IGC International Growth Centre

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

James Berry[a], Greg Fischer[b] and Raymond Guiteras[c]

[a] Cornell University (contact: jimberry@cornell.edu), [b] London School of Economics (contact: g.fischer@lse.ac.uk), [c]University of Maryland (contact: quiteras@econ.umd.edu)

Working Paper 12/0188
March 2012

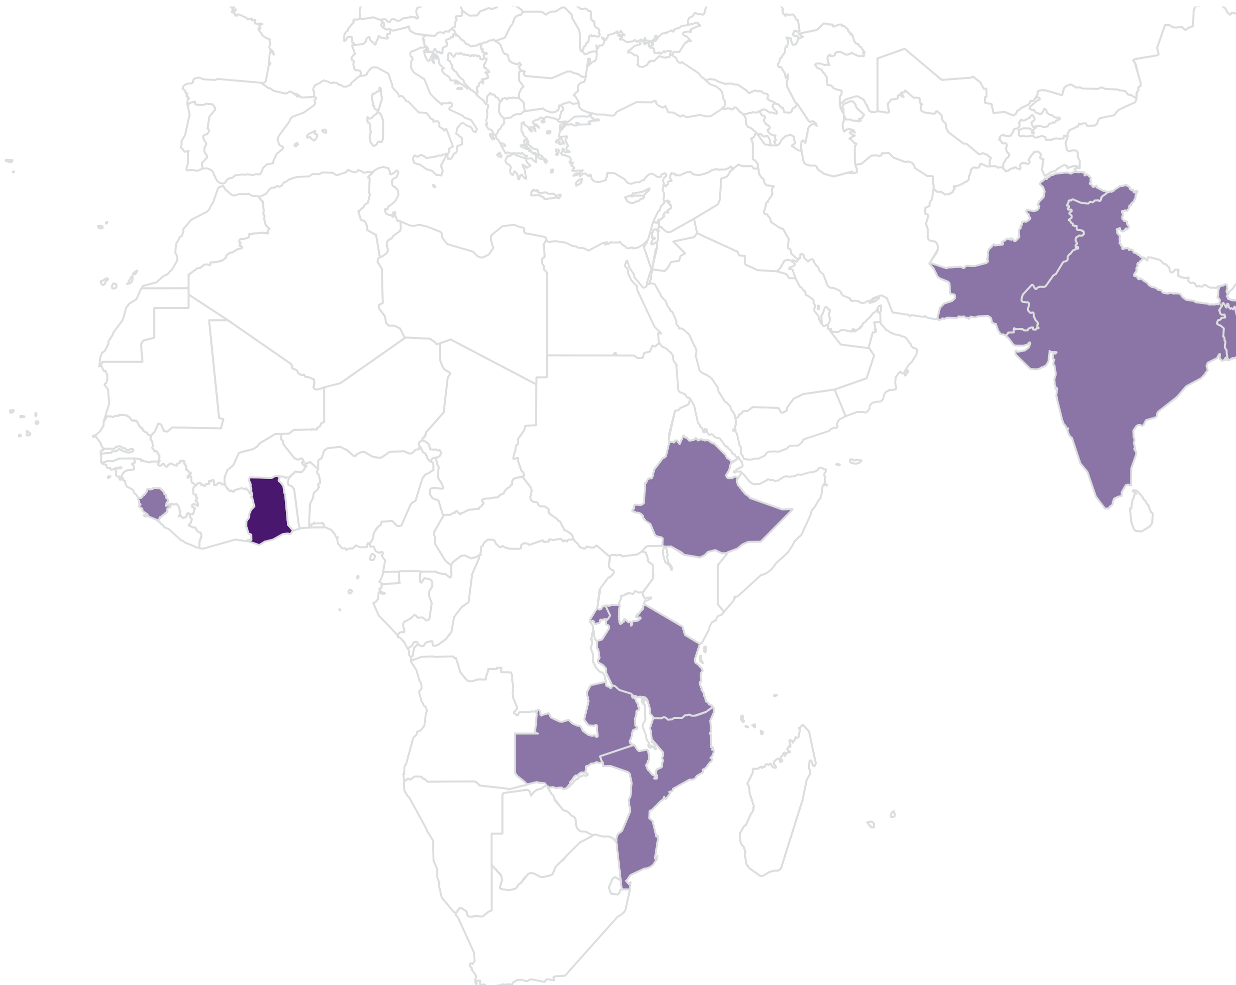

International Growth Centre
London School of Economics and Political Science
4th Floor, Tower Two
Houghton Street
London WC2A 2AE
United Kingdom

For media or communications enquiries, please contact Mazida Khatun
mazida.khatun@theigc.org

Directed and Organised by

LSE

UNIVERSITY OF OXFORD

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

James Berry, Cornell University

Greg Fischer, London School of Economics

Raymond Guiteras*†, University of Maryland

December 2011

Abstract

We utilize the Becker-Degroot-Marschak (1964) mechanism to estimate the willingness to pay for clean drinking water technology in northern Ghana. Under certain conditions, the BDM mechanism has attractive properties for empirical research, allowing us to directly estimate demand, compute heterogeneous treatment effects, and study the direct and screening effect of prices with minor modifications to a standard field experiment setting. We demonstrate the implementation of BDM along these three dimensions, compare it to the standard take-it-or-leave it method for eliciting willingness-to-pay, and discuss practical issues for implementing the mechanism in true field settings.

JEL Classifications: C93, D12, D82, L11, L31

*jimberry@cornell.edu, g.fischer@lse.ac.uk, guiteras@econ.umd.edu

†We thank Alex Nisichenko and Abdul-Salaam Sayuti for superb work in coordinating the field activities in Ghana, as well as Kelly Bidwell and Jessica Kiessel at Innovations for Poverty Action, and Mary-Kay Jackson and Susan Murcott at Pure Home Water. We have received helpful comments and suggestions from Douglas Almond, Abhijit Banerjee, Sylvain Chassang, Vivian Hoffman, Leigh Linden, Dean Karlan, Sheila Olmstead, Gerard Padro-i-Miquel, Giordano Palloni, Jesse Shapiro, Chris Udry, Sergio Urzua and seminar participants at Cambridge University, Columbia University, Cornell University, HECER, the London School of Economics, Queens University Belfast, Resources for the Future, Tufts University, Universitat Pompeu Fabra, the University of California Santa Barbara, the University of Maryland and the University of Virginia. S. Aqil Esmail, Kate Loose and Ben Zou provided excellent research assistance. This research was supported by the Improving Institutions for Pro-poor Growth research consortium, the International Growth Centre and by the University of Maryland Population Research Center.

1 Introduction

Economists, firms, and policy-makers are often interested in knowing how much an individual is willing to pay for some item. If an individual believes that his answer to the question "How much are you willing to pay?" will affect the actual price, there will be an incentive to answer strategically. Economists have considered a range of techniques to elicit a truthful answer, including revealed preference methods such as take-it-or-leave-it offers at randomized prices, Vickrey auctions, nth-price auctions and stated preference methods such as contingent valuation and conjoint analysis. Understanding willingness to pay is of particular importance in analyzing pricing policy for health products in the developing world. In this context, precise data on willingness to pay (WTP) can be used to study the screening effect of prices, that is, whether those who buy at a particular price are more likely to use a product, as well as their direct causal effect: whether higher prices cause higher use through psychological effects (Cohen and Dupas 2010; Ashraf, Berry and Shapiro 2010). In addition, precise data on willingness to pay can be used to identify heterogeneous treatment effects, as in Heckman and Vytlacil (2005) and Chassang, Padro i Miquel and Snowberg (2010). To date, however, research has typically relied on purchase decisions of products at a single fixed price. These data establish a ceiling or a floor on each individual's WTP, rather than the exact value.

The Becker-DeGroot-Marschak (BDM) mechanism (Becker, DeGroot, and Marschak 1964) has been widely utilized in experimental settings to induce revelation of an individual's willingness to pay. Under the BDM mechanism, an individual states her maximum willingness to pay, b, to receive an object. Then a random price, p, is drawn from a distribution. If p is less than or equal to b, then the individual receives the object and pays the random price p. If p is greater than b, then the individual pays nothing and receives nothing. Bidding one's true maximum willingness to pay is a dominant strategy for expected utility maximizers.

We implement and assess BDM in the context of studying WTP for and the impact of the *Kosim* filter, a ceramic water filter, among households in rural northern Ghana. While there is a substantial literature dealing with the implementation and behavior of BDM in university economics labs,[1] this is to our knowledge the first evaluation of BDM involving a non-trivial good, and additionally the first in a developing country context. We randomly assign respondents to be offered a water filter using either BDM or a take-it-or-leave-it (TIOLI) offer representing a more typical market transaction. We find strong evidence that the willingness to pay implied by BDM is consistently less than that implied by the TIOLI mechanism.

There are a number of potential reasons why the BDM mechanism may yield lower willingness to pay than the TIOLI mechanism. Our experiment focuses on two of these potential reasons. First, people may feel that they can influence the future price of the item by bidding low. Second, the TIOLI offer may anchor people to a valuation that is higher than what they would bid in the BDM mechanism. People may either feel that the stated price carries some information about a product, or that the stated price may make them focus on that price and resolve that their valuation is at least that much.

In order to test these potential channels, we varied the form of the BDM and TIOLI offers. We varied the extent to which individuals felt that they could influence the future price of the filter by either (1) telling them that their responses would be used to price the product in the future, or (2) telling them that the price of the filter was fixed. In order to limit the amount of anchoring in the TOILI mechanism, in a subgroup of TIOLI subjects we had the subject draw the randomized price herself. Further, the BDM treatment in which the individual was told that the filter was being sold elsewhere at a fixed price was intended to minimize anchoring effects.

