

Pricing People into the Market: Targeting through Mechanism Design

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

November 30, 2017

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 1.7 percentage points more likely to purchase a mechanical desludging, leading to an increase of market share of mechanical desludging of 5.1 percentage points. The increased probability of purchasing a mechanical desludging among those with the largest subsidies was 3 percentage points. The health impacts among the poorest were large: high subsidy households saw a decrease in the probability that one of their children had diarrhea of 7.1 percentage points.

JEL Classification Codes: L1, D4, C7, D8

Keywords: Procurement Auctions, Mechanism Design

*Johnson is assistant professor of economics at the University of Notre Dame. Lipscomb is assistant 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 Taryn Dinkelmann, Pascaline Dupas, Joe Kaboski, Isaac Mbiti, Laura Schechter, and conference participants at the Gates Foundation Workshop, Midwest International Development Conference, NBER Dev Summer meetings, University of Notre Dame, 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-pay 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 5.1 percentage points among households with access to our market relative to a control group. The market share goes up by 9.1 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 1.7 percentage points more likely to purchase a mechanical desludging, and households in the low price treatment group are 3 percentage points more likely to purchase a mechanical desludging. Next, we show that these results are not driven by flaws in the targeting method by demonstrating that the households receiving higher price offers use mechanical desludging at a higher rate and are more likely to

be wealthy along a number of dimensions that were not included in our targeting method. This establishes that the targeting scheme successfully uses observable data to pick out poor households and provides evidence that these kinds of data-driven targeting strategies can be used to increase the impact of limited subsidy budgets.

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 (Cohen and Dupas, 2010; Dupas, 2014; Kremer and Miguel, 2007). 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.

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: Coady et al. (2004) targeting programs were found to 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 which the household must submit coupons or undergo an application process (Alatas et al., 2012, 2016; Dupas et al., 2016)¹; and community-based targeting in which members of the local community or local

¹See Olken (2016) for a review.

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. (2012) in particular explore different methods of selling a new farming technology, and find that community voting delivers higher utilization and diffusion rates than auction-type mechanisms based on competing monetary or non-pecuniary bids. 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-pay, 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 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 of not purchasing a mechanical desludging but do, 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 constrains the prices that can be charged to households with higher willingnesses-to-pay. 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. Consequently, Section 3 also shows that the optimal mechanism can be implemented with a simple posted price scheme, where prices vary with the households' observable characteristics. Section 4 uses these insights and data from the first stage of the intervention to solve for the optimal price schedule. Thus, it 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

²This 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.

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 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 governments 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 household head 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 5.1 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 (94.2%) 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 58.8%, market share among the poor households in the treatment group is 68.2%. This switching effect is also seen when we focus on purchases of mechanical and manual desludging at the household level: households in the treatment group were 1.7 percentage points more likely to purchase a mechanical

desludging, and the increased probability of purchasing a mechanical desludging among those with the largest subsidies was 3 percentage points. This improvement in the sanitation conditions also led to a decrease in diarrhea in children in the poorest, most subsidized group: there was a 7.1 percentage point decrease in the probability that households reported that a child had diarrhea over the past week when they were in the treatment neighborhoods.

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.

One of our most surprising results is not on the demand side, but on the supply side. Throughout the second stage, we purchased a large number of desludgings, and started with the same kind of lowest-rejected-bid (LRB) auctions used in the first stage. Clearing prices in these auctions quickly converged to the average market price of a desludging. While we adopted the LRB auction with the objective of exploiting the fact that it is a weakly dominant strategy to bid honestly, the similarity between prevailing market prices and firm bids suggested that it was unlikely that firms were playing this equilibrium strategy. We addressed this by instituting an alternative system — which we call a structured negotiation — where we start by ranking desludgers by neighbor-

hood, with those desludgers with the lowest average prices in the auctions ranked highest. The highest-ranked desludger is the standing low bidder, and his average price the standing low bid. In these negotiations, we first approach lower-ranked desludgers and ask them if they are willing to undercut the standing low bidder and both get the current job as well as become the standing low bidder. If after a limited number of attempts we cannot find an interested desludger, we then offer the job to the standing low bidder at his most recent price. This resembles an auction, since prices paid today determine who receives work tomorrow, but potentially sacrifices efficiency by soliciting fewer bids as well as capping the amount by which the price can fall during a single transaction. On average, however, prices from structured negotiations were 9.7% less than auctions, and continued to trend downward as the study period ended. This shows that while auctions are a theoretically attractive solution for platform procurement, settings like the kind considered here might benefit from other, less traditional market designs.

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 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.

³14.4% of households that manually desludged at endline had used family members to do so for free.

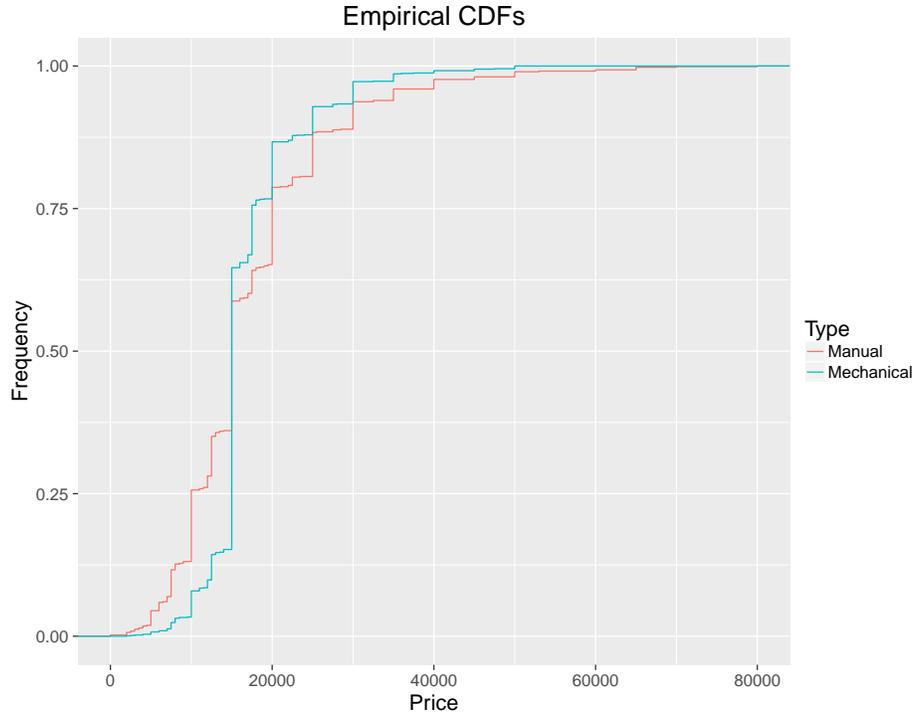


Figure 1: Cumulative Distributions of manual and mechanical prices

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 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, the “targeting problem” in development⁴ refers to using publicly available information about households to determine which are selected as the beneficiaries for a social program. Households, however, hold important private information about their willingness-to-pay or participate, and eliciting this information can help program or market designers to better distribute resources to foster more privately and socially beneficial outcomes. Households, however, will only reveal their private information if the program incentivizes them to be honest, and they can withdraw from the program at any time if they prefer the prevailing market. This section formalizes the targeting problem using the mechanism design framework, in which a platform competes against a prevailing market to maximize take-up of a socially beneficial health product, subject to a budget constraint that its losses not exceed a given subsidy level, incentive compatibility, and individual rationality. 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 studying which arrangements are optimal within some restricted class of designs. Thus, the contribution of this section is to provide a formalization of the targeting problem as a mechanism design problem, and derive an explicit solution.

Two features of the environment differentiate our problem from most applications in the mechanism design literature. First, our goal is not to target those households with the highest willingnesses-to-pay, as in a standard auction (Myerson, 1981) or non-linear pricing problem (Mussa and Rosen, 1978), but those households who are most likely to switch from manual to mechanical. Most applications of mechanism design theory focus on maximizing efficiency or total revenue, which are closely related because the agents with the highest willingness-to-pay are those from whom the seller can extract the highest price. In contrast, we wish to trade with precisely the people who are on the margin: we seek the households who, if the price were only a bit lower or the odds of success in finding a mechanical desludger a bit higher, they would be willing to switch from manual to mechanical. Second, the presence of the prevailing market gives households

⁴Alatas et al. (2016) define targeting as the method by which beneficiaries are selected for social programs.

a non-trivial outside option, and if they reject the platform’s offer, it could be because they cannot afford a desludging at that price or because they can get a more attractive deal from the prevailing market. Most classical mechanism design problems set the outside option to a constant that does not vary with the household’s willingness-to-pay, and analysis of general versions of such problems are typically quite complicated.

The crux of the design problem is that if the platform promises low prices, households who can already afford the service will pose as poor households in order to receive the service at a low price without actually increasing take-up. While this problem is obvious for the case of a simple program like a flat subsidy, it persists (and might even be worse) if standard tools like an auction are used: relatively rich households who would have purchased the good can outbid poorer households on subsidized products, leading to a zero net change in utilization. To address this challenge, we conceptualize a household’s information as being composed of a privately known willingness-to-pay, and publicly observable characteristics which the platform or existing markets can take into account when deciding how to price goods or services. Given the observable information, the platform can then make better decisions about on which households to spend subsidy dollars and from which to try to raise revenue.

This approach allows us to solve for optimal mechanisms, but also understand the value of private information in this market. In our case, the platform acts as a “profit-minded social planner”, who maximizes a weighted sum of total consumption of mechanical services and profits. We express the profits in terms of the virtual valuations of the households, which reflect the net benefit of selling to that household as well as the cost of providing incentives for it to report its privately held information honestly. Our analysis clarifies that in the targeting problem, the most attractive households are those with the highest willingness to pay and those who are just excluded from the search market, while the least attractive households are those with low willingnesses to pay and do purchase in the search market. These types neither contribute much to revenue nor substantially increase the overall consumption of mechanical desludging services. The platform would prefer to exclude them, but if it makes a more attractive offer to lower types, these households must receive terms that are at least as generous.

This analysis leads to two main results. First, the optimal mechanism is always implementable by making take-it-or-leave-it offers to the households conditional on observables. This is the only incentive compatible way to elicit their information, since higher types will always be tempted to “impersonate” lower types in order to get a subsidized deal, conditional on the same observables. Eliciting this information is, however, valuable: the platform can adjust prices to subsidize observable types of households that are most likely to switch while charging higher prices to those who are most likely to get the service anyway, but might agree and relax our budget constraint. Our second main result is to show that households who are more likely to have a high willingness-to-pay given their observables should be charged higher prices in any optimal mechanism. This provides guidance on how to use information to design the optimal pricing rules in Section 4.2.

