to be adjusted (unless, of course, $\bar{\alpha} = 1$). Taubinsky and Rees-Jones (2016) use the parameter θ to denote this wedge between the actual price and the effective price the household responds to (with the effective price lower due to the distortion). Thus, mapping our notation to theirs, we can define $\theta \equiv \frac{1}{2}(1 + \bar{\alpha})$, with a household of type θ thus responding to a corrective price $p + \tau$ as if it were $\theta(p + \tau)$.

To quantify θ for the optimal tax calculation, we use the estimated price elasticities from our experiment. We first use the amount shared in the dictator game to categorize households as either high or low θ (based on whether they have above- or below-median sharing in the game). We then use our regression estimates of the price elasticity in each of these two subgroups to quantify θ . The reason we do not just use the proportion shared in the dictator game to calculate θ is that, in the tax calculation, θ is specifically a measure of how much the friction (intra-household free-riding in our case) dampens the price elasticity of demand. Note that the elasticity estimates for the two subgroups tell us the relative distortion between them, but do not pin down the absolute level of θ ; pinning down the absolute level requires an estimate of the price elasticity absent any distortion. Thus, we make the additional assumption that $\theta = 1$ for the subgroup with high intra-household efficiency. In other words, we assume that our high θ households are free of intra-household distortions; this assumption, which we choose because it seems less arbitrary than choosing a specific value less than 1, means that in our calculations below, we are underestimating the average distortion. The θ for less efficient household types can then be calculated as the ratio of their price elasticity and the elasticity of high θ households.⁴² Note that this requires an additional assumption that households' distortion-free price elasticity is uncorrelated with our measure of intra-household efficiency.

The price elasticities for each subgroup implied by the quantity responses to our price intervention as well as the resulting θ parameters are shown in Table 11. The planner observes an empirical demand curve associated with the population average $\varepsilon_p = -0.27$.

⁴²As discussed in Section 4.2, we interpret the treatment effects of the price treatment as $\frac{\partial \ln(q)}{\partial p} = \frac{\partial q}{\partial p} \times \frac{1}{q}$ so that if we multiply by p (i.e., the average price prior to the intervention), we recover the elasticity. Here, instead, the demand function is $q = q(\theta p)$ and not just $q(p)$. As long as the conditional distribution of θ is independent of p , then $\frac{\partial \ln(q)}{\partial p} = q'(p) \times \theta \times \frac{1}{q}$. Multiplying both sides by p gives the elasticity on the left-hand-side in terms of θ . We interpret our point estimates for the two types of households as representing different responses to an identical exogenous price shock.

Planner's objective function

The optimal price can be derived by specifying the planner's objective function. Let the indicator \mathcal{I}_e denote conservation effort, v represent the constant marginal utility from water use, and D be the aggregate demand for water. In keeping with Taubinsky and Rees-Jones (2016), we make two further simplifying assumptions: (1) that terms of order $\tau^3 D_{pp}$ and higher are negligible; and (2) that χ , v , and θ are mutually independent. In addition, we assume that the value of public funds equals 1, which allows us to focus on correcting the environmental externality χ , and that there exist no other distortions within the household, which means that we can derive the optimal price for water by only considering the household's water use. The planner's objective function is then given by:

$$W(\tau) = \int_j [Y - (p + \tau)(\mathcal{I}_e \underline{w} + (1 - \mathcal{I}_e) \bar{w}) + (v - \chi)(\mathcal{I}_e \underline{w} + (1 - \mathcal{I}_e) \bar{w}) - c\mathcal{I}_e] + \tau D. \quad (5)$$

In the absence of any intrahousehold distortion, the planner sets $\tau = \chi$ and the optimal water price is $p^* = p + \chi$. However, in the presence of an intrahousehold friction θ , which varies across households, the optimal price, can be shown to be as follows:⁴³

$$p^* = (p + \chi) \frac{E(\theta)}{E(\theta)^2 + Var(\theta)}. \quad (6)$$

The optimal price given in equation (6) is increasing in $E(\theta)$ and decreasing in $Var(\theta)$.⁴⁴ In our two-type example, the effect of greater variance in θ can be illustrated by comparing the optimal adjustment term for each type. For high $\bar{\alpha}$ types, the optimal adjustment is just χ , so the optimal price based on the average demand distortion is too high, leading to a loss in consumer surplus. For low $\bar{\alpha}$ types, the optimal price based on the average distortion is too low; the adjustment is insufficient to fully correct their intrahousehold inefficiency. The more different the types, the more "off" the average adjustment is for either type.

⁴³This result matches Corollary 4 to Proposition 11 of Taubinsky and Rees-Jones (2016) Appendix B.1, except that we set the first term to zero (no deadweight loss of taxation), and the adjustment in our case is applied to the price and the tax, not just to the tax. See Taubinsky and Rees-Jones (2016) for the full derivation of these sufficient statistics for the welfare cost of heterogeneity.

⁴⁴If the cost of a higher water bill is unevenly split across individuals within the household, then the individual-level variance will be even higher than the household-level variance. Unfortunately, we lack the data needed to calibrate the variance at the individual level.

Calibration

Calibrating the optimal price in our setting is simply a matter of substituting the parameters in Table 11 into equation (6). The optimal price assuming homogenous consumers is:

$$p_{homogenous}^* = \frac{1}{E(\theta)}(p + \chi) = 1.64(p + \chi),$$

while the optimal price allowing for heterogeneous consumers is

$$p_{heterogeneous}^* = \frac{E(\theta)}{E(\theta)^2 + Var(\theta)}(p + \chi) = 1.17(p + \chi).$$

In other words, the standard corrective price will fail to address the additional inefficiency associated with intrahousehold frictions, while the optimal price based on the average intrahousehold friction will over-correct by roughly 40 percent. In this particular exercise, where we consider just two types of households, the optimal price is closer to what the planner would choose ignoring intrahousehold frictions than if she were to wrongly ignore the heterogeneity across households in the degree of intrahousehold frictions.

6 Conclusion

This paper highlights the importance of intrahousehold decision-making in the realm of consumption externalities that extend beyond the household. To analyze this problem, we combine water billing data with a measure of intrahousehold altruism and randomized variation in the effective price of water. Consistent with a simple model of household water consumption — in which an intrahousehold free-riding problem arises and households overuse water from their own perspective — we show that households in which spouses are less altruistic toward one another are also less responsive to an increase in the price of water; we generate the price increase by offering a financial prize for reductions in household water use. Also consistent with our model, targeting the conservation incentive to the individual within the household who normally has the least incentive to conserve water leads to a larger response.

This free-riding problem among spouses would exist even if men and women were perfect equals (just as moral hazard in teams exists with identical workers). However, the problem is exacerbated when there are traditional gender roles, with women doing most of the chores and men having more control over household finances. The husband-wife dynamic is likely more problematic in developing countries, with their greater degree of gender inequality (Jayachandran 2015). The welfare implications to households from overuse of water might

also be especially large in poor countries because utility bills constitute a larger share of total income. At the same time, our focus on husbands and wives illustrates a broader intrahousehold problem, for example with children being wasteful of water and energy, that is likely equally applicable in rich and poor countries.

