THE UNIVERSITY OF CHICAGO

ESSAYS IN ECONOMIC DEVELOPMENT

A DISSERTATION SUBMITTED TO

THE FACULTY OF THE DIVISION OF THE SOCIAL SCIENCES

IN CANDIDACY FOR THE DEGREE OF

DOCTOR OF PHILOSOPHY

KENNETH C. GRIFFIN DEPARTMENT OF ECONOMICS

BY

SAROJINI RAO

CHICAGO, ILLINOIS

JUNE 2020

Contents

Environmental Externalities and Free-riding in the Household

List of Figures

List of Tables

Acknowledgments

I could not have done this without the love and support of my family and remain eternally grateful to my parents Vanitha and Vasudev, my husband Colin S. Sargent, his parents Drs. Nancy and Colin W. Sargent, my aunt, Dr. Nandini Patel, and my uncle S. Govindan. My dear friend Maureen Barradas held my hand through the darkest hours, and made space for me even when I couldn't. I consider myself incredibly lucky to have crossed paths with so many wonderful friends and colleagues who lent an ear, buoyed my spirits, challenged me to do better, and celebrated every milestone on this journey. Last but certainly not least, I remain indebted to my advisors, Professors Anup Malani, Alessandra Voena, and Michael Dinerstein for their patience, encouragement, and thoughtful feedback. Acknowledgments specific to each research paper presented here are below.

For Highway Access and the Utilization of Public Health Insurance: Evidence from RSBY in Karnataka I am grateful for the Karnataka RSBY Society and Suvarna Arogya Suraksha Trust for sharing health insurance data, to the University of Chicago's Center in Delhi for funding, to Gaurav Khanna for sharing the shapefile for the Golden Quadrilateral highway, and to researchers at the Center for Global Health Research for providing access to 2011 census village and town point locations. Monisha Mason, Rohan Parakh, Vishal Talasani, Raghunath Kadamangudi, and Zhida Gui provided excellent research assistance. Conference participants and discussants at the APPAM DC regional student conference and the University of Michigan's H2D2 conference provided valuable feedback on preliminary work. Finally, none of this would have come about but for the Health Insurance Experiment in Karnataka, conducted by Anup Malani, Alessandra Voena, and others. All errors are my own.

For Environmental externalities and free-riding in the household This paper was coauthored with Kelsey Jack (University of California Santa Barbara) and Seema Jayachandran (Northwestern University) to whom I remain grateful for providing this opportunity to collaborate. 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. All errors are our own.

Highway Access and the Utilization of Public Health Insurance: Evidence from RSBY in Karnataka

Abstract

This paper looks at the effect of road access on the utilization of public health insurance in Karnataka, India. Although public health insurance for low-income households was introduced in 2008, insurance utilization and claims submissions remained lower than projected even five years later. The vast majority of private hospitals are located in urban areas, including small towns, making transportation access an important determinant of health insurance utilization. This paper measures the effect of physical access to hospitals, particularly the proximity to a major new highway, on beneficiaries' utilization of the health insurance. This paper uses the construction of a section of the Golden Quadrilateral in Karnataka to evaluate the role of road access on the number of hospitalizations claimed by beneficiaries in a given sub-district. Because the highway connects historically significant cities and was completed prior to the announcement of the insurance scheme, a straight-line approximation to the highway is a suitable instrument for the distance to the highway in a sub-district. Although some recent work has looked at the effect of last-mile road building on education, the evaluation of major highway projects has traditionally focused on macroeconomic outcomes like aggregate productivity and growth. This paper makes a contribution to a small but growing literature on healthcare utilization and access in developing countries, as well as to the infrastructure development literature by considering a unique outcome like hospital access.

1 Introduction

Measuring the demand for medical treatment remains a challenge for economists, in part due to the subjective nature of consumers' perception of their own health and the efficacy of medical interventions, and in part due to the unique nature of medical care when compared to most consumption goods. Self reported health status is notoriously subjective and even very poor people without adequate nutrition, or with chronic conditions or disabilities will often rate their own health as being very good (Banerjee et al. (2004), Sen (2002)). Sen (2002) further points out that self reported morbidity is often more indicative of knowledge of healthcare and morbidity than of actual health status. Expenditure surveys and retroactive surveys of healthcare utilization have been found to provide estimates that are extremely sensitive to the number of categories being surveyed, the reference period, and the choice of respondent within the household (Xu et al. (2009), Lu et al. (2009), Das (2011)). Further, the supply of healthcare (and lack thereof) affects the observed level of expenditure on medical interventions. Mortality rates provide an objective view of population health, but even when researchers can reliably identify a cause of death (Gomes et al. (2017)), estimating the demand for medical care that could have averted death remains a challenge. Approaches from other fields have informed the ways in which data on morbidity is collected in surveys (Murray and Chen (1992)) but although this provides more reliable information than self-reported measures, it remains limited in estimating a demand function for medical treatment.

Health insurance further complicates this landscape. One the one hand, health insurance data provides direct evidence of utilization and cost-sharing experiments (Newhouse (1986)) have shown that insured populations utilize less healthcare when they face higher copays. On the other hand, the observed utilization of health insurance is necessarily a function of which physicians accept the health insurance and why (Fossett and Peterson (1989), Perloff et al. (1995), and Perloff et al. (1997)) as well as non-price determinants including accessibility and transportation costs (Acton (1975), Coffey (1983)). Banerjee et al. (2014) document the low demand for health insurance in rural Rajasthan; however, once people receive health insurance through schemes like RSBY, they do in fact use it to obtain treatment. Further, since RSBY households do not have to pay a copay the observed utilization sheds light on the role of non-price determinants of the demand for medical care such as transportation costs, as shown in this paper.

As developing countries have undertaken mammoth infrastructure development projects, a small but growing literature has focused on evaluating the effect of transportation access in general, and the effect of transport access on healthcare in particular. Recent work by Khanna (2016) and Banerjee et al. (2012) evaluate highway construction projects in particular and its effects on growth and productivity. Related work by Gertler et al. (2016) and Rothenberg (2013) evaluates highway surface quality improvements in Indonesia on labor market outcomes and firm location choices, respectively. In terms of outcomes of interest, the literature on the effects of highways in India (Khanna (2016), Ghani et al. (2016) and Datta (2012)) have focused on growth and productivity, manufacturing, rural employment, and migration. Adukia et al. (2017) measures the effect of road building on school enrollment, focusing not on highways, but rather new roads that connect previously unconnected villages.[1] More generally, the transportation literature has looked at evaluating the effects of highway construction, largely on growth and productivity. (See Allen and Arkolakis (2019), Chandra and Thompson (2000), Michaels (2008), and Fernald (1998) for studies regarding the impacts of inter-state highways in the US.) This paper is unique in considering the effect of one government project on the effectiveness of another: namely of infrastructure projects on a social welfare scheme in a developing country context.

The relationship between transportation and health has been studied in the context of Medicare and Medicaid in the US; for example recent work has focused on the effect of hospital closures on mortality outcomes for Medicare patients (Gujral and Basu (2019).) However, Adukia et al. (2017), Bell (2012), and Bell and van Dillen (2018) are among a small

[1] The construction of new roads that connect previously unconnected villages could also play an important role in the utilization of RSBY in India. However, none of the new roads built under that scheme were in Karnataka. This also addresses any concerns that my results may be driven by "last-mile" road connections and not the Golden Quadrilateral highway.

handful of papers to evaluate the impact of rural road construction in India on education or healthcare outcomes. Independently, recent work by medical researchers has served to document the adverse outcomes that can be attributed to weak transportation systems (Varela (2019), Shrime et al. (2017), and Alkire and Raykar (2015)) as well as highlight the growing demand for essential surgery in developing countries and its role in both health and economic policy.

This paper makes a contribution to this small but growing literature by quantifying the effect of proximity to a major road highway on the utilization of public health insurance by low income households. I provide empirical evidence showing that transportation access imposes significant costs over and above that of medical care alone, and hence affects the utilization and effectiveness of the public health insurance program. Specifically, using an instrumental variables framework, I show that each kilometer added to the distance between a subdistrict and the express highway leads to roughly four fewer claims filed by beneficiaries in that subdistrict. The magnitude of the effect of highway proximity is much higher when the sample excludes subdistricts more than 125 km of the highway, for whom the highway is likely too far to factor into their RSBY utilization behavior. For this restricted sample of subdistricts, I show that each kilometer added to the distance between a subdistrict and the express highway results in 16 fewer claims by beneficiaries in that subdistrict.

The rest of this paper is organized as follows: the following section provides background information on the health insurance program, the highway project and the specific details that inform the data analysis. Section 3 presents a simple heuristic framework for the demand for medical treatment, section 4describes the different data sources used, section 5 maps the variables constructed from the data to the variables in the model. Section 6 summarizes the main findings and presents estimates from alternative specifications as robustness checks and section 7 includes a broader discussion of the underlying mechanisms and implications for healthcare policy in developing country contexts and section 8 concludes.

2 Background and Context

In recent years, the Indian government has shifted its approach to providing healthcare to poor households, supplanting efforts to directly provide free or subsidized treatment at public hospitals and clinics with a public health insurance product targeted to low income households. Independently, prior to the introduction of the health insurance program, the government undertook a major express highway construction project.

2.1 RSBY Health Insurance

In 2008, the Indian government created its first large-scale public health insurance program called Rashtriya Swasthya Bima Yojana (RSBY) or, in English, the National Health Insurance Scheme. RSBY resembles Medicaid in the US in that it provides free health insurance targeted at the very poor and is part of a larger global trend towards public provision of health insurance. RSBY provides coverage for hospital treatment with an annual limit of Rs. 30,000 (approx. USD 500) per annum per household, and aims to cover all 300 million people who fall below the Indian government's official poverty line (BPL), including 9 million individuals in Karnataka alone. The government is now considering expanding the program in an effort to move towards universal health insurance coverage.

The rationale behind introducing public health insurance hinges on the growth of private hospitals: if even one private hospital accepts the public health insurance, low income households are presumed to now have greater choice and financial access to relatively better quality services. It is common knowledge that public hospitals lack adequate capacity to meet the needs of the rural population, and widely perceived as having poor quality of facilities and services. Das and Hammer (2007)and Das et al. (2016) use audit studies to describe the treatment provided by doctors in a primary health setting, and find that private providers (even those with little to no medical training) put in more time and effort than public providers. Banerjee et al. (2004) document the absenteeism of staff at public health facilities, and the high travel and time costs of accessing formal medical care in Rajasthan.

Although private hospitals outnumber public hospitals 4-to-1 in Karnataka, they are much more likely to be located in towns and cities. Figure 3shows the spatial distribution of private hospitals in Karnataka. Hospital geocodes were acquired using publicly accessible geolocation data (such as the Google Maps API) for roughly 1329 private hospitals in Karnataka. The hospital listing, which contained 1651 private hospitals, was provided by Karnataka RSBY in 2011. The listing data included the hospital name and address, and the number of beds. Roughly 20% of the private hospitals from the list were not locatable using publicly available data. 30% of the missing locations are for hospitals located in Bangalore city and the remaining are located in towns that contain other hospitals. Hence the locations that contain one or more hospital is generally well represented by the map in Figure 3 and shows that hospitals are not overwhelmingly located only in towns along the highway, but rather in towns all over the state. Although the bed density per 1000 people for the entire country was still low as of 2010, at 1.30 compared to the WHO guideline of 3.5, 70\% of the growth in the number of beds over the previous decade had occurred in the private sector (Gudwani et al. (2012)). A 2013 report on healthcare access by a private consulting firm finds that although an increasing share of the population utilizes private rather than public hospitals, transportation to private hospitals remains limited in rural areas (Aitken et al. (2013)).