3.1 Model, Mechanisms, and Incentive Compatibility

Each household has a privately known willingness-to-pay w , and publicly known observables x . 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, like water or electricity bill expenditure. The observable type x takes values in a finite⁵ set X , where the proportion of households of type $x \in X$ is given by $\mu(x)$, with $\sum_{x \in X} \mu(x) = 1$. The willingness-to-pay of a household with observables x has cumulative distribution function $F[w|x]$, which is differentiable with probability density function $f[w|x]$, satisfying $f[w|x] > 0$ for all $w \in [\underline{w}, \bar{w}]$. 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.

In the prevailing market, the household can incur a cost t_x^0 that yields a probability $p_x^0 \in [0, 1)$ of getting a mechanical desludging for each $x \in X$. The cost t_x^0 includes foregone wages or the value of leisure time used to search, as well as the expected transfer to a desludger providing the household a mechanical desludging. In the absence of the platform, a household searches for a desludging in the prevailing market only if its willingness-to-pay is sufficiently high, so that

⁵Appendix 7 extends this to a countable measure space, (X, \mathcal{F}, μ) to allow for continuous covariates.

$p_x^0 w - t_x^0 \geq 0$. Define

$$w_x^0 = \frac{t_x^0}{p_x^0},$$

which is the marginal type that is indifferent between searching for a mechanical desludging and not. Thus, the indirect utility of a household with observables x and willingness-to-pay w in the prevailing market is given by

$$U^0(w, x) = \mathbb{I}\{w \geq w_x^0\} w p_x^0 - t_x^0 = \mathbb{I}\{w \geq w_x^0\} (w - w_x^0) p_x^0. \quad (1)$$

The platform is a competing market, layered on top of the prevailing market. Since our sample includes a small number of households relative to the market overall, it does not create general equilibrium effects that affect the expected probability of trade or payment in the prevailing market. Project subjects, however, can opt out of our market and search instead, providing them with a non-trivial outside option. We now seek to answer the question, “among all possible methods of arranging trade, which increases the take-up of mechanical desludging the most, subject to a limit on the level of subsidies available to facilitate trade?”

The Revelation Principle guarantees that any equilibrium of a game of incomplete information can be converted into an alternative game, 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, $\{p(w, x), t(w, x)\}_{w \in [\underline{w}, \bar{w}], x \in X}$, in which a household with observables x reports a type \hat{w} and trade occurs with probability $p(\hat{w}, x)$ at an expected price of $t(\hat{w}, x)$, where households find it in their best interests to participate and report their types honestly. Once the optimal direct mechanism is characterized, we can then determine what kinds of practical economic institutions implement the same outcome without relying on the abstract thought experiment of the direct mechanism.

Since the households can always withdraw from the platform and trade in the search market, their incentives are equivalent to maximizing the net benefit of participating on the platform, where their outside option equals zero. The household’s type report \hat{w} need not be truthful, it

selects \hat{w} strategically to maximize its direct utility function,

$$U(\hat{w}, w, x) = p(\hat{w}, x)(w - U_0(w, x)) - t(\hat{w}, x), \quad (2)$$

where \hat{w} is the report, w is the true type, and x are the observables. To insure that agents do not find it in their best interests to deviate from honesty nor withdraw from the platform entirely, we impose additional constraints on what kinds of direct mechanisms the platform can select. The direct mechanism is *incentive compatible* if for all w , w' , and x , $U(w, w, x) \geq U(w', w, x)$, so that honestly reporting one's type gives a weakly higher net benefit than lying, and *individually rational* if, for all w and x , $U(w, w, x) \geq 0$, so that the net benefit from participating for all types is at least as high as withdrawing.

The individual rationality and incentive compatibility constraints have a very strong implication: the households' payments and payoffs are determined entirely by the probability of trade. In particular, the envelope theorem⁶ prescribes the rate at which a household's net benefit of participating grows in its own type:

$$\frac{dU(w, w, x)}{dw} = p(w, x)(1 - \mathbb{I}\{w \geq w_x^0\}p_x^0). \quad (3)$$

To see the intuition for this, assume for the moment that $p(w, x)$ is differentiable. Then the total derivative of $U(w, w, x)$ with respect to w is

$$U_1(w, w, x) + U_2(w, w, x) = \left\{ \frac{\partial p(w, x)}{\partial \hat{w}}(w - U^0(w, x)) - \frac{\partial t(w, x)}{\partial \hat{w}} \right\} + p(w, x)(1 - \mathbb{I}\{w \geq w_x^0\}p_x^0).$$

The first effect, $U_1(w, w, x)$ is the change in net payoffs as a consequence of varying the report, and corresponds to the household's first order necessary condition⁷ in its report, \hat{w} . Incentive

⁶Milgrom and Segal (2002), Corollary 2

⁷If $p(\hat{w}, x)$ is differentiable, then the necessary condition for a report to be optimal for the household is that

$$\frac{\partial p(\hat{w}, x)}{\partial \hat{w}}(w - U^0(w, x)) - \frac{\partial t(\hat{w}, x)}{\partial \hat{w}} \leq 0$$

with equality at any interior report that is optimal. Imposing incentive compatibility implies that $\hat{w} = w$ is optimal.

compatibility implies that honest reporting is optimal, so the term in braces must be zero. Since the increase in the report is publicly observable, the increase in surplus attributable to the change in the report can be captured by the platform by adjusting the transfer accordingly. The second effect, $U_2(w, w, x)$, is the direct effect on surplus of an increase in the household's type, and accrues to the household. The envelope condition (3) thus prescribes the rate at which the household's net payoff changes with the allocation, $p(w, x)$. With the rate of change of $U(w, w, x)$ in w determined, we can integrate (3) with respect to w to get

$$U(w, w, x) = U(w_x^*, w_x^*, x) + \int_{w_x^*}^w p(z, x)(1 - \mathbb{I}\{z \geq w_x^0\}p_x^0)dz.$$

Since the marginal type w_x^* is indifferent between participating and not, the net benefit to this type is zero. This gives us an expression for the payoff of the (w, x) type in any incentive compatible mechanism as a function of the allocation alone:

$$U(w, w, x) = \int_{w_x^*}^w p(z, x)(1 - \mathbb{I}\{z \geq w_x^0\}p_x^0)dz. \quad (4)$$

Combining (4) with (2) then yields a formula for the expected transfer to the platform in terms of the allocation $p(w, x)$ alone:

$$\begin{aligned} t(w, x) &= p(w, x)(w - U^0(w, x)) - U(w, w, x) \\ &= p(w, x)(w - U^0(w, x)) - \int_{w_x^*}^w p(z, x)(1 - \mathbb{I}\{z \geq w_x^0\}p_x^0)dz \end{aligned} \quad (5)$$

Taking the expectation of this formula with respect to w conditional on x and simplifying through an integration by parts yields the next result:

Proposition 1 *Let c_x be the average cost of serving a household with observable type x . In any incentive compatible mechanism, expected platform profits equal*

$$\sum_{x \in X} \mu(x) \int_{\underline{w}}^{\bar{w}} p(w, x) \{\psi(w, x) - c_x\} f[w|x]dw,$$

where

$$\psi(w, x) = w - U^0(w, x) - \frac{1 - F[w|x]}{f[w|x]}(1 - \mathbb{I}\{w \geq w_x^0\}p_x^0).$$

This result associates each type (w, x) with an index of profitability, known in the mechanism design literature as the *virtual valuation*,

$$\psi(w, x) = \underbrace{w - U^0(w, x)}_{\text{Net change in welfare}} - \underbrace{\frac{1 - F[w|x]}{f[w|x]}(1 - \mathbb{I}\{w \geq w_x^0\}p_x^0)}_{\text{Informational rent}}.$$

This function answers the question⁸, “if you sold a mechanical desludging to a household with type (w, x) , what is the marginal revenue that you expect to earn from it?” The virtual valuation in this setting is composed of two parts. First, the net change in welfare represents the benefits that the platform creates for the household by ensuring that it receives a mechanical desludging for sure instead of its expected payoff in the prevailing search market. Since higher willingness-to-pay types can always report lower types and get at least as attractive an outcome, however, higher types must be incentivized to report honestly. Second, the informational rent captures the cost of incentive provision. If the platform reduces the lowest type that trades by a differential amount, it results in $f[w|x]dw$ more trades, but reduces the revenue the platform can extract from higher types by $-(1 - F[w|x])$, so that the marginal cost of including w is $-(1 - F[w|x])/f[w|x]$. The

⁸This footnote provides a clearer connection between the non-linear pricing problem and mechanism design; see Bulow and Klemperer (1996) for further details.

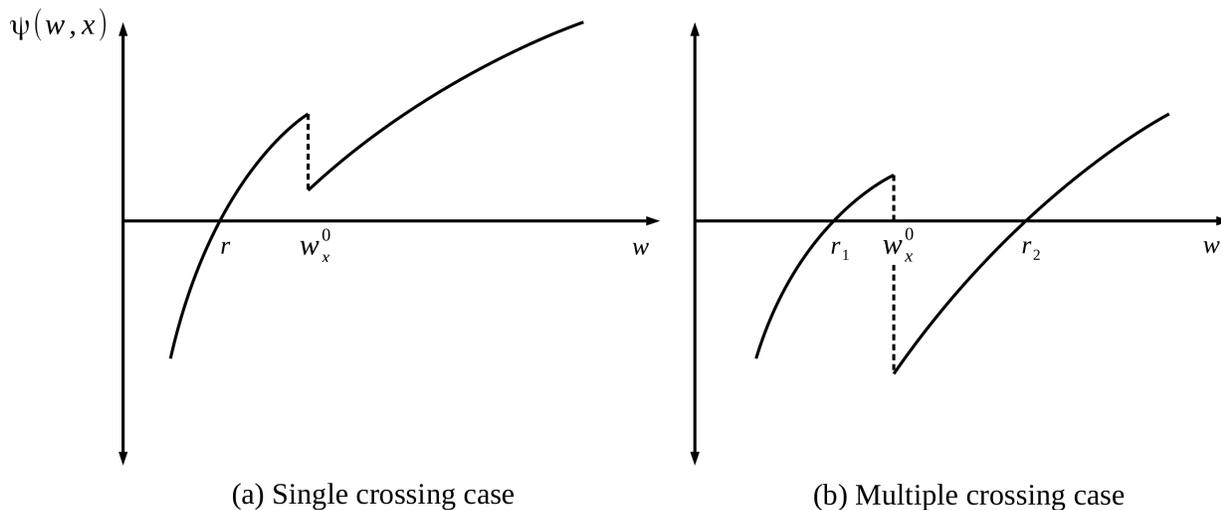
Consider a monopolist who faces demand curve $D(t) = 1 - F(t)$ and solves $\max_t (1 - F(t))t$, where consumers have an outside option of zero. His first-order necessary condition satisfies $1 - F(t) - f(t)t = 0$. If $1 - F(t)$ is log-concave, then $(1 - F(t))/f(t)$ is a decreasing function and his problem is globally concave, so any solution to the FONC is a global maximizer. Therefore, the optimal price satisfies $t^* = (1 - F(t^*))/f(t^*)$.

