Pricing People into the Market: Targeting through Mechanism Design

Terence R. Johnson*
University of Notre Dame
and
Molly Lipscomb
University of Virginia

January 30, 2019

Abstract

Subsidy programs are typically accompanied by large costs due to the difficulty of screening those who should receive the program from those who would have purchased the good anyway. We design and implement a platform intended to increase the take-up of improved sanitation services by targeting the poorest households for subsidies. The project proceeds in two stages: we first create a demand model based on market data and a demand elicitation experiment, and use the model to predict prices that will maximize take-up subject to an expected budget constraint. We then test the modeled prices on a new sample of households. A main feature of the platform is that prices are designed to exclude or raise revenue from households that would likely have otherwise purchased the improved service, while channelling subsidies to households that might otherwise be unable to pay. We provide evidence that the targeting strategy successfully identified households who would otherwise have failed to purchase improved services. Households in the treatment group were 4.4 percentage points more likely to purchase a mechanical desludging, leading to an increase in market share of mechanical desludging of 6.6 percentage points. The decreased probability of purchasing a manual desludging among those with the largest subsidies was 12.4-12.7 percentage points leading to a market share increase of mechanical desludging of 10.9 percentage points in that group. The health impacts among neighborhoods with many poor households were large: a 10% increase in the number of poor households in a treatment neighborhood meant that there was a 2.2 percentage point larger decrease in diarrhea.

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

*Johnson is assistant professor of economics at the University of Notre Dame. Lipscomb is associate professor of economics and public policy at University of Virginia's Batten School and the Department of Economics. 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 Vivek Bhattacharya, Taryn Dinkelman, Pascaline Dupas, 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 Notre Dame, University of Stockholm, and William and Mary.

1 Introduction

When households exhibit low willingness or ability-to-pay for products that reduce negative externalities, subsidies and other financial assistance can be an important factor in fostering public health. Subsidies, however, often end up in the hands of households who would purchase the healthy product even in the absence of assistance, reducing the share of aid that reaches those who most benefit from it. While one strategy is to try to exclude these households from participation in the program or market, we take a different approach: first develop a model that predicts whether a given household is likely to purchase the healthy product, and then use the model to select prices that maximize expected take-up subject to a budget constraint. Rather than exclude richer households, we invite them to participate at relatively high prices, and then use the revenue from their purchases together with the subsidy budget to offer relatively low prices to poorer households.

We consider the removal of human fecal sludge from residential compounds in Burkina Faso either by mechanical means, in which a truck vacuums the sludge from a pit and minimizes exposure to the waste, or manual means, in which pits are cleaned out by hand. First, we survey one sample of households on their past experiences in the market and elicit their willingness-to-switch through a demand revelation game similar to a second-price auction. Then, we use this information to design a schedule of prices based on the kind of information available to a local municipal authority and deploy it on a second sample of comparable households. The market share for mechanical desludging increases by 6.6 percentage points among households with access to our treatment relative to a control group. The market share goes up by 10.9 percentage points for those households receiving the lowest price offered by our platform that targets the poorest households, while there is no change among those households receiving higher prices, despite the fact that they still make purchases in our market. Treatment households are 4.4 percentage points more likely to purchase a mechanical desludging, and households in the low price treatment group are 12.4-12.7 percentage points less likely to purchase a manual desludging. Mechanical desludgings increase at similar rates.

Manual desludging can have severe health consequences: rates of diarrhea are extremely high in developing countries, in large part due to lack of access to sanitation. 1.8 billion people globally use a source of drinking water with fecal contamination, and 2.4 billion people lack access to basic sanitation services (WHO and UNICEF, 2015), which can result in stunting and other developmental disadvantages (Spears, 2013). These issues have been recognized by the development community, and sanitation and water access form the sixth of the Sustainable Development Goals. While large subsidies have been effective at increasing take-up of health and sanitation goods, even when partially subsidized, the demand for health and sanitation products often remains low (Kremer and Miguel, 2007; Cohen and Dupas, 2010; Dupas, 2014). Since wealthy households who would purchase the good anyway claim a share of the subsidy budget, poor households who only purchase at subsidized prices are robbed of the opportunity to buy the product. Thus, only a fraction of a given subsidy dollar reaches its target. Even for governments seeking to maximize welfare, the large expansions of subsidy budgets necessary to ensure the poorest households receive sufficient aid might not be politically or fiscally feasible. We find that the treatment is particularly effective at lowering diarrhea rates in neighborhoods in which there are more poor households.

An attractive alternative to budget expansion is to find ways of differentiating relatively poor from relatively wealthy households, increasing the impact of subsidy dollars. Targeting in existing programs has been found to be only moderately successful: for example, Coady et al. (2004) 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); ordeal mechanisms under whipch the household must submit coupons or undergo an application process (Alatas et al., 2012; Dupas et al., 2016; Alatas et al., 2016)[1]; and community-based targeting in which members of the local community or local government select which people should receive the program (Basurto et al., 2017). Jack (2013) shows that households sometimes have private information which they can be induced to reveal through auctions for improved targeting. Chassang et al. (2017) explore different methods of eliciting private information about propensity to experiment with a new technology.

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

While these mechanisms may work well when the government has the resources to devote to a large anti-poverty program, in cases where the transfer is limited to a subsidy on a particular product, it may be possible to cross-subsidize between households by keeping the wealthier households engaged in purchasing through the platform. In this paper, we analyze whether screening for eligibility for subsidies based on limited information about households can be used to increase the take-up of a sanitation product with substantial externalities.

We approach the problem of targeting from a mechanism design perspective: if we can induce a representative group of households and firms to honestly report their willingness-to-switch, then we can design a market to maximize take-up of mechanical services subject to a budget constraint and test the market on a second group. In particular, we treat the household's willingness-to-pay and the price it expects in the prevailing market as private information, and design pricing rules that discriminate on the basis of information that is either observable or readily available to a local municipal authority. Section 3 uses a theoretical approach that analyzes the mechanism design problem when households have private information and the outside option of purchasing in the search market, and Section 4 operationalizes these insights by constructing an optimal price schedule empirically. In particular, Section 3 shows that the types we most wish to exclude are those who are on the margin between purchasing manual and mechanical desludgings but choose to purchase mechanical even without the platform, since they exercise considerable bargaining power against the platform due to their outside option. Thus, since they already purchase a mechanical desludging, they do not significantly raise take-up of the healthy product, but including them reduces the subsidy budget available to low willingnesses- or ability-to-pay households who otherwise would purchase a manual desludging. Similarly, the households we most wish to include are those who are on the margin of purchasing a mechanical desludging but do not, since a small subsidy can induce them to choose the healthy product. At the optimum, however, higher willingness-to-pay households can always adopt the strategies of lower willingness-to-pay households, so the terms of trade offered to high types must be at least as attractive to those offered to low types. We then show that the optimal mechanism can be implemented with a simple posted price scheme, where prices vary with the households' observable characteristics, optimally incorporating the likelihood of switching the household from manual to mechanical and the shadow benefit of an additional subsidy dollar. Section 4 uses these insights and data from the first stage of the intervention to solve for the optimal price schedule. Thus, the optimized posted price schedule is incentive compatible and respects the households' outside option to purchase in the existing, decentralized market in which they typically face price discrimination from mechanical desludgers.

We break the market design task into two stages. In the first stage, we invite firms to participate in neighborhood-by-neighborhood auctions in which the lowest bidders win and are paid the lowest rejected bid.[2] This gives the firms a weakly dominant strategy to bid honestly, providing us with cost estimates. On the other side of the market, we use a similar demand elicitation game, asking households to make offers for a desludging. The households making the highest offers are selected to win, but only have to pay the highest rejected bid. We then combine these household-level willingness-to-pay data with a model of mechanical and manual-price determination and market selection and survey data to derive household-level demand curves. That households can opt out of our market in favor of a prevailing decentralized market is a novel feature of our environment: most previous demand studies focus on introducing a new good or expanding demand for a health or sanitation product that is not already widely consumed. In fact, we hope to induce wealthier households to either decline our offer in favor of buying a mechanical desludging at a more attractive price in the existing market, or to purchase a desludging from us and thereby provide revenue that can be used to cross-subsidize poorer households. By purchasing desludgings in bulk at low prices through competitive mechanisms, we can undercut the high prices offered to richer households in the existing search market. This kind of targeting has uses beyond sanitation services since it explicitly explores the demand curve below prevailing prices, providing policy-makers with information about the impact and sustainability of different subsidy levels.

In the second stage, we design a pricing rule based on limited observables that are known or easily verified by a local governmental authority, like the Burkina Faso Office of Sanitation (ONEA), apply the rule to a new, comparable sample of households, and test take-up in this

[2] This is an multiple-unit procurement auction. In the case of a single unit, this would be a second price procurement auction, ie the lowest price would win and be paid the second lowest price. When there are n units, the lowest n prices will be accepted, and the price paid will be the $n+1$ lowest price.

targeted price group relative to take-up in a third sample of comparable households that serves as a control group. The pricing rule maximizes the number of households who select mechanical desludging, subject to a budget constraint that the platform's expected loss not be more than a given subsidy level. In order to determine the price level for each household, we use variables accessible to a government and easily observable at the time of our survey: water and electricity expenditure; house type (precarious, concrete structure, or rooming house); whether the house is owned or rented, number of members in the household, number of women in the household, number of other households in the compound; desludging frequency; distance from the pit to the road; and whether the respondent has a high education level. The household is made a take-it-or-leave-it-offer at the time of the baseline survey, allowing us to tailor prices to each household's individual characteristics. This is similar to recent work by Chassang et al. (2012) and Chassang et al. (2017), who apply mechanism design concepts to the design of randomized controlled trials, particularly in a development context. The approach is also similar to Wolak (2016), who uses observable information about households to design water tariffs in California. The use of demand elicitation games to measure willingness-to-pay for health products has also been used by Berry et al. (2015), who estimate the demand for water purifiers in Ghana.

We test the impact of the platform with targeted prices using a randomized controlled trial, and we find that neighborhoods with the targeting treatment have 6.6 percentage points higher market share for the improved sanitation service than neighborhoods in the control group. 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. The treatment effect acts entirely on the poorest households receiving the lowest, below market average targeted prices: while market share of mechanical desludging among the poor households in the control group is 58.9%, market share among the poor households in the treatment group is 68.2%. This substitution effect among the poor, highly subsidized households is also seen when we focus on purchases of mechanical and manual desludging at the household level: households in the treatment group were 4.4 percentage points more likely to purchase a mechanical desludging, and 2.6-3.7 percentage points less likely to purchase manual. Those with the largest subsidies were 12.4 percentage points less likely to purchase a manual desludging. This improvement in the sanitation conditions also led to a decrease in diarrhea in children in neighborhoods with more poor households: a 10% increase in the number of poor households in treated neighborhoods corresponded with a 2.2 percentage point lower probability of a report of diarrhea among children in a household relative to similar neighborhoods in the control group.

To ensure that these effects are driven by successful targeting of poor households, we compare observable assets not used in the targeting equation across households in the different price groups. We find that households that received the subsidized prices were poorer based on their ownership of typical household assets, and they were less likely to have purchased the improved sanitation good in the past as well as more likely to plan to purchase unimproved sanitation services in the future.

Finally, the question of platform sustainability is important. The simplest metric of success in this dimension is whether the platform's realized expenditure was close to the subsidies budgeted. Using the most pessimistic cost estimates based on the time series of auction prices, we find a loss of less than a dollar after the subsidy allowance of $3.00 per household. Using slightly more optimistic estimates that use negotiated prices rather than auction prices, we find a slight profit of $2 to $5 dollars (again including the budgeted $3.00 per household). Taken together, our results imply that the platform prices were surprisingly realistic. Our main design mistakes were to ignore cluster-level correlations in household demand when designing the pricing rule and not imposing higher prices on households with slightly larger pits.

This paper proceeds as follows: in section 2 we explain the sanitation problem faced by peri-urban households in Ouagadougou, Burkina Faso, the existing regulatory environment, and the trade-offs between manual and mechanical desludging. In section 3 we discuss the model and the groups that it pinpoints as pivotal in the targeting problem. In section 4 we discuss the design of the experiment: the first stage in which we collect information on a random sample of households in peri-urban Ouagadougou and collect data from a demand elicitation experiment, the design of the price targeting model, and the second stage in which we provide access to targeted subsidies to a random sample of the neighborhoods. In section 5 we provide our estimate of the total effect of the targeting program, show that the effect is on the poorest households in the low price group, and show that the reduced levels of manual desludging lead to decreased levels of children's diarrhea. We then show that the targeted subsidies were effective in constraining the amount of budget spent on the program, and we compare procurement methods. Finally, in section 6 we conclude and discuss the public policy implications of a platform on which regulators are able to provide suppliers with incentives to cooperate with regulations through providing them access to increased business.

2 Background

Lack of adequate sanitation is a primary cause of approximately 10% of global diseases, primarily through diarrheal diseases (Mara et al., 2010). While there has been substantial attention to the issue of increasing access to toilets for households in rural areas where there is not full coverage of latrines (Guiteras et al., 2015; Kar and Pasteur, 2005), there has been much less attention to sanitation issues in urban environments where the impact of inadequate sanitation may be particularly high (Coffey et al., 2014). While the coverage of latrines in urban environments is high, latrines fill between every 6 months and every 4 years, and without adequate management of the fecal sludge from the latrine, the sludge becomes a health hazard to the neighborhood.

Households can choose between two services for latrine emptying: a mechanical emptying service in which a vacuum truck comes to the household, pumps the latrine sewage into the truck's tank, and empties the tank at a treatment center, and a manual emptying service in which a worker digs a trench in the road next to the household's compound and uses buckets to transfer the sewage from the latrine into the hole in the road.

The externalities associated with manual desludging are substantial: the sewage dries over time in the street, but attracts bugs and parasites, affecting both the household itself and its neighbors. Mechanical desludging is more expensive than manual desludging (the median price of both manual and mechanical desludging is 15,000 CFA (approximately $30), but the price of mechanical second order stochastically dominates the price of manual), see Figure 1. Rather than pay for either service, the poorest households often manually desludge their own latrine pits, compounding the potential for adverse health outcomes.[3]

While many countries have instituted programs to reduce open defecation in rural areas, there has been little attention to the urban problem of inadequate disposal of latrine waste, which has very similar consequences densely populated areas. Attempts by NGOs to improve the sanitation issues caused by manual desludging have focused on heavily subsidizing as many mechanical desludgings as possible, but these programs typically run out of budget quickly.

