Pricing People into the Market: Targeting through Price Discrimination

Terence R. Johnson
and
Molly Lipscomb*

May 21, 2022

Abstract

Maximizing the impact of aid — particularly in the presence of externalities — requires focusing benefits on those most likely to respond. We use mechanism design tools to provide a theory of targeting on observables, collect data to operationalize the solution, then test it in a randomized controlled trial. The treatment led to an increase in the market share of a socially beneficial good of 12.7 percentage points in the group likely to purchase the inferior product, 5.3 percentage points overall. We estimate and use a model of household behavior to compare alternative market designs, information structures, and market sustainability.

JEL Classification Codes: L1, D4, C7, D8
Keywords: Targeting, Subsidies, Mechanism Design

*Johnson is visiting assistant professor of economics at the University of Virginia (trj2j@virginia.edu). Lipscomb is associate professor of economics and public policy at University of Virginia's Batten School and the Department of Economics (molly.lipscomb@virginia.edu). We thank Shoshana Griffith, Adrien Pawlik and Adama Sankoudouma for excellent research assistance. We are grateful to the Bill and Melinda Gates Foundation for funding this project and research. We are grateful for helpful comments from Radu Ban, Jim Berry, Vivek Bhattacharya, Diane Coffey, Joshua Deutschmann, Taryn Dinkelman, Pascaline Dupas, Jessica Goldberg, Joe Kaboski, Isaac Mbiti, Dilip Mookherjee, Laura Schechter, and conference participants at the Gates Foundation Workshop, Midwest International Development Conference, NBER Development Summer meetings, NYU, NYU Abu Dhabi, the North Carolina TREE seminar, University of Stockholm, and William and Mary.

1 Introduction

When the budget of a social aid program that subsidizes consumption is tight, the poorest or most deserving groups can be targeted to receive assistance based on observable information about individual households through proxy means testing or other information revelation mechanisms (Alatas et al., 2012; Hanna and Olken, 2018). This ensures that aid ends up in the hands of relatively poor households who are most likely to benefit (Coady et al., 2004b; Karlan and Thuysbaert, 2019). On the other hand, when the objective of a subsidy program is to increase consumption of a good with positive externalities such as sanitation products, education, or vaccines, the problem becomes one of not only helping the poorest, but also determining the necessary level of subsidy to those with low willingness or ability to pay without "wasting" the budget on subsidies for households who would have purchased the good in the absence of assistance (Berry et al., 2020; Kremer and Miguel, 2007a). How should subsidy levels be set for these goods conditional on household observables in order to maximize the impact of a limited budget, particularly when households have private information about their ability and willingness to pay for the good? This paper approaches this problem through the lens of price discrimination, and uses a field experiment to design a subsidization schedule and test its performance using a randomized controlled trial in the market for sanitation services in Burkina Faso.

We provide a theory of targeting and show how the optimal subsidization schedule uses observable and verifiable information to set prices that strategically distribute subsidy aid and cross-subsidize poor households with profits made from sales to wealthier ones. We then gather the data necessary to operationalize this solution and apply it to a particular problem: the removal of human fecal sludge in peri-urban neighborhoods of Ouagadougou. To rigorously test the impact of the intervention, we use a randomized controlled trial, comparing a treatment group of households with access to our market to a control group that was excluded. Our data-driven cross-subsidization scheme and competitive procurement practices stretched a subsidy budget that allowed \$3.00 per household to \$5.43 for the poorest households, increasing the market share of sanitary removal services among the poorest households by 12.7 percentage points and increasing the overall market share of mechanical desludging by 5.3 percentage points. We conclude by providing a number of counterfactual experiments on the market design, the kind of observable information used to design the platform, and the trade-off between subsidization and treatment effect.

Our setting is the handling and disposal of human fecal sludge in residential compounds in Burkina Faso, which is a challenge in urban sanitation throughout the developing world. Lack of adequate sanitation is a primary cause of approximately 10% of global diseases, primarily diarrheal (Mara et al., 2010). While there has been substantial attention to increasing access to toilets for households (Guiteras et al., 2015; Kar and Pasteur, 2005) and reducing open defecation in rural areas (Gertler et al., 2015), there has been less attention to sanitation issues in urban environments where the impact of inadequate sanitation may be particularly large (Coffey et al., 2014). Rates of diarrhea are extremely high in developing countries: 1.8 billion people globally use a source of drinking water with fecal contamination, and 2.4 billion people lack access to safely managed sanitation services (WHO and UNICEF, 2015), which can result in stunting and other developmental disadvantages for young or vulnerable household members (Spears, 2013). Improving sanitation and water access is one of the Sustainable Development Goals, and while subsidies have been effective at increasing take-up of health and sanitation goods, the demand for such products often remains low (Kremer and Miguel, 2007b; Cohen and Dupas, 2010; Dupas, 2014) suggesting that larger subsidies for these goods may be necessary, particularly for the poorest households.

Household pits can be emptied mechanically, where a crew of two to four workers uses a truck-mounted vacuum to remove fecal sludge, or manually, where pits are cleaned out by family members or hired workers using trowels and buckets. Mechanical desludging minimizes exposure to fecal sludge and ensures its removal from the immediate neighborhood, while manual desludging typically ends with the disposal of the sludge in the street near the dwelling, resulting in negative externalities for nearby households through exposure to parasites and pests. Pits fill approximately every 6 to 12 months, at which point a household must find a service provider. At baseline 29% of households report having last used a manual desludging, and of these 14.4% reported manually desludging their own pit for free. The most common ways to find desludgers are calling the desludger that they used last time (44%), going to a parking lot (14%), and asking family or friends for a desludger phone number (8.5%); 12% of households that desludged manually began by searching for a mechanical desludger but eventually gave up. Prices are 1,700 CFA higher on average for households that use an intermediary, 945 CFA higher when they call a number that they saw on a truck, and 500 CFA higher when they ask a desludger that they know lives nearby, suggesting there is some price discrimination occurring in the market. The median household reports looking for their last mechanical desludger for 12 days and having searched for a mechanical desludger for 24 days or more on at least one occasion in the past. Over 60% of those who searched for a mechanical desludger but ultimately used a manual desludger report searching for a week or more. Overall, this paints a picture of a decentralized and inefficient market where service providers can extract rents from households who know that turning down a high price now will lead to a new search that will likely end with a similarly high price quote.

We adopt the perspective of a municipality, government agency, or NGO addressing these market failures by creating a platform to maximize utilization of mechanical desludging services, subject to a limited subsidy budget. Even when conditioning on observable or verifiable information about households in general — such as number of rooms in the dwelling, the size of the compound, or past municipal water or electricity bills — there is still uncertainty about a particular household's willingness-to-pay and the expected price it would face in the decentralized search market. An ideal toolbox for this kind of targeting problem is mechanism design, which studies what outcomes are achievable when participants have private information and individual agency. Because of the externalities of inferior sanitation methods, a budget-constrained social planner like a municipality or NGO is concerned not only with whether a household purchases from its platform like a profit maximizer, but whether the household would purchase a manual or mechanical desludging in the absence of the intervention. Thus, the objective of targeting is to strategically distribute aid to bring the improved product within the financial reach of relatively poor households, but also avoid giving aid to relatively rich households that would purchase in the absence of assistance.

The optimal platform uses observable information about households to formulate inferences about their propensity to purchase the improved product in the absence of the intervention, quoting high prices to households that will likely purchase anyway and low prices to households that are not just poor, but likely on the margin of purchasing[1]. This raises profits from relatively rich households and then distributes those profits and subsidy dollars to the households who are unlikely to purchase on their own but are most cost effective to assist in expectation. This approach can be applied whenever a socially motivated actor like a government or NGO seeks to increase total market share of an improved product, including goods that generate externalities like vaccines, cookstoves or toilets, as well as ones that do not such as agricultural technologies or solar lights.

We then gather the data necessary to build the optimal platform in the context of desludging services in Ouagadougou. We use a highest-rejected bid auction to procure services from mechanical desludgers, and a lowest-rejected bid auction to sell those services to households. These generalizations of the second-price auction give the firms and households a weakly dominant strategy to bid honestly, providing us with unbiased estimates of the firms' costs and the households' willingness-to-pay for a desludging from the platform. Combining the willingness-to-pay data with market data on the desludging purchases that households most recently made and the prices they paid allows us to predict which households are likely to purchase on their own and which are the cheapest for the platform to convert to mechanical desludging through an attractive price offer. Maximizing the *mechanical market share* then delivers the optimal pricing rule for the platform. This is an important distinction, because the goal of the platform is not to maximize platform sales, profit, or its own market share, but instead to maximize the total market share of mechanical desludging. That households can opt out of our market in favor of the prevailing, decentralized market differentiates our setting from many studies that focus on increasing demand for novel products that face little or no competition.

We then use a randomized controlled trial to measure the impact of the optimal platform. Access to the platform is provided to a second random sample of households, and their outcomes are compared with a third random sample of households that serves as a control group who do not receive access to the platform. During the baseline interview for these participants, each treatment household receives the optimal

[1] Compare this with proxy-means testing, which uses data to develop a predictive model of whether households are below a poverty line or not, and then offers aid to everyone judged to be below the poverty line, subject to a budget constraint.

take-it-or-leave-it offer based on the observable information and the enumerator's subjective assessments. We find that the price targeting treatment generates a 12.7 percentage point increase in market share of desludging among the poorest households. There is no impact from the treatment on the wealthy households who have a high (99.3%) use of mechanized desludging services, even without the treatment. This leads to an overall pooled treatment impact on market share of 5.3 percentage points. The health impact of these effects, particularly on the poorest households, is large. We find that reported diarrhea rates among children decrease by 5.2-7.0 percentage points among the poorest households following treatment.

By offering a new approach to targeting that relies on data to decide how much and to whom to distribute assistance, this paper contributes to a variety of literatures. Targeting in existing programs has been found to be only moderately successful: for example, Coady et al. (2004a) find that some targeting programs transfer only 25% more than random or universal allocation to poor households, with 27% of programs found to be regressive. Several methods of targeting aid and subsidies have been proposed and evaluated: proxy means tests based on the household's ownership of a basket of assets (Kidd and Wylde, 2011; Narayan and Yoshida, 2005; Banerjee et al., 2020); ordeal mechanisms in which the household must collect and submit coupons or undergo an application process (Alatas et al., 2012; Dupas et al., 2016; Alatas et al., 2016)[2]; and community-based targeting in which members of the local community or local government select which people should receive the program (Basurto et al., 2020). Allcott et al. (2015) show that restricting access to energy efficiency subsidies to those who are unlikely to otherwise purchase on their own can substantially increase the efficiency of the program. Our focus on effective pricing policies for improving the efficiency of utility use is related to McRae and Wolak (2021); Wolak (2016); Szabo (2015); Borenstein (2012), and our use of mechanism design tools to improve the effectiveness of policy and targeting is related to Pathak (2011); Dur et al. (2020).

The randomized controlled trial rigorously tests the effectiveness of the proposed platform, but a number of questions remain about how other designs would have performed. In Section 4, we estimate a model that predicts the purchasing behavior of households in response to counterfactual price quotes, allowing us to predict how

[2]See Olken (2016) for a review.

alternative policy designs and information structures would have performed, as well as explore the trade-off between the treatment effect and subsidization level. We show that auctions and proxy-means testing deliver slightly lower average treatment effects, but much lower treatment effects among the poorest households. Auctions select on willingness-to-pay and essentially act as a transfer to relatively wealthy households who would purchase anyway, attracting them through a modest discount on their expected market price achieved by putting suppliers in price competition. Proxy-means testing targets on noisy measures of wealth that do not correlate well with actual behavior, nor use cross-subsidization or adjust assistance to ensure poor households are sufficiently subsidized to buy the expensive service. A price ceiling typically benefits richer households who otherwise face price discrimination, without providing assistance to relatively poor households. To investigate what would happen if different kinds of information were used to design the platform, we restrict the variables used to determine prices to those an NGO and a municipal authority would have available, re-solve for the optimal platform, and predict counterfactual performance of the alternative design in the RCT. In general, the treatment effects remain relatively large, but as it becomes more difficult to target poor households the platform becomes unintentionally more generous, leading to larger violations of the expected budget balance constraint. Finally, we investigate the financial sustainability of the platform and the trade-off between the treatment effect and budget by estimating counterfactual outcomes under alternative values of the subsidy. We find that the platform can run a profit, and even without subsidization, there can be a positive treatment effect because access to the market and competition among service providers in the procurement auctions can be enough to increase welfare. Reaching the very poorest households, however, can be very expensive, requiring that subsidized prices fall to 60% of the average market price to induce them to take up the healthier service.

2 Platform Design

In general terms, "targeting" refers to the method by which beneficiaries are selected to receive aid from social programs (Alatas et al., 2016), including proxy-means testing, ordeal mechanisms, community-based selection games, and auctions. But among all the possible ways to encourage take-up of a socially beneficial good subject to a subsidy budget and incentive constraints for participants, which intervention has the largest impact? This section describes our solution to the targeting problem[3] and how market and experimental data were gathered and used to design the randomized controlled trial in Section 3.

Consider a budget-constrained platform that operates alongside a prevailing search market to maximize market share of a socially beneficial health product, mechanical desludging. Each household has *observables* associated with it that can be easily verified and measured, x, such as its neighborhood or the quality of its dwelling, or known to a government authority like ONEA through census or billing records, such as water or electricity bill expenditures. A household must decide whether to purchase the manual or mechanical service, and has two sources of private information: its willingness-to-pay for mechanical desludging, w, with distribution $F_w[w|x]$, and the price it expects to pay for a mechanical desludging in the search market, r, with distribution $F_r[r|x]$. Households are risk-neutral, have quasi-linear utility, and the disutility of the outside option of manual desludging is normalized to zero for each x. Let c_x be the cost to the platform of procuring a desludging for a household with observables x, and $\bar{s}$ be the available subsidy budget–the average amount of funding for each household.

Since w and r are known only to the household, relatively rich households that would purchase the service on their own satisfy $w \geq r$, but wish to misrepresent themselves to the platform as low willingness-to-pay households for whom $r > w$ in order to receive aid or lower prices. Consequently, relatively poor households for whom it is actually true that $r > w$ and who do not purchase the service on their own cannot be subsidized more generously without also allocating subsidy dollars to those relatively rich households impersonating them as a consequence of the private information. This is the essence of the targeting problem.

To solve for the optimal prices when w and r are unobserved, consider how a household with observables x responds to a quoted price of t_x from the platform. There are four cases. If $w \geq r$, so that the household would have purchased anyway, then the price is only attractive if $t_x < r$, and the household joins the platform and becomes a *participating buyer*; otherwise, $r < t_x$, and the household would reject the

[3]See Appendix A for the details of a standard mechanism design analysis. The rest of this section provide a discussion of the optimal solution.

platform's price and instead buy in the market, making it a *non-participating buyer*. On the other hand, if $w<r$, the household would be a *non-buyer* in the absence of the intervention, and only accepts the platform's offer and becomes a *switcher* if $t_x \leq w$. Figure 2 illustrates how a particular price quote of t_x divides the households with different (w,r) pairs into these four groups. Thus, only w is relevant to a relatively poor household that doesn't buy since $r>w$, and only r is relevant to a relatively rich household that does buy since $r<w$; only the quantity $\eta=\min\{w,r\}$ is relevant for determining selection onto the platform, which we call the *willingness-to-switch*.

Switching Behavior. Consider quoting a price t_x to a household with observables x. For some realizations of that household's private information w and r, the household will be indifferent between buying on the platform or taking their outside option, and so satisfy $\eta = t_x$. This includes switchers for whom $r \geq w = t_x$ and participating buyers for whom $w \geq r = t_x$. The hazard rate of transitioning from non-buyer to switcher at price of t_x is

$$h_w[t_x|x] = \frac{f_w[t_x|x]}{1-F_w[t_x|x]},$$

and the hazard rate of transitioning from non-participating buyer to participating buyer is

$$h_r[t_x|x] = \frac{f_r[t_x|x]}{1-F_r[t_x|x]},$$

and these quantities correspond approximately to the masses of agents along the indicated line segments in Figure 2. The probability that a household of type x who is indifferent between joining the platform or not at a price of t_x is actually a switcher is then given by the conditional probability

$$\sigma(t_x,x) = \frac{h_w[t_x|x]}{h_w[t_x|x]+h_r[t_x|x]},$$

which corresponds to the marginal benefit of reducing a price for a household with observables x, since the platform is trying to maximize take-up of the socially beneficial good.

Optimal Prices. Let λ^* be the Lagrange multiplier on the platform's budget constraint at the optimum. Given the platform's budget constraint and inference about the likelihood of a household being a switcher, the optimal platform prices t_x^* satisfy

$$\underbrace{t_x^*}_{\text{Price}} = \underbrace{c_x}_{\text{Service cost}} + \underbrace{\frac{1-F_\eta[t_x^*|x]}{f_\eta[t_x^*|x]}}_{\text{Informational Rent}} - \underbrace{\frac{\sigma(t_x^*,x)}{\lambda^*}}_{\text{Social Discount}}$$

where c_x is the expected cost to the platform of serving a household with observables x. This illustrates how the platform is a "profit-minded social planner," who places some weight on profits and some on consumption of improved services, where the weight is endogenously determined by balancing the budget needed for an additional or larger subsidy with the relative likelihood of the household switching to mechanical. A standard monopolist[4] would set its price equal to service cost plus the informational rent, but the platform is instead maximizing take-up of the mechanical service, reflected in the social discount term. Each observable type x has some probability of switching on the margin, which is then deflated by the shadow value of a subsidy dollar λ^*, representing the opportunity cost of providing assistance to this type x over some other type x'.

While participating buyers might be contributing profits that relax the platform's budget constraint, they do not increase the share of mechanical services purchased overall and increase the objective function: only the set of switchers corresponds to increased social welfare. The platform thus acts as a "profit-minded social planner," optimally charging relatively high prices to households who would likely purchase in the absence of the intervention and relatively low prices to households who will likely be unable to afford the mechanical service on their own but are the most likely to switch. This analysis provides a novel formalization of targeting as a mechanism design problem and illustrates the difference between classical price discrimination and targeting: the platform cares not only about making a profitable sale, but also what the household would likely have done in the absence of the intervention.

[4]Faced with the problem $\max_{\{t_x\}_{x\in X}} \mathbb{E}_x[(1-F_\eta[t_x|x])(t_x-c_x)]$, the monopolist's first-order conditions satisfy $-f_\eta[t_x^*](t_x^*-c_x)+(1-F_\eta[t_x^*|x])=0$ for each $x\in X$, which is equivalent to the standard "MR=MC" condition, $t_x^*-(1-F_\eta[t_x^*|x])/f_\eta[t_x^*|x]=c_x$, or the mechanism design version in terms of the informational rent, $t_x^*=c_x+(1-F_\eta[t_x^*|x])/f_\eta[t_x^*|x]$. See Bulow and Klemperer (1996).

Alternative Designs. How would alternative designs like auctions or proxy-means testing allocate subsidies? For the household's decision, what matters is the price quoted, so selection onto the platform is similar. What changes across designs is how information is used and where the budget constraint binds.

Proxy-means testing identifies households that are poor using a poverty line measure based on observables x, and then picks a subsidization level that balances the budget in expectation. Clearly there is a trade-off between the generosity of the subsidy and how many households can receive assistance, but typically the same subsidy is offered to all eligible households and an official poverty line measure is used to determine eligibility. This approach excludes relatively rich households from participation entirely, rather than quoting relatively high prices as targeted pricing does, thereby eliminating the possibility of cross-subsidization.

Similarly, standard auctions would elicit bids from households and then distribute desludgings on the basis of the highest *ability*-to-pay, indifferent to the observables x. This ultimately benefits relatively rich households who can already afford the service by subsidizing their search for a provider and using supply-side auctions to minimize cost, but does not necessarily assist relatively poor households or those unlikely to purchase in the traditional search market. Our design combines the best elements of both commonly used systems into a superior alternative.

Platform Demand. To operationalize the platform, we convert the theoretical taxonomy of household types conditional on x into an empirical one. To collect the necessary demand side data to estimate this model, we administered a Market Survey and Demand Elicitation Game in December 2014, with 2,088 participant households selected based on their proximity to 67 randomly selected grid points from 450 grid points evenly spaced across Ouagadougou[5]. Prior to randomization, grid points falling in the wealthiest neighborhoods, neighborhoods that were connected to the sewer system, and neighborhoods in which property rights are not well-defined were omitted. Enumerators were sent to map the households closest to the grid points prior to the survey, and households were randomly selected for participation in the survey from the mapped households near the 67 selected gridpoints. We refer to this set of

[5]Another 52 were reserved for the Treatment group and 40 for the Control group.

households as the *Demand Elicitation* group. During the Market Survey, we gathered household characteristics x_i that would be available to a local government authority or NGO, as well as information on their most recent desludging, including whether they purchased mechanical or manual; the mechanical price if they purchased mechanical, $r_{mech,i}$, and the manual price if they purchased manual, $r_{man,i}$; and asked them to participate in a demand elicitation game to measure their willingnesses-to-switch.

To elicit the household's willingness-to-switch, we use a *highest-rejected (or,* $K+1$*-st price) bid auction*[6]:

i. Each household i is told it is facing N competitors, but only $K < N$ will be selected to win a desludging.

ii. Each household i is asked to make an offer, η_i, for a desludging.

iii. The highest K offers are accepted, and all winners are asked to pay the $K+1$-st (highest losing) price when they purchase a desludging.

Since honest reporting is a weakly dominant strategy in the highest rejected bid auction, the offer η_i provides an unbiased estimate of the household's willingness-to-switch, η. A histogram of the offers received and summary statistics are given in Figure 4 and Table 1.

The variables included in x_i must be observable or verifiable, or else households could profitably manipulate the information to reduce the prices they face. For example, the quality of the dwelling is a particularly good candidate for an NGO since it is costly to change and can easily be determined by an enumerator during a baseline interview. For a government authority, water and electricity bills are even better candidates, since these are objective, impose no additional data collection costs, and are highly correlated with household income. The observable variables we use include:

i. Information gathered by the enumerator during a household interview: housing type (precarious, concrete, or rooming house); whether other households lived in the compound; the distance from the latrine pit to the road in meters; the number of people living in the household; the number of women; whether the respondent finished high school

ii. Information available to a government authority: whether the water bill was

[6] The script is provided in Appendix A.

more than 5,000 CFA last month; the previous month's electricity bill in CFA; and whether the household owns the dwelling

iii. Information available to a continuously operating platform that can keep its own records: average number of months between desludgings; whether the last service episode required more than one trip because of the large size of the pit

In Section 4.2, we provide a counterfactual analysis of the information structure using subsets of these variables to determine the consequences of using less or different kinds of information: in our environment, treatment effects remain similar as coarser information sets are used, but it becomes harder to correctly predict the budget surplus or deficit that will result as more subsidies end up in the hands of relatively rich households.

With these data, we turn to estimating household-level demand curves[7]. Figure 3 shows how the theoretical cases split into switcher, non-buyer, participating buyer, and non-participating buyer cases based on their willingness-to-pay and outside option. To turn this theoretical taxonomy into an empirical model, we substitute in a discrete choice for manual or mechanical using a probit model with a latent variable y_i that depends on observables x_i and a shock. We then predict the distribution of prices that the household would have faced conditional on its observables x_i, controlling for selection into mechanical and manual.

Combining the predicted distribution of prices with the empirical distribution of household willingness-to-switch values then allows us to take the expectation of the indicator functions in the theoretical model, replacing them with the probabilities of the household making different choices conditional on the price quoted, t_i, and the household's observables, x_i. The key group are the switchers: those for whom the latent variable $y_i < 0$, so they would not purchase on their own, but for whom the platform price t_i is both attractive relative to the market price, $t_i \leq r_{mech,i}$, and less than their willingness-to-switch, $t_i \leq \eta_i$, so that $r_{mech,i} \geq \eta_i \geq t_i$.

Estimated demand is illustrated in Figure 5: panel (a) plots a demand curve for each household in the Demand Elicitation group while panel (b) plots the average curve as well as the average curves by the prices they are ultimately quoted by the platform. Similarly, panels (c) and (d) plot household demand and average demand for platform

[7] A more formal discussion is given in Appendix A.

services, respectively. Of note is how quickly demand drops off after a price of 15,000 CFA for most households on the platform, despite demand remaining relatively high at the market level, thereby illustrating how elastic demand for platform services is.