Using the mechanism design formalism to analyze the same problem, the consumer’s direct utility function is $U(w, w) = p(w)w - \int_{w^*}^w p(z)dz - U(w^*)$, where $U(w^*)$ is the payoff of the worst-off type who trades with the monopolist, and can be set equal to the outside option, 0; otherwise, the monopolist could raise the payment made by w^* until they were indifferent between trading and not. Then the expected transfer is

$$\mathbb{E}[t(w)] = \int_{\underline{w}}^{\bar{w}} \left(p(w)w - \int_{w^*}^w p(z)dz \right) f(w)dw = \int_{\underline{w}}^{\bar{w}} p(w) \left(\underbrace{w - \frac{1 - F(w)}{f(w)}}_{\psi(w)} \right) f(w)dw.$$

The solution is to set $p(w) = 1$ for all w such that $\psi(w) \geq 0$. The worst-off type that trades with the monopolist is then selected by solving $\max_{w^*} \int_{w^*}^{\bar{w}} \psi(w)f(w)dw$, which has first-order necessary condition $-\psi(w^*)f(w^*) = 0$. Note that $\psi(w^*) = 0$ yields the same outcome as the first-order necessary condition $(1 - F(t^*)) - f(t^*)t^* = 0$. Thus, $\psi(w)$ captures the marginal revenue that accrues to the platform from adjusting the cutoff downwards.

Figure 2: Virtual Valuation



term $1 - \mathbb{I}\{w \geq w_x^0\}p_x^0$ then discounts the informational rent by the fact that the household can potentially purchase a mechanical desludging in the outside market. In other words, the household only receives an informational rent net of their outside option. The virtual valuation $\psi(w, x)$ thus captures the motivations of the platform to raise funds in order to relax its budget constraints.

To further understand the platform's motivations for raising revenue, Figure 2 plots the possible virtual valuations $\psi(w, x)$ for each unobservable type w . The presence of the outside option⁹, $U^0(w, x)$, means that the household's bargaining power with respect to the platform depends on whether they would purchase a desludging in the search market or not. At the indifferent type, w_x^0 , there is a discontinuous drop in the expected marginal revenue that would accrue to the platform from serving that type. This occurs because types less than w_x^0 do not have the bargaining power afforded by the prevailing search market, but types above w_x^0 can leverage their outside option to extract a better deal from the platform.

In the single crossing case, (a), once the virtual valuation becomes positive, all higher types $[r, \bar{w}]$ represent a profit to the platform, and serving them raises revenue. To the extent that the platform subsidizes trade, it is for types less than r , and the marginal benefit of serving them

⁹In most applications in the literature, the outside option $U^0(w, x)$ is normalized to zero. For the general analysis with type-dependent outside options, see Giovanni Rodriguez-Clare, Julien

is continuous in their type. In the multiple crossing case, (b), however, there is a set of types in $[w_x^0, r_2)$ which the platform would prefer to exclude from participating in the platform. The intuition is that these types would have purchased a desludging on their own, but benefit from low prices that are intended for households who would otherwise not have purchased, the set $[r_1, w_x^0)$. Thus, these types contribute neither to higher revenues nor a substantially higher probability of take-up of mechanical desludgings. The platform would rather offer a low price to the set $[r_1, w_x^0)$ and leave the set $[w_x^0, r_2)$ to purchase in the search market. But because their willingnesses-to-pay are private, the set $[w_x^0, r_2)$ must be served if $[r_1, w_x^0)$ is, and on terms that are at least as favorable. The platform then faces a dilemma: include the entire set $[r_1, r_2)$, or exclude them all¹⁰. This precisely expresses the problem of targeting, and clarifies that the set of types whose participation is undesirable to the platform is not the highest willingness-to-pay households, but those that constrain the prices that can be charged to high willingness-to-pay households without purchasing a mechanical desludging with a substantially higher probability.

3.2 Optimal Targeting

As discussed in the introduction, there are significant negative externalities from the collection and disposal of human fecal sludge, especially for young children for whom exposure to human waste can lead to stunting and other developmental disadvantages. Consequently, we model the platform's objective as maximizing the utilization¹¹ of improved sanitation services. In addition to the individual rationality and incentive compatibility constraints, the platform must also ensure that its total profits plus subsidies, S , are not negative, or

$$\sum_{x \in X} \mu(x) \int_{w \in [\underline{w}, \bar{w}]} (t(w, x) - c_x) f[w|x] dw + S \geq 0, \quad (6)$$

¹⁰Analytically, this involves ironing the non-monotonicity in the virtual valuation (Fudenberg and Tirole 1991, Noldeke and Samuleson 2011), and we show that in this context the optimal probability of trade in the pooled interval is 0 or 1 so that deterministic mechanisms are still optimal; see Appendix 7.

¹¹We focus on the population proportion opting for mechanical desludging rather than social welfare directly since it is difficult to gather credible data on the size of externalities. During a number of pilots, participants expressed that they found it unnatural to pay additional amounts to ensure that a neighbor received the service. Finding accurate and robust ways of measuring these externalities remains an open and important problem.

where c_x is the average cost faced by the platform in procuring a desludging for a household with observables x . Call (6) the expected budget balance constraint.

This leads to the *quantity-maximization problem*: pick the expected payments $t(w, x)$ and probabilities of trade $p(w, x)$ on the platform to solve

$$\max_{\{p, t\}} \sum_{x \in X} \mu(x) \int_{w \in [\underline{w}, \bar{w}]} \{p(w, x) + (1 - p(w, x))p_x^0\} f[w|x] dw$$

subject to individual rationality, incentive compatibility, and expected budget balance constraints.

Exploiting the virtual surplus characterization of platform revenue, we can use the Karush-Kuhn-Tucker theorem to reduce the quantity maximization problem to maximization of the Lagrangian,

$$\begin{aligned} \mathcal{L}(p, \lambda) = \sum_{x \in X} \mu(x) \left\{ \int_{\underline{w}}^{\bar{w}} (p(w, x)1 + (1 - p(w, x))\mathbb{I}\{w \geq w_x^0\}p_x^0) f[w|x] dw \right. \\ \left. + \lambda \int_{\underline{w}}^{\bar{w}} p(w, x) (\psi(w, x) - c_x) f[w|x] dw \right\} + \lambda S, \quad (7) \end{aligned}$$

where λ represents the multiplier on the expected budget balance constraint. The first line captures the quantity-maximization motive, while the second incorporates the platform's profit-maximization motive. If λ is small, the platform will neglect profit maximization in favor of distributing mechanical desludgings as widely as possible, while if λ is large, its budget constraint is particularly binding and it will behave more like a purely profit-maximizing platform. The combination of quantity maximization and the expected budget balance constraint thus results in a "profit-minded social planner": raising an additional dollar of revenue from some observable type x can be used to cross-subsidize consumption by some other observable type, x' , so that the profitability of types plays a key role in the design of the optimal platform. Analysis of (7) characterizes¹² optimal mechanisms in this environment:

¹²Appendix 7 provides a comprehensive analysis of the mechanism design problem, including existence of an optimal mechanism, necessary and sufficient conditions for a direct mechanism to be incentive compatible, verification of individual rationality, analysis of a "relaxed problem" that generates Figure 2, and analysis of the problem allowing for pooling at the bottom. These results are summarized in the statement of this Theorem to streamline the exposition.

Theorem 2 *Assume the standard regularity condition that $1 - F[w|x]$ is log-concave¹³ and let λ^* be the multiplier on the expected budget balance constraint at the optimum.*

i. For all x , there is a type w_x^ that satisfies*

$$1 - \mathbb{I}\{w_x^* \geq w_x^0\}p_x^0 + \lambda^*(\psi(w_x^*, x) - c_x) = 0,$$

and in the optimal mechanism, all types $w \geq w_x^$ trade on the platform and all types $w < w_x^*$ either purchase a mechanical desludging in the prevailing search market or get a manual desludging.*

ii. The optimal direct mechanism can be implemented by making take-it-or-leave-it offers conditional on each observable type $x \in X$, where the optimal price satisfies

$$t_x^* = c_x + \left(\frac{1 - F[w_x^*|x]}{f[w_x^*|x]} - \frac{1}{\lambda^*} \right) (1 - \mathbb{I}\{w_x^* \geq w_x^0\}p_x^0). \quad (8)$$

The intuition for part i is that if the platform offers an attractive offer to a type (w, x) household, any household of type (w', x) with $w' > w$ can behave as if it was a type w household and receive the same attractive offer. In a standard third degree price discrimination problem with increasing and strictly concave utility over quality or quantity, the platform could introduce distortions away from the optimal quality or quantity level to extract more rents. Here, however, the surplus is linear in the probability of trade (the “quantity” of trade) and such distortions are undesirable: it is optimal to trade with a given type (w, x) with probability one or zero. This leads to a cutoff rule, where types below some threshold w_x^* are excluded from the platform while those above are all offered the chance to buy a mechanical desludging. This can then be implemented by charging the price in (31), using part ii of the Theorem. If the platform were simply maximizing profits,

¹³If the demand curve $D(t) = 1 - F(t)$ is log-concave, then $\max_t D(t)t$ is quasi-concave, and the first-order necessary condition is sufficient to characterize a solution which satisfies $t^* = (1 - F(t^*))/f(t^*)$. For other uses in the mechanism design literature, see Mussa and Rosen (1978), Myerson (1981), or Bergstrom and Bagnoli (2005). The uniform, normal, and exponential distributions, for example, all satisfy this condition.

the optimal price (31) would be given by

$$t_x^* = c_x + \frac{1 - F[w_x^*|x]}{f[w_x^*|x]}(1 - \mathbb{I}\{w_x^* \geq w_x^0\}p_x^0).$$

To get to the optimal price, the $1/\lambda^*$ term is deducted from the household's informational rent $1 - F[w_x^*|x]/f[w_x^*|x]$, resulting in a lower price charged to observable type x . When the shadow benefit of an additional dollar is high, this deduction will be low, and the platform will behave more like a profit-maximizer. As the subsidy level S grows and more dollars are available, λ will decrease, and prices will fall. The optimal λ^* is determined at the optimum by finding the point at which the realized profits plus the fixed value of S equals zero. Thus, Theorem 2 provides guidance for designing an optimal platform, namely that posted prices that vary with household observables to balance the budget are optimal.