Initial reports on RSBY have focused on the design and initial implementation and adoption of the health insurance (La Forgia and Nagpal (2012), Das and Leino (2011), Krishnaswamy and Ruchismita (2011), and Shoree et al. (2014)) and were critical of the scheme for having lower than expected utilization rates. Further, in the context of Medicaid in the United States, the RAND health insurance experiment as well the more recent Oregon health insurance experiment have shown that utilization increases when people who were not previously insured receive insurance, and decreases in response to increased cost-sharing measures (Newhouse (1986), Finkelstein et al. (2012).) In contrast, RSBY utilization in 2015-

2016 was so low that insurance companies' medical loss ratio was about 50%, as opposed to the 95% threshold required by most state health plans in the US.

This paper explores the extent to which the low health insurance utilization can be attributed to transportation access to hospitals. There are two facets to the physical constraints to hospital access: first, private hospitals are overwhelmingly located in major cities and towns and hence require RSBY beneficiaries to travel a significant distance to get to a private hospital; second, that traveling even a modest distance is especially difficult due to the lack of good roads and highways, and adds a substantial travel time cost. The construction of high quality highways and subsequent increase in private transport providers (typically bus and van operators) is likely to play a significant role in determining the use of inpatient facilities for non-emergency treatments covered by RSBY.

Lastly, hospitals in India are not obligated to provide treatment to indigent patients and even empanelled hospitals can turn away patients. Hence, even (literally) card-carrying beneficiaries are incurring a risk when they try to use their insurance. In a series of experiments introducing Health Savings Accounts in Kenya, Dupas and Robinson (2013) discuss the unique challenges faced by low income households in developing countries who try to accumulate cash savings. Thus further supports the hypothesis that transportation costs make a significant difference for a BPL family trying to utilize their health insurance. In this context, the proximity to a highway not only makes it cheaper to access inpatient care in terms of the time cost of travel, but potentially also reduces search costs if they have to try their luck at multiple hospitals before obtaining treatment.

2.2 Golden Quadrilateral Highway

The Golden Quadrilateral (GQ) started to be built in 2001 and continued through 2012, although the bulk of the project was completed by 2006, with delays largely due to land acquisition conflicts. The goal of the project was to connect the major metropolitan areas of New Delhi, Mumbai, Kolkata, and Chennai. Of these, all but New Delhi are port cities that date back to the late 16th century. As a result, although the GQ construction included the upgrading of a number of previously built national highways, these cities had never before been connected by direct roadways. As part of the planning, a number of other historically significant cities were selected to be included in the network, including the city of Bangalore, which is the state capital of Karnataka, which will be significant for my application.

A naive approach to evaluating the effect of highway construction would be to compare the utilization rates for households near the highway to households away from the highway, but this approach runs into obvious challenges in picking the right "control group." Recent work in development economics has focused on evaluating the construction of major highways between ancient and historically significant cities in India and China by using spatial instruments for the location of the highway. Banerjee et al. (2012) and Faber (2014) use spatial instruments for Chinese highways; the former by using straight line connections between historically significant cities, and the latter by using a least-cost path based on terrain, water bodies, and other topological considerations. Khanna (2016) uses a straight line approximation to the Golden Quadrilateral project, to identify its effects on economic activity as measured by night-time satellite images. I use the same approach, but focus on road construction in Karnataka to identify the effect of proximity to the highway on the likelihood of a hospitalization as well as the number of hospitalizations claimed by beneficiaries in a given sub-district.

Transportation and physical access to services in general is especially relevant in a context where beneficiaries may not be familiar with the services or the service delivery process. In a survey of administrators at RSBY-empanelled hospitals, one of the main reasons given for the low utilization was that people would forget to bring their RSBY card. Doctors who have to send patients home to get their cards and come back and probably much more likely to see those patients again if the patients can travel to and from the hospital with relative ease. [2]

[2]The fact that patients may make multiple trips further suggests that RSBY is not generally utilized in an emergency situation, but rather to cover elective procedures.

Although the design of the health insurance program is not explicitly intended to create hassles or use an ordeal mechanism to dissuade over-utilization, travel and search costs could inadvertently act as a screening mechanism. For example, Deshpande and Li (2019) demonstrate the regressive effect of hassles in the context of Social Service Administration field office closings and disability insurance claims. In the context of this paper, proximity to the highway would benefit those with the highest opportunity cost of time, who may otherwise have put off seeking medical care.

3 Model

Based on the model for healthcare demand developed by Grossman (1972) and subsequently adapted by Acton (1975) to examine the effect of non-monetary costs on healthcare demand, I derive individual demand for inpatient medical care as a function of both the price of treatment as well as the travel cost of obtaining treatment. This model setup provides a clear prediction for the effect of transportation costs, including time costs, on health insurance utilization.

Individual i gains utility from their consumption of goods and services and suffers a loss of utility from illness. For each individual i, let $\bar{H}_i$ represent an individual-specific threshold that represents "perfect health" whereby the individual is able to function without significant physical or mental suffering. Let the individual's actual health status ("stock") at any given moment be represented by $H_i \leq \bar{H}_i$. Let $\delta_i = H_i - \bar{H}_i$ represent deviations from "perfect health" such that negative values of δ with higher magnitudes denote increasingly severe or prolonged morbidity.

As used here, an individual's δ_i at any given point in time represents the accumulation of "insults" over her lifetime. The concept of insult accumulation is borrowed from economic history; insults are said to gradually accumulate through episodes of illness and injury, adverse environmental conditions, and nutritional deficiencies. (Riley (1989)) Hence, for individual i at time t, δ_i can be expressed as $\delta_i = \sum_{j=1}^{t} \delta_{i,j}$, where $\delta_{i,j}$ is a random process. Abstracting from intertemporal considerations, δ_i generally measures the extent of cumulative damage to one's stock of health capital as the number, duration or severity of exposure to various types of insults increase over time, and as the body ages and becomes less able to repair itself and recover.

Individuals can lower their cumulative insults, δ_i, by consuming costly medical treatment M_i. For the purpose of the model, the endogenous variableM_i is continuous, so the consumer can optimally choose fractional quantities of treatment. However, it can also be thought of as a latent variable underlying a binary choice, namely whether or not the individual undergoes a medical procedure in a hospital setting. The magnitude of M_i is not restricted below or above by the magnitude of δ_i; hence it represents medical treatment that is generally considered to be elective, which results in an improved quality of life (like cataract surgery or a hip replacement) as compared to emergency life-saving interventions.[3] Thus individual i's utility function can be expressed as

$$U(X_i, M_i) = f(X_i) + g(M_i - \delta_i)$$

where X_i represents the composite consumption good, $f'(.) > 0$, $f''(.) \leq 0$, $g'(.) > 0$, $g''(.) < 0$. Using this notation helps narrow the interpretation of M_i as a secondary healthcare (i.e. in-hospital) intervention that counteracts any existing morbidity, as opposed to the consumption and production of health in general. In this model, genetic differences as well as intentional health-promoting behaviors may lead to individuals having a higher "perfect" level of health ($\bar{H}_i$), or the curvature of the subutility function (g), or both, but would not affect the distribution of health shocks ($\delta_{i,j}$). Based on insults, these shocks reflect population measures of health, such as childhood nutrition or the likelihood of vaccination for a certain

[3]In this framework, death occurs if δ_i falls below some threshold. If $\delta_i$were large enough to be arbitrarily close to the life-or-death threshold then the level of $M_i$would have to be at least large enough to prevent that threshold from being crossed. Conversely, M_i could be so large as to keep δ_i from crossing that threshold indefinitely; for e.g. an undefined amount of life support. Hence, in the extreme, death is endogenous to the extent that medical treatment is affordable and accessible to the person.

cohort of individuals, and location-based environmental health risks.

Suppose the cost of M_i to the individual consists of two components: a direct cost in terms of the medical fees and an indirect cost arising from accessing hospital services. This indirect cost includes the time cost of receiving medical treatment, but more importantly, the time cost of traveling to and from the hospital. Thus, the overall cost of medical treatment can be written as $(p_M + p_T)M_i$, where p_M is the price of the treatment and p_T is the total cost of accessing treatment. [4] Further, the individual's optimization problem is given by

$$\mathscr{L} = f(X_i) + g(M_i - \delta_i) + \lambda[Y - p_X X_i - (p_M + p_T)M_i]$$

where Y represents the individual's income, p_X represents the price of the composite consumption good, and λ is the Lagrange multiplier. The first order conditions are as follows:

$$\{\text{w.r.t. } M_i\} \qquad g'(M_i - \delta_i) - \lambda(p_M + p_T) = 0 \tag{1}$$

$$\{\text{w.r.t. } X_i\} \qquad f'(X_i) - \lambda p_X = 0 \tag{2}$$

$$\{\text{w.r.t. } \lambda\} \qquad Y - p_X X_i - (p_M + p_T)M_i = 0 \tag{3}$$

Thus,

$$g'(M_i^* - \delta_i) = \frac{f'(X_i)}{p_X}(p_M + p_T) \tag{4}$$

Since $g'(M_i^* - \delta)$ is decreasing in M_i^*, a reduction in p_M, or p_T, or both would increase the optimal amount of medical treatment, holding income and p_X constant. For the purposes of this paper, since I only observe medical treatment for individuals with public health insurance, I can let $p_M = 0$ while p_T varies depending on the distance to the nearest town or city with a hospital. RSBY beneficiaries are not charged copays, but even if that were to change,

[4] There are many ways of writing the cost function such that the optimal amount of medical treatment M_i^* is a function of both p_M and p_T; however, this is the simplest and is consistent with previous work on the nonmonetary costs of medical care. See Acton (1975).

any constant value of p_M would yield the same testable implication for the relationship between M_i^* and p_T, given below.

Proposition: The optimal quantity of medical treatment is decreasing in travel costs, i.e. $\frac{\partial M_i^*}{\partial p_T} < 0$.

Proof: Differentiating both sides of equation 4 with respect to p_T gives:

$$g''(M_i^* - \delta_i)\frac{\partial M_i^*}{\partial p_T} = \frac{f'(X_i)}{p_X}$$
$$\therefore \frac{\partial M_i^*}{\partial p_T} = \frac{f'(X_i)}{p_X} . \frac{1}{g''(M_i^* - \delta_i)}$$

where $\frac{f'(X_i)}{p_X} > 0$ since $f'(.) > 0$ and $\frac{1}{g''(M_i^* - \delta_i)} < 0$ since $g''(.) < 0$. ■

4 Data

Figure 1 depicts the location of the Golden Quadrilateral in India, along with the straight line connecting Mumbai and Bangalore that forms the instrument for this analysis. Coordinates for subdistrict centroids. Publicly available location data provided the coordinates for subdistrict centroids. Karnataka has 177 subdistricts; the distance from a subdistrict centroid to the highway reflects the variation in p_T, i.e. the time cost of accessing healthcare. The 2011 census village and towns amenities dataset includes data on village characteristics such as the presence of a national or state highway, the presence of public health facilities (community health centers, primary health centers, sub-centers, etc.), and the presence of a train station. These variables are aggregated to the sub-district level and used to control for alternative modes of transportation to the nearest town and access to free or subsidized healthcare locally. The 2011 census data also provides subdistrict-level population counts as well as counts for Scheduled Caste (SC) and Scheduled Tribe (ST) sub-populations and the "marginal worker" subpopulation, defined as those having worked for fewer than six months in the preceding year. Since the census does not count BPL individuals as a category, I use these subpopulation counts to control for RSBY eligibility in a given subdistrict. Table 1 summarizes the subdistrict level data used in the analysis.

RSBY claims data: For the 2011 to 2013 period the data does not include sub-district identifiers and is used for preliminary district-level analysis. However, there are only thirty districts. The scheme was inactive in 2014 and reintroduced in 2015 following a major re-enrollment campaign. The 2015-2016 claims data contains sub-district identifiers; the number of claims aggregated to the subdistrict level reflects the variation in M_i across subdistricts.