2.1 Market Failures: Search and Market Power

The problems in the desludging market are not limited to the low ability or willingness-to-pay of households. When negotiating with a desludger, the household has to weigh the likelihood of finding someone else to do the job at a lower price with the costs of further search and the burdens of a full latrine pit, giving the desludger a limited amount of price power over the household. As a consequence, the equilibrium number of mechanical desludgings done in the market is particularly low. The high prices in the market that result from market power together with the positive externalities from use of improved sanitation means that equilibrium levels of improved sanitation are too low.

Households report several facts consistent with a market in which search and market power are important. 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. The search typically begins before the pit is completely full, but in many cases the pit fills while the household is looking for a desludger: the median household takes two days to find a desludger once their latrine pit is full, and there is a long right tail: 20% of households take 10 days or more to find a desludger once their latrine pit is already full. 30% of households report waiting because they had trouble finding a desludger or because the desludger with whom they had negotiated was not available. Financial constraints also play a factor in delays: 42% of households report waiting because they had to collect funds to pay for the desludging. On the

[3] 14.4% of households that manually desludged at endline had used family members to do so for free.

margin, households that cannot afford to pay or wait then turn to manual desludging.

Households report several different methods of finding a desludger: asking friends and family for phone numbers, using desludgers that they have used in the past, using an agent to find a desludger, going to a garage, or calling a number that they saw on a truck. 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%). Prices tend to be higher for households that use an agent (1700 CFA higher on average), call a number that they saw on a truck (945 CFA higher), or ask a desludger that they know lives nearby (500 CFA higher).

Households using manual desludging typically would have preferred to use mechanical desludging, but choose manual because of the price. At baseline, 65% of households that last used a manual desludging state that they plan to use a mechanical desludging next time, 12% had searched for a mechanical desludger prior to getting their last manual desludging, and over 60% of those who searched for a mechanical desludger and used a manual desludger report searching for a week or more before going with a manual desludger.

We seek to address these market imperfections — rents accruing to desludgers through market power arising from search frictions, adverse selection due to households' private information about their willingness-to-pay, and firms' private information about their costs of provision — by acting as an intermediary. By centralizing trade through a two-sided platform, we can eliminate search frictions and leverage competition to reduce costs, thereby increasing household welfare and consumption of the mechanical service.

3 Targeting through Mechanism Design

In general terms, "targeting" refers to the method by which beneficiaries are selected to receive aid from social programs (Alatas et al. (2016)). This encompasses a variety of approaches, from auctions to the proxy-means testing to auctions to social voting to ordeal mechanisms. Participants hold important private information about their ability- or willingness-to-pay, and eliciting this information can help market designers better utilize their limited budgets. Households, however, will only reveal their private information honestly if the program incents them to do so, leading naturally to a mechanism design analysis. This section poses and solves the mechanism design problem of a platform that competes alongside a prevailing market to maximize take-up of a socially beneficial health product, subject to standard incentive compatibility and individual rationality constraints, as well as a budget constraint requiring that its total losses not exceed a given subsidy level. The strength of this approach is that it considers all incentive compatible methods of using publicly available and privately known information to arrange trade, rather than which arrangements are optimal within a restricted class of designs, such as proxy means testing or auctions. Ultimately, we show that the platform acts as a "profit-minded social planner," optimally charging relatively high prices to households who would purchase otherwise and relatively low prices to households who will likely be unable to afford the mechanical service on their own.

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. The willingness-to-pay is the maximum price at which a household would be willing to switch from manual to mechanical desludging. The outside price 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. 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 limited price 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[4].

[4]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).

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 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[5].

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 a type $(\hat{w}, \hat{r})$ and trade occurs with probability $p(\hat{w}, \hat{r}, x)$ at a price of $t(\hat{w}, \hat{r}, x)$, where $(\hat{w}, \hat{r})$ need not be truthful. A direct mechanism is *incentive compatible* if households find it in their best interests report their types honestly, and *individually rational* if it is better to participate rather than refuse.

More formally, a direct mechanism is incentive compatible if[6] for all w, r, x, $\hat{w}$, and $\hat{r}$,

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

[5]The conclusion discusses some of the issues concerning selection onto the platform, which is an issue at scale.

[6]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.

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

$$p(w, r, x)(\min\{w, r\} - t(w, r, x)) \geq p(\hat{w}, \hat{r}, x)(\min\{w, r\} - t(\hat{w}, \hat{r}, x)). \tag{1}$$

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,

$$p(w, r, x)(\min\{w, r\} - t(w, r, x)) \geq 0. \tag{2}$$

In addition to the individual rationality and incentive compatibility constraints, the platform must also ensure that its total profits plus subsidies, s, are non-negative, or

$$\mathbb{E}_{(w,r,x)}\left[p(w, r, x)(t(w, r, x) - c_x)\right] + s \geq 0 \tag{3}$$

where c_x is the expected cost of serving a household with observables x. Call (3) 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 social benefit of a household of type x consuming the mechanical service. 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 (1), individual rationality (2), and expected budget balance (3). Due to the challenges of eliciting households' preferences over their neighbors' consumption of the mechanical service[7], we focus on maximizing the share of mechanical services, setting $b_x = 1$.

What would happen in the full information benchmark where w and r are observable by the platform? With w and r known, the platform engages in perfect price discrimination against those households who would purchase anyway and are profitable to serve, for whom $w \geq r$ and $r \geq c_x$. These profits $w - c_x$ relax the platform's budget constraint but fail to increase the share of mechanical consumption. The platform then redistributes the proceeds plus the subsidies to those households who fail to purchase mechanical, for whom $w < r$. It maximizes mechanical take-up by prioritizing those households who are the least costly to convert, in the sense that $w - c_x$ is the largest, adding marginal households until the budget constraint binds.

Once incomplete information about w and r is introduced, however, this scheme of perfect price discrimination and cross-subsidization is not possible. Households that would purchase anyway will misrepresents themselves as households that would fail to purchase, and all households would strategically understate their willingness-to-pay. This is the essence of the targeting problem: scarce subsidy dollars can end up in the hands of relatively rich households instead of relatively poor ones, and this diversion of scarce subsidy dollars can make it impossible to provide low enough prices to relatively poor households to induce them to switch.

To solve the more demanding problem with incomplete information, recall the incentive compatibility constraints (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. 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

[7] 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.

compute[8] the platform's profits in terms of the probabilities of trade in any incentive compatible direct mechanism:

$$\mathbb{E}_{(\eta,x)}\left[p(\eta,x)(t(\eta,x)-c_x)\right]=\mathbb{E}_{(\eta,x)}\left[p(\eta,x)\left\{\eta-\frac{1-F_\eta[\eta|x]}{f_\eta[\eta|x]}-c_x\right\}\right]. \tag{4}$$

The quantity

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

is called the virtual value, and represents the expected marginal revenue[9] generated by providing a desludging to a household reporting η given x. It can be understood as the total surplus η, less an informational rent that accrues to the household due to the presence of private information, $(1-F_\eta[\eta|x])/f_\eta[\eta|x]$, that captures the cost to the platform of providing incentives for honest reporting.

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

$$\begin{aligned}\mathbb{E}_{(w,r,x)}&\left[p(\min\{w,r\},x)b_x+(1-p(\min\{w,r\},x))\mathbb{I}\{w\geq r\}b_x\right]\\&=\mathbb{E}_{(\eta,x)}\left[p(\eta,x)\frac{h_w[\eta|x]}{h_w[\eta|x]+h_r[\eta|x]}b_x+\frac{h_r[\eta|x]}{h_w[\eta|x]+h_r[\eta|x]}b_x\right]\end{aligned} \tag{5}$$

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: given a report that the minimum of w and r is η, what

[8] Appendix C provides a full analysis of the mechanism design problem, starting with the standard characterization of incentive compatibility in terms of the envelope theorem and monotonicity condition, computation of profits, derivation of a relaxed solution, and verification that under some set of reasonable regularity conditions, the monotonicity condition fails to bind at the optimum.

[9] To see this connection more explicitly, consider the monopolist's problem $\max_t D(t)(t-c)$, where demand is $D(t)=1-F(t)$. The first-order necessary condition satisfies $t^*-(1-F(t^*))/f(t^*)=c$. Thus, $\psi(t)=t-(1-F(t))/f(t)$ can be interpreted as marginal revenue. For more on this interpretation, see Bulow and Klemperer (1996).

is the probability that the household is on the margin from switching from manual to mechanical, rather than purchasing in the prevailing market to purchasing on the platform? This quantifies the social marginal benefit of serving a particular household.

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

$$\mathcal{L}(p, \lambda) = \mathbb{E}_{(\eta,x)} \left[p(\eta, x) \left\{ \sigma(\eta, x) + \lambda \left(\psi_\eta[\eta|x] - c_x \right) \right\} \right] \tag{6}$$

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,

$$\beta(w, r, x, \lambda) = \underbrace{\sigma(\eta, x)}_{\text{Marginal propensity to switch at } (\eta,x)} b_x + \underbrace{\lambda}_{\text{Shadow value of profit}} \underbrace{(\psi_\eta[\eta|x] - c_x)}_{\text{Marginal profit from } (\eta,x)} . \tag{7}$$

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[10] 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*.*

[10]See Appendix C. The results are summarized in the statement of this Theorem to streamline the exposition.

For all x, there is a type $\eta_x^*(\lambda^*)$ that satisfies $\sigma(\eta_x^*(\lambda^*), x)b_x + \lambda^*(\psi_\eta[\eta_x^*(\lambda^*)|x] - c_x) = 0$, and 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 $\mathbb{E}_{(\eta,x)}\left[(1 - F_\eta[\eta_x^*(\lambda)|x])(\eta_x^*(\lambda) - c_x)\right] + s = 0$. 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

$$t_x^* = c_x + \frac{1 - F_\eta[t_x^*|x]}{f_\eta[t_x^*|x]} - \frac{\sigma(t_x^*, x)b_x}{\lambda^*} \tag{8}$$

If the platform offers an attractive offer to a type (w, r, x) household, any household of type (w', r', x) with $\min\{w', r'\} > \min\{w, r\}$ can behave as if it was a type (w, r) household and receive the same attractive offer. Thus, the terms of trade for all households with observables x must be the same as the household with the lowest η with observables x that buys through the platform, which is implementable by posted prices[11]. This implies posted prices can implement the optimal separation of types, and the price in (8) can be decomposed as

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

A standard monopolist would set its price equal to marginal 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 and contributing a social benefit b_x, which is then deflated by λ^*, representing the opportunity cost of providing assistance to this household over some other household. Households who are certain to

[11] If the virtual value $\psi_\eta[\eta|x]$ or switch function $\sigma(\eta, x)$ failed to satisfy the single-crossing property in η, it is possible that the standard monotonicity condition that $p(\eta, x)$ be non-decreasing in η would bind. This can introduce the usual pooling and randomization into the optimal mechanism, where a household of type x would be presented with a schedule of prices and probabilities of service and asked to select one. Similarly, if there were a finite number of households instead of a unit mass, the problem would more closely resemble an auction rather than non-linear pricing, and would result in more complicate pricing schemes.

purchase a mechanical desludging would have $h_w[t_x^*|x]$ close to zero, and the optimal price would be the monopoly one. Conversely, households who are most likely to fail to purchase a mechanical desludging will have $h_w[t_x^*|x]$ close to zero, and their discount will be largest.

The operation of the optimal mechanism for a given $x \in X$ is illustrated in Figure 2. In the absence of the platform, only the households for whom $w \geq r$ would make a purchase, corresponding to the set of types below the diagonal. If the platform begins offering a price t_x^*, this creates four groups:

i. *Non-buyers*: those who find neither the market nor the platform price attractive, and purchase a manual desludging ($r > w$ and $t_x^* > w$)

ii. *Non-participating buyers*: those who prefer the market price to the platform price, and purchase in the prevailing market ($w > r$ and $t_x^* > r$)

iii. *Participating buyers*: those who prefer the platform price to the market price, and purchase on the platform but would have purchased in the market ($w > r$ and $t_x^* < r$)

iv. *Switchers*: those who prefer the platform price to the market price, and would not have purchased in the prevailing market ($r > w$ and $t_x^* < w$)

While participating buyers might be contributing profits that relax the platform's budget constraint, they do not increase the share of mechanical services purchased in the market: only the set of switchers corresponds to increased social welfare. The figure also illustrates what the switch function is: the measure of households along the boundary between switchers and non-buyers, divided by the measure of households along the boundaries between switchers and non-buyers and participating buyers and non-participating buyers. This taxonomy will form the basis of our empirical design strategy in Section 4.

In the optimal mechanism, some observable types x are paying into the platform, while others are beneficiaries of it. An observable type x is *subsidized* if

$$t_x^* - c_x = \frac{1 - F_\eta[t_x^*|x]}{f_\eta[t_x^*|x]} - \frac{\sigma(t_x^*, x) b_x}{\lambda^*} \leq 0,$$

and *profitable* otherwise. Subsidization for a particular type x occurs when the shadow price of the subsidy is greater than the hazard rate of the household's willingness to pay — recalling that $b_x = 1$ — so that

$$\lambda^* < \frac{f_\eta[t_x^*|x]}{1 - F_\eta[t_x^*|x]}\sigma(t_x^*, x) = h_w[t_x^*|x],$$

and otherwise x pays into the system. If the hazard rate at t_x^* is greater than the shadow price of the subsidy, the platform's social motive dominates and every such x receives a price below the cost of procurement. If the inequality is reversed, the profit motive dominates and every such x pays into the system and relaxes the budget constraint. This concept of subsidization will play a key role in evaluating the efficacy of different market design in the counterfactual analysis of Section **??**.

