Eliciting and Utilizing Willingness to Pay: Evidence from Field Trials in Northern Ghana

Online Appendices

James Berry, Greg Fischer and Raymond Guiteras*

March 2019

Contents

*Berry: University of Delaware, jimberry@udel.edu; Fischer: London School of Economics, g.fischer@lse.ac.uk; Guiteras: North Carolina State University, rpguiter@ncsu.edu.

A BDM Script

Section numbers refer to survey instrument. For full text of all sales treatments, see the Supplemental Materials.

J. REGULAR_BDM

READ EXACTLY FROM SCRIPT. DO NOT SAY <u>ANYTHING</u> THAT IS NOT IN SCRIPT.

READ:

- We would like to sell you a filter but the price is not yet fixed. It will be determined by chance in a game we are about to play.
- You will not have to spend any more for the filter than you really want to.
- You may even be able to buy it for less.

Here is how the promotion works:

- I will ask you to tell me the <u>maximum price (***dan kuli***) you are willing to pay (***ka a ni sagi dali***)</u> for the *Kosim* filter (koterigu di mali lokorigu).
- In this cup, I have many different balls with different numbers on them.
- The numbers represent prices for the filter.
- Then I will ask you to pick a ball from the cup, and we will look at the price together.
- If the number you pick is less than or equal to your bid, you will buy (***ani too dali***) the filter and you will pay the price you pick from the cup.
- If the number you pick is greater than your bid, then you cannot buy the filter.
- You will only have one chance to play for the filter.
- You cannot change your bid after you draw from the cup.
- You must state a price that you are actually able to pay now.
- We will practice in one moment, but for now, do you have any questions?

Answer any questions respondent has.

J.1 REGULAR_BDM PRACTICE

REMEMBER: Get respondent to state HIGHEST price they are WILLING AND ABLE to pay right now.

NOTE: Refer to p.23 for correct Dagbani translation of Cedi amounts.

- Before we play for the filter, let's practice the game. We'll play the same game, but instead of playing for the filter, we will play for this bar of soap. ***Show respondent soap.***

1) What is the maximum amount (dan kuli) that you are willing to pay for this soap? **[*Respondent states price X*]**

2) Now, if you pick a number that is less than or equal to X, you will buy the soap at the price you pick. If you pick a number greater than X, you will not be able to purchase the soap, even if you are willing to pay the greater number. You cannot change your bid after you pick a price. Do you understand?

3) Please, tell me - if you pick [X+5 peswas] now, what happens? ***If respondent does not give correct answer, explain the rules again and then ask question again.***

4) And if you pick [X-5 peswas] now, what happens? ***If respondent does not give correct answer, explain the rules again and then ask question again.***

5) If you draw [X+5], will you want to purchase the soap for [X+5]?
IF YES: **→ 5)**
IF NO: **→ 6)**

6) Do you want to change your bid to [X+5]?
IF YES: OK, your new bid is [X+5]. **→ 2) [*use* X+5 *as new* X]**
IF NO: **→ 6)**

7) So, is X truly the most you would want to pay?
IF YES: **→ 7)**
IF NO: **→ 1)**

8) If you pick X, you must be able to pay X. Are you able to pay X now?
IF YES: **→ J.1.1**
IF NO: What is the maximum price you are willing **and able** to pay now? → **2) [*use new X*]**

→ ***Record respondent's Final Bid*** (J.1.1, page 29)

9) Could you please fetch the amount you have stated you are willing to pay and show it to me?
Wait for respondent to fetch money and check to see she has enough funds for Final Bid.

10) Now you will pick a price from the cup. If you pick X or less, you will buy the soap at the price you pick. If you pick more than X, you will not be able to buy the soap. Are you ready to pick a ball?

Mix balls in cup, hold cup above eye level of respondent and have her pick a ball without looking.

11) Now you can draw a ball from the cup. ***Let respondent draw ball. Together, look at the ball and read the price picked.*** **[*Drawn price is Y*]**

→ ***Record Drawn Price*** (J.1.2, page 29)

12) Let us look at the ball together.

→ ***Record if Drawn Price is lower/equal to or higher than Final Bid Survey*** (J.1.3, page 29)

a. **[*If Y <= X*]:** The price is Y which is [less than/equal to] the amount you said you would be willing and able to pay for this soap. You can now buy the item at this price.

→ ***Exchange payment for soap*.**

b. **[*If Y > X*]:** The price is Y, which is greater than the amount you said you would be willing to spend. You cannot purchase the soap.

13) Do you have any questions about the game?

Address any questions or concerns respondent has. Make sure she understands rules of game.

J.2 REGULAR_BDM FILTER SALE

REMEMBER: Get respondent to state HIGHEST price they are WILLING AND ABLE to pay right now.

NOTE: Refer to p.23 for correct Dagbani translation of Cedi amounts.

Read:

- Now you will play to buy the filter
- Recall the community meeting on [day of community meeting]
- Have you thought about how much you are willing to pay for the filter?
- Do you have the funds available now?

Let's begin:

1) What is the maximum amount (dan kuli) that you are willing to pay for this filter?
[*Respondent states price X*]

2) Now, if you pick a number that is less than or equal to X, you will buy the soap at the price you pick. If you pick a number greater than X, you will not be able to purchase the soap, even if you are willing to pay the greater number. You cannot change your bid after you pick a price. Do you understand?

3) Please, tell me - if you pick [X+1 cedis] now, what happens? ***If respondent does not give correct answer, explain the rules again and then ask question again.***

4) And if you pick [X-1 cedis] now, what happens? ***If respondent does not give correct answer, explain the rules again and then ask question again.***

5) If you draw [X+1], will you want to purchase the filter for [X+1]?
IF YES: → **5)**
IF NO: → **6)**

6) Do you want to change your bid to [X+1]?
IF YES: OK, your new bid is [X+1]. → **2) [*use* X+1 *as new* X]**
IF NO: → **6)**

7) So, is X truly the most you would want to pay?
IF YES: → **7)**
IF NO: → **1)**

8) If you pick X, you must be able to pay X. Are you able to pay X now?
IF YES: → **J.2.1**
IF NO: What is the maximum price you are willing **and able** to pay now?
→ **2) [*use new X*]**

→ ***Record respondent's Final Bid*** (J.2.1, page 29)

9) Could you please fetch the amount you have stated you are willing to pay and show it to me?
Wait for respondent to fetch money and check to see she has enough funds for Final Bid.

10) Now you will pick a price from the cup. If you pick X or less, you will buy the filter at the price you pick. If you pick more than X, you will not be able to buy the filter. Are you ready to pick a ball?

Mix balls in cup, hold cup above eye level of respondent and have her pick a ball without looking.

11) Now you can draw a ball from the cup. ***Let respondent draw ball. Together, look at the ball and read the price picked.* [*Drawn price is Y*]**

→ ***Record Drawn Price*** (J.2.2, page 29)

12) Let us look at the ball together.

→ ***Record if Drawn Price is lower/equal to or higher than Final Bid*** (J.2.3, page 29)

a. **[*If Y* <= *X*]:** The price is Y which is [less than/equal to] the amount you said you would be willing and able to pay for this filter. You can now buy the filter at this price.

→***Receive payment for filter. Record filter tracking code on survey*** (I.2.5, page 29). ***Record filter tracking code on receipt and give it to respondent. Inform her of where and when she can pick up the filter.***

b. **[*If Y* > *X*]:** The price is Y, which is greater than the amount you said you would be willing to spend. You cannot purchase the filter.

→***Go to Household Survey question*** J.24, page 29

B Measurement of Risk and Ambiguity Aversion

This section provides additional detail on the hypothetical gambles used to measure risk and ambiguity aversion in the one-year follow-up surveys.

To measure of risk aversion, we presented subjects with a series of choices between (a) a 50-50 gamble for a gain of 8 GHS and (b) a certain gain of X. The certain gain X began at 0.5 GHS and increased by 0.5 GHS until the subject chose the certain sum over the risky gamble. We create an integer variable to indicate the switching point and reverse the scale to yield a measure increasing in risk aversion. For example, for a subject who chose the certain 0.5 GHS over the risky gamble–the most risk-averse choice–the variable takes on a value 11, while a switching point of GHS 1 corresponds to a value of 10. The median switching point was GHS 2, corresponding to an integer value of 8. We then repeated this exercise in the loss domain, in which we measured the minimum payment at which the subject would choose a 50-50 gamble for a loss of 8 GHS over a certain payment to the experimenter. Finally, we conducted the exercise in the gain-loss domain, in which we measured the minimum sum the subject would be willing to pay to avoid a 50-50 gamble for winning 4 GHS vs. losing 4 GHS, or, if the subject were risk-loving, how much the subject would need to be compensated to forgo such a gamble. In our analysis, we use the first principal component of these three measures, but the results in Section 6.2 are robust to other methods of combining them.

To measure ambiguity aversion, we presented subjects with a version of the game posed by Ellsberg (1961). Subjects were presented with one bag that contained 5 black balls and 5 white balls, and another bag that contained 10 black and white balls in unknown proportions. The subject would choose the winning color and draw from a bag. Subjects were asked to choose between the first bag with a payout of 4 GHS and a second bag with varying payouts. The payout of the second bag started at GHS 0.5 and increased by 0.5 GHS until the subject chose the second bag. We identify subjects as ambiguity averse if they required at least 4.5 GHS to choose the second bag. By this measure, 41.6 percent percent of subjects are classified as ambiguity averse. We also create an integer measure of ambiguity aversion that corresponds to point at which the subject chose the second bag.

C Attrition

In this section, we discuss attrition from the follow-up surveys. The overall attrition rate was 12.9 percent in the one-month survey and 9.5 percent in the one-year survey. Table A3 shows that attrition from the one-month survey was fairly well-balanced on assignment to BDM vs. TIOLI, the BDM bid, the BDM draw, the TIOLI price, and most observable characteristics. Households that attritted were somewhat more likely to have a young child than households that were captured (7.9 pp, $p < 0.05$). In the one-year follow-up, attrition was again largely balanced on observable variables. Attritted households had significantly more young children ($p < 0.05$) and reported more young children having diarrhea in the two weeks ($p < 0.1$) prior to the baseline survey. We also find that attritters in the BDM treatment had lower WTP for the filter than non-attritters (GHS 1.0, $p < 0.01$).

While attritters in the one-year survey had lower WTP, on average, than non-attritters, our heterogeneous treatment effects are estimated *across* the distribution of WTP. The most relevant test in this case thus whether treatment is correlated with attrition at different levels of WTP. To implement this test, we estimate the following equation at different levels of WTP:

$$y_{ic} = \beta_0 + \beta_1 T_{ic} + \varepsilon_{ic}, \tag{1}$$

In this equation, y_{ic} is an indicator for whether subject i in compound c attritted from the follow-up survey, T_{ic} is an indicator for treatment (subject i's BDM bid was greater than her draw). To condition on WTP, we estimate equation (1) using a kernel (local linear) regression. As in Section 5.2, we estimate at each GHS 0.1 step from GHS 1 to GHS 6, which correspond approximately to the 0.1 and 0.9 quantiles of WTP in the BDM sample. We use an Epanechnikov kernel and Silverman's rule of thumb to choose the bandwidth. Following our analysis of heterogeneous treatment effects, we restrict the sample to BDM subjects with one or more children age 0 to 5 in one-year follow-up villages.

The results are plotted in Figure A5. As shown in the figure, there is no significant difference in attrition between treated (BDM winners) and untreated (BDM losers) once we condition on WTP. While we cannot test whether attrition is balanced on unobservables, this null result may mitigate the potential concern regarding the correlation between WTP and attrition shown in Table A4.

D Heterogeneous Treatment Effects, Detail

D.1 Heterogeneous Treatment Effects: Theory, Detail

This section provides a more detailed treatment of the theory introduced in Section 4.2 and provides greater detail on the LIV estimator of Heckman et al. (2006).

We begin with the generalized treatment effects model of Equation (3) in the main text:

$$y = \beta_0 + \beta_1(w)\,T + \varepsilon. \tag{2}$$

As in Section 4.2, suppose the product is offered at two random TIOLI prices, $Z \in \{P_L, P_H\}$. If there is differential take-up at the two prices, $\Pr(T|P_L) > \Pr(T|P_H)$, then Z is correlated with T, so the instrument is relevant. For the instrument to be valid, it is necessary that $E[Zu] = 0$. Expanding u as in Equation (5) in the main text, we require

$$E\left[Z\left(\tilde{\beta}_1(w)\,T + \varepsilon\right)\right] = 0. \tag{3}$$

As in Section 4.2, we consider levels and gains separately. By randomization,

$$E[Z(\varepsilon)] = 0, \tag{4}$$

so the instrument solves the problem of selection on levels. However, we must also consider the selection-on-gains term

$$E[Zu] = E\left[Z\tilde{\beta}_1(w)\,T\right], \tag{5}$$

which need not be zero. Even though Z is *unconditionally* random, it may not be independent of $\tilde{\beta}_1(w)\,T$: since $T = 1\{\text{WTP} > Z\}$, if there is a relationship between WTP and gains then (5) will be nonzero. As a simple example, suppose $\tilde{\beta}_1(w)$ is positively related to w. Then when $Z = p_H$, the population selecting into treatment will have, on average, high values of $\tilde{\beta}_1(w)$ relative to the population treated when $Z = p_L$. As discussed by Heckman et al. (2006), (5) is only zero if (a) there is no heterogeneity in gains ($\beta_1(w) = \overline{\beta}_1$ for all w, or, equivalently,$\tilde{\beta}_1(w) = 0$ for all w,) or (b) agents either have no information on $\tilde{\beta}_1(w)$ or, if they do have such information, they cannot or do not act on it.

As described in Section 4.2, rather than estimating either $\bar{\beta}_1$ or $\beta_1(w)$, IV estimation using TIOLI estimates:

$$\beta_1^{IV}(P_L \leq \text{WTP} \leq P_H) = \int_{P_L}^{P_H} \beta_1(w)\,dF_{\text{WTP}}(w).$$

In order to provide estimates of $\beta_1(w)$, one could add more randomized prices $P_1, \ldots, P_M$, and, using instrumental variables as above, estimate treatment effects piecewise:

$$\beta_1^{IV}\left(P_1 \leq \text{WTP} \leq P_2\right), \ldots, \beta_1^{IV}\left(P_{M-1} \leq \text{WTP} \leq P_M\right).$$

As in the case of estimating a full demand curve using randomized TIOLI prices, this will require a relatively large sample.

A second strategy to estimate $\beta_1(w)$ is provided in the marginal treatment effects literature (Heckman and Vytlacil 2007). Given an instrument Z, the marginal treatment effect, $\Delta^{\text{MTE}}(z)$, is defined as the treatment effect on those just on the margin of indifference between being treated or not when the instrument has value z. When Z is a randomized price, $\Delta^{\text{MTE}}(w)$ is equivalent to $\beta_1(z)$, since by definition someone with WTP $= z$ is indifferent between purchasing and not purchasing at a price of w. Heckman et al. (2006) show that $\Delta^{\text{MTE}}(w)$ can be estimated even though WTP is typically not observed. Heckman et al. (2006) show that the marginal treatment effect is equal to the *local instrumental variables* parameter

$$\Delta^{\text{MTE}}(w) = \Delta^{\text{LIV}}(w) = \left.\frac{\partial E\left[y \middle| \Pr(z) = \Pr\right]}{\partial \Pr}\right|_{w=z},$$

where $\Pr(Z)$ is the propensity score with respect to the instrument, representing the probability of treatment among those facing (random) price Z. The marginal treatment effect at z, then, is the change in the outcome of interest on those brought into treatment by small changes in Z around z, $\partial E\left[y \middle| \Pr(z)\right] / \partial \Pr(z)$.[1] Heckman et al. (2006) provide a local instrumental variables estimator, which estimates the propensity score $\Pr(z)$ in a first step and then regresses the outcome of interest on the propensity score. As with the first strategy, this will require a large sample with a broad range of prices, since the MTE is only identified on the support of $\Pr(z)$, and the precision of the estimate depends on the precision of the estimated propensity score.[2]

BDM can estimate $\beta_1(w)$ with greater precision than these two alternatives. In the case of piecewise randomized prices, the reason is straightforward – as in the case of estimating demand curves, each BDM observation provides much more information on

[1]For intuition, note that this is a differential analogue of the traditional Wald estimator $\left(E\left[Y|Z=1\right] - E\left[Y|Z=0\right]\right) / \left(\Pr\left[T|Z=1\right] - \Pr\left[T|Z=0\right]\right)$ in the case of a binary instrument.