Data from the 60th round of the National Family Health Survey, conducted January-June 2004, are used to construct an index of morbidity indicators to capture the variation in δ_i as well as index of hospital utilization indicators to control for hospital utilization patterns that predated the introduction of RSBY. However, this variation is only available at the district level. Table 2 summarizes the variables used to control for morbidity and utilization. These variables are highly collinear, as shown in Tables 10 and 11. Thus, constructing indices based on the principal component factors for each set of variables helps isolate the variation across districts that results directly from the covariance between these variables. As a result, the indices are more likely to represents actual morbidity and utilization, as compared to the raw variables, which are likely correlated with the other control variables such as SC/ST subpopulations, and village characteristics. Tables 8 and 9 show the results from principal components analysis of the variables used in constructing each index, and present pairwise correlations for the variables used in each index, respectively.

5 Empirical strategy

The main outcome of interest is the number of health insurance claims that were filed by RSBY beneficiaries in a given sub-district, including the sub-districts that filed no claims. (Only three sub-districts did not file any claims.) The primary explanatory variable of interest is the distance from the subdistrict (centroid) to the south-west segment of the Golden Quadrilateral highway. To address the endogenous placement of the highway, I use the distance from the sub-district (centroid) to a straight line Mumbai and Bangalore as an instrument for highway location.

RSBY utilization in general could also be affected by variation in the propensity to become ill and use hospital services across subdistricts that is independent of the introduction of RSBY. To control for this variation in latent morbidity and utilization, I use an index of survey-based demographic and population health measures. Because highway access can directly affect morbidity (through accidents or the spread of infectious diseases) I use survey data from the 60th round of the National Sample Survey, conducted in 2004 instead of more recent years to control for latent morbidity and pre-existing levels of hospital utilization. If all eligible individuals, i.e. the entire BPL population were automatically enrolled, then the number of claims could be normalized by the number of enrollees in each subdistrict. However, RSBY enrollment is carried out by the contracted insurance companies who conduct a campaign to physically verify the identity of beneficiaries and sign them up in person. This process likely results in endogenous enrollment rates, whereby subdistricts that are closer to the highway or generally more accessible are more likely to be visited during the sign-up campaign and have longer duration visits enabling better coverage.[5] Thus, in lieu of normalizing the number of claims, I control for the variation in the eligible population using 2011 census data, as mentioned previously.

In addition to these concerns, households in general (and BPL households in particular) may have migrated to areas closer to the highway in order to better access urban centers for reasons including access to healthcare. Since highway construction in Karnataka had concluded before the introduction of RSBY, the relationship between highway proximity and utilization being measured here includes the effect of any migration towards the highway and potentially higher enrollment rates closer to the highway. That is, I cannot "net out"

[5] Differential enrollment rates based on distance to the highway could not be verified due to a lack of subdistrict level enrollment data.

the effect of highway proximity on RSBY utilization that is directly caused by migration instead of lowered transportation times from more remote sub-districts, especially by BPL households that supply informal construction labor.

Hence, the first-stage regression is given by

$$DistanceToGQ_i = \alpha_1 + \beta_1 DistanceToLine + \phi_1 MorbidityControls_i + \gamma_1 LocalControls + \epsilon_{i1} \tag{5}$$

and the second-stage regression equation is given by

$$Claims_i = \alpha_2 + \beta Distan\hat{c}eToGQ_i + \phi_2 MorbidityControls_i + \gamma_2 LocalControls + \epsilon_{i2} \tag{6}$$

where $Distan\hat{c}eToGQ_i$ contains the instrumented (predicted) values of the endogenous variable $DistanceToGQ_i$ in the second stage, following a first stage regression on the instrument, $DistancetoLine_i$ and the other covariates. $DistancetoLine_i$ is the distance to the straight-line connecting Mumbai and Bangalore. $MorbidityControls$ is an index of survey-based population health indicators and $LocalControls$ is the number of marginal workers and schedule caste population, both intended to capture the potential extent of RSBY enrollment in a given sub-district. $DistancetoLine$ meets the requirements for instrument validity since it is highly correlated with $DistancetoGQ$, but is completely independent from the number of claims, for any given value of $DistancetoGQ$. Formally, the exclusion restriction is given by:

$$Claims \perp\!\!\!\perp DistancetoLine | DistancetoGQ.$$

6 Results

Table 1 provides summary statistics for the 176 subdistricts used in the analysis. Of note, while the average distance to the GQ highway is about 99 km, the average distance to the district capital is closer to half that distance at 55 km. Since district capitals are home to the largest public hospital within the district, and BPL households are typically eligible to receive treatment at no direct cost at public hospitals, traveling to the district hospital presents a viable alternative to taking the highway to a larger city for hospital treatment. Hence, the very fact that individuals are willing to incur higher travel costs (even with the GQ highway) to seek treatment at a private hospital reflects the willingness to pay for presumably higher quality care.

6.1 First stage estimates

Table 3 presents regression results for the first-stage regression given by equation 5 above. Column 1 presents results for the full sample, containing all 176 subdistricts. Although, as expected, the distance to the straight line is closely related to the distance to the actual highway, the first stage results here also show pre-existing trends in hospital utilization. Locations close to the highway appear to have higher rates of (non-RSBY) hospital utilization as measured in 2004, suggesting that highway construction (which had begun in 2001) could have already started to influence the demand for medical care. The migration of laborers towards the highway also helps explain the significant coefficient on the size of the SC population in 2011 also suggests that a large number of highway construction laborers may have migrated towards the highway during this time. Column 2 presents the first stage results for a sub-sample excluding 24 subdistricts that lie between the highway and the straight line. For this sub-sample, the distance from the subdistrict centroid to the straight line is necessarily increasing with distance to the highway, thereby meeting the monotonicity requirement for the second stage IV estimate to be interpreted as a local average treatment effect (LATE).

Column 3 3 presents the first-stage results when the sample is restricted to subdistricts within 125 km of the GQ highway, resulting in 124 observations. [6]For this subsample, the coefficient on the distance to the straight line is much smaller since the distance between the highway and the straight line is now a larger proportion of the overall distance from the subdistrict centroid to the highway. Column 4 restricts the sample further by dropping the 24 subdistricts that lie between the highway and the straight line, leaving only 100 observations. As with column 2, the distance from the subdistrict centroid to the straight line is necessarily increasing with distance to the highway, thereby meeting the monotonicity requirement for the second stage estimate to be interpreted as a LATE. Lastly, prior hospital utilization appears to be more homogenous with respect to highway proximity within the subsamples presented in columns 3 and 4, which makes the second stage results easier to interpret.

6.2 Second stage IV estimates

Table 4 presents second stage results, with each column corresponding to the same sample restrictions used in the first stage, respectively. For the full sample, an additional kilometer of distance to the GQ highway is associated with roughly 4 fewer claims. Restricting the sample by removing subdistricts between the highway and the straight line reduces the estimated coefficient, but the coefficients in columns 1 and 2 are not statistically different. Columns 3 and 4, with the restricted sample, yield much larger estimates, suggesting that for subdistricts located within 125 km of the highway, each addition kilometer away from the highway is associated with roughly 16 fewer claims. By dropping subdistricts away from the highway that still have hospitals that accept RSBY, I effectively exclude people who are largely unaffected by the highway in their decision to seek medical treatment and utilize RSBY and keep the population for whom highway proximity is more relevant and salient.

[6]A distance greater than 125 km is likely too far to travel just to get to the highway; distances less than 125 km led to dropping too many observations and losing power.

Among the other control variables, the size of the SC subpopulation appears positively correlated with the number of claims, indicating that in addition to being closer to highway construction areas, enrollment is probably also higher within this group. In the full sample the morbidity and utilization indices have large coefficients in opposite directions: I interpret these as jointly showing that subdistricts that were *ex ante* more likely to utilize hospital services are also more likely to utilize RSBY, whereas subdistricts that had higher morbidity rates (and low utilization) are less likely to utilize RSBY. Coefficients on the indices are harder to interpret for the restricted samples since they index itself captures variation at the district level, which in some ways obfuscates the effect of restricting the sample, especially for the morbidity index. However, the large and significant coefficient on the utilization index in columns 3 and 4 suggests that within the restricted samples, RSBY claims are more likely to arise from a substitution towards using insurance by people who were already likely to seek out medical care, as opposed to arising from the external margin.

6.2.1 Elasticities

The estimated coefficients can be converted into elasticities using the average number of claims and the average distance to the GQ highway as follows:

$$\epsilon = \frac{\partial Cl\hat{a}ims}{\partial DistanceToGQ} \times \frac{Distan\bar{c}eToGQ}{Cla\bar{i}ms}$$

where $\frac{\partial Cl\hat{a}ims}{\partial DistanceToGQ}$ is the estimated second stage coefficient on $DistanceToGQ$, $Distan\bar{c}eToGQ$is the average distance to the GQ highway, and $Cla\bar{i}ms$is the average number of claims in the sample.

Using the LATE estimate for the full sample (-3.885) (Table 4, column 2) and the sample averages for distance to the GQ (99.24) and number of claims (586.61) as reported in Table 1, the elasticity of demand for RSBY utilization with respect to distance is -0.66.

For the restricted sample, I plug in the estimated coefficient (-16.475) (Table 4, column 4) and use the sub-sample averages of distance to the GQ (64.23) and number of claims

(614.31), which gives an elasticity of 1.72. Hence, for subdistricts within 125 km of the GQ highway a 1% decrease in highway proximity leads to a roughly 1.72% increase in the number of claims filed.

6.3 IV estimates with additional control variables

Column 1 in Table 5 presents the baseline results for the full sample (same as column 1 in Table 4) and each subsequent column adds variables controlling for potentially relevant local characteristics such as the numbers of public and private hospitals in the subdistrict as well as the distance to the subdistrict headquarters, district capital, and nearest census designated town, and shows results for the full sample. The main coefficient of interest is generally in keeping with the previous results. However, the significant negative coefficient on the distance to subdistrict HQ suggests that each additional kilometer away from the subdistrict headquarter is associated with 32 fewer claims. This has important implications for last mile accessibility. Access to the subdistrict HQ is often a necessary precondition for finding a private minibus or other private mass transit that uses the highway. It also serves as a conduit for public transit to the district capital and other towns that may have hospitals that accept RSBY.

Table 6 reports coefficients for the same specifications as Table 5, but for the restricted sample of subdistricts that are within 125 km of the GQ highway but not between the highway and the straight line. As with the previous table, the main coefficient of interest is not statistically different across the specifications. However, within this smaller sample, proximity to the subdistrict HQ does not appear to be as important, especially as compared to the hospital utilization index.

6.4 Robustness checks for morbidity and utilization

Tables 7 presents results using variables used to construct the morbidity and utilization indices rather than the indices without and with the additional controls for local characteristics respectively. While the coefficient on distance to GQ is slightly lower in magnitude compared to the most parsimonious version, an additional kilometer to the highway is still associated with about 3.5 fewer claims. The morbidity and utilization variables are largely insignificant or difficult to interpret because of their collinearity, except that immunization rates appear to be a robust predictor of the number of claims at the district level. This may reflect the role of education or access to health-specific knowledge which is likely endogenous with RSBY utilization. Thus although Tables 4, 5 and 7 yield coefficients in the same ballpark for $DistanceToGQ$, the specifications with all the morbidity and utilization variables instead of the indices are also likely to introduce endogeneity with the outcome variable.

7 Discussion

7.1 Mechanisms

On the whole, there are at least four likely mechanisms through which proximity to the highway results in more claims:

1. Lower cost to consumers: As illustrated in the model presented here, proximity to the highway lowers the transportation cost to consumers and leads to higher demand for healthcare in general, and for inpatient procedures covered by RSBY in particular. Specifically, using the highway makes it easier for beneficiaries to find a provider that (a) can perform the inpatient service or procedure that the patient needs and (b) accepts RSBY coverage. Highway access also makes it easier for family members to visit the patient or take turns staying with them, allows patients to return for follow-up visits as needed, and generally reduces the nonmonetary burden of undergoing surgery.