Our empirical results rule out these explanations for the difference between BDM and TIOLI. We therefore propose several alternative hypotheses for these differences, including risk, loss and ambiguity aversion. We then discuss future research to examine these effects.

We also study the health impact of the filter, as measured by reductions in cases of diarrhea in children age 5 and under. In particular, we study whether these health impacts vary by willingness to pay, as revealed by the participant's bid. Existence of such heterogenous treatment effects could

[1] See, e.g., Noussair et al. (2004), Smith (1982), Irwin et al. (1998) and Kellar, Segal, and Wang (1993)

shed light on the efficacy of the price mechanism in the allocation of health goods. We demonstrate how BDM can allow researchers to estimate local marginal treatment effects in a variety of field settings. In our particular context, we find some evidence for an inverted U-shaped pattern for these effects, although these estimates are not very precise. Importantly, we find evidence of a similar pattern for filter usage. This raises the possibility that the standard empirical technique of using price as an instrument for usage may violate the monotonicity assumption necessary for interpreting local average treatment effects (Imbens & Angrist, 1994).

Finally, our design allows us to study the screening and causal effects of prices in several ways. Because the BDM mechanism both provides an incentive-compatible measure of willingness to pay and because it randomizes the price paid conditional on willingness to pay, we are able to measure the screening and sunk-cost effects of prices, as in Ashraf, Berry and Shapiro (2010) and Cohen and Dupas (2010). We do not find consistent evidence of either effect on usage behaviors two weeks after the initial distribution of the filter. We also find no relationship between willingness to pay and use in our TIOLI treatments, although the TIOLI design does not separate causal effects from our estimation.

The rest of the paper proceeds as follows. Section 2 describes the background of the BDM mechanism. Section 3 explains the experimental setting and implementation. Section 4 describes the results of our comparison of willingness to pay under BDM and TIOLI, and outlines a model of BDM bidding behavior consistent with these results. Section 5 explores the relationship between willingness to pay and baseline characteristics and health behaviors. Section 6 explores the screening and causal effects of prices. Section 7 illustrates the use of BDM to estimate the relationship between willingness to pay and treatment effects. Section 8 concludes and presents a roadmap for future work.

2 Background

2.1 Uses of BDM

There are three major appealing qualities of BDM: precision in demand estimation; quasi-experimental variation in treatment; and random variation in price paid, conditional on willingness to pay. In this subsection, we discuss each in turn.

First, BDM provides an exact number for WTP, with the precision of this number limited only by the desired granularity of the researcher. In principle, it would be possible to measure exact maximum willingness to pay down to the smallest available denomination. In practice, we limited our prompts to units of GHS 0.50,[2] although nothing prevented respondents from bidding, for example, GHS 1.2 (as one respondent did). Regardless, the technique provides more precision than a randomized take-it-or-leave-it offer, which provides only a bound. For example, if a respondent accepted a TIOLI offer of GHS 4, we could only conclude that her WTP was at least GHS 4. Similarly, if she rejected an offer of GHS 6, we could only conclude that her WTP was less than GHS 6. Under BDM, though, we would obtain a number, e.g., 3, 5, 5.5, 8, etc. This precision aids in the estimation of demand, e.g., correlation with observables.

Second, BDM provides a quasi-experiment that allows for instrumental variables estimation of treatment effects. Because the respondent's draw is purely random and correlated with take-up, one can estimate the treatment effect of the good in question by instrumenting for take-up with the draw. However, this is also true of take-it-or-leave it with a randomized price – the researcher can instrument for take-up with the randomly assigned price. What BDM allows that TIOLI does not is *separate* IV estimates of treatment effects by levels of willingness to pay.[3] The intuition is as follows: because BDM reveals the respondent's WTP, the researcher could separate the sample into two groups, high WTP (i.e., above the median) and low WTP (i.e., below the median). The researcher could then analyze the two samples separately and obtain separate

[2] At the time of the study, the exchange rate was approximately GHS 1.5 per US$ 1.

[3] This point is emphasized by Chassang, Padro-i-Miquel and Snowberg (2010), who describe BDM as an example of a "selective trial."

estimates of treatment effects for the high WTP group and the low WTP group. More generally, the researcher could interact the BDM draw with WTP and powers of WTP and obtain a distribution of estimated treatment effects across different levels of WTP. This allows the researcher to address a fundamental economic question: are those who stand to benefit the most also those who are willing to pay the most?

Third, BDM provides random variation in price paid among those with equal WTP, which allows the researcher to estimate, for example, the causal effect of price paid on use, as in as in Ashraf, Berry and Shapiro (2010) and Cohen and Dupas (2010). That is, BDM allows the following experiment: two subjects, both willing to pay GHS 6 for the filter, both obtain the filter, but one pays 6 GHS and the other pays 2 GHS, with the difference in price paid randomly assigned. The researcher could then study whether the amount paid (or, if zero prices are included in the BDM distribution, the act of paying a positive amount) influences the recipient's behavior with respect to the product. BDM is particularly useful for this sort of two-stage pricing, because, unlike other mechanisms that use a surprise discount, the random variation in price paid does not need to be kept a secret for incentive compatibility. This enables its use in contexts, such as ours, where there is close communication among subjects and the possibility of a surprise discount could never be sustained. Hoffmann, Barrett and Just (2009) use BDM to measure the willingness to pay / willingness to accept compensation gap for bednets in Uganda, and find that subjects are less willing to sell a bednet received for free than purchase one at the same price. Coverage of young children is higher when nets are received for free, due to this endowment effect and possibly a psychological effect of free receipt (Hoffmann 2009).

2.2 Criticisms of BDM

Several studies have examined theoretical conditions under which BDM is not incentive compatible. Karni and Safra (1987) find that BDM is not incentive compatible over lotteries (or uncertainty over the good's value) when individuals do not maximize expected utility. In more recent work, Horowitz (2006a) shows that BDM is not incentive compatible for non-expected utility maximizers even when there is no uncertainty over the good's value.

2.3 Tests of BDM

The few empirical tests of BDM are ably described by Horowitz (2006b). We provide just a brief summary here, and update with one paper written since Horowitz conducted his survey. We note that none of these are tests when the object of interest is of great importance, none are in developing countries, and few are conducted in field settings.

Three papers compare BDM bids to bids in a Vickery auction (Rutstram 1998; Shogren et al. 2001; Noussair, Robin, and Ruffieux 2004). In all three cases, BDM bids are below Vickery.

Two studies assess the rationality of behavior under BDM by comparing bids under different distributions of prices. Since this should be irrelevant, a test of rationality is whether the bids are affected by the distribution. In both cases, they reject rationality (Bohm, Lindén, and Sonnegård 1997; Mazar, Koszegi, and Ariely 2010).

3 Experimental Design

3.1 Point of Use Water Treatment in Rural Northern Ghana

Lack of access to clean water is one of the most significant threats to health and welfare in the developing world, particularly rural Africa. Forty percent of Africans—and 54% of rural Africans—lack access to improved sources of drinking water. This has serious health consequences: diarrheal disease causes nearly 1.8 million deaths worldwide each year, and is responsible for 17% of deaths of African children under five years of age. Poor water quality also has large impacts on the health of the living, contributing to other debilitating diseases such as malaria, schistosomiasis, trachoma and worms (WHO 2011; WHO 2005).

Our study tests willingness to pay for a particular household water treatment: the *Kosim* filter, a ceramic filter sold by Pure Home Water, a Ghana-based non-profit organization (see Figure 1). This simple product is highly effective at improving water quality and is appropriate for the region, since it removes particulates and pathogens from water without the use of chemicals or electricity.). This product provides an excellent platform studying willingness to pay for health products because (1) similar technologies are sold through social marketing organizations, and (2) Pure Home Water itself is interested in learning about demand for its own product and is therefore a responsive partner in the evaluation.

3.2 Village Selection

Our study is based on sales of the *Kosim* filter to households in 15 villages in northern Ghana. Villages were selected according to several criteria. First, we targeted villages in which Pure Home Water had not conducted sales in that village previously, nor had there been a giveaway by UNICEF or any other organization. Second, we selected villages within 2 hours' travel time from our headquarters in Tamale. Third, we had to receive the initial cooperation the village chief and health liaison in conducting the study in the village. Forth, we selected villages which contained between 30 and 60 compounds (extended-family households). Finally, we chose villages in which residents did not already have easy access to clean water. In particular, we excluded villages with access to piped water and villages that were saturated with functioning bore-holes. These criteria produced a list of 24 villages, from which we visited 15 in random order.[4]

3.3 Unit of observation

Because the focus of most clean water interventions is child health, our unit of observation is defined as any person with primary caregiving responsibilities for children age 12 or under. In the vast majority of cases, this was the child's mother, although there were some instances of wid-

[4]Our initial assessment produced a list of 22 villages. After surveying had begun in the 4th village, we determined that 1) our initial estimates of village size were too small, and 2) we had the capacity to survey larger villages. Therefore, we added 10 villages with between 40 and 60 compounds to our list, dropped 7 villages from the original list with less than 40 compounds and 1 village which our initial visit revealed to have 98 compounds. We then re-randomized the village order. Since our treatments were randomized within villages and stratified by compound size, we do not view this as a problem for our study.

ower fathers, grandparents, aunts and uncles and foster parents caring for children whose parents were absent. For simplicity of prose, we will refer to these caretakers as “respondents” from this point forward, to distinguish them from other village residents potentially affected by the study or otherwise appearing in our data, who we will refer to generally as “subjects.”

Most subjects live in extended patrilineal family compounds, which are small clusters of individual huts, usually enclosed by a wall. Many resources are shared within the compound, although in most cases each mother is responsible for providing water for her own husband and children. As described below, the treatments were randomized at the compound level and all inference is robust to clustering at the compound level.