Our next result provides a way of predicting how optimal pricing patterns vary across different values of observables x . In practice, x will be a vector of covariates, such as water or electricity bills, the number of floors of the dwelling, the material from which the dwelling is built, the number of rooms, and so on. Consequently, we seek a way of comparing the distributions of willingness-to-pay under alternative observables x and x' . The most common stochastic ordering in economics is first-order stochastic dominance, in which x first-order stochastically dominates x' if $F[w|x] \leq F[w|x']$ for all $w \in [\underline{w}, \bar{w})$. The intuition for this is that under $F[w|x]$, lower values are less likely, implying that $F[w|x]$ must deliver higher values with higher probability. Unfortunately, first-order stochastic dominance implies very little structure that is useful for our purposes. Instead, say that x *hazard-rate dominates* x' if for all $w \in [\underline{w}, \bar{w})$,

$$\frac{f[w|x]}{1 - F[w|x]} \leq \frac{f[w|x']}{1 - F[w|x']}.$$

The intuition for this is as follows: conditional on knowing that a household's willingness-to-pay was at least as large as w , the probability that the household's value was exactly w is $f[w|x]dw/(1 - F[w|x])$. Now, if x hazard rate dominates x' , it follows that for every w , there is a lower probability

for x than x' that w is the household's true value, conditional on both having at least a willingness to pay of w . This implies that x has uniformly higher probabilities of being a high type and puts more structure on the form of the virtual surplus in x , since the inverse of the hazard rate appears explicitly. Note that hazard-rate dominance implies first-order stochastic dominance, but not necessarily the converse.

Proposition 3 *Suppose $\bar{c}_x = \bar{c}_{x'}$, $p_x^0 = \bar{p}_{x'}$, $t_x^0 = t_{x'}^0$, $\psi(w_x^0, x)$ and $\psi(w_{x'}^0, x')$ are non-negative. If x hazard rate dominates x' , then $t_x^* \geq t_{x'}^*$.*

This provides useful guidance for our empirical quantity-maximization exercise: if a type x is correlated with attributes that make it more likely to have a higher willingness-to-pay than x' , we should charge that type a higher price. Again, the intuition follows from the fact that while the platform is maximizing quantity, it is also profit-minded: by charging higher prices to households who probably have high values and will purchase in the search market anyway, the platform can relax its budget constraint in order to make more attractive price offers to relatively poorer households. In fact, this feature will be empirically verified in Section 4.2 for the rule that is deployed in the second stage of the project.

4 Experimental Design and Data

The project takes place in two stages. In the first stage, we gather market data on households' most recent transactions in the decentralized market and measure their willingness-to-pay for improved sanitation services through a demand elicitation game based on the second-price auction. These data allow us to construct a demand model for improved sanitation services that rationalizes observed selection into manual and mechanical desludging, and the pivotal price at which a given household would have switched services. This model is based on household characteristics that are either easily verified or gathered by a local municipal authority like ONEA. In the second stage, we use this demand model to determine the optimal prices that an intermediary platform would quote to households to maximize take-up of desludging services across the original sample, subject to a

budget constraint that expected profit not be less than a given subsidy level. Since the first sample is random, it is also an optimal pricing rule with respect to the population overall. We deploy this pricing rule on a new set of households, allowing us to both test whether this approach can increase take-up of mechanized services in practice and diagnose the shortcomings of the framework.

4.1 First-stage: Demand estimation

There is already an existing market for mechanized desludging services, which means that we not only have to understand household willingness to pay for mechanized desludging services, we also need to understand the trade-offs that households already face in the market. This allows us to estimate the combination of This consists of combining two sets of data: information on desludging transactions that have been made in the market, and information on the willingness to pay of consumers, many of whom have not purchased mechanical desludgings in the past. While market data can provide us with estimates of how demand responds to changes in price for households who have purchased desludgings, it can provide us with no information on the behavior of consumers who have not purchased mechanized desludgings in the past when faced with lower prices than they have seen in the market. We use an incentive compatible demand elicitation experiment to supplement the market data with information on the willingness to pay of households which have not purchased mechanized desludgings in the past.

The demand elicitation experiment and market survey took place in December 2014, with 2088 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 there are not well-defined property rights were omitted.

During the survey, we gathered information that would likely be available to a local municipal authority, like ONEA — given in Table 1 — as well as information on their most recent desludging. This information includes the mechanical price if they purchased mechanical $p_{mech,i}$; the manual price if they purchased manual, $p_{man,i}$; and whether they purchased mechanical, $y_i = 1$, or manual, $y_i = 0$. We model the determination of the manual and mechanical prices in the market and the

household's decision for a given set of characteristics x_i as a two part Tobit model (Cragg, 1971):

$$\tilde{y}_i = x_i\delta + \varepsilon_i \quad (9)$$

$$p_{mech,i} = \begin{cases} z_i\beta_{mech} + \varepsilon_{mech,i}, & \tilde{y}_i \geq 0 \\ \emptyset, & \tilde{y}_i < 0 \end{cases} \quad (10)$$

$$p_{man,i} = \begin{cases} \emptyset, & \tilde{y}_i \geq 0 \\ z_i\beta_{man} + \varepsilon_{man,i}, & \tilde{y}_i < 0 \end{cases}, \quad (11)$$

where the latent index, \tilde{y}_i , determines selection into manual or mechanical. Consequently, only the transaction price for the kind of desludging selected is observed, not the counterfactual price they would have been charged had they selected the other kind of service. The shock $\varepsilon = (\varepsilon_i, \varepsilon_{mech,i}, \varepsilon_{man,i})$ is jointly normally distributed, and we estimate $(\delta, \beta_{mech}, \beta_{man})$ by maximum likelihood. By restricting the selection of x_i to variables that would be available or reasonably easily observable to a local governmental entity, we are deliberately handicapping the model so that it can only operate on information that is observable or would be costly to manipulate. On the other hand, this means the fit will be inferior to a competing model that uses more covariates and excluded first-stage variables. However, the goal of this estimation is to provide a predictive tool for determining household demand subject to information constraints, not to provide the best possible econometric estimates of $(\delta, \beta_{mech}, \beta_{man})$.

To ensure that the model is not identified entirely off of 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 (2) and (3). Our argument that the exclusion restriction is satisfied is based on price discrimination: at the time of contracting, the desludger might observe many characteristics about the household, especially related to water consumption and sanitation, and potentially adjust the price to extract rents. These 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 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 externality damages from poor removal of sanitation.

While the demand 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 project, it is a reduced-form model and cannot be used to determine how a household would respond to a counterfactual price offer. Instead, we supplement the market data with information from a willingness-to-pay elicitation experiment based on the second-price auction.¹⁴ The rules of the game are as follows:

- 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, w_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.

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. The enumerators stated that 99.5% of households understood by the end of the exercise, though 10.5% of households required multiple explanations.

Since honest reporting is a weakly dominant strategy in the $K + 1$ -st price auction, the offer w_i provides an estimate of the minimum of the household's willingness-to-pay when it cannot afford

¹⁴The script is provided in appendix 8.

a desludging in the decentralized market, and the price the household faces in the decentralized market when it is willing to purchase a desludging at prevailing prices. 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.¹⁵ 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. Thus, our estimate w_i is a lower bound on the household’s true willingness-to-pay which we use to determine how a household would respond to a price quoted by the platform. A histogram of the offers received is given in Figure 2.

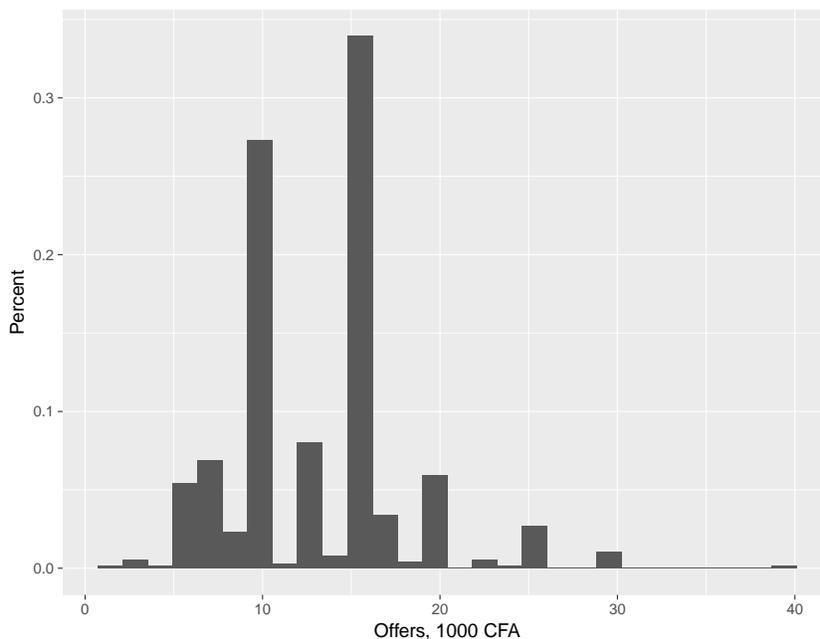


Figure 3: Histogram of offers received

For the platform’s purposes, it helps to provide a taxonomy of potential customers. When a household with characteristics x_i and the vector of shocks ε is quoted a price t_i by the platform, there are four possibilities:

- i. *Never-buyers*: The household finds neither its market price $p_{mech,i}$ nor the quoted price t_i attractive, and does not purchase a mechanical desludging.

¹⁵see, for example, Yishay et al. (2017)

- ii. *Switchers*: The household does not find its market price $p_{mech,i}$ attractive, but does purchase a mechanical desludging at the quoted price t_i .
- iii. *Participating always-buyers*: The household would purchase a mechanical desludging at its market price $p_{mech,i}$, but prefers the quoted price t_i .
- iv. *Non-participating always-buyers*: The household finds its market price $p_{mech,i}$ more attractive than the quoted price t_i , and purchases a mechanical desludging from the outside market.