One puzzle about our findings is why households are not adopting what seems like an obvious improvement: having women (or more precisely, the bigger water users) effectively pay most of the water bill so that they receive the financial upside of conservation. In our sample, only a third of households had made the bigger water user the effective bill payer. Thus, even given the constraints on observability of water use, households could do better than they are doing. To probe this issue, we conducted follow-up discussions with 40 households. Most stated that it had never occurred to them to implement this arrangement. Meanwhile, the husband often is giving the wife an allowance for groceries and other household spending. One conjecture for why households have not applied the same idea to water is that piped water is a new phenomenon for most of them. When women fetched water from water sources, they were the “effective bill payers”; wasting water meant more time spent fetching water. One possible intervention is simply to suggest to people this alternative financial arrangement for covering the utility bills. Indeed, in our qualitative interviews, when we asked respondents if they had ever thought of this arrangement, several who responded “no” then volunteered that it sounded like a good idea.

However, if there are information frictions, even the constrained first best outcome will entail water use above the household optimum and, in turn, further above the socially optimal level. The standard solution of corrective pricing remains applicable, but will need to correct both the environmental externality and this intrahousehold “internality” (Allcott et al. 2014). Moreover, following Taubinsky and Rees-Jones (2016), we show that the variance across households in the degree of intrahousehold inefficiency significantly dampens the welfare gains from a homogenous corrective price.

A different policy tack would be to try to ameliorate the intrahousehold constraints directly. For example, giving households better information about their household usage through smartphone apps with real-time data (as are available in many developed countries) would enable better monitoring; knowing total 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 water 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*.
- Ashour, M., T. Bernard, M. Hidrobo, and J. B. Hoel (2017). Productive inefficiencies in dairy farming and cooperation between spouses: Evidence from Senegal. *Working Paper*.
- Ashraf, N. (2009). Spousal control and intra-household decision making: An experimental study in the Philippines. *American Economic Review* 99(4), 1245–1277.
- Castilla, C. and T. Walker (2013). Is ignorance bliss? The effect of asymmetric information between spouses on intra-household allocations. *American Economic Review* 103(3), 263–268.
- Crowne, D. P. and D. Marlowe (1960). A new scale of social desirability independent of psychopathology. *Journal of Consulting Psychology* 24(4), 349.
- 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.
- Diamond, P. A. (1973). Consumption externalities and imperfect corrective pricing. *The Bell Journal of Economics and Management Science*, 526–538.
- Donni, O. and P.-A. Chiappori (2011). Nonunitary models of household behavior: A survey of the literature. In *Household Economic Behaviors*, pp. 1–40. Springer.
- Duflo, E. and C. Udry (2004). Intrahousehold resource allocation in Cote d'Ivoire: Social norms, separate accounts and consumption choices. *NBER Working Paper Series w10498*.
- 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.

- Fiala, N. (2017). Business is tough, but family is worse: Household bargaining and investment in microenterprises in Uganda. *Working Paper*.
- Gillingham, K., M. Harding, and D. Rapson (2012). Split incentives in residential energy consumption. *The Energy Journal* 33(2), 37.
- Gneezy, U., E. Haruvy, and H. Yafe (2004). The inefficiency of splitting the bill. *The Economic Journal* 114(495), 265–280.
- Hoel, J. B. (2015). Heterogeneous households: A within-subject test of asymmetric information between spouses in Kenya. *Journal of Economic Behavior and Organization* 118, 123–135.
- Ito, K. (2014). Do consumers respond to marginal or average price? Evidence from nonlinear electricity pricing. *American Economic Review* 104(2), 537–563.
- Jacobsen, M., M. Webster, and K. Vairavamoorth (2013). The future of water in African cities. Report, The World Bank.
- Jayachandran, S. (2015). The roots of gender inequality in developing countries. *Annual Review of Economics* 7(1).
- 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*.
- Kebede, B., M. Tarazona, A. Munro, and A. Verschoor (2013). Intra-household efficiency: An experimental study from Ethiopia. *Journal of African Economies*.
- Levinson, A. and S. Niemann (2004). Energy use by apartment tenants when landlords pay for utilities. *Resource and Energy Economics* 26(1), 51–75.
- Lundberg, S. and R. A. Pollak (1994). Noncooperative bargaining models of marriage. *American Economic Review* 84(2), 132–137.
- Mani, A. (2011). Mine, your or ours? The efficiency of household investment decision: An experimental approach. *Working Paper*.
- Mazzocco, M. (2007). Household intertemporal behaviour: A collective characterization and a test of commitment. *Review of Economic Studies* 74(3), 857–895.

- McRae, S. and R. Meeks (2016). Price perception and electricity demand with nonlinear tariffs. *Working Paper*.
- Myers, E. (2015). Asymmetric information in residential rental markets: Implications for the energy efficiency gap. *Working Paper* (246R).
- National Water Supply and Sanitation Council, Zambia (2015). Strategic plan: 2016-2020. Technical report, NWASCO.
- NWASCO (2014). Guidelines on tariff setting. Technical report, National Water Supply and Sanitation Council.
- Robinson, J. (2012). Limited insurance within the household: Evidence from a field experiment in Kenya. *American Economic Journal: Applied Economics* 4(4), 140–164.
- Rungie, C., R. Scarpa, and M. Thiene (2014). The influence of individuals in forming collective household preferences for water quality. *Journal of Environmental Economics and Management* 68(1), 161–174.
- Schaner, S. (2015). Do opposites detract? Intrahousehold preference heterogeneity and inefficient strategic savings. *American Economic Journal: Applied Economics* 7(2), 135–174.
- Taubinsky, D. and A. Rees-Jones (2016). Attention variation and welfare: Theory and evidence from a tax salience experiment. *NBER Working Paper Series*.
- Udry, C. (1996). Gender, agricultural production, and the theory of the household. *Journal of Political Economy* 104, 1010–1046.
- Van der Bruggen, B., K. Borghgraef, and C. Vinckier (2010). Causes of water supply problems in urbanised regions in developing countries. *Water Resources Management* 24, 1885–1902.
- Worthington, A. C. and M. Hoffman (2008). An empirical survey of residential water demand modelling. *Journal of Economic Surveys* 22(5), 842–871.

