Liquid markets: Allocating clean water to the rural poor

Pre-analysis plan

Fiona Burlig, Amir Jina, and Anant Sudarshan*

September 8, 2023

Abstract

How should scarce resources be allocated in developing countries? Weitzman (1977) highlights a trade-off between prices, which generate allocative efficiency, and quotas, which might have desirable distributional consequences. In partnership with a private company supplying clean water to rural Odisha, India, we plan to run an experiment to measure the relative effectiveness of different allocation mechanisms. We will measure the price elasticity of demand for clean water, health effects from consuming clean water, and the extent to which liquidity constraints and intra-household inefficiencies reduce consumption. To do so, we implement a cluster-randomized trial, where 160 villages are randomized into a pure control group and multiple treatment arms: (i) discounts; (ii) a monthly quota; and (iii) an exchangeable quota, where unused allocation can be exchanged for cash. We subsequently randomize which households within each treatment village will receive treatment. We plan to measure effects of treatment on water consumption and health outcomes using a combination of survey and administrative data. Following the main experiment we will also provide information on water quality to a subset of households and examine whether this changes other behaviors associated with keeping water clean.

Key words: Water; resource allocation; quotas
JEL Codes: D45; O13; Q25; Q53

*Burlig: Harris School of Public Policy and Energy Policy Institute (EPIC), University of Chicago, and NBER. Email: burlig@uchicago.edu. Mailing address: Keller Center, 1307 E 60th St., Chicago, IL 60637. Jina: Harris School of Public Policy and EPIC, University of Chicago, and NBER. Email: amirjina@uchicago.edu. Mailing address: Keller Center, 1307 E 60th St., Chicago, IL 60637. Sudarshan: Department of Economics, University of Warwick. Email: anant.sudarshan@warwick.ac.uk. Mailing address: S2.92 Department of Economics, Social Science, University of Warwick, Coventry, CV4 7AL, UK. We thank Kishan Nanavati and Krupa Paltasingh of Spring Health Water India Pvt. Ltd. for their partnership on this project. We thank Kashif Ahmed, Zoya Khan, Ambalika Khanna, Nisha Koppa, Yashaswi Mohanty, and Chinmaya Sahoo for excellent research assistance. Nadia Lucas provided key research design development assistance. We gratefully acknowledge financial support from the Oak Foundation, the Templeton World Charity Foundation, and the Becker Friedman Institute Development Economics Research Fund. All remaining errors are our own. This project received IRB approval from the University of Chicago (Protocol No. IRB22-0036), and is registered on the AEA RCT registry (Identification No. 10545). The pre-analysis plan for this project can be found at https://www.socialscienceregistry.org/trials/10545. **This version of the PAP updates the version posted to the AEA registry on December 5, 2022 to detail our use of the endline survey. We did not have access to the endline data prior to this update.**

Contents

1 Introduction

In 2022, we began a cluster-randomized experiment in Odisha, India, in partnership with Spring Health India Pvt. Ltd — a private company that sells clean water to rural households in Odisha, India.

The clean water sold by Spring Health comes originally from local sources and is then treated using electrochlorination, in a treatment plant powered using decentralized solar electricity. For most Spring Health water treatment plants, the input source of water is groundwater. In most cases there is one treatment facility per village, although a small share of villages share a treatment plant. Spring Health trains a local entrepreneur who operates and maintains the treatment facility and is normally also the owner of the well providing the input water.

The electro-chlorination process is intended to remove coliform and other organic contaminants, not contamination due to arsenic compounds, cynanide etc. However, we are aware of no evidence of arsenic contamination of groundwater in the areas of Odisha where SH functions. In this plan, unless explicitly specified, we use 'treated water' or 'clean water' to indicate water that has been treated to remove coliform, or that has passed a coliform water test.

In the status quo business model, any household in a village served by Spring Health can enrol in clean water deliveries. Payments are made against orders fulfilled and households may purchase varying amounts of water in the month and may discontinue purchases at any point. Spring Health water is packaged in 20-liter bottles each priced at INR 25 (about 30 cents) or 10-liter bottles priced at INR 14 (different villages use different bottles so households cannot normally mix and match). Bottles are delivered directly to households. This model differs from many previous studies that have studied the economics of clean water in developing countries in two important ways. First, the method of obtaining clean water does not involve chlorination at home. Second, there is little risk of contamination between the source of clean water and the household because transport is handled by Spring Health in bottles designed for this purpose.

2 Experimental design

In partnership with Spring Health, we will conduct a field experiment designed to test different ways of allocating clean water to households. These include (i) offering water for sale at different prices with households able to purchase as many bottles as they desire, (ii) providing a water quota for free, or (iii) an entitlement model where households are provided an entitlement equal to the water quota but can choose not to consume up to the entitlement limit in exchange for cash rebates. Figure 1 displays the design graphically, and we describe the experiment in more detail below.