This analysis of the optimal design also reveals the deficiencies of commonly used mechanisms. An auction allocates selects households entirely on the basis of willingness-to-switch, η, setting prices as $t(\eta)$ and ignoring x. This will to pass on cost savings and subsidies to relatively rich households at the expense of relatively poor households. Similarly, proxy means testing fixes a flat price t and seeks to exclude relatively wealthy households on the basis of x, rather than offering some contract to all households. This eliminates the possibility of exploiting profitable types and engaging in cross-subsidization. The optimal mechanism synthesizes features of both of these schemes to reach a superior alternative, targeted pricing.

4 Empirical Platform Design

In this section, we describe how market and experimental data were gathered and used to design the optimal platform derived in the previous section. The process has two steps.

In the first step, we gather market data on each household's most recent transactions in the decentralized market and measure each household's willingness-to-switch through a demand elicitation game similar to the second-price auction. The market data provide the platform's beliefs about the distribution of prices that a household of observable type x_i faces, and the likelihood that such a household would fail to purchase mechanical. One could then posit and estimate a structural model of decision-making to predict how a household might respond to a counterfactual change in the mechanical price it faces, based on the difference between the manual and mechanical prices. We adopt a simpler approach: we ask households their willingness-to-switch, η, in an incentive compatible way and exploit its correlation with past decisions, prices, and observables. This is accomplished through a simple game in which it is a weakly dominant strategy for households to report the price at which they are indifferent between buying from the platform and taking their outside option of finding the manual or mechanical service on their own. This allows us to assign a probability to any given household of being a non-buyer, participating buyer, non-participating buyer, or switcher.

In the second step, we use these probabilities to select prices that maximize the share of mechanical purchases subject to an expected budget balance constraint. Since the sample on which the prices are designed is large and random, the prices are also approximately optimal with respect to the population overall. This provides us with the optimal platform and project intervention, with results reported in Section 5.

The market survey and demand elicitation game were administered 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. 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. Because the demand elicitation game is a generalization of the second-price auction to allow for multiple winners, we call this the *Auction Treatment* group.

During the market survey, we gathered household characteristics x_i that would be available to a local municipal authority like ONEA — given in Table 2 — as well as information on their most recent desludging[12]. This information includes whether they purchased mechanical, $y_i = 1$,

[12] In Section **??**, we provide a counterfactual analysis of the information structure using subsets of these variables to determine which variables are the most useful and the consequences of using less information.

or manual, $y_i = 0$; the mechanical price if they purchased mechanical $r_{mech,i}$; the manual price if they purchased manual, $r_{man,i}$. We model the determination of the manual and mechanical prices in the market and the household's decision as a Type V Tobit (Amemiya, 1985) or a endogeneous regime switching regression (Maddala, 1983) :

$$\tilde{y}_i = x_i\delta + \varepsilon_{0i} \tag{9}$$

$$r_{mech,i} = \begin{cases} z_i\beta_{mech} + \varepsilon_{mech,i}, & \tilde{y}_i \geq 0 \\ \varnothing, & \tilde{y}_i < 0 \end{cases} \tag{10}$$

$$r_{man,i} = \begin{cases} \varnothing, & \tilde{y}_i \geq 0 \\ z_i\beta_{man} + \varepsilon_{man,i}, & \tilde{y}_i < 0 \end{cases}, \tag{11}$$

where the latent index, $\tilde{y}_i$, determines selection into manual[13] or mechanical, and the shock $\varepsilon_i = (\varepsilon_{0i}, \varepsilon_{mech,i}, \varepsilon_{man,i})$ is trivariate normal, so that

$$\begin{pmatrix} \varepsilon_{0i} \\ \varepsilon_{mech,i} \\ \varepsilon_{man,i} \end{pmatrix} \sim \text{Normal}\left(\mu = \begin{pmatrix} 0 \\ 0 \\ 0 \end{pmatrix}, \Sigma = \begin{bmatrix} 1 & \rho_{0,mech} & \rho_{0,man} \\ \rho_{0,mech} & \sigma^2_{mech} & \rho_{mech,man} \\ \rho_{0,man} & \rho_{mech,man} & \sigma^2_{man} \end{bmatrix} \right).$$

Only the transaction price for the kind of desludging selected is observed, not the counterfactual price that would have been charged had the household selected the other kind of service[14]. We estimate $(\delta, \beta_{mech}, \beta_{man})$ by maximum likelihood, and results are reported in Table 2.

To ensure that the model is not identified solely from the functional form of the errors, we exclude electricity expenditure, the number of people in the household, the number of women in the household, and whether or not the household head is highly educated from the price equations (10) and (11). Our argument that the exclusion restriction is satisfied is based on price discrimination:

[13]Why estimate the manual price equation? It exploits more data, allowing the selection equation to better rationalize observed choices.

[14]This is also why we did not use a standard multinomial logit model of demand model: for the vast majority of households, the alternative price is not observed and there is no centralized market with stable prices, and a standard version of such a model does not accommodate selection.

at the time of contracting, a desludger in the prevailing market might observe many characteristics about the household, especially related to water consumption and sanitation, and potentially adjust the price to extract rents. The variables excluded from the second-stage, however, are not observable to the desludger, but do shift the likelihood the household will prefer improved sanitation services: more highly educated household heads are more likely to understand the importance of health and sanitation, women typically value sanitation services at higher rates than men, electricity expenditure is unobserved by a one-time visitor, and larger households incur greater disutility from poor removal of sanitation.

This is a simple model that admits many potential improvements. By restricting the selection of x_i to variables that would be available or easily observable to a local governmental entity or NGO, we are deliberately handicapping the model so that it can only operate on information that is observable or would be costly to manipulate. This means the coefficient estimates suffer omitted variables bias, and would be inferior to a competing model that uses more covariates, particularly those that are highly correlated with wealth but difficult to observe or measure. However, the goal of this estimation is to provide a predictive tool for determining mechanical prices subject to information constraints conditional on viable observable variables, not to provide the unbiased estimates of $(\delta, \beta_{mech}, \beta_{man})$. With the rapid advances in machine learning tools and data-gathering methods over recent years, we expect there are variety of improvements that could be made on this approach.

While the model predicts the distributions of prices households would face in the decentralized market and how they would select into manual or mechanical desludging in the absence of the platform, it is a reduced-form model and cannot be used to determine how a household would respond to a price offer from the platform (at least, without extensive additional assumptions). In particular, such a model cannot predict the price at which a household that has selected a manual desludging would switch to mechanical. Instead, we supplement the market data with information from a willingness-to-switch elicitation experiment based on the second-price auction.[15] The rules of the highest-reject bid auction are as follows:

[15] The script is provided in appendix A.

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 come forward to purchase a desludging.

Since honest reporting is a weakly dominant strategy in the $K+1$-st price auction, the offer η_i provides an estimate of the minimum of the household's willingness-to-switch, the minimum of their willingness-to-pay and the price they expect to face in the prevailing market[16]. A histogram of the offers received and summary statistics are given in Figure 4 and Table 3.

Do these offers accurately reflect a household's true willingness-to-switch? We conduct a variety of thought experiments to check whether households understand their incentives to report honestly and the potential for regret if they submit a dishonest bid and lose. In particular, households were asked to confirm that they would want to purchase a desludging at a price 5% lower than their bid if that was the highest rejected bid; 2% of the households said no. They were also asked to confirm that they would not regret losing the ability to purchase a desludging at a price 5% higher than their offer if the other households were to bid higher than them and they were the highest rejected bid. 18% of households stated that they would regret losing the ability to purchase. Households stating that they would regret their bid were then allowed to revise their bids, before learning the clearing price they faced. The enumerators stated that 99.5% of households understood by the end of the exercise, though 10.5% of households required multiple explanations.

Conditional on the observables x_i, a given household might be a non-buyer, switcher, participating buyer, or non-participating buyer depending on the realization of the shocks ε_i. In particular, if the platform quotes a lower price to a household that was already planning on purchasing a mechanical desludging in the decentralized market ($\tilde{y}_i \geq 0$), it will be a participating buyer. On

[16] Households may have an intrinsically higher willingness to pay, but make lower offers because of credit constraints that constrain their access to funds in the short run. [17] While the distinction between willingness- and ability-to-pay is important for understanding potential desludging demand absent these market constraints, we argue that for our purposes, the minimum of the two is what is relevant for maximizing short-run demand.

the other hand, if the household was not planning on buying a mechanical desludging ($\tilde{y}_i < 0$), then it only switches if $t(x_i) < \eta_i = w_i < r_{mech,i}$, which is determined by the joint distribution of (η_i, x_i). These relationships are summarized in Figure 3, which corresponds to the theoretical framework given in 2. The total demand for mechanical desludgings can be posed more formally as[18]:

$$D(t_i, x_i) = \mathbb{E}_\varepsilon \Bigg[\underbrace{\mathbb{I}\{\tilde{y}_i \geq 0 \cap t_i < r_{mech,i}\}}_{\text{Participating buyers}} + \underbrace{\mathbb{I}\{\tilde{y}_i \geq 0 \cap t_i > r_{mech,i}\}}_{\text{Non-participating buyers}} + \underbrace{\mathbb{I}\{\tilde{y}_i < 0 \cap t_i < r_{mech,i} \cap t_i \leq \eta_i\}}_{\text{Switchers}} \Bigg| x_i \Bigg], \quad (12)$$

and platform demand as

$$D^P(t_i, x_i) = \mathbb{E}_\varepsilon \Bigg[\underbrace{\mathbb{I}\{\tilde{y}_i \geq 0 \cap t_i < r_{mech,i}\}}_{\text{Participating buyers}} + \underbrace{\mathbb{I}\{\tilde{y}_i < 0 \cap t_i < r_{mech,i} \cap t_i \leq \eta_i\}}_{\text{Switchers}} \Bigg| x_i \Bigg]. \quad (13)$$

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. Equations (12) and (13) highlight the importance of combining the Tobit estimation with the demand elicitation data, because the prices at which buyers are willing to participate versus not can be estimated from the Tobit model while switchers would be out of sample: those who typically purchase manual would not have market data on their mechanical purchases, and so their prices must be calculated from the auction. Estimated demand is illustrated in Figure 5: panel (a) plots a demand curve for each household in the Auction Treatment group — illustrating (12) — while panel (b) plots the average curve as well as the average curves by the price bin they are ultimately assigned to. Similarly, panels (c) and (d) plot household demand — illustrating (13) — and average demand for platform services, respectively.

To procure services and determine the cost of a desludging c_x, we used lowest reject bid auctions

[18] Appendix D shows how these quantities are computed.

on the seller side 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. We selected $K > 1$ and 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 for 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 reject bid for each job completed. These auctions were conducted monthly, and the neighborhood-average clearing price is illustrated in Figure 6. In particular, the average price in the small tank market was 17,500 CFA at the time the Targeted Pricing treatment began, requiring us to make a decision about the price, which we take as our average cost, c_x.

We can now exploit Theorem 1: the optimal mechanism is posted prices, selected to maximize the mechanical market share subject to an expected budget constraint. More formally, the platform takes the sample $X = \{x_i\}_{i=1}^N$, the subsidy s, and the average cost of procuring a desludging c_x as given, and maximizes market demand,

$$\max_{t=(t_1,\dots,t_n)} \frac{1}{N} \sum_{i=1}^{N} D(t_i, x_i) \tag{14}$$

subject to

$$\begin{aligned} 0 &\leq \frac{1}{N} \sum_{i=1}^{N} D_i^P(t_i, x_i)(t_i - c_x) + s && (15) \\ t_i &\in T = \{8{,}000, \quad 10{,}000, \quad 12{,}500, \quad 15{,}000, \quad 17{,}500, \quad 20{,}000\}. && (16) \end{aligned}$$

The set of prices T spans the observed transaction prices in the market data, and are the most commonly used denominations for payment[19]. The average subsidy per household was 1,750 CFA or about \$3.00, and we used an expected procurement cost of 17,500 CFA. Imposing the constraint (16) converts the maximization problem into a linear programming problem where each household i is assigned to a price t_i.

7 4

Since the original sample, $X = \{x_i\}_{i=1}^{N}$, is 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 maximizing adoption of mechanical desludging across the population, by a weak law of large numbers:

$$\max_{\{t^*(x)\}_{x \in X}} \mathbb{E}_x[D(t^*(x), x)] \tag{17}$$

subject to

$$0 \leq \mathbb{E}_x[D_i^P(t^*(x), x)(t^*(x) - c_x)] + s \tag{18}$$

$$t_i \in T = \{8{,}000, \quad 10{,}000, \quad 12{,}500, \quad 15{,}000, \quad 17{,}500, \quad 20{,}000\}. \tag{19}$$

While the solution to the linear program (14) — (16) is in terms of individual households, $\{(t_i, x_i)\}_{i=1}^{N}$, the solution to (17) — (19) is a mapping from observables x_i to prices, $t^*(x_i)$. Because we wish to apply the optimal rule to a new set of households, we must convert the first kind of solution into the second. To do this, we use an ordered logit model, mapping household characteristics to a latent index, and then using the index to assign households to price bins. We removed the $8,000$ and $12,500$ bins, since less than X% of the households allocated to them and

[19] From an experimental perspective, having more prices results in more treatment arms, creating a trade-off between power and treatment complexity.

created multiple modes in the distribution of prices[20] and fit the model

$$\tilde{t}_i = x_i \gamma + \varepsilon_{t,i},$$

by maximum likelihood; results are reported in Table 5. The assignment is then

$$t^*(x_i) = \begin{cases} 10,000, & x_i\hat{\gamma} < 10,000 \\ 15,000, & 10,000 \leq x_i\hat{\gamma} < 15,000 \\ 17,500, & 15,000 \leq x_i\hat{\gamma} < 17,500 \\ 20,000, & x_i\hat{\gamma} > 17,500. \end{cases} \tag{20}$$

This resembles proxy means testing, but is constructed not by maximizing a classification target like the fraction of households in a training data set below the poverty line, but instead by approximating our optimal pricing schedule. 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, and average deviations of the ordered logit from the linear programming price given in Table 4. The dark bars represent the optimal pricing rule and the light bars represent the proportion of these price quotes in the ordered logit approximation. The optimal pricing rule assigns very few households to the 8,000 CFA bin, making it very hard for the model to fit them, so we shift these households to the 10,000 CFA bin. It turns out that it is never optimal to offer 12,500: this is too high a price to induce a household to switch to mechanical, and too low to relax the budget constraint. The ordered logit tends to make too few 10,000 offers and too many 17,500 offers, but is exactly approximately 79% of the time, and within 2,500 CFA of the correct bin 92% of the time. Conditional on the bin, Table 4 shows that the ordered logit rule tends to overcharge households assigned to the 10,000 bin by 440 CFA on average, while it tends to undercharge households assigned to the 20,000 bin by 2,000 CFA on average. Consequently, mis-classifications by the ordered logit rule should be expected to attenuate the treatment effect and lead to higher budget deficits

[20] In practice, ordered logit and random forest models struggle to fit sparsely populated bins, creating more error overall than if they are simply eliminated.

than expected. We consider alternative methods of assigning observables to prices in Section **??** and Appendix D.