Note that all four of these cases are possible for a given household with characteristics x_i , depending on the realization of the shocks $\varepsilon = (\varepsilon_i, \varepsilon_{mech,i}, \varepsilon_{man,i})$. The goal of the platform, however, is to maximize the number of switchers, and only sell to always-buyers when the sale raises money that can be used to relax its budget constraint. In particular, the WTP experiment provides an estimate of the pivotal price at which a given household would switch from manual to mechanical. This is illustrated in Figure 3.

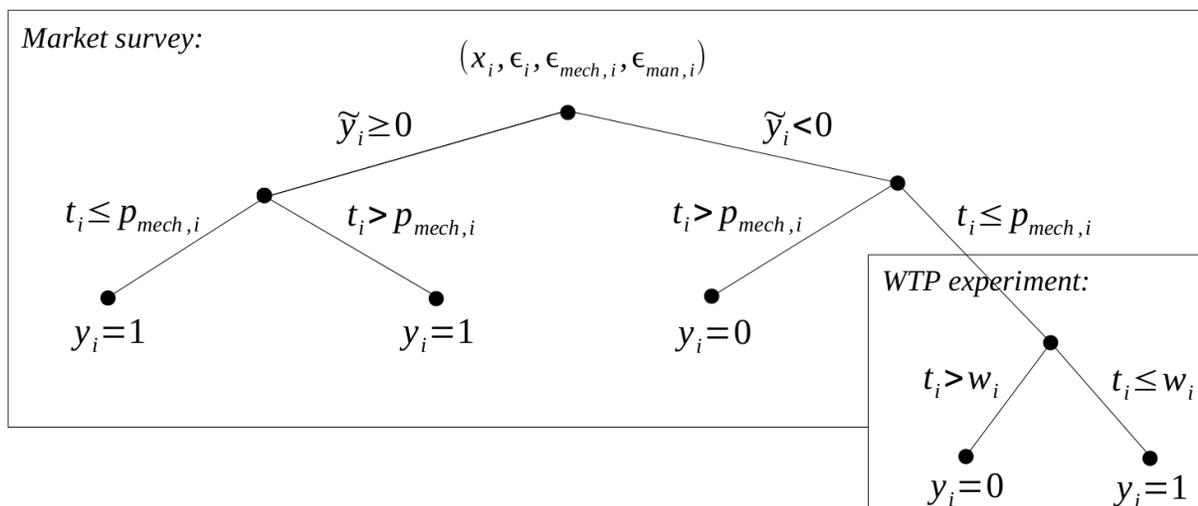


Figure 4: Demand model

The demand model thus provides predictions for a given x_i about the likelihood of selection into manual or mechanical and the likely price, and the willingness-to-pay experiment provides predictions about how households who would otherwise not purchase mechanical might switch when presented with a lower price.

Using the two sources of data together, we can derive a household-level demand curve for mechanical services, illustrated in Figure 4. The upper-left panel provides the unconditional estimate of each household's probability of accepting a given price, t_i . For prices well below the market average at 8,000 CFA, it approaches one, while for prices near 20,000 CFA, the probability tends to zero. The upper-right panel provides the predicted probability of mechanical given a price offer t_i , which will form the basis of the platform's optimization problem. The bottom two panels plot the probability of mechanical, conditional on accepting or rejecting a quoted price t_i . These reflect the platform's beliefs that a given household will purchase mechanical services, conditional on whether a given price is accepted or rejected. These household-level probabilities correspond to their demand curves, and summing yields the aggregate, market-level demand curve for improved sanitation services.

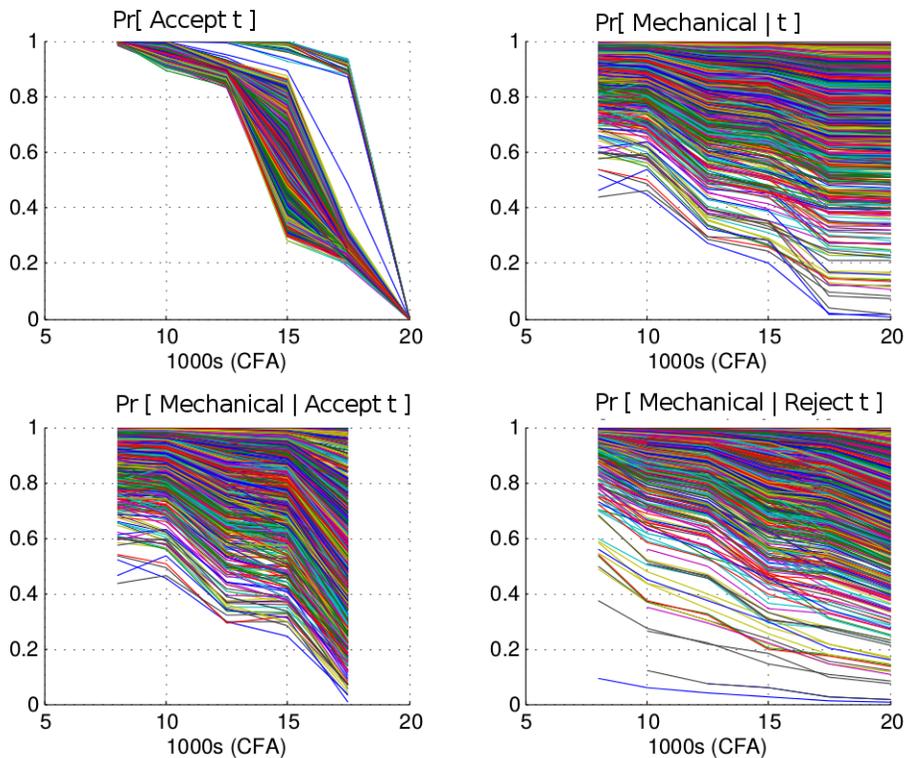


Figure 5: Estimated demands

If a low price like 8,000 CFA or 10,000 CFA is accepted, the likelihood of purchasing mechanical is relatively high (the household could still fail to call in to the center and purchase a manual instead), and these prices are predicted to be accepted with high likelihood. Given a rejection at

low prices like 8,000 CFA or 10,000 CFA, the conditional likelihood of purchase for many households goes down, while for others it increases. This reflects the fact that a rejection can occur for two reasons: the household cannot afford a mechanical desludging even at very low prices, or the household can get a much better price from the market.

The demand for mechanical desludgings can be posed more formally as¹⁶:

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

Similarly, define platform demand as

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

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.

4.2 Second-stage: Optimal Pricing

Using the demand model from stage one, we now invoke Theorem 3: the optimal mechanism involves posted prices, selected to maximize quantity subject to a budget constraint that profits (losses) not exceed a fixed subsidy level. More formally, the platform takes the sample $X = \{x_i\}_{i=1}^N$, the available subsidy budget S , and the average cost of procuring a desludging \bar{c} as given, and maximizes market demand,

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

¹⁶The indicator function $\mathbb{I}\{A(x_i)\}$ takes the value 1 when $A(x_i)$ is true, and 0 when $A(x_i)$ is false. These quantities are computed using some closed form results for the tri-variate normal, and then Monte Carlo integration.

subject to

$$S \geq \sum_{i=1}^N D_i^P(t_i, x_i)(t_i - \bar{c}) \quad (15)$$

$$t_i \in T = \{8,000, 10,000, 12,500, 15,000, 17,500, 20,000\}. \quad (16)$$

The set of prices T spans the observed transaction prices in the market data, and are the commonly used denominations for payment. The total subsidy per household was 1,750 CFA (about \$3.00), and we used an expected procurement cost¹⁷ of 17,500 CFA. Imposing the constraint (5) converts the maximization problem into a linear programming problem where each household i is assigned to a price t_i .

To appreciate how the pricing rule works, consider what the platform learns when a price is rejected, given that this price was offered to a household under the optimal pricing rule; these updated beliefs are illustrated as cumulative distribution functions in Figure 5. In the lowest price bin, the distribution associated with an acceptance first-order stochastically dominates the distribution associated with rejection: rejecting households are much more likely to get a manual desludging, and indeed more likely than the households in all the higher price bins. Thus, rejection conveys bad information about these households' likelihood of purchasing the healthy service. This pattern reverses in the 10,000 CFA price bin, where the posterior distributions are relatively close. In the 15,000 CFA bin, the rejection distribution first-order stochastically dominates the accept distribution: rejection has become a good signal as households may already have a desludger at that price. This pattern only becomes stronger in the higher price bins, until a rejection at 20,000 CFA conveys a very strong posterior that the household will purchase mechanical.

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:

$$\int_{x \in X} D(t^*(x), x) d\mu(x) \quad (17)$$

¹⁷See section 5 for more details on the supply side.

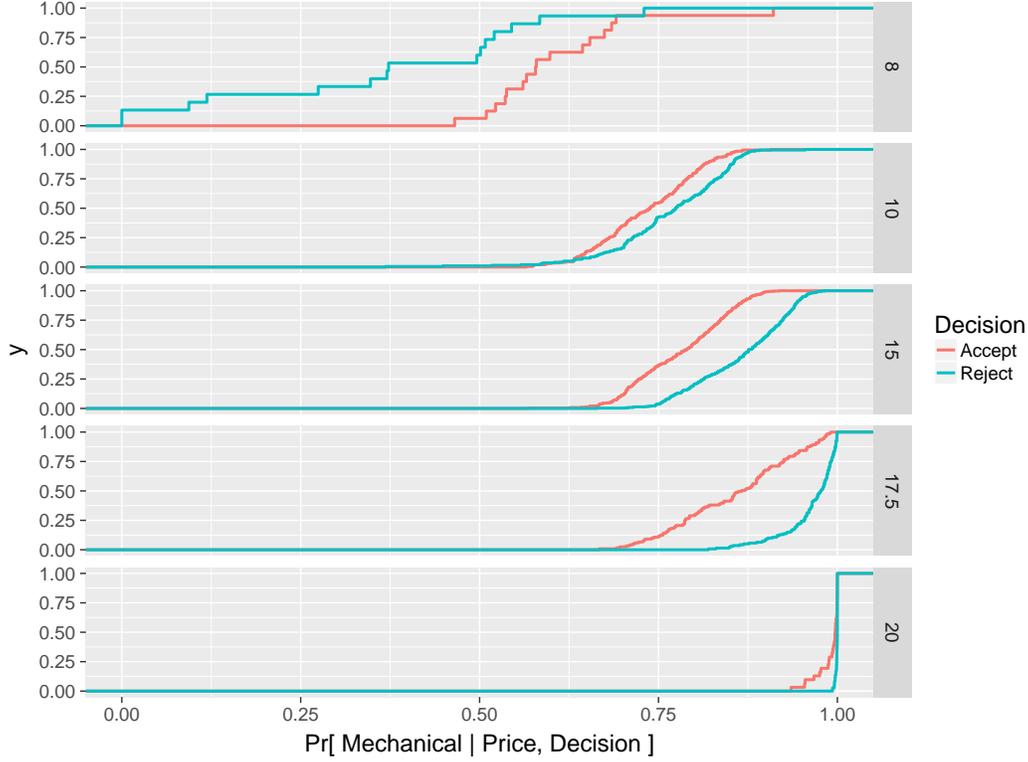


Figure 6: Updated beliefs

subject to

$$S \geq \int_{x \in X} D_i^P(t^*(x), x)(t^*(x) - \bar{c})d\mu(x) \quad (18)$$

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

While the solution to the linear program (6) — (8) is in terms of individual households, $\{(t_i, x_i)\}_{i=1}^N$, the solution to (9) — (11) is a mapping from characteristics to prices. To convert the first kind of solution into the second, we use a modification of an ordered logit model where the latent index, \tilde{t}_i , is given by