2.1 Sample Selection

Working with Spring Health, we selected 160 villages as the site of the experiment. These villages were selected because they have a sustained Spring Health presence. More precisely, all villages in the experiment had been served by Spring Health for at least 24 months prior to the beginning of the study, and Spring Health did not expect operational challenges in carrying out the experiment. These villages were then randomly assigned to four treatment arms to create a cluster-randomized design as described below.

Pure control (40 villages) Villages in this arm are completely 'business as usual' Throughout the experiment, all households in these villages will be able to purchase Spring Health water at the prevailing market price. No surveys will be conducted in these villages. Thus our only source of data on these villages comes from Spring Health's administrative data, which is generated through the normal course of business.

Free quotas (40 villages) Each household in villages assigned to the 'free quota' condition will be randomly assigned to receive either an offer of 400 litres of water per month per month for the duration of the experiment or nothing[1]. Households who do not receive a free quota continue to have the option of purchasing clean water from Spring Health at full price. Likewise, households who receive a free quota can purchase additional water at full price if they exhaust their quota. Households can also opt not to use some or all of their quota.

Household level randomization will be conducted by distributing scratch cards drawn from a shuffled deck to all homes in the village. The deck for every village will include 39 treatment cards (a non-zero quota). However, since the total number of distributed cards varies by the estimated number of households in the village (obtained through a listing exercise), the probability of being treated is not constant. Figure 2 shows an example scratch card. Note that all scratch cards look identical so the field operators distributing cards do not know whether they contain an offer until the household uses them. Each card has a unique ID number but a mapping from this number to the offer on the card was only available to the research team. This mapping was used by the research team to check the number of offers in each deck before delivering them to field operators.

Price discounts (40 villages) In price discount villages, 13 households will be randomly assigned to receive an offer of a 10% discount for the duration of the experiment, 13 will be assigned to receive an offer of a 50% discount, and 13 will be offered a 90% discount. The remaining households will receive no discount. Households who do not receive a discount can continue to purchase Spring Health water at the market price. These discounts will also be implemented by distributing scratch cards from a shuffled deck.

1. The mean consumption for existing consumers in a few pilot villages for which Spring Health shared data was 360 litres per month

Figure 1: Experimental design diagram

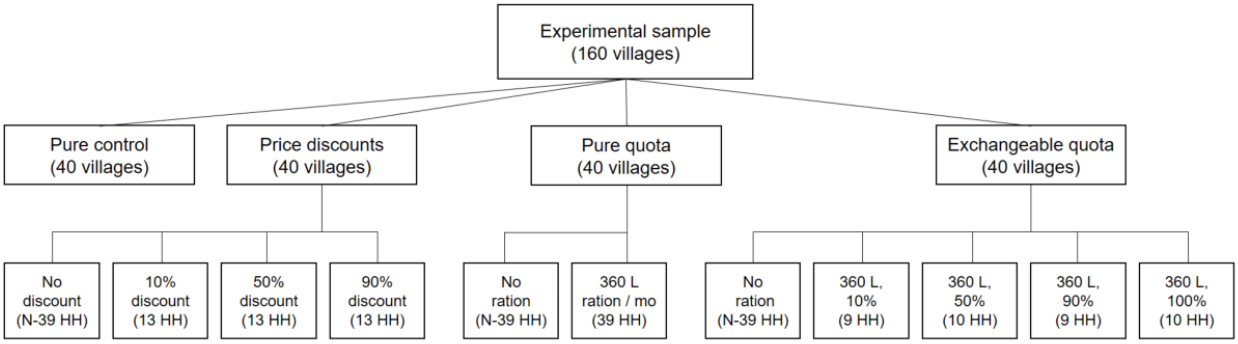

Notes: This figure displays the details of our experimental design: we randomly assign 160 villages to a pure control group, price discount group, pure quota group, or exchangeable quota group. Within the three treatment groups, we randomly assign households to sub-treatments. In price discount villages, every household is offered either no discount, a 10% discount, a 50% discount, or a 90% discount. In pure quota villages, every household is offered either no quota or a quota of 400 litres per month. In exchangeable quota villages, every household is offered no quota or a 400 litre per month quota with a cash-back value of 10%, 50%, 90%, or 100% of the market price.

Exchangeable quota (40 villages) In exchangeable quota villages, 38 households in each village will be given the offer of a 400 litre quota just as in the free quota condition. The difference in this treatment arm was that households could redeem unused water below this quota for cash. Specifically, for every unclaimed bottle of water, households were entitled to receive a payment equal to 10% of the market price (9/38 households), 50% of the market price (10/38), 90% of the market price (9/38 households), or 100% of the market price (10/38 households). As in the other arms, households who did not receive a non-zero exchangeable quota remained eligible to buy Spring Health water at full price. Note that in all treatment arms – including the exchangeable quota arm – to receive water, households needed to place order(s) with Spring Health, so they are not "endowed" with a physical quantity of water.

Free one-time 100L (all treatment villages) In addition to the treatments described above, 5 households in each treatment village will be randomized to receive a one-time offer of 100 free litres of water. Households who chose not to take this offer will receive nothing. These households could otherwise purchase water at the market price. We include this arm to understand the extent to which Spring Health water is an experience good, by comparing the number of bottles consumed by these households after this offer to control households – any difference greater than the initial free trial, obtained through regular purchases, will suggest that exposure to Spring Health water drives future purchases. These households will not otherwise play a role in our main analyses.