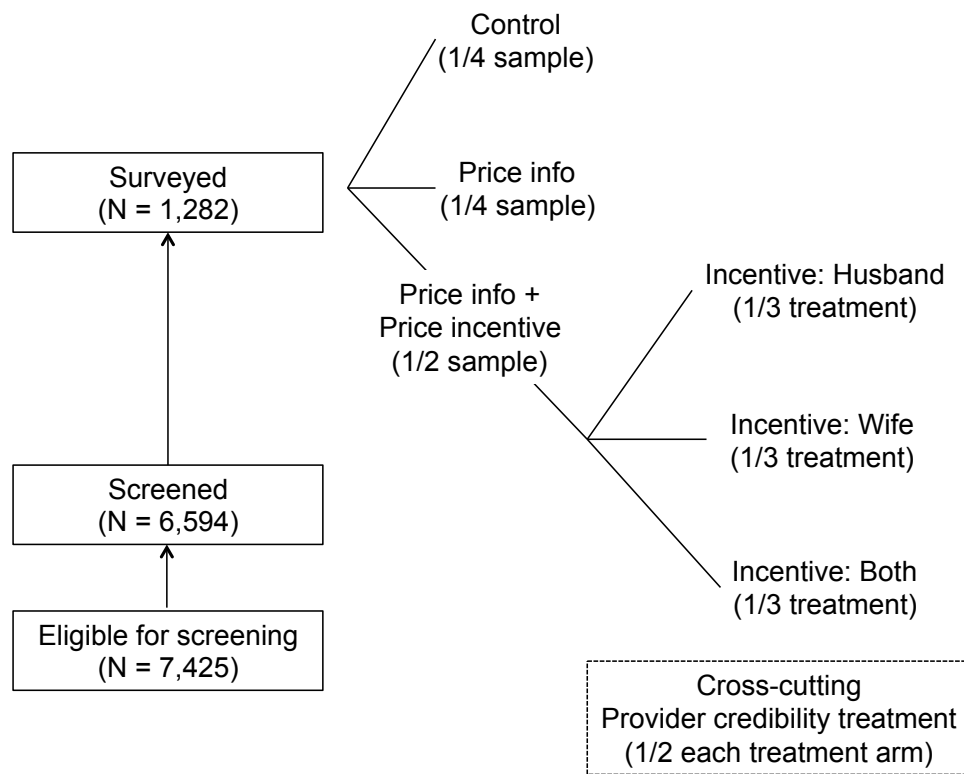


Figure 1: Experimental design

Notes: Experimental design and sampling flow. Treatment was assigned on a rolling basis to accommodate the high rate of ineligibility that led screened households to be disqualified from the survey sample.




Figure 2: Water consumption, relative to incentive reference month

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

Table 1: Sample statistics & balance

	Screened only (1)	No incentive (2)	Incentive (3)	P-val (2)=(3) (4)
Quantity consumed	20.940 (14.525)	18.995 (12.097)	18.247 (10.515)	0.239
Any payment	0.738 (0.195)	0.764 (0.166)	0.769 (0.166)	0.566
Missing meter reading	0.137 (0.188)	0.100 (0.157)	0.112 (0.170)	0.210
Total monthly bill	99.848 (88.152)	92.925 (69.044)	87.309 (60.949)	0.124
Household size		5.860 (2.286)	5.888 (2.218)	0.822
HH has maid		0.169 (0.375)	0.149 (0.356)	0.333
Owns home		0.512 (0.500)	0.495 (0.500)	0.546
Rooms in home		3.529 (1.264)	3.553 (1.444)	0.751
English fluency		0.768 (0.422)	0.772 (0.420)	0.873
Either underestimated price		0.619 (0.486)	0.637 (0.481)	0.551
Blame SWSC for high bill		0.440 (0.497)	0.414 (0.493)	0.356
Both know bill quantity		0.104 (0.305)	0.142 (0.350)	0.036
Both know bill charge		0.678 (0.468)	0.699 (0.459)	0.411
W: Effective bill payer		0.307 (0.462)	0.316 (0.465)	0.749
W: Bigger user		0.795 (0.404)	0.838 (0.369)	0.047
Share sent to spouse by husband		0.702 (0.269)	0.690 (0.254)	0.398
Share sent to spouse by wife		0.520 (0.262)	0.513 (0.260)	0.597
H: Share NGO		0.312 (0.253)	0.303 (0.232)	0.522
W: Share NGO		0.275 (0.222)	0.276 (0.220)	0.923
H: SDB score		19.938 (2.607)	20.010 (2.857)	0.640
W: SDB score		19.836 (2.838)	19.906 (2.999)	0.666
Households	5312	664	618	

Notes: Pre-treatment means for all households (top panel), and for surveyed households (middle and bottom panels). Column 1 is restricted to households screened out of the survey sample, column 2 to the sample that did not receive the incentive treatment, and column 3 to the sample that did receive the incentive treatment. Column 4 reports the p-value for a test of equal means between columns 2 and 3. The quantity consumed is measured in cubic meters per month. H and W refer to husband and wife. The share sent to the spouse is measured as a fraction of the respondent's endowment.

Table 2: Correlates of dictator game sharing

	Husband share sent (1)	Wife share sent (2)	Sent above median (3)
Quantity consumed	0.002*** (0.001)	0.002*** (0.001)	0.004*** (0.001)
Prob(payment)	-0.044 (0.045)	-0.113** (0.045)	-0.187** (0.086)
Prob(missing)	-0.064 (0.047)	0.079* (0.047)	0.051 (0.090)
Total charge	0.025** (0.013)	0.023* (0.012)	0.063*** (0.024)
Household size	-0.004 (0.003)	-0.009*** (0.003)	-0.015** (0.006)
HH has maid	0.019 (0.020)	0.068*** (0.020)	0.086** (0.038)
HH assets	0.008*** (0.003)	0.018*** (0.003)	0.030*** (0.005)
HH owns home	-0.015 (0.015)	-0.034** (0.015)	-0.060** (0.028)
HH rooms in home	0.014*** (0.005)	0.015*** (0.005)	0.034*** (0.010)
HH English fluency	0.022 (0.017)	0.082*** (0.017)	0.112*** (0.033)
Either underestimated price	0.009 (0.017)	0.017 (0.017)	0.033 (0.032)
Both blame high bill on SWSC	0.012 (0.015)	0.016 (0.015)	0.000 (0.028)
Both know bill quantity	0.005 (0.022)	0.014 (0.022)	0.046 (0.043)
Both know bill charge	0.010 (0.016)	-0.004 (0.016)	-0.004 (0.030)
W: Effective bill payer	0.003 (0.016)	0.001 (0.016)	0.013 (0.030)
W: Bigger water user	0.003 (0.019)	-0.006 (0.019)	-0.000 (0.036)
H: Share NGO	0.192*** (0.030)	0.079*** (0.030)	0.262*** (0.057)
W: Share NGO	0.034 (0.033)	0.198*** (0.033)	0.269*** (0.063)
H SDB score	0.004 (0.003)	0.002 (0.003)	0.006 (0.005)
W SDB score	0.003 (0.003)	-0.003 (0.003)	0.001 (0.005)

Notes: Each cell reports the coefficient from a separate regression of the dictator game measures (indicated in column headings) on a household characteristic (indicated in row headings). The share sent to the spouse and share sent to the NGO are measured as a fraction of the respondent's endowment. The SDB score measures social desirability bias using an adapted Crowne-Marlowe (1964) instrument. H and W refer to husband and wife.

Table 3: Average effects of all treatments

	log (quantity) (1)	log (quantity) (2)	log (quantity) (3)	log (quantity) (4)
Incentive x Post	-0.067** [0.030]	-0.062** [0.029]	-0.064** [0.029]	-0.062** [0.028]
Information x Post	-0.019 [0.034]	-0.011 [0.033]	-0.007 [0.033]	0.005 [0.032]
Credibility x Post	0.027 [0.026]	0.025 [0.025]	0.023 [0.024]	0.029 [0.024]
Surveyed x Post	0.194*** [0.028]	0.189*** [0.027]	-0.069** [0.035]	-0.104** [0.044]
Incentive treatment	-0.013 [0.041]			
Information treatment	0.007 [0.047]			
Provider credibility treatment	-0.011 [0.033]			
HH FE		x	x	x
Month-Year FE			x	
Zone-Month-Year FE				x
Observations (HH)	1,282	1,282	1,282	1,282
Observations (HH-months)	26,246	26,246	26,246	26,246

Notes: Regressions of log monthly quantity of water billed on treatment indicators. The panel begins 20 months (billing cycles) prior to the month of the survey and ends 4 billing cycles after the survey. Since treatment was provided at the time of the survey, we use the recorded survey date to define the treatment variables. The HH assignment indicator is time-invariant, while the post-survey indicator switches from 0 to 1 in the first full billing cycle after the date the household was surveyed; observations for billing cycles that contain the survey date are dropped. All households in the incentive treatment also received the information treatment. Standard errors are clustered at the household level.