[2]We focus on price as an instrument for comparability with our application. However, the method of Heckman et al. (2006) applies more broadly. For example, in their empirical example, they estimate the effect of high school graduation on wages using mother's graduation status and number of siblings as instruments. Note that continuous, many-valued, or multiple instruments will be required to estimate $\Pr(z)$ flexibly. Furthermore, the interpretation of the MTE is more subtle with non-price instruments: what is estimated is $\Delta^{\text{MTE}}(u_D)$, the effect on those with unobservables $u_D \in [0, 1]$ such that they are indifferent between treatment and non-treatment when the value of the instrument Z is z such that $\Pr(z) = u_D$. See Brinch et al. (2017) for progress on estimating MTEs with a discrete instrument and Kowalski (2016) for the interpretation of MTEs as a function of unobservables.

WTP than TIOLI. In the case of marginal treatment effects / local instrumental variables, BDM allows us to observe this dimension of heterogeneity directly, rather than obtaining it indirectly through the first-step propensity score estimation.

D.2 Comparison with Local Instrumental Variables

In this section, we compare estimated treatment effects using the BDM-IV method with the Local Instrumental Variables (LIV) methods of Heckman, Urzúa and Vytlacil (2006, herafter HUV). We compare estimates on the primary outcome of interest in the main text: long-term (one-year followup) cases of diarrhea among children age five and younger, based on caretaker recall over the previous two weeks.

In the first LIV step, we estimate the propensity score $\Pr(z) = \Pr[T = 1|Z = z]$, where z is the BDM draw. Following HUV, we estimate $\Pr(z)$ using locally linear regression. The estimated propensity score $\widehat{\Pr}(z)$, with a 95 percent confidence band, is plotted in Figure A6a.

In the second step, we estimate $\partial E[y \mid \Pr(z)] / \partial \Pr$ by regressing the outcome y on the estimated propensity score $\widehat{\Pr}(z)$. For comparability with our BDM-IV estimates, again we use local linear regression. The results are plotted in Figure A6b. As in the main text, we have flipped the sign of the dependent variable so benefits (reductions in diarrhea) correspond to positive point estimates.

Third, by Equation (16) in HUV (pg. 397), this derivative is the treatment effect for those at the margin of indifference when $\widehat{\Pr}(z) = \text{pr}$. That is, $\partial E[y \mid \Pr(z)] / \partial \Pr = E[y_1 - y_0 \mid \Pr(z) = \text{pr}]$, which, in turn, is equal to $\Delta^{\text{MTE}}(\text{pr})$.

Fourth, in an "inversion step," we use the fact that $\Pr(z)$ is strictly monotonic (decreasing) in z to translate $\Delta^{\text{MTE}}(\text{pr})$, effects plotted as a function of pr as in Figure A6b, into $\Delta^{\text{MTE}}(z) = \Delta^{\text{MTE}}(z : \Pr(z) = \text{pr})$, effects as a function of the price draw z. This marginal treatment effect, $\Delta^{\text{MTE}}(z)$, is plotted in Figure A6c. Note that, perhaps counterintuitively, relatively *low* values of the propensity score in Figure A6b correspond to relatively *high* values of the draw in Figure A6c, since the probability of treatment is low when the draw is high.

Fifth, since z is a price, $\Delta^{\text{MTE}}(z)$ represents the effect those on the margin of indifference at a price of z, and this is exactly $\beta(w)$, the effect on those with WTP $= w$. That is, we can simply re-label the x-axis of Figure A6c as WTP rather than the price draw Z. Comparing Figure A6c with Figure 2.b, we observe that the pattern of estimated treatment effects is similar, in that they are increasing with respect to WTP.

Finally, in Figure A6d we compare the precision of the estimates by plotting the width of the 95 percent confidence intervals. In the case of LIV, since the regressors in the second step are estimates from the first step, we bootstrap the entire process, resampling by compound with replacement. The confidence intervals for BDM-IV are narrower over most of the range of WTP (GHS 1 to 5).

E Sunk Cost Effects

BDM embeds a double randomization that allows researchers to separately identify two factors that may be important for understanding the relationship between prices and use: the causal effect of price paid conditional on WTP (a *sunk-cost* effect), and the correlation between WTP and use (a *screening* effect). In Section 4.4, we analyze screening effects, showing that there is evidence for a positive association between WTP and use in the long-term follow-up survey.

Because the price draw is random, we can test for causal effects of price paid by comparing measures of use for subjects with the same WTP but who paid different prices. For example, BDM generates the following experiment: consider three subjects, each willing to pay GHS 6 for a filter; one doesn't receive the filter; another pays GHS 6; and the other pays GHS 2. Thus, at every level of WTP above the minimum price, there is variation in both allocation and the price paid conditional on allocation.

Following the analysis of WTP and use in Section 4.4, we use three indicators of use: presence of an undamaged filter, presence of water in the storage reservoir, and presence of water in the clay pot. We estimate the impact of price paid on each measure separately and on an index following Kling, Liebman and Katz (2007).

Specifically, we estimate

$$\text{use}_{ic} = \alpha_0 + \alpha_1 D_{ic} + \alpha_2 f(WTP_{ic}) + \varepsilon_{ic}, \tag{6}$$

where use_{ic} represents the use measure, D_{ic} is the respondent's draw, and $f(WTP_{ic})$ is a cubic polynomial of bid. It is important to control adequately for WTP since, although the price draw was unconditionally random, conditional on receiving the filter it is positively correlated with WTP.

Table A5 presents results from OLS estimation of Equation (6). Panel A shows that there is little evidence for an effect of the price paid on use in the one-month follow-up. Panel B shows a similar null result in the one-year follow-up data. Taken together, this suggests there are no significant sunk-cost effects.

F Framework for Compensating Behavior

This section provides additional detail for the discussion of possible mechanisms for the detrimental long-run impacts of the filter described in Section 4.5. We begin with a simple model of household health production to frame the issue. Consider a world where households maximize additively separable utility over children's health (h) and all other consumption (x) subject to a budget y. Children's health is a function of both general health behavior s, which we label "sanitation" to fix ideas,and the consumption of clean water (w): $h(s, w)$, where both inputs are continuous and non-negative with unit costs p_s and p_w. We make the usual assumptions: $h_s > 0$, $h_{ss} < 0$, $h_w > 0$, and $h_{ww} < 0$. The household's maximization problem is then to pick a vector of inputs, (s, w, x), that maximize utility, $h(s, w) + x$, subject to the budget constraint, $p_s s + p_w w + p_x x \leq y$.The filter reduces the per unit cost of clean water, p_w. If clean water and sanitation are substitutes in the health production function, $h_{sw} < 0$, a reduction in the price of clean water will reduce sanitation and other investments in children's health. The substitution between s and w alone could explain a muted or even zero impact from the filter; however, it could not generate the perverse effects that we observe.

In Section 4.5, we outline the three factors that we believe are mostly likely to have combined with compensatory behavior to generate detrimental effects in our context: sporadic reoptimization in response to gradual declines in use of the filter or the filter's effectiveness; intrahousehold allocation decisions that limited children's access to filtered water; and non-convexities in the alternative health technologies. Combined with compensatory behavior, each can produce negative treatment effects.

First, upon receipt of the filter – a large shock to their health production function – households may have reoptimized, engaging in compensatory behavior. Then, in response to a gradual decrease in use or the filter's effectiveness over time, they may have failed to reoptimize again, either due to rational inattention (Tobin 1982; Reis 2006; Da et al. 2014) or simple mistakes: households may have misperceived the benefits of maintaining or using the filter. If households that value the filter more also tend to be more attentive, we would expect more failures to reoptimize among those with low WTP.

Second, even in households that purchased the filter, some children may not have had access to the filtered water. The filter produces a limited supply of drinking water, but this water comprises multiple goods. Most importantly, we consider children's health and better tasting water for adults. Before receiving the filter, households made health investments (such as traveling to cleaner water sources or boiling their water) that jointly produced both goods. The filter can decouple this production. With the filter, adults can obtain better tasting water with less effort devoted to activities that improve water quality for the entire household. In particular, field reports indicated that some children were not allowed to drink filtered water because of concerns that they might damage the filter or that there would be insufficient "sweet tasting water" for the male head of the household.[3] The pattern of treatment effects we observe is consistent with this mechanism. Households with a low value for children's health would be less likely to provide filtered

[3] In response to the field reports, we added survey questions regarding children's access to filtered water, but subjects' answers proved unreliable.

water for their children and, all else equal, tend to have a lower WTP for the bundle of goods produced by the filter.

Finally, compensatory behavior can worsen the targeted outcome while improving utility if the alternative health production technology is non-convex. Many health behaviors have a fixed cost component. For example, suppose a household can either obtain its water at low cost from a dirty source or at ha higher cost from a cleaner source. Without the filter, the household chooses to incur the higher cost and drink relatively clean water. The filter improves the quality of the dirty water sufficiently that, if the household has the filter, it optimally chooses not to incur the cost of obtaining clean water. If filtered water with low other investment produces less health than unfiltered water with high other investment, purchasing the filter can increase utility but reduce health.

Consider a setting in which the individual can (1) either choose to obtain water from a clean source (cw) or a dirty source (dw) and (2) either use the filter or not. Note that although we describe this as clean vs. dirty water, this could be any of a set of sanitation practices. Without the filter, the individual has utility as follows:

$$U^{NF} = max\left\{h^{cw} - c^{cw},\ h^{dw} - c^{dw}\right\},$$

and with the filter, the individual has utility:

$$U^{F} = max\left\{h^{cw} + F^{cw} - c^{cw},\ h^{dw} + F^{dw} - c^{dw}\right\},$$

where h^s represents the utility of health when water is obtained from source $s \in \{cw, dw\}$, F^s is the effect of the filter on the utility of water obtained from s, and c^s is the cost of obtaining water from s.

Assumption 1. *The clean water is better for your health:* $h^{cw} > h^{dw}$,

and

Assumption 2. *The filter works:* $F^{cw} > 0$; $F^{dw} > 0$.

Consider the case where without the filter, individuals use the clean water source:

$$h^{cw} - c^{cw} > h^{dw} - c^{dw} \tag{7}$$

Based on the assumption that the filter works, receiving the filter could lead to worse health if with the filter the individual switches to the dirty water source:

$$h^{dw} - c^{dw} + F^{dw} > h^{cw} - c^{cw} + F^{cw}. \tag{8}$$

If health worsens after receiving the filter, this implies:

$$h^{cw} > h^{dw} + F^{dw}. \tag{9}$$

For notational compactness, we define $\Delta h = h^{cw} - h^{dw}$ and the analogous variables similarly. Combining the preceding equations implies:

$$\Delta c - \Delta F > \Delta h > F^{dw}$$

$$\Delta c > F^{cw}.$$

That is: (1) the health benefit of clean vs. dirty water needs to be greater than the health benefit of using the filter with dirty water (thus health goes down if households with the filter switch to dirty water) and (2) the cost of obtaining clean vs. dirty water needs to be greater than the health benefit of using the filter with clean water (thus households improve utility by switching to dirty water once they have the filter even though this reduces health). Note this also implies that $F^{cw} < F^{dw}$, that is, the filter has less of a health benefit with clean water than with dirty. Further, the willingness to pay for the filter equals $U^F - U^{NF}$, so those who have perverse effects – and hence all else equal get less net benefit from the filter – will also have relatively low willingness to pay. As in Peltzman (1973), the non-convexity of the health production technology generates the possibility of perverse effects and differentiates this setting from benchmark models of health investment such as described in Greenstone and Jack (2015).

We also considered alternative mechanisms, such as improper use of the filter or sanitation externalities (as discussed in Bennett 2012). While either could, in principle, produce detrimental impacts, both are unlikely in our setting. We find no evidence of improper use causing detrimental effects. In fact, conditional on use, nearly all filters were in good condition and well maintained. As for externalities, our heterogeneity analysis finds that those with a low WTP who receive the filter have worse one-year outcomes than households in the same village with the same WTP who were not randomly assigned – via the BDM price draw or randomized price – to receive the filter. Since the treatment status of one's neighbors, who could be generating the negative externalities, is independent of one's own treatment status, sanitation externalities are unlikely to explain the observed pattern of effects.

G Policy Counterfactuals, Detail

This section provides further detail on the policy counterfactuals described in Section 5.1. As outlined in that section, we consider a social planner who values DALYs at B. The filter costs C_F, inclusive of production, marketing and delivery. For simplicity, we treat these costs as variable, although in reality there is likely to be a substantial fixed cost at the village level. We also abstract from time costs of use. The planner chooses the sales price P. Given a price P, we find $Q^D(P)$, the share of households purchasing the filter from our analysis in Section 3. The total cost of filters is $C_F \cdot Q^D(P)$, the cost of the filter times demand.

We compute the reduction in cases of diarrhea per household, given by $\Delta_H(w)$, under two assumptions on treatment effects. In the first scenario, we use the treatment effect $\hat{\beta}_1^{1M}(w)$ from the one-month follow-up survey. Since there is little evidence of heterogeneous treatment effects in the short run, we restrict $\hat{\beta}_1^{1M}(w)$ to be constant with respect to WTP. Formally, this is given by:

$$\Delta_H^{1M}(w) = \hat{\beta}_1^{1M} \cdot 26 \cdot n_k,$$

where $\hat{\beta}_1^{1M}$ is the average reduction in children's diarrhea in each 2-week period, which we estimated in Section 5.2 to be 0.049, and n_k is the number of children in the household.

In the second scenario, we take the average of the short-term and one-year effects. That is, we compute the total effect of the filter over the first year as if the effect changed smoothly over the course of the year. We again assume the short-term effects are constant with respect to WTP, and impose a linear functional form on the one-year effects:

$$\Delta_H^{1Y}(w) = \left(\left(\hat{\beta}_1^{1M} + \hat{\beta}_1^{1Y}(w)\right)/2\right) \cdot 26 \cdot n_k.$$

Finally, let $F_{\text{WTP}}(w)$ be the CDF of WTP in the population. Since households with WTP $\geq P$ purchase when the price is P, the reduction in cases of diarrhea when the price is P is given by

$$\mathrm{H}(P) = \int_{w \geq P} \Delta_H(w)\, dF_{\text{WTP}}(w),$$

where $\Delta_H(w)$ is either $\Delta_H^{1M}(w)$ or $\Delta_H^{1Y}(w)$ depending on the scenario.

Following Kremer et al. (2011), we assume that the gain in DALYs is proportional to the reduction in cases of diarrhea:

$$\text{DALY}(P) = 0.028 \cdot \mathrm{H}(P),$$

where 0.028 is the ratio of DALYs to diarrhea incidence for Ghana in 2010 from Global Burden of Disease Collaborative Network (2017).

We then compute the average costs per DALY gained as the ratio of the total cost of filters divided by the total gain in DALYs:

$$\mathrm{AC}\,(P) = \frac{C_F \cdot Q^D\,(P)}{\mathrm{DALY}\,(P)}.$$

where C_F equals USD 15, as described in Section 2.

In order to avoid parameterization of the cost function, we compute the marginal cost per DALY in terms of a discrete price change from $P + 0.5$. to $P - 0.5$. This reflects the cost per DALY of reducing the price from $P + 0.5$ to $P - 0.5$:

$$MC(P) = \frac{C_F \cdot Q^D\,(P - 0.5) - C_F \cdot Q^D\,(P + 0.5)}{DALY(P - 0.5) - DALY(P + 0.5)}.$$

This function equals the increase in costs resulting from increased demand at a lower price, divided by the increased DALYs from including the additional households purchasing at the lower price.

Table A6 displays diarrhea cases averted, DALYs averted, and average and marginal costs per DALY averted under two different assumptions about treatment effects. In Panel A, we assume constant treatment effects using data from our one-month follow-up survey. As the price increases (across columns), coverage decreases. Since we have assumed a constant treatment effect, cases reduced conditional on purchase are constant, and total cases reduced per household in the population decrease proportionally with demand. The same holds for DALYs gained conditional on purchase and total DALYs gained per household in the population. Because the treatment effect is constant, both average and marginal costs per DALY are also constant at USD 361.