The Supply Side. With the demand side modeled and estimated, we now turn to the supply side. To procure services and determine the cost of a desludging c_x, we used *lowest rejected bid auctions* in each neighborhood:

i. Each firm j is told it is facing N competitors, but only $K<N$ will be selected to win the right to provide services in this neighborhood for this service period.

ii. Each desludger j is asked to submit an ask, a_j for each desludging they provide to the platform during this service period.

iii. The lowest K asks are accepted, and all winners receive the $K+1$-st (lowest losing) price for each desludging they provide the platform during this service period.

Whenever a household called in to claim a desludging, we randomly selected one of the K winners to receive the job, with K typically equal to two or three, so that if the first desludger was unavailable, there were other service providers who could take the job. As with the households, it is a weakly dominant strategy for firms to bid their expected marginal costs, since all of the winners receive an equal share of the work and are paid the lowest rejected bid for each job completed. These auctions were conducted monthly. The average clearing price is illustrated in Figure 6. In particular, the average price was 17,500 CFA at the time the Targeted Pricing treatment began, which we ultimately take as our average cost of procurement, $\bar{c}$.

Setting Prices. To design the optimal platform, let the set of prices be $\mathcal{P} = \{8,10,12.5,15,17.5,20\} \times 1{,}000$ CFA and q_{it} be the probability of quoting household i a particular price $t \in \mathcal{P}$. Then the designer maximizes mechanical market share

$$\max_{q=\{q_{it}\}_{i=1,t\in\mathcal{P}}^{I}} \sum_{i=1}^{I}\sum_{t\in\mathcal{P}} q_{it}\left\{\underbrace{\hat{pr}[y_i \leq 0, t \leq \eta_i | x_i]}_{\text{Switcher}} + \underbrace{\hat{pr}[y_i \geq 0, t \leq \eta_i | x_i]}_{\text{Non-participating Buyer}} + \underbrace{\hat{pr}[y_i \geq 0, t > \eta_i | x_i]}_{\text{Participating Buyer}}\right\},$$

subject to expected budget balance

$$\sum_{i=1}^{I}\sum_{t\in\mathcal{P}} q_{it}\left\{\underbrace{\hat{pr}[y_i\leq 0,t\leq\eta_i|x_i]}_{\text{Switcher}}+\underbrace{\hat{pr}[y_i\geq 0,t\leq\eta_i|x_i]}_{\text{Participating Buyer}}\right\}(t-\bar{c})+\bar{s}\geq 0$$

and the linear programming constraints that for all $i=1,2,...,I$, $\sum_{t\in\mathcal{P}} q_{it}=1$, and for all $i=1,2,...,I$ and $t\in\mathcal{P}$, $q_{it}\geq 0$. The optimal prices are presented in Figure 7, panel (a). It turns out that it is never optimal to offer 12,500: this is too high a price to induce a poorer household to switch to mechanical, and too low to generate profits that relax the budget constraint. Table 2 provides predicted treatment effects, using the model presented here: it predicts a 10% average treatment effect, and a 28.9 percentage point increase in the 10,000 CFA price bin. These numbers are too high relative to the experiment, but we assumed that all households would require a desludging within the experimental window, when only about 41% did. Weighting the effects by the realized desludging rate reduces the figures to 4.2 percentage points on average and 8.9 percentage points in the 10,000 CFA bin, which are very close to the empirical results of the RCT[8]. Similarly, budget balance predictions from this model were largely satisfied: the realized average loss per household was 102 CFA ($0.18), against a 1,650 CFA ($3.00) subsidy.

Deployment on a new sample. The solution to the linear program is in terms of the households in the original sample, $\{(t_i,x_i)\}_{i=1}^{N}$, but we now wish to test the optimality of the platform design on a new sample of households, $X'=\{x_{i'}\}_{i'=1}^{N'}$. This requires a pricing rule that maps observables x_i to prices $t^*(x_i)$. To do this, we use a simple ordered logit model[9], to predict how a new household with characteristics $x_{i'}$ would have been assigned to prices $t_{i'}$ in the original sample, where households with higher latent index values are assigned to higher price bins. Since the original sample, $X=\{x_i\}_{i=1}^{N}$, is large and random, the platform can replace the personalized prices for each household t_i^* with a function that maps characteristics x_i into prices, $t_i^*=t^*(x_i)$, and the same pricing rule should also maximize adoption of mechanical

[8]In Table 3, the estimated treatment effects are 0.042 on average and 0.097 effect on 10,000 CFA households.

[9]We consider alternative, more algorithmic and automatic methods of assigning observables to prices in Section 4 particularly random forests.

desludging across the population and balance the budget in expectation.

Scale up. Before proceeding to the randomized controlled trial, a key question is whether or how this process might be done at scale. Of particular concern is that a household who anticipates that information it provides about η_i and the market prices ($r_{mech,i}$,$r_{man,i}$) will determine its future payoffs has an incentive to lie. There is a straightforward way to avoid this. For a given period of the program, a small, randomly selected subset of a large population can be drawn and offered the chance to participate in the market survey and game, just like the Demand Elicitation group. This period's market will be designed using information from this subset, but the subset will not be subject to the terms of the program this period, making their payoffs independent of their reports. This simple adjustment maintains incentives for honest reporting. Indeed this is how, say, an educational organization could conduct market research and design its quoted tuition price and determine scholarships to maximize the talent of its pool of matriculating students.

3 Experimental Results

We now seek to test the empirical performance of the platform. Our intervention is to offer the pricing rule $t = t^*(x_{i'})$ to a new random sample, $X' = \{x_{i'}\}_{i'=1}^{N'}$, called the *Targeted Pricing* group, and compare their outcomes with a third random sample who are excluded from the platform, the *Control* group. This section presents the results of the RCT.

The Data. We returned to Ouagadougou in August and September of 2015 in order to test the Targeted Prices treatment designed in Section 2.[10] We offered the pricing treatment during a baseline survey which included 1,660 pricing treatment households and 1,284 control group households in addition to gathering information about each household's latrine pit, sanitation practices, search process for a desludging operator, and their level of wealth. Summary statistics and balance tests are provided in Appendix B. Participation in the platform was offered to the treatment households at the end of the baseline survey. The enumerators' tablet computers used for the

[10]This trial was pre-registered, AEARCTR-0000834.

survey calculated the platform price that the household should be offered according to the pricing algorithm. If a treatment household wished to accept this price, they were asked to pay a deposit of 500 CFA (which was the same as the participation payment for the survey). They were then able to call in to the center at any time in the following 15 months in order to claim their desludging at the targeted price offered.

We returned to the households interviewed in both the demand elicitation survey and the baseline survey in December of 2016 to conduct an endline survey. Households were asked about any desludgings that they had done over the period and the diarrhea-related health of their children. At endline we also asked households that deposited but failed to call the call center why they had not used their fixed-price desludging: more than half stated that they had not needed a desludging during the study period; many of the others had forgotten about it or found a better outside option.

Deposit rates by price offered and use of the call center are shown in Appendix Table 18. Use of the call center was somewhat lower than predicted, but among those who purchased a desludging in the first 6 months and deposited, use of the call center was quite close to the level expected from the model (and somewhat higher among the 20k price group).

3.1 Impact on Desludging Choice

Outcome Measures. For sanitation goods in general, social benefits accrue to the neighborhood when a household chooses an improved service over the unimproved, traditional service; in this case mechanical over manual desludging. The objective is therefore not to increase the overall number of desludgings — which depends on biological processes that determine the rate of fill of the tank — but to convince households to switch from manual desludging to mechanical. This is in contrast to many previously studied settings where the objective is to increase the overall level of purchases of a new or under-utilized good (for example, water purifiers, chlorine tablets or insecticide-treated mosquito nets). To capture this impact on mechanical utilization at both the neighborhood and household levels, we use three measures of impact: the market share of mechanical services in the neighborhood, the percent of household purchases which are mechanical (share at the household level), and whether a household purchased any mechanical desludgings over the period of the study.[11] We discuss each in turn below.

First, given that the externalities from manual desludging are a local phenomenon that impacts nearby households, our primary measure of the success of the program is based on the amount of switching between manual and mechanical desludgings at the neighborhood level: the *market share of mechanical services*. Market share[12] is defined as

$$Share_n = \frac{Mechanical_n}{Mechanical_n + Manual_n} \tag{1}$$

where $Mechanical_n$ and $Manual_n$ are the numbers of mechanical and manual desludgings done in neighborhood n, for each of 92 neighborhoods of 25-40 households during the intervention period. Each household that switches from manual to mechanical represents a reduction in fecal sludge in the environment, so $Share_n$ is the best measure of the impact of the intervention at a neighborhood level, where the effects of the negative externalities are largest. Since effects may differ across income levels and price groups, we also calculate the shares at the neighborhood-price group level, adding the total number of mechanical desludgings done by households within a price group (10k, 15k, 17.5k or 20k) within a neighborhood and dividing by total number of desludgings done by households in the price group-neighborhood.

In order to observe the effects at the household level, we next calculate the share at the household level ($PctMechanical_i$) by substituting the neighborhood index n for a household index i, providing us with the change in a household's propensity to substitute between mechanical and manual.

We then provide results for whether a household i purchased *any mechanical* desludgings:

$$AnyMechanical_i = \mathbb{I}\{Mechanical_i > 0\}, \tag{2}$$

for each of the 1,199 households during the intervention period who purchased at least one desludging. Mechanical desludgings were the product for which the subsidy

[11] Whether the household purchased any manual desludgings is close in magnitude and opposite in sign, but not statistically significant in most specifications. A comparison of the manual and mechanical effects are provided in Appendix C

[12] Market share is a common outcome variable in papers estimating market effects, particularly when estimating the coverage of a certain product (see, for example, Jensen and Miller (2018) or Nevo (2001)).

was received, but the subsidy was valid only for one desludging, so we would expect the treatment to directly impact the first desludging the household needed during the 15 month period.

Regression Specifications. For each of these utilization measures $Share_n$, $PctMechanical_i$, $AnyMechanical_i$, we provide two sets of regression results: pooled results across the full sample and disaggregated results separating impacts by estimated treatment price group. Impacts on market share $Share_n$ are estimated at the neighborhood level, while impacts on $PctMechanical_i$ and $AnyMechanical_i$ are calculated at the household level. We also estimate the effect on health at the household level ($Diarrhea_i$).

First, we estimate the overall impact of the call center using a pooled average treatment effect β with the specification

$$y_i = \alpha + \beta TargetedPricesTreatment_i + \gamma' X_i + \varepsilon_i, \tag{3}$$

where $TargetedPricesTreatment_i$ takes the value 1 if i received the Targeted Prices treatment and 0 if it is in the Control group, X_i is a vector of control variables including the variables not balanced at baseline, the stratification variable, and the baseline values of past desludging decision variables,[13] and ε_i is a disturbance, clustered at the neighborhood level for $PctMechanical_i$ and $AnyMechanical_i$.[14]

Second, we estimate the effect of the call center on mechanical utilization by each price bin k taking values in the set $\mathcal{P} = \{10000, 15000, 17500, 20000\}$ with the

[13]The control variables include: the distance from the latrine pit to the road, whether more than one trip was required to desludge the last time a desludging was done, whether the water bill cost more than 5,000 CFA (approximately $9), whether the household has more than one pit, whether there are unrelated households living in the compound, the stratification variable–whether the neighborhood had an above-median number of low walls, and the controls for desludging behavior prior to the baseline for ANCOVA specification (McKenzie, 2012)–the percent of desludgings that the household recalls at baseline that were mechanical.

[14]In the market share regressions, the specification of interest is:

$$y_n = \alpha + \beta TargetedPricesTreatment_n + \gamma' X_n + \varepsilon_n, \tag{4}$$

the targeted prices treatment indicator $TargetedPricesTreatment_n$ takes the value of 1 if the neighborhood n is one of the neighborhoods randomly selected for treatment. Control variables are averaged at the neighborhood level.

specification

$$y_i = \sum_{k \in \mathcal{P}} \beta_k TargetedPricesTreatment_i \times PriceGroup_{ki} + \sum_{k \in \mathcal{P}} \alpha_k PriceGroup_{ki} + \gamma' X_i + \varepsilon_i \quad (5)$$

where $PriceGroup_{ki}$ takes the value 1 if i is in the Targeted Prices treatment and is offered a price of k, or if i is in the Control group and *would have been offered* a price of k. Each coefficient β_k measures the average difference in utilization between observational units in the Targeted Price group who received the price k quote versus those in the control group who would have received the price k quote, conditional on the controls.

In the neighborhood level regressions, the dependent variable is the market share for a price group within a neighborhood cluster: the market share is calculated as the number of mechanical desludgings purchased by households of that price group in that neighborhood (k equals 10000, 15000, 17500, or 20000) divided by the total number of desludgings purchased by households of that price group in that neighborhood. In the household level regressions, the dependent variable is percent mechanical or an indicator for whether the household purchased any manual desludgings. We omit the constant in order to include indicator variables for each price group.

Health impacts are estimated with similar regression specifications (equations 3 and 5) at the household level, with an indicator for any child in the household was reported as having had diarrhea in the past week. Health regressions include all households with children in the sample and control for baseline indicators of whether a child was reported to have diarrhea in the household in the past week. Desludging use among other households in the neighborhood may impact a household's health status as much as their own use, so for the health regressions we add a specification that estimates differential impacts by the percentage of households in the neighborhood that qualifies for each price level.

$$y_i = \sum_{k \in \mathcal{P}} \beta_k TargetedPricesTreatment_i \times PctNbhd{-}PriceGroup_{ki} + \sum_{k \in \mathcal{P}} \alpha_k PctNbhd{-}PriceGroup_{ki} + \gamma' X_i + \varepsilon_i \quad (6)$$

Estimation Methods. We estimate an OLS model with controls for unbalanced variables at baseline. We also exploit the panel nature of our data by estimating a household fixed-effects model. This model controls for all time-invariant differences across households, thereby controlling for most pre-existing differences between the households. It is of note that the desludging variables are somewhat time-sensitive: households that use manual desludging tend to desludge less frequently overall, so at endline we expect to see both an increase in the percentage of desludgings that are mechanical and an increase in the households using at least one mechanical desludging since the time frame between baseline and endline is restricted relative to the time frame at baseline. This is, however, true of both the treatment and control groups, and the estimate is based on the *differential* increase in mechanical desludging between baseline and endline for the treatment group relative to the control group, so the estimate is still the unbiased effect of being in the treatment group on the outcome variable of interest.

Main Results. Estimates are presented in Table 3 for $Share_n$, $PctMechanical_i$ and $AnyMechanical_i$. The qualitative patterns are consistent across the specifications: there is a modest average treatment effect,[15] but this effect is driven by much larger effects in the lowest priced 10,000 CFA bin, where potential switchers are concentrated. This pattern shows that the Targeted Prices intervention is working as intended by providing aid to the poor households who are unlikely to purchase mechanical on their own. Control group households that would have been quoted

[15] Households in the control group used more desludgings total over the course of the study than households in the treatment group–on average control group households did 0.92 desludgings per household while households in the treatment group did 0.8 desludgings per household, which biases our estimates toward zero since households which desludge more often use mechanical at a higher rate than manual. In part this difference in use of desludgings is a result of a balance issue between the control group and the treatment group. We include controls for variables unbalanced at baseline in order to address this and we include estimates with fixed effects by household. We discuss this further in Appendix B.

relatively high prices are already purchasing mechanical desludgings at high rates, so there is little scope for the intervention to increase utilization of the healthy service: 99% of those households in the control group who would have been quoted a price of 20,000 CFA purchased mechanical services anyway (91% for 17,500 CFA, and 85% for 15,000 CFA). The majority of the potential switchers are in the 10,000 CFA bin, where control group mechanical utilization is only 59%.

The Targeted Pricing treatment generates an increase in the mechanical market share of 5.3 percentage points in the pooled treatment (column 2), but an increase of 12.7 percentage points among the households in the lowest priced treatment bin (column 3). Market share for the 15k price bin increases by 8.8 percentage points and does not have a statistically significant increase in the higher price bins. The p-value that the effect is the same in the 10k bin as the 17.5k bin is 0.04, and the p-value that the effect is the same between the 10k bin and the 20k bin is 0.09, rejecting the null that the effects are the same between groups.

Estimated at the household level, the results are similar. The pooled effect across all price bins is 3.7-6.4 percentage points (columns 4 and 5), but the impact among the lowest price bins is in most cases largest, though not statistically significantly different from the other groups at the household level. Because we subsidized only the first desludging, we may think that the results for the first desludging used by the household are most important. Here we find that the results are similar–overall treatment increases the probability that a household will buy at least one mechanical over the period by 4.0-5.2 percentage points (columns 8 and 9), while the lowest price treatment group has an increase in the probability of purchase of a mechanical desludging of 5.7-10.1 percentage points (columns 10 and 11).

Who Receives Subsidies? Did the targeting strategy work? Did relatively poor households receive lower prices? We can observe the extent to which the model is targeting relatively poor households who are more likely to get manual desludgings in Table 5, where *non*-targeted measures of affluence are shown to increase along with the price quote. For example, households that receive a price quote of 10,000 CFA spend an average of 2,200 CFA per week on phone credit, while households that receive a price of 17,500 CFA or 20,000 CFA spend 4,512 CFA and 5,631 CFA on average, respectively: more than twice as much. The same pattern repeats for refrigerators, motorcycles, cars, televisions, mobile phones, and air conditioners. Taken together, this evidence from non-targeted variables implies that the intervention did target poor households for assistance, and was effective in determining the right amount of aid to provide in order to induce switching.

Indirect Mechanisms. One potential effect of the platform system is that by giving households an additional outside option, we improved their bargaining power against desludging operators. Treatment households who purchased a desludging outside of the system paid about $2 less than the prices paid by control households on average. Households may have stayed with the desludging operator that they knew, but they were able to purchase at better prices and avoid switching to manual desludging.

3.2 Health Impacts

The primary goal of the program was to reduce the use of manual desludging in order to improve local health and sanitation conditions. We test the impact of the targeted prices treatment on the rates of child diarrhea. We focus on children for two reasons. First, severe diarrhea in children may have important long-term consequences, so the welfare impact of health problems arising from poor sanitation is largest among children. Second, children are more sensitive to poor sanitation conditions than adults; if the intervention improves the health of children, we can infer that adults likely also benefited from the intervention.

Estimation proceeds as in (3) for the pooled effect and (5) for the separate price bins. Children contract illnesses based on the sanitation conditions of the neighborhood, not just the sanitation decisions made by the household. This motivates the inclusion of an additional analysis of the spillover effects of sanitation decisions within the neighborhood based on the percent of households in the neighborhood that fall into each of the price groups. We use the following specification to estimate the differential effect of the treatment in neighborhoods with larger numbers of each of the price groups, with price k taking values in the set $\mathcal{P}=\{10000,15000,17500,20000\}$:

$$AnyChildrenDiarrhea_i=\sum_{k\in\mathcal{P}}\alpha_k PriceGroup_{ki}+\sum_{k\in\mathcal{P}}\Theta_k PctPriceGroup_{ki}$$
$$+\sum_{k\in\mathcal{P}}\beta_k TargetedPricesTreatment_i * PctPriceGroup_{ki}+\gamma' X_i+\varepsilon_i, \quad (7)$$

we use the same control variables and clustering as in the main regression.

Observations are at the household level and standard errors are clustered at the neighborhood-cluster level (92 clusters in total), and the vector of controls X_i includes the same variables as in our main specifications: variables unbalanced across neighborhoods at baseline, the stratification variable, and a control for whether the household had a child suffering from diarrhea at baseline following the ANCOVA specification in McKenzie (2012) to increase efficiency.

When separated by price bin, the treatment effect on health is largest among households that qualify for the lowest prices. Diarrhea rates are higher among children in these households–17 percent of households qualifying for the lowest price reported their children having had diarrhea in the past week in the control group, as compared to the overall mean rate of diarrhea among children of 13 percent. Results are shown in Table 4 columns (3) and (4). Among the poorest households there is a 5.2 to 7.0 percentage point reduction in diarrhea. This is a 40 percent reduction in the incidence of diarrhea among the poorest households relative to the control group at the mean. The average impact estimated across the sample is somewhat lower, as it includes wealthier households which already had a lower baseline diarrhea rate. Columns (1) and (2) provide the results for the pooled specification. We are under-powered to find an effect in the pooled regression, but the point estimate on children's diarrhea overall is a 1.2-3.2 percentage point decrease. At the mean of 13% of households reporting that at least one of their children had diarrhea in the last week at baseline, this would be a 9.2 percent effect.

Neighborhood sanitation is strongly impacted by local subsidies: in columns (5) and (6) we estimate that if 10% more of the neighborhood is in the 10k group, treatment reduces diarrhea rates by an additional 2.2-2.5 percentage points. This group of targeted households is more likely to use manual desludging in the absence of a subsidy, and therefore neighborhoods with more of these households are less healthy. Among households in neighborhoods with an above median percentage (27%) of households qualifying for the lowest price group 17.8 percent report that at least one child in their household had diarrhea over the past week while the rate is 13 percent overall. Providing the poorest households with the 10k price reduces diarrhea rates substantially in neighborhoods with more households qualifying for the lower price: columns (5) and (6) show that if 10% more of the neighborhood is in the 10k group, treatment reduces diarrhea rates by an additional 2.2-2.5 percentage points.

Comparison with Epidemiological Literature. We compare the effects found in this paper to those reported in Fewtrell et al. (2005), a large epidemiology meta-study of the impacts of water and sanitation interventions on diarrhea rates in children. They compare relative risks of falling ill with a specified disease for the treatment group versus the control group for all four available sanitation studies and report an average relative risk ratio[16] following sanitation treatments of 0.68. The relative risk ratio for our pooled sample is 0.82, but when the sample is constrained to the households which would receive a price of 10,000, the relative risk ratio for this group is 0.60. This is a large effect, but makes sense: restoring a toilet that would otherwise likely remain broken to working order and avoiding an instance of manual desludging has comparable impacts to providing poor households with toilets.

4 Counterfactual Platform Designs

While the intervention in Section 3 demonstrates that the Targeted Pricing treatment had a variety of statistically significant impacts on behavioral and health outcomes, at least three questions remain: How would other, alternative market designs have fared? If different observable design variables had been used, what would the consequences have been? How do treatment effects respond to the subsidy level, $\bar{s}$? This section uses a model and all the sources of available data to address these questions.

In the intervention, household i was quoted a price t_i; household i then decided whether or not to deposit, $Deposit_i \in \{0,1\}$, and whether or not to arrange a desludging through the platform, $Arrange_i \in \{0,1\}$; which resulted in the percentage

[16] The relative risk ratio is the ratio of the rates in the control and treatment groups: $\frac{Outcome_{Treated}}{Outcome_{Control}}$, and numbers closer to 0 suggest larger impacts.

of mechanical desludgings it purchased given that the household did or did not deposit, $m_i^1 \in [0,1]$ or $m_i^0 \in [0,1]$, respectively. To measure the overall impact of the platform and its financial prospects under alternative designs, we want to predict what would have happened at new prices $\{t_i'\}_{i=1}^I$ and a potentially different subsidy level s' with cost of service c_i, and compute for each household i in the Targeted Pricing group:

- *Expected mechanical share*[17]: $\text{Share}_i' = \text{Deposit}_i' \times m_i^{1'} + (1 - \text{Deposit}_i') \times m_i^{0'}$.
- *Profit*: $\text{Profit}_i' = \text{Arrange}_i' \times (t_i' - c_i)$
- *Budget balance*: $\text{BB}_i' = \text{Arrange}_i' \times (t_i' - c_i + s')$
- *Subsidization rate*: $\text{SR}_i' = \text{Arrange}_i' \times \dfrac{(c_i - t_i')}{c_i}$

We then average over i to compute expected outcomes across the whole sample. Expected share, profit, and budget balance are all self-explanatory, but the subsidization rate is a "reverse Lerner index" measuring the proportion of the cost that is covered by the platform rather than the household: positive values correspond to profit losses to the platform on the sale, while negative values correspond to profits. The subsidization rate is particularly effective at identifying how different designs achieve their effects, since it provides a dimensionless measure of the assistance received by each household.

We estimate[18] the $Deposit_i$ as a probit discrete choice, followed by two linear probability model regressions for m_i^0 and m_i^1 using the semi-parametric method described in Powell (1994) and Newey et al. (1990) for controlling for selection, with the enumerator's subjective assessment of the reliability of the household's responses as an instrument. The $Arrange_i$ model is estimated as a simple probit discrete choice, since there are no unobserved product characteristics and the exact pricing rule is known.

One notable change from the original platform design is that, rather than using the ordered logit model to assign prices to design variables, we use the more automatic random forest algorithm for classification when necessary[19].