Table 4: Heterogeneous effects of price information and provider credibility treatments

	log (quantity) (1)	log (quantity) (2)
Info treatment	-0.018 [0.050]	
Info x Underestimated price	-0.002 [0.060]	
Provider credibility treatment		0.011 [0.033]
Provider credibility x Distrust billing		0.042 [0.048]
Treatment + Interaction	-0.020 [0.044]	0.053 [0.035]
HH FE	x	x
Zone-Month-Year FE	x	x
Observations (HH)	1,282	1,282
Observations (HH-months)	26,246	26,246

Notes: Regressions include the post-survey indicator interacted with the heterogeneity variables. The incentive treatment indicator is excluded (treatments are pooled). 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. The bottom panel reports the linear combination of the treatment effect and the interaction term. Standard errors are clustered at the household level. Price beliefs are imputed for 257 households.

Table 5: Heterogeneous effects of incentive treatment by intrahousehold efficiency

	log (quantity) (1)	log (quantity) (2)	log (quantity) (3)
Incentive treatment	-0.022 [0.031]	-0.043 [0.033]	-0.058 [0.043]
Incentive x Sent above median on average	-0.080* [0.048]		
Incentive x Husband sent above median		-0.023 [0.050]	-0.031 [0.064]
Incentive x Wife sent above median		-0.024 [0.056]	-0.107 [0.073]
Treatment + Interaction	-0.102*** [0.037]		
Treatment + Husband Interaction		-0.066 [0.041]	-0.089* [0.054]
Treatment + Wife Interaction		-0.067 [0.053]	-0.165** [0.071]
Sample	Full	Full	Restricted
HH FE	x	x	x
Zone-Month-Year FE	x	x	x
Observations (HH)	1,275	1,275	726
Observations (HH-months)	26,122	26,122	14,799

Notes: Regressions include the post-survey indicator interacted with the heterogeneity variables. The information and credibility treatment indicators are excluded (treatments are pooled). *Shared above median* equals one if the share of the endowment transferred in the dictator game was above the median. Column 3 restricts the sample to households in which the woman is the larger water user and the man is the effective bill payer. The bottom panel reports the linear combination of the treatment effect and the interaction term. Standard errors are clustered at the household level. Dictator game outcomes are missing for at least one member of the couple in 7 households.

Table 6: Price incentives directed toward the wife, husband, or couple

	log (quantity) (1)	log (quantity) (2)
Couple incentive	-0.049 [0.039]	-0.046 [0.036]
Husband incentive	-0.046 [0.037]	-0.042 [0.034]
Wife incentive	-0.090** [0.036]	-0.089*** [0.033]
Couple = Wife (p-val)	0.350	0.337
Husband = Wife (p-val)	0.291	0.266
HH FE	x	x
Zone-Month-Year FE	x	x
Observations (HH)	1,282	1,282
Observations (HH-months)	26,246	26,246

Notes: Regressions include the post-survey indicator. Column 1 includes the information and credibility treatment indicators. Column 2 excludes them (pools them with the incentive and control conditions). The bottom panel reports the p-value for a test of equal coefficients for the incentive sub-treatments. Standard errors are clustered at the household level.

Table 7: Price incentive effects, based on whether recipient is effective bill payer

	log (quantity) (1)	log (quantity) (2)	log (quantity) (3)	log (quantity) (4)
Individual incentive	-0.022 [0.033]	-0.047 [0.032]	-0.019 [0.035]	-0.019 [0.036]
Incentive to non-bill-payer	-0.091** [0.041]		-0.087* [0.045]	-0.086* [0.044]
Incentive to bigger user		-0.044 [0.042]	-0.012 [0.045]	
Wife incentive				-0.011 [0.044]
Treatment + Payer Interaction	-0.113*** [0.034]		-0.106** [0.044]	-0.105** [0.046]
Treatment + User Interaction		-0.091*** [0.035]	-0.031 [0.046]	
HH FE	x	x	x	x
Zone-Month-Year FE	x	x	x	x
Observations (HH)	1,100	1,100	1,100	1,100
Observations (HH-months)	22,483	22,483	22,483	22,483

Notes: Regressions include the post-survey indicator interacted with the heterogeneity variables. The bottom panel reports the linear combination of the treatment effect and the interaction term. Standard errors are clustered at the household level. The couple incentive treatment arm is excluded.

Table 8: Heterogeneity by survey measures of intrahousehold decisions

	log (quantity) (1)	log (quantity) (2)	log (quantity) (3)
Incentive x Sent above median	-0.080* (0.048)	-0.083* (0.049)	-0.083* (0.049)
Incentive x Decide budget together	0.020 (0.048)	0.008 (0.052)	
Incentive x Decide extra spending together	-0.022 (0.062)	-0.058 (0.066)	
Incentive x Never disagree	0.048 (0.053)	0.044 (0.056)	
Incentive x Make plans together	0.068 (0.051)	0.073 (0.053)	
Incentive x Stick to plans	-0.027 (0.056)	-0.048 (0.059)	
Incentive x Respondent never deviates	0.001 (0.051)	-0.092 (0.068)	
Incentive x Spouse never deviates	0.105** (0.050)	0.151** (0.068)	
Incentive x Know if spouse deviates	-0.066 (0.050)	-0.100* (0.057)	
Incentive x Can prevent spouse from deviating	0.028 (0.049)	0.067 (0.054)	
Incentive x Does things spouse wants	0.011 (0.048)	0.046 (0.057)	
Incentive x Spouse does things respondent wants	-0.025 (0.047)	-0.013 (0.057)	
Incentive x Never hide income	0.064 (0.048)	0.058 (0.052)	
Incentive x Hard to hide income	-0.007 (0.050)	-0.007 (0.051)	
Incentive x 1st principal component	0.012 (0.015)		0.012 (0.014)
Incentive x 2nd principal component	0.004 (0.018)		0.009 (0.019)
Observations (HH)	1,275	1,275	1,275
Observations (HH-months)	26,122	26,122	26,122

Notes: Column 1 shows separate regressions in each cell, where each of the household level characteristics is interacted with treatment and the post-survey variable. Columns 2 and 3 each correspond to a single regression. Standard errors are clustered at the household level. The 7 households with missing dictator game outcomes are excluded.