Figure 2: Sample scratch card

Notes: This figure shows an example promotional scratch card used to randomize households into treatment groups. The left shows the front of the card, which is common among all offers. The right shows the back of the card, which differs across offers.

2.2 Timeline

For logistical reasons, we will implement the experiment in phases. We will randomly divide villages into phases, preserving equal treatment and control group sizes in each phase. Figure 3 shows the implementation timeline, including scratch card distribution, water distribution, and survey data collection (described in more detail in Section 3 below).

3 Data collection and outcome variables

We will bring together several sources of data to conduct our analysis.[2]

3.1 Record of scratch-card distribution

While distributing scratch cards to each household, we record the village name, the name of the household head, whether the household is an existing Spring Health customer, contact and address details, the type of offer and the identification number of the scratch card, allowing us to link a household to its treatment offer. We will use these data to estimate intent-to-treat effects, as described in Section 4 below. We will also use this data set to carry out a set of randomization checks.

When a customer contacts Spring Health to claim their offer, their staff record their address and phone number if they consent, as well as their scratch card identification number.

2. The PI team had no access to these data when this PAP was first uploaded to the AEA registry. This document was later augmented with additional specifications using the end-line survey, before the survey was completed. Note that none of the original specifications were changed.

Figure 3: Tentative experimental timeline

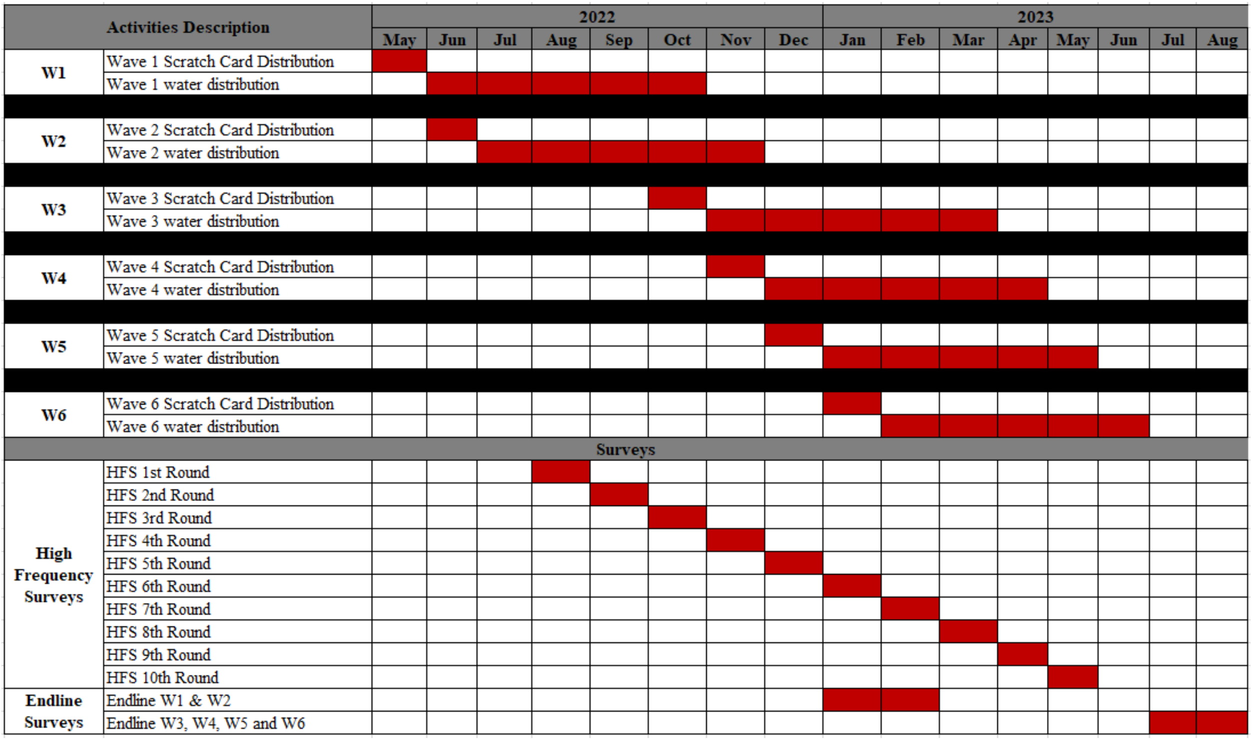

Notes: This figure shows the experimental timeline in the form of a Gantt chart.

3.2 Administrative data

One of our main outcomes of interest is the demand for treated water. To measure this, we will collect administrative water sales records from Spring Health. For every household that purchases water from Spring Health, including both households who received a treatment offer and those who did not, we will collect information on the daily number of bottles of water purchased from Spring Health, and at what price. For exchangeable quota households, we will use these data to calculate how much money they are owed at the end of each reporting period.