In Panel B, we assume treatment effects are an average of the effects estimated from the one-month and one-year surveys. The one-month effect is assumed to be constant, while the one-year effect is assumed to be linear in WTP. Now, as price increases, negative-gains purchasers – those with low WTP – no longer purchase the filter, and diarrhea cases reduced conditional on purchase increase. For small positive prices, total gains in the population increase as well. Above a price of GHS 4, the decrease in coverage outweighs the increasing gain per household and total gains decline. We see a similar pattern in DALYs gained, both conditional on purchase (monotonically increasing with price) and total DALYs gained in the population (increasing, then decreasing, with a maximum at GHS 4). Because the treatment effect is increasing in WTP, higher prices screen out those with lower treatment effects, and average and marginal costs per DALY decrease with price.

H Valuing Health, Detail

We calculate WTP_H, the household's observed WTP to avoid a case of diarrhea, as w, the household's WTP for the filter, divided by Δd, the number of cases avoided over the anticipated life of the filter:

$$\text{WTP}_H = \frac{w}{\Delta_H(w)}. \tag{10}$$

We obtain the numerator of Equation 10, w, directly from our WTP data. As per the discussion in Appendix F, we compute the denominator, $\Delta_H(w)$, under two scenarios about the filter's impact on child health. In the first scenario, we assume the one-month treatment effects $\hat{\beta}_1^{1M}(w)$ are constant. That is, the household correctly anticipates the average benefit, but not necessarily its own benefit.

$$\text{WTP}_H = \frac{w}{\hat{\beta}_1^{1M} \cdot 26 \cdot n_k},$$

where $\hat{\beta}_1^{1M}$ defined as in Appendix F.

In the second scenario, we assume the treatment effects evolve smoothly between the short-term and one-year effects. We again assume the short-term effects are constant and that the one-year effects are linear in WTP. A linear functional form for $\hat{\beta}_1^{1Y}(w)$ implies households' beliefs and valuations are, on average, consistent: households with WTP $=$ w believe that they will receive a health benefit given by the best linear approximation of $\beta_1^{1Y}(w)$ and, on average, households are correct in this belief. The household's WTP to avoid a case of children's diarrhea, then, is

$$\text{WTP}_H = \frac{w}{\left(\left(\hat{\beta}_1^{1M} + \hat{\beta}_1^{1Y}(w)\right)/2\right) \cdot 26 \cdot n_k}.$$

I Correlates of WTP, Detail

I.1 Feature Selection

As described in Section 6.1, we find that a model of demand using an a priori list of covariates such as wealth, education, and health status has limited predictive power for both TIOLI purchase decisions and WTP elicited directly through BDM. This reflects a common pattern for studies of health goods in low-income countries (Ashraf et al. 2010; Cohen and Dupas 2010) and the consumer behavior literature more generally (Browning and Carro 2007; Nevo 2011). In this appendix, we describe the use of LASSO for covariate selection.

The LASSO (Tibshirani 1996), common in the machine learning literature, is a penalized regression approach to variable subset (model) selection in which the data determine the set of covariates. It solves a similar minimization problem to ordinary least squares, but with a penalty for model complexity. This produces something similar to a linear regression in which only a small number of predictors have non-zero coefficients. The parameter estimates are given by

$$argmin\, L(\beta|x) + \lambda \sum_{j=1}^{p} \left|\beta_j\right| ,$$

where $L(\beta|x)$ is the loss function, usually a quantity proportional to the negative log likelihood. In our setting, we use the residual sum of squared errors for the BDM data, where WTP is directly observed, and the negative log likelihood of the logistic function for the TIOLI data, where WTP is a latent variable. The term $\lambda \sum_{j=1}^{p} \left|\beta_j\right|$ penalizes the inclusion of additional regressors, and λ is a tuning parameter that determines the extent of this penalty. For small λ, the penalty is minor and LASSO recovers the OLS regression coefficients. When λ is sufficiently large, some of the coefficients will be set to zero and the LASSO performs variable selection.

In order to perform the LASSO, we standardize all predictors to have variance one. We then randomly allocate half of the sample to a training set, which will be used to select the tuning parameter with the best out-of-sample prediction properties, and half to a hold-out set, on which we will estimate the model using the selected tuning parameter. We choose a grid of λ values for which we compute the out-of-sample prediction error using 10-fold cross validation.

To construct the cross-validation error, we divide our training sample, for example, the BDM observations allocated to the training set, into 10 groups, or folds, of approximately equal size. We reserve the observations in the first fold as a validation set and fit the model, for each value of λ in the grid, on the observations in the other nine folds. We then calculate the prediction error for the observations in the first (validation) fold. This is calculated as the mean squared error in the BDM sample and the classification error in the TIOLI sample. We repeat this procedure ten times, using each fold as a validation set. This produces ten estimates of the cross-validation error for each value of λ. We select then refit the model using all available observations and the largest value of λ (the most parsimonious model) that produced the lowest cross-validation error. This determines the set of covariates. We then estimate the model using the hold-out sample.

We include the following set of baseline features in the models for both BDM and TIOLI purchase decisions: number of adult females in the compound; number of adult males in the compound; number of children in the compound; number of children aged 5 or less; marital status; whether the respondent is the primary caregiver; indicators and counts of household assets (bicycle, bucket, chair, sewing machine, cooking pot, cutlass, lantern, light bulb, mattress, mobile phone, motorcycle, radio, refrigerator, sewing machine, television, torches, video player); the first principle component of all assets; educational attainment; acres farmed; acres owned; shared land farmed; total loans outstanding; primary occupation; pregnant; beliefs about actions that prevent diarrhea (boiling water, clean clothes, clean dishes, eating clean food, cooking food, drink clean water, filter water, good hygiene, using a latrine, medication, prayer or God, treating water, washing hands, nothing, does not know); respondent has primary responsibility for water collection in the compound; water source in the dry season (well; dugout; dam; borehole; rainwater; private standpipe; public standpipe; public dug well; river, stream or pond); water source in the rainy season (same categories as dry); status of water source in dry season; protection status of water source in rainy season; water treatment activities (boiling, ceramic filter, chemicals, cloth filter, pipe filter, letting settle); and village fixed effects.

We also consider measures of preferences and numeracy that were elicited for a subset of households at the one-year follow-up: risk aversion in the gains domain, risk aversion in the loss domain, and risk aversion in the gain/loss domain, forward digit span, backwards digit span, total digit span, ambiguity aversion categories, and whether the respondent indicated that she felt lucky in games of chance.

For the BDM subjects, we include the expressed willingness to pay for soap in the BDM practice round. For TIOLI subjects, we include each subject's purchase decision for soap in the TIOLI practice round. We do not include the TIOLI price for soap, which was randomly assigned and orthogonal to both individual characteristics and the TIOLI price for the filter.

Table A7 reports the selected features and estimated coefficients in the hold-out samples for both the BDM and TIOLI groups. In both samples, the dominant feature relates to the respondent's purchase decision for soap in the practice round. Those with a higher valuation for soap were more likely to purchase the filter or value it highly. This predictive power holds even when including all other household characteristics for which we would imagine soap purchase behavior might serve as a proxy, such as education, wealth, income, health beliefs, practices and status. Only village fixed effects, in the case of TIOLI, have comparable predictive power. In contrast, other household characteristics explain little of purchasing behavior. Characteristics related to education and asset ownership, which are often considered predictive of demand for health (Ashraf, Berry and Shapiro 2010; Cohen and Dupas 2010), appear in the regularized model for BDM demand but carry relatively limited explanatory power; they do not appear in the model for TIOLI purchase decisions.

We then expand the set of features to include preferences and numeracy and estimate the model on the smaller sample for which this data is available. Digit span (0.130), risk aversion in the gains domain (-0.074), and the indicator for whether the respondent indicated that she felt lucky in games of chance (0.046) appear in the regularized model for BDM subjects. None of these features appears in the TIOLI model.

I.2 Cross Validation

In this appendix we describe the k-fold cross validation procedure used to assess the performance of the BDM and TIOLI mechanisms in predicting TIOLI responses. We begin by randomly dividing the TIOLI sample into 10 groups of approximately equal size. The first group is treated as a validation set, and we fit the latent demand model in Equation (8) from the main text on the remaining nine groups via probit. We repeat this procedure ten times, treating each group as a validation set in turn. We denote by $\hat{r}^{(-k)}(X_{ic}; p)$ the predicted probability of purchasing at price p for an individual with characteristics X_{ic}, computed with the k^{th} part of the data removed. We then form a predicted binary purchase decision, $\hat{b}_i = 1\left(\hat{r}^{(-k)}(X_{ic}; P_{ic}) \geq 0.5\right)$ for each observation in the validation set, where P_{ic} is the randomized TIOLI price actually faced by household i in compound c. We repeat this procedure for all ten folds. We then estimate the accuracy of TIOLI for out-of-sample prediction of behavior under the TIOLI mechanism itself based on the share of correct predictions in the full TIOLI sample. The resulting accuracy rate is 76.1 percent.

To calculate the analogous accuracy rate of prediction based on the BDM mechanism, we randomly divide the BDM sample into 10 groups of approximately equal size. Since the validation set is drawn from the TIOLI sample, this procedure serves to replicate the sampling variability and sample size effects of the cross-validation procedure within the TIOLI sample. We estimate Equation (7) from the main text for the test set via ordinary-least-squares and then estimate $w\hat{t}p^{(-k)}(X_{ic})$ for each observation in the corresponding validation set from the TIOLI sample. Based on this estimation, we form a predicted binary purchase decision $\hat{b}_i = 1\left(w\hat{t}p^{(-k)}(X_{ic}) \geq P_{ic}\right)$, for each observation in the validation set, where again P_{ic} is the randomized TIOLI price actually faced by household i in compound c. We repeat this procedure for all ten folds and estimate the accuracy of BDM for out-of-sample prediction of behavior under the TIOLI mechanism. The share of correct predictions in the full TIOLI sample is 73.9 percent. These accuracy rates compare to a base rate—the accuracy of trivially predicting the most-frequent decision within each validation set—of 56.2 percent. Consistent with the pattern of demand estimated by two mechanisms, TIOLI more accurately predicts affirmative purchase decisions while BDM performs better when predicting refusals.

To explore the relative performance of the two mechanisms in greater depth, we construct ROC curves for both mechanisms and compare model accuracy via their respective areas under the curve (AUCs). The ROC curve plots the sensitivity of the predictive model (the rate of true positives) on the y-axis against the specificity (the rate of true negatives) on the x-axis as we vary the cutoff for predicting a purchase. The simple comparison above is equivalent to setting the cutoff at a 50 percent probability of purchase. The AUC is a commonly used measure to summarize the performance of a classifier over all possible thresholds. Figure A8 displays the AUCs for the BDM and TIOLI models. The diagonal represents the performance of a model that randomly classified each observation. For TIOLI and BDM, the AUCs are 84 percent and 79 percent respectively. While TIOLI outperforms BDM in predicting TIOLI behavior, their performance is remarkably close. We consider this encouraging evidence that, at least in this setting, the noise generated by the BDM mechanism is outweighed by the additional information it provides.

J Mechanism Effects, Detail

In this section we extend the discussion in Section 6.2 by providing further analysis of the magnitude and potential sources of the differences between BDM and TIOLI-elicited WTP.

J.1 Comparing Demand Under BDM and TIOLI

This subsection presents regression estimates of the differences in demand between BDM and TIOLI at the three TIOLI price points, as displayed in Figure 1a of the main text.

In order to perform the comparison, we run a similar regression to that presented in Section 6.2. We estimate

$$\text{buy}_{icp} = \alpha_p + \beta_p \text{BDM}_{ic} + x_{ic}'\gamma + \varepsilon_{icp}, \tag{11}$$

where buy_{icp} indicates whether subject i in compound c purchased at price p (under the TIOLI mechanism), or would have purchased at price p given her bid (under the BDM mechanism), and BDM_{ic} is an indicator for whether subject i was assigned to the BDM mechanism. For each price p, α_p represents the share purchasing under TIOLI and β_p represents the difference in shares between BDM and TIOLI at price p.[4]

The regression results are presented in Table A8. As shown in Columns 1, 3, and 5, the difference between the two mechanisms is significant at the 5 percent level or greater for each of the three prices. The test of joint significance of all three differences yields a p-value of less than 0.001. While the absolute (percentage point) differences are declining with each price, we cannot reject that all three differences are equal ($p = 0.239$ without controls; $p = 0.354$ with controls), and there is no such pattern in relative (percentage) differences. As shown in Columns 2, 4, and 6, the results are virtually unchanged with the inclusion of controls.

J.2 Correlation Between BDM-TIOLI Gap and Risk Aversion

This sub-section presents details on the comparison of the BDM-TIOLI gap across terciles of risk aversion, discussed in Section 6.2 of the main text. In order to implement the comparison between BDM and TIOLI, we collapse the more precise individual WTP information from BDM to the binary purchase indicators generated by TIOLI. Our outcome variable is $\text{buy}_{i,p}$, which represents subject i's purchase decision when facing a price $p \in \{2, 4, 6\}$. For TIOLI subjects, this is just whether they agreed to purchase at the offer price. For BDM subjects, $\text{buy}_{i,p} = 1\{\text{WTP}_i \geq p\}$, where WTP_i is subject i's BDM

[4] Since each BDM subject's bid can be used to simulate purchase behavior at all three prices, each regression contains about three times as many BDM observations as TIOLI observations. We estimate the system (one equation for each of $p = 2, 4, 6$) via seemingly unrelated regression to account for correlation of errors across equations and to conduct cross-equation tests.

bid. We create the variables RA_i^1, RA_i^2, RA_i^3 to indicate that subject i is in the first (most risk-averse), second, or third (least risk-averse) tercile, respectively. We then estimate

$$\text{buy}_{icp} = \sum_{t=1}^{3} \alpha_p^t \text{RA}_i^t + \sum_{t=1}^{3} \beta_p^t \left(\text{RA}_i^t \times \text{BDM}_i\right) + x_{ic}'\gamma + \varepsilon_{icp}, \tag{12}$$

where BDM_i is an indicator for whether subject i was assigned to the BDM mechanism. For each price p, α_p^t represents the purchase probability for TIOLI subjects in the t-th tercile, while β_p^t represents the "BDM effect" in the t-th tercile. The differences without controls are presented in Figure A9. The top panel plots the estimated coefficients $\hat{\beta}_2^1$, $\hat{\beta}_4^1$, $\hat{\beta}_6^1$, with 90 percent confidence intervals, for tercile 1 of risk aversion (the most risk-averse subjects), while the middle and bottom panels plot the same set of coefficients for terciles 2 and 3 (the least risk-averse subjects), respectively. As Figure A9 makes clear, the BDM-TIOLI gap is largest among the most risk-averse subjects (mean BDM effect -0.200, $p = 0.000$), and has largely closed among the least risk-averse subjects (mean BDM effect -0.051, $p = 0.425$). These results are unconditional, but they are robust to controlling for a large set of household controls (see Figure A10) and when testing multiple possible determinants of the BDM-TIOLI gap jointly (see Table A9).

J.3 Correlation Between BDM-TIOLI Gap and Observables

To supplement the analysis of risk aversion presented in Section 6.2, this section presents an exploratory analysis of the correlation between the BDM-TIOLI gap and other relevant observables. For binary observables, we compare the BDM-TIOLI gap between the two levels of the variable; for continuous observables, we break the sample into terciles and compare the top to the bottom tercile. Similar to Equation (12), we estimate

$$\text{buy}_{icp} = \alpha_{0p} + \alpha_{1p} D_i + \beta_{0p}\text{BDM}_i + \beta_{1p}\left(D_i \times \text{BDM}_i\right) + \varepsilon_{icp}, \tag{13}$$

where D_i is an indicator for the subgroup of interest and the other variables are as in Equation (12). For each price p, β_{0p} is the BDM-TIOLI gap (the difference in purchase probabilities) for subjects with $D_1 = 0$, $\beta_{0p} + \beta_{1p}$ is the BDM-TIOLI gap for subjects with $D_1 = 1$, and β_{1p} is the difference between the two sub-groups. We then average the coefficients over the three TIOLI prices to obtain $\overline{\beta}_0$, $\overline{\beta}_1$ and $\overline{\beta}_0 + \overline{\beta}_1$.