[17] This linear probability model rarely predicts values below 0 which are rounded up to zero, but does predicts values above 1. These are rounded to 0 and 1, respectively, in simulations. In practice, the Targeted Pricing statistics would disproportionately benefit at the expense of the other designs from not rounding these values.

[18] Appendix E provides the model and estimation in more detail.

[19] Neither the ordered logit nor the random forest were uniformly better in terms of fit across all exercises (Figure 12 graphs all three for comparison), but the random forest algorithm is automatic and more robust, which is especially important when bootstrapping.

We present results for the pooled Targeted Pricing group as well as by each pricing bin, which relies on the endogenous assignment of households to prices. The results in Table 5 from Section 3 show that the Targeted Pricing bins accurately target relatively poor households based on non-targeted measures of wealth. We therefore maintain comparisons across these bins in order to have a parsimonious comparison of the treatment effects across different wealth levels and market designs.

4.1 Alternative Market Designs

While the randomized controlled trial provides causal estimates of the treatment effect of access to the platform for households in the Targeted Pricing group, it does not address the question of how alternative designs would have performed. This section addresses this question and illustrates the channels through which Targeted Pricing operates. We consider three alternative designs:

i. *Procurement Auctions*: the clearing prices in the procurement auctions are passed directly to households who call in from those neighborhoods, who then either accept or reject.

ii. *Proxy-means testing* (PMT): We predict the income per household member in the demand elicitation group using a LASSO regression with 113 potential control variables[20], and use the results to predict the income per household member of the targeted price treatment households. We then create an indicator variable for whether the household per capita income falls above or below Burkina Faso's urban central region per capita poverty line, and 150 percent of the poverty line. We quote households classified as poor with the auction clearing price in their neighborhood minus the subsidy.

iii. *Market Average*: We offer the average price available in the decentralized market, 16,833 CFA, to all households, as if a municipal authority were able to fix the price and sufficiently penalize firms that deviated from it.

The proxy-means testing counterfactual is constructed assuming the same subsidization level as our experiment, \$3.00. We also provide results for *Subsidized Procurement Auctions* and a *Subsidized Market Average*, where the clearing price or average price

[20] Variables directly related to income were removed, which is why there are fewer variables included here than in the main regressions.

are reduced by the same subsidy level, to provide a more fair comparison to Targeted Pricing. The Market Average can be thought of as a price stabilization or simple market centralization policy that requires all service providers to offer the current price to any customer.

Table 6 provides counterfactual mechanical shares[21], and Figure 8 provides a visualization of the main results. The average mechanical share for all the designs is between 78 and 80 percent, but each design works differently. The treatment effect on the 10,000 CFA group is 10.6 percentage points for Targeted Pricing, 4.2 for PMT, 2.2 for the Auction, and 1.5 for the Market Average. So Targeted Pricing achieves its average effect by increasing the mechanical share among the poorest households, while the other designs achieve their effects by inducing switching in relatively wealthier groups.

How does Targeted Pricing outperform the other counterfactual designs, on average and particularly among households in the 10,000 CFA price bin[22]? Figure 10 illustrates subsidization rates for the households in the Targeted Pricing group and Table 8 provides the averages. The Targeted Pricing treatment achieves average subsidization rates of around 11.8% for the 10,000 CFA group, 1.3% for the 15,000 CFA group, -1.2% for the 17,500 CFA group, and -3.1% for the 20,000 CFA group. In contrast, no other treatment achieves a subsidization rate larger than 2.0% for the 10,000 CFA group — or indeed, any household — and the only other design that achieves negative (profitable) subsidization rates is the Market Average, where a small number of households are purchasing at a relatively high price.

Budget Balance. One objection might be that Targeted Pricing achieves these results by running larger budget deficits, but this is not the case. Budget balance calculations are reported in Table 7. Surprisingly, Targeted Pricing only loses about 100 CFA on average per household, despite delivering an average transfer of 1,334 CFA

[21] These results are comparable to the more pessimistic measure of impact from the RCT, $PctMechanical_i$, in Table 3.

[22] For the reader that disagrees with the use of the pricing bins as a way of "cutting the data" for this analysis, there are two things to note. First, the intervention is designed to maximize overall take-up and not just the impact on the 10,000 CFA bin, and within that bin, the baseline utilization of mechanical desludging is about 40%. This suggests that even if other partitions based on measures of poverty or utilization were used, the results would be similar. Second, no other treatment achieves subsidization rates on the order of 10% for *any households* — see Figure 10 — so that even if other subgroups were of particular interest, the other designs considered would not achieve similar impacts on them.

to households in the 10,000 CFA pricing group, far more than any other design. The Auctions and Proxy-Means have average loses of zero by construction[23], the Market Average raises 78 CFA per household, and Targeted Pricing loses 100 CFA per household ($.18) on average. While $.18 might not be an insignificant amount at scale, it is also not so large that the differences in performance between Targeted Pricing and the other designs can be explained solely by violations of the expected budget constraint.

Design Variations. An alternative and perhaps more direct way of comparing designs is to answer the question, "How large a subsidy is required for Proxy-Means Testing, Auctions, and the Market Average to achieve the same performance as Targeted Pricing?" We call this quantity the *design variation*, due to its similarities with the compensating or equivalent variation in consumer theory. Table 9 provides the design variations for Auctions, Proxy-Means Testing, and the Market Average, on average and for the 10,000 CFA group. To achieve the same average treatment effect, an Auction requires an additional subsidy of 1,704 CFA, Proxy-Means Testing requires 1,643 CFA at the 100% level and 350 CFA at the 150% level, and the Market Average requires 2,478 CFA. In order to match the average treatment effect on the 10,000 CFA group, the numbers are much larger: 6,132 CFA for Auctions, 5,388 CFA for Proxy-Means at 100% of the poverty line and 4,518 at 150% of the poverty line, and 6,825 for the Market Average, all values in excess of $10. This, again, illustrates the importance of targeting and cross-subsidization in delivering aid effectively.

Comparison with PMT. Why does Targeted Pricing consistently outperform Proxy-Means Testing, in particular? Proxy-Means Testing takes a training data set and attempts to predict poverty using wealth, then subsidize those who are diagnosed as poor using simple poverty line measures. There are two drawbacks:

[23] These calculations for the alternative designs are based on *ex post* costs, so the subsidy plus the price paid minus the cost equal zero. For the Targeted Pricing treatment, the costs were unknown at the time of contracting, so the budget does not balance exactly. Using the realized costs here advantages the counterfactual designs. We do this in order to make the alternative designs as competitive as possible, rather than speculate on what a designer might have predicted would happen at the time of design. In practice, predicting the exact amount that would be spent on Proxy Means or Subsidized Auctions would be as difficult as predicting the cost of the Targeted Pricing program, and would also run a deficit or surplus in practice.

first, being above or below the poverty line does not necessarily correlate strongly with the desludging decision, and, second, the PMT design does not predict whether a particular household will actually purchase the product conditional on a given level of subsidy aid; indeed, the most impoverished households might be unable to purchase, even given very large levels of subsidization. Our approach addresses both issues by modeling the desludging decision in the absence of the intervention and then strategically distributing aid to maximize take-up, resulting in better targeting and better redemption by targeted households.

4.2 Information Structures

A key choice in the design of the platform is selection of observables used to estimate household level demand curves. These variables must be observable or verifiable, so that households cannot manipulate them to their advantage. The version of the model which we took to the field used a variety of variables based on enumerators' subjective assessments, municipal records, and platform records[24]. Now, we wish to consider more conservative versions of this model, giving special attention to the case of an NGO that does not have access to municipal records and a municipal authority that lacks a budget for household surveying.

Figure 11 succinctly illustrates the relationships between the original and alternative information structures, and Figure 12 illustrates the optimal Linear Programming, Ordered Logit, and Random Forest rules side-by-side. The Conservative version removes some of the more difficult to gather variables. The Municipal and Municipal 2 versions rely on public and billing records, foregoing the information that an enumerator might gather during a visit to the household. The NGO and NGO 2 structures drop public and billing records in favor of data observable during a

[24]Recall from page 12: information gathered by the enumerator during a household interview including housing type (precarious, concrete, or rooming house), whether other households lived in the compound, the pit's distance to the road in meters, the number of people living in the household, the number of women living in the household, and whether the respondent finished high school; information available to a municipal authority, including last month's water and electricity bills, whether the household owns the dwelling; and information available to a continuously operating platform that can keep its own records, including average months between desludgings, whether the last service episode required more than one trip because of the large size of the pit. Instruments included electricity expenditure, the number of people in the household, the number of women in the household, and whether or not the respondent completed high school.

household visit. The NGO 2 version, in particular, uses no instruments, identifying the demand model entirely off of the structural assumption of normality of the shocks.

Counterfactual shares are presented in Table 10 and visualized in Figure 8 (b). Note that, compared with the ordered logit version (80.8 on average and 68.6 in the 10,000 CFA bin), the predicted treatment effects are slightly lower using the random forest rule (80.6 and 66.3). On average, the Original and Conservative information structures achieve the same market share of 80.6, and the Conservative version exhibits a decline of only .04 percentage points in the 10,000 CFA bin.

The Municipal information structures exhibit higher average market shares at 81.5 and 82.1, respectively, compared to the original version, which achieves an 80.8 mechanical share, but lower values in the 10,000 CFA bin, at 66 and 65.6, respectively, compared to 66.3. This can be attributed to larger violations of the budget balance constraint, given in Table 11: Municipal and Municipal 2 lose 177 and 306 CFA per household on average compared with 24 CFA in the Original structure; the loss in the 10,000 CFA bin is 894 CFA for the Original structure, but only 808 and 775 CFA for the Municipal structures, respectively. This illustrates that as information is restricted, effective targeting becomes more difficult, leading to interventions that can be more generous than intended and target the "wrong" households. The NGO information structure exhibits a similar pattern to the Municipal ones, with a comparable average market share of 80.7 versus 80.6, but a smaller 10,000 CFA bin share of 64.6 versus 66.3.

For the second NGO structure — the one information structure without any instruments in the Tobit V estimation, so that the model is identified off the structural assumption of normality — the results are markedly worse. The average and 10,000 CFA bin market shares are indistinguishable from Predicted Control at 77.6 and 63.3, respectively. The use of instruments in the Tobit V ensures that estimation is robust to violations of normality, and here the pricing rule inherits this robustness.

One by-product of the construction of a random forest pricing rule is a measure of variable importance in translating household specific prices into a mapping from observables to prices[25], given in Table 12. Across versions, the most informative

[25] A rough explanation of this measure for those unfamiliar with random forests is: A decision tree is a sequence of binary decision rules — e.g., "is the electricity bill greater or less than 16,000 CFA?", "is the distance from the pit to the road greater or less than 8 meters?" — that terminates by assigning a price prediction to any observables. A random forest is a collection of decision trees

variables are electricity bill, desludging frequency, and housing type dummies, which are all either observable on a visit from an enumerator or included in municipal records. Other variables tend to become relatively more important as the information sets become more restricted, but the most informative information happened to be the easiest to gather. Overall, this analysis provides one explanation for why the mapping from design variables to prices is fairly robust across information structures: some strongly predictive information about price and behavior goes a long way in accurately assessing the needs of households.

4.3 Sustainability

If a design requires continual financial assistance to operate, it may not be sustainable in the long run. This section studies the trade-off between the level of subsidization required and the treatment effect that can be achieved.

We vary the subsidy per household $\bar{s}$ from a profitable -750 CFA (-\$1.36) to an unprofitable 13,000 CFA (\$23.64), solving for the prices that would have been quoted, and then using the counterfactual model to predict what the treatment effects would have been[26]. This exploits the duality between profit maximization subject to a quantity constraint and quantity maximization subject to a budget balance condition to construct estimates of the trade-off between subsidization and treatment effect. In addition to using a random forest algorithm to map observables to prices, we (a) use a cost of 13,750 CFA, which is the average of the lowest costs we achieved across neighborhoods at the time the project concluded and the desludgers had started to compete, rather than 17,500 CFA, which is the average clearing price at the time the Targeted Pricing intervention began and the desludgers were coordinating on high prices, and (b) allow greater freedom in platform design by expanding the set

constructed on bootstrapped samples using different, random selections of explanatory variables that assigns the modal price prediction made by all of its trees. When adding a new binary rule to the tree, the algorithm scans the tree and finds the addition that minimizes the heterogeneity in final predictions as measured by the Gini coefficient. The variable importance measure of a variable takes the average reduction in the Gini coefficient across all of the decision trees in the forest in which that variable appears. With a large number of trees, this is a measure of the reduction in misclassification attributable to the variable.

[26]The pricing model predicts that at any $\bar{s}$ less than -750 CFA, the linear programming problem (23) — (25) has no feasible points, and therefore -750 CFA corresponds to the profit-maximizing average profit per household.

of prices quotes to $\mathcal{P}'=\{0,5000,10000,12500,15000,17500,20000,22500\}$. This gives a better estimate of what an optimally designed platform could achieve, given the information learned through the intervention.

Panel (a) of Figure 9 shows how the average Mechanical share and the share for the 10,000 CFA price bin households vary with the subsidy. Even at a profitable subsidy value of -750 CFA, there is a 3.3 percentage point increase, and 5.2 percentage points in the 10,000 CFA price group. At a subsidy value of 0 CFA, there is a 4.0 percentage point increase in mechanical share and 8.0 points in the 10,000 CFA price group. This shows that access to a centralized market can improve welfare even without subsidization, due to supplier competition and cross-subsidization. As the subsidy increases to approximately 4500 CFA, the impact on the 10,000 CFA households flattens out, increasing again at a much higher subsidy level of 10,000 CFA. The reason for this is illustrated in Panel (b) of Figure 9: the budget constraint prevents the platform from offering the 0 price offer until it has a substantial amount of funds available at approximately 60% of the average price of the mechanical service at 10,000 CFA, and the mechanical share begins increasing again. Once it does, the platform begins offering free desludgings, and the treatment effect on the poorest group begins increasing again; indeed, kinks in panel (a) correspond to subsidization levels in panel (b) at which the set of prices offered changes. This is another way to illustrate how important it is to distinguish between marginal households who are relatively cheap to switch from manual to mechanical, and those households who are so poor they will require extensive assistance.

5 Conclusion

Increasing the take-up of health and sanitation goods may require large subsidies, and the use of targeting on observables and cross-subsidization can stretch budgets further to foster take-up among vulnerable populations. We formalize this problem using a mechanism design approach where households' private information about their willingness-to-pay and their outside option are private information, solve for the optimal mechanism, gather the data necessary to implement it in a market for sanitation services, and test the impact of the solution using a randomized controlled trial. In contrast to much of the current literature on increasing take-up of health and sanitation goods, this is accomplished in the presence of a prevailing decentralized market. Consumers can opt out of purchasing through our platform, as is common in developing countries (for example, water purification tablets and basic mosquito nets are often available in local markets).

The existing research on sanitation shows that it is difficult to have significant impacts on health only through behavioral approaches to changing the sanitation choices made by households. The evidence on encouragement campaigns and shaming is mixed, and CLTS campaigns have been found to have little impact except for the most intensive programs (see, for example, Gertler et al. (2015)). For comparison, in a randomized controlled trial of water and sanitation and nutrition programs of over 8,000 households, Null et al. (2018) find no impact of the interventions on diarrhea in children, and small impacts in year two on height only for the treatment arm which included sanitation, hand washing, and nutrition. McIntosh and Zeitlin (2021) find that a USAID program aimed at nutrition and water and sanitation and costing $142 per household had little impact on health indicators and served only to somewhat increase savings levels in the households. Many health and sanitation campaigns have been focused on rural areas where open defecation is common, yet because of the high population density in urban areas, the externalities from improper disposal of fecal waste may be much larger (**?**).

Rather than relying on behavioral change interventions and information treatments, we focus on improving take up through pricing policies. Our approach to platform design and pricing is different from the standard literature on demand estimation, because we are able to adopt an experimental approach. A standard structural model of household demand would be sensitive to modeling assumptions about how the household searches for service operators, how negotiations proceed, and how the levels and differences between mechanical and manual prices determine household utility. We adopt a simpler approach: ask households their willingness-to-switch, η, in an incentive compatible game similar to a second-price auction, and exploit its correlation with past decisions, prices, and observables to estimate a household's probability of purchasing on the platform or the search market. Particularly because counterfactual prices for manual or mechanical are not typically observed for households conditional on their choice, the use of the Tobit model and willingness-to-switch data allows the platform to explore the shape of the demand curve at prices that are not observed in the market, which is a serious challenge for policymakers when offering novel goods or assistance to poor populations in a way that, say, multinomial logit does not[27].

There are many potential improvements to this overall approach to modeling and estimating household behavior. Picking observable variables to target is an important decision, and they must be available or easily observable to a local governmental entity or NGO, since otherwise the resulting mechanism would not be incentive compatible: households would generally be able to guess how to misrepresent themselves as poorer than they are. Ideally, variables used would be highly correlated with wealth and sanitation decision-making, be cheap to gather, and robust to measurement error. In principle, machine learning methods like penalized regression can be adapted to the log-likelihood of the Tobit model, (19)–(22), to avoid over-fitting and reduce prediction error. Bootstrapping methods can be used to simulate the distribution of expected losses. With the rapid advances in data-gathering methods and machine learning tools over recent years, we expect there are variety of other improvements that could be made.

While this paper focuses on a platform operated by a local government or NGO, the general methodology employed could be useful for a variety of other actors. In particular, NGOs often face questions of impact and sustainability. The approach used here answers both questions, by first gathering exactly the kind of data required to predict how much impact a market intervention may have, and then testing the optimal design. By further refining this kind of methodology, pilot studies and small grants might be made more effective in channeling limited public and international aid dollars into well-designed programs with impact. Which households benefit then depends on the ability of the intermediary to target and cross-subsidize effectively. Many non-state actors or financially constrained governments cannot perpetually subsidize goods or services, but can build platforms that minimize procurement costs, engage in cross-subsidization, and potentially even return a modest profit.

In addition to allowing the government to encourage take-up on the demand side, platforms such as this could be useful for regulating the supply side of the market. Operators engaged with the platform have the incentive to make sure that they have the correct licenses from the government and that they are providing the mandated quality of service so that they can continue to operate with the platform. Using the

[27]In addition, the asymptotics for standard demand estimation models are in the number of varieties of goods, which is fixed here at three: manual, mechanical, and household member.

“carrot” of additional business through engagement with the platform might allow government regulators to oversee the operations of suppliers more effectively than if they had to rely on costly “sticks” like fines or other punishments. This raises the possibility of general equilibrium effects — which our data and experiments do not address — that might further benefit both households and firms.

References

Alatas, V., Banerjee, A., Hanna, R., Olken, B., and Tobias, J. (2012). Targeting the poor: Evidence from a field experiment in Indonesia. *American Economic Review*, 104(2):1206–1204.

Alatas, V., Banerjee, A., Hanna, R., Olken, B., Wai-Poi, M., and Purnamasari, R. (2016). Self targeting: Evidence from a field experiment in Indonesia. *Journal of Political Economy*, 124(2):371–427.

Allcott, H., Knittel, C., and Taubinsky, D. (2015). Tagging and targeting of energy efficiency subsidies. *American Economic Review*, 105(5):187–91.

Amemiya, T. (1985). *Advanced Econometrics*. Harvard University Press.

Bagnoli, M. and Bergstrom, T. (2005). Log-concave probability and its applications. *Economic Theory*, 26:445–469.

Banerjee, A., Hanna, R., Olken, B. A., and Sumarto, S. (2020). The (lack of) distortionary effects of proxy-means tests: Results from a nationwide experiment in indonesia. *Journal of Public Economics Plus*, 1:100001.

Basurto, M. P., Dupas, P., and Robinson, J. (2020). Decentralization and efficiency of subsidy targeting: Evidence from chiefs in rural malawi. *Journal of public economics*, 185:104047.

Berry, J., Fischer, G., and Guiteras, R. (2020). Eliciting and utilizing willingness to pay: Evidence from field trials in northern ghana. *Journal of Political Economy*, 128(4):1436–1473.

Berry, S., Levinsohn, J., and Pakes, A. (1995). Automobile prices in market equilibrium. *Econometrica*, 63(4):841 – 890.

Borenstein, S. (2012). The redistributional impact of nonlinear electricity pricing. *American Economic Journal: Economic Policy*, 4(3):56–90.

Bulow, J. and Klemperer, P. (1996). Auctions versus negotiations. *American*

Economic Review, 86(1):180–194.

Coady, D., Grosh, M., and Hoddinott, J. (2004a). Targeting outcomes redux. *World Bank Research Observer*, 19(1):61–85.

Coady, D., Grosh, M. E., and Hoddinott, J. (2004b). Targeting of transfers in developing countries: Review of lessons and experience.

Coffey, D., Haque, S., Hathi, P., Pant, L., and Spears, D. (2014). Place and child health: The interaction of population density and sanitation behavior in developing countries. *Demography*, pages 1–24.

Cohen, J. and Dupas, P. (2010). Free distribution or cost-sharing? Evidence from a randomized malaria prevention experiment. *Quarterly Journal of Economics*, 125(1):1–45.

Dupas, P. (2014). Short-run subsidies and long-run adoption of new health products: Evidence from a field experiment. *Econometrica*, 82(1):197–228.

Dupas, P., Hoffman, V., Kremer, M., and Zwane, A. P. (2016). Targeting health subsidies through a nonprice mechanism: A randomized controlled trial in kenya. *Science*, 353(6302):889–895.

Dur, U., Pathak, P. A., and Sönmez, T. (2020). Explicit vs. statistical targeting in affirmative action: Theory and evidence from chicago's exam schools. *Journal of Economic Theory*, 187:104996.

Fewtrell, L., Kaufmann, R., Kay, D., Enanoria, W., Haller, L., and Colford, J. M. (2005). Water, sanitation, and hygiene interventions to reduce diarrhoea in less developed countries: a systematic review and meta-analysis. *Lancet infectious diseases*, 5(1):42–52.

Gertler, P., Shah, M., Alzua, M. L., Cameron, L., Martinez, S., and Patil, S. (2015). How does health promotion work? evidence from the dirty business of eliminating open defecation. *NBER Working paper*, (20997).

Guiteras, R., Levinsohn, J., and Mobarak, A. M. (2015). Encouraging sanitation investment in the developing world: A cluster-randomized controlled trial. *Science*, 348(6237):903–906.

Hanna, R. and Olken, B. (2018). Universal basic incomes vs targeted transfers: Anti-poverty programs in developing countries. *Journal of Economic Perspectives*, 32(4):201–226.

Jensen, R. and Miller, N. H. (2018). Market integration, demand, and the growth of firms: Evidence from a natural experiment in india. *American Economic Review*, 108(12):3583–3625.

Kar, K. and Pasteur, K. (2005). *Subsidy or self-respect? Community led total sanitation. An update on recent developments*. Institute of Development Studies.

Karlan, D. and Thuysbaert, B. (2019). Targeting ultra-poor households in honduras and peru. *The World Bank Economic Review*, 33(1):63–94.

Kidd, S. and Wylde, E. (2011). Targeting the poorest: An assessment of the proxy means test methodology. *Australian Aid*.

Kremer, M. and Miguel, E. (2007a). The illusion of sustainability. *The Quarterly journal of economics*, 122(3):1007–1065.

Kremer, M. and Miguel, E. (2007b). The illusion of sustainability. *Quarterly Journal of Economics*, 122(3):1007–1065.

Maddala, G. (1983). *Limited Dependent and Qualitative Variables in Econometrics*. Econometric Society Monographs. Cambridge University Press.

Mara, D., Lane, J., Scott, B., and Trouba, D. (2010). Sanitation and health. *PLoS Medicine*, 7(11).

McIntosh, C. and Zeitlin, A. (2021). Benchmarking a child nutrition program against cash: Experimental evidence from rwanda. *Journal of Development Economics*.

McKenzie, D. (2012). Beyond baseline and follow up: the case for more T in experiments. *Journal of Development Economics*, 99(2):210–221.

McRae, S. D. and Wolak, F. A. (2021). Retail pricing in colombia to support the efficient deployment of distributed generation and electric stoves. *Journal of Environmental Economics and Management*, page 102541.

Milgrom, P. and Segal, U. (2002). Envelope theorems for arbitrary choice sets. *Econometrica*, 70:583–601.

Mussa, M. and Rosen, S. (1978). Monopoly and product quality. *Journal of Economic theory*, 18(2):301–317.

Myerson, R. (1981). Optimal auction design. *Mathematics of Operations Research*, 6(1):58–73.

Narayan, A. and Yoshida, N. (2005). Proxy means test for targeting welfare benefits in Sri Lanka. *World Bank PREM Working Paper Series*, (SASPR-7).