3.4 Sales Process

3.4.1 Preliminary activities & household survey

INITIAL VISIT AND CENSUS. For each village, we conducted an initial village meeting, during which we provided a demonstration of the filter and the two sales mechanisms. Two surveyors performed a mock version of both BDM and TIOLI for a token item, such as chocolates or a bar of soap. The surveyors also practiced the sales mechanisms with volunteers from the attendees, again for a token item. We informed villagers that a filter would be installed at the health liaison’s home and encouraged them to visit the liaison to see the filter working, taste the water and ask questions. We instructed the attendees that we would visit their households in approximately two weeks to offer them an opportunity to purchase the filter via one of the mechanisms we performed. Attendees were encouraged to discuss with their families what they were willing to pay for the filter. The two-week interim period was chosen to allow families time to try the filter and determine their willingness to pay, as well as to obtain the desired level of cash, either through occasional work or by selling assets such as grain.

After the meeting, we conducted a comprehensive census of all residents of the village. With this information, we were able to identify the study subjects as defined above and perform random assignment of the subsequent treatments.

Water quality testing and health education treatment. In the roughly two-week period between the village presentation and the sale, we visited each household to remind them of the upcoming sales visit and to answer any questions they had about the treatment. During this reminder visit, we took a 100 ml baseline sample of their water for testing in the lab. In addition, we conducted one of two health education treatments in randomly selected households. The first treatment is a general message describing the link between untreated water and health and explaining how the filter helps prevent diseases such as diarrhea. This treatment discusses child health but it is not a particular emphasis. In contrast, the second treatment emphasizes child health. The substantive elements of the treatment are largely similar but the emphasis is placed on the dangers to children of untreated water and the potential benefits of the filter to children. The impacts of these treatments will be analyzed in a separate paper.

Household survey. Roughly one week after the reminder visit (and, for treatment households, the education treatment), we conducted a survey and sales visit with each respondent. Respondents were compensated with a 1.00 GHS cash gift, awarded at the beginning of the survey.[5]

All respondents were administered a basic survey covering income, assets, education and health status. We also included a few specialized modules, including a module on water collection and treatment practices and a module asking about basic health knowledge, in particular asking respondents to list causes of diarrhea in children. This latter module was performed as an unprompted question to avoid suggestion bias. At the end of the survey, we conducted the sales experiment. By conducting the sale at the end of a survey on water and health, we may have primed the respondent's demand for the filter. However, it was not feasible to conduct the sale first, because respondents, and especially respondents who were not able to purchase the filter, would quickly lose interest.

Sale. Respondents were randomly assigned to either a BDM or TIOLI sales treatment. Within

[5]This was awarded in small denomination coins to ensure that respondents could submit reasonably fine-scale bids in the practice WTP game described below. It is possible that a cash gift influenced willingness to pay for the filter by inducing goodwill toward the surveyor. However, because of the length of the survey there was always at least 30 minutes between the gift and the sales offer, which would ameliorate any "house money" effect. Also, we would not expect the gift to affect WTP differentially between the two mechanisms.

each of the two broad categories described above, there are three sub-treatments designed to disentangle potential reasons for different WTP responses in the two treatments. We describe first the basic BDM and TIOLI treatments, then the refinements and the rationale for them. Sales treatments were stratified by number of respondents in the compound.

Regardless of treatment, the scripts for the sales were designed to be as similar as possible across treatments. Each sale began with a practice round for a token item, in most cases a bar of soap with retail value approximately 1.00 GHS. The respondent was then given the opportunity to purchase the soap using the mechanism corresponding to her treatment category. After the practice round was complete, the respondent was given the opportunity to purchase the *Kosim* filter using the same mechanism.

3.4.2 BDM

First, the surveyor reads a brief description of the BDM procedure. We emphasize that the respondent will have only one chance to play, cannot change her bid after drawing from the cup, and must be able to pay that day. The surveyor then plays a practice round for the bar of soap. The respondent is asked for her maximum willingness to pay for the bar of soap. The surveyor reminds her that if she draws slightly more than her bid, she will not be able to purchase the soap. She is then allowed to adjust her bid. This process repeats until the final bid is established. Before she draws, the surveyor then reviews various hypothetical outcomes to test her understanding. Once the final bid is established, the subject draws a ball from the cup and purchases or does not purchase the soap. The balls in the cup are distributed in increments of 5 pesewas (0.05 GHS), with amounts 0,. . . , 45 receiving double weight (i.e., there are two balls marked with each of these numbers, versus only one ball for each of 50,. . . ,100).

The procedure for the filter is similar. As with the practice round, the respondent is reminded that she will not be allowed to purchase the filter if she draws a price slightly higher than her bid, and is then allowed to adjust her bid upwards if she wishes. The distribution of prices is 0, 0.5, 1, 1.5, 2, 2.5, 3, 3.5, 4, 5, 6, 7, 8, 9, 10, 11, 12 in equal proportions. At the completion of the game, the respondent, if successful, pays for the filter and receives a receipt that can be redeemed for a filter at the village's health liaison's home.[6]

We do not require respondents to present the amount of cash they are willing to bid before the draw is made. Rather, we permit the household to gather the money by the end of that day. Before the draw is made, we ask multiple times whether the respondent will have access to the funds. We do this to maintain realism: households routinely make small loans to each other for purchases. In a few cases, households drawing less than their bid were unable to come up with sufficient cash by the end of the day and as a result could not purchase the filter. This could be interpreted in one of two ways. The first is that they did not understand the sales mechanism and were trying to game it in some way. This would call into question whether their bid was a true reflection of their willingness to pay. The second is that they truly were willing to pay the amount they bid, and anticipated having access to the necessary resources through a loan from a friend or family member, but for some idiosyncratic reason were not able to access funds that day. We include some checks in the procedure to try to distinguish between these two reasons, but ultimately we are not able to distinguish cleanly between them.

We also track whether a losing respondent attempts after the fact to purchase at the price drawn (i.e., above her final offer) and ask all losing respondents whether they wish they had bid more. Either of these occurrences could suggest that the respondent did not understand the game. An alternative interpretation is that the act of drawing a higher price actually alters the respondent's willingness to pay. That is, someone who thought her willingness to pay was 5 GHS and drew 5.5 GHS might experience regret at missing by only 0.5 GHS and, upon further introspection, realize that she would in fact have been willing to pay 5.5 GHS. We attempt to distinguish between these two interpretations by looking at the correlation between this ex-post regret and the gap between the bid and the draw. If regret is associated with just missing, i.e., a small gap between this bid and

[6]Perhaps the most salient difference between the procedure for the filter and for the practice round is that the filter is not physically present in front of the respondent during the bidding. We chose not to have surveyors bring the filters to compounds: first because they are bulky and could break; second because there is some instruction on assembly that should be given at the time the household receives the filter. This instruction is most efficiently provided at a central location.

the draw, this is evidence in favor of the latter interpretation.

3.4.3 TIOLI

The standard take-it-or-leave it treatment is a simple sales offer at a randomized price. We randomized at three prices, 2, 4, and 6 GHS. These prices were chosen as roughly the 25th, 50th and 75th percentiles of BDM bids in pilot exercises in similar villages. Before revealing to the respondents, we emphasized that there would be no bargaining. If they accepted the offer price, respondents were allowed until the end of the day to obtain the necessary cash. If the respondent initially agreed to the purchase but is ultimately unable to obtain the funds, we code her as not purchasing but note the attempt to purchase. As discussed below, this is important for the comparison between TIOLI and BDM.

The price is randomized at the level of the compound. It was determined that in a TIOLI setting it would not be acceptable to offer the filter to different members of the same compound at different prices. All analysis is clustered at the household level.

3.4.4 Sub-treatments

Each the broad categories of BDM and TIOLI contains three sub-treatments. These sub-treatments were included to help disentangle the possible channels through which purchasing behavior under BDM and TIOLI may differ. Pilots conducted in the summer of 2009 indicated that BDM underpredicted willingness to pay relative to TIOLI (consistent with what we find below), and each sub-treatment is designed to test the potential mechanisms. The two hypothesized mechanisms and the associated sub-treatments are described in the remainder of this section. We denote the standard treatments, described above and without modification, as the "standard BDM" and "standard TIOLI" sub-treatments.

First, we hypothesize that respondents in the BDM treatment could act strategically in hopes of influencing the future price of the filter. To examine whether this is in fact the case, we include two sub-treatments. In the first sub-treatment (the "anchoring BDM" sub-treatment), we tell the respondent that the filter is sold in shops for 20 GHS. In the second sub-treatment (the "market study BDM" sub-treatment), we tell the respondent that we are using the information from the study to help decide on the future price of the filter. If strategic bidding is indeed the case, then one would expect that respondents in the anchoring sub-treatment will have higher bids than those in the standard sub-treatment, and that those in the market study sub-treatment will have lower bids than those in the standard sub-treatment.

Second, we hypothesize that the given price in the TIOLI treatment causes respondents to anchor their own valuations to that price. If that is the case, then the "anchoring BDM" sub-treatment may induce a similar form of anchoring and thereby raise bids relative to the standard BDM sub-treatment. However, if anchoring is the key factor influencing bids rather than strategic bidding, one would not expect a difference between the "market study BDM" sub-treatment and the standard BDM sub-treatment.

We were also concerned that the anchoring induced by the BDM anchoring treatment would not increase valuations above those in the standard TIOLI sub-treatment. We therefore include an anchoring sub-treatment (the "anchoring TIOLI" sub-treatment) using the TIOLI mechanism where the respondents are informed of the retail price in town.