2. Lower costs to producers: Because the highway also reduces the costs of traveling from the city to rural areas, it could also lower the operating costs of running a hospital in a rural area. Physicians who may not have wanted to relocate to such areas may find it

viable to commute to rural hospitals, while continuing to live closer to the city where they have access to better amenities. This would lead to an increase in the number of hospitals opening up in rural areas, which does not appear to be the case.

3. Migration: Employment in highway construction leads a significant share of BPL households (i.e. would-be RSBY beneficiaries) to move closer to the highway.

4. Enrollment: Insurance companies may enroll more BPL households in areas closer to the highway since they are easier to access and enroll fewer households in more remote areas. This is harder to rule out: the government requires insurance companies to meet a minimum enrollment rate in every subdistrict but may not monitor enrollment as closely as would be necessary to enforce the requirement. Thus spatial patterns in claims filing may reflect underlying spatial patterns in enrollment into RSBY.

Highway access most likely affects claims behavior because it allows RSBY beneficiaries to substitute private hospital care for local public hospital care. Every sub-district in this sample has at least one inpatient public facility. Data from the Post Health Event Survey shows that 75 percent of those experiencing a health event go to a private hospital.[7] Local private hospitals are likely a closer substitute for larger private hospitals in major cities, but the location of private hospitals is likely also endogenous based on proximity to the highway.

Investigating any of these channels would require multiple years of data, not only on the number of claims, but also for measures of private healthcare provision at the subdistrict-level and migration among BPL households. The hospital survey indicated that private hospitals in rural subdistricts chose to accept RSBY primarily as a way to increase their capacity-utilization; i.e. to increase bed occupancy rates. RSBY reimbursement rates are far below the prices charged by private hospitals to non-RSBY patients who typically pay out of pocket. To the extent that private hospitals in rural subdistricts had excess/underutilized capacity, it is unlikely that mechanism 2 in the list above actually results in increased supply

[7] This survey was conducted in Karnataka over two waves in 2016 and 2017 with a total of roughly 750 respondents who had experienced a health event for which they sought treatment.

of private healthcare in those subdistricts. Even if the cost of operating a facility in a rural area is reduced by the presence of the highway, doctors from large cities would not find it worthwhile to compete with the rural private hospitals just for RSBY patients who don't bring in much revenue for them. They would need a critical mass of uninsured patients paying the higher out-of-pocket rates to justify the high fixed cost of building new private facilities in rural subdistricts simply because it is close to the GQ and would have lower operating costs in the future. It would be cheaper for them to stay in the city but advertise in rural areas, clearly indicating that they accept RSBY, thereby drawing RSBY beneficiaries from rural subdistricts to travel to the city for treatment.

7.2 Healthcare outcomes

Although the results presented here shed light on the relationship between the cost of access to secondary healthcare and the demand for the same as reflected by the number of claims, the ultimate outcome of interest, at least for policymakers, is the actual health status of the BPL population. Higher rates of RSBY utilization would not necessarily imply improved health outcomes for two possible reasons. First, if the demand for medical interventions were "induced" and not medically necessary, then patients' baseline health status may not be improved by utilizing RSBY and may even be worsened if, for example, they were exposed to hospital-acquired infections. However, RSBY reimbursement rates are far below market rates for the same procedures, thereby providing very little incentive for providers to "induce demand" or "upsell" inpatient procedures.

RSBY utilization could also be ineffective in improving health conditions in the BPL population if the treatment provided were ineffective. Given the concerns raised in the literature about the quality of private outpatient care, it would not be unreasonable to expect similar inefficacy for inpatient medical treatment as well.

RSBY claims data from the initial two years of the program show that the most common procedures undertaken by the newly insured population were cataract surgeries and the removal of non-cancerous tumors and cysts. There had likely been pent-up demand for elective surgeries to resolve conditions that may not have added significantly to an individual's mortality risk but likely affected their quality of life by making activities of daily living more onerous. Thus, cost-benefit analyses of the public health insurance program are likely to underestimate the benefits of the program if they focus exclusively on mortality reductions as opposed to the economic benefits of reduced morbidity or improved quality of life. Such economic benefits could include labor market participation and greater productivity in market- and non-market activities. Whereas Banerjee et al. (2014) find that poor households in Rajasthan do not value health insurance enough to pay for it, this paper suggests that conditional on having health insurance, poor households in Karnataka not only utilize the health insurance, but also respond to travel costs as theory predicts.

7.3 Health systems design

Between the low reimbursement rates and the differences in quality of life between rural and urban areas, it is unlikely that RSBY will induce significant entry of hospitals into rural markets. However, the presence of strong transportation networks suggests that a better approach might be to develop a wheel-and-spokes model of healthcare delivery such that patients from rural areas can be referred to specialists located in cities. Larger specialized hospitals in urban areas could increase their "footprint" by actively building such referral networks with mobile testing and diagnostic services and potentially providing transportation (like a shuttle service) along highways.

Finally, the analysis and results presented here all hinge on the elective nature of RSBY utilization. In the context of a global pandemic such as COVID, this analysis also points to the limitations of RSBY. To the extent that highway access and use is enabled by an expansion of public transit, this is immediately curtailed by a lockdown. Private hospitals are more likely to have Intensive Care Units (ICU) as well as equipment such as ventilators.[8]

[8]Data comes from a survey of 420 hospitals conducted by the author in Karnataka in 2014.

However, even hospitals that accept RSBY can choose to allocate scarce ICU beds and ventilators to patients who are able to pay out-of-pocket, leaving RSBY beneficiaries no better off than if they were uninsured.

8 Conclusion

In conclusion, this paper shows that regardless of self-reported health status, individuals utilize health insurance once they have it and obtain elective surgeries that likely contribute to their quality of life and productivity. Highway access facilitates the utilization of public health insurance by significantly reducing travel costs and enabling the insured to seek out higher quality treatment. This paper makes a contribution to this small but growing literature by quantifying the effect of proximity to a major road highway on the utilization of public health insurance by low income households.

Using an instrumental variables framework, this paper shows that each kilometer added to the distance between a subdistrict and the GQ highway leads to about four fewer claims filed by beneficiaries in that subdistrict. When the sample is restricted to subdistricts that are within 125 km of the highway, the magnitude of the impact of highway proximity is much higher: each kilometer added to the distance between a subdistrict and the GQ highway results in 16 fewer claims by beneficiaries in that subdistrict. Thus, this paper provides empirical evidence showing that transportation access imposes significant costs over and above that of medical care alone, which affects the utilization and effectiveness of the public health insurance program.

Figures

Figure 1: The Golden Quadrilateral Project (Karnataka highlighted)

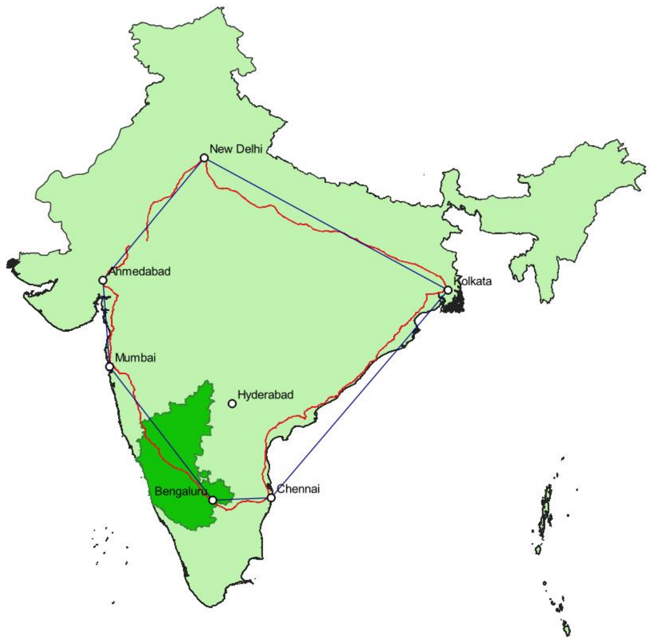

Source: GADM (Map), Khanna (GQ)

Figure 2: Golden Quadrilateral and straight-line in Karnataka

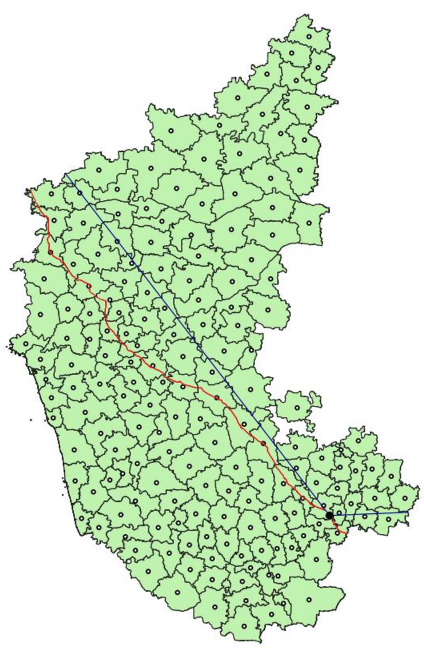

Source: GADM (Map) and Khanna (GQ)

Figure 3: Private hospital location in Karnataka

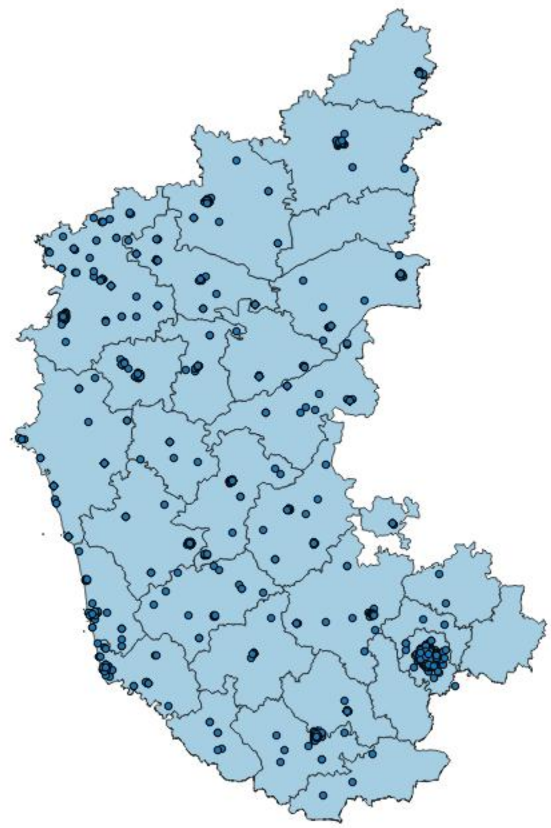

Source: GADM (Map), Karnataka RSBY Society (Hospital names and addresses), Google Maps API (Hospital locations)

Tables

Table 1: Summary statistics

	Mean	St Dev	Min	Max
Number of claims	586.61	850.956	0	5980
Distance to GQ (km)	99.24	81.806	0	373
Distance to straight line (km)	100.24	73.627	1	326
Marginal workers	23333.74	1.3e+04	1384	69351
Schedule Caste population	53463.10	3.3e+04	3271	181961
Schedule Tribe population	23055.70	2.4e+04	226	132166
Morbidity index	-0.07	0.789	-1	4
Hospital utilization index	0.01	0.937	-2	2
Private inpatient facilities	4.18	4.224	0	26
Empanelled hospitals (district)	27.45	13.130	2	53
Distance to subdistrict HQ (km)	19.33	4.747	8	42
Distance to district capital (km)	54.85	30.899	13	155
Distance to nearest town (km)	16.76	4.131	7	36
Observations	176			

Notes: Summary statistics are provided for the full sample of 176 subdistricts. Only 3 subdistricts had 0 claims.