Table 9: Heterogeneity by knowledge and monitoring of water use

	log (quantity) (1)	log (quantity) (2)	log (quantity) (3)	log (quantity) (4)
Incentive treatment	-0.056** [0.027]	-0.056 [0.045]	-0.039 [0.030]	-0.013 [0.042]
Incentive x Know bill quantity	-0.080 [0.068]			
Incentive x Know bill charge		-0.017 [0.053]		
Incentive x Know spouse's water use			-0.077 [0.051]	
Incentive x Knowledge sum				-0.040 [0.030]
Treatment + Interaction	-0.136** [0.063]	-0.072** [0.029]	-0.116*** [0.041]	-0.053** [0.024]
HH FE	x	x	x	x
Zone-Month-Year FE	x	x	x	x
Observations (HH)	1,282	1,282	1,282	1,282
Observations (HH-months)	26,246	26,246	26,246	26,246

Notes: Regressions include the post-survey indicator interacted with the heterogeneity variables. *Look at meter* equals one if both spouses report looking at their water meter. *Know bill quantity* and *Know bill charge* equal one if both spouses can identify the quantity and total amount owed on the bill, respectively. *Know spouse's water use* equals one if both spouses know above the median share of their spouses primary water using activities. The bottom panel reports the linear combination of the treatment effect and the interaction term. Standard errors are clustered at the household level.

Table 10: Robustness check: Controlling for household characteristics

	log (quantity) (1)	log (quantity) (2)	log (quantity) (3)
Incentive x Sent above median	-0.080* (0.048)	-0.091* (0.050)	-0.075 (0.049)
Incentive x Quantity	-0.002 (0.002)	0.002 (0.005)	
Incentive x Prob(payment)	0.050 (0.170)	-0.074 (0.163)	
Incentive x Prob(missing)	-0.167 (0.198)	-0.200 (0.219)	
Incentive x Total charge	-0.023 (0.040)	-0.011 (0.092)	
Incentive x Household size	-0.007 (0.010)	-0.006 (0.011)	
Incentive x Maid	-0.039 (0.066)	-0.001 (0.066)	
Incentive x Assets	-0.016* (0.009)	-0.017* (0.010)	
Incentive x Owns home	0.003 (0.048)	0.022 (0.053)	
Incentive x Rooms in home	-0.034* (0.019)	-0.033 (0.022)	
Incentive x English fluency	0.027 (0.059)	0.088 (0.068)	
Incentive x Underestimated price	-0.010 (0.046)	0.030 (0.045)	
Incentive x Blame SWSC	0.010 (0.048)	0.016 (0.048)	
Incentive x Know quantity	-0.060 (0.068)	-0.067 (0.070)	
Incentive x Know charge	0.012 (0.052)	0.039 (0.051)	
Incentive x W is effective bill payer	0.080 (0.049)	0.086* (0.049)	
Incentive x W is bigger user	-0.110* (0.066)	-0.123* (0.065)	
Incentive x Sent above median to NGO	-0.065 (0.048)	-0.038 (0.049)	
Incentive x Above median SDB score	0.044 (0.048)	0.052 (0.048)	
Incentive x 1st principal component	-0.016 (0.014)		-0.014 (0.014)
Incentive x 2nd principal component	0.008 (0.018)		0.004 (0.018)
Observations (HH)	1,274	1,274	1,274
Observations (HH-months)	26,105	26,105	26,105

Notes: Column 1 shows separate regressions in each cell, where each of the household level characteristics is interacted with treatment and the post-survey variable. Columns 2 and 3 each correspond to a single regression. Standard errors are clustered at the household level. The 7 households with missing dictator game outcomes are excluded.

Table 11: Parameters for calibration of socially optimal price

	ε_p	θ
average $\bar{\alpha}$	-0.27	0.61
high $\bar{\alpha}$	-0.44	1
low $\bar{\alpha}$	-0.10	0.23

Notes: Inputs for the welfare calibration for the average $\bar{\alpha}$ household, and for above and below median $\bar{\alpha}$ households. The price elasticities are calculated based on the observed response to the incentive treatment and the θ parameter is a ratio of elasticities, normalized by the observed elasticity for above-median $\bar{\alpha}$ households.

A.1 Model appendix

Derivations

We show the derivation of our three main results from the model, and the results both in terms of e_i^* and w_i^* .

Individual utility for person i is given by $u_i = \lambda_i(Y - pW) - e_i^2 + \alpha_i[(1 - \lambda_i)(Y - pW) - e_j^2]$ where $W = 2\bar{w}(1 - \frac{e_i + e_j}{2})$.

Substituting in for W gives

$$u_i = \lambda_i(Y - 2p\bar{w}(1 - \frac{e_i + e_j}{2})) - e_i^2 + \alpha_i[(1 - \lambda_i)(Y - 2p\bar{w}(1 - \frac{e_i + e_j}{2})) - e_j^2]$$

The first order condition with respect to e_i is

$$\begin{aligned} \lambda_i \cdot (-2p\bar{w}) \cdot (-\frac{1}{2}) - 2e_i + \alpha_i(1 - \lambda_i) \cdot (-2p\bar{w}) \cdot (-\frac{1}{2}) &= 0 \\ \implies e_i^* &= \frac{1}{2}\lambda_i p\bar{w} + \alpha_i(1 - \lambda_i)p\bar{w} \text{ or} \\ w_i^* &= \bar{w}[1 - \frac{1}{2}p\bar{w}(\lambda_i + \alpha_i(1 - \lambda_i))] \end{aligned}$$

Result 1

Our first result signs the cross partial with respect to α_i and p .

$$\begin{aligned} \frac{\partial w_i^*}{\partial p} &= -\frac{1}{2}\bar{w}^2(\lambda_i + \alpha_i(1 - \lambda_i)) \\ \frac{\partial}{\partial \alpha_i}[-\frac{1}{2}\bar{w}^2(\lambda_i + \alpha_i(1 - \lambda_i))] &= -\frac{1}{2}\bar{w}^2(1 - \lambda_i). \end{aligned}$$

Since $0 < \lambda_i < 1$ and $\bar{w} > 0$, $\frac{\partial^2 w_i^*}{\partial p \partial \alpha_i} < 0$.

Result 2

Our second result signs the cross partial with respect to λ_i and p .

$$\begin{aligned} \frac{\partial w_i^*}{\partial p} &= -\frac{1}{2}\bar{w}^2(\lambda_i + \alpha_i(1 - \lambda_i)) \\ \frac{\partial}{\partial \lambda_i}[-\frac{1}{2}\bar{w}^2(\lambda_i + \alpha_i(1 - \lambda_i))] &= -\frac{1}{2}\bar{w}^2(1 - \alpha_i). \end{aligned}$$

Since $0 \leq \alpha_i \leq 1$ and $\bar{w} > 0$, $\frac{\partial^2 w_i^*}{\partial p \partial \alpha_i} < 0$.

Note that the marginal effect of α_i is largest when λ_i is small. For $\lambda_i = 1$ and/or $\alpha_i = 1$, $e_i^* = \frac{1}{2}p\bar{w}$ and $w_i^* = \bar{w}(1 - \frac{1}{2}p\bar{w})$. As $\lambda_i \rightarrow 0$, $e_i^* \rightarrow \frac{1}{2}p\bar{w}\alpha_i$; as $\lambda_i \rightarrow \frac{1}{2}$, $e_i^* \rightarrow \frac{1}{4}p\bar{w}(1 + \alpha_i)$; and as $\lambda_i \rightarrow 1$, $e_i^* \rightarrow \frac{1}{2}p\bar{w}$.