As a final test for anchoring, we include a TIOLI sub-treatment (the "random TIOLI" treatment) in which the randomization of the price was carried out in front of the respondent. In this sub-treatment, the respondent drew a price (2, 4 or 6 GHS) from a cup and stated whether she wanted to purchase at this price. We include this sub-treatment to confirm to the respondents that the price at which the filter was offered did not serve as a signal of quality. Furthermore, if respondents are generally uncomfortable with randomness, this "random TIOLI" sub-treatment should match BDM more closely.

The following table summarizes the predictions of bidding behavior in the six sub-treatments, if each of the channels is a factor:

Channel	WTP implied by bidding/purchase behavior
Strategic Behavior	Anchoring BDM > Standard BDM Market BDM < Standard BDM
Anchoring	Anchoring BDM > Standard BDM Anchoring BDM = Anchoring TIOLI Random TIOLI < Standard TIOLI

4 Comparison of Demand under BDM and TIOLI

4.1 Sample characteristics

Table 1 displays the number of respondents in each of the 6 categories of sale treatments. Treatment was randomized at the compound level. The first two columns show the number of compounds and individuals assigned to each treatment. To the extent permitted by the preliminary census, we attempted to stratify on the number of subjects in each compound. However, there is still some residual imbalance in the number of respondents assigned to each treatment category.

The second two columns display the number of compounds and subjects for which the survey and sales offer were actually completed. We have not detected any clear differences in attrition rates by treatment, which is reasonable since most attrition occurred when respondents had travelled away or were spending nights at their agricultural plots. In either case, the respondents did not know their treatment at the time of attrition.

4.2 Comparison of BDM and TIOLI Treatments

Figure 2 displays a histogram of the BDM bids. The average bid among all 603 BDM respondents was 3.1 GHS, median 2.5 GHS.

Figure 3 displays the acceptance behavior of the TIOLI treatments at each price. Overall, 90.2% accepted at 2 GHS (N=224), 47.8% accepted at 4 GHS (N=216), and 22.4% accepted at 6 GHS (N=199). Figure 3 also overlays the acceptance frequencies with the demand curve implied by the BDM bids. As shown by the figure, at each price, the fraction of respondents in the TIOLI treatments that purchased the filter was higher than the height of the BDM demand curve at that price. In other words, Figure 3 suggests that the fraction of respondents for whom the mechanism revealed a willingness to pay of more than 2, 4, or 6 GHS was higher in the TIOLI treatments than in the BDM treatments.

To estimate the sizes of these differences via regression, we compare the BDM bids with TIOLI acceptance behavior at each price by determining whether the BDM bid implies a purchase at that price (i.e., a willingness to pay of greater than or equal to that price). For example, in order to compare a BDM bid with TIOLI purchase at a price of 2 GHS, we generate a new variable that indicates purchase if the BDM bid is greater than or equal to 2 GHS. At each price, we compare purchase behavior using the following equation:

$$buy_{icp} = \alpha_0 + \alpha_1 BDM_c + \sigma_{icp} \tag{1}$$

where buy_{icp} indicates whether person 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 $BDM_c$is an indicator for whether compound c was assigned to the BDM mechanism.

We estimate the equation separately for each TIOLI price. Note that because the BDM bid gives a point estimate of willingness to pay, an observation assigned the BDM treatment can be used for multiple prices. Therefore, each regression contains about three times as many BDM observations as TIOLI observations.

The regression results are presented in Table 2. The difference between the two mechanisms is significant at the 5% 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-value = 0.18), and there is no such pattern in relative (percentage) differences

4.3 Comparison of BDM and TIOLI Sub-treatments

We now turn to the differences between the BDM and TIOLI sub-treatments. Starting with the three BDM treatments, Figure 4 displays the BDM demand curves for each of the three BDM sub-treatments. The graph shows slightly higher demand for the market study treatment relative to the standard treatment. The anchoring treatment yields a similar demand curve to the standard treatment.

Table 3 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. In both tests, the distribution under the marketing treatment is significantly different from the standard treatment at the 10% level, while neither test rejects the null that the anchoring and standard treatments are the same.[7]

Turning to the three TIOLI sub-treatments, Figure 5 displays the percentage of respondents who accepted in each treatment and at each price. Our hypothesis was that demand would be lowest for the random TIOLI treatment and highest for anchoring TIOLI, with standard TIOLI in between. However, over all three prices, there is no clear pattern. Demand under the random treatment is lower than under the standard treatment at a price of 2, but is higher at prices of 4 and 6. Demand under the anchoring treatment is higher than under the standard treatment at prices of 4 and 6, but not at a price of 2. In all cases, the 95-percent confidence intervals are overlapping (within each price).

To compare the mean acceptance behavior across treatments, we run the following regression for each price:

$$buy_{icp} = \alpha_0 + \alpha_1 TIOL\_ANCH_c + \alpha_2 TIOLI\_RAND_c + \sigma_{icp} \tag{2}$$

[7] Cluster-robust significance levels for the distributional tests are constructed via a bootstrap percentile method, in which 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 for each bootstrap sample. By sampling compounds and assigning placebo treatments by compound, we preserve the clustering structure in the data. We repeat this for 999 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.

where $TIOL\_ANCH_c$ is a dummy variable indicating the anchoring treatment, and $TIOLI\_RAND_c$ is a dummy variable indicating the random treatment. Table 4 presents the results of these regressions. Among all six of the individual comparisons, the only significant difference is the difference between the anchoring and the standard treatments at a price of 4. However, the joint test that all three differences between the anchoring and standard treatments equals zero is rejected at the 10% level (p-value = 0.081).

In sum, the analysis of sub-treatments provides no support for the hypotheses put forward in Section 3. In fact, two of the relationships go in the wrong direction: the market study treatment increases demand relative to the standard BDM treatment, and the TIOLI anchoring sub-treatment seems to depress demand relative to the standard TIOLI sub-treatment. In what follows, we propose a model of BDM bidding consistent with these surprising facts, which we hope to test in future laboratory and field work.

We caution against interpreting these results as supporting a conclusion that BDM does not or cannot "work," in a broadly defined sense. First, the maintained hypothesis that TIOLI yields a true measure of willingness to pay is tenuous.[8] Second, we chose a context and designed our experiment in such a way as to provide a severe – perhaps extreme – test of BDM. For example: numeracy among our subject pool was low; superstitious beliefs about probabilities are commonplace; we did not tell subjects the distribution of prices; the good in question was unfamiliar to respondents (although they were provided an opportunity to familiarize themselves with it) and yields benefits – reduced episodes of water-related disease – that are uncertain and difficult to quantify. All of these can work against BDM. Future work will investigate which of these factors, among others, are significant and inform improved versions of BDM.

[8]For example, Jack (2010) provides evidence that subjects respond less rationally to a take-it-or-leave-it offer of a payments for environmental services contract than when they participate in an auction.

4.4 A Model of BDM Bidding

While the BDM mechanism has long been described as unconditionally incentive compatible, Karni and Safra (1987) and Horowitz (2006a) show that this feature depends critically on the assumption of expected utility maximization. When preferences do not satisfy the von Neumann-Morgenstern independence axiom, the BDM procedure is not necessarily incentive compatible. More generally, an individual's willingness to pay for a product may depend on any number of the features of the mechanism for the potential purchase.

The following model serves to compare the optimal bidding strategy under BDM to the "true" willingness-to-pay, which we define as the highest price an individual would accept in a single take-it-or-leave-it offer. Note that the concept of a "true" value is somewhat arbitrary as the structure of the take it or leave it mechanism may itself impact an individual's willingness to pay. For example, as described above, the offer price may serve as an anchor or a signal of quality.

Suppose an individual has utility $u(c,w)$ where w is the value of possessing a certain item,W, and c is the monetary value of all other consumption. The quantity of W consumed is $q \in \{0,1\}$—that is, the individual can either consume one unit of W or not—and we can normalize w to q without loss of generality. Her income is Y, and if she purchases the good she pays a price p. We can then write her utility as $u(q,Y-qp)$.

Under BDM, the individual states her maximum bid, b, for the item. The price, p, is then drawn from a distribution $F(\cdot)$. If $b \geq p$ then she buys the item at price p. Otherwise, she pays nothing and does not receive the item. With expected utility, an individual's optimal bid solves:

$$\max_b \int_0^b u(1,Y-x)\,dF(x) + u(0,Y)(1-F(b)) \tag{3}$$

Thus, the optimal bid, $b*$, satisfies $u(1,Y-b*) = u(0,Y)$, which is precisely the condition for the maximum take-it-or-leave-it offer to which she would agree. BDM elicits the "true" value of the item.

Under expected utility, this result also easily extends to the case of uncertain value to the purchase item. Assume, again without loss of generality, that w is distributed according to $G(\cdot)$on the interval $[0,1]$. The individual's optimal bid now solves:

$$\max_b \left\{ \int_0^b \int_0^1 u(w, Y-x)\, dG(w)\, dF(x) + u(0,Y)(1-F(b)) \right\} \tag{4}$$

and the optimal bid, $b*$, now satisfies $\int_0^1 u(w, Y-b*)\, dG(w) = u(0,Y)$. Again, $b*$ is also the point at which the individual would be indifferent to a take-it-or-leave it offer, and BDM elicits the items "true" value. Note, however, that this value is not necessarily the expected value of the item. In particular, if $u(c,w)$ is concave in its first argument, the optimal BDM bid and the maximum TIOLI acceptance will be identical but less than the expected value of w

As shown in Machina (1982) and applied directly in Horowitz (2006a), when preferences do not necessarily conform to expected utility, the utility function is no longer independent of the distribution of prices. However, when preferences are smooth, the individual acts as a standard utility maximizer with a "local utility function" $u(w,c;F_b,G)$, where F_b is the induced distribution of q and c based on b.