Table 2: Summary statistics: Variables in morbidity and utilization indices

	Mean	St Dev	Min	Max
Whether ailing-last 15 days	0.30	0.142	0.09	0.75
Whether ailing-on the day before	0.19	0.114	0.01	0.64
Number of ailment episodes in last 15 days	0.32	0.151	0.11	0.82
Days ill in last 15 days	3.07	1.834	0.83	10.51
Days on restricted activity in last 15 days	1.13	0.923	0.19	5.20
Days on bedrest in last 15 days	0.27	0.260	0.01	0.98
Mortality	0.01	0.012	0.00	0.05
Seniors	0.26	0.065	0.17	0.49
OBGYN cases	0.81	0.078	0.64	0.95
Number of times treatment was sought	0.26	0.149	0.04	0.82
Number of times treatment was sought at govt facility	0.08	0.057	0.01	0.21
Whether hospitalised	0.11	0.040	0.04	0.19
No. of times hospitalized	0.12	0.045	0.04	0.19
Inpatient stays at pvt hospital in last year	0.07	0.044	0.02	0.18
Inpatient stays at public hospital in last year	0.05	0.027	0.01	0.10
Total inpatient days in last year	1.07	0.553	0.22	2.45
Underwent surgery	0.03	0.016	0.00	0.07
Any pre-natal care	0.08	0.043	0.01	0.20
Post natal care	0.05	0.032	0.00	0.15
Immunisation	0.42	0.128	0.23	0.66
Any health scheme	0.02	0.022	0.00	0.06
Observations	27			

Notes: These variables represent proportions or averages across households within a given district; this table presents the average values of these variables across all districts. These variables were constructed using data from the 60th round of the National Sample Survey, which was conducted in 2004. At the time Karnataka only had 27 districts, 3 more were created in 2011.

Table 3: First-stage regression

	Full Sample	Monotonic IV	Within 125 km	Restricted sample
	Distance to GQ (1)	Distance to GQ (2)	Distance to GQ (3)	Distance to GQ (4)
Distance to straight line (km)	0.992***	0.995***	0.525***	0.472**
	[0.055]	[0.062]	[0.140]	[0.168]
Marginal workers	-0.000	0.000	-0.001	-0.000
	[0.000]	[0.000]	[0.001]	[0.001]
Schedule Caste population	0.001***	0.000***	0.000**	0.000
	[0.000]	[0.000]	[0.000]	[0.000]
Schedule Tribe population	-0.000	-0.000	0.000	0.000
	[0.000]	[0.000]	[0.000]	[0.000]
Morbidity index	10.432	9.358	12.756	12.672
	[6.841]	[6.988]	[8.584]	[9.911]
Hospital utilization index	-27.048***	-24.988***	-7.596	-3.810
	[7.200]	[7.058]	[10.559]	[10.751]
Constant	-17.882*	-23.092**	20.290	20.645
	[10.509]	[10.213]	[16.162]	[18.074]
Observations	176	152	124	100
F-Stat	64.892	70.292	16.502	10.830

Notes: The sample in column 1 contains all subdistricts. Column 2 drops subdistricts that lie in between the highway and the straight line; the instrument meets the monotonicity requirement for the remaining subdistricts, relevant to the second stage. Column 3 drops subdistricts located more than 125 km from the GQ highway, and Column 4 combines sample restrictions for columns 2 and 3. All standard errors are clustered by district; there are 30 clusters for the samples in columns 1 and 2 and 22 clusters in for the samples in columns 3 and 4.

Table 4: IV regression results

	Full Sample	Monotonic IV	Within 125 km	Restricted sample
	Claims (1)	Claims (2)	Claims (3)	Claims (4)
Distance to GQ (km)	-4.081***	-3.885***	-14.778**	-16.475*
	[1.010]	[1.042]	[6.047]	[8.980]
Marginal workers	0.009	0.015**	0.006	0.017
	[0.009]	[0.007]	[0.013]	[0.011]
Schedule Caste population	0.008**	0.007*	0.008	0.004
	[0.004]	[0.004]	[0.005]	[0.005]
Schedule Tribe population	-0.004	-0.004	-0.002	0.000
	[0.004]	[0.004]	[0.004]	[0.004]
Morbidity index	-175.762	-85.592	-367.556**	-173.030
	[108.738]	[91.107]	[174.816]	[197.308]
Hospital utilization index	196.219*	155.147	396.974**	379.703**
	[115.086]	[109.588]	[161.066]	[188.852]
Constant	407.464**	338.652*	949.659**	1009.052
	[175.335]	[191.905]	[467.252]	[654.979]
Observations	176	152	124	100

Notes: The sample in column 1 contains all subdistricts. Column 2 drops subdistricts that lie in between the highway and the straight line; the instrument meets the monotonicity requirement for the remaining subdistricts. Column 3 drops subdistricts located more than 125 km from the GQ highway, and Column 4 combines sample restrictions for columns 2 and 3. All standard errors are clustered by district; there are 30 clusters for the samples in columns 1 and 2 and 22 clusters in for the samples in columns 3 and 4.

Table 5: Robustness check: Varying and adding control variables, full sample

	Claims (1)	Claims (2)	Claims (3)	Claims (4)
Distance to GQ (km)	-4.081***	-4.210***	-3.743***	-3.786***
	[1.010]	[1.099]	[0.930]	[1.048]
Marginal workers	0.009	0.004	0.011	0.008
	[0.009]	[0.009]	[0.008]	[0.009]
Schedule Caste population	0.008**	0.009**	0.008*	0.009*
	[0.004]	[0.004]	[0.004]	[0.004]
Schedule Tribe population	-0.004	-0.003	-0.003	-0.003
	[0.004]	[0.004]	[0.005]	[0.004]
Morbidity index	-175.762	-211.638*	-166.494*	-196.071*
	[108.738]	[110.922]	[100.736]	[106.750]
Hospital utilization index	196.219*	199.719*	207.299**	208.292**
	[115.086]	[111.134]	[103.926]	[104.425]
Public inpatient facilities		5.952		4.250
		[3.969]		[4.058]
Private inpatient facilities		-10.913		-8.668
		[17.860]		[15.292]
Empanelled hospitals (district)		3.068		3.010
		[7.122]		[7.129]
Distance to subdistrict HQ (km)			-34.757**	-32.217**
			[16.229]	[15.968]
Distance to district capital (km)			-2.833**	-2.485*
			[1.439]	[1.363]
Distance to nearest town (km)			17.683	14.145
			[21.728]	[20.869]
Constant	407.464**	174.693	859.266**	662.512
	[175.335]	[282.264]	[356.178]	[421.220]
Observations	176	176	176	176

Notes: All subdistricts are included for all columns. All standard errors are clustered by district; there are 30 clusters in the full sample.

Table 6: Robustness check: Varying and adding controls, restricted sample

	Claims (1)	Claims (2)	Claims (3)	Claims (4)	Claims
Distance to GQ (km)	-16.475*	-17.316*	-16.839*	-17.761	
	[8.980]	[10.323]	[9.353]	[11.050]	
Marginal workers	0.017	0.011	0.017	0.013	
	[0.011]	[0.017]	[0.011]	[0.019]	
Schedule Caste population	0.004	0.005	0.004	0.004	
	[0.005]	[0.005]	[0.005]	[0.005]	
Schedule Tribe population	0.000	0.002	0.003	0.004	
	[0.004]	[0.005]	[0.004]	[0.005]	
Morbidity index	-173.030	-155.000	-130.312	-108.751	
	[197.308]	[222.348]	[211.866]	[243.297]	
Hospital utilization index	379.703**	401.993**	405.722**	428.698**	
	[188.852]	[203.159]	[180.802]	[200.114]	
Public inpatient facilities		5.487		4.588	
		[8.132]		[8.434]	
Private inpatient facilities		-2.641		-4.504	
		[26.020]		[24.630]	
Empanelled hospitals (district)		0.210		-0.678	
		[11.783]		[13.242]	
Distance to subdistrict HQ (km)			-47.834	-46.816	
			[33.374]	[35.802]	
Distance to district capital (km)			-1.142	-1.060	
			[2.946]	[3.089]	
Distance to nearest town (km)			17.223	16.133	
			[31.583]	[33.720]	
Constant	1009.052	937.417	1655.577**	1623.755*	
	[654.979]	[853.295]	[655.671]	[886.380]	
Observations	100	100	100	100	

Notes: The sample is restricted to subdistricts within 125 km of the GQ highway that are not located between the highway and the straight line. All standard errors are clustered by district; there are 22 clusters in the restricted sample.

Table 7: Robustness check: Morbidity and utilization variables

	claims	claims	claims	claims
Distance to GQ (km)	-3.503***	-3.235**	-3.787***	-3.535***
	[1.300]	[1.336]	[1.193]	[1.260]
Marginal workers	0.004	-0.003	0.004	-0.003
	[0.007]	[0.008]	[0.007]	[0.008]
Schedule Caste population	0.007**	0.007**	0.008**	0.008**
	[0.003]	[0.003]	[0.003]	[0.003]
Schedule Tribe population	-0.003	-0.002	-0.003	-0.002
	[0.003]	[0.003]	[0.003]	[0.003]
Whether ailing-last 15 days	-5048.094	-4774.303	-889.456	385.950
	[4253.660]	[4861.944]	[4199.857]	[4650.665]
Whether ailing-on the day before	-50.231	1353.234	1185.301	2394.524
	[1940.667]	[2897.736]	[1815.264]	[2694.586]
Number of ailment episodes in last 15 days	4846.058	4039.617	1455.989	-260.180
	[5268.474]	[5680.825]	[5090.964]	[5388.197]
Days ill in last 15 days	-23.907	-118.291	-65.877	-149.071
	[199.203]	[238.270]	[199.419]	[237.549]
Days on restricted activity in last 15 days	110.938	133.156	110.530	120.711
	[298.412]	[300.250]	[313.447]	[310.060]
Days on bedrest in last 15 days	844.365***	754.332**	683.523*	572.491
	[306.491]	[330.377]	[349.661]	[377.776]
Mortality	-19966.526	-20800.840	-18914.527	-18703.723
	[13628.238]	[14733.445]	[12834.646]	[14204.147]
Seniors	1037.935	1293.869	-575.363	-168.755
	[1714.730]	[2224.688]	[1790.260]	[2230.757]
OBGYN cases	2895.301	3560.585*	2162.316	2743.152
	[1852.508]	[2051.335]	[1728.019]	[1930.340]
Number of times treatment was sought	-3443.346	-3204.588	-3674.020	-3487.020
	[2950.264]	[2851.224]	[3018.430]	[2943.589]
Underwent surgery	-13243.683	-14130.427	-14245.567*	-15449.240*
	[8263.639]	[8805.513]	[7760.455]	[8505.033]
Whether hospitalised	51346.640***	46592.210***	49931.949***	45300.229***
	[8235.244]	[9363.308]	[8263.177]	[9528.645]
No. of times hospitalized	-32261.360***	-26571.845***	-31690.308***	-26314.047***
	[6270.368]	[8198.143]	[6220.413]	[7979.726]
Total inpatient days in last year	319.853	264.617	398.343	352.226
	[262.116]	[253.345]	[255.985]	[250.528]
Any pre-natal care	2070.172	1840.755	1765.802	1358.467
	[3758.386]	[4124.541]	[3665.942]	[4039.795]
Post natal care	-2804.915	-4491.529	-1135.889	-2396.337
	[6617.912]	[7474.704]	[6419.803]	[7196.339]
Immunisation	2221.826***	2411.227***	2447.407***	2618.032***
	[600.118]	[604.720]	[628.303]	[636.488]
Public inpatient facilities		4.599		5.247*
		[3.241]		[3.126]
Private inpatient facilities		9.328		10.401
		[13.436]		[12.358]
Empanelled hospitals (district)		8.695		8.168
		[7.681]		[7.796]
Distance to subdistrict HQ (km)			-38.324**	-37.266**
			[15.013]	[14.480]
Distance to district capital (km)			-2.170	-2.374
			[1.765]	[1.713]
Distance to nearest town (km)			23.627	19.198
			[15.206]	[15.703]
Constant	-4065.984**	-4935.400**	-2887.078	-3628.491*
	[1846.102]	[1991.626]	[1775.194]	[1919.484]
Observations	176	176	176	176

Notes: All subdistricts are included for all columns. All standard errors are clustered by district; there are 30 clusters in total.