Nevo, A. (2001). Measuring market power in the ready to eat cereal industry. *Econometrica*, 69(2):307–342.

Newey, W. (2009). Two-step estimation of sample selection models. *The Econometrics Journal*, 12(1):217–229.

Newey, W., Powell, J., and Walker, J. (1990). Semiparametric estimation of selection models: Some empirical results. *American Economic Review*, 80(2):324.

Null, C., Stewart, C. P., Pickering, A. J., Dentz, H. N., Arnold, B. F., Arnold, C. D., Benjamin-Chung, J., Clasen, T., Dewey, K. G., Fernald, L. C., et al. (2018). Effects of water quality, sanitation, handwashing, and nutritional interventions on diarrhoea and child growth in rural kenya: a cluster-randomised controlled trial. *The Lancet Global Health*, 6(3):e316–e329.

Olken, B. (2016). Hassles versus prices. *Science*, 6302:864–865.

Pathak, P. A. (2011). The mechanism design approach to student assignment. *Annu. Rev. Econ.*, 3(1):513–536.

Petrin, A. and Train, K. (2010). A control function approach to endogeneity in consumer choice models. *Journal of Marketing Research*, 47(1):3 – 13.

Powell, J. (1994). Estimation of semiparametric models. *Handbook of Econometrics*, 4:2443 – 2521.

Rochet, J. C. (1987). A necessary and sufficient condition for rationalizability in a quasi-linear context. *Journal of Mathematical Economics*, 16:191–200.

Spears, D. (2013). How much international variation in child height can sanitaiton explain? *Working Paper*.

Szabo, A. (2015). The value of free water: Analyzing south africa's free basic water policy. *Econometrica*, 83:1913–1961.

WHO and UNICEF (2015). 25 years progress on sanitation and drinking water: 2015 update and MDG assessment.

Wolak, F. A. (2016). Designing nonlinear price schedules for urban water to balance revenue and conservation goals. *NBER Working Paper*, (22503).

Table 1: Offers, Summary Statistics

	Min	1Q	Median	Mean	3Q	Max
Offer, η:	2,000	10,000	12,500	12,995.47	15,000	40,000

Table 2: Model-Based Predicted Treatment Effects

	Control	Treatment Effect	Rate	Adjusted Effect
Average	0.725	0.102	0.407	0.042
10,000	0.429	0.290	0.309	0.089
15,000	0.725	0.076	0.433	0.033
17,500	0.932	0.002	0.474	0.001
20,000	0.938	0.000	0.424	0.000

Control level and treatment effect computed based on (19) model and optimal Linear Programming prices. Rate is the proportion of households who got any desludging at all during the experiment, and the adjusted effect weights for the relatively low desludging rates, getting numbers strikingly close to the main results (see Table 3).

Table 3: Market Share Effects of Treatment

	Market Share Control Grp EL	Market Share		Percent Mechanical				Any Mechanical			
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)
Targeted Price		0.053**		0.037	0.064**			0.040*	0.052**		
		(0.024)		(0.024)	(0.029)			(0.024)	(0.026)		
Treat X 10k grp			0.127***			0.062	0.064			0.057	0.101*
			(0.040)			(0.057)	(0.049)			(0.058)	(0.052)
Treat X 15k grp			0.088**			0.015	0.078**			0.024	0.033
			(0.038)			(0.031)	(0.034)			(0.032)	(0.029)
Treat X 17.5k grp			0.022			0.054*	0.047			0.047	0.051
			(0.040)			(0.030)	(0.038)			(0.031)	(0.031)
Treat X 20k grp			0.007			0.017	0.003			0.046	0.052
			(0.061)			(0.046)	(0.079)			(0.043)	(0.055)
10k group	0.59					0.536***				0.543***	
						(0.053)				(0.053)	
15k group	0.85					0.687***				0.690***	
						(0.050)				(0.051)	
17.5k group	0.91					0.733***				0.746***	
						(0.049)				(0.051)	
20k group	0.99					0.776***				0.767***	
						(0.059)				(0.059)	
Endline					0.208***		0.208***		0.069***		0.069***
					(0.021)		(0.021)		(0.017)		(0.017)
N		92	301	1199	2390	1199	2390	1199	2390	1199	2390
R^2		0.286	0.935	0.109	0.203	0.857	0.204	0.105	0.047	0.862	0.049
mean				0.825		0.825		0.841		0.841	

Market share is defined as $\frac{\#mechanicaldesludgings}{\#mechanical+\#manualdesludgings}$. Column (1) gives the average neighborhood market share for each price group of mechanical desludging at endline for the control group. Column (2) provides the OLS for the Average Treatment Effect on market share of mechanical (3) provides the OLS estimate of the market share effect for each price group in a neighborhood cluster (level of observation is neighborhood-price group, not all neighborhoods include households from each price group). Column (4) provides the OLS and (5) the FE estimate of the pooled effect with observations at the household level. Column (6) provides the OLS and (7) the FE estimate of the effect of the treatment on the percent of desludgings that a household purchases that are mechanical by price group with observations at the household level. Columns (8) and (9) provide the OLS and FE estimates respectively of the pooled effect of treatment on the likelihood of a household purchasing at least one mechanical desludging. Columns (10) and (11) provide the OLS and FE effects of treatment on households belonging to each price group. Controls are included in the OLS regressions for the neighborhood means (in column 2) or neighborhood-price group means (in column 3) or household level values (columns 4, 6, 8 and 10) of the variables not balanced at baseline (water bill more than 5,000 CFA, latrine pit distance to road, two tanks used last desludging, other households in compound, compound has 2 or more pits). Baseline percent of desludgings for which the household purchased mechanical included as a control for pre-baseline desludging behavior of the household in the OLS specification. A control has also been included for the average of the stratification variable–above median number of households in neighborhood have high compound walls in the OLS specification. Fixed effects for each of the 1199 households which purchased desludgings are included in the fixed effects regressions. Standard errors are clustered at the neighborhood level in household level regressions.

Table 4: Impact of Treatment on Children's Diarrhea

	Any Child Diarrhea					
	(1)	(2)	(3)	(4)	(5)	(6)
	OLS	FE	OLS	FE	OLS	FE
Targeted price group	-0.012	-0.032				
	(0.017)	(0.025)				
Treatment*10k price group			-0.052*	-0.070**		
			(0.030)	(0.030)		
Treatment*15k price group			0.004	-0.026		
			(0.021)	(0.028)		
Treatment*17.5k price group			0.007	-0.001		
			(0.035)	(0.044)		
Treatment*20k price group			0.002	0.047		
			(0.059)	(0.062)		
Treatment*Pct Nbhd in 10k group					-0.216**	-0.249***
					(0.104)	(0.094)
Treatment*Pct Nbhd in 15k group					0.070	0.044
					(0.078)	(0.080)
Treatment*Pct Nbhd in 17k group					0.007	0.048
					(0.157)	(0.153)
Treatment*Pct Nbhd in 20k group					0.217	0.290
					(0.321)	(0.400)
		0.001		0.001		0.001
		(0.018)		(0.018)		(0.018)
N	1995	3758	1995	3758	1995	3758
R^2	0.020	0.003	0.142	0.005	0.144	0.007
mean	0.123	0.130	0.123	0.130	0.123	0.130

The dependent variable is an indicator for the participant reports that a household child had diarrhea in the past 7 days at endline in the cross sectional regressions, in the 7 days prior to the survey in the fixed effects regressions. Price-group and percent neighborhood in price-group controls are included in specifications (3) and (5) but not shown for brevity. The OLS specifications include controls for the variables not balanced at baseline, pre-baseline desludging variables, and baseline diarrhea in household children, and a control for the stratification variable.

Table 5: Mean Baseline Characteristics by Price Group

	10000	15000	17500	20000	Pooled
Phone Credit use over past week	1107	1754	4078	4882	2157
	(1929)	(2930)	(5660)	(10195)	(4152)
Number of Refrigerators	0.168	0.507	0.927	1.371	0.530
	(0.430)	(0.665)	(0.771)	(0.726)	(0.706)
Number of Cars	0.061	0.298	0.671	1.529	0.357
	(0.248)	(0.577)	(0.838)	(1.073)	(0.683)
Number of Air Conditioners	0.016	0.081	0.477	1.486	0.199
	(0.157)	(0.359)	(1.004)	(1.909)	(0.722)
Ever Desludged Mech	0.357	0.571	0.621	0.686	0.524
	(0.479)	(0.495)	(0.486)	(0.468)	(0.499)
Expected Price Mechanical (CFA)	12792	14103	15847	16716	14243
	(4717)	(4743)	(5550)	(7120)	(5173)
Last used Manual	0.510	0.219	0.153	0.030	0.263
	(0.500)	(0.414)	(0.361)	(0.171)	(0.440)

This table provides means for each variable at baseline by the price group to which they were assigned by the pricing model. Standard deviations are in parentheses.

Table 6: Counterfactual Shares (%): Alternative Market Designs

	Realized	$\widehat{\text{Control}}$	$\widehat{\text{Treatment}}$	Auction	Auction(S)	Proxy-Means(100)	Proxy-Means(150)	Market	Market (S)
Average	81.1	76.7	80.8	78.8	80.7	79.3	80.4	78	79.8
	(79.5,82.8)	(76.1,78)	(79.1,82.3)	(75.1,80.8)	(78.6,82.2)	(77.6,80.8)	(77.9,81.8)	(73.3,80.4)	(77.04,81.4)
10,000	68.7	57.9	68.6	60.2	62	61.6	62.2	59.5	61.2
	(65.38,72.4)	(57,61.1)	(65.8,71.8)	(50,65.7)	(53.4,66.5)	(54.5,65.9)	(53.94,66.6)	(49,65.4)	(51.94,66.1)
15,000	80.7	77.9	80.4	79	81.3	80.1	80.9	78	80.2
	(78.68,82.7)	(77,79.1)	(78.3,82.3)	(75.6,81.2)	(79.5,83.36)	(78.8,81.96)	(78.9,82.8)	(73.5,80.7)	(77.9,82)
17,500	89.7	86.5	89.3	90.7	92.2	89.2	91.8	89.9	91.4
	(87.3,91.9)	(85.3,88.1)	(87.3,90.7)	(88.5,92.8)	(89.4,94.4)	(87.8,91)	(89.2,93.5)	(87.9,91.6)	(89,93.7)
20,000	96.8	95.9	95.9	98.1	98.4	97.2	98	97.9	98.3
	(94.8,98.8)	(94.9,98.2)	(93.74,97.4)	(95.9,98.9)	(96.4,99.1)	(95.8,98.5)	(95.5,98.8)	(95.54,98.7)	(96.2,99)

Gives market shares for alternative platform designs. Realized is the experimental value, $\widehat{\text{Control}}$ is the lasso-predicted value for the targeted pricing group using control group data, $\widehat{\text{Treatment}}$ is the model predicted value, Auction is the predicted outcome when households are given the procurement auction prices, Proxy-Means(100) and Proxy-Means(150) refer to predicted outcomes using proxy-means testing at 100% and 150% of the poverty line, Market is a policy that constrains prices to be the average price in the search market, and (S) denotes an additional subsidy of $3.00. Bootstrapped 90% confidence intervals reported below point estimate.

Table 7: Counterfactual Budget Balance (CFA): Alternative Market Designs

	Realized	$\widehat{\text{Treatment}}$	Auction	Auction (S)	Proxy-Means(100)	Proxy-Means(150)	Market	Market(S)
Average	-102	-100	0	0	0	-230	78	144
	(-150,-53)	(-147,-52)	(0,0)	(0,0)	(0,0)	(-251,-208)	(64,98)	(127,167)
10,000	-1334	-1404	0	0	0	-275	82	120
	(-1477,-1189)	(-1551,-1267)	(0,0)	(0,0)	(0,0)	(-343,-216)	(54,128)	(89,170)
15,000	116	116	0	0	0	-252	86	150
	(91,144)	(98,138)	(0,0)	(0,0)	(0,0)	(-275,-228)	(70,107)	(130,174)
17,500	334	404	0	0	0	-140	53	143
	(273,401)	(352,473)	(0,0)	(0,0)	(0,0)	(-171,-106)	(42,65)	(113,179)
20,000	689	696	0	0	0	-226	98	191
	(308,1064)	(516,861)	(0,0)	(0,0)	(0,0)	(-309,-129)	(59,131)	(120,259)

Bootstrapped 90% confidence intervals reported below point estimate.

Table 8: Counterfactual Subsidization Rates (%): Alternative Market Designs

	Realized	$\widehat{\text{Treatment}}$	Auction	Auction (S)	Proxy-Means(100)	Proxy-Means(150)	Market	Market(S)
Average	2.66	2.66	0.00	2.12	1.43	1.79	-0.76	1.02
	(2.31,3.01)	(2.31,3.00)	(0.00,0.00)	(1.93,2.33)	(1.3,1.57)	(1.58,1.95)	(-0.93,-0.64)	(0.87,1.15)
10,000	11.35	11.82	0.00	2.04	1.70	1.97	-0.63	1.12
	(10.17,12.58)	(10.67,13.03)	(0.00,0.00)	(1.63,2.55)	(1.34,2.13)	(-1.54,2.38)	(-0.98,-0.42)	(0.83,1.44)
15,000	1.41	1.31	0.00	2.20	1.57	1.95	-0.79	1.06
	(1.16,1.62)	(1.14,1.47)	(0.00,0.00)	(2.00,2.38)	(1.42,1.71)	(1.72,2.12)	(-0.99,-0.65)	(0.9,1.2)
17,500	-0.98	-1.22	0.00	1.93	0.88	1.23	-0.75	0.83
	(-1.19,-0.78)	(-1.45,-1.06)	(0.00,0.00)	(1.51,2.42)	(0.66,1.07)	(0.9,1.51)	(-0.91,-0.62)	(0.66,1.03)
20,000	-3.15	-3.12	0.00	2.42	1.43	1.97	-1.02	0.98
	(-4.86,-1.47)	(-3.85,-2.31)	(0.00,0.00)	(1.49,3.32)	(0.81,1.95)	(1.19,2.69)	(-1.33,-0.68)	(0.61,1.40)

Bootstrapped 90% confidence intervals reported below point estimate.

Table 9: Design Variations (CFA)

	π^*_{ATE}	$\pi^*_{\text{ATE, 10k group}}$
Auction	1,704	6,132
	(1,530, 2,254)	(6,017, 6,257)
PMT(100)	1,665	5,407
	(907, 4,176)	(4,747, 7,338)
PMT(150)	350	4518
	(44, 1,292)	(4,331, 4,818)
Market	2478	6825
	(2,321, 2,973)	(6,807, 6,875)

Gives the additional subsidy requires for each alternative market design to match the Targeted Pricing treatment effect on average (column 1) and in the 10,000 CFA price bin (column 2).

Table 10: Counterfactual Shares (CFA): Alternative Information Designs

	Realized	$\widehat{\text{Control}}$	Original	Conservative	Municipal	Municipal 2	NGO	NGO 2
Average	81.1	76.7	80.6	80.6	81.5	82.1	80.7	77.6
	(79.5,82.8)	(76.1,78)	(78.6,82)	(78.6,82.1)	(79.6,82.8)	(80.1,83.5)	(78.6,82)	(72.44,80.3)
10,000	68.7	57.9	66.3	65.9	66	65.6	64.6	63.3
	(65.38,72.4)	(57,61.1)	(62.2,69.4)	(61.64,69)	(61.6,69.1)	(60.9,68.8)	(59.04,67.9)	(58.34,67.3)
15,000	80.7	77.9	80.4	80.4	81.9	82.4	80.5	76.8
	(78.68,82.7)	(77,79.1)	(78.2,82.2)	(78.2,82.3)	(80,83.6)	(80.4,84.3)	(78.4,82.3)	(70.44,80)
17,500	89.7	86.5	90	90.4	90.8	92.4	91.4	88.1
	(87.3,91.9)	(85.3,88.1)	(88,91.5)	(88.3,92.1)	(88.6,92.6)	(89.6,94.46)	(89,93.6)	(84.64,89.7)
20,000	96.8	95.9	97.7	97.7	97.7	98.3	98.1	96.4
	(94.8,98.8)	(94.9,98.2)	(95.3,98.6)	(95.3,98.6)	(95.3,98.6)	(96.2,99)	(95.9,98.9)	(94.5,97.8)

Gives market shares for alternative information structures, defined in Figure 11. Bootstrapped 90% confidence intervals reported below point estimate.

Table 11: Budget Balance (CFA): Alternative Information Designs

	Realized	Original	Conservative	Municipal	Municipal 2	NGO	NGO 2
Average	-102	-24	-22	-177	-306	-18	176
	(-150,-53)	(-62,18)	(-55,14)	(-224,-120)	(-383,-225)	(-43,14)	(23,406)
10,000	-1334	-894	-808	-849	-776	-553	-432
	(-1477,-1189)	(-1000,-780)	(-910,-699)	(-965,-724)	(-884,-652)	(-629,-452)	(-577,-149)
15,000	116	109	100	-158	-242	94	325
	(91,144)	(84,139)	(73,128)	(-211,-87)	(-325,-150)	(74,117)	(131,601)
17,500	334	331	283	225	-132	137	363
	(273,401)	(288,390)	(245,333)	(191,271)	(-227,-58)	(109,170)	(248,518)
20,000	689	530	530	520	155	300	292
	(308,1064)	(372,677)	(372,677)	(364,670)	(94,213)	(191,409)	(86,494)

Gives market shares for alternative information structures, defined in figure 11. Bootstrapped 90% confidence intervals reported below point estimate.

Table 12: Variable Importance of Design Variables (Gini coefficient)

	Original	Conservative	Municipal	Municipal 2	NGO	NGO 2
Average Months Between Desludgings	68	65.728				
Water Bill More Than 5,000 CFA	9.018	10.646	14.748	16.617		
House Type: Precarious	29.491	38.374			97.735	52.681
House Type: Concrete	23.124	29.911			81.786	13.392
House Type: Rooming House	4.947	5.382			12.559	2.759
Other Households in Compound	16.264	14.336			6.632	18.71
Own House	16.399	14.96	39.702			
Pit Meters From Road	19.208	16.077			5.229	57.635
More than 1 Trip Last Desludging	11.521	10.462				
Electricity Bill	151.488	174.222	236.082	189.703		
Household Size	25.631	22.908			11.336	
Number of Women in Household	21.412					
Respondent Finished High School	19.257					

Variable Importance averages the amount by which adding the variable to one of the decision trees in the random forest reduced misclassification as measured by the Gini coefficient at terminal nodes of the decision tree.

A Figures

Figure 1: Baseline Prices of Mechanical and Manual Services

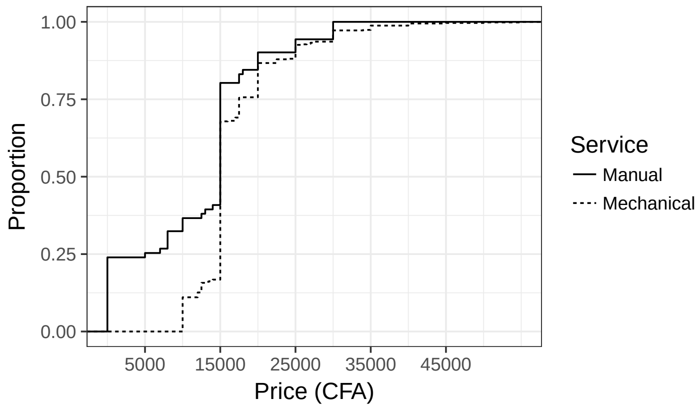

Households that provide manual desludging services themselves pay nothing, accounting for the high value for Manual at a price of 0. The modal price in the search market for manual and mechanical services is 15,000 CFA, but the mechanical price distribution first-order stochastically dominates the manual price distribution on all of the support.