Figure A11 presents the results, with levels ($\overline{\beta}_0$ and $\overline{\beta}_0 + \overline{\beta}_1$) in the top panel (Figure A11a) and differences ($\overline{\beta}_1$) in the bottom panel (Figure A11b). First, household wealth is associated with a smaller BDM-TIOLI gap, but not to the same extent as risk aversion (see Figure A9 and discussion above). Second, the gap among subjects who have attended school is approximately zero, although this is imprecisely estimated since only nine percent of subjects have ever attended school. On the other hand, the gap is wider among subjects who scored in the top tercile of the digit span test. Third, the gap is narrower among subjects who have a child age 0 to 5, and narrower still (with a point estimate close to zero) if one or more children has had a case of diarrhea in the previous two weeks. This may suggest that respondents with more at stake took the exercise more seriously. On the other hand, the gap is significantly wider among respondents whose water samples were in the top tercile in terms of *E. coli* (those with the poorest water quality). This is somewhat surprising, especially since the gap is largely unaffected by turbidity, which, unlike *E. coli*, is visible.

The results in Figure A11 test one covariate at a time, but the results are generally similar when we test several covariates jointly, as shown in Table A9. Note that water quality is not included in this comparison because of sample size limitations – we collected the risk aversion measure only in the one-year sample of villages (8 of 15 villages), and tested water only in a 50 percent subsample of households. Table A10 repeats this exercise for the full sample with just the variables available for all households household survey, and shows that the coefficients and statistical significance of these variables are similar in single and joint tests.

J.4 BDM and TIOLI Experimental Sub-Treatments

This section describes the experimental sub-treatments designed to test mechanisms behind the BDM-TIOLI gap. We test the "standard" presentation of the BDM and TIOLI mechanisms against four sub-treatments that incorporated modifications to the sales scripts. Descriptive graphs demand across treatments are provided in Figures A12 and A13, with formal statistical tests reported in Tables A11 and A12.[5]

The first two sub-treatments were designed to test the hypothesis that the stated prices in the TIOLI treatment could cause respondents to anchor their valuations to those prices. In the "anchoring" treatments for both BDM and TIOLI, we informed subjects that the price of the filter in the Tamale town market (the nearest market town) was GHS 20. Based on our pilot results, we believed this information would dominate any conveyed in the TIOLI price, placing both mechanisms on equal footing and allowing us to estimate any anchoring or signaling effects from the offer price. However, these anchoring treatments did not produce any consistent effect on BDM bids or TIOLI purchase behavior. There was a significant effect on TIOLI demand at GHS 4 (-0.233, $p < 0.05$), but there was no effect at the other TIOLI prices or in BDM bids.

We also included a "random TIOLI" sub-treatment, in which the TIOLI offer price was drawn by the respondent from a cup of numbered wooden beads, the same mechanism used to determine the BDM price. The aim was to make salient the arbitrariness of the

[5]Table A11 presents the results of two tests that compare the distributions of the BDM sub-treatments using both the Wilcoxon-Mann-Whitney rank-sum and Kolmogorov-Smirnov tests. Cluster-robust significance levels for the distributional tests are constructed via a bootstrap percentile method. We pool data from the two treatments being compared, draw block-bootstrap samples, where the compound is the block, and then randomly assign placebo treatments by compound and run the distributional test in question. Since the placebo treatments are randomly generated, the null hypothesis of equality of distribution is true by construction. By sampling compounds and assigning placebo treatments by compound, we preserve the clustering structure in the data. We repeat this for 500 bootstrap repetitions, and then obtain a p-value for our test by finding where the original test statistic falls in the distribution of bootstrap test statistics.

TIOLI prices and reduce the likelihood that they were serving as signals of quality. Based on our pilot results and the evidence that in some settings BDM bids are sensitive to the underlying price distribution (Bohm et al. 1997; Mazar et al. 2014), we hypothesized that the randomness in the price draw may contribute to the BDM-TIOLI gap, through a failure to reduce compound lotteries, subjects' general discomfort with randomness and ambiguity, or other departures from expected utility maximization. However, demand under the random TIOLI treatment was statistically indistinguishable from standard TIOLI, indicating that our efforts to equate the perceived randomness in the two mechanisms had no effect on subjects' purchase behavior. We note, however, that our modifications were designed to increase the perceived randomness of the TIOLI mechanism, and as we speculate in the text, reducing the perceived randomness of the BDM mechanism may narrow the gap.

Unrelated to explaining potential differences between BDM and TIOLI, we also conducted a "market study BDM" treatment in which we told respondents that we were using the information from the study to help decide on the future price of the filter in similar villages. If strategic bidding was important, then this sub-treatment could lead to enhanced strategic bidding and decrease BDM bids. However, we found that the market study treatment increased valuations, with marginal statistical significance.

J.5 Comparing Demand for Soap

As we argue in Section 6.2, the gap in elicited WTP between BDM and TIOLI also does not appear to be driven by lack of familiarity with the filter and uncertainty of its benefits. Although the sale of soap was primarily intended to be a practice round for the elicitation mechanism, the data provides suggestive evidence of the BDM-TIOLI gap for a more familiar product. Using these data, we find a similar difference in elicited WTP between the mechanisms: as shown in Table A13, BDM predicts between 8 and 45 percentage points lower purchase at the TIOLI price points for the soap.

J.6 Ex Post Regret

As discussed in Section 6.2, 19.2 percent of BDM respondents stated that they wished they had bid more. As shown in table A14, the proportion expressing regret is highest for those who narrowly missed winning in BDM: roughly 40 percent of those who missed by GHS 1 or less wished that they had bid more, with this percentage declining to approximately 12 percent among those who missed by GHS 5 to 10. To estimate the influence of regret on elicited WTP, we calculate what the adjustment to BDM bids would have been if all respondents who wished they had bid more had actually bid the value of the draw. Because those whose bid exceeded their draw cannot express regret, we apply an adjustment to this group that equals the average adjustment of BDM losers who have similar bids. Calculated in this manner, the average adjustment across all subjects equals GHS 0.6, or about 60 percent of the gap between BDM and TIOLI. Note that this likely represents an upper bound on underbidding due to regret because responses were not tied to an actual purchase decision. The share of respondents who actually offered to pay more than their final bid is substantially less than those who stated they wished they had more, at 5.4 percent. If we adjust the bids of those who offered to pay more up to the value of the draw (and apply an adjustment for those whose bid exceeded the draw following the procedure described above) this would result in an average increase in WTP of GHS 0.07, which would account for little of the BDM-TIOLI gap.

Although the upwards revision of bids after the price draw could result from respondents misunderstanding the BDM mechanism, it is also consistent with non-expected utility maximization in which a respondent revises her reference point upwards when the price is revealed. Further, a substantial share of TIOLI subjects, 17.0 percent, attempted to bargain with surveyors over the randomly drawn price. As noted in the text, we take this as evidence that both mechanisms may have seemed unusual to respondents who are unaccustomed to fixed prices.

K Using BDM in the Field

This section offers additional discussion of the practical tradeoffs between BDM and TIOLI for researchers considering using one or the other method. The key advantages of BDM are precision in measuring WTP, the ability to separately identify selection by WTP and the impacts of price paid, and the ability to estimate heterogeneous treatment effects with respect to WTP. The key disadvantage is complexity, which carries both fixed costs – time to tailor BDM to local context and train enumerators – and variable costs – time to explain BDM to subjects, conduct practice rounds, etc. Which method is preferable will depend on context and the questions the researcher is asking, but the relative advantages and disadvantages just mentioned offer some general guidelines.

First, the number of prices at which the researcher would like to measure demand affects the choice. The more prices of interest there are, the more advantageous BDM is likely to be, since more prices will require ever greater TIOLI sample sizes. Second, if the causal effect of price paid is of interest and a surprise randomized discount is not feasible, then BDM becomes attractive, since TIOLI cannot separately identify selection by WTP from the effect of price paid. Third, the extent to which it is plausible that treatment effects vary by WTP affects the choice. If there is strong prior evidence that treatment effects are constant, or constant with respect to WTP, then this tips the balance towards TIOLI. Fourth, developing a context-specific BDM protocol is a significant investment, and it is important to spend time explaining to and practicing with subjects. Based on our experience, multiple demonstration rounds with a different product or products, emphasis on the bid as the subject's optimal response, training the subjects to understand their bid as their maximum WTP, and the understanding check after respondents stated their bids are essential to successful implementation. These procedures have been emphasized elsewhere in the laboratory literature as important for eliciting accurate WTP through BDM (Plott and Zeiler 2005), although more research is needed on how each detail may contribute to subjects' understanding. Finally, the cost of each observation (including the cost of the item itself, the cost of collecting follow-up data on use or the outcome of interest, etc.) affects the tradeoff. If each observation is very cheap, then the burden of increased sample size from TIOLI is less of a concern. If each observation is relatively expensive, it becomes more important to obtain as much information as possible from each subject and the balance tilts towards BDM.

References

Ashraf, N., J. Berry, and J. M. Shapiro, "Can Higher Prices Stimulate Product Use? Evidence from a Field Experiment in Zambia," *American Economic Review*, 2010, *100* (5), 2383–2413.

Bennett, D., "Does Clean Water Make You Dirty? Water Supply and Sanitation in the Philippines," *Journal of Human Resources*, 2012, *47* (1), 146–173.

Bohm, P., J. Linden, and J. Sonnegård, "Eliciting Reservation Prices: Becker-DeGroot-Marschak Mechanisms vs. Markets," *Economic Journal*, 1997, *107* (443), 1079–1089.

Brinch, C. N., M. Mogstad, and M. Wiswall, "Beyond LATE with a Discrete Instrument," *Journal of Political Economy*, 2017, *125* (4), 985–1039.

Browning, M. and J. Carro, "Heterogeneity and Microeconometrics Modeling," in T. P. R. Blundell, W. Newey, ed., *Advances in Theory and Econometrics, Vol. 3*, Cambridge, England: Cambridge University Press, 2007.

Cohen, J. and P. Dupas, "Free Distribution or Cost-Sharing? Evidence from a Randomized Malaria Prevention Experiment," *Quarterly Journal of Economics*, 2010, *125* (1), 1–45.

Da, Z., U. G. Gurun, and M. Warachka, "Frog in the Pan: Continuous Information and Momentum," *Review of Financial Studies*, 2014, *27* (7), 2171–2218.

Ellsberg, D., "Risk, Ambiguity, and the Savage Axioms," *Quarterly Journal of Economics*, 1961, *75* (4), 643–669.

Global Burden of Disease Collaborative Network, "Global Burden of Disease Study 2016 Results," 2017. `http://ghdx.healthdata.org/gbd-results-tool`.

Greenstone, M. and B. K. Jack, "Envirodevonomics: A Research Agenda for an Emerging Field," *Journal of Economic Literature*, 2015, *53* (1), 5–42.

Heckman, J. J. and E. J. Vytlacil, "Econometric Evaluation of Social Programs, Part II: Using the Marginal Treatment Effect to Organize Alternative Econometric Estimators to Evaluate Social Programs, and to Forecast their Effects in New Environments," in J. Heckman and E. Leamer, eds., *Handbook of Econometrics*, Vol. 6B 2007, pp. 4875–5143.

_ , S. Urzúa, and E. J. Vytlacil, "Understanding Instrumental Variables in Models with Essential Heterogeneity," *Review of Economics and Statistics*, 2006, *88* (3), 389–432.

Kling, J. R., J. B. Liebman, and L. F. Katz, "Experimental Analysis of Neighborhood Effects," *Econometrica*, 2007, *75 (1)*, 83–119.

Kowalski, A. E., "Doing More When You're Running LATE: Applying Marginal Treatment Effect Methods to Examine Treatment Effect Heterogeneity in Experiments," 2016. NBER Working Paper 22363.

Kremer, M., J. Leino, E. Miguel, and A. P. Zwane, "Spring Cleaning: Rural Water Impacts, Valuation, and Property Rights Institutions," *Quarterly Journal of Economics*, 2011, *126*, 145–205.

Mazar, N., B. Koszegi, and D. Ariely, "True Context-dependent Preferences? The Causes of Market-Dependent Valuations," *Journal of Behavioral Decision Making*, 2014, *27* (3), 200–208.

Nevo, A., "Empirical Models of Consumer Behavior," *Annual Review of Economics*, 2011, *3* (1), 51–75.

Peltzman, S., "The Effect of Government Subsidies-in-Kind on Private Expenditures: The Case of Higher Education," *Journal of Political Economy*, 1973, *81* (1), 1–27.

Plott, C. R. and K. Zeiler, "The Willingness to Pay-Willingness to Accept Gap, the Endowment Effect, Subject Misconceptions, and Experimental Procedures for Eliciting Valuations," *American Economic Review*, 2005, *95* (3), 530–545.

Reis, R., "Inattentive Consumers," *Journal of Monetary Economics*, 2006, *53* (8), 1761–1800.

Tibshirani, R., "Regression Shrinkage and Selection via the Lasso," *Journal of the Royal Statistical Society. Series B (Methodological)*, 1996, *58* (1), 267–288.

Tobin, J., "Money and finance in the macroeconomic process," *Journal of Money, Credit and Banking*, 1982, *14* (2), 171–204.

Table A1: Constant-Effects Instrumental Variables: Flexible Demand Curve
Dependent Variable: Child age 0 to 5 has had diarrhea over previous two weeks

	Combined all subjects		TIOLI subjects		BDM subjects	
	(1)	(2)	(3)	(4)	(5)	(6)
A. One-month followup						
Bought Filter	-0.057*	-0.066**	-0.083	-0.085*	-0.036	-0.048
	(0.033)	(0.032)	(0.051)	(0.049)	(0.047)	(0.040)
Mean dependent variable	0.145	0.145	0.149	0.149	0.142	0.142
First-stage F-statistic	441.0	228.4	111.2	96.1	504.0	338.8
Number of compounds	472	472	244	244	229	229
Number of subjects	786	786	418	418	368	368
Number of children	1244	1244	665	665	579	579
B. One-year followup						
Bought Filter	0.116*	0.138**	0.142	0.211**	0.115	0.127
	(0.067)	(0.067)	(0.100)	(0.100)	(0.089)	(0.085)
Mean dependent variable	0.241	0.241	0.215	0.215	0.262	0.262
First-stage F-statistic	132.3	80.2	58.8	36.0	179.3	170.0
Number of compounds	247	247	121	121	126	126
Number of subjects	387	387	197	197	190	190
Number of children	539	539	266	266	273	273
Controls	No	Yes	No	Yes	No	Yes
Village FEs	No	Yes	No	Yes	No	Yes

Notes: Each column displays the results of a linear two-stage least squares regression of child diarrhea status at the child level on filter purchase. For TIOLI subjects, filter purchase is instrumented dummies for each level of the randomly assigned TIOLI price (GHS 2, 4, 6). For BDM subjects, filter purchase is instrumented by a quadratic in the random BDM price draw. Controls include all variables (other than BDM bid) listed in Table 1. Missing values of control variables are set to 0, and dummy variables are included to indicate missing values. Standard errors clustered at the compound (extended family) level in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table A2: Relationship between Use and Willingness to Pay

	Filter present and unbroken (1)	Storage vessel contains water (2)	Clay pot contains water (3)	Usage index (4)
A. Short-term effects				
Bid (GHS)	-0.010	-0.008	-0.009	-0.022
	(0.010)	(0.012)	(0.013)	(0.021)
Mean dep. var.	0.877	0.753	0.728	-0.003
Adj. R-sqd.	0.002	-0.002	-0.002	0.002
Num. Obs.	235	235	235	235
B. One-year effects				
Bid (GHS)	0.013	0.027*	-0.013	0.018
	(0.014)	(0.014)	(0.012)	(0.021)
Mean dep. var.	0.641	0.486	0.380	0.066
Adj. R-sqd.	-0.002	0.016	-0.002	-0.003
Num. Obs.	142	142	142	142

Notes: The sample includes those subjects in the BDM treatment who purchased the filter, i.e., drew a price less than or equal to their bid. Each column presents the results of a separate regression of the depend variable, listed in the column heading, on the willingness to pay, i.e, the subject's bid in BDM. Usage index is the average of the normalized values of the three individual usage measures. Usage measures are observed by the enumerator at indicated follow-up survey. Standard errors clustered at the compound (extended family) level in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table A3: Attrition -- 1-month survey