Because Spring Health's administrative data are complete, and our treatment offer dataset enumerates every household in each village, we can also infer that households who do not appear in the Spring Health administrative dataset must consume 0 litres of their water (they may consume untreated water or treat water in other ways which we will investigate in our surveys). This yields a household-by-date panel of Spring Health water purchases.

3.3 Survey

In addition to collecting administrative data from Spring Health, we plan to conduct a series of household surveys. We will not conduct surveys in pure control villages. In each treatment village, we will sample 15 randomly selected households, stratified by treatment arm including the control group.[3]

We plan that each household surveyed five times: The first survey will be held one month after scratch card distribution. Three shorter 'high-frequency' check-ins will be held approximately two,three, and four months after scratch card distribution. An endline survey will be conducted five months after scratch card distribution. We will also conduct household water testing at this time. These test results will form part of an 'information treatment' experiment following the main design (see Section 3.3.4). Figure 3 presents a rough experimental timeline. We describe each survey in more detail below. Through each round of the survey, households who dropped out of the survey were replaced by selecting a similar backup at random from the corresponding treatment arm.

3.3.1 First Survey

We will conduct a brief survey, consisting of basic household demographic information, including total number of household members, number of children, and number of infants; main source of water for cooking and handwashing, laundry and cleaning, gardening and farming, and drinking; drinking water treatment behavior; water purchase behavior; information on earnings in the past month; information on missed school in the past week; and data on household illness, including symptoms, clinic visits, and expenses. Several variables in this survey are time-invariant and may therefore be useful baseline controls in some empirical specifications.

3.3.2 High-frequency survey

In the high-frequency surveys, we will ask a subset of the questions in the baseline questionnaire: we ask for information on earnings in the past month; information on missed school in the past week; and data on household illness, including symptoms, clinic visits, and expenses. This allows us to measure any treatment effects on work and school attendance and health, which may be relatively short-lived and therefore difficult to capture in the endline survey.

3.3.3 Endline and Water Quality Measurement

This portion of the pre-analysis plan was completed as an addendum prior to completing the endline survey but after implementation of the main experiment had begun. At the end of the 5-month period of water distribution, an endline survey is being conducted for

3. In price discount villages, we will survey 4 control households and 3 households in each discount level; in pure quota villages we will survey 6 control households and 7 quota households; and in exchangeable quota villages, we will survey 3 control households, 2 10% exchange households, 3 50% exchange households, 2 90% exchange households, and 3 100% exchange households.

all households who were part of prior survey rounds. During the Endline, drinking water samples from the containers in the households were collected (if the household had more than 4 containers, the samples were restricted using the following logic: 1 sample per container type defined by material of the container i.e. 4 plastic bottles become 1 sample and so on). In addition to questions asked during the high frequency survey the endline contains additional questions about household use of alternative treatment methods such as chlorine and boiling.

The water quality tests for individual households will be conducted using an H2S test vial manufactured by Transchem Agritech- the test kit consists of a collection vial and an H2S strip to determine the presence or absence of any form of coliform bacteria. This test provides a binary result with a change in colour in the presence of coliform that exceeds safe limits.

We will also conduct tests of community sources in our experiment villages (that is, wells, municipal pipes etc). These will use a different technology, namely a continuous test using a bactaslyde. This test provides a visual measure of the extent of contamination.

Figure 4 shows a filled collection vial from each of these kits.

Figure 4: Test vials corresponding to household drinking water samples (left) and community drinking water sources (right). Images show contaminated water in both instances.

3.3.4 Information Treatment

The household water quality samples collected during the end of the study will be used as part of a follow-up intervention intended to identify the impact of information on household water quality, as well as information on methods to treat and store water. Our interest is in examining whether such information changes (i) the demand for clean water, (ii) household behaviours associated with storing and treating water, and (iii) household water quality.

To accomplish this, 1368 households who had been part of the survey will form the full experiment sample. These households will be given a two-month extension to the original five-month treatment in order to provide time for us to carry out the information treatment and then evaluate its (short-run) impacts.

This sample is smaller than the full end-line sample of 2000 households because we will exclude the first two implementation waves for whom the Spring Health offers expired over a month before the end-line began. This exclusion is due to logistical difficulties in resuming their treatments after a long hiatus. Therefore, we focus only on households that continued to be eligible for Spring Health water at the time of the end-line (or for whom the offer expired very recently).

Of these 1368 households, 456 households will receive personalised information in the form of water quality test results corresponding to samples from their homes. In addition, they will be provided with a pamphlet with information on proper storage and/or purification of drinking water at home. A second group of 456 will only receive this general information but not their test results (we intend to provide these households with their test results after the experiment is complete). Finally, 456 households form the control and will be given no information at all, although they will still be part of the Spring Health extension offers.

4 Analysis

4.1 Main effect 1: Household-level treatment effects on water consumption

4.1.1 Intent to treat effects

We begin by estimating the intent-to-treat effects of each of our treatments on clean water consumption.[4]

Spring Health water: Extensive margin We first estimate a cross-sectional extensive-margin regression to understand whether our treatments changed households' use of Spring

4. Note that when referring to the Spring Health sales data in this section, when we write "consumption" we mean "purchases of water", as we cannot directly observe household usage in Spring Health's sales data.