Our intervention is then to offer the pricing rule $t = t^*(x_{i'})$ to a new sample, $X' = \{x_{i'}\}_{i'=1}^{N'}$, under the same circumstances: a household survey is conducted, the results are recorded on a tablet computer, and in the background $x_{i'}$ is used to compute a price $t_{i'} = t^*(x_{i'})$. Take up of this targeted price group is then compared to the take up of a control group in a randomized controlled trial. We refer to this group as the *Targeting Pricing* treatment group.

How might the process we have just described be done at scale? Of particular concern is that a household who anticipates that its reports of η_i and $(r_{mech,i}, r_{man,i})$ will determine its future payoffs has an incentive to lie. There are straightforward ways 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. 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. This maintains incentives for honest reporting, both within each period and across periods.

5 Experimental Results

5.1 Data and Experimental Design

5.1.1 Sampling and Data

We sampled for the three treatments (auction survey, targeted price treatment, and control), by evenly spacing 450 grid points across Ouagadougou. 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 remaining grid points prior to the survey, maintaining a maximum distance of the center grid point to the household in order to avoid any overlap between neighborhoods. Grid point-centered neighborhoods were then randomly selected for each treatment arm from the remaining points. Households were then randomly selected to participate from among the mapped households.

The market survey and demand elicitation game were run in December 2014 with 2088 households selected near 67 grid points. In the baseline survey, run from July through September 2015, neighborhoods were randomized into either the targeted prices group (1,660 households in 52 neighborhoods) or the control group (1,284 households in 40 neighborhoods). Randomization was done at the neighborhood level, with stratification by number of households in the neighborhood with low walls (a proxy for low income).

We show that the targeted prices treatment group and the control group are similar on a variety of household level and neighborhood-averaged level observables in tables 6 and 7. There are five variables on which the two groups are not balanced in the household data: the treatment group had a lower water bill on average, a smaller distance from the latrine pit to the road, was more likely to have more than one latrine pit, was more likely to have other families living in the compound, and was less likely to be wealthy based on a first principal components index. In the neighborhood averages (seen in table 7) treatment households were also less likely to have needed multiple truck loads to do their last mechanical desludging.[21] We control for these variables in the main regressions.

In addition, treatment and control households are similar in terms of their past choices of mechanical versus manual desludging. Fifty-four percent of households have ever used manual desludging in the past, and 79 percent have ever used a mechanical desludging in the past. Households who had ever purchased a desludging estimated that across all their past desludgings they had used mechanical 88 percent of the time on average, but based on their last desludging, they were using mechanical approximately 73 percent of the time. However, desludging practices vary substantially across the income distribution. Among the poorest households which we target for the largest subsidies, at baseline only 45 percent of households used mechanical for their last desludging. 77 percent of those in the group which received only small subsidies had last used

[21] With the exception of the fact that they were more likely to have multiple latrine pits, the lack of balance in these variables suggests that the households in the control group may have been somewhat more wealthy than those in the treatment group and the auctions group, which would tend to bias our estimates toward 0.

mechanical, and 86 percent of those which were offered desludgings at the platform's expected cost had last used mechanical. 94 percent of the wealthiest households which we assigned to the highest price bin had last used mechanical.

5.1.2 Targeted Prices Treatment

All households were given a participation gift of 500 CFA at the end of the baseline survey. Targeted price treatment households were asked to use the participation gift as a deposit on their desludging if they wanted to reserve a desludging at the target price they were offered. They were then told to call in when they were ready for their desludging, and by providing their member number they could receive the desludging at the price that they were offered during the survey. 763 households (49%) agreed to leave the 500 CFA ($0.80) deposit at baseline. Most households declining to deposit at baseline stated that the primary reason they rejected the offer was that they did not expect to need a desludging, though at high price levels, some participants stated that they thought that they could get better prices elsewhere. In table 8, we see that households with the high subsidies were more likely to leave deposits, but some who received high prices still left a deposit. There was no mention of the call center made to control group households. At endline, no control group households stated that they called the call center when looking for a desludger.

Of the 404 treatment group households that needed a desludging over the year between the baseline and endline survey and paid the deposit, 147 reported calling the call center at endline [22]. Table 8 shows the use of the call center by price group from among those who deposited and purchased a desludging during the period within the first 6 months and over the period as a whole. 62% of the households who paid deposits in the lowest price group called the call center for a desludging when they needed it, while the rates among the higher price groups were lower: 47% among the 15,000 CFA price group purchased a desludging from the call center, and 38% in the 17,500 CFA and 47% in the 20,000 CFA price groups purchased a desludging from the call center.[23]

[22] 173 households from the targeted prices treatment actually called the call center based on administrative data– we attribute the difference to different members of the household answering the endline survey in cases where the person who arranged the desludging was not available at endline

[23] Most households who paid the deposit but failed to use the service did not end up needing a desludging. Reasons given by households that failed to call the center are given in table 9. Some households who called the call center

Use of the call center in the first 6 months of operation among those who needed a desludging early in the program was much higher, suggesting that increased advertising could have increased recall of the availability of the call center services and improved take-up.

While the rate of use of the call center was somewhat lower than expected, there are several potential mechanisms in addition to households purchasing desludgings through the call center that may have increased the use of mechanical desludging. Households may have used the price offered by the call center in order to negotiate with desludgers in the market. There is some evidence of this; at endline, households in the treatment group who did not purchase their desludging through the call center report paying a price of 1,120 CFA (approximately $2 less than the control group on average (this is statistically significant at the 5% level). Households in the treatment group which haven't desludged in the past may also have updated their beliefs on the attainability or the importance of mechanical desludging following the interview even if they did not deposit with us at the time of the baseline if they didn't expect to need to desludge. Finally, the program may provide neighborhoods with a stronger community interest in keeping the neighborhood clean including more peer pressure to take up mechanized desludging. This final explanation can not be tested with our data since prices within communities were allocated by our program rather than being randomized.

5.2 Overall Impact of the Targeted Prices Treatment

The objective of a subsidy program is to incent households to switch to the subsidized good over the commonly used alternative. We can observe the impact of a subsidy program by observing the change in the market share of the subsidized good. It is calculated as:

$$MarketShare = \frac{NumberMechanicalDesludgings}{NumberMechanical + NumberManual} \tag{21}$$

Market share is estimated at the level of each of 92 neighborhoods of 25-40 households. It is calculated as the sum of all mechanical desludgings done in the neighborhood divided by the

did not end up using it.

sum of all desludgings done in the neighborhood.[24] An increase in the market share of mechanical desludgings indicates that households are substituting from manual desludgings toward mechanical desludgings. Market share has several nice properties that we are interested in for the purposes of this paper. First, desludgings are something that households purchase with low frequency. While the median household stated at baseline that they need a desludging every 12 months, some households need desludgings considerably less frequently: the 75th percentile household stated that they needed a desludging approximately every 3 years.[25] During our sample period of 15 months, only 40% of households purchased a desludging, with the relatively dry weather that year being the most likely explanation for the decrease in frequency. Survey evidence suggests that it is unlikely that households strategically time desludgings; seventy-eight percent of households stated that they got a desludging within one week of noticing that their latrine pit was full, 96% stated that they had gotten a desludging within one month of noticing that their latrine pit was full. Households avoid getting desludgings unless they are necessary: in order to open the pit, they typically have to crack the cement cover, the sludge smells, and some sludge gets around the courtyard even when the household has hired a mechanical truck. Yet when they do need a desludging, it comes up as an urgent necessity–once it is full or nearly full, the latrine smells and the household can not use their latrine until it has been emptied. It is therefore unlikely that our project would have had any impact on desludgings besides through substitution between manual and mechanical. The number of desludgings that a household purchases are therefore censored because we have the decisions over a limited time period, but market share provides us with the full picture of the substitution of manual for mechanical even with censoring.

In addition, market share is a useful indicator as there is a pre-existing market for both mechanical desludgings and manual desludgings. The goal of our subsidy program was to reduce the number of households getting manual desludgings by providing the households most at risk of getting a manual desludging with as large of a subsidy as possible. In contrast to much of the

[24]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 (2017) or Nevo (2001).

[25]Variation in desludging frequency typically is related to the type of pit (whether it is cracked so that it seeps into the ground, whether it is connected to a septic system, the size of pit), and the number of people using the pit.

other literature on subsidies of purchases of chlorinated water, water filters, or new medicines, we are *not* trying to convince households to purchase more desludgings–only to substitute a mechanical desludging for any manual desludging that they would have purchased. The market share of mechanical desludgings is therefore the best indicator of the extent to which this has occurred.

In this section we test the effect of the targeted prices on the market share of mechanical desludging, both at the full neighborhood level, and by the price level that the household *would have received* if it were in the treatment group. We then run household level regressions on the impact of the program on the household level correlate of market share–the percentage of desludgings they purchased over the 15 months that were mechanical.

In order to look at the effects by price level, we also calculate market share by price group: we sum all of the mechanical desludgings done by households receiving a given price in a neighborhood and divide by total number of desludgings done by households in the neighborhood in that price group (in the case of the control group, we sum all of the mechanical desludgings done by households who *would have received* the price in the neighborhood if the neighborhood had received the fixed-price treatment and divide by the sum of the desludgings done by households who *would have received* the price).

We use the following empirical specification for estimation of the overall effect of the targeted prices treatment on market share. Estimates are shown for the neighborhood level regression in specifications (2) and (3) in table 10 and for the household level in specifications (1) and (2) in table 11:

$$MarketShareMechanical_i = \alpha + \beta TargetedPricesTreatment_i + \gamma' X_i + \varepsilon_i \qquad (22)$$

X_i includes each of the 6 variables on which the sample was not well balanced at the cluster level with neighborhood means used in the neighborhood level regression and the household level variables used in the household regression. These include water bill more than 5,000 CFA, latrine pit distance to road, other households in compound, two tanks used last desludging, compound has one pit only, and wealth index). Also included are whether the neighborhood had an above-median number of low walls during the mapping phase (the variable on which the clusters were stratified). Finally, we follow the advice of McKenzie (2012) for ANCOVA estimation in order to improve the efficiency of the estimates, and include controls for past desludging choices (the percent of desludgings prior to the baseline that were mechanical, whether the household has ever deslduged, and whether the household used mechanical for their last desludging). A programming error on the enumerators' tablets led to some households being offered prices higher at baseline than the model had predicted.[26] We include controls for the percent of households in each price group who received an incorrect price in the neighborhood level regression and a dummy for the incorrect price in the household level regression.

We also estimate the model using post-double-selection lasso (Belloni et al., 2014; Ahrens et al., 2018) in order to allow the model to flexibly control for any pre-existing differences at baseline between the control and treatment neighborhoods and avoid any concern of the possibility of specification search. The LASSO (Least Absolute Shrinkage and Selection Operator, (Tibshirani, 1996)) regression selects control variables from the full set of potential controls in order to minimize the potential for either over or underestimating the effect size. We include 115 potential control variables[27] from the baseline survey as potential controls, from which the LASSO algorithm selected five in the neighborhood level regression and 12 in the household level regression.

Estimates are shown in table 10 for neighborhoods and table 11. For reference, we also present the endline market shares in the control group by the price group the households *would have received* if they were in the treatment. Households targeted with high prices are overwhelmingly those who would have purchased mechanical in the market anyway: the market share at endline in the control group for the highest price group is 99.3% compared with 58.9% for the lowest price group. This limits the potential treatment effect in the highest price bins, and motivates our

[26]The error occurred at 10.9% of households, and 27 of the 52 treatment neighborhoods. Most 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.

[27]In cases in which an observation of a variable is missing, either because the respondent declined to answer or the respondent did not know the answer, the missing observation was replaced with the mean value and an indicator variable was included which takes a value of 1 when the observation is missing and 0 otherwise.Each of the non-binary control variables has been standardized by subtracting the mean and dividing by its the standard deviation.

focus on the lower price bins. We expect to see the largest effects on the heavily subsidized group which has a substantially lower market share for mechanical desludging. This is the group that contained the highest share of potential *switchers*, and therefore had the most capacity to change market share. We expect the pooled effects to be relatively small since only 27% of the sample were targeted for the lowest subsidized prices.

The price targeting treatment generates an increase in the neighborhood market share of the improved desludging service of 6.6% percentage points in the OLS regression (significant at the 5% level) and 5.2 percentage points in the LASSO regression. This is a 7.7% effect at the mean mechanical desludging market share of 85%. The estimated effects at the household level are somewhat smaller: 4 percentage points in the OLS specification and 3.7 percentage points in the LASSO specification.

5.3 Market Share Effects by Price Group

One key implication of the pricing model is that because take-up among the wealthiest households is already high, and some households (the $non-buyers$) will not take up mechanical desludging even when offered large subsidies, any impact of the system must take place through increases in take-up among the *switchers*. We observe the treatment effects by group by comparing mean market shares by neighborhood and price group for low price households versus high price households. To construct counterfactuals for the treatment group, we calculate the price that the households in the control group *would have received* through the platform given their characteristics. This allows us to construct market shares for households who would have received the same prices in treatment and control neighborhoods. While the market share of mechanical desludging for high price households is 99.3% in the control group at endline, the market share of mechanical desludging among households receiving the most subsidized prices was 58.9% and the market share among households receiving the second most subsidized prices was 85.2%.