	Baseline (1)	Surveyed (2)	Not Surveyed (3)	Difference (4)
Share of households		0.871	0.129	
Assigned to BDM treatment	0.480	0.479	0.485	0.006
	[0.500]	[0.500]	[0.501]	(0.043)
Number of respondents in compound	3.593	3.580	3.681	0.101
	[2.323]	[2.378]	[1.914]	(0.207)
Respondent's husband lives in compound	0.794	0.804	0.730	-0.074**
	[0.404]	[0.397]	[0.445]	(0.037)
One or more children age 0-5 in household	0.723	0.713	0.791	0.078**
	[0.448]	[0.452]	[0.408]	(0.035)
Number of children age 0-5 in household	1.569	1.561	1.620	0.059
(conditional on positive)	[0.801]	[0.760]	[1.017]	(0.111)
Num. children age 0-5 w. diarrhea in prev. 2 wks.	0.337	0.328	0.388	0.059
(among households with children age 0-5)	[0.592]	[0.597]	[0.563]	(0.055)
Num. children age 6-17 w. diarrhea in prev. 2 wks.	0.075	0.081	0.033	-0.047**
(among households with children age 0-5)	[0.335]	[0.349]	[0.181]	(0.023)
Respondent ever attended school	0.090	0.088	0.104	0.016
	[0.286]	[0.283]	[0.307]	(0.025)
Respondent's spouse ever attended school	0.233	0.225	0.300	0.075
	[0.423]	[0.418]	[0.464]	(0.078)
Wealth index	0.132	0.115	0.245	0.130
	[1.555]	[1.556]	[1.549]	(0.131)
Improved water source, all year	0.187	0.185	0.202	0.017
	[0.390]	[0.389]	[0.403]	(0.036)
Treats water with an effective method	0.115	0.107	0.166	0.059*
	[0.319]	[0.309]	[0.373]	(0.035)
Water quality: E. coli (MPN, z-score)	-0.052	-0.050	-0.067	-0.018
	[0.949]	[0.958]	[0.877]	(0.108)
Water quality: Turbidity (index, z-score)	-0.065	-0.080	0.072	0.151
	[0.997]	[0.985]	[1.100]	(0.141)
Bid for filter (GHS)	3.051	3.022	3.243	0.221
(among BDM respondents)	[2.268]	[2.247]	[2.414]	(0.276)
Filter draw (GHS)	4.650	4.621	4.842	0.221
(among BDM respondents)	[3.663]	[3.669]	[3.641]	(0.434)
Filter offer price (GHS)	3.824	3.864	3.548	-0.316
(among TIOLI respondents)	[1.616]	[1.604]	[1.682]	(0.199)

Notes: Standard deviations in brackets. Standard errors clustered at the compound (extended family) level in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table A4: Attrition -- 1-year survey

	Baseline (1)	Surveyed (2)	Not Surveyed (3)	Difference (4)
Share of households		0.904	0.096	
Assigned to BDM treatment	0.479	0.488	0.400	-0.088
	[0.500]	[0.500]	[0.494]	(0.084)
Number of respondents in compound	3.676	3.620	4.215	0.596
	[2.552]	[2.523]	[2.781]	(0.378)
Respondent's husband lives in compound	0.800	0.816	0.646	-0.170**
	[0.400]	[0.388]	[0.482]	(0.070)
One or more children age 0-5 in household	0.722	0.725	0.692	-0.033
	[0.448]	[0.447]	[0.465]	(0.062)
Number of children age 0-5 in household	1.556	1.578	1.333	-0.245**
(conditional on positive)	[0.763]	[0.771]	[0.640]	(0.114)
Num. children age 0-5 w. diarrhea in prev. 2 wks.	0.367	0.379	0.244	-0.134*
(among households with children age 0-5)	[0.606]	[0.620]	[0.435]	(0.072)
Num. children age 6-17 w. diarrhea in prev. 2 wks.	0.080	0.079	0.088	0.009
(among households with children age 0-5)	[0.288]	[0.288]	[0.288]	(0.052)
Respondent ever attended school	0.099	0.093	0.154	0.061
	[0.298]	[0.290]	[0.364]	(0.047)
Respondent's spouse ever attended school	0.303	0.311	0.200	-0.111
	[0.461]	[0.464]	[0.414]	(0.103)
Wealth index	0.030	0.070	-0.341	-0.411
	[1.574]	[1.553]	[1.726]	(0.269)
Improved water source, all year	0.216	0.220	0.185	-0.035
	[0.412]	[0.414]	[0.391]	(0.069)
Treats water with an effective method	0.107	0.106	0.123	0.017
	[0.310]	[0.308]	[0.331]	(0.043)
Water quality: E. coli (MPN, z-score)	-0.132	-0.113	-0.313	-0.200
	[0.928]	[0.951]	[0.664]	(0.131)
Water quality: Turbidity (index, z-score)	-0.348	-0.353	-0.298	0.055
	[0.532]	[0.536]	[0.495]	(0.088)
Bid for filter (GHS)	3.068	3.150	2.115	-1.035***
(among BDM respondents)	[2.383]	[2.428]	[1.519]	(0.369)
Filter draw (GHS)	4.606	4.632	4.308	-0.324
(among BDM respondents)	[3.585]	[3.614]	[3.290]	(0.631)
Filter offer price (GHS)	3.768	3.778	3.692	-0.085
(among TIOLI respondents)	[1.636]	[1.648]	[1.559]	(0.298)

Notes: Standard deviations in brackets. Standard errors clustered at the compound (extended family) level in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table A5: Casual Effect of Prices

	Filter Present and Undamaged (1)	Storage Vessel Contains Water (2)	Clay Pot Contains Water (3)	Usage Index (4)
A. Short-term effects				
Draw	0.017	0.037*	-0.003	0.043
	(0.018)	(0.022)	(0.024)	(0.038)
Mean Dependent Variable	0.877	0.753	0.728	-0.003
R-squared	0.020	0.033	0.010	0.025
Observations	235	235	235	235
B. One-year effects				
Draw	-0.013	0.021	0.019	0.018
	(0.034)	(0.033)	(0.033)	(0.051)
Mean Dependent Variable	0.641	0.486	0.380	0.066
R-squared	0.029	0.033	0.010	0.015
Observations	142	142	142	142

Notes: The sample includes those subjects in the BDM treatment who purchased the filter, i.e., drew a price less than or equal to their bid. Each column presents the results of a separate regression of the dependent variable, listed in the column heading, on BDM draw and BDM bid. See Section 5 for discussion of data. Usage index is the average of the normalized values of the three individual usage measures. Usage measures are observed by enumerator at indicated follow-up survey. Standard errors clustered at the compound (extended family) level in parentheses.

Table A6: Estimated Impacts of Pricing Policy

	Price (GHS)						
	0	1	2	3	4	5	6
Share Purchasing	1.00	0.94	0.73	0.46	0.31	0.19	0.11
A. Constant one-month effects							
Diarrhea cases averted per household (conditional on purchase)	1.43	1.43	1.43	1.43	1.43	1.43	1.43
Diarrhea cases averted per household (unconditional)	1.43	1.35	1.05	0.66	0.44	0.28	0.15
DALYs averted per household (conditional on purchase)	0.041	0.041	0.041	0.041	0.041	0.041	0.041
DALYs averted per household (unconditional)	0.041	0.038	0.030	0.019	0.013	0.008	0.004
Average social cost per DALY (USD)	369	369	369	369	369	369	369
Marginal cost per DALY (USD)		369	369	369	369	369	369
B. Average of one-month effects and one-year effects							
Diarrhea cases averted per household (conditional on purchase)	-1.09	-0.72	0.62	2.73	4.29	5.73	6.81
Diarrhea cases averted per household (unconditional)	-1.09	-0.68	0.46	1.26	1.33	1.10	0.73
DALYs averted per household (conditional on purchase)	-0.031	-0.021	0.018	0.077	0.121	0.162	0.193
DALYs averted per household (unconditional)	-0.031	-0.019	0.013	0.036	0.038	0.031	0.021
Average social cost per DALY (USD)	–	–	849	194	123	92	78
Marginal cost per DALY (USD)		–	–	–	361	128	79

Notes: In Panel A, short-term impacts on diarrhea are assumed to be constant and last for one year. Panel B assumes the average of short- and long-term impacts last for one year. In Panel B, the short-term impacts are constant and the long-term impacts are linear in willingness-to-pay. Diarrhea incidence is converted to DALYs at the rate of 35.3 cases per year to one DALY, using data from the Global Burden of Disease Collaborative Network (2017). The average social cost does not account for revenue generated from sales. The marginal cost per DALY is computed as the difference in costs between price $P - 0.5$ and price $P + 0.5$ divided by the difference in DALYs averted between price $P - 0.5$ and price $P + 0.5$. Missing entries in the average and marginal cost rows indicate that costs cannot not be computed because treatment effects are negative for average or marginal households at the prices indicated.

Table A7: Correlates of Willingness to Pay, LASSO Regularization

	Regularized Coefficient	
Variable	BDM	TIOLI
BDM soap bid	0.322	-
Purchased TIOLI soap	-	0.232
Village fixed effect, V08	-0.007	-0.164
Village fixed effect, V09	-	-0.050
Village fixed effect, V12	-	0.230
Highest education attained, kindergarten	0.075	-
Highest education attained, other	0.001	-
Spouse occupation, animal tending	0.023	-
Primary occupation, non-ag wage labor	-	-0.010
Primary occupation, household enterprise	-	0.001
Household has mobile phone	0.019	-
Number of phones in household	0.016	-
Number of cutlasses in household	-0.002	-
Household has chair	-0.010	-
Primary water source, dry season: dugout	-0.005	-
Primary water source, dry season: dam	-	0.009
Household treats water by boiling	0.004	-
Believes hygiene prevents diarrhea	-	0.062

Notes: Regularized coefficients reported for all features with non-zero coefficients in test sample using regularization parameter (λ) with minimum out-of-sample error rate in training sample. See Appendix I.1 for details.

Table A8: Demand Comparison by Purchase Mechanism
Take-it-or-leave-it Price Points

	Dependent Variable: WTP $\geq$ Price (GHS)					
	Price = 2		Price = 4		Price = 6	
	(1)	(2)	(3)	(4)	(5)	(6)
Difference (BDM - TIOLI)	-0.182***	-0.171***	-0.163***	-0.152***	-0.100**	-0.103***
	(0.033)	(0.033)	(0.052)	(0.052)	(0.040)	(0.039)
Mean TIOLI Purchase	0.915	0.915	0.473	0.473	0.207	0.207
Controls:	No	Yes	No	Yes	No	Yes
Number of BDM Respondents	607	607	607	607	607	607
Number of TIOLI Respondents	246	246	224	224	188	188
Number of clusters	395	395	390	390	376	376

Notes: BDM acceptance rate calculated based on share of respondents bidding greater than or equal to the evaluated price. TIOLI acceptance rate equal to share of respondents offered the evaluated price who agreed to purchase. Controls include all variables listed in Table 1 (except BDM bid). Missing values of the control variables are set to 0, and dummy variables are included to indicate missing values. Standard errors clustered at the compound (extended family) level in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$. Joint p-value testing significance of BDM across all three equations (GHS 2, 4, 6): 0.000 without controls; 0.000 with controls. P-value testing equality of BDM across equations: 0.239 without controls; 0.354 with controls.

Table A9: Relationship between Household Observables and BDM-TIOLI Gap Risk Aversion Sample (One-Year Follow-Up Villages)

	Pairwise (1)	Joint (2)
Top vs. bottom tercile of risk aversion	-0.138^{*}	-0.150^{*}
	(0.081)	(0.083)
First principal component of durables ownership	0.030	0.030
	(0.028)	(0.030)
Respondent has ever attended school	0.141	0.180^{*}
	(0.113)	(0.107)
Has child age 0-5	-0.002	0.004
	(0.095)	(0.100)
Husband lives in compound	0.066	0.037
	(0.094)	(0.098)
All-year access to improved water source	-0.150	-0.104
	(0.106)	(0.097)
Currently treats water	-0.077	0.011
	(0.129)	(0.155)
Ambiguity aversion category (more is more AA)	0.003	0.003
	(0.009)	(0.010)
Total digit span score forward + backward	-0.028	-0.035^{*}
	(0.018)	(0.019)
Number of compounds	233	233
Number of households	399	399

Notes: This table presents estimates of the interaction between the mean BDM-TIOLI gap (the probability of purchase at 2, 4 or 6 GHS) and the household observable indicated. Column (1) shows the results of pairwise comparisons: an indicator for whether the household would agree to purchase the filter at the given price on an indicator for the BDM treatment, a level term for the indicated covariate, and the interaction between the two. Column (2) shows the estimated interaction terms in a joint regression. Coefficients are estimated for offer prices of 2, 4 and 6 and then averaged across the three prices, with standard errors calculated by SUR. The sample consists of households surveyed in the one-year followup (conducted in 8 of the 15 study villages) in the top or bottom tercile of risk aversion. Standard errors clustered at the compound (extended family) level in parentheses.

Table A10: Relationship between Household Observables and BDM-TIOLI Gap
All Villages; Household Survey Measures Only

	Pairwise (1)	Joint (2)
First principal component of durables ownership	0.007	0.001
	(0.016)	(0.016)
Respondent has ever attended school	0.137**	0.128**
	(0.061)	(0.059)
Has child age 0-5	0.138***	0.129***
	(0.045)	(0.045)
Husband lives in compound	0.050	0.036
	(0.053)	(0.053)
All-year access to improved water source	0.056	0.072
	(0.060)	(0.059)
Currently treats water	0.022	0.028
	(0.075)	(0.071)
Number of compounds	556	556
Number of households	1265	1265

Notes: This table presents estimates of the interaction between the mean BDM-TIOLI gap (the probability of purchase at 2, 4 or 6 GHS) and the household observable indicated. Column (1) shows the results of pairwise comparisons: an indicator for whether the household would agree to purchase the filter at the given price on an indicator for the BDM treatment, a level term for the indicated covariate, and the interaction between the two. Column (2) shows the estimated interaction terms in a joint regression. Coefficients are estimated for offer prices of 2, 4 and 6 and then averaged across the three prices, with standard errors calculated by SUR. Standard errors clustered at the compound (extended family) level in parentheses.

**Table A11: Equality of Bid Distributions
Comparison with Standard BDM**

	Market (1)	Anchor (2)
A. Wilcoxon		
Z-statistic	2.754	−0.900
P-value	0.022	0.748
Num. Obs.	411	408
B. Kolmogorov-Smirnov		
D-statistic	0.144	0.058
P-value	0.050	0.777
Num. Obs.	411	408

Notes: This table reports results of nonparametric tests for equality of bid distributions across BDM treatments. The anchoring and marketing BDM treatments (describe in the text) are separately compared to the standard BDM treatment. P-values robust to clustering at the compound level are calculated via randomization inference.

Table A12: Differences Between TIOLI sub-treatments

	Price=2	Price=4	Price=6
	(1)	(2)	(3)
Random TIOLI	0.018	-0.133	0.022
	(0.063)	(0.104)	(0.084)
Anchoring TIOLI	0.066	-0.233**	-0.087
	(0.059)	(0.114)	(0.075)
Constant	0.890	0.600	0.232
	(0.053)	(0.083)	(0.056)
Mean Dependent Variable	0.915	0.473	0.207
R-squared	0.009	0.036	0.014
Observations	246	224	188

Notes: This table reports results of a linear probability model for purchase of the filter at the TIOLI price indicated in the column header. The ommitted category is standard TIOLI. The p-values for joint tests across equations are calculated from SUR estimation. P-value for joint test that coefficient on Random TIOLI=0 in all three equations: 0.587. P-value for joint test that coefficient on Anchoring TIOLI=0 in all three equations: 0.077. Standard errors clustered at the compound (extended family) level in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table A13: Soap Demand Comparison by Purchase Mechanism
Soap Take-it-or-leave-it Price Points

	Dependent Variable: WTP $\geq$ Price (GHS)					
	Price = 0.3		Price = 0.5		Price = 0.7	
	(1)	(2)	(3)	(4)	(5)	(6)
Difference (BDM - TIOLI)	-0.077***	-0.067**	-0.169***	-0.171***	-0.454***	-0.449***
	(0.028)	(0.029)	(0.053)	(0.054)	(0.058)	(0.055)
Mean TIOLI Purchase	0.947	0.947	0.757	0.757	0.663	0.663
Controls:	No	Yes	No	Yes	No	Yes
Number of BDM Respondents	607	607	607	607	607	607
Number of TIOLI Respondents	189	189	148	148	172	172
Number of clusters	364	364	359	359	356	356

Notes: BDM acceptance rate calculated based on share of respondents bidding greater than or equal to the evaluated price. TIOLI acceptance rate equal to share of respondents offered the evaluated price who agreed to purchase. Controls include all variables listed in Table 1. Missing values of the control variables are set to 0, and dummy variables are included to indicate missing values. Standard errors clustered at the compound (extended family) level in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$. Joint p-value testing significance of BDM across all three equations (GHS 0.3, 0.5, 0.7): 0.000 without controls; 0.000 with controls. P-value testing equality of BDM across equations: 0.000 without controls; 0.000 with controls.