Result 3

Our third result signs how the difference between $\frac{\partial w_i^*}{\partial p}$ and $\frac{\partial w_i^*}{\partial P_i}$ varies with λ_i .

$$\begin{aligned}\frac{\partial w_i^*}{\partial p} &= -\frac{1}{2}\bar{w}^2(\lambda_i + \alpha_i(1 - \lambda_i)) \\ \frac{\partial w_i^*}{\partial P_i} &= -\frac{1}{2}\bar{w}^2\end{aligned}$$

Since $\lambda_i + \alpha_i(1 - \lambda_i)$ is decreasing in λ_i for all $\alpha_i < 1$, $\frac{\partial}{\partial \lambda_i} \left(\frac{\partial w_i^*}{\partial p} - \frac{\partial w_i^*}{\partial P_i} \right) < 0$.

Semi-elasticities

Our empirical estimates are semi-elasticities, $\partial \ln(w_i^*)/\partial p$, while our model is in levels, $\partial w_i^*/\partial p$. Here we show that the model results are unchanged if we solve in semi-elasticities instead.

Result: $\frac{\partial^2 w_i^*}{\partial p \partial \alpha_i} < 0$ also holds for $\frac{\partial \ln(w_i^*)}{\partial p}$

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

In semi-elasticities: $\frac{\partial}{\partial \alpha_i} \left[\frac{\partial \ln(w_i^*)}{\partial p} \right] = \frac{\partial}{\partial \alpha_i} \left[\frac{1}{w_i^*} \frac{\partial w_i^*}{\partial p} \right] = -\frac{1}{2} \frac{\bar{w}^2}{w_i^*} (1 - \lambda_i) \left[1 + \frac{1}{2} \frac{\bar{w}^2}{w_i^*} p (\lambda_i + \alpha_i(1 - \lambda_i)) \right]$. All terms in this expression are positive. The negative sign implies that the cross-partial derivative overall is negative.

Result: $\frac{\partial^2 w_i^*}{\partial p \partial \lambda_i} < 0$ also holds for $\frac{\partial \ln(w_i^*)}{\partial p}$

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

In semi-elasticities: $\frac{\partial}{\partial \lambda_i} \left[\frac{\partial \ln(w_i^*)}{\partial p} \right] = \frac{\partial}{\partial \lambda_i} \left[\frac{1}{w_i^*} \frac{\partial w_i^*}{\partial p} \right] = -\frac{1}{2} \frac{\bar{w}^2}{w_i^*} (1 - \alpha_i) \left[1 + \frac{1}{2} \frac{\bar{w}^2}{w_i^*} p (\lambda_i + \alpha_i(1 - \lambda_i)) \right]$. Again, all terms in this expression are positive. The negative implies that the cross partial derivative overall is negative.

Result: $\frac{\partial w_i^*}{\partial P_i} |_{\lambda_i \rightarrow 0} < \frac{\partial w_i^*}{\partial p}$; $\frac{\partial w_i^*}{\partial P_i} |_{\lambda_i \rightarrow 1} = \frac{\partial w_i^*}{\partial p}$ also holds for $\frac{\partial \ln(w_i^*)}{\partial p}$

These are just first derivatives so if they hold in levels, they will also hold in semi-elasticities.

A.2 Appendix figures and tables




Figure A.1: Observability of consumption

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




Figure A.2: 2015 tariff schedule

Notes: Increasing block tariff for residential piped water in Livingstone. The price is shown in 2015 USD per cubic meter and is increasing in cumulative consumption over the course of the month.




Figure A.3: Pre-treatment water outcomes

Notes: Average log monthly water consumption, in the months preceding the survey. The top figure splits the surveyed sample into those receiving the incentive treatment and those not receiving the incentive treatment. The middle figure splits the surveyed sample into households with a measured level of efficiency below (low share sent in the dictator game) and above the median (high share sent). The bottom figure compares sample households with households screened but not surveyed.




Figure A.4: Consumption, relative to incentive reference months

Notes: Density plots of post-survey monthly consumption relative to the average monthly consumption in the reference months used to determine price incentive eligibility. The top figure compares households with and without price information. The bottom figure compares households with and without the provider credibility treatment. Note that the incentive treatment is excluded from these plots. The dashed vertical line shows the 70 percent threshold for lottery eligibility.




Figure A.5: Robustness: Different pre-treatment panel lengths

Notes: Coefficients on the interaction between the incentive treatment and an indicator for above median average share of the endowment sent to spouse in the household, as in Column 1 of of Table 5. Our main specification includes 20 months pre-treatment.

Table A.1: Balance at each sampling stage

	<i>Coefficient on:</i>		
	Successfully screened	Screened in	Surveyed
	(1)	(2)	(3)
Price information treatment	-0.015 [0.013]	-0.005 [0.075]	0.000 [0.000]
SWSC trust treatment	-0.008 [0.015]	-0.020 [0.086]	0.000 [0.000]
Lottery treatment	-0.023 [0.015]	0.084 [0.084]	0.000 [0.000]
Lottery to couple	-0.030*** [0.011]	0.058 [0.048]	0.000 [0.000]
Lottery to woman	0.000 [0.011]	-0.005 [0.065]	0.000 [0.000]
Lottery to man	0.007 [0.012]	0.031 [0.060]	0.000 [0.000]
Mean of regressor	0.173	0.973	1.000
Sample	Eligible for screening	Successfully screened	Screened in
Observations	7,391	1,282	1,247

Notes: Each cell corresponds to one regression, with each column representing a binary regressor. The sample for regressions in each stage is conditional on surviving the previous stage. See text for further detail.

Table A.2: Dictator game correlations with other intrahousehold measures

	Husband share sent (1)	Wife share sent (2)	Sent above median (3)
Decide budget together	0.035** (0.015)	0.034** (0.015)	0.054* (0.028)
Decide extra spending together	0.065*** (0.020)	0.048** (0.020)	0.138*** (0.039)
Never disagree	0.008 (0.016)	-0.006 (0.016)	0.008 (0.030)
Make plans together	0.036** (0.016)	0.039** (0.016)	0.074** (0.031)
Stick to plans	0.041** (0.018)	0.009 (0.018)	0.038 (0.034)
Respondent never deviates	0.006 (0.016)	-0.022 (0.016)	-0.017 (0.031)
Spouse never deviates	0.010 (0.015)	-0.001 (0.015)	0.001 (0.030)
Know if spouse deviates	0.012 (0.015)	0.014 (0.015)	0.037 (0.028)
Can prevent spouse from deviating	0.029* (0.015)	0.029** (0.015)	0.090*** (0.029)
Does things spouse wants	0.006 (0.015)	0.044*** (0.015)	0.057** (0.028)
Spouse does things respondent wants	0.019 (0.015)	0.024 (0.015)	0.053* (0.028)
Never hide income	-0.007 (0.015)	-0.020 (0.015)	-0.027 (0.028)
Hard to hide income	-0.013 (0.015)	0.021 (0.015)	-0.004 (0.029)

Notes: Coefficients are from separate univariate regressions of each dictator game measure on each intrahousehold survey measure. The share sent to the spouse and share sent to the NGO are measured as a fraction of the respondent's endowment.