This highlights that the treatment effect must be coming from the change in market share for the low price groups, not a change in behavior by the high-price groups which already include primarily $always-takers$. By targeting the lowest prices to these households, we are able to target *switchers* and induce changes in take-up at high rates. We test the effect on market share by price group using the following specification for the effects across price groups:

$$
\begin{aligned}
MktShareMech\_PriceGrp_{ki} = & \sum_{k=1}^{4} \alpha_k PriceGroup_{ki} \\
& + \sum_{k=1}^{4} \beta_k TargetedPricesTreatment_{ki} * PriceGroup_{ki} + \gamma' X_{ki} + \varepsilon_{ki}
\end{aligned} \tag{23}
$$

For 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 $10,000$, $15,000$, $17,500$, or $20,000$) divided by the total number of desludgings purchased by households of that price group in that neighborhood. We omit the constant in order to include indicator variables for each price group, and we control for neighborhood-price group average values for the variables that were not well balanced at baseline, the stratification variable of above median number of high walls in the neighborhood. and the percentage of the households that initially erroneously received the incorrect price at baseline ($\gamma' X_{ki}$). We run the related household level regressions using the percentage of the desludgings that the household purchased which were mechanical over the 15 months as the dependent variable.

Results are shown in table 10. Our coefficients of interest are the β_k's, the estimates of the effect of the targeted prices treatment on the market share for each price group. We find that the market share for the $10,000$ price group increases by 10.9 percentage points in the OLS specification, 9 percentage points in the LASSO specification (significant at 5% level in the OLS regression, significant at the 10% level in the LASSO specification). The treatment in the other price bins has small coefficients, none of which are statistically significant. This differentially large effect among the 10,000 CFA price group was expected as they were receiving the subsidies and included more potential *switchers* (the market share of the sample in the control group which would have received a price of 10,000 CFA was 58.9%).

5.3.1 Impact on use of any mechanical and any manual desludgings

While market share provides a nice summary measure of the extent of switching toward the use of mechanical desludging, we may also be interested in understanding the extent to which the effect is coming from changes in manual versus mechanical desludgings. The treatment allowed for a subsidy only on the first mechanical desludging that the household purchased, so we measure whether the household purchased at least one mechanical desludging during the period: "Any Mechanical." This is a good measure of the impact of the program since the program could not directly affect the purchases of mechanical desludgings after the first. The objective of the program is to induce households not to use manual desludgings, so we also measure the effect on "Any Manual;" whether the household purchased at least one manual desludging between the baseline and the endline. In both the regressions on any mechanical and any manual, we restrict the sample to households which purchased at least one desludging: this is done both to maintain a consistent sample size with the regressions on percent mechanical and to avoid downward bias from the households which did not need a desludging over the time period.

The results are shown in table 12. The overall probability that a household purchases a mechanical desludging increases for the targeted prices treatment group by 4.2 percentage points in the OLS regression (significant at the 10% level) or 3.2 percentage points in the LASSO regression (not significant), and manual decreases at similar rates (2.6-3.3 percentage points, not significant). Similar to the results for the percentage of desludgings that were mechanical, effects are largest in the lowest price bin: the probability that a household purchases a mechanical desludging if they purchase any desludgings increases by 9.5 percentage points (not significant) in the OLS regression, 10.2 percentage points in the LASSO regression (significant at the 10% level). The probability that a household in the 10,000 CFA price bin purchases a manual desludging decreases by 12.5 percentage points (significant at the 10% level), 12.2 percentage points in the LASSO specification (significant at the 5% level).

5.4 Health Impacts

The ultimate goal of the program was to reduce the use of manual desludging in order to improve local sanitation. We test the impact of the targeted subsidies on the use of manual desludging and the effect on child diarrhea rates. We focus on children as they are the most sensitive to health issues in the environment, and their long term health is most likely to be affected by bouts of diarrhea. We use the following specification to estimate the pooled effect across price groups:

$$AnyChildrenDiarrhea_i = \alpha + \beta TargetedPricesTreatment_i + \gamma' X_i + \varepsilon_i \tag{24}$$

Observations are at the household level and standard errors are clustered at the neighborhood-cluster level (92 clusters in total). $TargetedPricesTreatment_i$ takes the value of 1 for all households in treatment clusters, 0 for households in control clusters. X_i is a vector of controls including the same variables in our main specifications: variables unbalanced across neighborhoods at baseline, the stratification variable, and an indicator for whether the household was offered an incorrect price in the baseline. We also control for whether the household had a child suffering from diarrhea at baseline, following McKenzie (2012).

Results are shown in table 13 columns (1) and (2). 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 percentage point decrease in the probability that a household uses a manual desludging in both the OLS and LASSO specifications (not significant). At the mean of 9.7% of households reporting that at least one of their children had diarrhea in the last week at baseline, this is a 12.4 percent effect at the mean (but not statistically significant).

We are also interested in the difference in effects between the price groups. We run the following specification to estimate the differential effects on the lowest price group:

$$\begin{aligned} AnyChildrenDiarrhea_i = & \sum_{k=1}^{4} \alpha_k PriceGroup_{ki} \\ & + \sum_{k=1}^{4} \beta_k TargetedPricesTreatment_i * PriceGroup_{ki} + \gamma' X_i + \varepsilon_i \end{aligned} \tag{25}$$

$PriceGroup_{ki}$ is an indicator for the price level to which household i was assigned (or, for the control group, the price level to which household i *would have been* assigned according to the pricing function used in the treatment to assign prices): $10,000$, $15,000$, $17,500$, or $20,000$. Our coefficients of interest are the β_k's, or the coefficients on the interaction between the price groups and the treatment indicator.

Results are shown in table 13, columns (3) and (4). The low price households in the treatment group are 5.9 percentage points less likely to report that a child in their household had diarrhea than a low price household in the control group (significant at the 10% level in the OLS regression. In the LASSO regression the coefficient is 5 percentage points, but not statistically significant). At the mean of 12.04 percent of households which would have received the low price reporting diarrhea at baseline, the OLS specification suggests a 47.6 percent effect. Note that there is an important timing difference: the question asks about diarrhea in the past week, while desludging by the household or households in the neighborhood could have been done at any time over the 15 months of the program. The effect of desludging choices made by the household and its neighbors just prior to the endline should have substantially more impact than the choices made in the beginning of the time period. We therefore expect any effects on children's diarrhea to be a lower bound.

Children contract illnesses based on the sanitation situation of the neighborhood, not just the sanitation decisions made by the household. We also look at the spillover effects of sanitation decisions within the neighborhood based on the percent of households in the neighborhood that falls 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:

$$
\begin{aligned}
AnyChildrenDiarrhea_i = &\sum_{k=1}^{4} \alpha_k PriceGroup_{ki} + \sum_{k=1}^{4} \Theta_k PctPriceGroup_{ki} \\
&+ \sum_{k=1}^{4} \beta_k TargetedPricesTreatment_i * PctPriceGroup_{ki} + \gamma' X_i + \varepsilon_i \quad (26)
\end{aligned}
$$

The results for the spillovers model are presented in table 13, columns (5) and (6). The diarrhea effects of the treatment are focused on neighborhoods with more poor households–households that would receive a price of 10,000 in the treatment. We can see that a 10% increase in households that would receive 10,000 (about 3 households with our neighborhood sizes of 25-40), leads to a 2.0-2.3 percentage point reduction in reports of a child with diarrhea among households in the neighborhood. Recall that at the baseline mean of 9.7 percent of households reporting an episode of child diarrhea, this is a substantial effect.

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. They could find only 4 sanitation studies, and report an average relative risk ratio[28] following sanitation treatments of 0.68. As expected from the point estimates, the relative risk ratio for our pooled sample is 0.89, which is close to 1 suggesting little impact in the pooled sample as a whole. However, when the sample is constrained to the households which would receive a price of $10,000$, the relative risk ratio for this group is 0.64. This is a large effect: the diarrhea rates are self reports of households over diarrhea in their children under 12 in the past week, while manual desludgings in their neighborhood could have taken place at any time during the treatment period. Fixing an existing toilet so that it can be used by the members of the household therefore has similar impacts to providing households with toilets at the low end of the income distribution.

5.5 Who Receives the Subsidies?

One potential concern is that since the demand model was estimated on a different sample than the sample used to test the model, the model fit on the new test sample could be poor, resulting in either inclusion errors in which wealthy households receive subsidies or exclusion errors in which poor households are not offered subsidies. In this subsection, we compare households which receive large subsidies through the pricing model to households which receive the highest prices from the model.

[28]The relative risk ratio is the ratio of the rates in the control and treatment groups: $\frac{Outcome_{Treated}}{Outcome_{Control}}$

We can observe the extent to which the model is targeting relatively poor households who are more likely to get manual desludgings in table 14. Households that receive a price of 10,000 CFA (approximately $20, and subsidized by approximately $10) spend an average of 2,200 CFA per week on phone credit while households that receive a price of 17,500 spend nearly twice that, an average of 4,512 CFA for those receiving 17,500 and 5,631 per week for those receiving 20,000 CFA. On average, approximately one quarter of the households receiving a price of 10,000 CFA have a refrigerator, while households receiving 20,000 CFA as their price have on average 1.5 refrigerators. Motorcycles are the most common type of transport in Ouagadougou, and we see that again the households receiving the largest subsidies have fewer motorcycles on average (1.8) than the households receiving no subsidies (2.2 on average for those receiving a price of 17,500 and 3.1 for those receiving a price of 20,000 CFA). We see very similar trends for other asset markers of wealth: cars, televisions, mobile phones, and air conditioners.

We see similar differences in terms of key summary statistics about the household's use of desludging services. Households in the highest subsidy group get desludgings the most infrequently (just under four years between desludgings, while households in the 20,000 CFA price group get desludgings just less than once per year).

One way that the platform could reduce the budget needed to subsidize the poorest households would be to cross-subsidize with desludgings done on wealthier households. This would be possible if the platform receives lower prices in procurement than the households. There are two potential ways that this could happen. First, because the platform buys in bulk, it may be able to bid down prices among the desludging operators. Second, wealthier households may face price discrimination as desludgers may take advantage of the lower price elasticity of demand of wealthier households and charge them higher prices. In table 14 we provide suggestive evidence that this second mechanism does occur: households in the 10,000 CFA group report expected prices of 13,850 for their next manual desludging and 14,300 for their next manual desludging while households in the 20,000 CFA price group report expected prices of 16,600 for manual and 16,200 for their next mechanical desludging.

We also see lower take-up of mechanical desludging among those in the most highly subsidized price group: 80% of those in the 10,000 CFA price group state that they expect their next desludging will be mechanical, while 89% in the highest price group state that they expect their next desludging will be mechanical. The differences are even larger when we compare the last desludging of each group: 69% of households in the lowest price group got a mechanical desludging for their last desludging while 94% of those in the highest price group purchased a mechanical desludging for their previous desludging. If we compare manual desludgings in the past, we see that the gap widens even further: 76% of households in the lowest price group have ever purchased a manual desludging, while only 40% of households in the highest price group have ever purchased a manual desludging.

6 Conclusion

We show how call center platforms can be designed and implemented to target subsidies to poor households, ensuring that aid reaches those who most need it. This is accomplished in the presence of a prevailing decentralized market, where consumers can opt out depending on the prices we offer. Subsidies can be more effectively employed to raise take-up of key products and services with externalities if a data-driven approach is adopted so that only those who would not have purchased the good are able to receive the subsidy.

In addition to allowing the government to encourage take-up on the demand side, platforms such as this are extremely 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 correct quality service so that they can continue to operate with the platform. Using the carrot of additional business through engagement with the platform allows government regulators to oversee the operations of suppliers much more effectively than if they need to find and police operators on their own.

While this paper focuses on a platform operated by a local government, 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 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.

We also show that simplified procurement policies can be the most cost effective. While we might expect auctions to be the best way to procure the services of decentralized suppliers, we find that a simple negotiation rule is more effective in decreasing the prices over time and maintaining competitive prices. Prices went down by an average of 9% when the platform used structured negotiations for procurement relative to both first and second price auctions.

We see model and implementation criticism as an important feature of a project such as this. While the demand model was deliberately selected to be simple and a workhorse economic model, we might have exploited other tools to deliberately focus on prediction rather than point estimation, drawing on the machine learning literature. A particularly difficult parameter to estimate is the correlation between mechanical and manual price shocks, which is typically unidentified in the Type V Tobit model. Our solution was to estimate this parameter off the subset of households that recalled both mechanical and manual prices for the last job, but there is obviously selection into this group, since "shoppers" will likely get lower prices. In addition, the linear program did not account for correlation in shocks between households in clusters. We see this as explaining a large proportion of deviations in stage two outcomes from the stage one estimates. Finally, the large pit market was generally more costly and prevalent than we initially thought. We designed parallel markets to separate out these costs, but did not separate out the households on the demand side. This is a short list of short-comings, but we hope to further investigate and refine the methodology by exploiting the availability of the datasets from the two stages, as was done in the counterfactual subsidy exercise.

References

Ahrens, A., Hansen, C., and Schaffer, M. (2018). Pdslasso: Stata module for post-selection and post-regularization ols or iv estimation and inference. *Statistical Software Components S458459*.

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.

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.

Basurto, P., Dupas, P., and Robinson, J. (2017). Decentralization and efficiency of subsidy targeting: Evidence from chiefs in rural Malawi.

Belloni, A., Chernozhukov, V., and Hansen, C. (2014). Inference on treatment effects after selection among high-dimensional controls. *The Review of Economic Studies*, 81(2):608–650.

Berry, J., Fischer, G., and Guiteras, R. (2015). Eliciting and utilizing willingness to pay: Evidence from field trials in northern Ghana. *Working Paper*.

Bulow, J. and Klemperer, P. (1996). Auctions versus negotiations. *American Economic Review*, 86(1):180–194.

Chassang, S., Dupas, P., and Snowberg, E. (2017). Mechanism design meets development: Selective trials for technology diffusion. *Working Paper*.

Chassang, S., Miquel, G. P. I., and Snowberg, E. (2012). Selective trials: A principal-agent approach to randomized controlled experiments. *American Economic Review*, 102(4):1279–1309.

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

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.

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.

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.

Jack, K. (2013). Private information and the allocation of land use subsidies in Malawi. *American Economic Journal: Applied*, 5(3):113–135.