Now, the optimal bid under BDM satisfies $u(1,Y-b*;F_b,G)$. However, the maximum TIOLI acceptance price, $b*$, now satisfies $u(1,Y-b*;F_1,G) = u(0,Y;F_0,G)$, where F_1 is the degenerate distribution at t. Thus, the model predicts divergence in optimal BDM and TIOLI behavior. The size and direction of the divergence between BDM and TIOLI depends on the nature of the deviation from expected utility and the shape of underlying preferences. For example, consider a special case of Gul's (1991) disappointment aversion where individuals with the opportunity to purchase a filter under BDM expect to "get a good deal" and are disappointed if the mechanism does not generate significant savings. This will lead individuals to underbid in BDM relative to TIOLI, and the extent of this underbidding will increase with greater risk aversion. If, however, individuals are relatively more disappointed by not being able to purchase the filter when they have a chance, then they will tend to overbid in BDM. Again, increased risk aversion exacerbates this deviation.

Continuing work seeks to uncover preferences that predict the deviation between BDM and

TIOLI. We are collecting additional data on preference measures including risk aversion, loss aversion, and ambiguity aversion. The direction and magnitude of any divergence will depend on the interaction of these preference parameters and the nature of the independence assumption in expected utility. Rather than make significant assumptions about this functional form, we intend to remain agnostic and estimate reduced-form relationships between preference measures and bidding behavior in hopes of recovering either ordinal or cardinal meaning from BDM bids.

5 Baseline Characteristics and Willingness to Pay

While the analysis above presents strong evidence that BDM measures lower willingness to pay than TIOLI, it is important to explore what information each measure contains. In this section, we analyze the relationship between willingness to pay and baseline characteristics. In the next two sections, we test whether willingness to pay conveys information about a household's propensity to use the filter or its treatment effect from the filter.

The relationship between willingness to pay and baseline characteristics is important to understand how pricing targets different types of households. Previous studies have found little evidence of relationships between willingness to pay for health goods and baseline health characteristics or wealth (Ashraf, Berry, and Shapiro 2010; Cohen and Dupas 2010Ashraf et al. (2010); Cohen and Dupas (2010)). To create a parallel comparison between willingness to pay and baseline characteristics for both BDM and TIOLI, it is necessary to reduce the BDM data to simulate the data generated under a TIOLI procedure. We do this by randomly assigning each BDM observation to a synthetic price of 2, 4 or 6 and generating a simulated purchase variable based on that price and the BDM bid. For TIOLI, we simply use the actual purchase behavior at the randomly assigned price. We model the relationship between willingness to pay and baseline characteristics and behaviors as follows:

$$WTP_1 = \alpha_0 + X_i^{'}\beta + \sigma_i \tag{5}$$

where X_i is a vector of the characteristics of interest and σ_i is a standard normally distributed error term. For households assigned to the TIOLI treatment, we do not observe WTP, but a binary indicator for whether WTP is greater than a certain value. That is, we observe

$$1\{WTP_i \geq p_i\} = 1\left\{\alpha_0 + X_i'\beta + \sigma_i - p_i \geq 0\right\} \tag{6}$$

where p_i is the price assigned to that household. We therefore run a probit model on the TIOLI data and simulated BDM purchase data. In this case, because the latent variable WTP_i is of primary interest, we focus on the estimated β coefficients rather than on marginal effects. That is, by normalizing the coefficient on price to be equal to -1, the coefficients obtained via probit are directly interpretable in terms of willingness to pay.

Table 5 shows the results of estimating the probit equation above on the TIOLI (Column (1)) and BDM (Column (2) samples. Column (3) displays the differences and Column (4) uses the bid information to estimate the WTP equation via OLS. The most interesting contrast is between columns 1 and 4. Column (1), which uses only the bound on willingness to pay provided by the TIOLI decision, appears to contain less plausible coefficients than those in Column (4), which use the prices WTP data provided by BDM. For example, in Column (1), both schooling and a recent episode of diarrhea in the household are negatively associated with willingness to pay, although the latter is not significant. In contrast, in Column (4), wealth indicators are more strongly (and positively) associated with willingness to pay, as is recent diarrheal illness among young children. We view these results as support for the idea that while BDM may reduce the overall level of willingness to pay (as shown in our previous results), it nevertheless provides useful and precise information.

6 Screening and Sunk-cost effects

As noted in the introduction, the BDM mechanism allows us to identify screening and causal effects of prices on use. That is, we are able to separate the extent to which usage of the filter varies by willingness-to-pay (independent of price paid) and by price paid (independent of willingness-to-pay). This occurs because the BDM elicits willingness-to-pay before randomly assigning price paid. Because filter usage is only measurable among buyers (i.e., those with bids above their draw), we restrict our estimates to the sample of buyers. Among this group, price paid is still random conditional on willingness-to-pay. We utilize four measures of usage, all measured as of the two-week followup survey. First, we include two subjective measures based on the water level in the plastic bucket. Water can only be drawn from the bucket once the water level in the bucket reaches the level of the spigot (after about one-half liter of filtering). Water below the level of the spigot indicates that the bucket has been emptied or has been in disuse for some time. We therefore use a variable indicating the water level at or above the level of the spigot, and one indicating the water level strictly above the level of the spigot. Second, we use one subjective measure indicating that there is water in the clay filter pot. This indicates that the filter is actively filtering water. Finally, we include self-reports of how many times per day the respondent fills the filter.

Table 6 estimates screening and causal effects of prices among respondents who received the BDM treatment. The table presents ordinary-least-squares estimates of the regression of usage on the BDM bid (willingness-to-pay) and draw (price paid):

$$use_{ic} = \alpha_0 + \alpha_1 WTP_{ic} + \alpha_2 D_{ic} + \varepsilon_{ic}$$

where use_{ic} represents the usage measure, $WTP_{ic}$is the respondent's BDM bid, D_{ic} and is her draw. Provided that the effects of each of these on usage are linear, the coefficients separately identify screening and causal effects. Across all usage measures, there is little evidence of either effect. The coefficients on both bid and draw are small and inconsistently signed. Households who bid more reported filling their filter slightly more often (significant at the 10% level), but given that the evidence is inconsistent across usage measures we cannot conclude that there is a causal effect of prices in this context.

We also compare the relationship between prices and use for both the BDM and TIOLI treatments. There are several issues to confront in making this comparison. First, because the TIOLI treatment did not include a second-stage randomization of the price paid, the relationship between prices and use could reflect both screening and sunk-cost effects.[9] We generate an equivalent comparison using BDM by not including price paid as a control, as in BDM those with higher willingness-to-pay paid higher prices for the filter, on average. Nonetheless, we recognize that this is not a perfect comparison, as BDM bids did not equal actual prices paid as in the case of TIOLI. Second, we must reduce the data from BDM in a way that makes it comparable to the data from the TIOLI experiment. We do this by constructing 3 quasi-observations for each actual BDM observation. We then assign hypothetical prices of 2, 4, or 6 to each observation and use the BDM bids to simulate TIOLI purchase behavior at each price. We then include in our regressions actual BDM buyers who would have hypothetically bought at the generated price ("quasi-buyers"). Thus, each actual BDM observation results in 3 quasi-observations for the regressions. This generates a group of actual buyers with $WTP > X, X = \{2,4,6\}$, as in the TIOLI observations. We run the following OLS regression, among TIOLI buyers and BDM quasi-buyers:

$$use_{ic} = \alpha_0 + \alpha_1 1\{WTP_{ic} \geq X\} + \varepsilon_{ic}$$

The results of this exercise are presented in Table 7. Overall, the table presents little evidence for a relationship between price and use for either the BDM or TIOLI treatments. In the BDM treatment, there is a negative relationship between price and use for water above the spigot, but this relationship is reversed when the frequency of filling is the measure of use. In no case is there a significant difference between the relationship under the BDM and TIOLI treatments, but these differences are not very precisely estimated.

[9] Because we cannot separately identify screening and causal effects in this section, if the true screening and causal effects are different from zero but have opposite signs, we may still find a null result.

7 Heterogenous Treatment Effects

The random nature of the BDM price draw creates a quasi-experiment that allows estimation of treatment effects, in which the random price draw provides an instrument for allocation of the filter. In this BDM is no different than TIOLI, which also provides an instrument: the random offer price. However, BDM additionally allows for the estimation of *heterogeneous* treatment effects, where treatment effects are allowed to vary by willingness to pay. This property is emphasized by Chassang, Padro-i-Miquel and Snowberg (2010).

In practice there are several different ways one could estimate heterogeneous treatment effects in the context of BDM. We present the simplest two here, but emphasize that BDM could be used in more sophisticated estimation strategies, such as those proposed by Heckman and Vytlacil (2005) and Heckman, Urzua and Vytlacil (2006).

First, we present IV estimates by quartile of WTP. We divide the sample into quartiles of willingness to pay, and in each quartile estimate the following outcome equation by linear two-stage least squares:

$$y_{ic} = \beta_0 + \beta_1 T_{ic} + x_{ic}\beta_2 + u_{ic} \tag{7}$$

Where y_{ic} is the outcome of interest (cases of diarrhea among children age 5 or below in the previous 2 weeks) for subject i in compound c, T_{ic} is a dummy for whether the subject possesses a filter, and T_{ic} is a vector of covariates. To instrument for the endogenous treatment variable , we use the following first-stage equation:

$$T_{ic} = \gamma_0 + \gamma_1 D_{ic} + x_{ic}\gamma_2 + v_{ic} \tag{8}$$

where D_{ic} is the subject's BDM draw. Since the draw is random, it is uncorrelated with u_{ic} and therefore it is a valid instrument for treatment.