Table 8: Principal Component Factors- Morbidity index variables

	Factor 1	Factor 2
Whether ailing-last 15 days	0.915	0.159
Whether ailing-on the day before	0.974	0.001
Number of ailment episodes in last 15 days	0.937	0.124
Days ill in last 15 days	0.966	-0.043
Days on restricted activity in last 15 days	0.834	-0.107
Days on bedrest in last 15 days	0.229	0.476
Mortality	0.256	0.763
Seniors	0.842	-0.218
OBGYN cases	-0.382	0.734
Eigenvalue	5.266	1.449

Table 9: Principal Component Factors- Hospital utilization index variables

	Factor 1	Factor 2
Number of times treatment was sought	0.713	-0.118
Underwent surgery	0.759	0.354
Whether hospitalised	0.964	0.071
No. of times hospitalized	0.953	0.093
Total inpatient days in last year	0.867	0.023
Any pre-natal care	0.111	0.916
Post natal care	0.215	0.918
Immunisation	-0.359	0.686
Eigenvalue	3.862	2.308

Table 10: Pairwise correlation matrix of morbidity control variables

	ailing_15days	ailing_daybefore	ailment_cases	ill_duration	rest_duration	bedrest_duration	dead	senior_60plus	maternal
ailing_15days	1								
ailing_daybefore	0.852***	1							
ailment_cases	0.995***	0.874***	1						
ill_duration	0.831***	0.952***	0.868***	1					
rest_duration	0.632***	0.835***	0.660***	0.830***	1				
bedrest_duration	0.341	0.165	0.312	0.0854	0.109	1			
dead	0.243	0.283	0.251	0.310	0.130	0.115	1		
senior_60plus	0.687***	0.827***	0.713***	0.798***	0.687***	0.0269	0.0385	1	
maternal	-0.235	-0.345	-0.280	-0.409*	-0.300	0.0657	0.336	-0.346	1
N	27								

Table 11: Pairwise correlation matrix of utilization control variables

	any_treatment	received_surgery	hospitalized	times_hospitalized	hosp_los	prenatal	postnatal	immunization
any_treatment	1							
received_surgery	0.406*	1						
hospitalized	0.665***	0.676***	1					
times_hospitalized	0.599***	0.672***	0.983***	1				
hosp_los	0.438*	0.675***	0.789***	0.826***	1			
prenatal	-0.0590	0.308	0.222	0.224	0.0555	1		
postnatal	0.0964	0.483*	0.241	0.246	0.174	0.866***	1	
immunization	-0.218	-0.0377	-0.294	-0.247	-0.206	0.429*	0.426*	1
N	27							

References

Acton, J. P. (1975). Nonmonetary factors in the demand for medical services: Some empirical evidence. *Journal of Political Economy 83*(3), 595–614.

Adukia, A., S. Asher, and P. Novosad (2017). Educational investment responses to economic opportunity: Evidence from indian road construction. *Becker-Friedman Institute Working Paper Series*.

Aitken, M., A. Backliwal, M. Chang, and A. Udeshi (2013). Understanding Healthcare Access in India: What is the current state? Technical report, IMS Institute for Healthcare Informatics.

Alkire, B. C. and N. P. e. a. Raykar (2015, June). Global access to surgical care: a modelling study. *Lancet Global Health e316-323*.

Allen, T. and C. Arkolakis (2019). The welfare effects of transportation infrastructure improvements. *NBER Working Paper Series WP 25487*.

Banerjee, A., A. Deaton, and E. Duflo (2004, May). Wealth, Health, and Health Services in Rural Rajasthan. *American Economic Review 94*(2), 326–330.

Banerjee, A., E. Duflo, and R. Hornbeck (2014). Bundling Health Insurance and Microfinance in India: There Cannot Be Adverse Selection If There Is No Demand. *American Economic Review 104*(5), 291–97.

Banerjee, A., E. Duflo, and N. Qian (2012). On the road: Access to transportation infrastructure and economic growth in china. *NBER Working Paper Series WP 17897*.

Bell, C. (2012). The benefits of india's rural roads program in the spheres of goods, education and health. *World Bank Policy Research Working Paper WP 6169*.

Bell, C. and S. van Dillen (2018). One the way to good health? rural roads and morbidity in upland orissa. *Journal of Transport and Health 10*, 369–380.

Chandra, A. and E. Thompson (2000). Does Public Infrastructure Affect Economic Activity? Evidence from the rural interstate highway system. *Regional Science and Urban Economics 30*, 457–490.

Coffey, R. M. (1983). The Effect of Time Price on the Demand for Medical-Care Services. *The Journal of Human Resources 18*(3), 407–424.

Das, J. and J. Hammer (2007). Money for nothing: The dire straits of medical practice in Delhi, India. *Journal of Development Economics 83*, 1–36.

Das, J., A. Holla, A. Mohpal, and K. Muralidharan (2016). Quality and Accountability in Healthcare Delivery: Audit-Study Evidence from Primary Care in India. *American Economic Review (Forthcoming)*.

Das, J. and J. Leino (2011, August). Evaluating the RSBY: Lessons from an Experimental Information Campaign. *Economic and Political Weekly 46*(32), 85–93.

Das, Jishnu; Hammer, J. S.-P. C. (2011). The Impact of Recall Periods on Reported Morbidity and Health Seeking Behavior. *World Bank Policy Research Working Paper*.

Datta, S. (2012). The impact of improved highways on Indian firms. *Journal of Development Economics 99*, 46–57.

Deshpande, M. and Y. Li (2019). Who Is Screened Out? Application Costs and the Targeting of Disability Programs. *American Economic Journal: Economic Policy 11*(4), 213–248.

Dupas, P. and J. Robinson (2013). Why Don't the Poor Save More? Evidence from Health Savings Experiments. *American Economic Review 103*(4), 1138–1171.

Faber, B. (2014). Trade Integration, Market Size, and Industrialization: Evidence from China's National Trunk Highway System. *Review of Economic Studies 81*(3), 1046–1070.

Fernald, J. G. (1998). Roads to Prosperity? Assessing the Link between Public Capital and Productivity. *The American Economic Review 89*, 619–638.

Finkelstein, A., S. Taubman, B. Wright, M. Bernstein, J. Gruber, J. P. Newhouse, H. Allen, K. Baicker, and O. H. S. Group (2012, 07). The Oregon Health Insurance Experiment: Evidence from the First Year. *The Quarterly Journal of Economics 127*(3), 1057–1106.

Fossett, J. W. and J. A. Peterson (1989, April). Physician Supply and Medicaid Participation: The Casues of Market Failure. *Medical Care 27*(4), 386–396.

Gertler, P. J., M. Gonzalez-Navarro, T. Gracner, and A. D. Rothenberg (2016). Road Quality, Local Economic Activity, and Welfare: Evidence from Indonesia's Highways. *CEGA Working Paper Series*.

Ghani, E., A. G. Goswami, and W. R. Kerr (2016). Highway to success: The impact of the golden quadrilateral project for the location and performance of indian manufacturing. *The Economic Journal 126*(591), 317–357.

Gomes, M., R. Begum, P. Sati, R. Dikshit, P. C. Gupta, R. Kumar, J. Sheth, A. Habib, and P. Jha (2017, November). Nationwide Mortality Studies To Quantify Causes of Death: Relevant Lessons From India's Million Death Study. *Health Affairs 36*(11), 1887–1895.

Grossman, M. (1972). *The Demand for Health: A Theoretical and Empirical Investigation*. NBER.

Gudwani, A., P. Mitra, A. Puri, and M. Vaidya (2012). India Healthcare: Inspiring Possibilities, Challenging Journey. Technical report, McKinsey&Company Healthcare Systems and Services.

Gujral, K. and A. Basu (2019). Impact of Rural and Urban Hospital Closures on Inpatient Mortality. *NBER Working Paper 26182*.

Khanna, G. (2016). Road oft taken: The route to spatial development. *Working Paper*.

Krishnaswamy, K. and R. Ruchismita (2011, September). Performance trends and policy recommendations. Technical report, RSBY Working Paper.

La Forgia, G. and S. Nagpal (2012). *Government-Sponsored Health Insurance in India - Are You Covered?* Directions in Development. Washington, DC: World Bank.

Lu, C., C. B, G. Li, and C. Murray (2009). Limitations of methods for measuring out-of-pocket and catastrophic private health expenditures. *Bulletin of the World Health Organization 87*, 238–244.

Michaels, G. (2008, November). The Effect of Trade on the Demand for Skill: Evidence from the Interstate Highway System. *The Review of Economics and Statistics 90*(4), 683–701.

Murray, C. J. L. and L. C. Chen (1992). Understanding Morbidity Change. *Population and Development Review 18*(3), 481–503.

Newhouse, J. P. e. a. (1986). Use of Medical Care in the Rand Health Insurance Experiment. Technical report, RAND.

Perloff, J. D., P. Kletke, and J. W. Fossett (1995, April). Which Physicians Limit Their Medicaid Participation, and Why. *Health Services Research*.

Perloff, J. D., P. R. Kletke, J. W. Fossett, and S. Banks (1997, February). Medicaid Participation among Urban Primary Care Physicians. *Medical Care 35*(2), 142–157.

Riley, J. C. (1989). *A History and Forecast of Ill Health.* Iowa City: University of Iowa Press.

Rothenberg, A. D. (2013). Transport Infrastructure and Firm Location Choice in Equilibrium: Evidence from Indonesia's Highways. *Working Paper*.

Sen, A. (2002). Health: Perception versus observation: Self reported morbidity has severe limitations and can be extremely misleading. *BMJ: British Medical Journal 324*(7342), 860–861.

Shoree, S., R. Ruchismita, and K. R. Desai (2014, July). Evaluation of RSBY's Key Performance Indicators. Research Paper 42, International Labour Office.

Shrime, M. G., M. Hamer, and S. e. a. Mukhopadhyay (2017). Effect of removing the barrier of transportation costs on surgical utilisation in guinea, madagascar and the republic of congo. *BMJ Global Health 2*(e000434).

Varela, C. e. a. (2019). Transportation Barriers to Access Health Care for Surgical Conditions in Malawi: A cross sectional nationwide household survey. *BMC Public Health*.

Xu, K., F. Ravndal, D. B. Evans, and G. Carrin (2009). Assessing the reliability of household expenditure data: Results of the World Health Survey. *Health Policy 91(3)*, 297–305.

Environmental Externalities and Free-riding in the Household[9]

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.

[9]Coauthored with Seema Jayachandran and Kelsey Jack. A published working paper version of this paper can be found at https://www.nber.org/papers/w24192.

9 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 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.[10] 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.[11] Each individual decides how much costly effort to put into water conservation, taking others' decisions as given. Because the individual

[10] Surface 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).

[11] As 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.

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.

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.[12] 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

[12]All 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.

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.[13]

Overall, we find an average decrease in monthly water use of 6.2 percent in response to the 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.[14] 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.

[13]We also investigate two other features of the billing environment that might affect households' response to our price incentive treatment, described in detail in Section 11.3. 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.

[14]This 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.

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 (2018), heterogeneity 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.[15] We thus add to literature on misperceptions of price (Ito 2014; McRae and Meeks 2016), lack of information about the price (Jessoe 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).[16] Intrahousehold inefficiencies might be especially important in poor countries, where gender roles are particularly imbalanced. However, while our empirical application 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

[15] Rungie 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.