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

and assignment is then

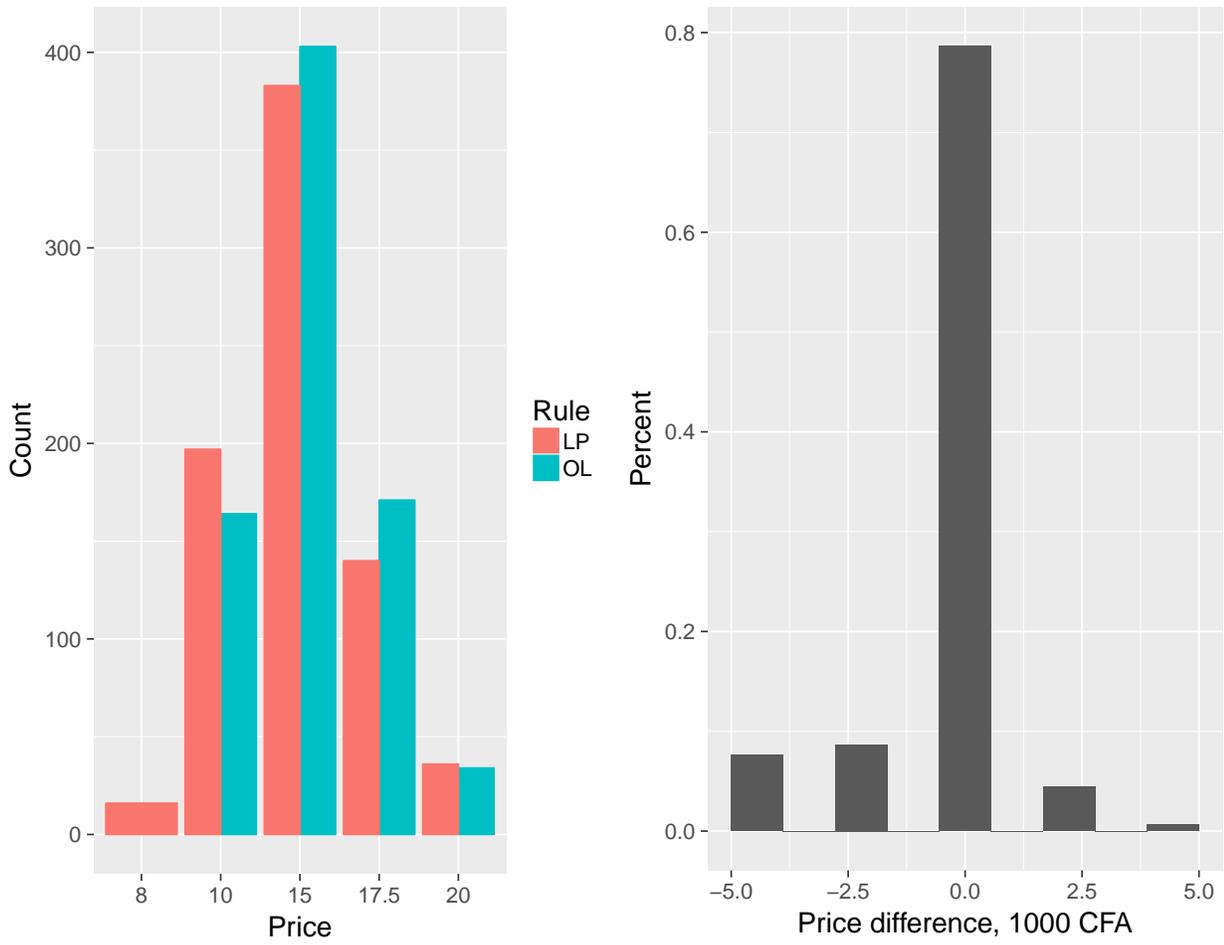
$$t^*(x_i) = \begin{cases} 10,000, & x_i \hat{\gamma} < 10,000 \\ 12,500, & 10,000 \leq x_i \hat{\gamma} < 12,500 \\ 15,000, & 12,500 \leq x_i \hat{\gamma} < 15,000 \\ 17,500, & 15,000 \leq x_i \hat{\gamma} < 17,500 \\ 20,000, & x_i \hat{\gamma} > 17,500. \end{cases} \quad (20)$$

This resembles proxy means-based testing, but is constructed not by maximizing a classification target like the fraction of households in a training data set below the poverty line that receive treatment according to the rule, but instead by approximating our optimal pricing schedule. The fit of the ordered logit rule to the linear programming solution on the original sample is illustrated in Figure 6.

The red bars represent the optimal pricing rule evaluated at the original data, X , and the blue bars represent the frequency of these price quotes in the ordered logit approximation. The optimal pricing rule assigns very few households to the 8,000 CFA bin, so we add 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 someone to switch to mechanical, and too low to relax the budget constraint. The ordered logit tends to make too many 10,000 offers and fewer 17,500 offers, but is correct approximately 79% of the time, and within 2,500 CFA of the correct bin 92% of the time. When a mis-classification does occur, the vast majority are by one price increment of 2,500. Figure 6 illustrates the linear programming and ordered logit pricing rules in the left panel, and the frequency of mis-classification in the right panel.

Our experiment of interest is then to offer the pricing rule $t = t^*(x')$ 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

Figure 7: Figure 6: Ordered logit fit



Panel A presents the count of people accepting at each price in thousands of CFA under the linear program and the ordered logit approximation. Panel B shows the price difference between the estimated price in the linear program and the estimated price in the ordered logit.

of this targeted price group is then compared to the take up of a control group in a randomized controlled trial.

4.3 Randomized Controlled Trial Data

We offered households one of three treatments in a baseline survey, run from July through September 2015: households were randomized into the targeted prices group (1,660 households in 52 neighborhoods) and a control group (1284 households in 40 neighborhoods). Each neighborhood included up to 40 households.¹⁸ Randomization was done at the neighborhood level, with stratification by number of households in the neighborhood with low walls (a proxy for low income). An initial household mapping was conducted prior to the baseline in order to select the households to be surveyed. Enumerators were told not to exceed a defined radius from the central gridpoint in order to avoid any overlap between treatment and control neighborhoods.

We show that the targeted prices treatment group and the control group are similar on a variety of household level and neighborhood-level observables in table 4. There are four variables on which the two groups are not balanced: the treatment group had a lower water bill on average, a further distance from the latrine pit to the road (which increases costs), was less likely to have needed multiple truck loads to do their last mechanical desludging, and was more likely to have more than one latrine pit.¹⁹ 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. Forty percent of households have ever used manual desludging in the past. Twenty-six percent of households state that the last desludging that they got was

¹⁸Prior to the baseline survey, lead enumerators did a mapping exercise to select households in order to make sure that households were chosen randomly for sampling in the baseline survey. Households without latrine pits and businesses which did not include a residence were omitted from the sample. Enumerators kept data on the number of households in each cluster that had low walls, which is a mark of poverty in Ouagadougou. Prior to randomization, clusters were stratified by whether they had above or below the median number of households with low walls. There was an additional treatment arm, a structured negotiations group (1,073 households in 34 neighborhoods) in which households were allowed to call in and be matched with a desludging operator at a price negotiated for them by the platform. This group was aimed at increasing the number of orders for procurement in order to observe the effects of the structured negotiations on procurement over time on the supply side.

¹⁹With 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, which would tend to bias our estimates toward 0.

manual.

All households were given a participation gift of 500 CFA at the end of the baseline survey. Treatment households were asked to use the participation gift as a deposit on their desludging if they wanted to reserve their 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 weren't accepting was because 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 5, we can see that households with the high subsidies were more likely to leave deposits, but some in the high price bins 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 endline²⁰. Table 5 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 desludging when they needed it, while the rates among the higher price groups were lower: 47% among the 15,000 CFA price group used the desludging, and 38% in the 17,500 and 47% in the 20,000 CFA price groups used the desludging.²¹ Use of the call center in the first 6 months of operation among those who needed a desludging was much higher, suggesting that more advertising was necessary among the treatment group in order to increase recall of the availability of the call center services.

While the rate of use of the call center was somewhat lower than expected, there are several

²⁰173 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

²¹Most 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 6. Some households who called the call center did not end up using it.

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 Main Results

In this section we test the effect of the targeted prices on the market share of mechanical desludging, we decompose the effect on market share into the effect on whether the household used any mechanical and any manual desludgings, and we test the effect on children's diarrhea. Market share is our primary variable of interest as the market share of mechanical desludging allows us to observe substitution from manual to mechanical. Market share is calculated as:

$$MarketShare = \frac{NumberMechanicalDesludgings}{NumberMechanical + NumberManual} \quad (21)$$

To the extent that the market share of mechanical desludgings increases, households are substituting from manual desludgings toward mechanical desludgings.²²

We then observe the direct effect of the treatment on households' purchases of mechanical and

²²Market 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).

manual desludgings. Because we subsidized only one mechanical desludging per household, we measure the impact of the program on whether the household purchased at least one mechanical desludging service during the project period. The improvement in community sanitation would come from households ending their use of manual desludging, so we also test the effect on whether households purchase any manual desludgings during the project period. Finally, we estimate the effect of the program on children’s diarrhea which has important welfare consequences.