Jensen, R. and Miller, N. (2017). Information, demand and the growth of firms: Evidence from a natural experiment in India. *Working Paper*.

Kar, K. and Pasteur, K. (2005). Subsidy or self respect? Community led total sanitation: an update on recent developments. *IDS Working Paper*, (257).

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

Kremer, M. and Miguel, E. (2007). 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).

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

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.

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

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*.

Tibshirani, R. (1996). Regression shrinkage via the lasso. *Journal of the Royal Statistical Society*, 58(1):267–288.

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. *Working Paper*.

Yishay, A. B., Fraker, A., Guiteras, R., Palloni, G., Shah, N. B., Shirrell, S., and Wang, P. (2017). Microcredit and willingness to pay for environmental quality: Evidence from a randomized-controlled trial of finance for sanitation in rural Cambodia. *Journal of Environmental Economics and Management*, 86:121–140.

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.

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 desludging.

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 Differences in number of desludgings used since baseline by households

A large percentage of desludgings (71% at baseline) are mechanical, which means that neighborhoods with more desludgings will mechanically have more mechanical desludgings if the total number of desludgings are not included as a control variable. At endline, we find that the number of desludgings 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), which is significant at the 10% level.

The difference in the mean number of desludgings between the treatment and control groups also demonstrates that we need to control for number of desludgings in regressions testing the effect of treatment on desludgings at the household level. 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, therefore this is our preferred specification.

At endline we asked households a number of questions about the state of their latrine pit and whether they delayed desludgings over the treatment period. Households in the treatment group reported no difference in the number of days it took to get a desludging relative to the control group. Mean number of days to desludging is 8, and treatment households take 0.39 days less to get a desludging (not significant–p value of 0.97). From among those who did delay their desludging by more than 7 days, we find that households in the treatment group were 8.8% less likely to delay their desludging due to lack of funds (significant at the 5% level), but 1.1% more likely to delay their desludging due to accessibility issues (significant at the 10% level) and 3.9% due to difficulties in coordinating with the desludger (significant at the 10% level).

When asked what pushes households to get more desludgings, households and desludgers typically respond that the frequency with which households need desludgings depends on factors about the latrine pit such as: the size and type of latrine pit that the household has; factors about the households such as: the frequency with which they use water and the number of people using the latrine pit, factors about the geography of the region (which can vary substantially across a city including: elevation, the height of the water table, and soil type. Many of these factors are not known to the household and are not readily available (few households are able to tell us the size of their latrine pit–only 467 of 2944 households gave an answer to the question, and many of the sizes reported are far outside standard sizes so are likely to be incorrect). We find that 31% of the variation in the number of desludgings that the household gets during the treatment period can be explained by a combination of the household baseline variables and geographic variables about the area. When we use these variables as controls as an alternative to controlling directly for the number of desludgings that the household purchased during the treatment period, the point estimate on the treatment effect increases but the standard errors also increase (which is to be expected with a less precise control).

C Theory Appendix

Proof of Theorem 1:

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

$$U(\hat{\eta}, \eta, x) = p(\hat{\eta}, x)(\eta - t(\hat{\eta}, x)), \tag{27}$$

and define the indirect utility function

$$V(\eta, x) = \max_{\hat{\eta}} p(\hat{\eta}, x)(\eta - t(\hat{\eta}, x)). \tag{28}$$

This characterization of the household's problem allows for a simple characterization of incentive compatibility:

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* $\hat{\eta}$*.*

Proof: Assume the direct mechanism is incentive compatible. From the Milgrom-Segal envelope theorem (Milgrom and Segal (2002)), $V_\eta(\eta,x)=p(\eta,x)$. Taking two revealed-preference constraints

$$p(\eta,x)\eta-t(\eta,x)\geq p(\eta',x)\eta-t(\eta',x),\quad p(\eta',x)\eta'-t(\eta',x)\geq p(\eta,x)\eta'-t(\eta,x)$$

and re-arranging them yields

$$(\eta-\eta')(p(\eta,x)-p(\eta',x))\geq 0,$$

so that if $\eta>\eta'$, $p(\eta,x)\geq p(\eta',x)$, and $p(\eta,x)$ is non-decreasing in η.

Now assume $V_\eta(\eta,x)=p(\eta,x)$ and $p(\eta,x)$ is non-decreasing in η. Then

$$\begin{aligned}
U(\eta,\eta,x)-U(\eta',\eta,x) &= U(\eta,\eta,x)-U(\eta',\eta',x)+U(\eta',\eta',x)-U(\eta',\eta,x)\\
&= \int_{\eta'}^{\eta}p(z,x)dz+\int_{\eta}^{\eta'}p(\eta',x)dz\\
&= \int_{\eta'}^{\eta}p(z,x)-p(\eta',x)dz,
\end{aligned}$$

where the second line follows from $V_\eta(\eta,x)=p(\eta,x)$. Now, since $p(\eta,x)$ is non-decreasing, the integrand on the third line is positive whenever $\eta>\eta'$, and negative whenever $\eta<\eta'$, so that the third line is always weakly positive. Therefore, the mechanism is incentive compatible. ■

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

$$p(\eta, x)t(\eta, x) = p(\eta, x)\eta - \int_{\eta_x^*}^{\eta} p(z, x)dz. \tag{29}$$

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 (4).

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.

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

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

subject to

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

Consider the term $\mathbb{E}_{(w,r,x)}\left[(1-p(\min\{w,r\},x))\mathbb{I}\{w \geq r\}\right]$. 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

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

or

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

expressing the problem entirely in terms of η, which is equation (6). The objective is linear in

$p(\eta, x)$, and collecting terms multiplied by $p(\eta, x)$ yields

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

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

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

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

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

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 $\phi(\lambda) = \mathbb{E}_x\left[(\eta_x^*(\lambda) - c_x)(1 - F_\eta(\eta_x^*(\lambda))\right] + s$ with derivative

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

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 $\lambda \to \infty$. Therefore, by the intermediate value theorem, there exists a λ^* that balances the budget and characterizes the optimal mechanism. ■

D Empirical Design Appendix

E Tables

Table 1: Baseline Prices of Mechanical and Manual Services

	Min	1Q	Median	Mean	3Q	Max
All:	5000	15000	15000	16262.613	20000	30000
Manual:	5000	10000	15000	15530.864	20000	30000
Mechanical:	10000	15000	15000	16598.113	17500	30000

Table 2: Demand model: Tobit

	(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)
Desludging frequency	-0.006***	0.005	0.017*
	(0.001)	(0.005)	(0.01)
Water Bill > 5k	0.054	0.453	-1.341
	(0.113)	(0.412)	(0.991)
House type 1: Precarious	-1.936***	0.315	6.194
	(0.541)	(0.986)	(6.413)
House type 2: Concrete	-1.521***	0.8	5.711
	(0.527)	(0.735)	(6.24)
House type 4: 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 distance to road	-0.005	0.038	-0.016
	(0.012)	(0.041)	(0.1)
Trips last > 1	0.452	5.91***	-4.643
	(0.359)	(0.88)	(3.853)
Electricity Bill	0.038***		
	(0.006)		
Hhd size	0.006		
	(0.013)		
Hhd women	0.045		
	(0.034)		
Respondent finished secondary edu	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: $-2\ln(\lambda)$	445.802***		

Selection equation estimated from households in the Auction Treatment group who purchased a desludging prior to our survey. Mechanical (Manual) Price equation estimated from households in the Auction Treatment 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 (9) to (11).

Table 3: Offers, Summary Statistics

	Min	1Q	Median	Mean	3Q	Max
Offer, η:	2	10	12.5	12.995	15	40

Table 4: Deviations of Ordered Logit from Linear Programming Pricing Rule

	Min	1Q	Median	Mean	3Q	Max
Average (CFA):	-5000	0	0	446.24	0	5000
10k (CFA):	0	0	0	1370.56	5000	5000
15k (CFA):	-5000	0	0	189.3	0	5000
17.5k (CFA):	-2500	0	0	303.57	0	2500
20k (CFA):	-2500	-2500	-2500	-2013.89	-2500	0

Table 5: Ordered Logit Pricing Rule

	Ordered Logit
Desludging frequency (months)	-0.020
Water greater than 5,000 CFA	0.457
Precarious House	-4.686
Concrete House, 1 story	-1.556
Rooming House	-1.489
Other households in compound	0.102
Own house	-1.375
Pit distance to road	0.037
Last trips greater than one	1.365
Electricity Bill	0.062
Number persons in household	0.023
Number women in household	0.087
Household head educated	1.269
Constant	15.062

Table 6: 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 7: 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)
Compound has 1 pit only	0.335	-0.062*
	(0.13)	(0.03)
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

Table 8: 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.

Table 9: Reasons Households did not Call the Call Center

	Targeted Price
Didn't need a desludging	368
Forgot about it	60
Better Outside option	59
Too Confusing/didn't understand	46
New to the compound	24
Not in charge of desludging	32
Other/refusal	20
Total	606

Note: Households were able to select multiple responses. Sample restricted to treatment households that paid a deposit at baseline but did not use the call center between baseline and endline.

Table 10: Market Share Effects of Treatment

Dependent Variable:	Market share	Market share		Mkt share, cluster-price level	
	Control Group, at EL	OLS	LASSO	OLS	LASSO
Overall	0.840	0.066**	0.052**		
		(0.029)	(0.024)		
10k group×TP Group				0.109**	0.090*
				(0.050)	(0.048)
15k group×TP Group				0.028	0.028
				(0.047)	(0.044)
17.5k group×TP Group				0.018	0.006
				(0.052)	(0.051)
20k group×TP Group				0.014	-0.001
				(0.080)	(0.078)
10k group	0.589			0.281***	0.000
				(0.098)	(.)
15k group	0.852			0.410***	0.127
				(0.116)	(0.094)
17.5k group	0.906			0.427***	0.093
				(0.128)	(0.143)
20k group	0.993			0.468***	0.109
				(0.148)	(0.217)
N		92	92	300	300
R^2		0.354		0.942	
mean		0.852	0.852	0.832	0.832

Notes: column (1) gives the average neighborhood market share for each price group of mechanical desludging at baseline where market share is defined as $\frac{nmechanicaldesludgings}{Nmechanical+manualdesludgings}$. Column (2) provides the OLS estimate of the pooled effect with observations at the neighborhood cluster level. Column (3) provides the LASSO estimate of the pooled effect with observations at the cluster level, with 115 potential control variables. Column (4) gives the OLS estimate for the market share effect for each price group in a neighborhood cluster (level of observation is neighborhood-price, but not all neighborhoods include households from each price group). Column (5) gives the LASSO estimate for the market share effect for each price group in a neighborhood cluster, with 113 potential control variables. Controls are included in the OLS regressions for the neighborhood means (in column 2) or neighborhood-price group means (in column 3) of the variables not balanced at baseline at the household level (water bill more than 5,000 CFA, latrine pit distance to road, two tanks used last desludging, compound only has 1 pit, other households in compound, wealth index). Percent of desludgings for which the household purchased mechanical, last desludging mechanical (dummy) and never desludged (dummy) are included as controls for past desludging behavior of the household. A control has also been included for the average of the stratification variable–above median number of households in neighborhood have high compound walls. Controls have also been included for the percent of households in each price group that mistakenly received the wrong price due to a programming error on enumerators' tablets.

Table 11: Household Level Market Share: Effects on the Percent of Mechanical at the Household level

	(1) Percent Mechanical OLS	(2) Percent Mechanical LASSO	(3) Percent Mechanical OLS	(4) Percent Mechanical LASSO
Targeted price group	0.040	0.038*		
	(0.024)	(0.022)		
Treatment*10k price group			0.102	0.111*
			(0.063)	(0.058)
Treatment*15k price group			0.011	0.022
			(0.031)	(0.031)
Treatment*17.5k price group			0.032	0.031
			(0.035)	(0.033)
Treatment*20k price group			0.029	0.047
			(0.059)	(0.056)
10k price group			0.428***	0.497***
			(0.057)	(0.183)
15k price group			0.506***	0.465***
			(0.058)	(0.172)
17k price group			0.573***	0.501***
			(0.062)	(0.176)
20k price group			0.576***	0.459**
			(0.078)	(0.187)
N	1199	1199	1199	1199
R^2	0.212		0.856	
mean	0.804	0.804	0.804	0.804

Notes: Observations are at the household level. The constant is dropped in specifications (3) and (4). Included in the OLS specifications, but not shown for brevity are controls for the variables not balanced at baseline: water bill more than 5,000 CFA, latrine pit distance to road, two tanks used last desludging, compound only has 1 pit, other households in compound, and wealth index. Percent of desludgings for which the household purchased mechanical, last desludging mechanical (dummy) and never desludged (dummy) are included as controls for past desludging behavior of the household. The lasso specifications have 115 potential control variables. A control for the stratification variable: less than half of compound walls in the neighborhood are high, is also included but not shown. Indicator variables for each of the price groups when a household was assigned an incorrect price in error are also included. Standard errors, clustered by neighborhood, are in parentheses.