Health water:

$$
\begin{aligned}
\mathbf{1}[Q_i > 0]_i = {} & \eta_1 \cdot 10\% \text{ discount}_i + \eta_2 \cdot 50\% \text{ discount}_i + \eta_3 \cdot 90\% \text{ discount}_i \\
& + \eta_4 \cdot \text{Pure quota}_i + \eta_5 \cdot 10\ \% \text{ exchange}_i + \eta_6 \cdot 50\ \% \text{ exchange}_i + \eta_7 \cdot 90\ \% \text{ exchange}_i \\
& + \eta_8 \cdot 100\ \% \text{ exchange}_i + \gamma_v + \varepsilon_i \qquad (1)
\end{aligned}
$$

where $\mathbf{1}[Q_i > 0]_i$ is an indicator equal to 1 if and only if household i purchased any Spring Health water during the course of the experiment, the treatment indicators are each listed, γ_v is a village fixed effect, and ε_i is an error term. As a baseline specification, we compare treated consumers to only untreated customers within treated villages. As robustness, we also run a version of this specification which includes pure control customers. For this, as well as all other household-level regressions in the Analysis section, we will estimate a version that includes a vector of pre-determined controls, $\mathbf{X}_i$, selected using LASSO.

We also estimate a pooled version, where we combine treatments of different levels:

$$
\mathbf{1}[Q_i > 0]_i = \eta_1 \cdot \text{Any discount}_i + \eta_2 \cdot \text{Pure quota}_i + \eta_3 \cdot \text{Any exchange}_i + \gamma_v + \varepsilon_i \qquad (2)
$$

Spring Health water: Quantity We next estimate the effect of our treatments on the quantity of Spring Health water purchased by households. We first use a cross-sectional specification:

$$
\begin{aligned}
Q_i = {} & \eta_1 \cdot 10\% \text{ discount}_i + \eta_2 \cdot 50\% \text{ discount}_i + \eta_3 \cdot 90\% \text{ discount}_i \\
& + \eta_4 \cdot \text{Pure quota}_i + \eta_5 \cdot 10\ \% \text{ exchange}_i + \eta_6 \cdot 50\ \% \text{ exchange}_i + \eta_7 \cdot 90\ \% \text{ exchange}_i \\
& + \eta_8 \cdot 100\ \% \text{ exchange}_i + \gamma_v + \varepsilon_i \qquad (3)
\end{aligned}
$$

where Q_i is the total amount of Spring Health water purchased by household i during the experimental time period, and all other variables are the same as in Equation (1) above. We also estimate a pooled version identical to Equation (2) with Q_i on the left-hand side.

In addition, we estimate a panel model that includes time fixed effects to soak up additional identifying variation and serve as a basis for our event study specification.

$$
\begin{aligned}
Q_{it} = {} & \beta_1 \cdot 10\% \text{ discount}_i + \beta_2 \cdot 50\% \text{ discount}_i + \beta_3 \cdot 90\% \text{ discount}_i \\
& + \beta_4 \cdot \text{Pure quota}_i + \beta_5 \cdot 10\ \% \text{ exchange}_i + \beta_6 \cdot 50\ \% \text{ exchange}_i + \beta_7 \cdot 90\ \% \text{ exchange}_i \\
& + \beta_8 \cdot 100\ \% \text{ exchange}_i + \gamma_v + \delta_t + \varepsilon_{it} \qquad (4)
\end{aligned}
$$

where Q_{it} is the quantity of Spring Health water purchased by household i in village v on date t, we include an indicator variable for households who were *offered* each treatment, γ_v is a village fixed effect, δ_t is a month fixed effect, and ε_{it} is an error term, clustered at the individual level. We will estimate Equation (4) both with and without the pure control villages. We will also estimate a version where we pool all treatments of the same type.

Dynamic effects Next, we estimate the extent to which our treatment effects change over time. To do this, we estimate an event study version of our panel model. In partcular, we regress:

$$\begin{aligned} Q_{it} = \sum_{m=0}^{6} \beta_1^m \cdot [\text{months from 10\% discount start} = m]_i + \beta_2^m \cdot [\text{months from 50\% discount start} = m]_i \\ + \beta_3^m \cdot [\text{months from 90\% discount start} = m]_i + \beta_4^m \cdot [\text{months from pure quota start} = m]_i \\ + \beta_5^m \cdot [\text{months from 10\% exchange start} = m]_i + \beta_6^m \cdot [\text{months from 50\% exchange start} = m]_i \\ + \beta_7^m \cdot [\text{months from 90\% exchange start} = m]_i + \beta_8^m \cdot [\text{months from 100\% exchange start} = m]_i \\ + \gamma_v + \delta_t + \varepsilon_{it} \qquad (5) \end{aligned}$$

where we now estimate m treatment effects of each type, one for each month after treatment begins. We will also estimate a version where we pool all treatments of the same type.