Table A14: BDM: Respondents Interested in Changing Their Bid Ex Post

Difference between draw and bid	Number whose draw exceeded bid (1)	Frac. who wished they bid more (2)	Frac. who tried to pay more (3)
Difference<1	20	0.45	0.30
1≤Difference<2	45	0.33	0.13
2≤Difference<3	44	0.25	0.07
3≤Difference<4	32	0.19	0.06
4≤Difference<5	33	0.12	0.03
Difference>5	159	0.12	0.00
Total	333	0.19	0.05

Notes: Column 1 presents the number of subjects in the BDM treatment whose draw exceeded their bid. Column 2 presents the fraction of those subjects who answered "Yes" to the question "Do you wish you had bid higher?" Column 3 presents the fraction of those subjects who attempted to pay more than their bid after the draw was realized.

Figure A1: The *Kosim* filter

Figure A2: Experimental Timeline for a Typical Village

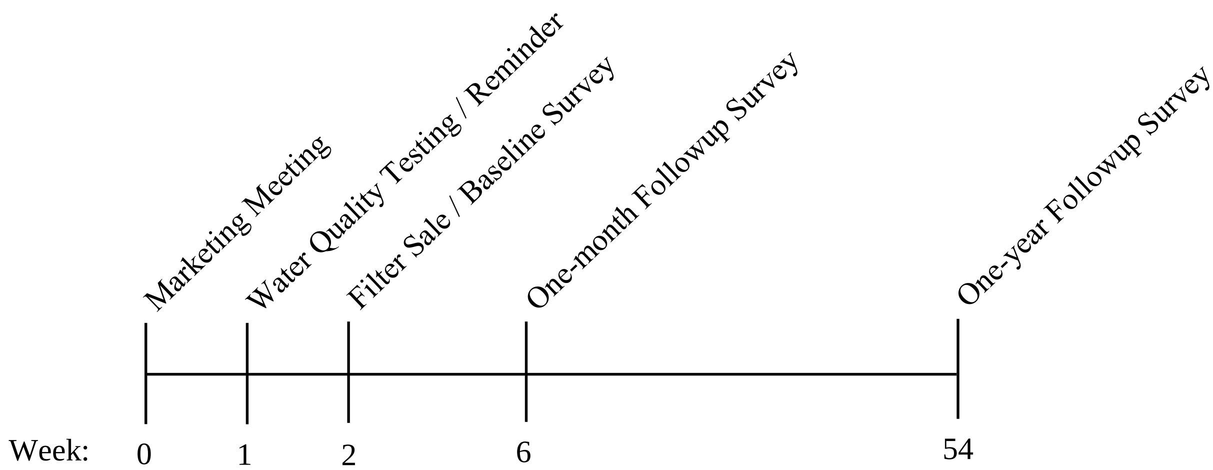

Figure A3: Participant Flow Diagram

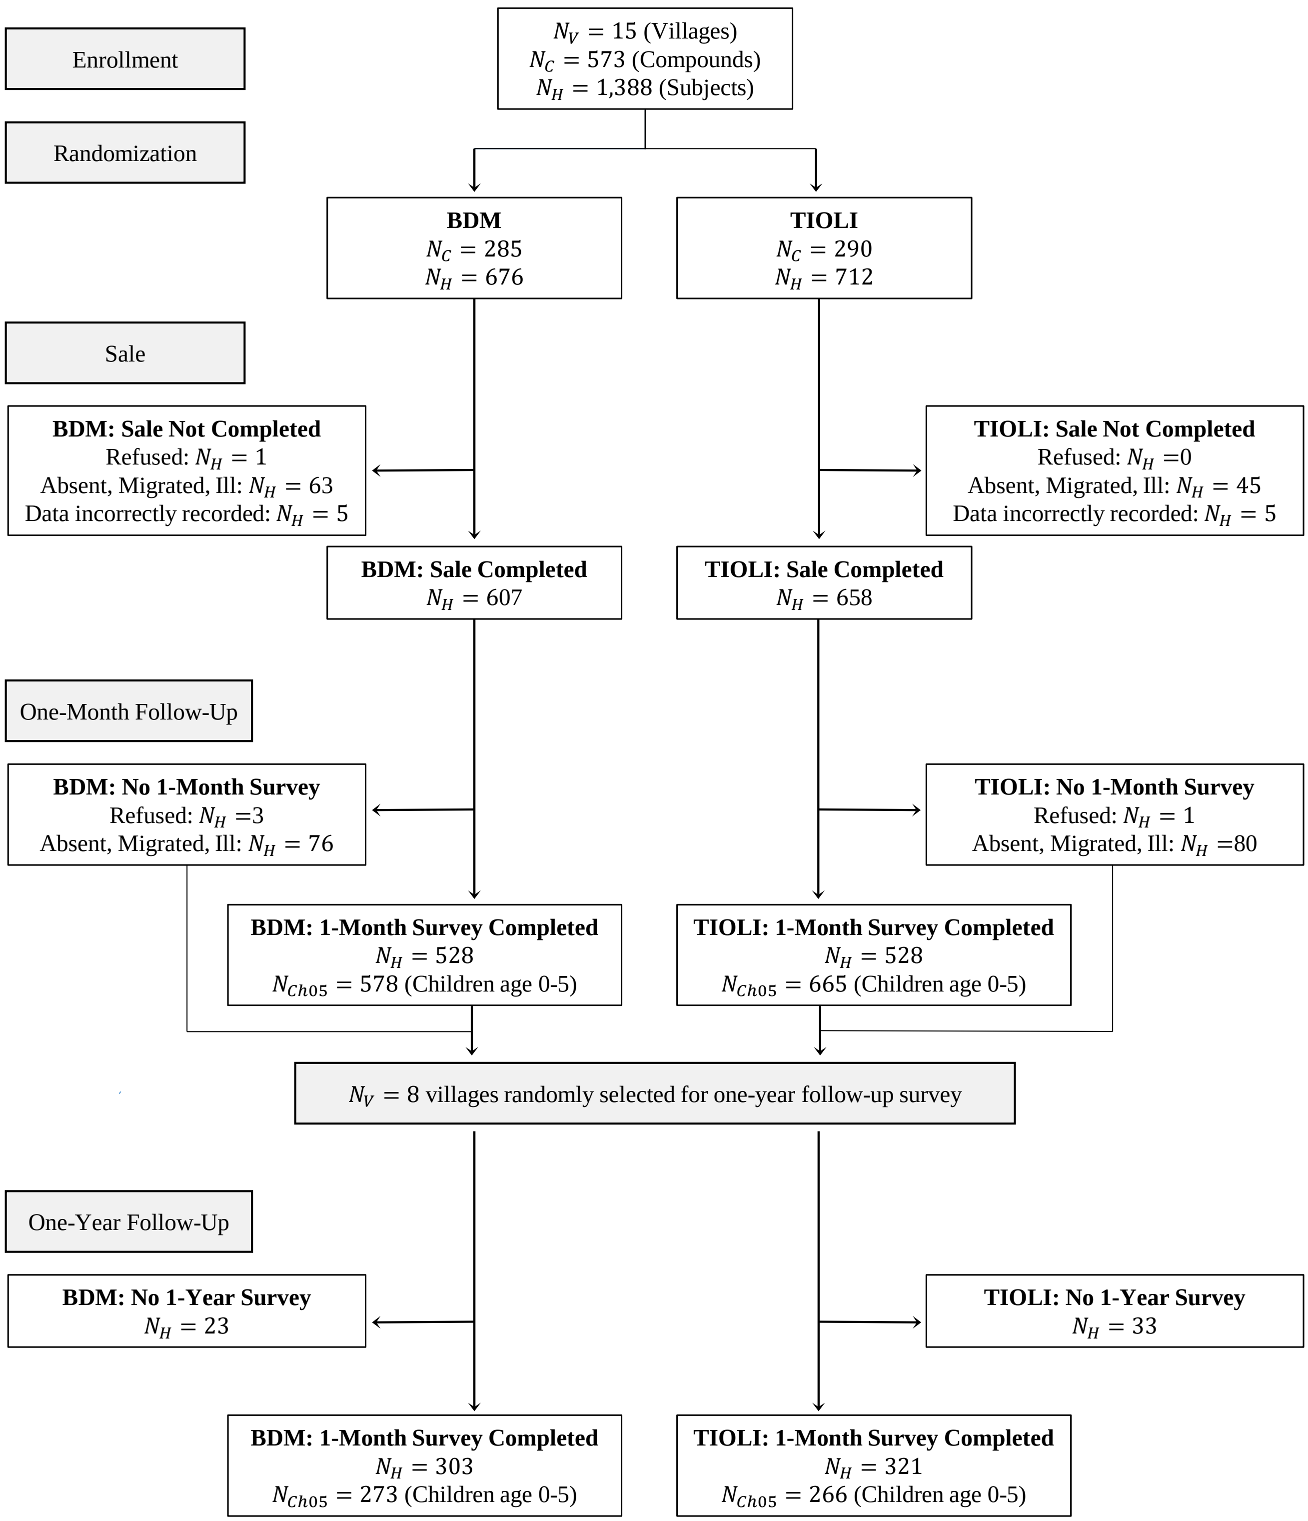

Figure A4: Kernel IV Estimates of Treatment Effects
Ancillary Statistics

(a) One-Month Follow-Up

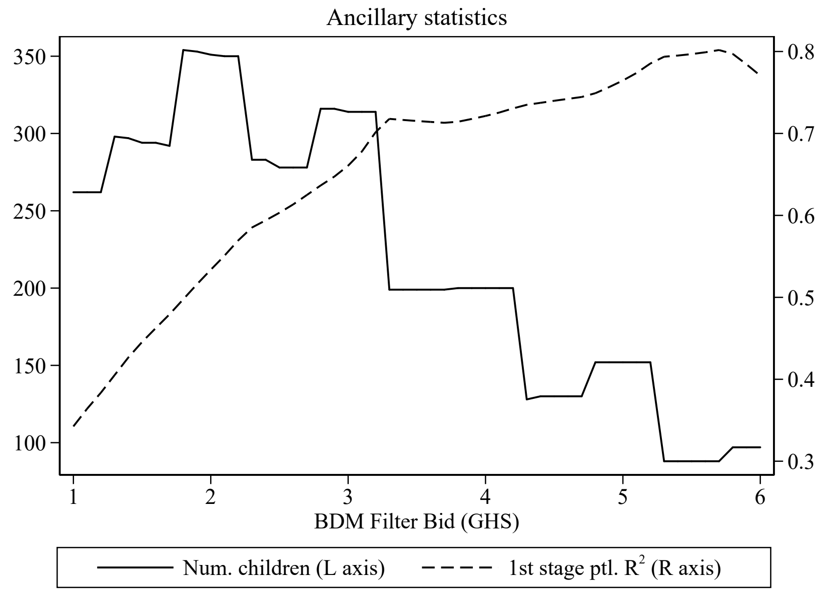

(b) One-Year Follow-Up

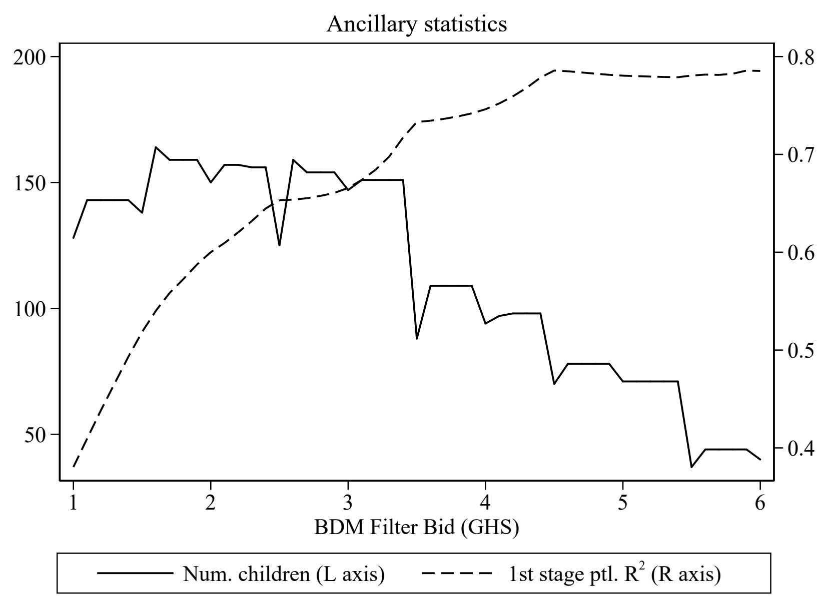

Notes: The solid line (left axis) plots the sample size, i.e., the number of children receiving positive weight in the kernel regression, at each evaluation point WTP = 1.0, 1.1, . . . , 6.0 (GHS). The dashed line (right axis) plots Shea's partial R-squared for the excluded instrument (the BDM price draw) in the first-stage regression at each evaluation point.

**Figure A5: Difference in attrition rates: BDM Winners vs. Losers
1-year follow-up survey; BDM participants with children 0 to 5 years old**

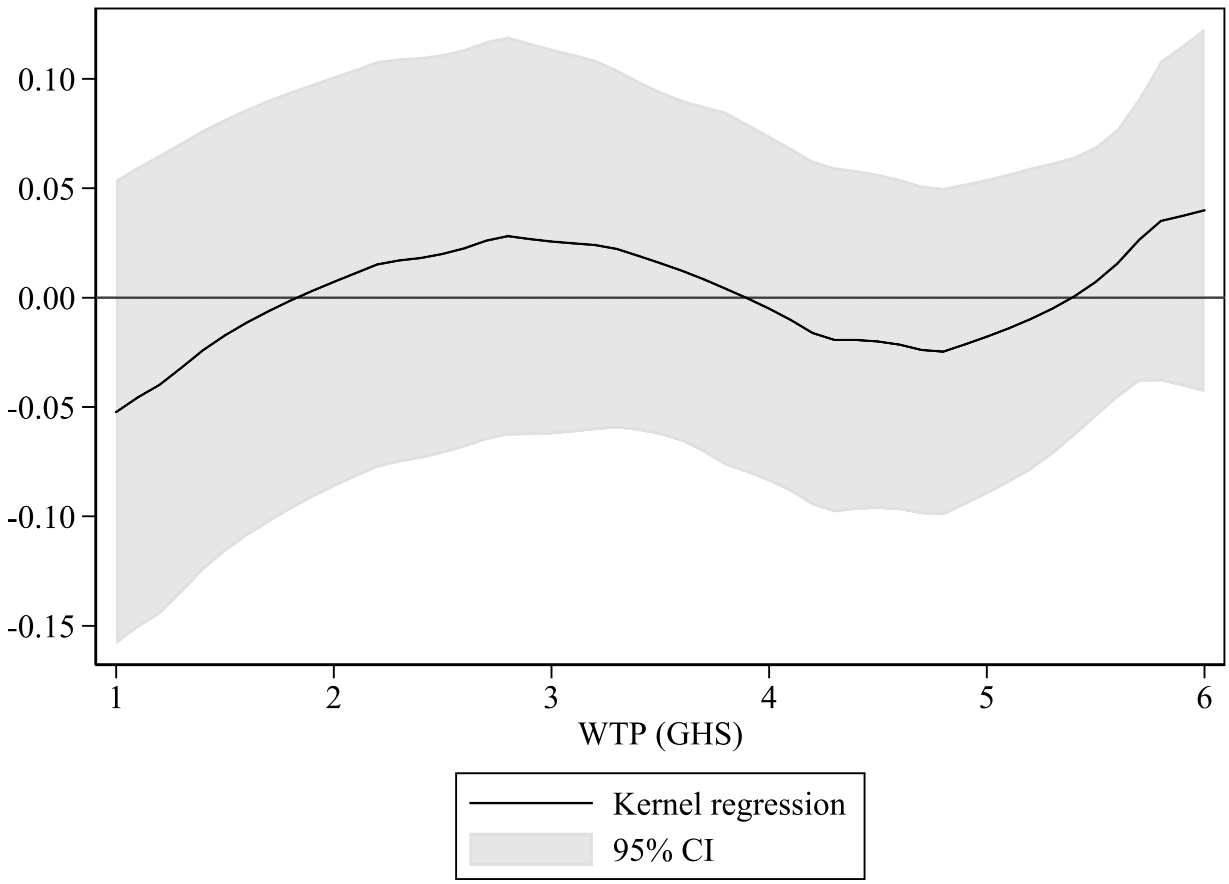

Notes: This figure plots estimated differences, with 95 percent confidence bands, in the rate of attrition from the one-year follow-up survey between BDM subjects who won the filter and subjects who did not win. The line plots estimates from kernel regressions of attrition on winning the filter, using Epanechnikov kernel with Silverman's rule-of-thumb bandwidth. Standard errors are robust to clustering at the compound (extended family) level.