Table 8 presents results from this estimation. (The numbers of subjects in each quartile is uneven because of lumpiness in the distribution of bids.) For each quartile, we present estimates with and without controls (number of adults in the compound, number of children in the compound, an indicator of whether the subject ever attended school, the first principal component of household durables, hectares of land owned by the compound and whether a child in the subject's household had a case of diarrhea in the two weeks prior to the sales visit) and village fixed effects. There is not a clear pattern in the estimated treatment effects, other than that the estimated treatment effect by group is inversely related to the mean of the untreated group.

The second method is to conduct a set of kernel IV regressions in the neighborhood of each level of WTP. We perform this exercise for each GHS 0.1 step between GHS 0.5 and 6 (inclusive), both without controls and with the same controls as in Table 5. We do not include village fixed effects in this set of regressions as some villages have zero or one subject in the domain of the kernel. Again, we do not see a clear pattern in the results, which are presented in Figure 6.

8 Conclusion

This paper provides an empirical test of whether BDM provides an accurate measure of an individual's willingness to pay through the sales of point-of-use water filters in Ghana. We provide strong evidence that BDM under-predicts willingness to pay relative to TIOLI. We also provide preliminary evidence against several potential explanations for this difference, including strategic bidding and anchoring.

We conclude by describing the analyses that form the focus of our continuing work. First, we will explore the alternative explanations driving the differences between willingness-to-pay estimates through BDM and TIOLI. As outlined in the model in the previous section, these differences could be driven by ambiguity aversion, disappointment aversion, in addition to strategic behavior. To distinguish among these potential channels, we will measure risk preferences among our original sample through measurement of risk, ambiguity, and disappointment aversion, as in e.g., Holt and Laury (2002).

Second, using our survey responses, we will examine whether BDM bids maintain ordinal information that can be useful in estimating screening and heterogeneous treatment effects even in the absence of a cardinal interpretation.

Third, we plan to use the results of the analyses described above to refine BDM methods so that it is more readily understood by participants and more amenable to implementation in the field, while retaining the desirable characteristics of BDM, in particular improved precision in measuring WTP, the ability to obtain distributional estimates of treatment effects, and to estimate selection and sunk cost effects.

References

Ashraf, Nava, James Berry, and Jesse M. Shapiro, "Can Higher Prices Stimulate Product Use? Evidence from a Field Experiment in Zambia," *American Economic Review*, 2010. Forthcoming.

Becker, G., M. DeGroot, and J. Marschak, "Measuring utility by a single-response sequential method.," *Behavioral Science*, 1964, *9*, 226–236.

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

Chassang, S., G. Padro i Miquel, and E. Snowberg, "A Principal-Agent Approach to Randomized Controlled Experiments," September 2010. Mimeo, London School of Economics.

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

Gul, Faruk, "A Theory of Disappointment Aversion," *Econometrica*, 1991, *59* (3), 667–86.

Heckman, James J. and Edward Vytlacil, "Structural Equations, Treatment Effects, and Econometric Policy Evaluation," *Econometrica*, 2005, *73* (3), 669–738.

Heckman, James, Sergio Urzua, and Edward Vytlacil, "Understanding Instrumental Variables in Models with Essential Heterogeneity," *The Review of Economics and Statistics*, 2006, *88* (3), 389–432.

Hoffmann, Vivian, "Intrahousehold Allocation of Free and Purchased Mosquito Nets," *American Economic Review*, 2009, *99* (2), 236–41.

Holt, Charles A. and Susan K. Laury, "Risk Aversion and Incentive Effects," *American Economic Review*, 2002, *92* (5), 1644–1655.

Horowitz, John K., "The Becker-DeGroot-Marschak mechanism is not necessarily incentive compatible, even for non-random goods," *Economics Letters*, 2006, *93* (1), 6–11.

_ , "What Do We Know About the Performance of the Becker-DeGroot-Marschak Mechanism?," 2006.

Imbens, Guido W. and Joshua D. Angrist, "Identification and Estimation of Local Average Treatment Effects," *Econometrica*, March 1994, *62* (2), 467–475.

Irwin, Julie R., Gary McClelland, Michael McKee, and William D. Schulze, "Payoff Dominance Vs. Cognitive Transparency in Decision Making," *Economic Inquiry*, 1998, *36* (2).

Karni, Edi and Zvi Safra, ""Preference Reversal" and the Observability of Preferences by Experimental Methods," *Econometrica*, May 1987, *55* (3), 675–685.

Keller, L. Robin, Uzi Segal, and Tan Wang, "The Becker-DeGroot-Marschak mechanism and generalized utility theories: Theoretical predictions and empirical observations," *Theory and Decision*, March 1993, *34* (2), 83–97.

Machina, Mark J, "Expected Utility Analysis without the Independence Axiom," *Econometrica*, 1982, *50* (2), 277–323.

Mazar, Nina, Botond Koszegi, and Dan Ariely, "Price-Sensitive Preferences," *SSRN eLibrary*, February 2010.

Noussair, Charles, Stephane Robin, and Bernard Ruffieux, "Revealing consumers' willingness-to-pay: A comparison of the BDM mechanism and the Vickrey auction," *Journal of Economic Psychology*, 2004, *25* (6), 725–741.

Rutstram, E., "Home-grown values and incentive compatible auction design," *International Journal of Game Theory*, 1998, *27* (3), 427–441.

Shogren, Jason F., Sungwon Cho, Cannon Koo, John List, Changwon Park, Pablo Polo, and Robert Wilhelmi, "Auction mechanisms and the measurement of WTP and WTA," *Resource and Energy Economics*, 2001, *23* (2), 97–109.

Smith, Vernon L., "Microeconomic Systems as an Experimental Science," *The American Economic Review*, December 1982, *72* (5), 923–955.

World Health Organization, "Water, sanitation and hygiene: links to health," http://www.who.int/water_sanitation_health/publications/facts2004/en/ 2004.

_ , *World Health Report 2005* 2005.

_ , *World health statistics 2011*, Geneva Switzerland: World Health Organization, 2011.

Table 1: Sample Composition

	Assigned		Sales Offer Made	
	Compounds	Subjects	Compounds	Subjects
BDM Standard	99	249	88	213
BDM Marketing	93	240	81	196
BDM Anchoring	97	242	83	196
Total BDM	289	731	252	605
TIOLI Standard	100	264	85	220
TIOLI Anchor	94	247	86	211
TIOLI Random	99	268	87	229
Total TIOLI	293	779	258	660
Total	582	1510	510	1265

Notes: Our unit of observation is any primary caretaker of one or more children age 12 and under. A compound is an extended patrilineal family of several sub-families living in a cluster of huts. Treatments were randomized at the compound level and all inference is robust to clustering at the compound level. Our activities were conducted in 14 villages, selected according to several criteria described in the text.

Table 2: BDM Effect

	price=2		price=4		price=6	
BDM	-0.185	***	-0.159	***	-0.091	**
	(0.033)		(0.053)		(0.039)	
Constant	0.916	***	0.470	***	0.199	***
	(0.024)		(0.046)		(0.035)	
Observations	841		822		789	
Num. clusters	391		385		374	

Standard errors in parentheses

* p<0.10, ** p<0.05, *** p<0.01

p-value for joint test that BDM=0 in all three equations: 0.000.

Notes: the dependent variable is a dummy variable indicating either that the individual did purchase the filter at the indicated price (for take-it-or-leave-it subjects) or that the individual's willingness to pay was as least as high as the indicated price (for BDM subjects). The intercept represents the share purchasing under TIOLI, while the coefficient on BDM indicated the difference in demand between BDM and TIOLI at that price. Each BDM subject appears in all three regressions, while each TIOLI subject appears only in the regression corresponding to her offer price. The p-value for a joint test that the BDM coefficient is zero in all three equations is calculated from SUR estimation. All standard errors account for clustering at the compound (extended family) level.

Table 3: Equality of Distributions of Bids
Comparison with Standard BDM

	Market	Anchor
Wilcoxon		
Z-statistic	2.82	-0.92
P-value	0.022	0.758
Num. Obs.	408	407
Kolmogorov-Smirnov		
	Market	Anchor
D-statistic	0.141	0.061
P-value	0.052	0.756
Num. Obs.	408	407

Notes: this table reports results of nonparametric tests for equality of bid distributions across BDM treatments. The anchoring and marketing treatments (described in the text) are separately compared to standard BDM. P-values are calculated via cluster-bootstrapping with resampling at the compound level.

Table 4: Differences between TIOLI subtreatments

	price=2	price=4	price=6
Random TIOLI	0.009	-0.135	0.024
	(0.063)	(0.105)	(0.084)
Anchoring TIOLI	0.053	-0.226 *	-0.073
	(0.060)	(0.118)	(0.079)
Constant	0.899	0.594	0.218
(Standard TIOLI)	(0.053)	(0.084)	(0.054)
Observations	238	219	186
Num. clusters	119	113	102

Standard errors in parentheses
* p<0.10, ** p<0.05, *** p<0.01
p-val for joint test that coefficient on Random TIOLI = 0 in all three equations: 0.595. p-val for joint test that coefficient on Anchoring TIOLI = 0 in all three equations: 0.081.

Notes: this table reports results of a linear probability model for purchase of the filter at the take-it-or-leave-it 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. All standard errors account for clustering at the compound (extended family) level.