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

consumption 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.[17] 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).[18] 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 11 describes the experimental design and implementation. Section 12 presents the results, and Section 13 discusses the implications for optimal pricing. Section 14 concludes.

[17] Respondents 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 **??**.

[18] Many 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 10.4. A much larger literature has examined, and often rejected, Pareto efficiency in investment and production (e.g., Udry (1996)).

10 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 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.

10.1 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.[19] The water utility charges the household pW, where $W \equiv \sum_i w_i$. The household has total income Y and we assume

[19]This 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.

$pW < Y$. 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:

$$\max_{e_i} Y - 2p\bar{w}(1 - e_i) - 2ce_i^2. \tag{7}$$

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

10.2 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 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,

$$w_i^* = \bar{w}[1 - \frac{p}{2c}\bar{w}(\lambda_i + \alpha_i(1 - \lambda_i))] \tag{8}$$

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 (8) shows that w_i^* is decreasing in p, α_i and λ_i.A higher price, more altruism toward 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 $\lambda_i = 1$.

10.3 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 results 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: $\frac{\partial^2 w_i^*}{\partial p \partial \alpha_i}|_{\lambda_i \to 0} < \frac{\partial^2 w_i^*}{\partial p \partial \alpha_i}|_{\lambda_i \to 1} < 0$. 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 $v_i = \lambda_i(Y - pW) - ce_i^2 - P_i W$. 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 $\lambda_i \to 0$, 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.

10.4 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.[20] Adding

[20] This improvement in intrahousehold observability may explain part of the decline in electricity use

to these observability challenges, knowing one's own consumption is often difficult.[21] 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: $\hat{w_{-i}} = W - (w_i + \epsilon)$. Moreover, the fact that some part of water consumption is a public good at the household level (e.g., washing the family's dinner dishes) further complicates the problem of quantifying others' effort toward conservation. (Note that even when water is used to produce public goods, there is still some "private"' consumption if, conditional on how clean you get the dishes, washing them in a manner that wastes less water requires more effort and hence higher private costs.)

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.

associated with the introduction of smart metering (e.g., Jessoe and Rapson 2014).

[21] The fact that even one's own consumption is difficult to gauge means that, even leaving aside the free-riding problem within a group, an individual might not consume the amount of water she is targeting. For example, if there were a prize for reducing water, a person living alone might 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.

11 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.

11.1 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).[22] 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.[23] 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 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.

[22]Tariffs 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).

[23]Households 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.

The tariff schedule for 2015, when our study took place, is shown in Appendix Figure ??. 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.[24] 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.[25] 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 12, which also tests for balance across the treatments, as discussed below.

11.2 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 4 summarizes the experimental design.[26]

The treatment was implemented in conjunction with a household survey, run between May and December 2015. Households in the *incentive treatment* are provided with

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

[25] Customers 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.

[26] The 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.

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, 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.[27]

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.[28] 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.[29] 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

[27]The 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.

[28]The 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.

[29]Note 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.

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 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 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.

11.3 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.

11.3.1 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?*"[30] Treated households received information that cutting back by 20 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.[31] Thus, all households enrolled in the incentive treatment also receive the price information treatment.

[30]The 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.

[31]Note 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.

11.3.2 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.[32]

11.4 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.

[32]The 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."

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.[33] Sharing with the spouse also reflects enforcement-based income sharing within the household, which may or

[33]For 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.

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 measured does not allow us to separate altruism and enforcement.[34]

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

[34]It 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.

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.

11.5 Sample construction and summary statistics

11.5.1 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),[35] 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

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

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 our screening criteria.[36] 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 **??** shows how sample characteristics evolve at each stage of sampling and randomization.[37]

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

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

[37]As a robustness check, presented in Section 12.5, 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).

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.

11.5.2 Sample statistics

Table 12 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 household 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.

11.5.3 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 12. 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 12 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 **??**). 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 13 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).[38] 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.

12 Results

12.1 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:[39]

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

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

[39]The theoretical predictions in Section 10 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 10 over water use levels while our empirical results test for effects on log water use. Rewriting the model in logs generates the same predictions.

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.

12.2 Estimation strategy

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

$$y_{it} = \beta_1 treatedhh_i + \beta_2 post_{it} + \beta_3 treatedhh_i \times post_{it} + \epsilon_{it} \tag{9}$$

where $treatedhh_i$ is a binary indicator for whether the household was assigned to the relevant treatment group and $post_{it}$ is a time-varying indicator that turns on for household i in the month after the survey. Note that $post_{it}$ 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 $post_{it}$ 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 $post_{it}$ 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.[40]

To improve precision, we include neighborhood by time and household fixed effects in our preferred estimates. Defining $treat_{it} \equiv treatedhh_i \times post_{it}$, we estimate:

$$y_{it} = \beta_1 treat_{it} + \beta_2 post_{it} + \tau_t + \eta_i + \epsilon_{it} \tag{10}$$

[40] Note 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.

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 $treat_{it}$ and $post_{it}$ with relevant household characteristics.

To illustrate magnitudes and as an input to our policy calibrations in Section 13, we use the estimates of β_1 associated with our incentive treatment in equation (10) to calculate short run price elasticities as follows.[41] First, with y_{it} equal to log of monthly water quantity, we 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.[42]

We show the exogeneity of treatment assignment to observable household characteristics in Table 12 for the incentive treatment and Appendix Tables **??** and **??** for the information and provider credibility treatments. We also plot average water use across treatment

[41] The elasticity calculation requires a number of assumptions: (1) that households respond similarly to a discrete price change as to a continuous price change, (2) that households respond similarly to a quantity target as to a continuous price change, and (3) that households respond similarly to a probabilistic payout as to a (smaller) certain payout from conservation.

[42] Note 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.

conditions and our dictator-game measure of intrahousehold efficiency in the months leading up to the survey (Appendix Figure ??). 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 ??) 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 13, 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 ??. 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 12.5.

12.3 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 treatment reference window. Figure 5 shows a decrease in consumption across most of the distribution.[43] It is worth noting that there is relatively little mass in either the treatment

[43] Appendix Figure ?? shows the distribution of consumption pre- and post-treatment for the price incentive treatment and the provider credibility 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 14 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 (9), without household or time fixed effects. Columns 2, 3 and 4 add household, month-year and zone-month-year fixed effects, respectively. The main effects 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 12.1): 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).[44] 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 (10).

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

[44] Our 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.

about water prices and about the correspondence between water use and bills, respectively. Thus, we supplement the analysis in Table 14 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 water provider. Both should lead to reductions in water use because they imply that the treatments increased the effective perceived water price. Table 15 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.[45]

12.4 Intrahousehold heterogeneity

12.4.1 Household price incentives

Table 16 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

[45] Our 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.

(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 **??** shows the robustness of these results to alternate approaches to aggregating the dictator game measure.[46]

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 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.

12.4.2 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

[46] It 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.

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 17 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 12, 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 17 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 18, aiming to disentangle whether the effect is due to gender per se or the recipients' 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 responsive to the conservation incentive. Overall, the results presented in Table 18 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 14.

12.4.3 Altruism versus enforcement

As discussed in Section 10.4, 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 11.4 and Appendix Table **??** for correlations of these measures with the dictator game outcomes). Table 19 shows the results from replacing the dictator game measure in equation (10) 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).[47] 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

[47] Appendix Table **??** shows how the variables contribute to the principal components.

whether altruism or enforcement matters more.

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 20 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.[48] 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.

12.5 Robustness checks

12.5.1 Interpretation of the heterogeneity results

As shown in Tables 13 and **??**, 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 21, 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).[49] While we do see some heterogeneity in price sensitivity by these

[48] See Section 11.4 for further details on these variables.

[49] Appendix Table **??** shows how the survey variables contribute to the principal components.

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 21 is therefore a conservative test of our main result.

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 18), 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).

12.5.2 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 **??** 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 **??** 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 **??**), but only a subset of our sample (those surveyed first) contribute to these effects. Finally, Appendix Table **??** 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 **??** 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 45, we pool treatments to improve power. Appendix Tables **??** and **??** 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 **??** 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 **??** 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.

13 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 (2018).[50] 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 (2018) 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 $\underline{\mathrm{w}} < \bar{w}$. Individual i chooses to exert effort if and only if $\frac{1}{2}(1+\alpha_i)(Y - p(\underline{\mathrm{w}} + E(w_j))) - c \geq \frac{1}{2}(1+\alpha_i)(Y - p(\bar{w} + E(w_j)))$. 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 $\underline{\mathrm{w}}$, 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

[50] To 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.

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 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 (intrahousehold 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 intrahousehold efficiency. In other words, we assume that our high θ households are free of intrahousehold 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.[51] Note that this requires an

[51] As discussed in Section 12.2, 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,

additional assumption that households' distortion-free price elasticity is uncorrelated with our measure of intrahousehold 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 22. The planner observes an empirical demand curve associated with the population average $\varepsilon_p = -0.27$.

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:

$$W(\tau) = \int_j [Y - (p+\tau)(\mathcal{I}_e \underline{\mathrm{w}} + (1-\mathcal{I}_e)\bar{w}) + (v-\chi)(\mathcal{I}_e \underline{\mathrm{w}} + (1-\mathcal{I}_e)\bar{w}) - c\mathcal{I}_e] + \tau D. \quad (11)$$

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:[52]

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.

[52]This 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.

$$p^* = (p+\chi)\frac{E(\theta)}{E(\theta)^2 + Var(\theta)}. \tag{12}$$

The optimal price given in equation (12) is increasing in $E(\theta)$ and decreasing in $Var(\theta)$.[53] 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.

Calibration

Calibrating the optimal price in our setting is simply a matter of substituting the parameters in Table 22 into equation (12). 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.

[53] If 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.

14 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 (2018), 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.

Figures

Figure 4: Experimental design

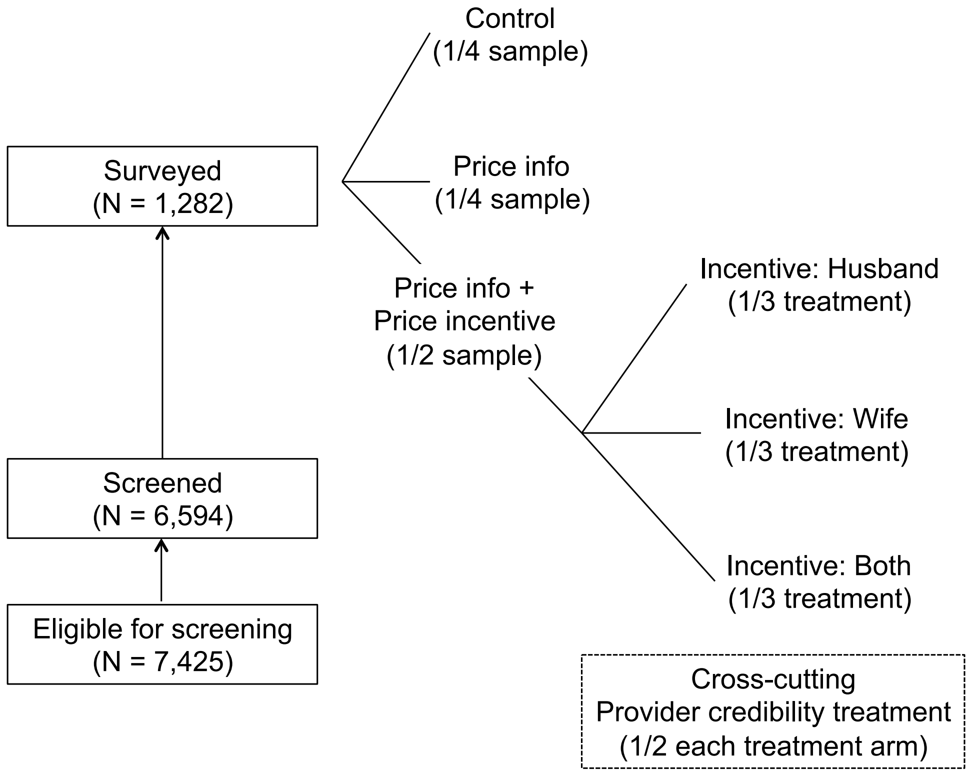

Notes: Experimental design and sampling flow. Treatment was assigned on a rolling basis to accommodate the high rate of ineligibility that led screened households to be disqualified from the survey sample.