Figure 2: Household decision-making

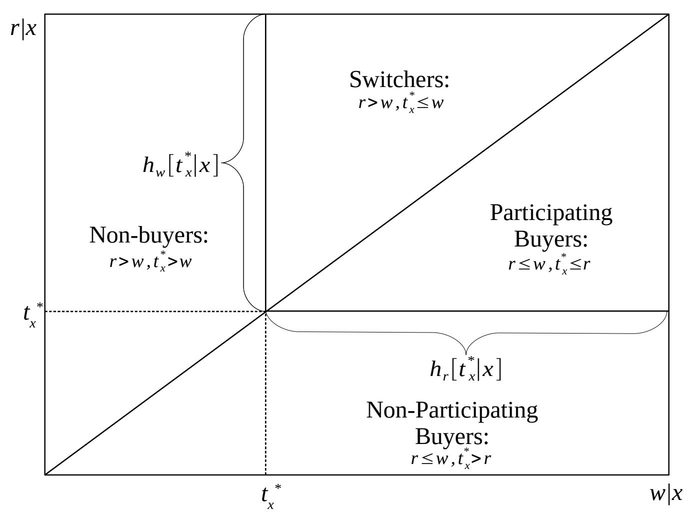

Figure 3: Theoretical and Empirical Taxonomies

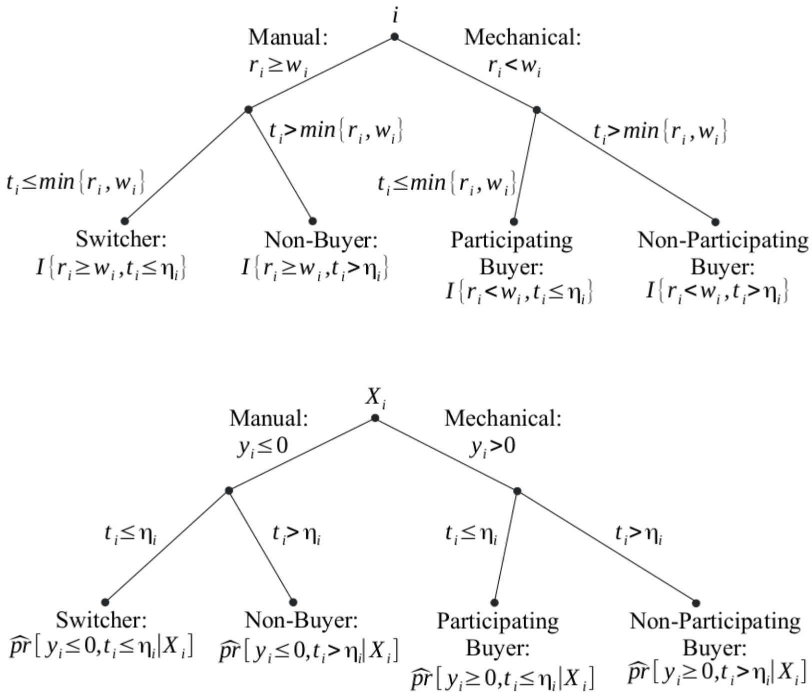

Figure 4: Offers

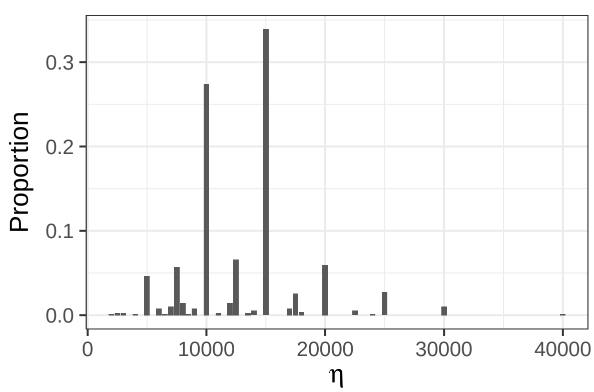

Figure 5: Demand

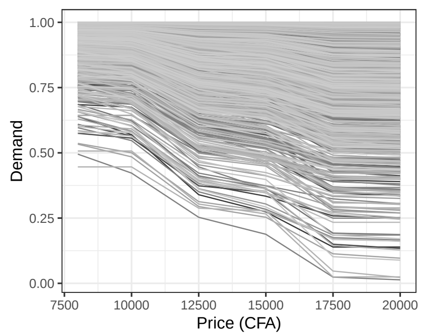

(a) Mechanical Demand by Household

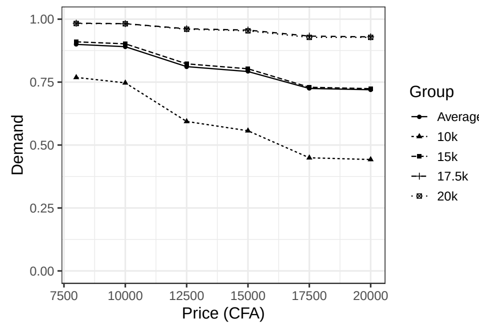

(b) Average Mechanical Demand

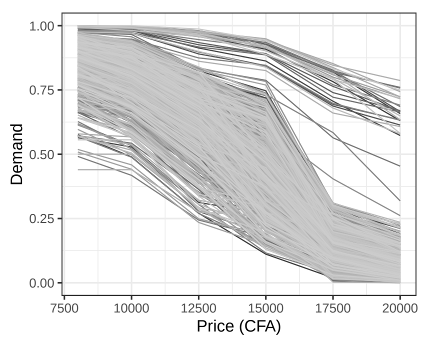

(c) Platform Demand by Household

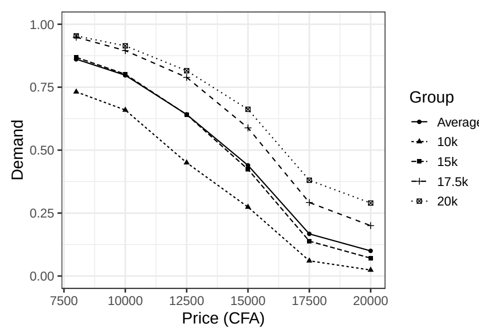

(d) Average Platform Demand

Figure 6: Supply-Side Auctions Average Clearing Prices by Round

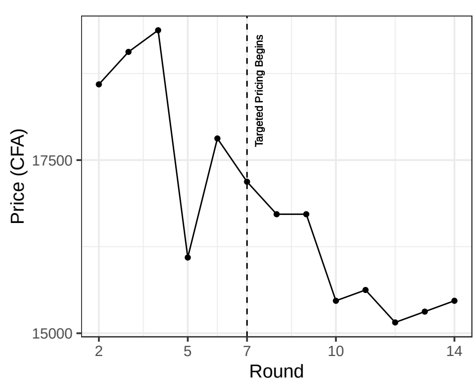

Figure 7: Pricing Rules

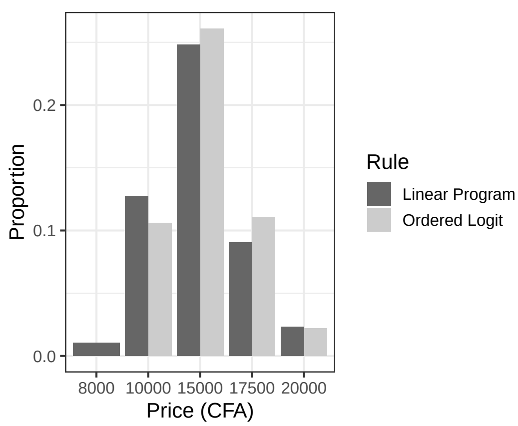

(a) Linear Programming and Ordered Logit Prices

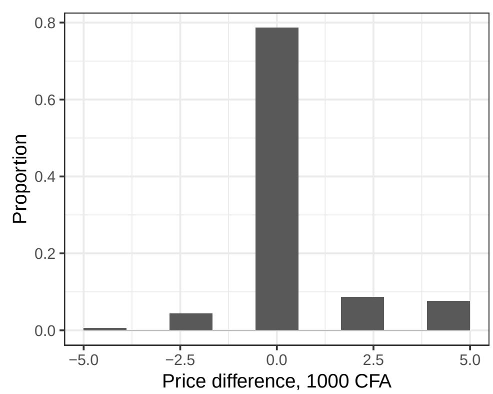

(b) Deviations of Ordered Logit from Linear Programming Price

Figure 8: Counterfactual Market Shares

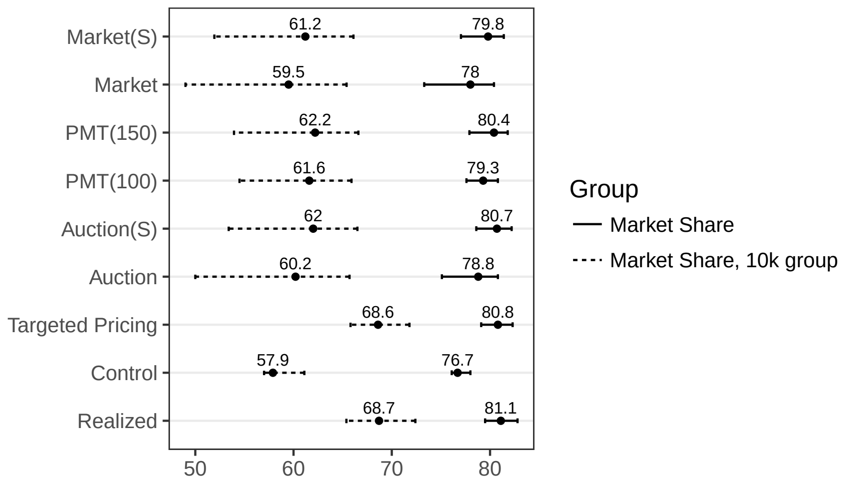

(a) Alternative Designs

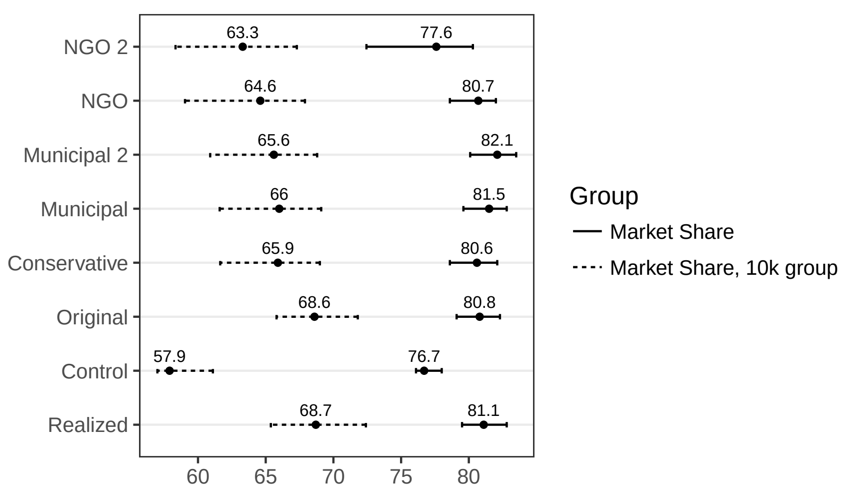

(b) Alternative Information

Figure 9: Sustainability Analysis

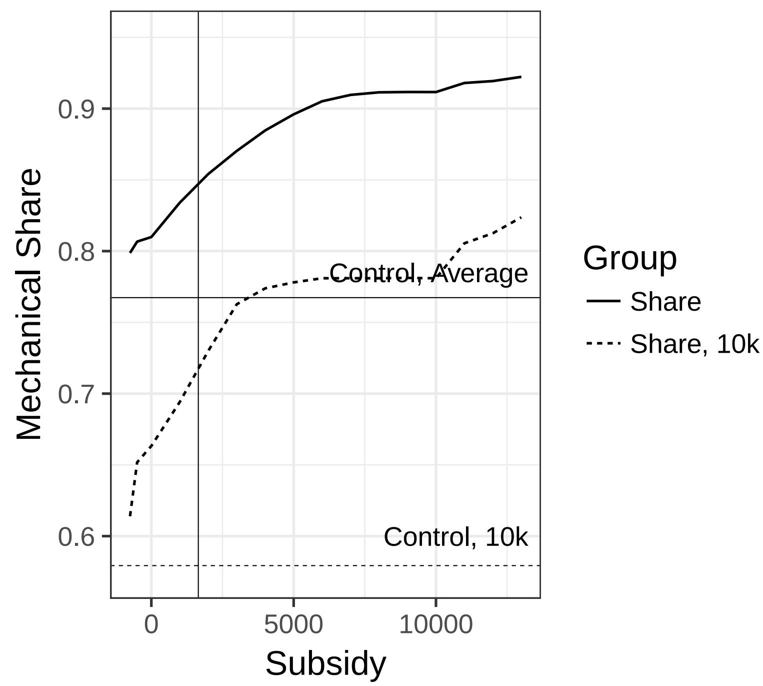

(a) Mechanical Market Share

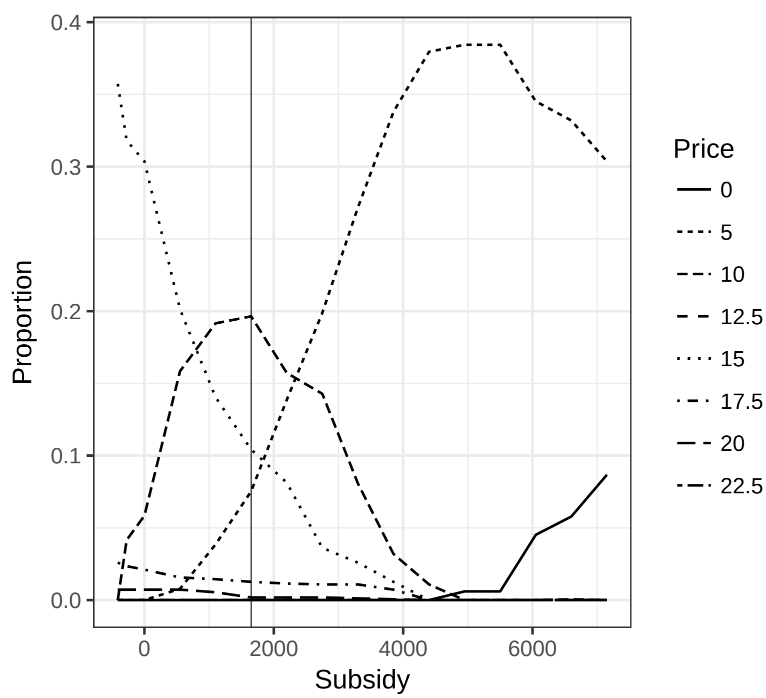

(b) Price Offers

Figure 10: Alternative Designs: Subsidization Rates

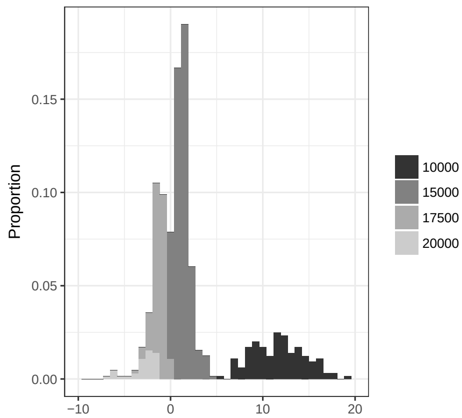

(a) Targeted Pricing

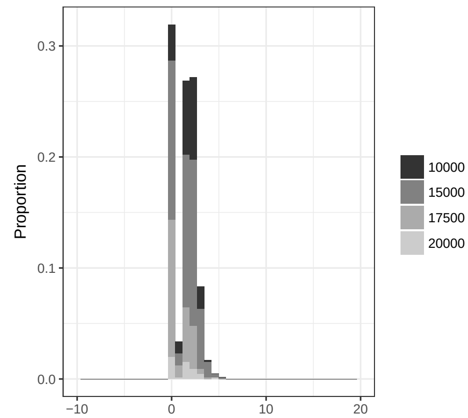

(b) Proxy-Means Testing

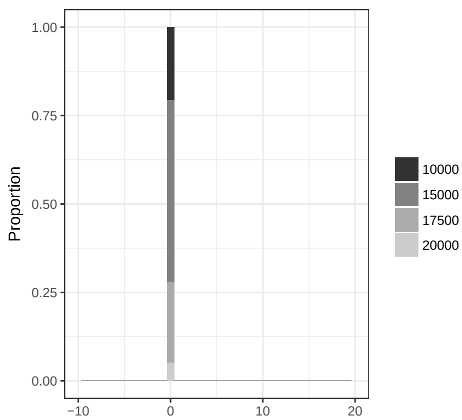

(c) Auction

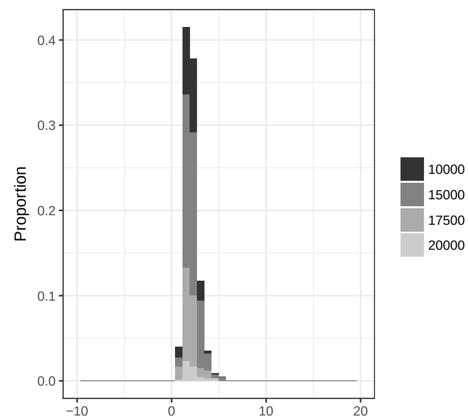

(d) Subsidized Auction

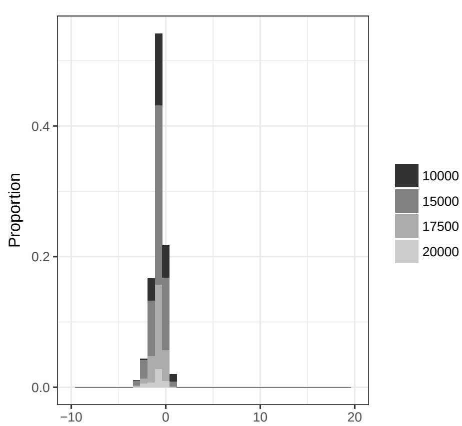

(e) Market Average

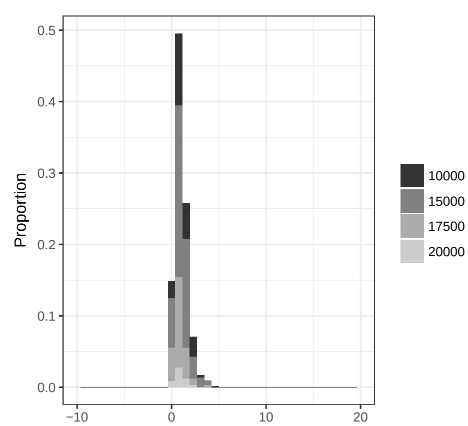

(f) Subsidized Market Average

Figure 11: Information Structures

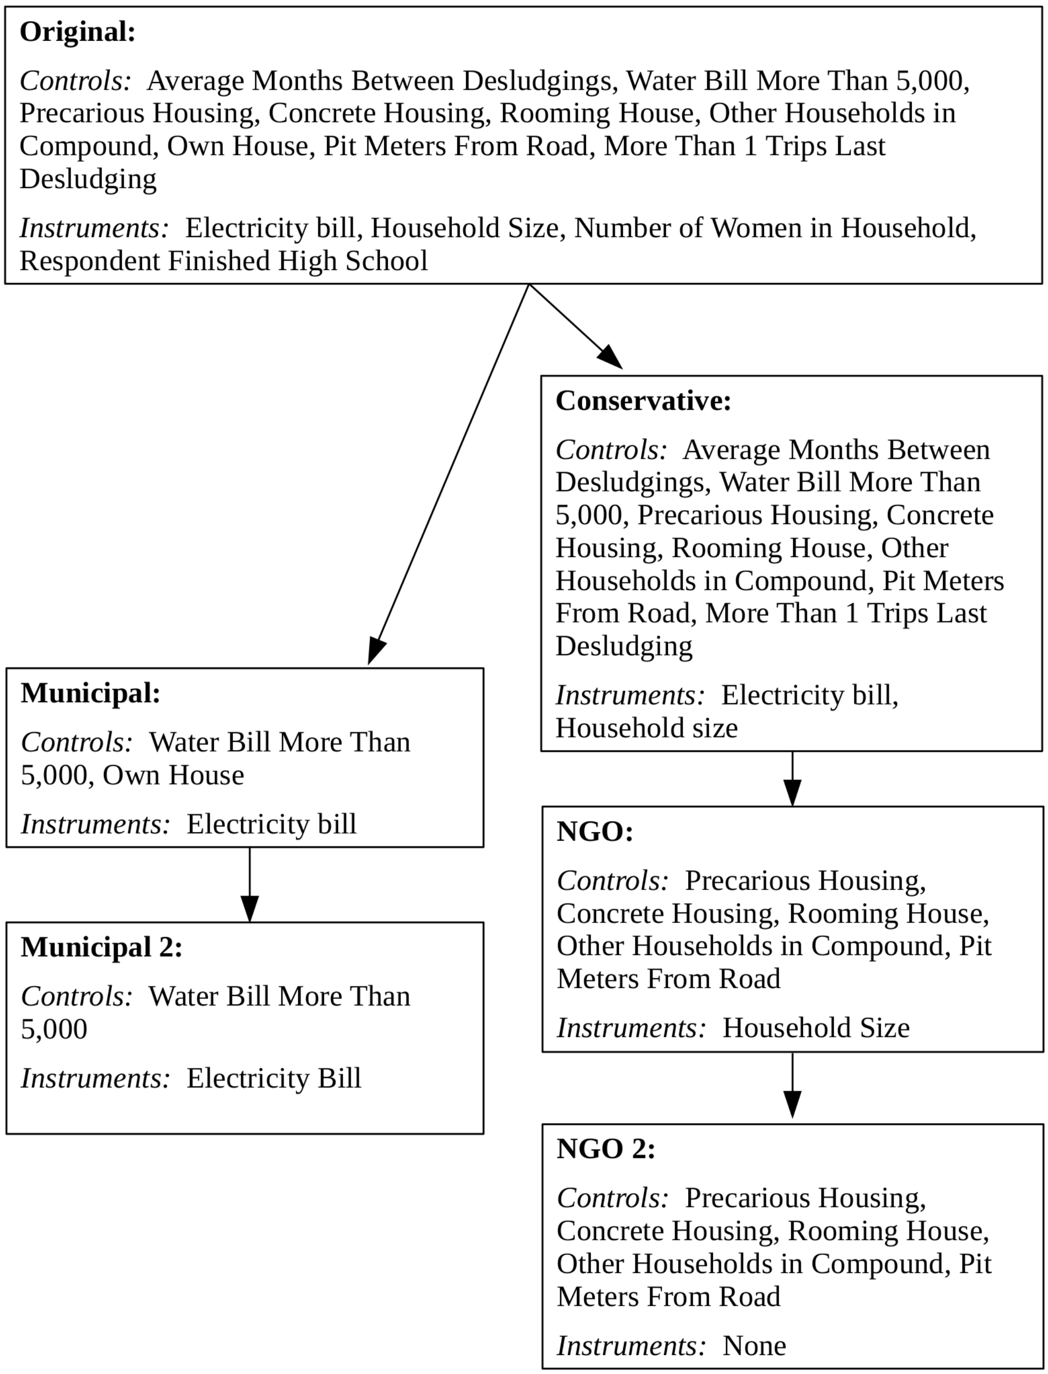

Figure 12: Alternative Information Design: Price Rules

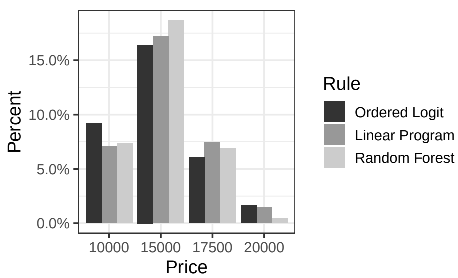

(a) Original

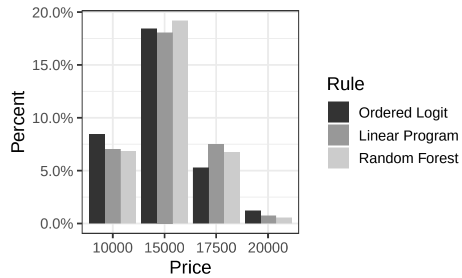

(b) Conservative

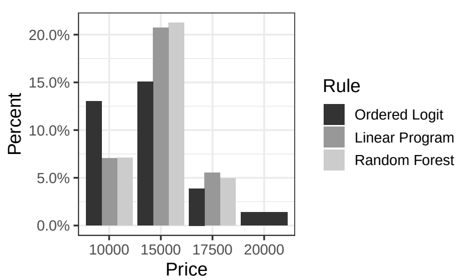

(c) Municipal

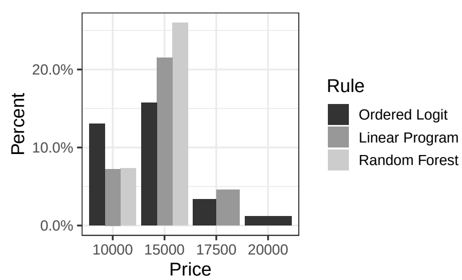

(d) Municipal 2

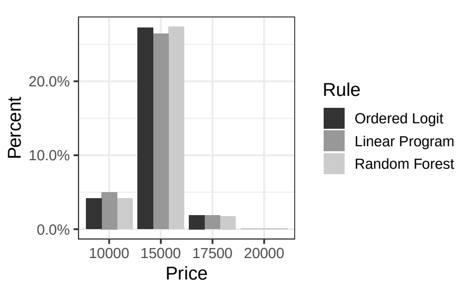

(e) NGO

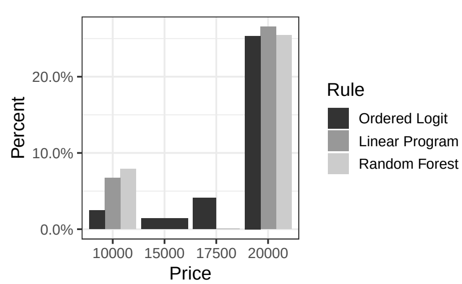

(f) NGO 2

B Web Appendices (Not for Publication)

A Mechanism Design Details, Theoretical Platform Design

There is a unit mass of households, all of whom must decide between purchasing the manual or mechanical service. Each household has a privately known willingness-to-pay w, privately known outside price r, and publicly known observables x from a set[28] X. The willingness-to-pay w is the maximum price at which a household would be willing to switch from manual to mechanical desludging. The outside price r is the amount the household anticipates paying in the prevailing decentralized market for a mechanical desludging. The observable type x corresponds to characteristics observable to the market, like the household's neighborhood or the quality of its dwelling, or observable to a municipal authority such as ONEA, such as water or electricity bill expenditures.

The willingness-to-pay w and outside price r are distributed $F_w[w|x]$ and $F_r[r|x]$, with support on $[\underline{w},\overline{w}]$ with densities $f_w[w|x]$ and $f_r[r|x]$, respectively. Note that conditional on x, r is independent of w since the market does not observe the household's private information[29]. The prevailing market can never reasonably charge more than the maximum willingness-to-pay of a household, and neither is it profitable for a firm with some pricing power to charge less than the minimum willingness-to-pay of a household. Thus, the support of r is also $[\underline{w},\overline{w}]$. Assume the standard regularity conditions that $1-F_w[w|x]$ and $1-F_r[r|x]$ are log-concave[30]. Households have quasi-linear utility, so that consuming a mechanical desludging at a price of t yields a payoff $w-t$, while the payoff of consuming a manual desludging is normalized to 0. A household procures the mechanical service in the prevailing market only if its willingness-to-pay is sufficiently high, so that $w-r \geq 0$. The platform competes alongside a prevailing, decentralized market for mechanical services. Since our sample

[28] In our applications, x will correspond to data including a mix of real-valued variables, integer-valued ones, dummy variables, and categorical variables or factors, so we place no additional structure on X beyond requiring it be a valid measure space and the distributions and densities of w and r and the mechanism be $(\Omega_X,\mathcal{F}_X,\mu_X)$-measurable.

[29] It if were otherwise, knowing r would reveal additional information about w about x, implying that the market must be using additional variables to determine a price.

[30] This implies that a profit-maximizing monopolist's second-order condition is satisfied. See (Mussa and Rosen, 1978), (Myerson, 1981), and (Bagnoli and Bergstrom, 2005).

includes a small number of households relative to the overall size of the market, we assume that the platform does not create general equilibrium effects that change the expected probability of trade or payment in the prevailing market.