Table 12: Decomposition of Market Share: Effects on purchases of Mechanical and Manual

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
	Any Mechanical	Any Mechanical	Any Mechanical	Any Mechanical	Any Manual	Any Manual	Any Manual	Any Manual
	OLS	LASSO	OLS	LASSO	OLS	LASSO	OLS	LASSO
Targeted price group	0.044*	0.039*			-0.036	-0.027		
	(0.024)	(0.023)			(0.023)	(0.022)		
Treatment*10k price group			0.094	0.102*			-0.124*	-0.127**
			(0.064)	(0.060)			(0.066)	(0.060)
Treatment*15k price group			0.023	0.031			-0.003	-0.013
			(0.032)	(0.032)			(0.030)	(0.030)
Treatment*17.5k price group			0.018	0.011			-0.039	-0.047
			(0.032)	(0.030)			(0.035)	(0.035)
Treatment*20k price group			0.065	0.086			0.053	0.042
			(0.055)	(0.053)			(0.066)	(0.065)
10k price group			0.441***	0.463**			0.595***	0.011
			(0.059)	(0.185)			(0.059)	(0.092)
15k price group			0.521***	0.442**			0.496***	0.022
			(0.060)	(0.173)			(0.058)	(0.071)
17k price group			0.601***	0.497***			0.445***	-0.007
			(0.064)	(0.174)			(0.065)	(0.050)
20k price group			0.579***	0.434**			0.400***	0.000
			(0.080)	(0.186)			(0.076)	(.)
N	1199	1199	1199	1199	1199	1199	1199	1199
R^2	0.201		0.861		0.182		0.350	
mean	0.804	0.804	0.804	0.804	0.804	0.196	0.804	0.196

Notes: Observations are at the household level. Included in the OLS specifications, but not shown for brevity are controls for the variables not balanced at baseline: water bill more than 5,000 CFA, latrine pit distance to road, two tanks used last desludging, compound only has 1 pit, other households in compound, and wealth index. Percent of desludgings for which the household purchased mechanical, last desludging mechanical (dummy) and never desludged (dummy) are included as controls for past desludging behavior of the household. A control for the stratification variable: less than half of compound walls in the neighborhood are high, is also included but not shown. The lasso specification has 115 potential control variables. Indicator variables for each of the price groups when a household was assigned an incorrect price in error are also included. Standard errors, clustered by neighborhood, are in parentheses.

Table 13: Impact of Treatment on Children's Diarrhea

	Dependent Variable: Any Child Diarrhea					
	Pooled		By Price Group		Spillovers	
	OLS	LASSO	OLS	LASSO	OLS	LASSO
	(1)	(2)	(3)	(4)	(5)	(6)
Targeted price group	-0.013	-0.014				
	(0.017)	(0.017)				
Treatment*10k price group			-0.059*	-0.046		
			(0.030)	(0.032)		
Treatment*15k price group			0.010	0.012		
			(0.021)	(0.020)		
Treatment*17.5k price group			-0.004	0.002		
			(0.036)	(0.035)		
Treatment*20k price group			0.016	0.009		
			(0.059)	(0.057)		
10k price group			0.115***	-0.072		
			(0.036)	(0.142)		
15k price group			0.077**	-0.087		
			(0.038)	(0.137)		
17k price group			0.110**	-0.030		
			(0.045)	(0.138)		
20k price group			0.102*	-0.024		
			(0.058)	(0.130)		
Treatment*Pct Nbhd in 10k group					-0.217**	-0.208**
					(0.096)	(0.095)
Treatment*Pct Nbhd in 15k group					0.085	0.091
					(0.075)	(0.074)
Treatment*Pct Nbhd in 17k group					-0.037	-0.059
					(0.144)	(0.137)
Treatment*Pct Nbhd in 20k group					0.241	0.297
					(0.313)	(0.297)
Pct Nbhd in 10k group					0.656	0.863
					(0.623)	(0.650)
Pct Nbhd in 15k group					0.507	0.717
					(0.632)	(0.666)
Pct Nbhd in 17.5k group					0.620	0.835
					(0.636)	(0.665)
Pct Nbhd in 20k group					0.276	0.435
					(0.652)	(0.671)
N	2017	2017	2017	2017	2017	2017
R^2	0.027		0.157		0.031	
mean	0.131	0.131	0.131	0.131	0.131	0.131

Notes: observations are at the household level, standard errors are clustered at the neighborhood cluster level. The dependent variable is whether the participant reports that a child in the household has had diarrhea in the past 7 days at endline. Specifications (1) and (2) are the pooled effect across all price groups. Odd specifications are estimated with OLS, even specifications with Lasso. Specifications (3) and (4) provide the estimates of the effect for each price group. Specifications (5) and (6) show the spillover effects of the treatment within the neighborhood: the regressions control for the percent of the neighborhood that would be assigned to each price group, and estimate the effect of the treatment for each price group. The sample includes only households with children. The diarrhea question is posed as follows: "In the past seven days, of the children in your household, how many had diarrhea, even once?" Children are defined in the survey as being 14 and younger, if the number is 1 or more, the variable is classified as a 1, otherwise 0. Included, but not shown for brevity are controls for the variables not balanced at baseline: water bill more than 5,000 CFA, latrine pit distance to road, two tanks used last desludging, compound only has 1 pit, other households in compound, and wealth index. Percent of desludgings for which the household purchased mechanical, last desludging mechanical (dummy) and never desludged (dummy) are included as controls for past desludging behavior of the household. A control for the stratification variable: less than half of compound walls in the neighborhood are high, is also included but not shown. Indicator variables for each of the price groups when a household was assigned an incorrect price in error are also included. Standard errors, clustered by neighborhood, are in parentheses.

Table 14: 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)

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

Table 15: Mechanical share fit

	Ex Post	Predicted	Control (Str)	Control (ML)
Average	0.811	0.798	0.783	0.78
10,000	0.687	0.663	0.612	0.606
15,000	0.807	0.804	0.791	0.797
17,500	0.897	0.876	0.89	0.861
20,000	0.968	0.918	0.905	0.932

Table 16: Mechanical share fit, given selection

	Ex Post	Predicted
Average	0.809	0.798
10,000	0.615	0.607
15,000	0.786	0.782
17,500	0.898	0.884
20,000	0.956	0.903

	Ex Post	Predicted
Average	0.813	0.809
10,000	0.717	0.711
15,000	0.821	0.824
17,500	0.895	0.873
20,000	0.98	0.95

Table 17: Counterfactual Mechanical Share

	Ex Post	Hat	Control (Str)	Market	Market(S)	Auction	Auction(S)	PMT
Average	0.811	0.798	0.709	0.734	0.76	0.72	0.745	0.744
10,000	0.687	0.663	0.544	0.553	0.582	0.539	0.565	0.578
15,000	0.807	0.804	0.745	0.773	0.802	0.757	0.785	0.787
17,500	0.897	0.876	0.838	0.871	0.889	0.86	0.88	0.855
20,000	0.968	0.918	0.86	0.934	0.943	0.926	0.938	0.908

Table 18: Subsidization Rates

	Ex Post	Hat	Market	Market(S)	Auction	Auction(S)	PMT
Average	-1.381	-1.452	-0.648	-2.221	0	-1.296	-1.356
10,000	-4.296	-4.668	-0.614	-1.999	0	-1.106	-1.596
15,000	-0.683	-0.625	-0.625	-2.147	0	-1.251	-1.404
17,500	0.53	0.652	-0.698	-2.481	0	-1.506	-0.834
20,000	1.528	1.108	-0.921	-3.424	0	-2.167	-1.473

Table 19: Profit in USD per Household

	Ex Post	Hat	Market	Market(S)	Auction	Auction(S)	PMT
Average	-0.407	-0.43	-0.192	-0.653	0	-0.378	-0.399
10,000	-1.26	-1.377	-0.183	-0.591	0	-0.325	-0.472
15,000	-0.203	-0.186	-0.186	-0.631	0	-0.365	-0.413
17,500	0.153	0.187	-0.207	-0.726	0	-0.437	-0.244
20,000	0.437	0.318	-0.271	-0.996	0	-0.626	-0.427

Table 20: η Regression

	Dependent variable:
	η
matrix(1, length(y), 1)	5.921*** (0.937)
desludging_frequency_tc	−0.004 (0.004)
water_bill_g5k	0.106 (0.225)
house_type_dum1	0.376 (0.549)
house_type_dum2	0.363 (0.466)
house_type_dum4	−0.123 (0.649)
othr_hhds_compound_tc	0.063 (0.069)
own_rent	−0.336 (0.276)
trips_last_greater_than1	−1.640** (0.746)
electricity_bill	0.006 (0.010)
household_size_tc	−0.065* (0.037)
household_women	0.099 (0.079)
pit_distance_to_road_tc	0.091*** (0.022)
total_rooms	0.034 (0.035)
ac	0.475*** (0.176)
mechanical_walk_away_price	0.309*** (0.028)
resp_age	−0.019** (0.008)
resp_education_level_high	−0.055 (0.259)
fridge	0.238 (0.243)
ever_desludged_mech	0.826*** (0.319)
ever_desludged_manual	0.406 (0.302)
desludging_manual_bl	0.378 (0.273)
ave_time_to_desl_bl_1yrless	−0.753*** (0.289)
car	−0.162 (0.284)
mechanical_max_paid_past	0.166*** (0.028)
reliable_responses	−0.623*** (0.207)
women_surveyed	−0.419* (0.215)
arranger	−0.183 (0.208)
hhd_head_respondent	−0.105 (0.240)
Observations	1,882
R^2	0.910
Adjusted R^2	0.908
F Statistic	643.475*** (df = 29; 1853)
Note:	*p<0.1; **p<0.05; ***p<0.01

Table 21: m^0 Regression

	Dependent variable:
	Y0
X0desludging_frequency_tc	−0.001 (0.001)
X0water_bill_g5k	−0.023 (0.055)
X0house_type_dum1	−0.282* (0.162)
X0house_type_dum2	−0.009 (0.120)
X0house_type_dum4	−0.012 (0.136)
X0othr_hhds_compound_tc	0.004 (0.010)
X0own_rent	0.009 (0.068)
X0trips_last_greater_than1	0.235 (0.185)
X0electricity_bill	0.001 (0.003)
X0household_size_tc	−0.008 (0.010)
X0household_women	0.022 (0.019)
X0pit_distance_to_road_tc	0.001 (0.008)
X0total_rooms	−0.002 (0.007)
X0ac	0.039 (0.060)
X0mechanical_walk_away_price	−0.003 (0.009)
X0resp_age	−0.002 (0.002)
X0resp_education_level_high	0.067 (0.060)
X0fridge	0.031 (0.039)
X0ever_desludged_mech	0.193*** (0.068)
X0ever_desludged_manual	−0.136** (0.061)
X0desludging_manual_bl	−0.050 (0.067)
X0ave_time_to_desl_bl_1yrless	0.123* (0.065)
X0car	0.055 (0.041)
X0wtp_hat_s	−0.151* (0.084)
X0eps_1	0.191 (0.213)
X0eps_2	0.479 (0.366)
X0eps_3	0.575 (0.446)
Constant	1.256*** (0.340)
Observations	278
R^2	0.295
Adjusted R^2	0.219
F Statistic	3.881*** (df = 27; 250)
Note:	*p<0.1; **p<0.05; ***p<0.01

Table 22: m^1 Regression

	Dependent variable:
	Y1
X1desludging_frequency_tc	0.002 (0.001)
X1water_bill_g5k	0.009 (0.047)
X1house_type_dum1	−0.131 (0.155)
X1house_type_dum2	0.053 (0.136)
X1house_type_dum4	0.060 (0.178)
X1othr_hhds_compound_tc	−0.018 (0.011)
X1own_rent	0.046 (0.068)
X1trips_last_greater_than1	−0.236 (0.172)
X1electricity_bill	−0.003 (0.003)
X1household_size_tc	−0.002 (0.008)
X1household_women	0.003 (0.015)
X1pit_distance_to_road_tc	0.006 (0.006)
X1total_rooms	0.008 (0.006)
X1ac	0.029 (0.058)
X1mechanical_walk_away_price	−0.005 (0.007)
X1resp_age	−0.0003 (0.001)
X1resp_education_level_high	−0.011 (0.052)
X1fridge	0.018 (0.034)
X1ever_desludged_mech	0.055 (0.059)
X1ever_desludged_manual	−0.027 (0.058)
X1desludging_manual_bl	−0.359*** (0.058)
X1ave_time_to_desl_bl_1yrless	0.074 (0.055)
X1car	0.033 (0.047)
X1wtp_hat_s	0.190 (0.161)
X1price_s	−0.103 (0.071)
X1eps_1	−0.513 (0.348)
X1eps_2	−0.139 (0.359)
X1eps_3	0.084 (0.528)
Constant	0.543 (0.356)
Observations	370
R^2	0.224
Adjusted R^2	0.160
F Statistic	3.507*** (df = 28; 341)
Note:	*p<0.1; **p<0.05; ***p<0.01

Table 23: Deposit regression

	Dependent variable:
	y
constant	-0.664^{***} (0.204)
price_s	-0.605^{**} (0.290)
price_2_s	0.067^{**} (0.031)
price_3_s	-0.003^{**} (0.001)
η_*hat*_s	0.421^{***} (0.092)
Observations	1,557
Log Likelihood	$-1{,}054.233$
Akaike Inf. Crit.	2,118.466
Note:	*p<0.1; **p<0.05; ***p<0.01

Table 24: Arrangement regression

	Dependent variable:	
	y	
	(1)	(2)
constant		
desludging_frequency_tc	0.001 (0.002)	
water_bill_g5k	−0.073 (0.158)	
house_type_dum1	−0.103 (0.447)	
house_type_dum2	−0.037 (0.387)	
house_type_dum4	−0.669 (0.572)	
othr_hhds_compound_tc	0.010 (0.032)	
own_rent	−0.019 (0.200)	
trips_last_greater_than1	0.366 (0.475)	
electricity_bill	0.003 (0.007)	
household_size_tc	−0.011 (0.035)	
household_women	0.042 (0.053)	
pit_distance_to_road_tc	−0.019 (0.030)	
total_rooms	0.024 (0.019)	
ac	0.038 (0.190)	
mechanical_walk_away_price	0.017 (0.035)	
resp_age	0.005 (0.005)	
resp_education_level_high	−0.045 (0.149)	
fridge	−0.020 (0.098)	
ever_desludged_mech	0.394^{***} (0.145)	
ever_desludged_manual	0.234 (0.302)	
desludging_manual_bl	−0.135 (0.155)	
ave_time_to_desl_bl_1yrless	0.413^{***} (0.149)	
car	−0.094 (0.174)	
mechanical_max_paid_past	0.055^{*} (0.030)	
reliable_responses	−0.103 (0.149)	
women_surveyed	0.152 (0.144)	
arranger	0.232 (0.160)	
hhd_head_respondent	0.071 (0.245)	
wtp_hat_s		0.473^{***} (0.112)
p_1	0.358 (1.506)	-1.027^{**} (0.407)
p_2	−0.022 (0.112)	0.093^{**} (0.045)
p_3	−0.0001 (0.003)	-0.003^{**} (0.001)
Constant	−3.374 (2.644)	-1.582^{***} (0.265)
Observations	1,659	1,659
Log Likelihood	−451.960	−507.862
Akaike Inf. Crit.	967.919	1,025.724
Note:		*p<0.1; **p<0.05; ***p<0.01