Households targeted with high prices are overwhelmingly *always – takers*: the market share at baseline for the highest price group is 94.2% compared with 45.5% for the lowest price group. This limits the potential treatment effect in the highest price bins, and motivates our focus on the lower price bins. We expect to see the largest effects on the heavily subsidized group which had 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. Similarly, diarrhea rates in the lowest price group were higher than in the other price groups. We expect the pooled effects to be relatively small since only 27% of the sample were given the lowest subsidized prices.

5.1 Overall Impact of the Call Center

Households needing a desludging can either purchase manual or mechanical services: our key dependent variable of interest is the market share of mechanical desludging in a neighborhood since this allows us to observe the substitution between the two products. We use the following empirical specification:

$$MarketShareMechanical_i = \alpha + \beta TargetedPricesTreatment_i + \gamma' X_i + \epsilon_i \quad (22)$$

We run this regression at the neighborhood cluster level, including the 92 treatment and control clusters. X_i includes neighborhood means for each of the 7 variables on which the sample was not well balanced at the cluster level (water bill more than 5,000 CFA, electricity bill, latrine pit distance to road, two tanks used last desludging, compound has one pit only, and needed extension

hose last time) and whether the neighborhood had an above median number of low walls during the mapping phase (the variable on which the clusters were stratified).

Estimates are shown in table 7. The price targeting treatment generates an increase in the neighborhood market share of the improved desludging service of 5.1% (significant at the 10% level). This is a 7.7% effect at the mean mechanical desludging market share of 71%.

5.2 Market Share Effects by Price Group

One key implication of the pricing model is that because take-up among the wealthiest households (*always – buyers*) is already high, and some households (the *never – buyers*) will not take up mechanical desludging even at 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 94% in the baseline, the market share of mechanical desludging among households receiving the most subsidized prices was 45.5% and the market share among households receiving the second most subsidized prices was 77%.

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:

$$\begin{aligned}
MktShareMech_PriceGrp_{ki} &= \sum_{k=1}^4 \alpha_k PriceGroup_{ki} \\
&+ \sum_{k=1}^4 \beta_k TargetedPricesTreatment_{ki} * PriceGroup_{ki} + \gamma' X_{ki} + \epsilon_{ki} \quad (23)
\end{aligned}$$

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 ($\gamma' X_{ki}$).

Results are shown in table 7. 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 9.1 percentage points (significant at the 10% level). The market share for the 15,000 and 17,000 groups increase, but not statistically significantly so, and the market share for the 20,000 group does not significantly change (although the point estimate is negative). This differentially large effect among the 10,000 CFA price group was expected as they were receiving the subsidies and included more potential *switchers* (with a baseline market share of 45.5). The lowest price bin in the control group also had a substantial increase in the market share of desludging (to 58%), so this 9.1 percentage point effect is actually an increase to a market share of 68.2% in the low price bin in the treatment group.

5.3 The Household Level Effect on Manual and Mechanical Desludging

The results on market share shown above demonstrate that across the neighborhoods, households substitute from manual to mechanical desludging when they have access to the targeted price treatment. However, we may 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 “Any Mechanical” 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” which measures whether the household purchased at least one manual desludging between the baseline and the endline. It is important to control for the number of desludgings that households had to purchase over the time period, as there is a large variance in number of desludgings purchased, but this comes through factors external to the household such as weather effects, flooding, and the height of the water table. These shocks are unequally distributed across neighborhoods as they are correlated in space, and therefore could bias the coefficient of interest in a limited-size sample. We show in appendix C that it is unlikely that treatment households are strategically delaying desludgings, and that results are quite similar (larger in magnitude, but somewhat less precise) if we control for baseline variables which predict the number of desludgings that a household is likely to need over the treatment period.

We estimate these regressions at the household level, using the regression equation:

$$\begin{aligned}
 \text{AnyMechanical}_h = & \beta_0 + \beta_1 \text{TargetedPriceTreatment}_h + \gamma X_h \\
 & + \phi \sum_{x=0}^5 \mathbb{I}\{N\text{desludgings} = x\}_h + \epsilon_h \quad (24)
 \end{aligned}$$

Where h indexes the households. X_h is a vector of controls for household factors that were unbalanced at the baseline: distance to the latrine pit from the road, two trips needed in last desludging, and water bill greater than 5,000 CFA. We also control for the stratification variable, an indicator for more than half of the households in the neighborhood have low walls. Finally, we follow the advice of McKenzie (2012) for ANCOVA estimation in order to improve the efficiency of the estimates, and include controls for whether the household last used mechanical desludging and whether the household has ever desludged. Standard errors are clustered at the neighborhood level.

The results are shown in table 8. The overall probability that a household purchases a me-

chanical desludging increases by 1.7 percentage points if they are in the targeted prices treatment group. Similarly, the probability that the household purchases a manual desludging decreases by 1.1 percentage points—this is not statistically significant, but if the control for the interaction between the treatment group and the number of desludgings is included then the effect on manual desludging is significant at the 10% level.

We also estimate the effect separately by price group, using the equation:

$$\begin{aligned} AnyMechanical_{kh} = & \sum_{k=1}^4 \alpha_k PriceGroup_{kh} + \sum_{k=1}^4 \beta_k TargetedPricesTreatment_{kh} * PriceGroup_{kh} \\ & + \gamma' X_{kh} + \phi \sum_{x=0}^5 \mathbb{I}\{Ndesludgings = x\}_h + \epsilon_{kh} \quad (25) \end{aligned}$$

This specification allows us to observe the effect of treatment across the price groups. h indexes the household, k indexes the price group to which the household belongs. For the control group, we estimate which price group the household would have been assigned to had they been allocated the treatment, assigning the households to $PriceGroup_{kh}$ according to the observables used in the treatment. Controls for variables unbalanced at baseline, the stratification variable, and dummies for the number of desludgings the household purchased over the period are included. Standard errors are clustered at the neighborhood level. As shown in the market share regressions, the effect on desludging choice is coming primarily from the low price group: Households in the low price treatment group are 3 percentage points more likely to purchase a mechanical desludging during the period than similar households in the control group.

We may be concerned that households strategically purchase desludgings when in the treatment group in order to take advantage of the discount before it ends, or strategically hold off purchasing desludgings if they believe that the discount may become larger in the future. Survey evidence suggests that it is unlikely that households strategically wait longer than a month to get their 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.

In order to directly test whether the treatment pushed some households into purchasing or

not purchasing desludgings, potentially biasing the effects on purchases of mechanical or manual desludgings, we control directly for the interaction between the treatment and the total number of desludgings purchased by a household during the treatment period. Results are shown in table 8 columns (3), (4), (7) and (8). The estimated coefficient on the interaction between the total number of desludgings and the treatment group is very close to 0 and not statistically significant. The estimated effect on whether the household purchased any mechanical desludgings is slightly smaller when controlling for the interaction effect, but not statistically significantly so, and the effect remains significant at the 10% level. The estimated effect on manual desludging has a slightly larger magnitude and becomes statistically significant when controlling for the interaction between treatment and total number of desludgings purchased.

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 using the following specification to estimate the pooled effect across price groups:

$$ManualDesludging_i = \alpha + \beta TargetedPricesTreatment_i + \gamma' X_i + \epsilon_i \quad (26)$$

Observations are at the household level and standard errors are clustered at the neighborhood-cluster level (92 clusters in total). $ManualDesludging_i$ is an indicator taking the value of 1 if household i reports having used a manual desludging between the baseline and the endline, and takes the value of 0 otherwise. $TargetedPricesTreatment_i$ takes the value of 1 for all households in treatment clusters, 0 for households in control clusters. X_i is a vector of controls for variables unbalanced across neighborhoods at baseline. We also control for whether the household had a child suffering from diarrhea at baseline, following McKenzie (2012).

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:

$$\begin{aligned}
Manual_i = & \sum_{k=1}^4 \alpha_k PriceGroup_{ki} \\
& + \sum_{k=1}^4 \beta_k TargetedPricesTreatment_i * PriceGroup_{ki} + \gamma' X_i + \epsilon_i \quad (27)
\end{aligned}$$

$PriceGroup_{ki}$ 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 observable variables 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 10. While we are under-powered to find an effect in the pooled regression, we find that the point estimate on manual desludging overall is a 1.5 percentage point decrease in the probability that a household uses a manual desludging. At the mean of 26.3% using manual desludging at baseline, this is a 5.7% effect.

We can also see that the use of manual desludging among those targeted with low prices is much higher than those receiving the highest targeted prices—these low price groups were the households whose use of manual desludging we were most interested in changing. The effect on the use of manual desludging for those in the lowest price group is large: we estimate that the probability that a household uses manual during the period since baseline decreases by 4 percentage points relative to the control group for the lowest price group. The lowest price group had an average probability of manual desludging of 55% at baseline, so this is a 7.3% effect at baseline.

We are particularly interested in the impact of the program on health. Sanitation has important effects, particularly on children's health since the sewage waste from manual desludging is commonly disposed of in the street where children play during the day. We run the same specification for children's diarrhea rates as we ran for manual desludging. We expect our estimates on child diarrhea rates to be lower bounds, because while the question was asked about diarrhea rates in the past 7 days, any manual desludging impacts could have occurred any time over the preceding 17 months of the program.

Results are shown in table 10 columns (3) and (4). We find that the overall impact on children's diarrhea rates from being in a treatment cluster is a 1.1 percentage point reduction in the probability of a diarrhea episode being reported, but the effect is not significant at standard levels. However, if we estimate the effects across the different price groups, we find that we did have a large and statistically significant impact on child diarrhea rates on the lowest price group. The targeted price treatment had caused a 6.1 percentage point decrease in the probability that a child in a household in the lowest price group had an episode of diarrhea in the past week (significant at the 10% level). The average diarrhea rate among this subgroup is 13%, so this is a very large 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 following sanitation treatments of 0.68. As expected from the point estimates, the relative risk ratio for our pooled sample is 0.96, which is close to 1 and suggests little impact in the group 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 0.68. 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.

We can observe the extent to which the model is targeting relatively poor households who are more likely to get manual desludgings in table 9. 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 9 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.