Figure A6: Local Instrumental Variables and Marginal Treatment Effects

(a) Step 1: Propensity Score

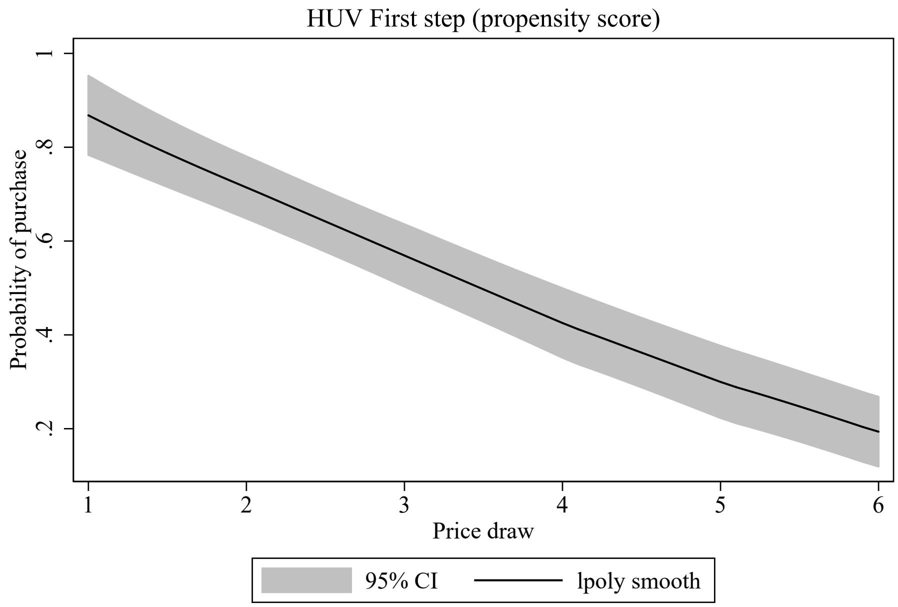

Propensity score estimated by local linear regression.

(b) Step 2: Local Instrumental Variables

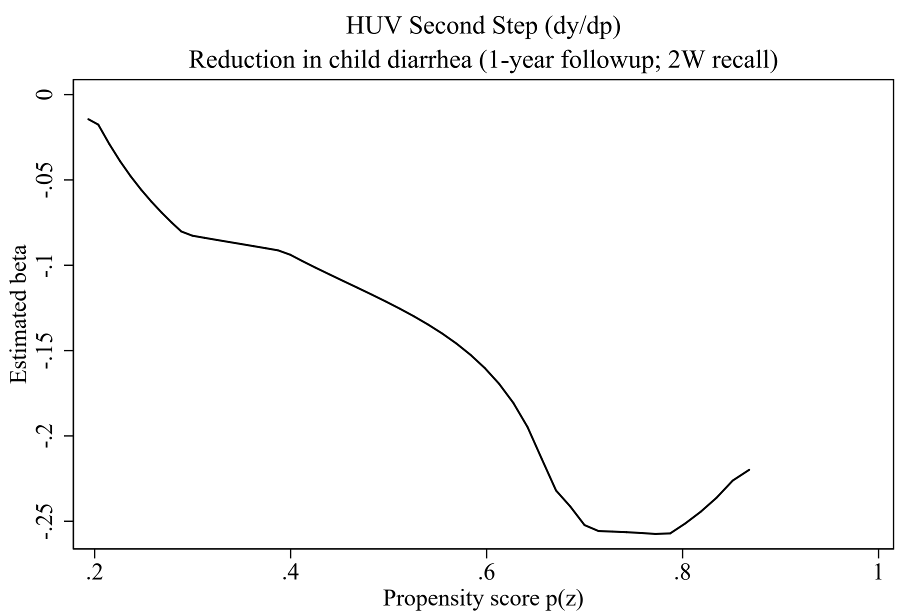

Second step (LIV): local linear regression of outcome on estimated propensity score.

(Continued next page)

Figure A6: Local Instrumental Variables (Continued)

(c) Marginal Treatment Effect

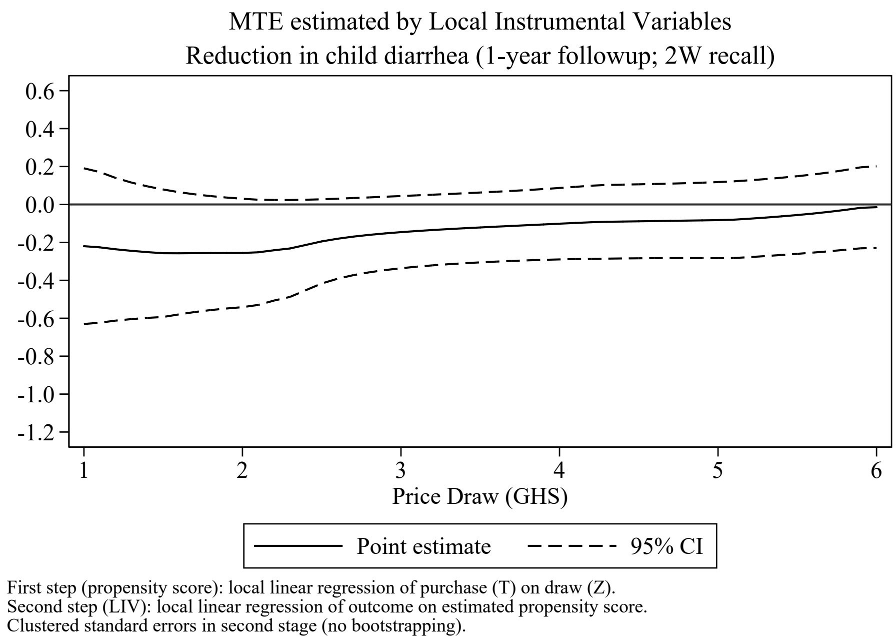

(d) Comparison of BDM-IV and LIV

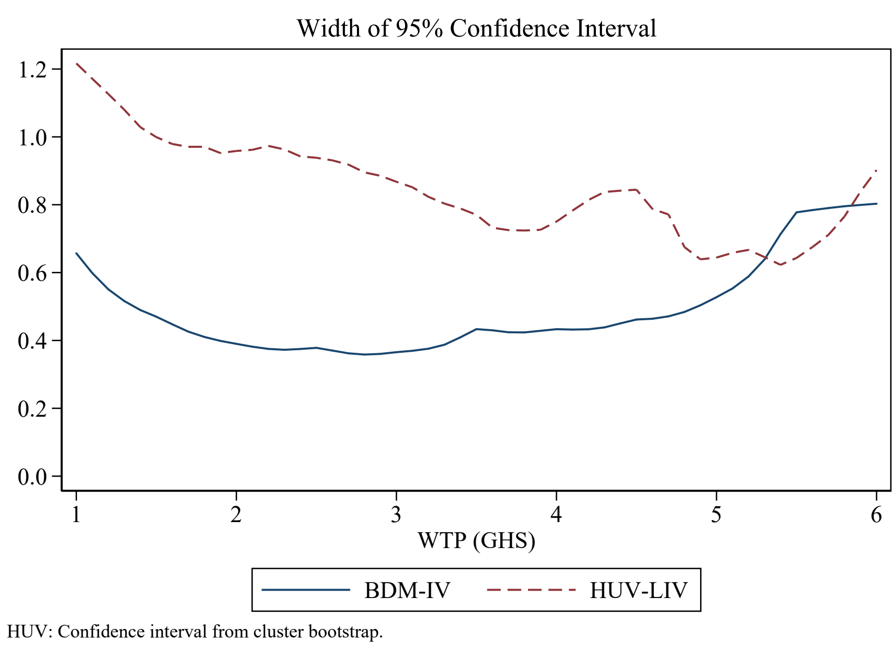

Figure A7: Health Outcomes for Long-Term Non-Users
All Subjects with Children 0 to 5

(a) Reported Cases (1 Year)

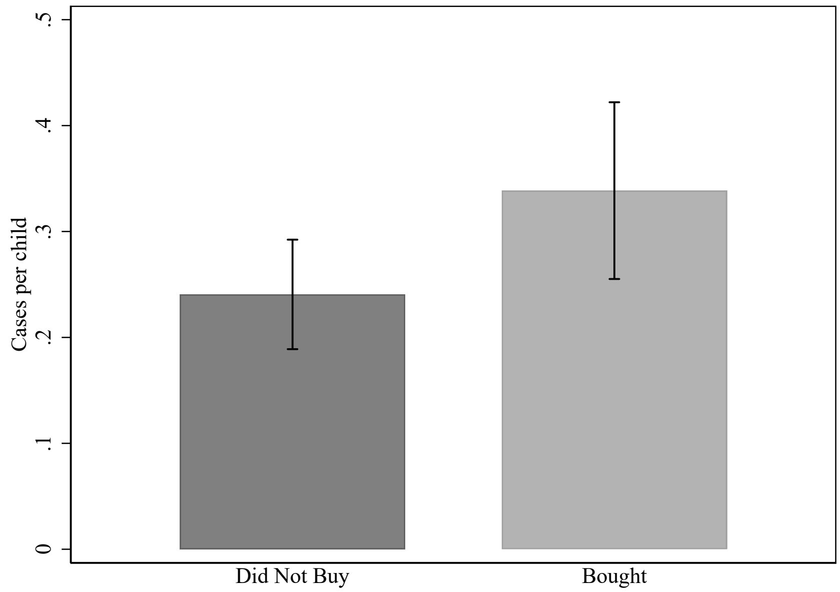

(b) Reported Cases (1 Year), By use History

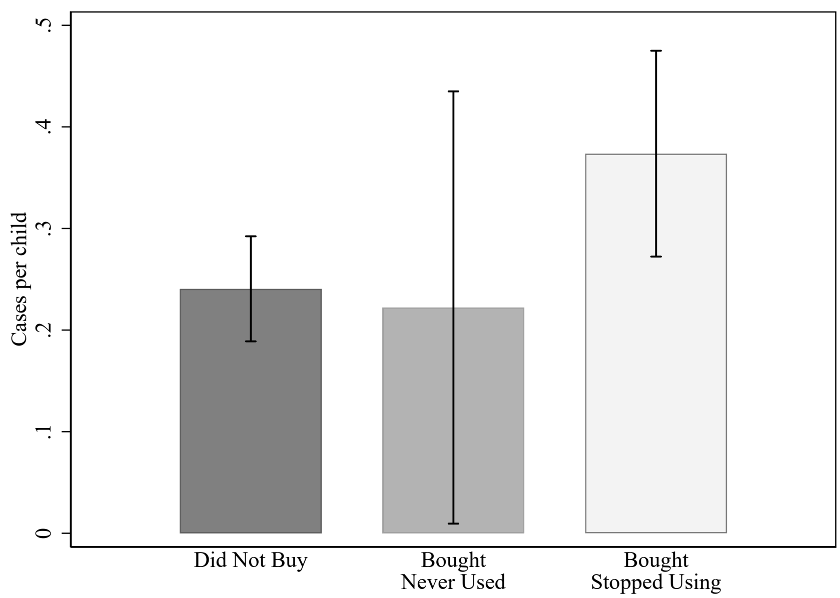

Notes: This figure displays incidence of diarrhea (lower is better) in the prior two weeks for children aged five or under among all households that are not using the filter at the one-year follow up. The second figure separates purchasers into those who were not using the filter after one month and those who were using the filter but stopped using at the one-year follow up. For all households, use is defined as an indicator for the filter being present, operational, and either containing water in the clay pot or storage vessel. Whiskers represent 95 percent confidence intervals.

Figure A8: ROC Comparison of BDM vs. TIOLI for Predicting TIOLI Purchase Behavior

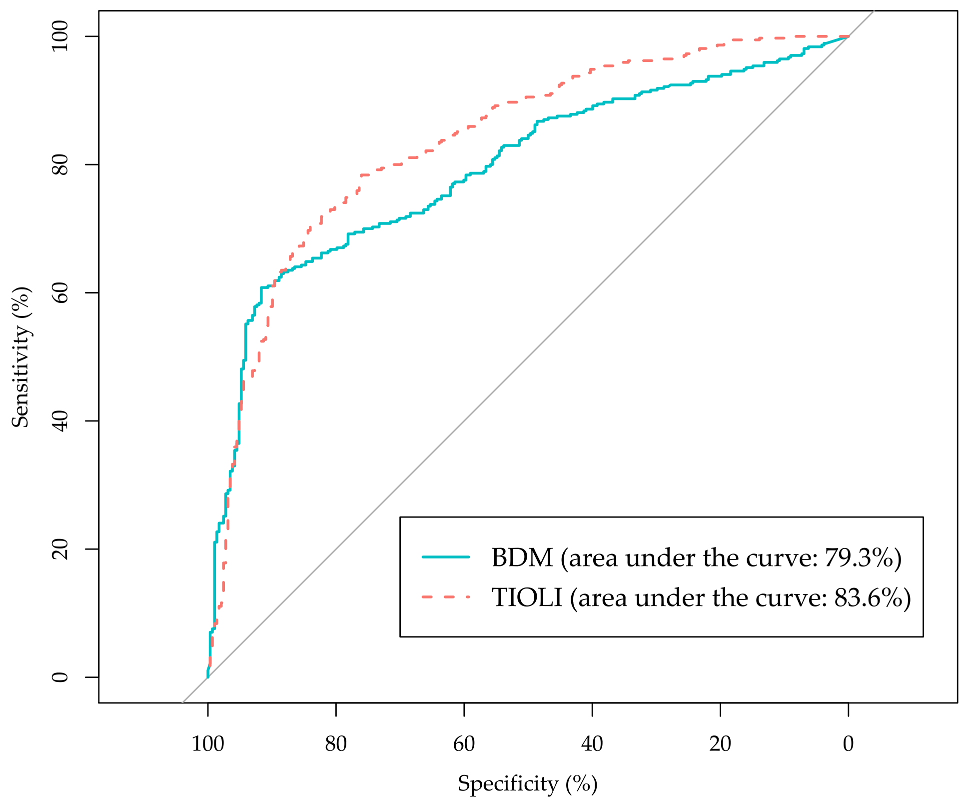

Notes: The target outcome is the TIOLI purchase decision (yes/no) in cross-validation sample. The ROC curves plot the sensitivity of each predictive model (the rate of true positives) vs. the specificity (the rate of true negatives) as we vary the threshold for predicting purchase. The 45-degree line represents the performance a model that randomly classified each observation.