Table A.3: Balance: Information treatment

	Screened only (1)	No info (2)	Info treatment (3)	P-val (2)=(3) (4)
Quantity consumed	22.471 (16.486)	19.913 (12.980)	19.336 (11.714)	0.405
Any payment	0.732 (0.193)	0.758 (0.164)	0.764 (0.164)	0.514
Missing meter reading	0.144 (0.181)	0.100 (0.153)	0.111 (0.158)	0.196
Total monthly bill	83.467 (81.532)	77.873 (61.084)	74.272 (55.673)	0.271
Household size		5.860 (2.286)	5.888 (2.218)	0.822
HH has maid		0.169 (0.375)	0.149 (0.356)	0.333
Owns home		0.512 (0.500)	0.495 (0.500)	0.546
Rooms in home		3.529 (1.264)	3.553 (1.444)	0.751
English fluency		0.768 (0.422)	0.772 (0.420)	0.873
Either underestimated price		0.619 (0.486)	0.637 (0.481)	0.551
Blame SWSC for high bill		0.440 (0.497)	0.414 (0.493)	0.356
Both know bill quantity		0.104 (0.305)	0.142 (0.350)	0.036
Both know bill charge		0.678 (0.468)	0.699 (0.459)	0.411
W: Effective bill payer		0.307 (0.462)	0.316 (0.465)	0.749
W: Bigger user		0.795 (0.404)	0.838 (0.369)	0.047
Share sent to spouse by husband		0.702 (0.269)	0.690 (0.254)	0.398
Share sent to spouse by wife		0.520 (0.262)	0.513 (0.260)	0.597
H: Share NGO		0.312 (0.253)	0.303 (0.232)	0.522
W: Share NGO		0.275 (0.222)	0.276 (0.220)	0.923
H: SDB score		19.938 (2.607)	20.010 (2.857)	0.640
W: SDB score		19.836 (2.838)	19.906 (2.999)	0.666
Households	5312	664	618	

Notes: See Table 1 for a description of the table set up. Column 1 shows averages for households screened out of the survey sample, column 2 for the survey sample that did not receive the information treatment, and column 3 for the survey sample that did receive the information treatment. Column 4 reports the p-value for a test of equal means between columns 2 and 3.

Table A.4: Balance: Provider credibility treatment

	Screened only (1)	No credibility (2)	Credibility treatment (3)	P-val (2)=(3) (4)
Quantity consumed	22.471 (16.486)	19.913 (12.980)	19.336 (11.714)	0.405
Any payment	0.732 (0.193)	0.758 (0.164)	0.764 (0.164)	0.514
Missing meter reading	0.144 (0.181)	0.100 (0.153)	0.111 (0.158)	0.196
Total monthly bill	83.467 (81.532)	77.873 (61.084)	74.272 (55.673)	0.271
Household size		5.860 (2.286)	5.888 (2.218)	0.822
HH has maid		0.169 (0.375)	0.149 (0.356)	0.333
Owens home		0.512 (0.500)	0.495 (0.500)	0.546
Rooms in home		3.529 (1.264)	3.553 (1.444)	0.751
English fluency		0.768 (0.422)	0.772 (0.420)	0.873
Either underestimated price		0.619 (0.486)	0.637 (0.481)	0.551
Blame SWSC for high bill		0.440 (0.497)	0.414 (0.493)	0.356
Both know bill quantity		0.104 (0.305)	0.142 (0.350)	0.036
Both know bill charge		0.678 (0.468)	0.699 (0.459)	0.411
W: Effective bill payer		0.307 (0.462)	0.316 (0.465)	0.749
W: Bigger user		0.795 (0.404)	0.838 (0.369)	0.047
Share sent to spouse by husband		0.702 (0.269)	0.690 (0.254)	0.398
Share sent to spouse by wife		0.520 (0.262)	0.513 (0.260)	0.597
H: Share NGO		0.312 (0.253)	0.303 (0.232)	0.522
W: Share NGO		0.275 (0.222)	0.276 (0.220)	0.923
H: SDB score		19.938 (2.607)	20.010 (2.857)	0.640
W: SDB score		19.836 (2.838)	19.906 (2.999)	0.666
Households	5312	664	618	

Notes: See Table 1 for a description of the table set up. Column 1 shows averages for households screened out of the survey sample, column 2 for the survey sample that did not receive the provider credibility treatment, and column 3 for the survey sample that did receive the credibility treatment. Column 4 reports the p-value for a test of equal means between columns 2 and 3.

Table A.5: Consumption correlates

	log (quantity) (1)	log (quantity) (2)
Sent above median to spouse	0.086*** (0.032)	0.003 (0.027)
Household size		0.029*** (0.006)
HH has maid		0.206*** (0.038)
HH assets		0.041*** (0.006)
HH owns home		0.038 (0.028)
HH rooms in home		0.146*** (0.011)
HH English fluency		0.073** (0.035)
Either underestimated price		0.002 (0.030)
Both blame high bill on SWSC		0.091*** (0.027)
Both know bill quantity		0.097** (0.041)
Both know bill charge		-0.031 (0.029)
W: Effective bill payer		0.012 (0.028)
W: Bigger water user		-0.033 (0.034)
Sent above median to NGO		0.026 (0.027)
Above median SDB score		0.018 (0.027)
Observations (HH)	1,275	1,274

Notes: Cross sectional regression of pre-intervention log monthly water use on an indicator for whether the households sent above median in the dictator game (column 1) and other observables (column 2).

Table A.6: Principal components: Intrahousehold survey measures

	Comp 1	Comp 2	Comp 3	Comp 4	Comp 5
Both say decide income use together	-0.031	0.440	-0.461	-0.065	-0.178
Both say decide extra money use together	0.008	0.405	-0.486	-0.132	-0.197
Both say they never disagree	0.328	0.067	0.009	0.204	-0.001
Both say they make plans together	-0.144	0.257	-0.260	-0.117	0.461
Both say it is easy to stick to plans	0.166	0.242	0.063	0.304	0.077
Both say they never deviate	0.424	-0.050	-0.047	0.224	0.364
Both say other never deviates	0.426	0.029	-0.162	0.132	0.373
Both say they know if spouse deviates	0.158	0.345	0.436	-0.412	0.028
Both say they can keep spouse from deviating	0.127	0.428	0.409	-0.328	0.139
Both say they do things for other	-0.387	0.242	0.164	0.425	0.209
Both say other does things for them	-0.391	0.290	0.177	0.384	0.130
Both say they never hide income	0.317	0.181	0.010	0.232	-0.129
Both say it is very difficult to hide income	0.193	0.180	0.189	0.322	-0.585
Eigenvalue	2.480	1.580	1.292	1.125	1.046

Notes: Factor loadings for principal component analysis of survey measures of intrahousehold decision-making.