The Revelation Principle guarantees that any game of incomplete information can, without loss of generality, be converted into an alternative game called a *direct mechanism*, in which agents report their types and types determine payoffs. In this setting, it ensures that any method the platform can use to arrange trade is equivalent to some direct mechanism,

$$\{p(w,r,x),t(w,r,x)\}_{w\in[\underline{w},\overline{w}],r\in[\underline{w},\overline{w}],x\in X}$$

in which a household with observables x reports — not necessarily honestly — some type $(\hat{w},\hat{r})$ and trade occurs with probability $p(\hat{w},\hat{r},x)$ at a price of $t(\hat{w},\hat{r},x)$. A direct mechanism is incentive compatible if[31] for all w, r, x, $\hat{w}$, and $\hat{r}$,

$$\underbrace{p(w,r,x)}_{\text{Pr[Trade on the platform]}}\underbrace{(w-t(w,r,x))}_{\text{Platform payoff}}+\underbrace{(1-p(w,r,x))}_{\text{Pr[Trade off the platform]}}\underbrace{\max\{w-r,0\}}_{\text{Outside option}}$$
$$\geq p(\hat{w},\hat{r},x)(w-t(\hat{w},\hat{r},x))+(1-p(\hat{w},\hat{r},x))\max\{w-r,0\}$$

or, converting to net quantities and noting that $w-\max\{w-r,0\}=\min\{w,r\}$,

$$p(w,r,x)(\min\{w,r\}-t(w,r,x))\geq p(\hat{w},\hat{r},x)(\min\{w,r\}-t(\hat{w},\hat{r},x)). \tag{8}$$

Similarly, a direct mechanism is individually rational if for all w, r, and x,

$$\underbrace{p(w,r,x)}_{\text{Pr[Trade on the platform]}}\underbrace{(w-t(w,r,x))}_{\text{Platform payoff}}+\underbrace{(1-p(w,r,x))}_{\text{Pr[Trade off the platform]}}\underbrace{\max\{w-r,0\}}_{\text{Outside option}}\geq\underbrace{\max\{w-r,0\}}_{\text{Outside option}},$$

or, again converting to net quantities,

$$p(w,r,x)(\min\{w,r\}-t(w,r,x))\geq 0. \tag{9}$$

In addition to the individual rationality and incentive compatibility constraints, the

[31] The more standard way of writing these constraints in the mechanism design literature is to subtract $t(w,r,x)$ from expected surplus without loss of generality, without multiplying it by $p(w,r,x)$. Readers more familiar with this approach can substitute $\tilde{t}(w,r,x)=p(w,r,x)t(w,r,x)$, and the analysis and results will be identical.

platform must also ensure that its total profits plus subsidies, s, are non-negative, or

$$\mathbb{E}_{(w,r,x)}[p(w,r,x)(t(w,r,x)-c_x)]+s\geq 0 \tag{10}$$

where c_x is the expected cost of serving a household with observables x. Call (10) the *expected budget balance* constraint.

As discussed in the introduction, there are significant negative externalities from the collection and disposal of human fecal sludge, especially on young children for whom exposure to human waste can lead to diarrhea, stunting, and death. Let b_x be the net social benefit of a household of type x consuming the mechanical service rather than manual. The platform seeks to solve the *targeting problem*: pick the payments $t(w,r,x)$ and probabilities of trade $p(w,r,x)$ to solve:

$$\max_{\{p,t\}}\mathbb{E}_{(w,r,x)}\left[\underbrace{p(w,r,x)b_x}_{\text{Platform purchases}}+\underbrace{(1-p(w,r,x))\mathbb{I}\{w\geq r\}b_x}_{\text{Market purchases}}\right]$$

subject to incentive compatibility (8), individual rationality (9), and expected budget balance (10). Due to the challenges of eliciting households' preferences over their neighbors' consumption of the mechanical service[32], we focus on maximizing the share of mechanical services, setting $b_x=1$.

The incentive compatibility constraints require the platform to select mechanisms which respect the households' agency in reporting, but honest revelation of both w and r cannot be incented: only the minimum of the two appears directly in the household's problem, so that the household will lie in the most advantageous way about the maximum of the two[33]. In order for a direct mechanism to be incentive compatible, it must then be a function only of the minimum of $\hat{w}$ and $\hat{r}$. Define $\eta=\min\{w,r\}$, and instead ask households to make a report of this value, $\hat{\eta}$; to distinguish this from the willingness-to-pay w, we refer to η as the household's *willingness-to-switch*. Transforming the problem in this way allows us to use standard tools to compute (steps included in the proof of Theorem 1) the platform's profits

[32]We piloted a variety of demand elicitation games that asked whether households would be willing to pay something if ensured their neighbors received mechanical services, but participants found this unnatural, given the political economy of their neighborhoods.

[33]If a household was going to purchase anyway, the relevant thing to lie about is the price it would have faced; if a household was not going to purchase mechanical services on its own, the relevant thing to lie about is its willingness-to-pay. So in either case, only w or r is payoff-relevant, not both.

in terms of the probabilities of trade in any incentive compatible direct mechanism:

$$\mathbb{E}_{(\eta,x)}[p(\eta,x)(t(\eta,x)-c_x)]=\mathbb{E}_{(\eta,x)}\left[p(\eta,x)\left\{\eta-\frac{1-F_\eta[\eta|x]}{f_\eta[\eta|x]}-c_x\right\}\right]. \tag{11}$$

The quantity

$$\psi_\eta[\eta|x]=\eta-\frac{1-F_\eta[\eta|x]}{f_\eta[\eta|x]}$$

is called the virtual value, and corresponds to the expected marginal revenue from an η type conditional on x.

Given that the mechanism must be a function of $\eta=\min\{w,r\}$ and not w and r separately, the objective function can similarly be simplified:

$$\begin{aligned}&\mathbb{E}_{(w,r,x)}[p(\min\{w,r\},x)b_x+(1-p(\min\{w,r\},x))\mathbb{I}\{w\geq r\}b_x]\\ &\qquad=\mathbb{E}_{(\eta,x)}\left[p(\eta,x)\frac{h_w[\eta|x]}{h_w[\eta|x]+h_r[\eta|x]}b_x+\frac{h_r[\eta|x]}{h_w[\eta|x]+h_r[\eta|x]}b_x\right]\end{aligned} \tag{12}$$

where $h_z[\eta|x]$ is the hazard rate of the random variable z at η given x: $f_z[\eta|x]/(1-F_z[\eta|x])$. The *switch function*

$$\sigma(\eta,x)=\frac{h_w[\eta|x]}{h_w[\eta|x]+h_r[\eta|x]}$$

is an odds ratio of hazard rates that captures the platform's inference about a household's propensity to switch from manual to mechanical. It answers the question, what is the probability that a household reporting η switches from manual to mechanical rather than from buying in the market to on the platform?

The optimal allocation rule $\{p(\eta,x)\}_{x\in X}$ necessarily satisfies the Kuhn-Tucker conditions, so that p maximizes the Lagrangian

$$\mathcal{L}(p,\lambda)=\mathbb{E}_{(\eta,x)}[p(\eta,x)\{\sigma(\eta,x)+\lambda(\psi_\eta[\eta|x]-c_x)\}] \tag{13}$$

where λ is the multiplier on the expected budget balance constraint. The term in braces represents the marginal benefit of serving a household with observables x reporting η,

$$\beta(w,r,x,\lambda)=\underbrace{\sigma(\eta,x)}_{\text{Marginal propensity to switch at }(\eta,x)}b_x+\underbrace{\lambda}_{\text{Shadow value of profit}}\underbrace{(\psi_\eta[\eta|x]-c_x)}_{\text{Marginal profit from }(\eta,x)}. \tag{14}$$

When this term is positive, the platform prefers to provide a desludging to the (η,x)

type and set $p(\min\{w,r\},x,\lambda)=1$, and otherwise set $p(\min\{w,r\},x,\lambda)=0$. The first term is the odds of a switch at η given x, capturing the social motive. The second term is the marginal profit generated by the sale to the (η,x) type weighted by the shadow value of the expected budget balance constraint, capturing the profit motive. If λ is small, the platform will widely distribute mechanical desludgings at low prices, while if λ is large, the budget constraint is relatively binding and it will behave more like a purely profit-maximizing platform. This shows how the platform is a "profit-minded social planner," who places some weight on profits and some on consumption of improved services, where the weight is endogenously determined by the balancing the budget with the relative likelihoods of the households to switch.

A full analysis of the problem characterizes the optimal mechanisms in this environment:

Theorem 1 *Suppose $\underline{w}-c_x \leq s$, so that the subsidy is not sufficiently large to provide every household in the market with a mechanical desludging, and that for all $x\in X$, $\sigma(\eta,x)$ is non-decreasing in η.*

For all x, there is a type $\eta_x^*(\lambda^*)$ that satisfies $\sigma(\eta_x^*(\lambda^*),x)b_x+\lambda^*(\psi_\eta[\eta_x^*(\lambda^*)|x]-c_x)=0$, such that in the optimal mechanism,

$$p^*(\eta,x,\lambda^*)=\begin{cases}1, & \eta\geq\eta_x^*(\lambda^*)\\ 0, & \textit{otherwise},\end{cases}$$

where λ^* exists and is a solution to $\mathbb{E}_{(\eta,x)}[(1-F_\eta[\eta_x^*(\lambda)|x])(\eta_x^*(\lambda)-c_x)]+s=0$. The optimal cut-offs $\{\eta_x^*(\lambda^*)\}_{x\in X}$ can be implemented by making take-it-or-leave-it price offers conditional on each observable type x, where the optimal price satisfies $t_x^*=\eta_x^*(\lambda^*)$, or

$$t_x^*=c_x+\frac{1-F_\eta[t_x^*|x]}{f_\eta[t_x^*|x]}-\frac{\sigma(t_x^*,x)b_x}{\lambda^*}. \tag{15}$$

Proof: We begin by providing a standard iff characterization of incentive compatibility in terms of the single-crossing property and the envelope theorem. We then solve the "relaxed problem" by dropping the monotonicity condition and investigating what sufficient conditions on primitives ensure that it fails to bind at the optimum.

Each household strategically chooses its report $\hat{\eta}$ to maximize

$$U(\hat{\eta},\eta,x)=p(\hat{\eta},x)(\eta-t(\hat{\eta},x)), \tag{16}$$

and define the indirect utility function

$$V(\eta,x)=\max_{\hat{\eta}} p(\hat{\eta},x)(\eta-t(\hat{\eta},x)). \tag{17}$$

This characterization of the household's problem allows for the standard characterization of incentive compatibility (see Myerson (1981) or Milgrom and Segal (2002)):

Proposition 2 *A direct mechanism* $\{p(\hat{\eta},x),t(\hat{\eta},x)\}$ *is incentive compatible iff* $\frac{\partial}{\partial\eta}V(\eta,x)=p(\eta,x)$ *and* $p(\hat{\eta},x)$ *is non-decreasing in* $\hat{\eta}$.

We drop the constraint that $p(\hat{\eta},x)$ be non-decreasing in $\hat{\eta}$ and solve the problem only requiring that $V_\eta(\eta,x)=p(\eta,x)$, and then determine sufficient conditions for $p(\hat{\eta},x)$ to be non-decreasing. The logic of the relaxed solution is that the monotonicity condition is mathematically difficult to handle (e.g. Mussa and Rosen (1978), Myerson (1981), and Rochet (1987)) and is often satisfied at the optimum if a mild regularity condition is imposed.

If $V_\eta(\eta,x)=p(\eta,x)$, then its expected payoff must satisfy $V(\eta,x)=\int_{\eta_x^*}^{\eta}p(z,x)dz$ where η_x^* is the lowest type who trades with positive probability; note that the worst-off type $\underline{w}$ is quoted a price of $\underline{w}$ with probability zero in the market and there aren't enough subsidies to cover the whole market, so that $V(\underline{w},x)=0$. In any incentive compatible mechanism, this implies $p(\eta,x)(\eta-t(\eta,x))=\int_{\eta_x^*}^{\eta}p(z,x)dz$, and a household of type (η,x) expects to pay

$$p(\eta,x)t(\eta,x)=p(\eta,x)\eta-\int_{\eta_x^*}^{\eta}p(z,x)dz. \tag{18}$$

Taking the expectation with respect to η and integrating by parts then yields

$$\int_{w_x^*}^{\overline{w}}p(\eta,x)t(\eta,x)dF_\eta[\eta|x]=\int_{w_x^*}^{\overline{w}}p(\eta,x)\eta-\int_{\eta_x^*}^{\eta}p(z,x)dzdF_\eta[\eta|x]=\int_{w_x^*}^{\overline{w}}p(\eta,x)\left\{\eta-\frac{1-F_\eta[\eta|x]}{f_\eta[\eta|x]}\right\}dF_\eta[\eta|x].$$

This expresses the expected revenue from an x-type of household entirely in terms of the probability of trade, $p(\eta,x)$. Taking the expectation over x then yields expected total revenue. The preceding arguments establish equation (11).

Dropping the monotonicity condition that $p(\eta,x)$ be non-decreasing in η, the simplified problem is to maximize quantity

$$\mathbb{E}_{(w,r,x)}[p(\min\{w,r\},x)b_x+(1-p(\min\{w,r\},x))\mathbb{I}\{w\geq r\}b_x]$$

subject to

$$\mathbb{E}_{(\eta,x)}[p(\eta,x)(\psi_\eta[\eta|x]-c_x)]+s\geq 0.$$

Consider the term $\mathbb{E}_{(w,r,x)}[(1-p(\min\{w,r\},x))\mathbb{I}\{w\geq r\}]$. Since $\eta=\min\{w,r\}$, the indicator function takes the value 1 only when $\eta=\min\{w,r\}=r$. Therefore, this term equals

$$\begin{aligned}
\int_r\int_w (1-p(\min\{w,r\},x))\mathbb{I}\{w\geq r\}dF_w[w|x]dF_r[r|x] &= \int_{\eta=\underline{w}}^{\eta=\overline{w}}\int_{w=\eta}^{\overline{w}}(1-p(\eta,x))dF_w[w|x]dF_r[\eta|x] \\
&= \int_{\eta=\underline{w}}^{\eta=\overline{w}}(1-F_w[\eta|x])(1-p(\eta,x))f_r[\eta|x]d\eta \\
&= \int_{\eta=\underline{w}}^{\eta=\overline{w}}\frac{(1-F_w[\eta|x])f_r[\eta|x]}{f_\eta[\eta|x]}(1-p(\eta,x))dF_\eta[\eta|x] \\
&= \mathbb{E}_{(\eta,x)}\left[\frac{(1-F_w[\eta|x])f_r[\eta|x]}{f_\eta[\eta|x]}(1-p(\eta,x))\right],
\end{aligned}$$

and note that

$$\frac{(1-F_w[\eta|x])f_r[\eta|x]}{f_\eta[\eta|x]}=\frac{(1-F_w[\eta|x])f_r[\eta|x]}{(1-F_w[\eta|x])f_r[\eta|x]+(1-F_r[\eta|x])f_w[\eta|x]}=\frac{h_r[\eta|x]}{h_w[\eta|x]+h_r[\eta|x]}.$$

The Lagrangian then is

$$\mathcal{L}(p,\lambda)=\mathbb{E}_{(\eta,x)}\left[p(\eta,x)b_x+(1-p(\eta,x))\frac{h_r[\eta|x]}{h_w[\eta|x]+h_r[\eta|x]}b_x\right]+\lambda\big(\mathbb{E}_{(\eta,x)}[p(\eta,x)(\psi_\eta[\eta|x]-c_x)]+s\big)$$

or

$$\begin{aligned}
\mathcal{L}(p,\lambda)=\mathbb{E}_{(\eta,x)}&\left[p(\eta,x)\frac{h_w[\eta|x]}{h_w[\eta|x]+h_r[\eta|x]}b_x+\frac{h_r[\eta|x]}{h_w[\eta|x]+h_r[\eta|x]}b_x\right] \\
&+\lambda\big(\mathbb{E}_{(\eta,x)}[p(\eta,x)(\psi_\eta[\eta|x]-c_x)]+s\big),
\end{aligned}$$

expressing the problem entirely in terms of η, which is equation (13). The objective is linear in $p(\eta,x)$, and collecting terms multiplied by $p(\eta,x)$ yields

$$\beta(\eta,x,\lambda)=\sigma(\eta,x)b_x+\lambda(\psi_\eta[\eta|x]-c_x).$$

Now if $\beta(\eta,x,\lambda)$ has the single-crossing property in η for all (x,λ) and the crossing point is increasing in η, the monotonicity condition will be satisfied. Sufficient conditions for this to hold are that $\psi_\eta[\eta|x]$ and $\sigma(\eta,x)$ both be non-decreasing in η. This means that households with higher η types are more profitable to serve on the margin, and whenever a household reports a higher type, the platform infers it is more likely to switch *conditional on* x — i.e., that households of similar socio-economic observables who report higher η are more likely to face high prices in the market and be a switcher, conditioning on x. If either of these sufficient conditions is violated, $\beta(\eta,x,\lambda)$ might still be non-decreasing in η or satisfy the single-crossing property in η. If $\beta(\eta,x,\lambda)$ exhibits violations of the single-crossing property, the monotonicity condition binds, and an optimal control approach is required; the optimal mechanism will then involve a deterministic contract for high reports of η, and a series of contracts with lower probability of service and lower prices.

To characterize λ^*, note that the optimal allocation is a cut-off rule for every x, where the cutoff is given by

$$\sigma(\eta_x^*(\lambda),x)b_x+\lambda(\psi_\eta[\eta_x^*(\lambda))|x]-c_x)=0.$$

By the implicit function theorem, $\eta_x^*(\lambda)$ is a continuous function, with derivative

$$\frac{d}{d\lambda}\eta_x^*(\lambda)=\frac{-(\psi_\eta[\eta_x^*(\lambda)|x]-c_x)}{\frac{d}{d\eta}\sigma(\eta_x^*(\lambda),x)b_x+\lambda\frac{d}{d\lambda}\psi_\eta[\eta_x^*(\lambda)|x]}$$

Now, since $\psi_\eta[\eta|x]$ is increasing in η and $w_x^*(\lambda)$ is weakly less than the monopoly cutoff where $\psi_\eta[\eta_x^m|x]-c_x=0$, the numerator is positive, and monotonicity assumptions ensure the denominator is positive. Therefore, $w_x^*(\lambda)$ is a continuous and non-decreasing function.

The budget is then given by $\phi(\lambda)=\mathbb{E}_x[(\eta_x^*(\lambda)-c_x)(1-F_\eta(\eta_x^*(\lambda))]+s$ with derivative

$$\phi'(\lambda)=\mathbb{E}_x\left[\{(1-F_\eta[\eta_x^*(\lambda)|x])-f_\eta[\eta_x^*(\lambda)|x])(\eta_x^*(\lambda)-c_x)\}\frac{d\eta_x^*(\lambda)}{d\lambda}\right]$$
$$=\mathbb{E}_x\left[-(\psi_\eta[\eta_x^*(\lambda)|x]-c_x)f_\eta[\eta_x^*(\lambda)|x]\frac{d\eta_x^*(\lambda)}{d\lambda}\right]\geq 0,$$

because, again, $\psi_\eta[\eta|x]$ is non-decreasing and $\eta_x^*(\lambda)$ is weakly less than the monopoly solution, where $\psi_\eta[\eta_x^m|x]-c_x=0$. Therefore, the budget is negative at $\lambda=0$ since $\underline{w}+s<c_x$ for all x, non-decreasing, continuous, and strictly positive as $\lambda\rightarrow\infty$. This in turn implies there exists a finite $\bar{\lambda}$ for which it is strictly positive, allowing us to restrict attention to a compact interval $[0,\bar{\lambda}]$. Therefore, by the intermediate value theorem, there exists a λ^* that balances the budget and characterizes the optimal mechanism. ■

This implies posted prices can implement the optimal separation of types, and the price in (15) can be decomposed as

$$\underbrace{t_x^*}_{\text{Price}} = \underbrace{c_x}_{\text{Marginal cost}} + \underbrace{\frac{1-F_\eta[t_x^*|x]}{f_\eta[t_x^*|x]}}_{\text{Informational Rent}} - \underbrace{\frac{\sigma(t_x^*,x)b_x}{\lambda^*}}_{\text{Social Discount}}.$$

as reported in the text. Note that the distribution of η is

$$F_\eta[\eta|x]=(1-F_w[\eta|x])F_r[\eta|x]+(1-F_r[\eta|x])F_w[\eta|x]+F_w[\eta|x]F_r[\eta|x],$$

with density

$$f_\eta[\eta|x]=(1-F_w[\eta|x])f_r[\eta|x]+(1-F_r[\eta|x])f_w[\eta|x],$$

and virtual valuation

$$\psi_\eta[\eta|x]=\eta-\frac{1-F_\eta[\eta|x]}{f_\eta[\eta|x]}=\eta-\frac{1}{\dfrac{f_w[\eta|x]}{1-F_w[\eta|x]}+\dfrac{f_r[\eta|x]}{1-F_r[\eta|x]}}=\eta-\frac{1}{h_w[\eta|x]+h_r[\eta|x]}.$$

So if the standard regularity condition that $1-F_w[w|x]$ and $1-F_r[r|x]$ are each log-concave, the associated hazard rates will be increasing, and $\psi_\eta[\eta|x]$ will be increasing in η.

C Empirical Platform Design

A key feature of the environment is that due to price dispersion as a consequence of search frictions, if a household purchases a mechanical desludging, the price it would have faced for the manual service is not observed, and vice versa, resulting in selection bias in price regressions[34]. To model the household's decision and the observed price,

[34]This is also why we did not use a standard multinomial logit model of demand: for the vast majority of households, the alternative price is not observed and there is no centralized market with stable prices. A structural approach would require instead estimating a search model that characterizes distributions of prices and the propensity to continue searching for a mechanical provider or take the option of manual desludging.

we use an endogeneous regime switching regression, or Tobit V model[35] :

$$\tilde{y}_i = x_i'\delta + \varepsilon_{0i} \tag{19}$$

$$y_i = \begin{cases} 1, & \tilde{y}_i \geq 0 \\ 0, & \tilde{y}_i < 0 \end{cases} \tag{20}$$

$$r_{mech,i} = \begin{cases} z_i'\beta_{mech} + \varepsilon_{mech,i}, & \tilde{y}_i \geq 0 \\ \varnothing, & \tilde{y}_i < 0 \end{cases} \tag{21}$$

$$r_{man,i} = \begin{cases} \varnothing, & \tilde{y}_i \geq 0 \\ z_i'\beta_{man} + \varepsilon_{man,i}, & \tilde{y}_i < 0 \end{cases}, \tag{22}$$

where the latent index, $\tilde{y}_i$, determines selection into manual[36], $y_i = 0$, or mechanical, $y_i = 1$, and the shock $\varepsilon_i = (\varepsilon_{0i}, \varepsilon_{mech,i}, \varepsilon_{man,i})$ is trivariate normal, with correlations $\rho_{x,y}$ between ε_{xi} and ε_{yi}.

To ensure that the model is not identified solely from the functional form of the structural errors, we use a restricted vector z_i of the data x_i to estimate the price equations (21) and (22), which excludes the electricity bill, the number of people in the household, the number of women in the household, and whether the respondent completed high school. Our argument that the exclusion restriction is satisfied is based on price discrimination: the excluded variables are not observable to a desludger over the phone while contracting or during a service visit, and therefore can not affect the negotiated price, but do shift the likelihood the household will purchase mechanical desludging because households that are larger, more educated, and have more female members will place more value on sanitation services on average, and electricty expenditure is not observable to a one-time visitor (but is for a municipal authority).

We estimate $(\delta, \beta_{mech}, \beta_{man})$ in equations (20), (21), and (22) by maximum likelihood, and results are reported in Table 14. Measures of wealth like electricity bill, quality of the dwelling, and respondent education have a positive and statistically significant effect on the likelihood of mechanical desludging, while households that

[35]See, for example, (Amemiya, 1985) or (Maddala, 1983).

[36]Why estimate the manual price equation at all? It exploits more decision-relevant data, allowing the selection equation to better rationalize observed choices by incorporating information about the household's perceived outside option, manual.

desludge more frequently or own their own dwelling are less likely to use mechanical. Similar patterns hold for the mechanical and manual price equations. A likelihood ratio test of a restricted model that drops the instruments against the unrestricted model rejects the hypothesis that the instruments are jointly insignificant ($\lambda = 126.56$).