Figure A9: BDM–TIOLI gap by tercile of risk aversion

(a) Tercile 1 (most risk-averse)

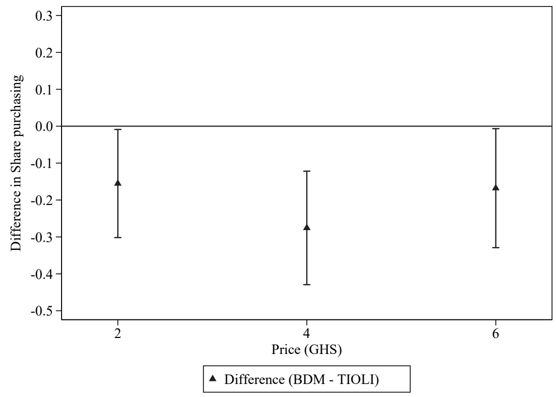

(b) Tercile 2

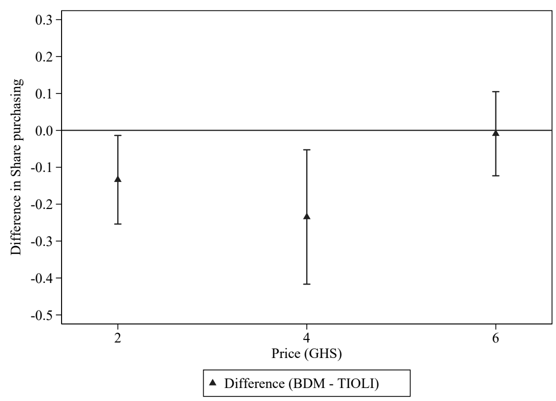

(c) Tercile 3 (least risk-averse)

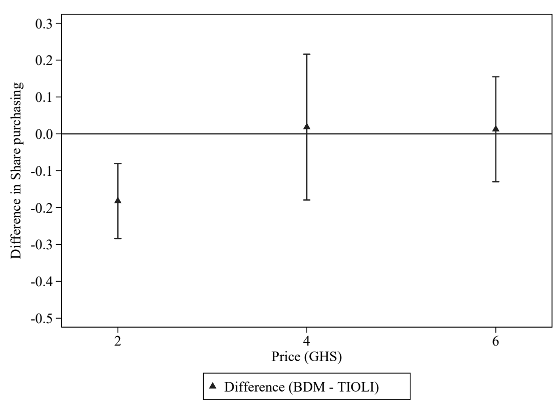

Notes: These figures plot estimated differences, with with 90 percent confidence intervals, between the share of BDM subjects and the share of TIOLI subjects agreeing to purchase at each TIOLI price (GHS 2, 4, 6), separately by tercile of risk aversion. The results here are unconditional, see Figure A10 for robustness checks with additional controls.

Figure A10: BDM–TIOLI gap by tercile of risk aversion
Robustness check: with controls, including ambiguity aversion

(a) Tercile 1 (most risk-averse)

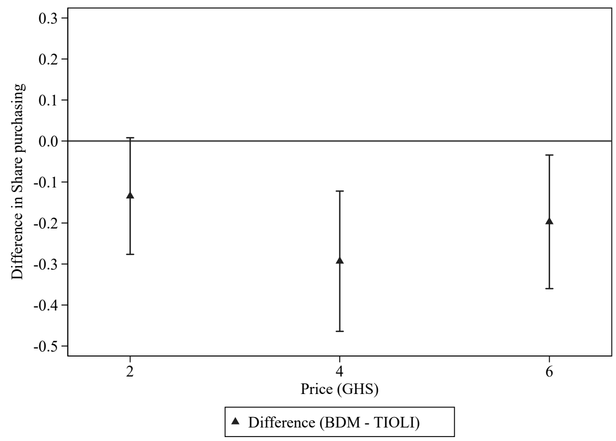

(b) Tercile 2

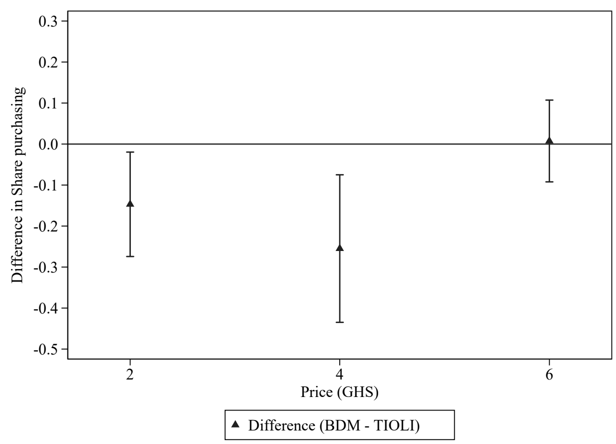

(c) Tercile 3 (least risk-averse)

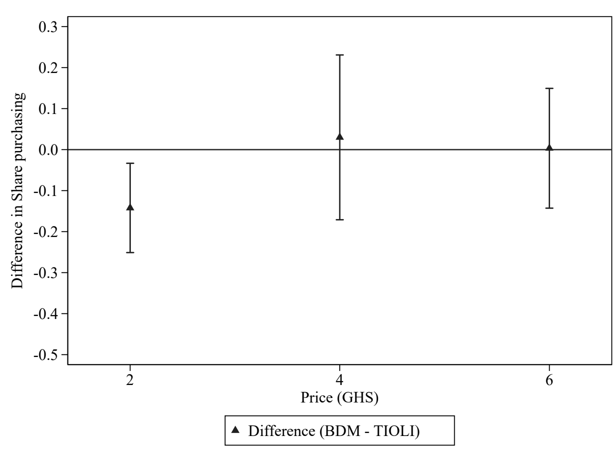

Notes: These figures plot the difference between the share of BDM subjects and the share of TIOLI subjects agreeing to purchase at each TIOLI price (GHS 2, 4, 6), separately by tercile of risk aversion. The regression includes the standard set of household controls and our measure of ambiguity aversion, described in Appendix B.

Figure A11: Heterogeneity in BDM–TIOLI gap across relevant observables

(a) Levels

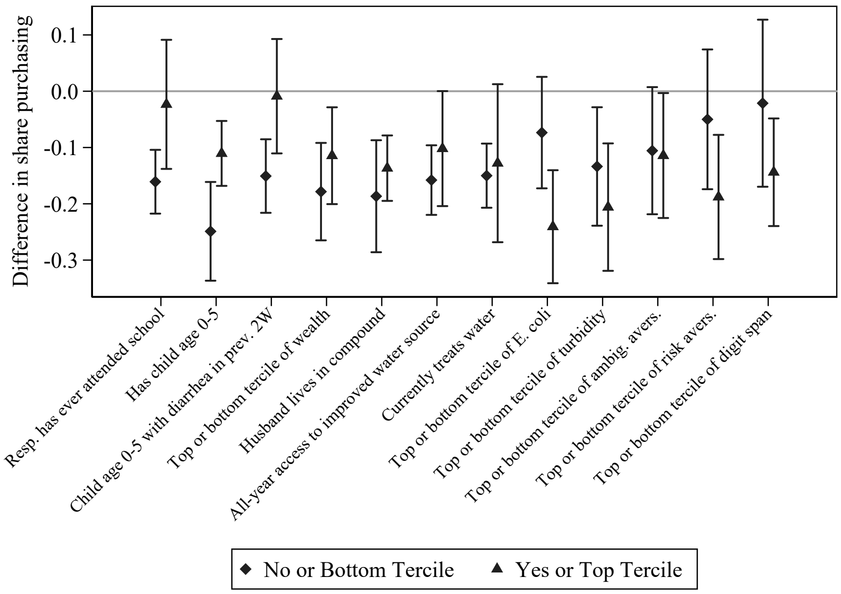

(b) Differences

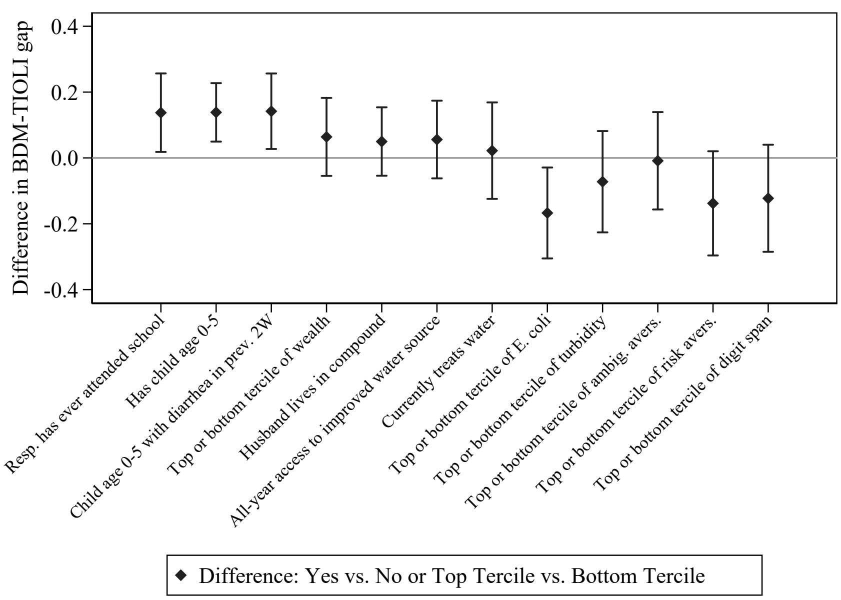

Notes: These figures compare the average BDM-TIOLI gap (percentage point difference in shares purchasing, averaged over the three TIOLI prices) in different sub-groups. For binary observables, we compare the two levels of the variable. For continuous observables, we divide the sample into terciles and compare the top and bottom terciles. The top panel shows the level of the gap for the two sub-groups; a more negative number indicates a larger BDM gap. The bottom panel shows the difference in the gap between the two groups; a positive number that the BDM-TIOLI gap is narrower among the "Yes" or "Top Tercile" subgroup than among the "No" or "Bottom Tercile" subgroup. "Child age 0-5 with diarrhea in prev. 2W" is limited to respondents with one or more children age 0-5. For "Top or bottom tercile of E. coli" and "Top or bottom tercile of turbidity" the top tercile category refers to the highest levels of E. coli and highest levels of turbidity, respectively, i.e., the poorest quality water.

Figure A12: Comparison of BDM Sub-treatments

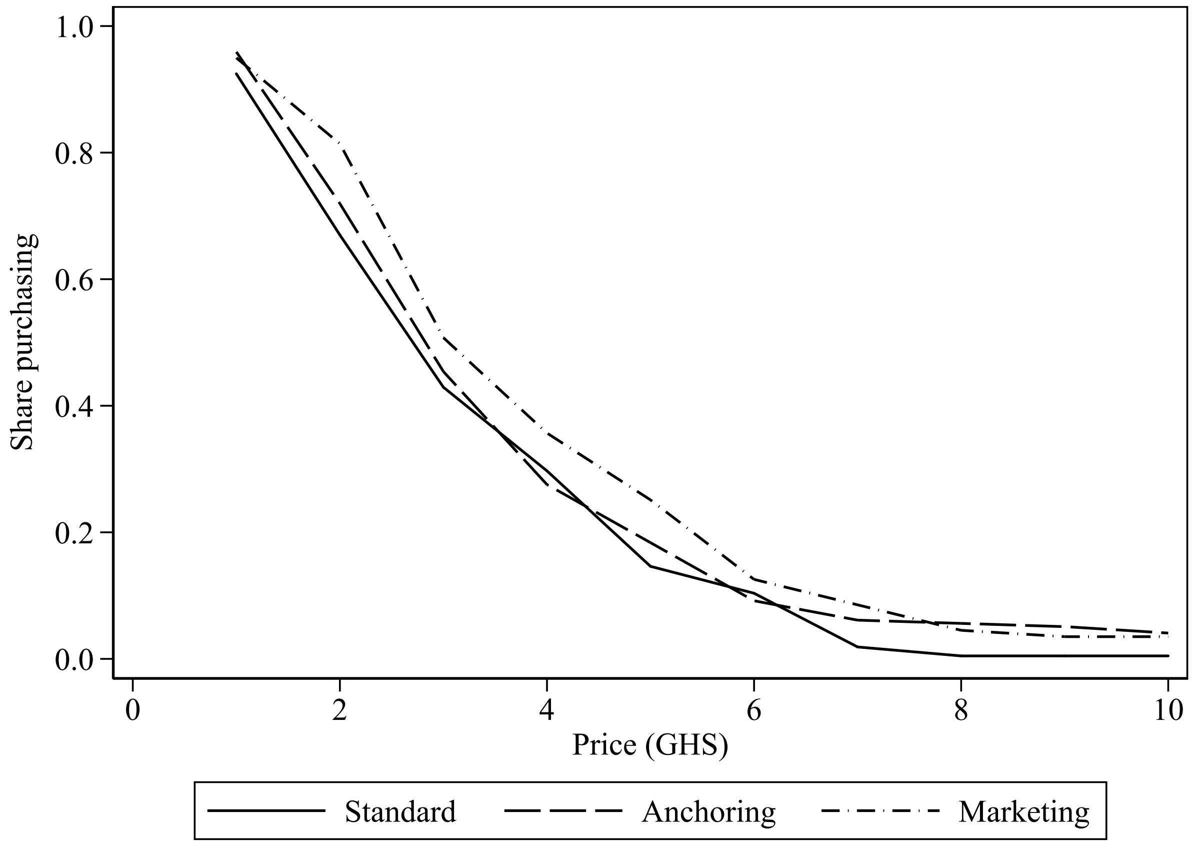

Notes: The standard, anchoring and marketing treatments are described in detail in the text. 607 observations total, of which 212 are standard BDM, 199 are marketing BDM and 196 are anchoring BDM. All treatments were randomized at the compound (extended family) level.

Figure A13: Comparison of TIOLI Sub-treatments

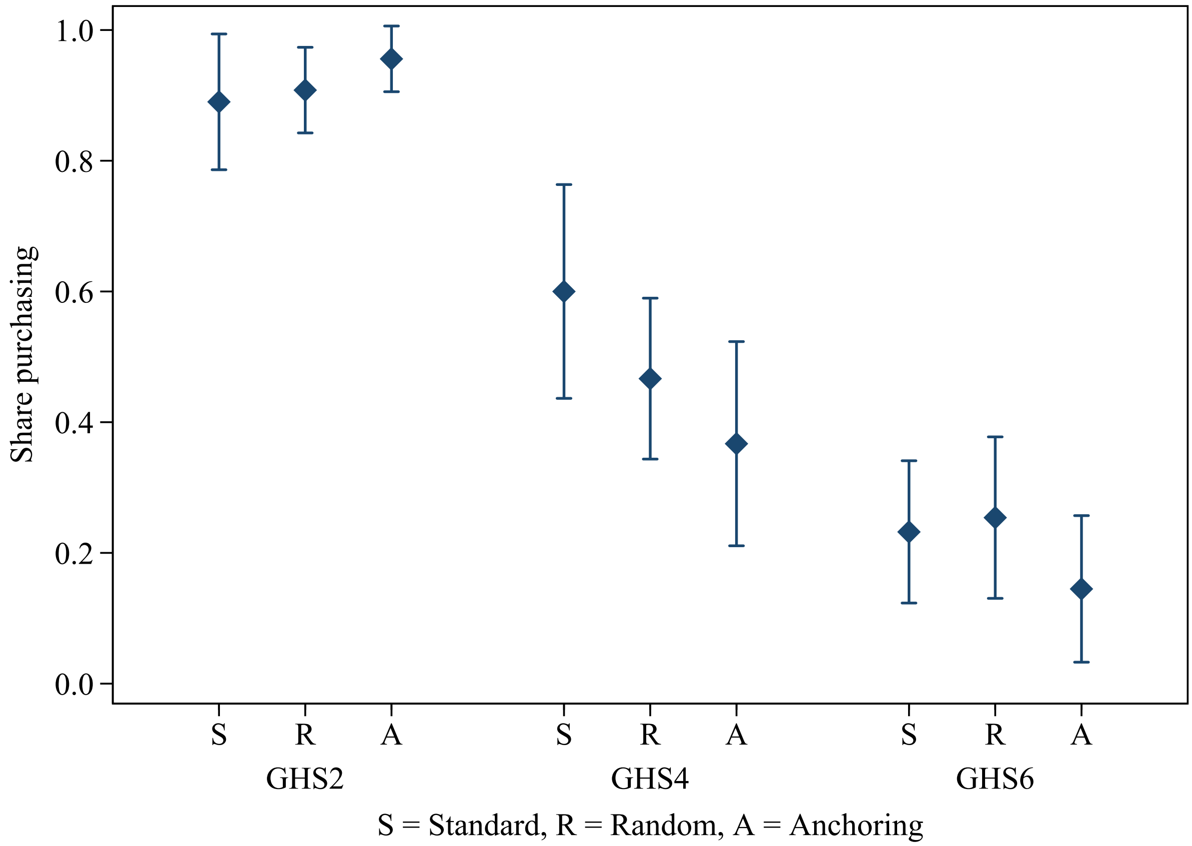

Notes: This graph plots demand for the filter at each take-it-or-leave-it price, for each TIOLI sub-treatment. The random, anchoring and standard sub-treatments are described in detail in the text. Each treatment was randomized at the compound level. For the standard and anchoring TIOLI treatments, the price was also randomized at the compound level. For the random TIOLI treatment, the price was drawn by individual respondents. 658 observations, of which: standard 217 (GHS2 91, GHS4 70, GHS6 56); random 225 (GHS2 87, GHS4 75, GHS6 63); anchoring 216 (GHS2 68, GHS4 79, GHS6 69).