4.2 Main effect 2: Price elasticity of demand

Price discount villages A key parameter is the price elasticity of demand for clean water. To estimate this, we first run the following specification, using price discount villages only:

$$\log(Q)_i = \beta \cdot \log(P)_i + \gamma_v + \varepsilon_i \qquad (6)$$

where P_i is the price of Spring Health water, and identifying variation comes from random assignment of households to discount levels. In addition, we will construct a figure that plots average consumption at each price level.

Exchangeable quota villages Households in our exchangeable quota villages face the same marginal incentive as households in the price discount villages: the opportunity cost of water for a household with a 10% exchangeable quota is 10% of the market price. As a result, we can use these villages to generate a second estimate of the price elasticity of demand. To estimate this, we estimate Equation (6) instead using exchangeable quota villages only, where we define P_i to be the opportunity cost of using a quota bottle. In this case, the identifying variation comes from random assignment of households to exchangeable quota levels.

4.2.1 Testing for liquidity constraints

Next, we compare households facing the same marginal incentive in the price discount villages and exchangeable quota villages, to test for the presence of liquidity constraints. While both sets of households face the same marginal incentive, exchangeable quota households begin with a water allocation, so any differences between these two demand curves should reflect liquidity constraints. To test for this, we estimate, using (i) the price discount villages, the exchangeable quota villages, and the control villages; and (ii) the price discount villages and exchangeable quota villages only:

$$\log(Q)_i = \beta_1 \cdot \log(P)_i + \beta_2 \cdot \log(P)_i \times \mathbf{1}[\text{Exchangeable quota}]_v + \gamma_v + \varepsilon_i \tag{7}$$

where the main coefficient of interest is β_2, which estimates the differential effect of being in an exchangeable quota village (denoted by the indicator function) on the price elasticity of demand.

We also estimate the difference in quantities at each price point for discount vs. exchangeable quota households with the following specification:

$$\begin{aligned} Q_i = {} & \beta_1 \cdot [10\% \text{ discount price}]_i + \beta_2 \cdot [10\% \text{ discount price}]_i \times \mathbf{1}[\text{Exchangeable quota}]_v \\ & + \beta_3 \cdot [50\% \text{ discount price}]_i + \beta_4 \cdot [50\% \text{ discount price}]_i \times \mathbf{1}[\text{Exchangeable quota}]_v \\ & + \beta_5 \cdot [90\% \text{ discount price}]_i + \beta_6 \cdot [90\% \text{ discount price}]_i \times \mathbf{1}[\text{Exchangeable quota}]_v \\ & + \gamma_v + \varepsilon_{it} \end{aligned} \tag{8}$$

where the difference between β_1 and β_2, β_3 and β_4, and β_5 and β_6 is the difference in demand at each point between the discount arm and the exchangeable quota arm.

4.3 Main effect 3: Household-level treatment effects on health

4.3.1 Short-run ITT effects

Next, we measure intent-to-treat effects of each of our treatments on health. To do this, we estimate the following specification:

$$\begin{aligned} \bar{H}_i = {} & \beta_1 \cdot 10\% \text{ discount}_i + \beta_2 \cdot 50\% \text{ discount}_i + \beta_3 \cdot 90\% \text{ discount}_i \\ & + \beta_4 \cdot \text{Pure quota}_i + \beta_5 \cdot 10\ \% \text{ exchange}_i + \beta_6 \cdot 50\ \% \text{ exchange}_i + \beta_7 \cdot 90\ \% \text{ exchange}_i \\ & + \beta_8 \cdot 100\ \% \text{ exchange}_i + \gamma_v + \varepsilon_i \end{aligned} \tag{9}$$

where $\bar{H}_i$ is the average answer to whether any household member was sick in the past week by household i across all high-frequency survey waves. We will additionally run this specification using the other health questions in our survey, including whether any household members had diarrhea, vomiting, or abdominal pain. We will also estimate this using childrens' outcomes on the left-hand side. Because we do not conduct high-frequency surveys in pure control villages, for these treatment effects, we exclude the pure control villages from the sample. We will also estimate a version where we pool all treatments of the same type.

4.3.2 Short-run LATEs

Our key objective with our health measures is to quantify the impacts of clean water consumption on health. In order to do this, we move from reduced-form estimates of intent-to-treat effects to 2SLS estimates of LATEs, instrumenting for water consumption with indicators of treatment assignment. We use the following system of equations:

First stage:

$$\begin{aligned} Q_i = {} & \beta_1 \cdot 10\% \text{ discount}_i + \beta_2 \cdot 50\% \text{ discount}_i + \beta_3 \cdot 90\% \text{ discount}_i \\ & + \beta_4 \cdot \text{Pure quota}_i + \beta_5 \cdot 10 \text{ \% exchange}_i + \beta_6 \cdot 50 \text{ \% exchange}_i + \beta_7 \cdot 90 \text{ \% exchange}_i \\ & + \beta_8 \cdot 100 \text{ \% exchange}_i + \gamma_v + \varepsilon_i \end{aligned} \tag{10}$$

Second stage:

$$\bar{H}_i = \pi \cdot \widehat{Q_i} + \gamma_v + \varepsilon_i \tag{11}$$

where π is the coefficient of interest, mapping out the relationship between consumption of clean water and health. We will also estimate a version where we pool all treatments of the same type.