Figure 5: Water consumption, relative to incentive reference month

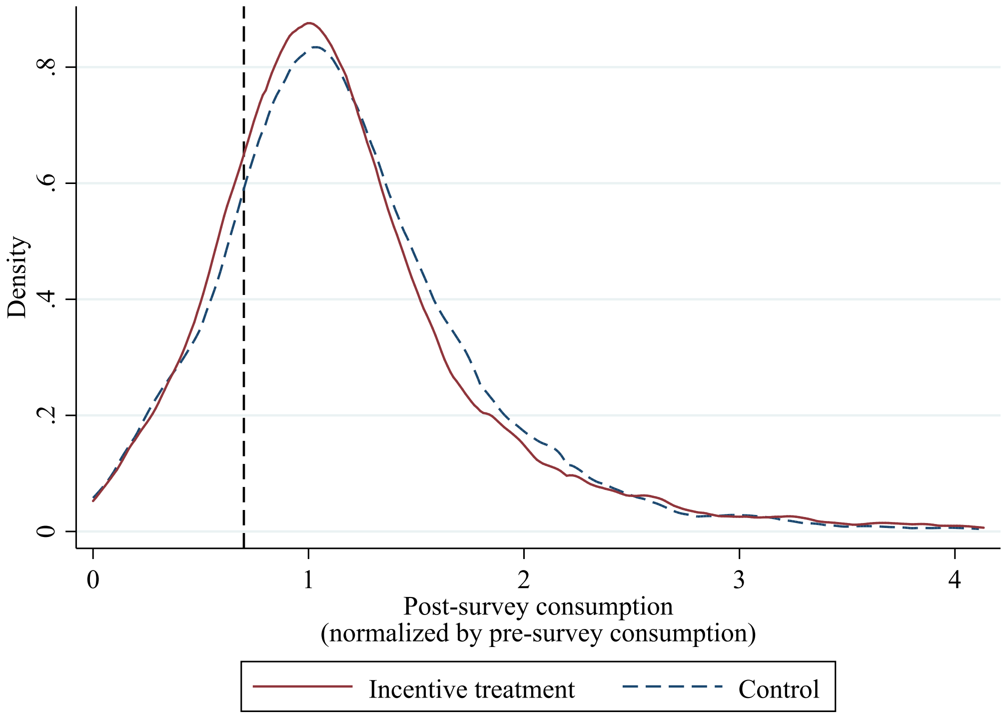

Notes: Density plots of post-survey monthly consumption relative to the average monthly consumption in the reference months used to determine price incentive eligibility. The control group includes all surveyed households not assigned to the incentive treatment. The dashed vertical line shows the 70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

Tables

Table 12: Sample statistics & balance

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

	Husband share sent (1)	Wife share sent (2)	Sent above median (3)
Quantity consumed	0.002***	0.002***	0.004***
	(0.001)	(0.001)	(0.001)
Prob(payment)	-0.044	-0.113**	-0.187**
	(0.045)	(0.045)	(0.086)
Prob(missing)	-0.064	0.079*	0.051
	(0.047)	(0.047)	(0.090)
Total charge	0.025**	0.023*	0.063***
	(0.013)	(0.012)	(0.024)
Household size	-0.004	-0.009***	-0.015**
	(0.003)	(0.003)	(0.006)
HH has maid	0.019	0.068***	0.086**
	(0.020)	(0.020)	(0.038)
HH assets	0.008***	0.018***	0.030***
	(0.003)	(0.003)	(0.005)
HH owns home	-0.015	-0.034**	-0.060**
	(0.015)	(0.015)	(0.028)
HH rooms in home	0.014***	0.015***	0.034***
	(0.005)	(0.005)	(0.010)
HH English fluency	0.022	0.082***	0.112***
	(0.017)	(0.017)	(0.033)
Either underestimated price	0.009	0.017	0.033
	(0.017)	(0.017)	(0.032)
Both blame high bill on SWSC	0.012	0.016	0.000
	(0.015)	(0.015)	(0.028)
Both know bill quantity	0.005	0.014	0.046
	(0.022)	(0.022)	(0.043)
Both know bill charge	0.010	-0.004	-0.004
	(0.016)	(0.016)	(0.030)
W: Effective bill payer	0.003	0.001	0.013
	(0.016)	(0.016)	(0.030)
W: Bigger water user	0.003	-0.006	-0.000
	(0.019)	(0.019)	(0.036)
H: Share NGO	0.192***	0.079***	0.262***
	(0.030)	(0.030)	(0.057)
W: Share NGO	0.034	0.198***	0.269***
	(0.033)	(0.033)	(0.063)
H SDB score	0.004	0.002	0.006
	(0.003)	(0.003)	(0.005)
W SDB score	0.003	-0.003	0.001
	(0.003)	(0.003)	(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 14: Average effects of all treatments

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

Notes: 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 15: Heterogeneous effects of price information and provider credibility treatments

	log (Quantity) (1)	log (Quantity) (2)
Price information treatment	0.025	
	[0.048]	
Info x Underestimated price	-0.063	
	[0.060]	
Provider credibility treatment		0.009
		[0.033]
Provider credibility x Distrust billing		0.048
		[0.048]
Observations (HH)	1,282	1,282
Observations (HH-months)	25,506	25,506

Notes: 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 16: Heterogeneous effects of incentive treatment by intrahousehold efficiency

	log(Quantity)				
	(1)	(2)	(3)	(4)	(5)
Incentive treatment	-0.034	-0.023	-0.036	-0.098***	-0.079
	[0.048]	[0.031]	[0.050]	[0.031]	[0.050]
Incentive x Sent > median on avg	-0.036	-0.067			
	[0.080]	[0.048]			
Incentive x Husband sent > median			0.019		
			[0.081]		
Incentive x Wife sent > median			-0.098		
			[0.102]		
Incentive x Bigger user is RC				0.054	-0.023
				[0.062]	[0.106]
Husband = Wife (p-val)			0.413		
Incentive	Couple	Pooled	Couple	Pooled	Couple
Sample	Full	Full	Full	Restricted	Restricted
Observations (HH)	843	1,275	843	1,007	658
Observations (HH-months)	16,893	25,388	16,893	19,998	13,111

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 17: Price incentives directed toward the wife, husband, or couple

	log (quantity) (1)
Couple incentive	-0.046
	[0.036]
Husband incentive	-0.042
	[0.034]
Wife incentive	-0.089***
	[0.033]
Couple = Wife (p-val)	0.337
Husband = Wife (p-val)	0.266
HH FE	x
Zone-Month-Year FE	x
Observations (HH)	1,282
Observations (HH-months)	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 18: 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	-0.047	-0.019	-0.019
	[0.033]	[0.032]	[0.035]	[0.036]
Incentive x Non-bill-payer	-0.091**		-0.087*	-0.086*
	[0.041]		[0.045]	[0.044]
Incentive x Bigger user		-0.044	-0.012	
		[0.042]	[0.045]	
Incentive x Wife				-0.011
				[0.044]
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 19: Heterogeneity by survey measures of intrahousehold decisions

	log (quantity) (1)	log (quantity) (2)	log (quantity) (3)
Incentive x Sent above median	-0.080*	-0.083*	-0.083*
	(0.048)	(0.049)	(0.049)
Incentive x Decide budget together	0.020	0.008	
	(0.048)	(0.052)	
Incentive x Decide extra spending together	-0.022	-0.058	
	(0.062)	(0.066)	
Incentive x Never disagree	0.048	0.044	
	(0.053)	(0.056)	
Incentive x Make plans together	0.068	0.073	
	(0.051)	(0.053)	
Incentive x Stick to plans	-0.027	-0.048	
	(0.056)	(0.059)	
Incentive x Respondent never deviates	0.001	-0.092	
	(0.051)	(0.068)	
Incentive x Spouse never deviates	0.105**	0.151**	
	(0.050)	(0.068)	
Incentive x Know if spouse deviates	-0.066	-0.100*	
	(0.050)	(0.057)	
Incentive x Can prevent spouse from deviating	0.028	0.067	
	(0.049)	(0.054)	
Incentive x Does things spouse wants	0.011	0.046	
	(0.048)	(0.057)	
Incentive x Spouse does things respondent wants	-0.025	-0.013	
	(0.047)	(0.057)	
Incentive x Never hide income	0.064	0.058	
	(0.048)	(0.052)	
Incentive x Hard to hide income	-0.007	-0.007	
	(0.050)	(0.051)	
Incentive x 1st principal component	0.012		0.012
	(0.015)		(0.014)
Incentive x 2nd principal component	0.004		0.009
	(0.018)		(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 20: Heterogeneity by knowledge and monitoring of water use

	log (quantity) (1)	log (quantity) (2)	log (quantity) (3)	log (quantity) (4)
Incentive treatment	-0.056**	-0.056	-0.039	-0.013
	[0.027]	[0.045]	[0.030]	[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 avg				-0.120
				[0.089]
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 more than the median share of their spouse's 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 21: Robustness check: Controlling for household characteristics

	log (quantity) (1)	log (quantity) (2)
Incentive x Sent above median	-0.080*	-0.090*
	(0.048)	(0.050)
Incentive x Quantity	-0.002	0.002
	(0.002)	(0.005)
Incentive x Prob(payment)	0.050	-0.081
	(0.170)	(0.163)
Incentive x Prob(missing)	-0.167	-0.202
	(0.198)	(0.220)
Incentive x Total charge	-0.023	-0.009
	(0.040)	(0.092)
Incentive x Household size	-0.007	-0.006
	(0.010)	(0.012)
Incentive x Maid	-0.039	0.000
	(0.066)	(0.066)
Incentive x Assets	-0.016*	-0.017
	(0.009)	(0.010)
Incentive x Owns home	0.003	0.024
	(0.048)	(0.052)
Incentive x Rooms in home	-0.034*	-0.033
	(0.019)	(0.022)
Incentive x English fluency	0.027	0.088
	(0.059)	(0.068)
Incentive x Underestimated price	-0.049	-0.012
	(0.045)	(0.044)
Incentive x Blame SWSC	0.010	0.015
	(0.048)	(0.048)
Incentive x Know quantity	-0.060	-0.068
	(0.068)	(0.070)
Incentive x Know charge	0.012	0.040
	(0.052)	(0.051)
Incentive x W is effective bill payer	0.080	0.087*
	(0.049)	(0.049)
Incentive x W is bigger user	-0.110*	-0.123*
	(0.066)	(0.064)
Incentive x Sent above median to NGO	-0.065	-0.035
	(0.048)	(0.049)
Incentive x Above median SDB score	0.044	0.050
	(0.048)	(0.049)
Observations (HH)	1,275	1,274
Observations (HH-months)	26,122	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 22: 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.

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. *National Bureau of Economic Research Working Paper*.

Elinder, M., S. Escobar, and I. Petré (2017). Consequences of a price incentive on free riding and electric energy consumption. *Proceedings of the National Academy of Sciences 114*(12), 3091–3096.

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), 37–62.

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), 63–88.

Jessoe, K. and D. Rapson (2014). Knowledge is (less) power: Experimental evidence from residential energy use. *American Economic Review 104*(4), 1417–1438.

Kahn, M. E. and F. A. Wolak (2013). Using information to improve the effectiveness of nonlinear pricing: Evidence from a field experiment. *Working Paper*.

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 (2018). Attention variation and welfare: Theory and evidence from a tax salience experiment. *Review of Economic Studies 85*, 2462–2496.

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.