Table A.7: Principal components: Control variables

	Comp 1	Comp 2	Comp 3	Comp 4	Comp 5	Comp 6	Comp 7	Comp 8
Quantity consumed	0.506	0.176	-0.070	-0.033	-0.121	-0.022	-0.065	0.134
Prob(payment)	0.024	0.253	-0.319	0.352	0.006	0.240	0.106	-0.419
Prob(missing)	-0.036	-0.080	0.626	-0.121	-0.275	-0.224	0.198	-0.037
Total charge	0.476	0.206	-0.261	-0.007	-0.043	0.076	-0.111	0.134
Household size	0.016	0.457	0.122	-0.324	0.138	0.213	0.200	0.102
HH has maid	0.326	-0.176	-0.033	-0.070	-0.024	-0.014	0.044	-0.050
Assets	0.365	-0.248	0.005	-0.123	-0.014	-0.037	0.067	-0.025
Owms home	-0.015	0.521	0.166	0.037	0.158	-0.276	-0.011	-0.062
Rooms in home	0.384	0.153	0.232	-0.101	-0.064	-0.168	0.001	0.100
English fluency	0.216	-0.424	-0.013	0.150	0.039	-0.083	-0.064	-0.067
Either underestimated price	0.016	-0.026	0.060	-0.337	0.566	0.185	-0.105	0.044
Blame SWSC for high bill	0.123	0.121	0.149	0.023	-0.368	0.239	0.330	-0.487
Know bill quantity	0.184	-0.075	0.246	0.426	0.212	0.090	0.050	0.050
Know bill charge	0.038	-0.033	0.354	0.295	0.110	0.535	-0.098	0.096
W: Effective bill payer	0.015	-0.051	-0.078	0.227	0.177	0.027	0.755	0.457
W: Bigger user	-0.066	-0.083	0.010	-0.297	-0.372	0.568	-0.093	0.268
Sent above median to NGO	0.162	-0.118	0.231	-0.140	0.409	0.135	-0.005	-0.413
Above median SDB score	-0.001	0.186	0.254	0.399	-0.084	-0.001	-0.415	0.230
Eigenvalue	2.849	1.581	1.255	1.192	1.151	1.060	1.017	0.971

Notes: Factor loadings for principal component analysis of control variables.

Table A.8: Robustness checks: Outcomes

<i>Outcome is:</i>	quantity		log (total bill)	
	(1)	(2)	(3)	(4)
Incentive treatment	-1.025** [0.470]	-0.375 [0.552]	-0.058** [0.023]	-0.020 [0.031]
Incentive x Sent above median		-1.422 [0.964]		-0.082* [0.047]
Treatment + Interaction		-1.797** [0.791]		-0.102*** [0.036]
HH FE	x	x	x	x
Zone-Month-Year FE	x	x	x	x
Observations (HH)	1,282	1,275	1,282	1,275
Observations (HH-months)	26,246	26,122	26,246	26,122

Notes: Regressions replicate the specification in column 3 of Table 3 and column 1 of Table 5. The outcome in columns 1 and 2 is the monthly quantity billed (in levels) and in columns 3 and 4 is the log of the total Zambian Kwacha owed on the bill.

Table A.9: Robustness checks: Sample

<i>Sample restriction:</i>	Including screened		All treated months		10 mo pre-treatment	
	(1)	(2)	(3)	(4)	(5)	(6)
Incentive treatment	-0.067*** [0.025]	-0.034 [0.032]	-0.057** [0.025]	-0.026 [0.032]	-0.042* [0.023]	-0.005 [0.031]
Incentive x Sent above median		-0.071 [0.050]		-0.064 [0.049]		-0.081* [0.047]
Treatment + Interaction		-0.105*** [0.038]		-0.089** [0.038]		-0.086** [0.036]
HH FE	x	x	x	x	x	x
Zone-Month-Year FE	x	x	x	x	x	x
Observations (HH)	6,594	6,587	1,282	1,275	1,282	1,275
Observations (HH-months)	129,899	129,775	27,828	27,697	16,235	16,157

Notes: Regressions replicate the specification in column 3 of Table 3 and column 1 of Table 5. Columns 1 and 2 include the screened out households. Columns 3 and 4 include up to 9 treated months per customer. Columns 5 and 6 restrict the sample to 10 months pre-treatment.

Table A.10: Robustness checks: Dictator game aggregation

	log (quantity) (1)	log (quantity) (2)	log (quantity) (3)	log (quantity) (4)
Incentive treatment	-0.022 [0.031]	0.031 [0.078]	-0.049 [0.054]	0.074 [0.079]
Incentive x Above median	-0.080* [0.048]			
Incentive x Avg sent		-0.150 [0.126]		
Incentive x Minimum sent			-0.022 [0.112]	
Incentive x Max sent				-0.176* [0.103]
HH FE	x	x	x	x
Zone-Month-Year FE	x	x	x	x
Observations (HH)	1,275	1,275	1,275	1,275
Observations (HH-months)	26,122	26,122	26,122	26,122

Notes: Regressions replicate the specification in column 1 of Table 5 showing different measures from the dictator game. Column 1 is the same as Table 5. Column 2 shows a continuous measure of the dictator game endowment share sent to spouse. Columns 3 and 4 show the minimum and maximum share of endowment sent by either spouse.

Table A.11: Robustness checks: Heterogeneous effects of incentive treatment, including other treatments

	log (quantity) (1)	log (quantity) (2)	log (quantity) (3)
Incentive treatment	-0.022 [0.031]	-0.024 [0.035]	-0.041 [0.037]
Incentive x Sent above median	-0.080* [0.048]	-0.081* [0.048]	-0.042 [0.056]
Information treatment		0.004 [0.032]	0.039 [0.042]
Info x Sent above median			-0.075 [0.064]
Credibility treatment		0.027 [0.024]	0.026 [0.031]
Credibility x Sent above median			-0.001 [0.048]
HH FE	x	x	x
Zone-Month-Year FE	x	x	x
Observations (HH)	1275	1275	1275
Observations (HH-months)	26,122	26,122	26,122

Notes: Regressions replicate the specification in column 1 of Table 5 adding the other treatments (column 2) and their interactions with the dictator game measure (column 3). Recall that the information treatment was also provided to all households that received the incentive treatment.

Table A.12: Robustness check: Other outcomes

	Any pay (1)	Missing quantity (2)
Incentive	0.011 [0.014]	-0.005 [0.007]
Price info	0.011 [0.016]	0.003 [0.008]
SWSC credibility	0.005 [0.011]	-0.008 [0.006]
Surveyed	-0.046** [0.022]	0.035*** [0.012]
HH FE	x	x
Zone-Month-Year FE	x	x
Observations (HH)	1,282	1,282
Observations (HH-months)	27,664	29,845

Notes: Regressions replicate column (3) of Table 3, with different outcomes. Column 1 reports the probability that a household made any payment during the month. Column 2 reports the probability that $\log(\text{quantity})$ is missing, conditional on the household having received a prior non-missing meter reading.

Table A.13: Robustness checks: Knowledge and and monitoring of water use, including other treatments

	log (quantity) (1)	log (quantity) (2)	log (quantity) (3)
Incentive treatment	-0.013 [0.042]	-0.013 [0.042]	-0.022 [0.051]
Incentive x Knowledge sum	-0.040 [0.030]	-0.040 [0.030]	-0.033 [0.036]
Credibility treatment		-0.036 [0.042]	-0.037 [0.042]
Credibility x Knowledge sum		0.055* [0.029]	0.055* [0.029]
Information treatment			0.018 [0.055]
Info x Knowledge sum			-0.013 [0.040]
HH FE	x	x	x
Zone-Month-Year FE	x	x	x
Observations (HH)	1,282	1,282	1,282
Observations (HH-months)	26,246	26,246	26,246

Notes: Regressions replicate the specification in column 1 of Table 9 adding the credibility treatment (column 2) and the information treatment (column 3).