4.3.3 Medium-run effects

We also measure medium-run ITT and LATE effects of our treatments on health using data from the endline survey. Our endline survey includes an identical set of health questions as the high-frequency surveys. We will therefore estimate the effect of Spring Health water on health using Equations 9, 10, and 11, replacing the outcome variables with their endline-survey equivalents. We will also include whether anyone in the household missed work and total health expenditures as outcome variables.

4.4 Secondary effects: Non-health benefits

In addition to measuring treatment effects on health, we expect that providing clean water at lower cost may lead households to (A) spend less time collecting water, (B) spend less money treating water, (C) running out of water less often, and (D) switching primary water sources. We test for these by estimating the following specification:

$$\begin{aligned}\bar{NH}_i = {} & \beta_1 \cdot 10\% \text{ discount}_i + \beta_2 \cdot 50\% \text{ discount}_i + \beta_3 \cdot 90\% \text{ discount}_i \\ & + \beta_4 \cdot \text{Pure quota}_i + \beta_5 \cdot 10 \text{ \% exchange}_i + \beta_6 \cdot 50 \text{ \% exchange}_i + \beta_7 \cdot 90 \text{ \% exchange}_i \\ & + \beta_8 \cdot 100 \text{ \% exchange}_i + \gamma_v + \varepsilon_i \end{aligned} \tag{12}$$

where $\bar{NH}_i$ is the average answer by household i to our high-frequency survey questions on the amount of time spent collecting water, amount of money spent purchasing water, amount of money spent treating water, chance of running out of water, and the chance that the household listed a particular source as its main source of water. Because we do not conduct high-frequency surveys in pure control villages, for these treatment effects, we exclude the pure control villages from the sample. We will also estimate a version where we pool all treatments of the same type.

4.5 Secondary effect: Spillovers

We test for the possibility that our treatments impact the demand for Spring Health water among control households by comparing control households in treatment villages against those in pure control villages.

Spillovers on Spring Health customer status First, we test whether treatments induced control households to sign up for Spring Health deliveries. Using a sample of only control and pure control households, we estimate:

$$\mathbf{1}[Q > 0]_v = \beta_1 \cdot \text{price discount}_v + \beta_2 \cdot \text{pure quota}_v + \beta_3 \cdot \text{exchangeable quota}_v + \varepsilon_v \tag{13}$$

where $\mathbf{1}[Q > 0]_v$ is the count of households purchasing at least 1 bottle of Spring Health water in each village, v. We include the entire sample of villages in this regression. For pure control villages, we include all households in the count; for treated villages, we include only control households in the count. In treated villages, to account for the fact that some households are treated, we scale the count variable by $\frac{1}{1-\text{treated household share}_v}$.

Spillovers on water purchases To test for spillover effects of treatment on water purchases, we estimate the following equation, on a sample that excludes all households who received a treatment offer:

$$Q_v = \beta_1 \cdot \text{price discount}_v + \beta_2 \cdot \text{pure quota}_v + \beta_3 \cdot \text{exchangeable quota}_v + \varepsilon_v \tag{14}$$

where Q_v is the sum of Spring Health water purchased over the course of the experiment in each village v. As above, we include the entire sample of villages in this regression, and scale the quantity for treated villages to account for the existence of treated households.

We do not estimate spillover effects on health, because we restricted our household surveys to treatment villages for budgetary and power reasons. However, if there are any spillover effects between treatment and control households, and if there is a positive effect of clean water on health, any health effects we estimate in Equations (9) and (11) can be viewed as lower bounds.

4.6 Heterogeneity

In addition to our average treatment effects (ITTs and LATEs), we are interested in heterogeneity along several dimensions. Here, we highlight specific dimensions of heterogeneity that we can explore in the absence of the endline survey, and describe our approach for estimating these effects.[5] We expect different dimensions of heterogeneity to matter for different outcome variables, so we segment our discussion by outcome. In this section, we discuss our main heterogeneous effects specifications, though we may estimate heterogeneity along additional dimensions in the final paper.

4.6.1 Water quantity effects

We will estimate heterogeneous treatment effects on clean water quantities using specifications of the following form:

$$Q_i = \sum_{k \in \text{treatments}} \pi_k \cdot \text{Treatment}_{ik} \cdot \text{Covariate}_i + \eta_k \cdot \text{Treatment}_{ik} + \beta \cdot \text{Covariate}_i + \gamma_v + \varepsilon_i \tag{15}$$

5. For implementation timing reasons, we plan to collect a variety of characteristics that are unlikely to change in the endline survey rather than in the baseline survey. We plan to test for heterogeneity using these exogenous characteristics. We will not test for heterogeneity along dimensions we expect to vary with treatment.

where the main covariates of interest are: household size (and separately number of children); household income; and the quality of existing drinking water options, all measured at baseline; above vs. below median liquidity constraints; an indicator for whether agriculture is the main source of income; above vs. below median savings; whether a household ever wanted to take a loan but was unable to; above vs. below median consumption; and above vs. below median assets.[6] We will estimate versions of this specification where we pool and separate individual sub-treatments.