Table 5: Relationship Between Willingness to Pay and Baseline Characteristics

	Purchase Decision Probit			WTP OLS
	TIOLI	BDM	Diff.	
	(1)	(2)	(3)	(4)
Num. adults in compound	-0.066 *	-0.038	0.029	-0.030
	(0.032)	(0.044)	(0.054)	(0.043)
Num. children in compound	0.012	0.043	0.031	0.033
	(0.032)	(0.034)	(0.047)	(0.033)
Ever attended school	-0.412 +	0.008	0.420	0.476
	(0.223)	(0.375)	(0.436)	(0.441)
Year-round improved water source	-0.172	-0.220	-0.048	-0.436 +
	(0.275)	(0.279)	(0.392)	(0.243)
First principal component of durables	0.075	0.084	0.009	0.127
	(0.069)	(0.076)	(0.103)	(0.078)
Land owned by household (Ha.)	0.031	0.209	0.178	0.260
	(0.222)	(0.253)	(0.337)	(0.240)
Child 0-5 with diarrhea at baseline	-0.334	0.432	0.766 *	0.642 *
	(0.227)	(0.278)	(0.359)	(0.291)
Price (GHS)	-1	-1	0	
	.	.	.	
Num. Obs.	660	602	1,262	

+ denotes significant at 0.10; * significant at 0.05; ** significant at 0.01

Notes: Columns (1) and (2) display coefficients from probit regressions of the household's purchase decision on key baseline characteristics. For the BDM subjects (column (2)), the purchase decision is synthetic, constructed as an indicator for the subject's bid being at least as high as a randomly generated synthetic TIOLI price in {2,4,6}. (See discussion in the text.) Column (3) displays differences between coefficients. Column (4) displays coefficients from a linear regression of WTP (BDM bid) on the same covariates. Because the coefficient on price in equations (1) and (2) is normalized to -1, the estimated coefficients can be interpreted as analogous to willingness to pay and are directly comparable to those in equation (4). Standard errors clustered at the compound (extended family) level in parentheses.

Table 6: Sunk costs and screening effects
Regression of use on bid and draw

	Dependent Variable			
	Water at or Above Spigot (1)	Water At Spigot (2)	Water in pot (3)	Fill Freq. (4)
Bid	-0.017	-0.015	0.003	0.053*
	(0.014)	(0.015)	(0.011)	(0.027)
Draw	0.034	-0.009	-0.033	-0.038
	(0.021)	(0.027)	(0.022)	(0.040)
Mean Dep. Var.	0.756	0.526	0.783	1.523
R-squared	0.013	0.009	0.012	0.023
Num. Obs.	234	234	217	217

The sample includes those in the BDM treatment who purchased the filter. Each column presents the results of a separate regression of use on BDM bid and BDM draw, among those whose draw was below their bid.
+ $p < 0.10$; * $p < 0.05$; ** $p < 0.01$

Table 7: Relationship between WTP and Use

	Combined	TIOLI	BDM	Difference
	(1)	(2)	(3)	(4)
Panel A: Dependent Variable: Water at or Above Spigot				
Price	-0.013	-0.001	-0.023	-0.021
	(0.011)	(0.017)	(0.015)	(0.022)
Mean Dep. Var.	0.743	0.749	0.737	
R-squared	0.002	0.000	0.005	
Observations	673	319	354	
Panel B: Dependent Variable: Water Above Spigot				
Price	-0.029*	-0.021	-0.037*	-0.016
	(0.013)	(0.022)	(0.016)	(0.027)
Mean Dep. Var.	0.478	0.476	0.480	
R-squared	0.007	0.003	0.011	
Observations	673	319	354	
Panel C: Filter Pot Contains Water				
Price	-0.013	-0.008	-0.016	-0.008
	(0.012)	(0.020)	(0.015)	(0.024)
Mean Dep. Var.	0.785	0.794	0.776	0.785
R-squared	0.002	0.001	0.003	
Observations	622	296	326	
Panel D: Dependent Variable: Frequency of filling				
Price	0.038	-0.005	0.074**	0.079
	(0.025)	(0.044)	(0.027)	(0.052)
Mean Dep. Var.	1.554	1.481	1.619	
R-squared	0.004	0.000	0.017	
Observations	614	289	325	

Columns 1-3 present the results of regressions of use on "price." For TIOLI treatments, price equals the assigned price. Price equals WTP≥X for X={2, 4, 6} for BDM observations. Each BDM observation is split into three observations (one for each price) for the purposes of these regressions. See text for additional details. Column 4 presents the difference in coefficients between the TIOLI sample (Column 2) and the BDM sample (Column 3).

+ $p < 0.10$; * $p < 0.05$; ** $p < 0.01$

Table 8: Treatment Effects by Quartile of WTP

	Quartile 1 0≤WTP<2		Quartile 2 2≤WTP<2.5		Quartile 3 2.5≤WTP<4		Quartile 4 4≤WTP<19	
	(1)	(2)	(3)	(4)	(5)	(6)	(5)	(6)
Filter Purchase	-0.142	-0.120	0.132	0.129	-0.168	-0.156	-0.067	-0.058
	(0.213)	(0.210)	(0.115)	(0.093)	(0.115)	(0.113)	(0.079)	(0.071)
Controls:	No	Yes	No	Yes	No	Yes	No	Yes
Village FEs:	No	No	No	No	No	No	No	No
Mean dep. var.:	0.194	0.197	0.048	0.026	0.257	0.257	0.184	0.184
r2	-0.003	0.148	0.070	0.289	0.020	0.097	0.003	0.120
N	87	86	72	67	79	79	122	120

+ denotes significant at 0.10; * significant at 0.05; ** significant at 0.01

Notes: Sample for each regression consists of BDM households with bids in the quartile given in the column header. Dependent variable is an indicator for one or more cases of diarrhea among children in the household aged 0-5 over the previous 2 weeks. Filter purchase is instrumented by TIOLI price for TIOLI subjects and BDM draw for BDM subjects. Standard errors clustered at the compound (extended family) level in parentheses. The mean of the dependent variable shown is calculated among non-treated households.

Figure 1: Kosim Filter

Figure 2: Histogram of BDM Bids

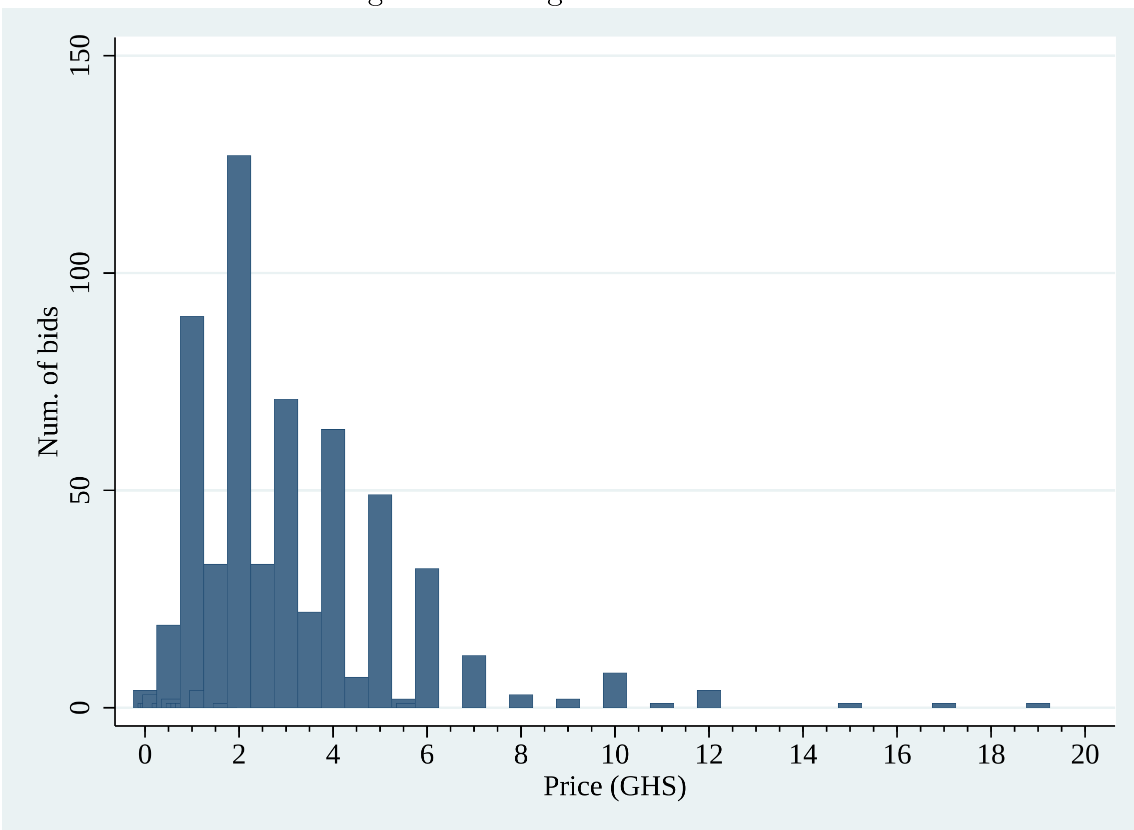

Notes: this figure plots a histogram of all bids for the filter under the BDM treatments (603 observations total). Approximately 1.5 GHS per USD.