The total demand for mechanical desludgings is then

$$\begin{aligned}
D(t_i,x_i) &= \mathbb{E}_{\varepsilon_i,\eta_i}\Bigg[\underbrace{\mathbb{I}\{\tilde{y}_i \geq 0 \cap t_i < r_{mech,i}\}}_{\text{Participating buyers}} + \underbrace{\mathbb{I}\{\tilde{y}_i \geq 0 \cap t_i > r_{mech,i}\}}_{\text{Non-participating buyers}} \\
&\qquad + \underbrace{\mathbb{I}\{\tilde{y}_i < 0 \cap t_i \leq \eta_i \leq r_{mech,i}\}}_{\text{Switchers}}\Bigg| x_i\Bigg] \\
&= \underbrace{\hat{pr}[y_i \geq 0, t > \eta_i | x_i]}_{\text{Participating Buyer}} + \underbrace{\hat{pr}[y_i \geq 0, t \leq \eta_i | x_i]}_{\text{Non-participating Buyer}} + \underbrace{\hat{pr}[y_i \leq 0, t \leq \eta_i | x_i]}_{\text{Switcher}}.
\end{aligned}$$

In order to derive a tractable way of computing this quantity[37], we assume the joint

[37] If X and Y are jointly normally distributed random variables with $\sigma_y = 1$, then $Y|X$ is distributed normally, with mean $\mu_y + \frac{\rho\sigma_x}{\sigma_x}(x - \mu_x)$ and variance $(1-\rho^2)$, yielding the conditional distribution $F[y|x] = \Phi\left([y - \mu_y - \frac{\rho}{\sigma_x}(x - \mu_x)]/\sqrt{1-\rho^2}\right)$. We then have, for example for the switchers,

$$\begin{aligned}
&\mathbb{E}_{\varepsilon_i,\eta_i}[\mathbb{I}\{\tilde{y}_i < 0 \cap t_i < r_{mech,i} \cap t_i \leq \eta_i \leq r_{mech,i}\}|x_i] \\
&\qquad = \mathbb{E}_{\varepsilon_i,\eta_i}[\mathbb{I}\{x_i\delta + \varepsilon_{i0} < 0 \cap t_i < r_{mech,i} \cap t_i \leq \eta_i \leq z_i\beta_{mech} + \varepsilon_{i,mech}\}|x_i] \\
&\qquad = \int_{t_i - z_i\beta_{mech}}^{\infty} \int_{-\infty}^{-x_i\delta} \int_{t_i}^{z_i\beta_{mech} + \varepsilon_{i,mech} \wedge t_i} f[\varepsilon_{i0}|\varepsilon_{i,mech}] f_\eta[\eta_i|x_i] d\eta_i d\varepsilon_{i,0} d\Phi\left(\frac{\varepsilon_{i,mech}}{\sigma_{mech}}\right) \\
&= \int_{t_i - z_i\beta_{mech}}^{\infty} \Phi\left(\frac{-x_i\delta - \frac{\rho_{0,mech}}{\sigma_{mech}}\varepsilon_{i,mech}}{\sqrt{1-\rho_{0,mech}^2}}\right) \max\{F_\eta[z_i\beta_{mech} + \varepsilon_{i,mech}|x_i] - F_\eta[t_i|x_i], 0\} d\Phi\left(\frac{\varepsilon_{i,mech}}{\sigma_{mech}}\right).
\end{aligned}$$

The details of the other calculations are similar (and simpler, because they do not include η_i). To compute this, we use a Monte Carlo approach and take a large number of draws from the distribution of residuals of $\varepsilon_{mech,i}$, use the given closed form solutions for the integrand, and average over the results from all the draws. One can think of the switchers' contribution to the sum as integrating the area under a demand curve between the prices t_i and $r_{mech,i}$ using the distribution of willingness-to-switch values F_η, then weighting the probability of switching by the probability by being a non-buyer, and finally integrating over the mechanical shock.

To estimate the conditional probability of a switch given the prices and observables, $F_\eta[r_{mech,i}|z_i] - F_\eta[t_i|z_i]$, we need only estimate $F_\eta[y|x_i] = Pr[\eta_i \leq y|x_i]$ for the set of relevant prices (those observed or those we intend to quote), then evaluate at $r_{mech,i}$ and t_i and take the difference. This only requires estimating the probability that η_i is above a given set of thresholds, for which we use a sequence of simple logit regressions, with results reported in Table 15.

density of (ε_i,η_i) takes the form $f[\varepsilon_{i0},\varepsilon_{i,mech},\eta_i|x_i] = f\varepsilon(\varepsilon_{i0},\varepsilon_{i,mech})f_\eta[\eta_i|x_i]$, so that the joint density is the product of a bivariate normal and a distribution that is independent of ε_i, conditional on x_i. Similarly, platform demand is given by

$$\begin{aligned} D^P(t_i,x_i) &= \mathbb{E}_\varepsilon\left[\underbrace{\mathbb{I}\{\tilde{y}_i \geq 0 \cap t_i < r_{mech,i}\}}_{\text{Participating buyers}} + \underbrace{\mathbb{I}\{\tilde{y}_i < 0 \cap t_i \leq \eta_i \leq r_{mech,i}\}}_{\text{Switchers}} \middle| x_i\right] \\ &= \underbrace{\hat{pr}[y_i \geq 0, t > \eta_i|x_i]}_{\text{Participating Buyer}} + \underbrace{\hat{pr}[y_i \leq 0, t \leq \eta_i|x_i]}_{\text{Switcher}}. \end{aligned}$$

computed similarly to (23). This corresponds to the probability that household i with characteristics x_i decides to purchase from the platform, and will play a key role in the constrained optimization problem that determines the prices we quote.

Putting the estimates together, the platform selects prices that maximize the share of households that purchase mechanical services subject to an expected budget balance constraint, taking as given Demand Elicitation group data, $X = \{x_i\}_{i=1}^N$, the average subsidy budget, $\bar{s} = 1{,}750$ (\$3.00), and the average cost of procuring a desludging given x, $\bar{c} = 17{,}500$ CFA:

$$\max_{t=(t_1,\ldots,t_n)} \frac{1}{N}\sum_{i=1}^N D(t_i,x_i) \tag{23}$$

subject to

$$\begin{aligned} 0 &\leq \frac{1}{N}\sum_{i=1}^N D_i^P(t_i,x_i)(t_i - \bar{c}) + \bar{s} && (24) \\ t_i &\in T = \{8{,}000, \quad 10{,}000, \quad 12{,}500, \quad 15{,}000, \quad 17{,}500, \quad 20{,}000\}. && (25) \end{aligned}$$

While the solution to the linear program (23) — (25) is in terms of individual households, $\{(t_i,x_i)\}_{i=1}^N$, we seek a pricing rule that maps observables x_i to prices, $t^*(x_i)$ for use on a new sample of households, $X' = \{x_{i'}\}_{i'=1}^{N'}$. Since the original sample, $X = \{x_i\}_{i=1}^N$, is large and random, the platform can replace the personalized prices for each household t_i^* with a function that maps characteristics x_i into prices, $t_i^* = t^*(x_i)$, and the same pricing rule should also maximize adoption of mechanical desludging across the population. To do this, we use an ordered logit model[38],

[38] We consider alternative, more algorithmic and automatic methods of assigning observables to prices in Section 4, particularly random forests.

mapping household characteristics x_i to a latent index $\tilde{t}_i$, and then use the index to assign households to price bins. Because no households were allocated to the 12,500 CFA bin and only 4% of households were allocated to the 8,000 CFA bin, we reassign them to the 10,000 CFA bin and fit the index

$$\tilde{t}_i = x_i'\gamma + \varepsilon_{t,i},$$

by maximum likelihood; results are reported in Table 13. Letting

$$\begin{aligned}
\pi_{10{,}000}(x_i) &= Pr[x_i'\hat{\gamma}+\varepsilon_{t,i}<10{,}000] = \frac{e^{10{,}000-x_i'\hat{\gamma}}}{1+e^{10{,}000-x_i'\hat{\gamma}}} \\
\pi_{15{,}000}(x_i) &= Pr[10{,}000 \leq x_i'\hat{\gamma}+\varepsilon_{t,i}<15{,}000] = \frac{e^{15{,}000-x_i'\hat{\gamma}}}{1+e^{15{,}000-x_i'\hat{\gamma}}} - \frac{e^{10{,}000-x_i'\hat{\gamma}}}{1+e^{10{,}000-x_i'\hat{\gamma}}} \\
\pi_{17{,}500}(x_i) &= Pr[15{,}000 \leq x_i'\hat{\gamma}+\varepsilon_{t,i}<17{,}500] = \frac{e^{17{,}500-x_i'\hat{\gamma}}}{1+e^{17{,}500-x_i'\hat{\gamma}}} - \frac{e^{15{,}000-x_i'\hat{\gamma}}}{1+e^{15{,}000-x_i'\hat{\gamma}}} \\
\pi_{20{,}000}(x_i) &= Pr[x_i'\hat{\gamma}+\varepsilon_{t,i}>17{,}500] = 1 - \frac{e^{17{,}500-x_i'\hat{\gamma}}}{1+e^{17{,}500-x_i'\hat{\gamma}}},
\end{aligned}$$

the assignment rule is

$$t^*(x_i) = \underset{t\in\{10000,15000,17500,20000\}}{\text{argmax}} \pi_t(x_i),$$

mapping x_i to the most likely bin under $\hat{\gamma}$. Figure 7 illustrates the linear programming and ordered logit pricing rules in the left panel, and the propensity for mis-classification in the right panel. The ordered logit rule is correct approximately 79% of the time, and within 2,500 CFA of the correct bin 92% of the time, tending to quote too many high prices; mis-classifications by the ordered logit rule should then attenuate the treatment effect.

A Demand Elicitation Script

At the end of the market survey, the enumerator reads the following script to the participant in their native language (Moore or Diola depending on the preference of the participant), and records the value that they state:

We had a study of desludging businesses in Ouagadougou, and we purchased some of their services.

We are selling the services of the desludgers that we purchased in your neighborhood and in a few other neighborhoods in Ouagadougou.

Table 13: Ordered Logit Pricing Rule

	Ordered Logit
Average months between desludgings	-0.020
Water Bill more than 5,000	0.457
House type: Precarious	-4.686
House type: Concrete	-1.556
House type: Rooming House	-1.489
Other households in compound	0.102
Own house	-1.375
Pit meters from road	0.037
More than 1 trip last desludging	1.365
Electricity Bill	0.062
Number persons in household	0.023
Number of women in household	0.087
Respondent finished high school	1.269
Constant	15.062

We are asking households for their price for the services and we will sell the services to the households that suggest the highest prices.

We would like to sell you a desludging service, but the price is not yet set.

The offer that you make for the desludging service will determine if you win and if you win the price that you pay will always be lower than what you have offered.

Here is the way we will determine who get the desludging services and how much they will pay:

I will ask you how much you are willing to pay for the desludging service.

We will leave a sticker here with the number that you can call to arrange the desludging.

When you call, the operator will compare your price to those of 8 other households who also need desludgings. There will be [randomized K number of winners] desludgings available.

The [randomized K number of winners] households that offer the highest prices will win, and each of the winners will pay the amount offered by the household that offered the highest amount but still lost.

The winners will pay for the desludging at the time that they get a desludg-

Table 14: Tobit 5 Pricing Model

	(1) Selection (δ)	(2) Mechanical Price (β_{mech})	(3) Manual Price (β_{man})
Constant	1.787***	16.601***	3.542
	(0.551)	(0.917)	(6.834)
Average months between desludgings	-0.006***	0.005	0.017*
	(0.001)	(0.005)	(0.01)
Water Bill more than 5,000	0.054	0.453	-1.341
	(0.113)	(0.412)	(0.991)
House type: Precarious	-1.936***	0.315	6.194
	(0.541)	(0.986)	(6.413)
House type: Concrete	-1.521***	0.8	5.711
	(0.527)	(0.735)	(6.241)
House type: Rooming House	-1.332**	0.92	6.595
	(0.592)	(1.152)	(6.804)
Other households in compound	0.051	-0.02	0.406
	(0.033)	(0.114)	(0.305)
Own house	-0.351**	-0.91*	-0.077
	(0.153)	(0.499)	(1.476)
Pit meters from road	-0.005	0.038	-0.016
	(0.012)	(0.041)	(0.1)
More than 1 trip last desludging	0.452	5.91***	-4.643
	(0.359)	(0.88)	(3.853)
Electricity bill	0.038***		
	(0.006)		
Household size	0.006		
	(0.013)		
Number of women in household	0.045		
	(0.034)		
Respondent finished high school	0.424***		
	(0.138)		
atanh(ρ_{mech})	-0.798***		
	(0.247)		
atanh(ρ_{man})	-0.498***		
	(0.122)		
log(σ_{mech})	8.022***		
	(0.091)		
log(σ_{man})	4.274***		
	(0.039)		
N	773	530	243
LR test statistic: $-2\ln(\lambda)$	309.438***		
LR test statistic, Instruments: $-2\ln(\lambda)$	126.563***		

Selection equation estimated from households in the Demand Elicitation group who purchased a desludging prior to our survey. Mechanical (Manual) Price equation estimated from households in the Demand Elicitation group who purchased a mechanical (manual) desludging for their most recent desluding. Omitted housing type is concrete, multi-level. Estimated by Maximum Likelihood, model given in equations (19) to (21). $\rho_{0,man}$ ($\rho_{0,mech}$) is the correlation between the manual (mechanical) price shock and the selection shock; higher prices lead to a lower likelihood of purchasing mechanical. σ_{man} (σ_{mech}) is the standard deviation of the manual (mechanical) price shock. One correlation, $\rho_{mech,man}$ is not identified by the Tobit 5 model; we compute $corr(\varepsilon_{mech,i},\varepsilon_{man,i})$ for a small number of households who did have prices for both services, and takes the value $-.01645$: the shocks are close to independent controlling for observables.

Table 15: Demand model: Logit Regressions of F_η on Price Grid T

	10,000	15,000	17,500	20,000
Constant	-8.609***	-0.037	1.712***	2.264
	(0.443)	(0.254)	(0.634)	(1.415)
Average months between desludgings	0.003	0.001	0	0.001
	(0.007)	(0.002)	(0.004)	(0.027)
Water Bill more than 5,000	-0.134	0.07	-0.193	-0.807
	(0.233)	(0.163)	(0.369)	(1.176)
House type: Precarious	7.15***	0.275	-0.094	-0.345
	(0.422)	(0.221)	(0.516)	(1.031)
House type: Concrete	7.043***	0.093	0.281	0.339
	(0.413)	(0.198)	(0.453)	(0.784)
House type: Rooming House	6.974***	0.023	0.534	0.115
	(0.86)	(0.263)	(0.638)	(1.025)
Other households in compound	-0.086	0.011	-0.002	0.117
	(0.079)	(0.049)	(0.12)	(0.402)
Own house	0.39	0.275	0.389	0.382
	(0.392)	(0.209)	(0.395)	(0.842)
Pit meters from road	-0.01	-0.032*	-0.07***	-0.076
	(0.046)	(0.019)	(0.024)	(0.075)
More than 1 trip last desludging	-6.702***	-0.438	-0.69	-0.913
	(2.087)	(0.303)	(0.541)	(0.878)
Electricity bill	-0.022	-0.013*	-0.01	0
	(0.025)	(0.007)	(0.012)	(0.02)
Household size	0.039	0.017	0	0.052
	(0.029)	(0.025)	(0.054)	(0.09)
Number of women in household	-0.12*	-0.021	0.131	0
	(0.071)	(0.06)	(0.169)	(0.322)
Respondent finished high school	-0.266	-0.12	-0.088	-0.151
	(0.439)	(0.192)	(0.375)	(0.614)
N	773	773	773	773
% above, actual	0.846	0.481	0.135	0.105
% above, predicted	0.846	0.481	0.135	0.105

Estimated by non-linear least squares, standard errors bootstrapped with 5,000 repetitions. Each column corresponds to a logit regression predicting whether the household's offer, conditional on observables, is above the value given in the first row of the table.

ing.
For example, suppose [8 minus randomized K] each offer 25,000 CFA and [randomized K minus 1] households offer 15,000 CFA.
If you were to offer more than 15,000 CFA, you would win and pay 15,000 CFA.
If you offered less than 15,000 CFA, then you would lose and you would not have access to the desludging.
Not read aloud: (If the respondent asks about ties, then the enumerator should explain that ties are resolved by randomization).
If you win, the price that you pay will always be less than the price that you offer.
You should never make an offer larger than what you would really want to pay, otherwise you could lose money.
You should never make an offer lower than what you would want to pay, because you would risk losing the opportunity to have a good price.
Is this clear to you, or would you like me to explain part of it again?
What offer would you like to make?
To be sure, if you win and the next household offers [households price minus 5%], would you want to purchase the desludging at that price?
If you lose, and you were to find out later that the price was [households price plus 5%], would you regret not having offered more?
If yes, what new offer would you like to make?

B Data Appendix

We ran three household surveys for this project: the demand elicitation survey, the baseline survey, and the endline survey. In addition, we collected data on the cost of desludgings provided through our call center, the bidding behavior of trucks in our call center, and the timing of calls by households. Each data source is described in turn below:

Sample Selection: We placed 450 evenly spaced grid points across Ouagadougou, and randomly selected 67 for the Demand Elicitation survey, 52 for the Treatment group and 40 for the Control group.[39] Prior to randomization, grid points falling in the wealthiest neighborhoods, neighborhoods that were connected to the sewer system, and neighborhoods in which property rights are not well-defined were omitted. Enumerators were sent to map the households within 100 meters of a grid point. Households were randomly selected from among mapped households for inclusion into the project. Households without latrines were excluded from the survey. Each neighborhood cluster point included approximately 30 households.

Demand Elicitation Survey: The demand elicitation survey and the incentive-compatible demand elicitation game were administered in December 2014, with 2,088 participant households selected based on their proximity to the 67 randomly selected grid points selected for the demand elicitation treatment. The survey collected data on household choices related to sanitation, and at the end of the survey households were asked to participate in a demand elicitation auction in which they were asked to bid on a desludging in a K+1 price auction developed to elicit their incentive-compatible willingness to pay for a mechanical desludging from the call center (the auction is further explained in section 2 and the script for the demand elicitation game is given in appendix A). The data collected for this survey was used primarily to inform the design of the pricing treatment.

Baseline Survey: The baseline survey took place in August and September of 2015, with 1,660 pricing treatment households and 1,284 control group households. During the survey, households were asked about their latrine pit, their sanitation practices, the process of search for a desludging operator, and their level of wealth.

Table 16 presents the balance of the treatment versus control groups, both at the household level and 17 presents the balance at the cluster level. Standard errors are clustered at the cluster level in the household balance tests. Balance is not perfect between the treatment and control groups – the control group is somewhat more wealthy than the treatment group (the average principal components index for the control group is significantly higher than both the treatment and auction group, and they are more likely to spend more than 5,000 CFA on water bills). We control for

[39]We selected more gridpoints for the auction survey in order to ensure that we would have adequate representation of different consumer types for generating the model. We increased the size of the treatment group by randomly selecting additional gridpoints after a bug in the algorithm on the enumerators' tablets was detected during the baseline.

baseline variables that are not well balanced in the OLS specifications in our main regressions. We can also see from column (3) in Table 16 that the balance between the control group and the demand eliciation group is also not perfect, and again the control group is wealthier than the demand elicitation group. To the extent that the demand elicitation group was not fully representative of the treatment and control groups for the randomized controlled trials, we would expect that the pricing model would not perform as well, biasing our estimates toward 0.

A programming error on the enumerators' tablets led to some households being offered prices higher at baseline than the model had predicted. The error occurred at 10.9% of households, and 27 of the 52 treatment neighborhoods. Nearly all of the households receiving incorrect prices received prices that were too high by 1 price bin. In cases in which the household received a price that was too high, we returned to the household to offer them the correct price, and if they had initially rejected the price offer they were given the opportunity to purchase. In the specifications in this paper, we use the price bin to which the household would have been assigned by the correct pricing system (the ITT). The results controlling for households given a different price in error are available on request.

Endline Survey: We returned to the households interviewed in both the Demand Elicitation survey and the Baseline survey in December of 2016. During the Endline survey, households were asked about their latrine pit, their sanitation practices since the Baseline or Demand Elicitation survey, their search process for a desludging operator for any desludgings that they had done over the period, the diarrhea related health of their children, and their level of wealth.

At endline, we find that the number of desludgings of any type procured in households in the control neighborhoods was 16% higher than the number of desludgings procured in households in the treatment neighborhoods (households in the control neighborhoods purchased on average 0.83 desludgings while households in the treatment neighborhoods purchased on average 0.70 desludgings). A large percentage of desludgings (71% of the most recent desludgings at baseline) are mechanical, which means that neighborhoods with more desludgings will mechanically have more mechanical desludgings. Disparities in number of desludgings across neighborhoods are directly controlled for in market share estimates which divide by the total number of desludgings at the neighborhood level thereby avoiding bias from differences in the overall use of desludgings.

One potential concern is whether the treatment may have induced treatment households to delay purchasing a desludging. We investigate this concern by analyzing survey responses. In the endline, households in the treatment group reported no difference in the number of days it took to get a desludging once they realized they needed one relative to the control group. Households and desludgers list the following factors in determining the frequency at which desludging is needed: the size and type of latrine pit; factors about the households such as the frequency with which they use water and the number of people using the latrine pit, and factors about the geography of the region including elevation, the height of the water table, and soil type.

Call Center Data: We also collected data in our call center. This included data on the supply side: the bids made by desludging operators and the winners in each month, and data on the demand side: which households called for a desludging and the date on which they called. Deposit rates by price offered and use of the call center are shown in table 18. Use of the call center was somewhat lower than predicted, but among those who purchased a desludging in the first 6 months and deposited, use of the call center was quite close to the level expected from the model (and somewhat higher among the 20k price group).

Table 16: Balance Tests: Household level

	Control(SD)	Diff Treat-Control (SE)	Diff Auct-Control (SE)
Household Size	7.862	-0.379	-0.070
	(4.66)	(0.32)	(0.30)
Number of Women in Household	2.742	-0.034	0.016
	(1.81)	(0.12)	(0.11)
Respondent Finished High School	0.278	-0.009	-0.048
	(0.45)	(0.04)	(0.35)
Precarious Housing	0.107	-0.010	0.010
	(0.31)	(0.02)	(0.026)
Concrete Building	0.795	-0.013	-0.028
	(0.40)	(0.03)	(0.029)
Rental Dormatories	0.051	0.031	0.004
	(0.22)	(0.02)	(0.01)
Own house	0.819	-0.035	-0.008
	(0.39)	(0.02)	(0.02)
Water bill more than 5,000 UGX	0.563	-0.106***	-0.132***
	(0.50)	(0.03)	(0.03)
Electricity Bill	13.808	-0.389	-0.787
	(14.19)	(0.98)	(0.96)
Pit meters from Road	5.281	-0.857**	0.371
	(3.71)	(0.30)	(0.34)
More than 1 trip last desludging	0.025	-0.010	0.001
	(0.16)	(0.01)	(0.01)
Average Months between desludgings	21.960	1.298	1.678
	(26.84)	(1.86)	(1.83)
Other households in compound	1.392	0.349*	-0.006
	(2.27)	(0.18)	(0.15)
Respondent Arranges Desludgings	0.610	-0.039	0.032
	(0.49)	(0.03)	(0.32)
Respondent is the Household Head	0.555	-0.004	0.005
	(0.50)	(0.04)	(0.04)
Years respondent lived in Compound	21.269	-1.570	0.181
	(13.92)	(1.33)	(1.24)
Number of households sharing pit	1.338	0.297	-0.126
	(2.21)	(0.17)	(0.13)
Compound has 1 pit only	0.262	-0.070**	0.485***
	(0.41)	(0.03)	(0.03)
Ever used Manual Desludging	0.544	0.037	-0.103***
	(0.50)	(0.03)	(0.03)
Ever used Mechanical Desludging	0.786	-0.048	-0.048
	(0.41)	(0.03)	(0.03)
Never desludged at this residence	0.114	0.045	0.030
	(0.32)	(0.02)	(0.02)
Percent of desludgings mech before BL	0.881	-0.015	-0.030
	(1.62)	(0.02)	(0.02)
Last Desludging was Mechanical	0.726	-0.008	-0.007
	(0.446)	(0.03)	(0.03)
Number of income earners	1.617	-0.071	-0.000
	(1.28)	(0.08)	(0.08)
Respondent Earns income	0.628	-0.026	-0.027
	(0.48)	(0.03)	(0.03)
Wealth Index (1st principal Component)	0.628	-0.224**	-0.372***
	(0.48)	(0.10)	(0.09)
N	551	648	840

Note: The first column provides the variable average and standard deviation in the control group. The second column provides the difference between the treatment group and the control group, with standard errors in parentheses. Standard errors are clustered at the neighborhood cluster level. Sample is restricted to the 1199 households that purchased any desludgings during the period to be consistent with the main regressions.