4.6.2 Demand elasticity

We will estimate heterogeneous elasticities of demand using the following specification:

$$\log(Q)_i = \pi \cdot \log(P)_i \cdot \text{Covariate}_i + \beta \cdot \log(P)_i + \eta \cdot \text{Covariate}_i + \gamma_v + \varepsilon_i \tag{16}$$

where the covariates of interest are: household size (and separately number of children); household income; and the quality of existing drinking water options. We will estimate this separately for price discount villages only and for exchangeable quota villages only.

4.6.3 Health effects

We also plan to estimate heterogeneity in health effects, using the following ITT specification:

$$\bar{H}_i = \sum_{k \in \text{treatments}} \pi_k \cdot \text{Treatment}_{ik} \cdot \text{Covariate}_i + \eta_k \cdot \text{Treatment}_{ik} + \beta \cdot \text{Covariate}_i + \gamma_v + \varepsilon_i \tag{17}$$

For health, we plan to examine heterogeneity by the number of children in the household, household income, quality of household drinking water measured at baseline, and whether the household treats their drinking water at baseline.

We will also test for heterogeneous health LATEs, where the first stage is Equation (15), and the second stage is:

$$\bar{H}_i = \pi \cdot \widehat{Q_i} \cdot \text{Covariate}_i + \eta \cdot \widehat{Q_i} + \beta \cdot \text{Covariate}_i + \gamma_v + \varepsilon_i$$

4.7 Multiple hypothesis testing

This study has a limited number of key outcome variables. We are primarily concerned with (i) consumption of clean water; and (ii) health. (i) is measured with one outcome variable.

6. We may conduct additional heterogeneity analysis as well.

For (ii), we will (a) focus on a primary outcome, whether any household member was sick in the past week, and (b) generate an index over the following health variables from our high-frequency survey: whether any household member was sick in the past week, whether any household members had diarrhoea, vomiting, or abdominal pain, and outcomes for children. We will also report the False Discovery Rate $q-$values following Anderson (2008) for health outcomes.

5 Information intervention

The second phase of this information is an experimental intervention to investigate whether households change consumption and water treatment behaviours, in response to information about the quality of water in their homes. The hypothesis we are looking to test is that when households are told that their drinking water *stored in their homes*, is not clean, they may be more willing to purchase alternative sources of water or change how they store and treat water at home. Note that households may have contaminated water in the house even if their original source is clean if they store water under unhygenic conditions, mix clean and contaminated sources, etc.

Our primary test of the impact of information (as defined in Section 3.3.4 will be a set of intent-to-treat regressions:

$$\bar{Q}_i = \beta_1 \cdot \text{Test}_i + \beta_2 \cdot \text{Pamphlet}_i + \gamma_v + \varepsilon_i \tag{18}$$

where $\bar{Q}_i$ is the quantity of Spring Health water purchased by household i in village v after the endline survey; Test_i is an indicator for being in the water testing results plus information pamphlet group; Pamphlet_i is an indicator for being in the generic information pamphlet group; γ_v are village fixed effects, and ε_i is an error term. We will estimate Equation (18) without pure control villages, because we only do this intervention in treatment villages.

The equation above uses the quantity of water consumed on the left hand side. In a similar manner we will run regressions investigating whether water is purchased or not (extensive margin), and health complaints in the preceding week (similar to the endline in the first phase of the experiment, and collected in a follow-up survey which will be fielded several months after the intervention).

We are also interested in changes in health behaviours relevant to how they store and use drinking water (whether from a clean source or not). To investigate these we will look at additional outcome variables within a similar framework. These include

(i) whether household use a point-of-use treatment technology (eg. Boil, Bleach/chlorine, Strain through cloth, Use water filter, RO system, Solar disinfection) and the causal impact of different types of information on these behaviours.

(ii) questions on storage practices including whether the container is closed (by a lid or a cork)? Does it have a tap or dispenser? Are there flies or other insects around it? Is it in a place that is lit or dark? How often do you wash the container?

In addition, we are interested in how the treatment effects described above vary by households' original intervention group. We will estimate separate impacts of the information and pamphlet interventions for households in the control group, discount group, free quota group, and exchangeable quota group, and test whether these effects are significantly different from one another.

We will run several specifications and adjust for multiple hypothesis testing *within category* to guard against false positives. That is to say, we will separately inflate the p-values for questions corresponding to treatment, storage, and health. We may also combine these questions into indices using simple averaging for binary responses, but we will also report separate results.

References

Anderson, Michael. 2008. "Multiple inference and gender differences in the effects of early intervention: A reevaluation of the Abecedarian, Perry Preschool, and Early Training Projects." *Journal of the American Statistical Association* 103 (484): 1481–1495.