F Figures

Figure 1: Baseline Prices of Mechanical and Manual Services

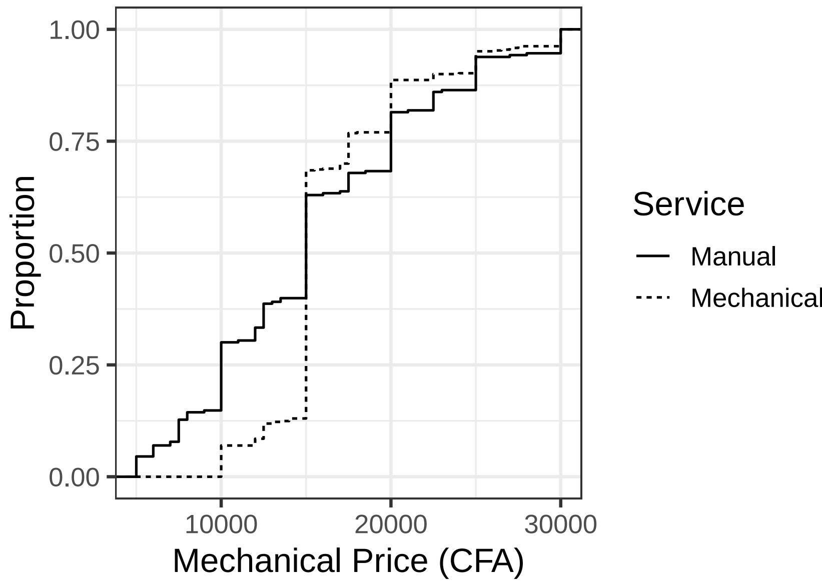

Figure 2: Taxonomy of Household Types, Theoretical

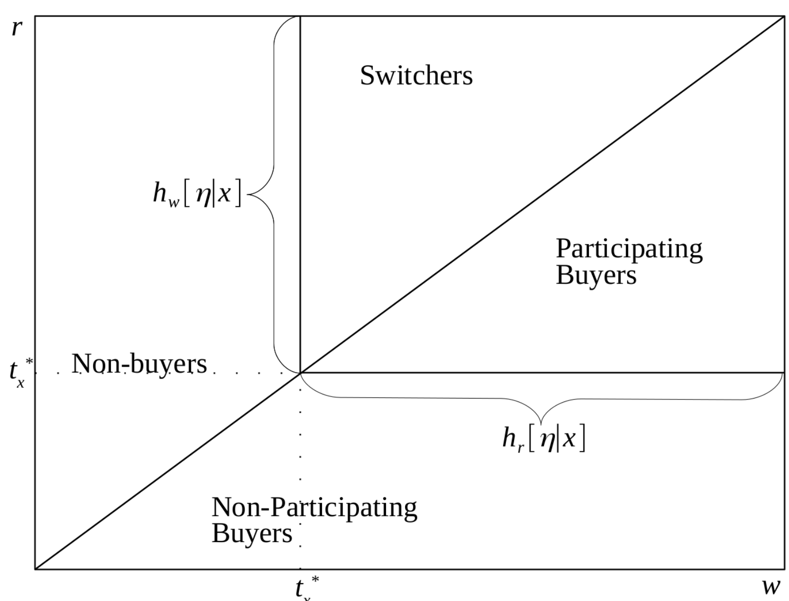

Figure 3: Taxonomy of Household Types, Empirical

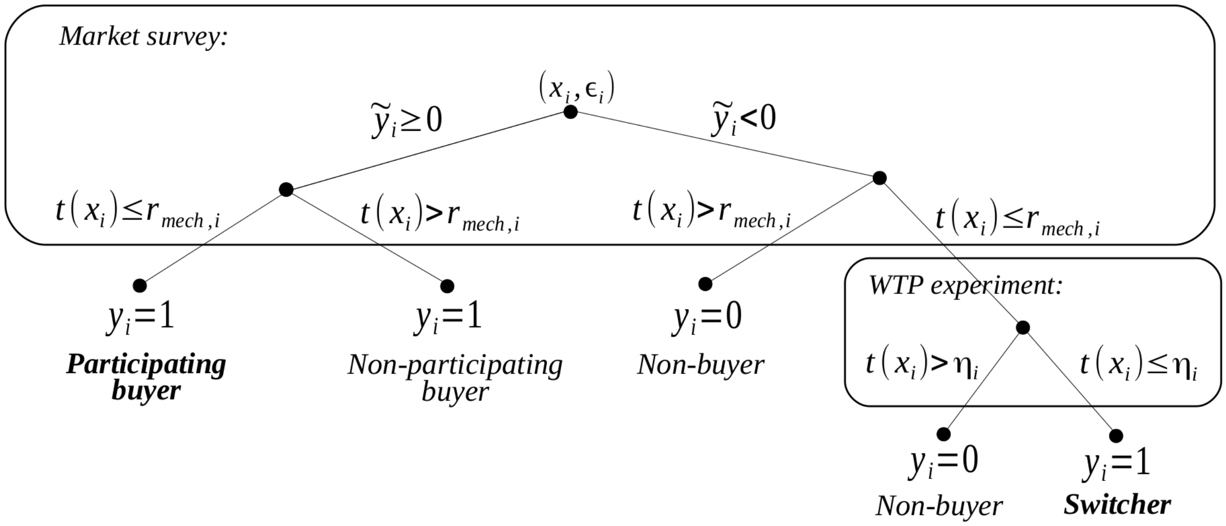

Figure 4: Offers

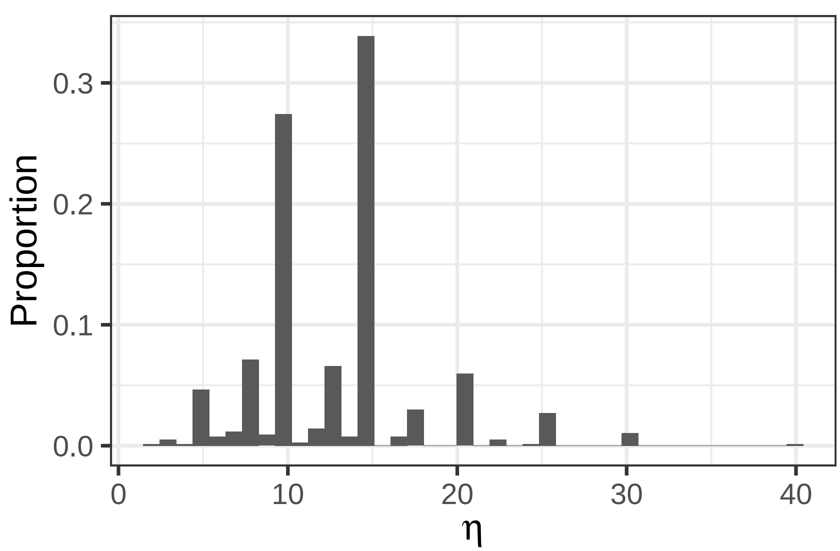

Figure 5: Demand

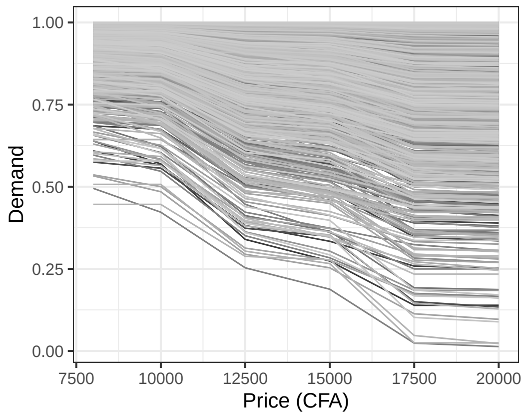

(a) Mechanical Demand by Household

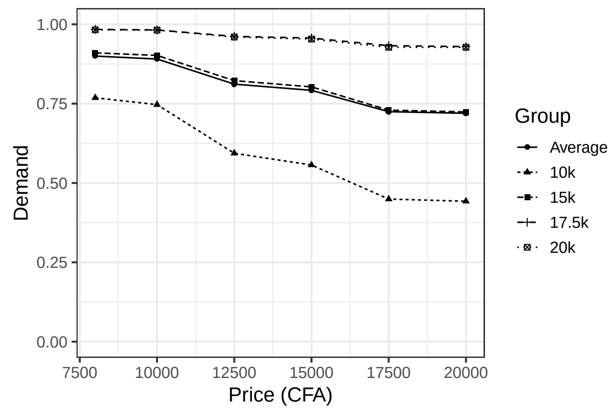

(b) Average Mechanical Demand

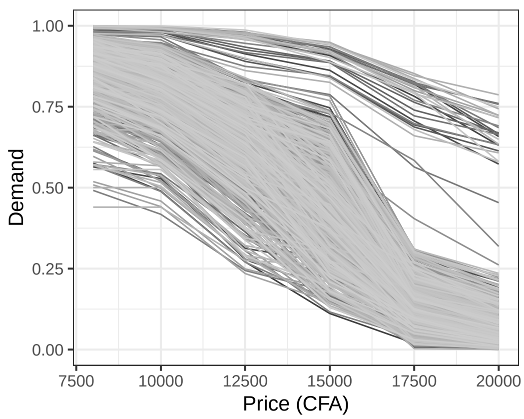

(c) Platform Demand by Household

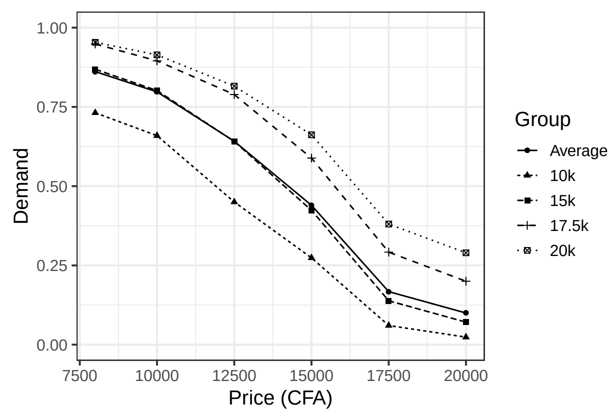

(d) Average Platform Demand

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

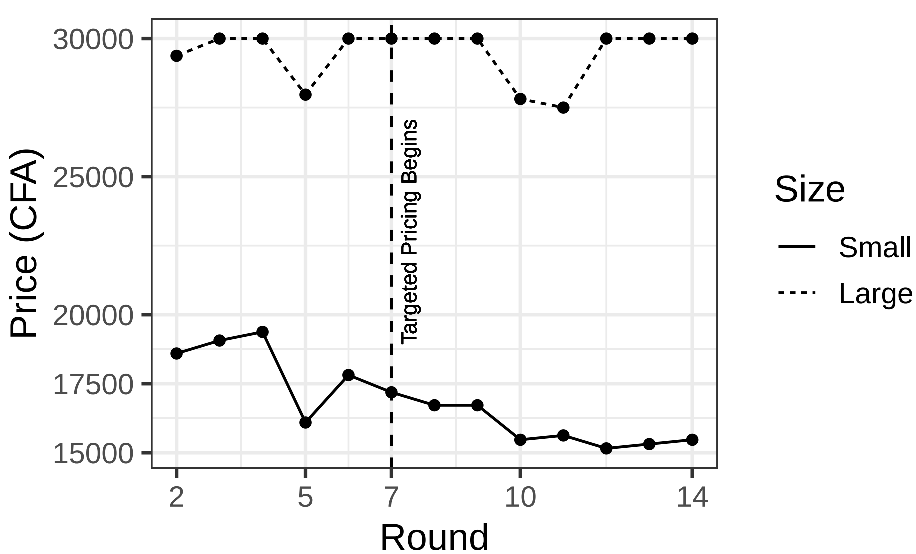

Figure 7: Pricing Rules

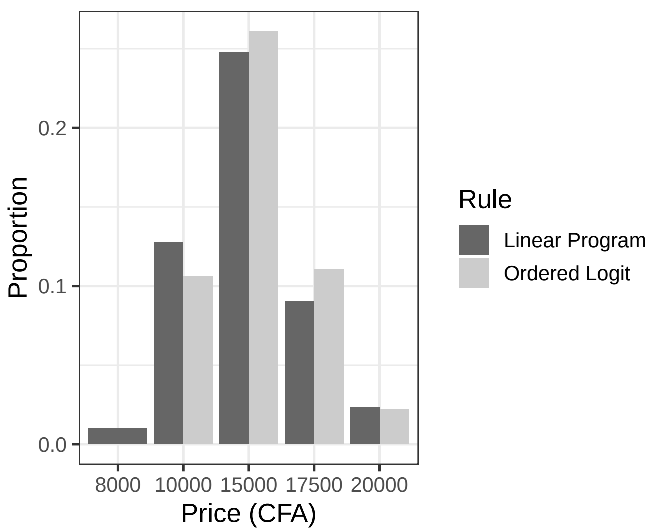

(a) Linear Programming and Ordered Logit Prices

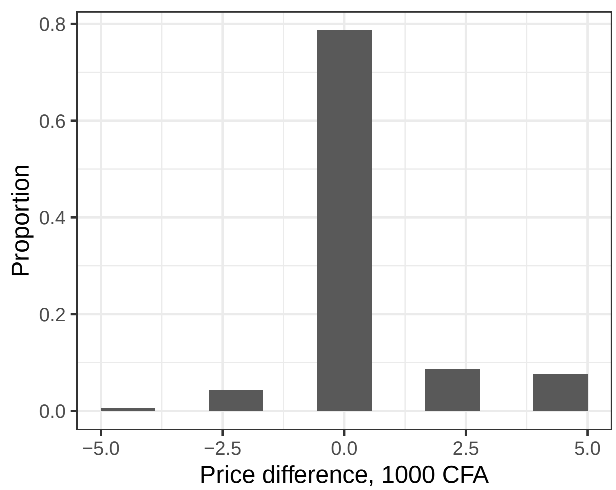

(b) Deviations of Ordered Logit from Linear Programming Price

Figure 8: Subsidization Rates