Table 17: Balance Tests: Cluster level

	Control(SD)	Diff Treat-Control (SE)
Household Size	6.798	-0.107
	(1.14)	(0.23)
Number of Women in Household	2.439	-0.045
	(0.37)	(0.08)
Respondent Finished High School	0.316	-0.037
	(0.15)	(0.03)
Precarious Housing	0.119	-0.001
	(0.11)	(0.02)
Concrete Building	0.769	0.014
	(0.12)	(0.02)
Rental Dormatories	0.046	0.014
	(0.05)	(0.01)
Own house	0.769	0.000
	(0.09)	(0.02)
water bill more than 5,000 UGX	0.490	-0.064*
	(0.13)	(0.03)
Electricity Bill	14.266	-1.744
	(5.58)	(1.03)
Pit meters from Road	5.581	-0.802*
	(1.60)	(0.34)
More than 1 trip last desludging	0.024	-0.016**
	(0.03)	(0.01)
Average Months between desludgings	27.301	2.799
	(7.98)	(1.59)
other households in compound	1.112	0.093
	(0.57)	(0.12)
Respondent is the Arranger for Desludgings	0.576	-0.018
	(0.14)	(0.03)
Respondent is the Household Head	0.549	-0.010
	(0.12)	(0.02)
Years respondent has lived in Compound	18.316	-0.315
	(5.48)	(1.06)
Number of households sharing pit	1.061	0.061
	(0.55)	(0.12)
Ever used manual desludging	0.660	0.034
	(0.11)	(0.03)
Ever used Mechanical Desludging	0.555	-0.028
	(0.17)	(0.04)
Never desludged at this residence	0.313	0.026
	(0.16)	(0.03)
Percent of desludgings mech before BL	0.884	-0.019
	(0.07)	(0.02)
Last Desludging was Mechanical	0.728	-0.034
	(0.16)	(0.04)
Number of income earners	1.502	-0.065
	(0.25)	(0.05)
Respondent Earns income	0.628	-0.031
	(0.11)	(0.02)
Wealth Index (1st principal component)	0.274	-0.255*
	(0.59)	(0.11)
$N clusters$	40	52

Note: The first column provides the mean of the variables averaged at the cluster level and standard deviation in the control group. The column provides the difference between the treatment group and the control group, with standard errors in parentheses. There are 92 neighborhoods: 40 control neighborhoods and 52 treatment neighborhoods.

Table 18: Call Center Take Up

Targeted Price Level	10000	15000	17500	20000	Total
Pct Offered Price	28	49	18	4	100
Deposited	55	52	35	38	49
Percent take-up through CC					
1st 6 months	100	58	78	75	70
Percent take-up through CC					
(from deposited and desludged)	62	47	38	47	50
Modeled Take up	94	59	33	0	58

Note: Shown are percentages of each group. "Percent offered price" is the percent of the treatment group that were offered each of the price levels in accordance with the price targeting model. "Deposited" is the percent of those offered each price who accepted the price offer and paid a deposit. "Percent take-up through Call Center 1st 6 months" is the percentage of people who called the call center from among those that ended up purchasing a desludging that called the call center at least once–separated between those who purchased a desludging in the first 6 months of the program and those that purchased a desludging at some point between baseline and endline. "Modeled take-up" is the expected level of take-up generated from the pricing model.

C Comparison of Mechanical and Manual Treatment Effects

The subsidy was directed toward mechanical desludgings so we test the direct impact of the treatment on mechanical desldugings. However, the ultimate objective of the program is to reduce manual desludgings, so we show in this appendix the corresponding impact on manual desludging. The impact of the targeted prices treatment on mechanical desludging is shown in table **??**. The magnitude of the effect on mechanical desludging is similar to the impact on manual desludging (though opposite in sign).

Targeted Price
Treat X 10k grp
Treat X 15k grp
Treat X 17.5k grp
Treat X 20k grp
10k group
15k group
17.5k group
20k group
Endline
N
R^2
mean

D Attrition

One may be concerned that we could ha

$$\begin{aligned} Deposit_i &= \mathbb{I}\{z_i\alpha_d + \beta_d t_i + \varepsilon_i\}, \\ PctMechanical_i &= \begin{cases} x_i\delta_1 + \pi t_i + \varepsilon_{i1}, & Deposit_i = 1 \\ x_i\delta_0 + \varepsilon_{i0}, & Deposit_i = 0 \end{cases} \end{aligned} \tag{26}$$

x_i is a subset of z_i and ε_i has a standard normal distribution, but the shocks $(\varepsilon_{i0}, \varepsilon_{i1})$ are correlated with ε_i but not necessarily normally distributed. Price is excluded from the $PctMechanical_i$ equation when $Deposit_i = 0$ because if the household fails to leave the deposit, it no longer has access to the platform, and the price it was quoted should no longer play a role in its service choice.

We control for variables related to wealth and past desludging behavior, including a wealth index based on a principal components analysis of the household's assets; whether the last desludging was mechanical or manual; whether the current residents ever desludged at that household; the share of desludgings in the previous five years that were mechanical; and the respondent's age. Also included is a predicted value of the household's willingness-to-switch, $\hat{\eta}$, created by estimating the Demand Elicitation group's reported willingness-to-switch values using the LASSO and then predicting values for the Targeted Pricing group; the value of the latent variable that determines the household's assignment to a pricing bin, called Weight; the enumerator's subjective assessment of whether the household's responses were believable or accurate, called Reliable Responses; and price.

We estimate α_d and β_d in (C) by maximum likelihood, with results reported in the first and second columns of Table 21, which are the coefficients and marginal effects at the mean, respectively. The marginal effect at the mean of price is $-.03$, statistically significant at the 10% level, so that households facing higher prices are indeed less likely to purchase. Respondent Age and $\hat{\eta}$ are statistically significant and positive, and a Likelihood Ratio test rejects the hypothesis that the model is jointly insignificant.

The deposit choice, however, potentially creates selection: the Targeted Pricing households who forwent the 500 CFA gift in order to secure a mechanical desludging at the price quoted might differ systematically from those households who failed to do so[40]. In order to provide unbiased estimates of the $PctMechanical_i$ equation (C), we adopt the semi-parametric approach described in (Powell, 1994) or (Newey, 2009) and evaluated empirically in (Newey et al., 1990). Using the estimated coefficients from the probit model, define

$$\hat{\xi}_i = -(z_i' \hat{\alpha}_d + \hat{\beta}_d t_i), \tag{27}$$

which can roughly be interpreted as the negative of the expected net utility of making a

[40]More formally, $\mathbb{E}[\varepsilon_{i1} | x_i, t_i, Deposit_i = 1] \neq 0$ and $\mathbb{E}[\varepsilon_{i0} | x_i, Deposit_i = 0] \neq 0$ because households who deposit are systematically more likely to purchase mechanical, so that OLS regression of $PctMechanical_i$ on x_i and t_i will result in biased estimates.

deposit, conditional on household characteristics z_i and price t_i. Plots of the empirical cumulative distribution function of $\hat{\xi}_i$ conditional on $Deposit_i=0$ and $Deposit_i=1$ are provided in Figure 13, and a Kolmogorov-Smirnoff test rejects the hypothesis of equality of the distributions from which the two samples were drawn at any conventional level of significance ($D=0.263$, p-value $=5.665\times10^{-10}$). We then regress $PctMechanical_i$ on x_i, t_i, and powers[41] of $\hat{\xi}_i$ for the group that deposited and the group that did not,

$$m_i^0 = x_i\delta_0+\sum_{\ell=1}^{L^0}\rho_\ell^0(\hat{\xi}_i)^\ell, \quad Deposit_i=0 \tag{28}$$

$$m_i^1 = x_i\delta_1+\sum_{\ell=1}^{L^1}\rho_\ell^1(\hat{\xi}_i)^\ell+\pi t_i, \quad Deposit_i=1. \tag{29}$$

Since $\hat{\xi}_i$ is a linear combination of variables in (z_i,t_i), some first-stage variables must be excluded from the second-stage in order to achieve identification. For a variable to be valid for exclusion in the second stage, it should shift the household's propensity to leave a deposit conditional on the price quote, but not their propensity to get a me-

[41] The intuition for this procedure can be seen by comparing it with the Heckman two-step approach, which begins with a discrete choice

$$d_i=\begin{cases}1, & \varepsilon_i\geq\xi_i\\ 0, & \varepsilon_i<\xi_i,\end{cases}$$

so that $Pr[d_i=1]=1-\Phi(\xi_i)$, and the second-stage equation

$$y_{ij}=x_i\delta+\varepsilon_{ij},$$

but $\mathbb{E}[\varepsilon_{ij}|d_i]\neq0$ due to selection, where $cov(\varepsilon_i,\varepsilon_{ij})\neq0$. The structural assumption of joint normality implies that this conditional expectation can be computed analytically, so that

$$\mathbb{E}[y_i|d_i]=x_i\delta+\underbrace{\rho_{\varepsilon_i,\varepsilon_{ij}}\sigma_{\varepsilon_{ij}}\lambda(\xi_i)}_{\text{Heckman Correction}},$$

where $\lambda(z)$ is the Mills ratio $\phi(z)/\Phi(z)$ or inverse Mills ratio $-\phi(z)/(1-\Phi(z))$, depending on whether the data are observed given $d_i=1$ or $d_i=0$. The semi-parametric approach "replaces" $\rho_{\varepsilon_i,\varepsilon_{ij}}\sigma_{\varepsilon_{ij}}\lambda(z)$ with a flexible polynomial,

$$\rho_{\varepsilon_i,\varepsilon_{ij}}\sigma_\varepsilon\lambda(\xi_i)\approx\sum_{\ell=1}^{L}\rho_\ell(\xi_i)^\ell\approx\mathbb{E}[\varepsilon_{ij}|d_i],$$

allowing a similar two-step approach to estimation that relaxes the structural assumption of normality and provides consistent estimates of δ.

chanical desludging conditional on the deposit choice and price quote. We use Reliable Responses — a dummy variable that takes the value 1 if the enumerator judged the household's responses to be dishonest or unreliable, thus indicating distrust or skepticism of the project and a diminished propensity to deposit — as the excluded variable.

We estimate (28) and (29) by OLS, and results are presented in Columns 3 and 4 of Table 21, where standard errors are bootstrapped[42] to account for the estimated regressors. We used LASSO to select the number of powers[43] of ξ_i. This results in the selection of two powers for m^0, although the first three powers are all significant, and only one power for m^1, although selecting zero was a possibility. The index $\hat{\xi}_i$ can be interpreted as the *negative* of the expected net benefit of depositing, so that the negative coefficients on $\hat{\xi}_i$ imply that households with higher benefit are more likely to purchase mechanical. The excluded variable, Reliable Responses, was not significant at conventional levels (p-value $=.190$), but has a large marginal effect at the mean of $-.181$. Indicators of past mechanical purchases increase the likelihood of purchasing mechanical, even if the household did not leave the deposit. Price in the m^1 regression has a negative coefficient, but it is small in magnitude and not statistically significant. We interpret these results as suggesting that households were price sensitive at the deposit stage, but conditional on depositing, were very likely to purchase mechanical even if they likely wouldn't have purchased otherwise.

Is modeling the deposit step really necessary? A simple probit regression of $PctMechanical_i$ on t_i and z_i is presented in Table 26 and the predicted counterfactual shares for different market designs in Table 23. The regression coefficients are similar to the deposit regression in Table 21, but the predicted market shares are not credible, with many values below the predicted control (compare with Table 7). Modeling the deposit step allows households to be more price elastic with respect to the platform than their demand for mechanical services, delivering more accurate counterfactuals.

Estimating $Arrange_i$ only requires estimating the mapping from prices and co-

[42]Some bootstrap samples resulted in positive price coefficients. These samples were discarded because they could not be consistent with the underlying data generating process.

[43]In practice, there is no definitive way to select the number of powers. (Newey et al., 1990) use generalized cross validation and other authors have used the Akaike or Bayesian Information Criteria. We used the LASSO to select which powers $\{\hat{\xi}_i^{\ell}\}_{\ell=1}^{25}$ to include to minimize cross-validated mean-squared error — allowing for the possibility of selecting none — which is similar to Newey et al's approach.

variates to arrangement decisions, so we use a simple probit model

$$Arrange_i = \mathbb{I}\{\alpha_a' z_i + \beta_a t_i + \varepsilon_i\} \tag{30}$$

where ε_i is distributed standard normal. In our setting, we can control directly for the latent index (Weight) that determines the group to which each household i is assigned: there are no unobserved household or product characteristics which might bias β_a, which is the usual problem in estimating demand models[44]. Results are reported in columns 5 and 6 of Table 21. Price is negative and statistically significant, with a marginal effect at the mean of -0.022 (statistically significant at the 10% level).

To provide an accurate prediction of what would have happened if the Targeted Pricing group was denied access to the platform, we estimate $PctMechanical_i$ by LASSO using the Control group data, and then predict what mechanical consumption for the Targeted Pricing group would have been. To leverage the data as much as possible, we include the same variables as in z_i and then use the LASSO to select from among 115 other control variables to minimize prediction error.

Do these models fit the data? Table 22 provides realized and predicted values for $Deposit_i$, m_i^0, m_i^1, and $Arrange_i$, with bootstrapped 90% confidence intervals[45]. For population averages, the largest deviations on average and in the 10,000 CFA bin are only 0.6 of a percentage point and 2 percentage points, respectively. The empirical and estimated mechanical market shares — our main outcome of interest — are given in Table 24, where all of the deviations are under a percentage point. The predicted average treatment effect is 4.0 percentage points and the predicted average treatment effect on the 10,000 CFA group is 10.6 percentage points, in line with estimates from the randomized controlled trial. Finally, Table 25 provides a financial statement for the platform. The platform is designed to balance the *expected* budget, but ex post losses or gains are possible. Here, the realized loss was 102 CFA on average, or about 0.18 USD per household, compared with a predicted loss of 116 CFA. Overall, the counterfactual model (C) – (29) credibly fits the quantitative and qualitative features

[44] See, for example, (Berry et al., 1995) or Petrin and Train (2010).

[45] Some bootstrap samples resulted in positive price coefficients in the m_i^1 regression. These samples were discarded because they could not be consistent with the underlying data generating process, and produced incoherent results (counterfactuals with systematically lower prices achieved lower mechanical market shares, for example). This leads to asymmetric confidence intervals around the estimated value, which are based on the variation in the data and not asymptotic formulas.

of the household decisions, market outcomes, and platform finances.

Another concern with this procedure is that the model might be extrapolating outside the domain of what the data can credibly explain. Consider Figure 14, which provides empirical cumulative distribution functions of the Demand Elicitation group's willingness-to-switch values: each group has essentially the same support, so there is substantial overlap between the groups' willingness-to-switch values. In our models, we use the Demand Elicitation data to predict willingness-to-switch values for the Targeted Pricing group households, $\hat{\eta}$, incorporating finer-grained information about their propensity to deposit or arrange into the estimation.

Figure 13: ξ Conditional on Deposit Decision

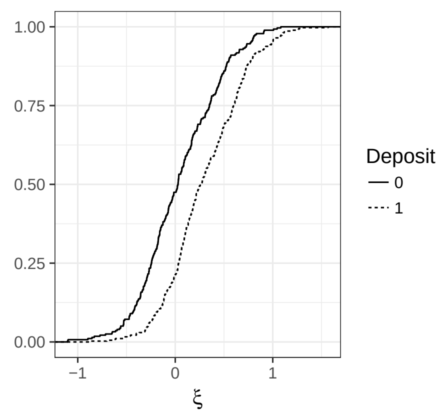

A Kolmogorov-Smirnoff test of equality of the two distributions is rejected at any conventional level of significance, with $D = .263$, corresponding to a p-value of 5.665×10^{-10}.

Figure 14: Willingness-to-switch values by price bin

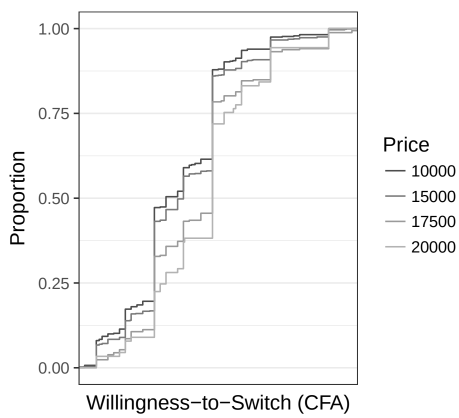

Note that while the supports of the willingness-to-switch values reported by the Demand Elicitation group have essentially the same support, the reports of the 20,000 CFA group first-order stochastically dominate those of the 17,500 CFA group, the 17,500 CFA group dominates the 15,000 CFA group, and the 15,000 CFA group dominates the 10,000 CFA group.

Table 21: Deposit, Mechanical, and Arrange Regressions

	$Deposit_i$		m_i^0	m_i^1	$Arrange_i$	
	Coefficient	M. Effect			Coefficient	M. Effect
Constant	-0.375		0.184	0.409	-1.698**	
	(0.698)		(0.211)	(0.768)	(0.742)	
Wealth Index	0.035	0.014	-0.001	0.023	-0.016	-0.004
	(0.035)	(0.014)	(0.012)	(0.026)	(0.039)	(0.011)
Last Desludging Mechanical	0.085	0.033	0.248***	0.363**	0.357*	0.092*
	(0.186)	(0.073)	(0.08)	(0.161)	(0.207)	(0.05)
Never Desludged	0.289	0.11	-0.043	0.355	0.07	0.02
	(0.186)	(0.068)	(0.091)	(0.316)	(0.213)	(0.06)
% Mechanical at Baseline	0.176	0.069	0.046	0.077	-0.169	-0.046
	(0.167)	(0.065)	(0.063)	(0.181)	(0.186)	(0.05)
Respondent Age	0.018***	0.007***	-0.002	0.004	0.011***	0.003***
	(0.004)	(0.001)	(0.003)	(0.016)	(0.004)	(0.001)
$\hat{\eta}$	0.129***	0.051***	-0.01	0.016	0.158***	0.043***
	(0.05)	(0.02)	(0.019)	(0.093)	(0.052)	(0.014)
Weight	-0.066	-0.026	0.043*	0.004	-0.045	-0.012
	(0.06)	(0.024)	(0.025)	(0.057)	(0.065)	(0.018)
Reliable Responses	-0.459	-0.181			-0.14	-0.036
	(0.351)	(0.135)			(0.424)	(0.101)
Price	-0.076*	-0.03*		-0.022	-0.08*	-0.022*
	(0.042)	(0.016)		(0.096)	(0.044)	(0.012)
$\hat{\xi}$			-0.001	-0.226		
			(0.156)	(0.835)		
$\hat{\xi}^2$			0.109			
			(0.11)			
N	648	648	278	370	648	648
R^2			0.455	0.71		
LR Stat.	63.119***				33.66***	

Columns 1 and 2 provide probit estimates and marginal effects at the mean for (C), column 3 provides OLS estimates for 28, column 4 provides OLS estimates for 29, and columns 5 and 6 provide probit estimates and marginal effects at the mean for 30. Variables m_i^0 and m_i^1 correspond to $PctMechanical_i$, conditional on depositing or not, respectively. Variable $\hat{\eta}$ is predicted willingness-to-switch, and ξ is the predicted latent variable from the deposit regression. Reliable Responses is the excluded instrument in specifications 3 and 4. Standard errors in columns 1, 2, 5, and 6 computed analytically, and in columns 3 and 4, by the bootstrap.

Table 22: Model Fit, Household Decisions

	Deposit	$\widehat{\text{Deposit}}$	m^0	$\hat{m}^0$	m^1	$\hat{m}^1$	Arrange	$\widehat{\text{Arrange}}$
Average	57.1	57	81.3	81.3	80.9	80.3	20.2	20.2
	(54.5,59.8)	(54.4,59.8)	(79.5,83.3)	(79.4,83.2)	(78.4,83.5)	(77.7,82.72)	(18.5,22)	(18.5,21.9)
10,000	70.7	72	71.7	71	61.5	63.5	30.1	31
	(67.3,74.4)	(68.4,75.8)	(67.8,75.82)	(67.7,74.8)	(54.8,68)	(57.8,69.1)	(27,33.3)	(28,34.3)
15,000	60.7	58.3	82.1	82.9	78.6	76.5	21	20
	(57.2,64.1)	(55.5,61)	(79.9,84)	(80.7,85)	(75.2,82.4)	(73.4,79.6)	(18.5,23.4)	(18.1,21.9)
17,500	38.5	43.3	89.5	88.2	89.8	90	11.5	12.8
	(34.3,43.12)	(40.2,47.4)	(85.7,93.8)	(85.9,90.4)	(86.9,92.5)	(87,91.9)	(9.3,14)	(11.1,14.9)
20,000	50	45.4	98	96.1	95.6	96	11.8	12.2
	(40.9,60)	(38.7,49.9)	(96.3,100)	(91.8,98.8)	(92.5,100)	(93.6,98.4)	(5.3,18.2)	(9,15.12)

Compares actual and predicted household desludging decisions based on counterfactual model. Variables m_i^0 and m_i^1 correspond to $PctMechanical_i$, conditional on depositing or not, respectively. Bootstrapped 90% confidence intervals reported below point estimate.

Table 23: Model Fit, Simple probit with no selection

	Realized	$\widehat{\text{Control}}$	$\widehat{\text{Treatment}}$	Auction	Auction (S)	PMT(100)	PMT(150)	Price Control	Price Control(S)
Average	81.1	76.8	80.4	76.8	79.7	52.7	65.8	75.4	78.4
10,000	68.7	58	67.3	52.1	56.3	46.4	53.7	50.4	54.6
15,000	80.7	77.9	80.6	78.5	81.7	59.1	72.2	77	80.3
17,500	89.7	86.6	88.4	90.3	92.1	42.2	59.4	89.3	91.2
20,000	96.8	96.1	95.2	97.5	98.1	60.8	77.9	97.1	97.8

Market shares for alternative platform designs, on average and by price group. Variable definitions given on page 4.1.

Table 24: Model Fit, Mechanical Shares

	Realized	$\widehat{\text{Treatment}}$	$\widehat{\text{Control}}$
Average	81.1	80.8	76.7
	(79.5,82.8)	(79.1,82.3)	(76.1,78)
10,000	68.7	68.6	57.9
	(65.38,72.4)	(65.8,71.8)	(57,61.1)
15,000	80.7	80.4	77.9
	(78.68,82.7)	(78.3,82.3)	(77,79.1)
17,500	89.7	89.3	86.5
	(87.3,91.9)	(87.3,90.7)	(85.3,88.1)
20,000	96.8	95.9	95.9
	(94.8,98.8)	(93.74,97.4)	(94.9,98.2)

Realized mechanical market shares in the intervention, predicted treatment based on counterfactual model, and predicted control based on LASSO regression on the control group predicted for the Targeted Pricing group. Bootstrapped 90% confidence intervals reported below point estimate.

Table 25: Model Fit, Finances

	Profit	$\widehat{\text{Profit}}$	Budget	$\widehat{\text{Budget}}$	SR	$\widehat{\text{SR}}$
Average	-433	-433	-102	-116	-3	-3
	(-489,-374)	(-488,-375)	(-150,-53)	(-146,-84)	(-3,-2)	(-3,-2)
10,000	-1830	-1917	-1334	-1043	-11	-12
	(-2027,-1635)	(-2114,-1732)	(-1477,-1189)	(-1179,-919)	(-13,-10)	(-13,-11)
15,000	-231	-214	116	70	-1	-1
	(-267,-189)	(-240,-185)	(91,144)	(59,85)	(-2,-1)	(-2,-1)
17,500	156	193	334	190	1	1
	(125,188)	(167,228)	(273,401)	(160,232)	(1,1)	(1,1)
20,000	495	495	689	343	3	3
	(222,764)	(367,612)	(308,1064)	(232,447)	(1,5)	(2,4)

Platform profit equals $Profit = \frac{1}{N}\sum_{i=1}^{N} Arrange_i \times (t_i - c_i)$, budget balance equals $BB = \frac{1}{N}\sum_{i=1}^{N} Arrange_i \times (t_i - c_i + s)$, and subsidization rate equals $SR = \frac{1}{N}\sum_{i=1}^{N} Arrange_i \times \frac{(t_i - c_i)}{c_i}$, where c_i is the cost of procurement for household i. Columns 1-4 are in CFA, columns 5-6 are percentages. Bootstrapped 90% confidence intervals reported below point estimate.

Table 26: Probit Regression Without Selection

	$PctMechanical_i$	M. Effect
Constant	-1.35	
	(0.946)	
Wealth Index	0.054	0.013
	(0.046)	(0.011)
Last Desludging Mechanical	0.99***	0.276***
	(0.207)	(0.064)
Never Desludged	0.398**	0.083**
	(0.189)	(0.034)
% Mechanical at Baseline	0.2	0.048
	(0.214)	(0.052)
Respondent Age	0	0
	(0.004)	(0.001)
$\hat{\eta}$	-0.012	-0.003
	(0.062)	(0.015)
Weight	0.195***	0.047***
	(0.072)	(0.017)
Reliable Responses	0.176	0.043
	(0.137)	(0.033)
Arranger	-0.17	-0.041
	(0.129)	(0.03)
Price	-0.079*	-0.019*
	(0.048)	(0.011)
N	648	648

Provides estimates of probit regression of $PctMechanical_i$ directly on covariates, without semi-parametric selection. Estimated by maximum likelihood. Variable $\hat{\eta}$ is predicted willingness-to-switch.