Figure 3: BDM-TIOLI Comparison

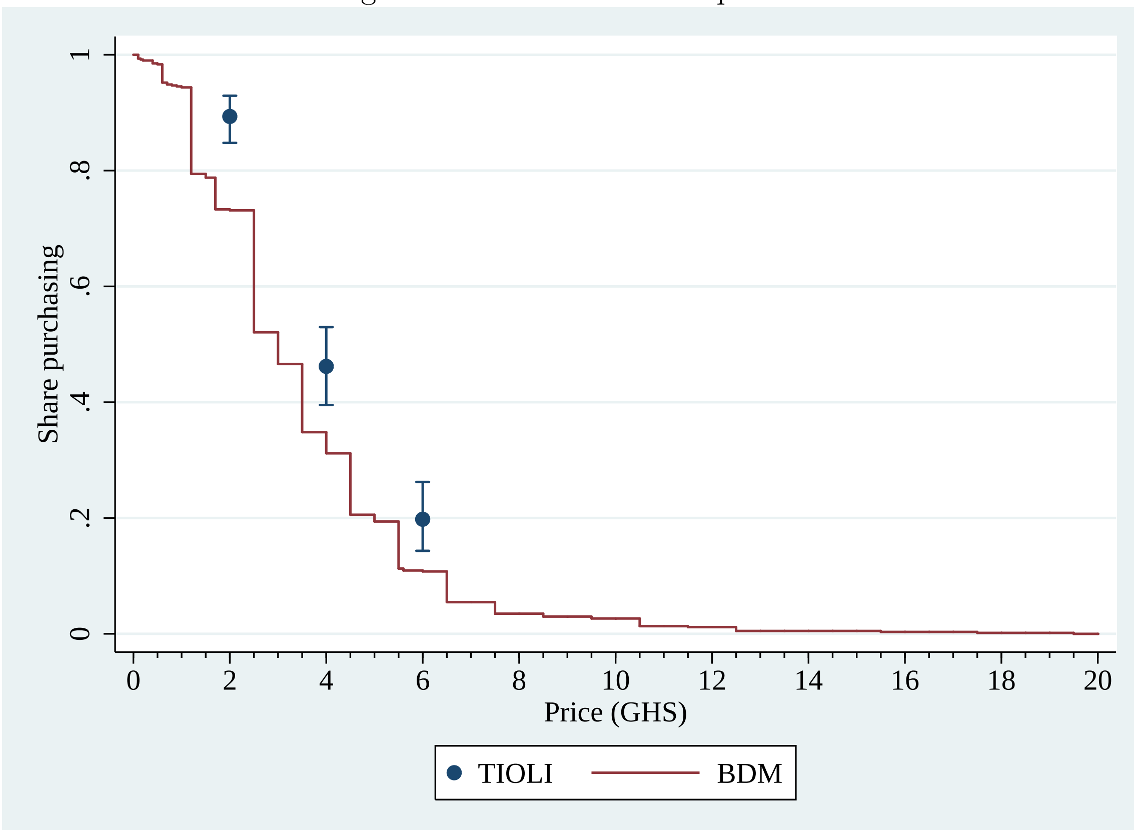

Notes: the solid line plots the share of BDM respondents with bids at or above the indicated price. The markers indicate the share of TIOLI respondents who purchased the filter at the offered price with 95% confidence intervals. 603 BDM observations. 645 TIOLI observations, of which 238 at a price of 2, 219 at a price of 4 and 186 at a price of 6.

Figure 4: Comparison of BDM Sub-treatments

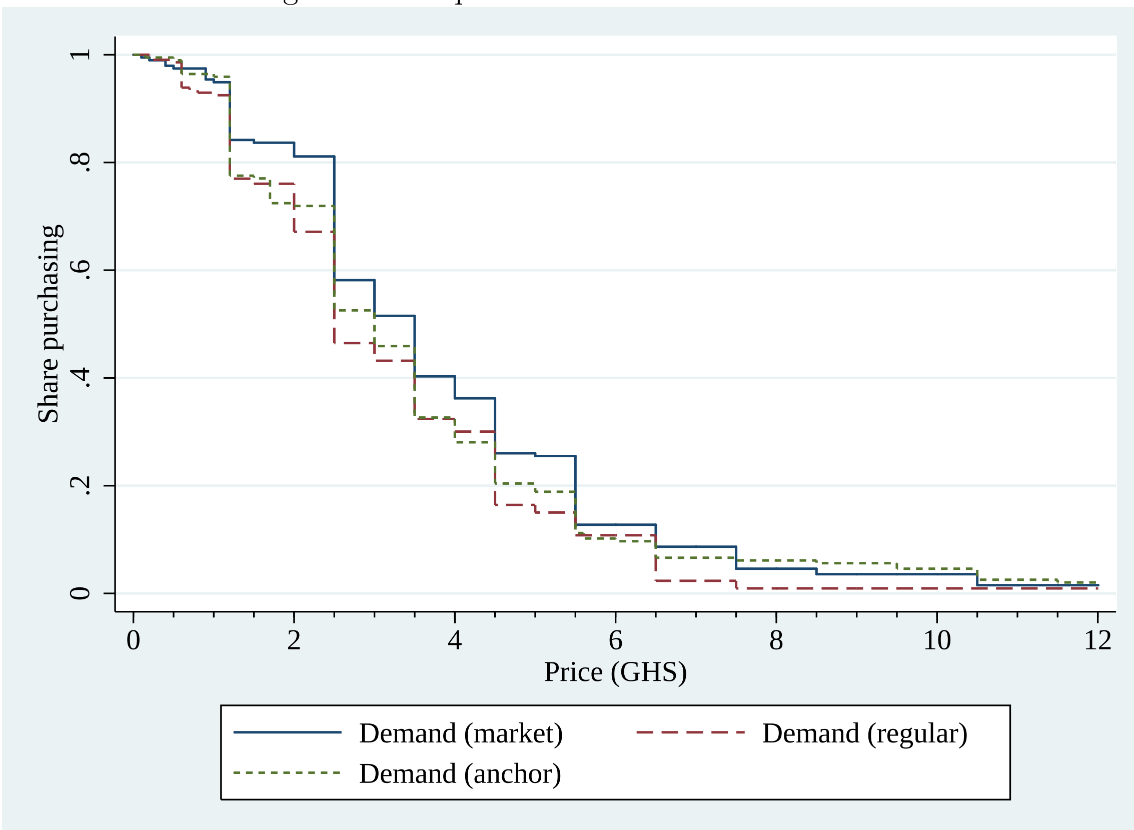

Note: the standard, anchoring and marketing treatments are described in detail in the text. Bids are truncated at 12 GHS because there are no appreciable differences among the three treatments at prices above 12. 603 observations total, of which 211 are standard BDM, 196 are marketing BDM and 196 are anchoring BDM. All treatments were randomized at the compound (extended family level); prices were drawn by individual respondents.

Figure 5: Comparison of TIOLI Sub-treatments

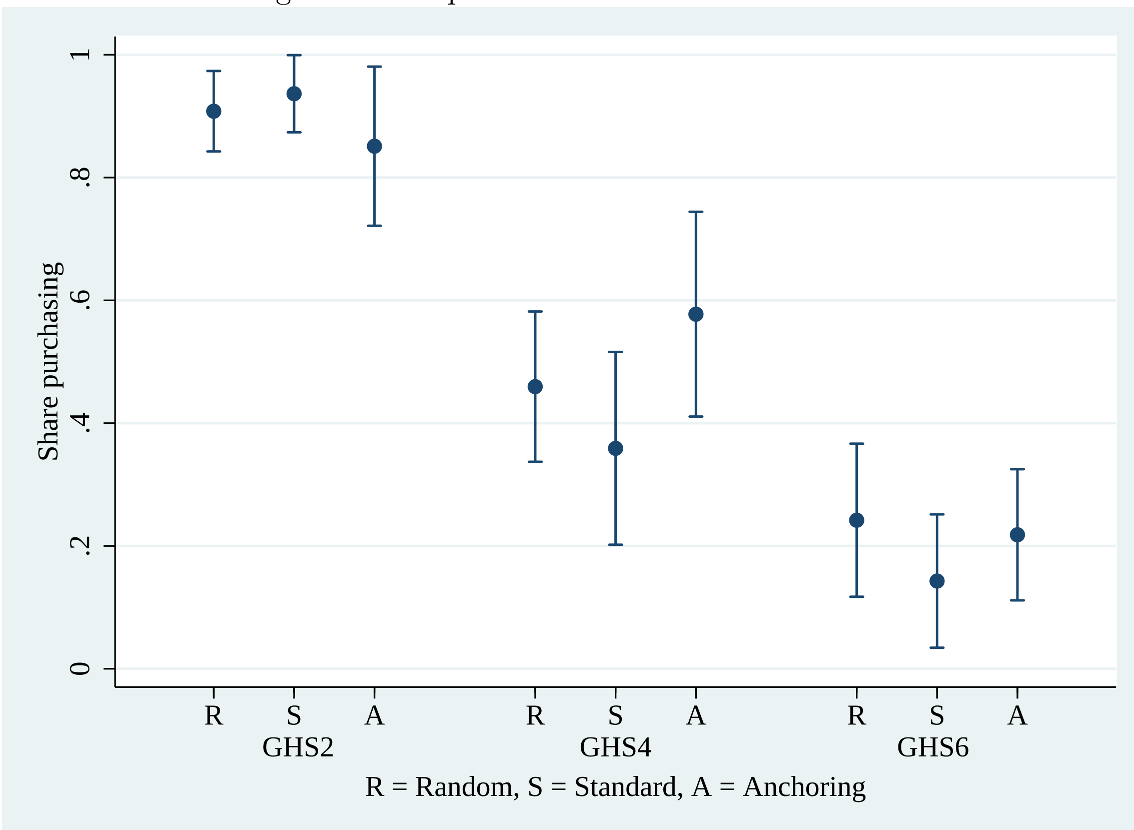

Notes: this graph plots demand for the filter at each take-it-or-leave-it price, for each TIOLI sub-treatment. The random TIOLI, regular TIOLI and anchoring TIOLI treatments are described in detail in the text. Each treatment was randomized at the compound level. For the regular and anchoring TIOLI treatments, the price was also randomized at the compound level. For the random TIOLI treatment, the price was drawn by individual respondents.

Figure 6: Local Estimates of Heterogeneous Treatment Effects