5.6 Comparison With a Subsidy Program

While the previous results establish that the program had a number of significant effects, simpler options are available that might have delivered similar outcomes. This section exploits the availability of multiple datasets to train a predictive model and provide estimates of the impact of a subsidy program that was subject to the same budget constraints as the demand-maximizing platform. We show that this counterfactual program fails to deliver a statistically significant impact on the lowest price group, suggesting that our targeting strategy was an important component in achieving the program’s results.

The platform was designed using the assumption that the average cost of a mechanical desludging would be approximately 17,500 CFA and that a subsidy of 1,935 CFA could be raised for each participating household, resulting in a subsidized price of approximately 15,500 CFA. We exploit household data from the first phase and additional data from households included in the platform to boost demand but not exploited in other results, but set aside the data for our targeted price and control groups as a test set. The quantity of interest is the probability a household i selects mechanical, conditional on its observable characteristics X_i and the price quoted by the platform, t_i , $pr[Mech|X_i, t_i]$. Note that the quantity of interest is not whether the household purchases through

the platform, since many households reject our offer but still purchase a mechanical desludging. Our outcome of interest is the market share of mechanical services. Similarly, the household is treated simply by having the platform available, since it provides a means of negotiating against desludgers even if the household decides not to purchase a desludging from the platform. We employ a probit model to predict the likelihood a household purchases a mechanical desludging, conditional on its characteristics and a given platform price²³.

We then evaluate the predictive model at a price of 15,500 CFA for all households, using the covariates in the actual treatment group. This provides a predicted probability that each such household selects a mechanical desludging given our quote. These estimates are used to construct market share outcome variables exactly as in Section 4.1 and 4.2 for the counterfactual treatment group and the actual control group. These counterfactual results are reported in Table 11.

The counterfactual subsidies have an undetectable pooled effect and a positive but statistically insignificant effect of 1.5% on the subset of households that would have received a price quote of 10,000 CFA from the optimal platform. The intuition comes from figure 6: the poorer households who exhibit a low market share for mechanical find a price even of 12,500 CFA unattractive, and certainly opt out at a price of 15,500 CFA. Consequently, the subsidies only benefit richer households, who are already consuming the mechanical service at rates approaching 86% and 95%. Thus, a standard subsidy program in this environment seems to face exactly the problems described in the introduction.

5.7 The Supply Side and Cost-Effectiveness

Previous results have focused on the demand side, but we now turn to the supply side and issues of cost minimization. The traditional answer is an auction (Myerson (1981)), which is the first market design we adopted. This section shows that while auctions tend to achieve the mean price in the market, a slightly different design — structured negotiations — can substantially further

²³We considered a variety of machine learning models including ridge regression and the Lasso, but find that a standard probit model provides the best fit on both the training data and the test set. We plan to further explore this class of models, but found that at the optimum cross-validated penalties, ridge regression and the Lasso over-predicted mechanical desludging by approximately 10%, even on the training set.

reduce prices.

We surveyed and worked with 33 desludging operators to procure 477 desludgings (this comprised a census of the desludging operators that we were able to locate in the region of Ouagadougou, Burkina Faso). We conducted procurement in two different ways: we started with monthly procurement auctions, and after

For the first fourteen rounds, we utilized the following design in each round, often called a lowest-rejected bid (LRB) auction, or a $W+1$ -st price auction:

- i. For each neighborhood $k = 1, \dots, K$, there are a maximum of W_k winners.
- ii. Each firm i is asked to submit a bid for each neighborhood, $b_i = (b_{i1}, \dots, b_{iK})$. The W_k lowest bids in each neighborhood are selected as winners for that round, and are each paid the W_k+1 -st bid whenever they complete a job for the platform.

Since desludgers are typically not capacity constrained nor subject to other dynamic constraints — during the rainy season, they report doing up to 10 jobs a day, while rarely coming close to that number in the dry season — it is a weakly dominant strategy to bid the expected cost of a job for each desludger in each neighborhood. This was illustrated to desludgers during training sessions through simple, one-unit examples, then larger games that approximated the actual LRB auctions, and the information was repeated again when bids were solicited from round to round. This was intended to allow clean identification of firms’ cost structures across neighborhoods.

We doubt, however, that the received bids reflect the firm’s true costs. The left panel of Figure 7 illustrates the clearing prices in each neighborhood over time. Clearing prices quickly converged to 15,000 CFA, the average market price for a desludging, and firms rarely bid below 15,000 CFA despite prices lower than 15,000 CFA being observed in market pricing data²⁴. Faced with such stubborn and persistent convergence of clearing prices, we considered a number of possibilities: that 15,000 CFA is the average expected cost of a job once risk aversion is incorporated, that the firms were explicitly colluding, or that the lowest-rejected-bid auction was a potentially unnatural form of competition for this environment. Hoping to encourage more competition, we switched to

²⁴This motivated our choice of a relatively high expected cost of $\bar{c} = 17,500$ when solving for the optimal prices: if competition softened further and prices increased, we could face significantly higher prices than expected.

a paid-as-bid auction with the same number of winners starting in round 15. This led to some experimentation and temporarily lower prices in a few neighborhoods, but convergence back to a price of 15,000 CFA within a few rounds.



Figure 7: Price paths

Once it became clear that auctions could only achieve the average market price, we adopted a different approach: structured negotiations. The protocol is simple:

- i. Each neighborhood b has a standing low bidder with a price of p_{slb} . All desludgers are initially ranked by the average of their past bids in that neighborhood.
- ii. When a desludging must be procured, we first call the *second*-lowest-cost firm, and ask them to do the job at $p_{slb} - 500$ CFA. If that firm accepts, the job is offered to the household at that price and that desludger becomes the new standing low bidder; if the desludger rejects, we repeat this offer with the second-least cost, third-least cost, and a randomly chosen desludger in that neighborhood.
- iii. If step 2 fails to provide a new standing low bidder, the standing low bidder is offered the job at his standing low bid.
- iv. If the standing low bidder rejects or is unavailable, we offer the job to randomly selected desludgers at the standing low bid plus 500 CFA until a firm agrees.

Prices quickly fell below the market price of 15,000 CFA in most neighborhoods, and by the end of the project, the mean price in some neighborhoods was below 10,000 CFA. This provides strong evidence that cost uncertainty was not the main concern; otherwise, the fact that some uncertainty was resolved about the jobs should yield a mean price of 15,000 CFA on average, but with higher and lower bids realized depending on idiosyncratic shocks like traffic or weather. Similarly, it suggests explicit collusion is not a compelling explanation, because whatever methods the desludgers previously used to enforce the cartel in the auctions should have carried over to this new environment. Thus, common causes for artificially high prices do not seem to entirely explain the convergence of auctions to the average market price, while such a simple system typically yields lower prices.

However, negotiations mimic auctions: the price a desludger charges today is a bid for future business. There are potential losses in efficiency by only considering a small number of competitors and bounding the amount by which the price can fall in a given negotiation, but we appear to gain

much more by asking them to individually outbid one another on a specific, concrete job. This leads to significantly lower costs of procurement.

Since the prices paid by the demand side in the fixed-price treatment are fixed, the method by which the jobs were procured is essentially independent of their decisions and participation. Thus, we can consider four potential cost estimates:

- i. Round-by-neighborhood auction prices: For each neighborhood and each round, use the clearing price realized in the auctions.
- ii. Average negotiation prices by neighborhood: For each neighborhood, use the mean of the standing low bids over time.
- iii. Last observed auction prices by neighborhood: For each neighborhood, use the last auction price observed.
- iv. Last observed negotiation price by neighborhood: For each neighborhood, use the last standing low bid observed.

The round-by-neighborhood auction price estimate is the most pessimistic — it includes early rounds when prices were higher than the market average — but also most closely measures the financial position of the platform over the course of the study. Since negotiations began relatively late, observations of prices in early rounds are not available, motivating the use of average prices for an estimate of how the platform would have fared if negotiation prices were used exclusively. These two estimates, however, ignore the decreasing nature of prices over time, motivating the use of the last observed prices under auctions and under negotiations as a method of determining whether or not such a platform is sustainable in the long-run at current subsidy levels.

To compute these financial estimates, we take the set of households who actually purchase, subtract the price paid to the desludger from the price paid by the household under each of these scenarios, and then add the subsidy of about 3.00. Thus, for each participating household i whose desludging is provided by firm j , we compute

$$t_i - b_j + S/N,$$

which is the difference between the price paid by household i and the price paid to the winning firm b_j , and S/N is the average subsidy, equal to about \$3.00 USD or 1,750 CFA. Averaging over all the households yields a budget estimate for the platform, shown in table 12.

This shows that our goal to satisfy the expected budget constraint was largely realized: the budget surplus or deficit is on the order of a few dollars. This does, however, ignore that we expected prices to be higher, around 17,500 CFA or \$30 USD. Under this much more pessimistic cost estimate, the platform would have lost approximately 167,875 CFA, or \$305.23 dollars^{25,26}.

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 competing search 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

²⁵Our conjecture is that correlations between shocks ε_i at the neighborhood level explain this, since it is consistent with expected losses of \$300 to \$500 based on Monte Carlo simulations using the original model and cluster-based re-sampling.

²⁶Sustainability counterfactual goes here.

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

- 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.
- Bagnoli, M. and Bergstrom, T. (2005). Log-concave probability and its applications.
- Basurto, P., Dupas, P., and Robinson, J. (2017). Decentralization and efficiency of subsidy targeting: Evidence from chiefs in rural Malawi.
- 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).
- 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.

- Cragg, J. (1971). Some statistical models for limited dependent variables with application to the demand for durable goods. *Econometrica*.
- 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*.
- 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.
- 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*.
- Mussa, M. and Rosen, S. (1978). Monopoly and product quality. *Journal of Economic theory*.
- Myerson, R. (1981). Optimal auction design. *Mathematics of Operations Research*.
- 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*.
- Olken, B. (2016). Hassles versus prices. *Science*, 6302:864–865.
- Spears, D. (2013). How much international variation in child height can sanitation explain? *Working Paper*.
- 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*.

7 Appendix A: Section 3 Proofs

In the appendix, we analyze a slightly more general model. There is a unit mass of households, who each have observable covariates, $x \in X$, and a privately known willingness-to-pay, $w \in [\underline{w}, \bar{w}]$. The observables x are drawn from a probability space (X, Σ, μ) where X is the set of potential observables, Σ is the Borel σ -algebra over X , and μ is a countably additive probability measure over observables. This formulation allows observable characteristics of households to be continuous, like the amount of water or electricity purchased in a month, as well as discrete, like the number of children in the household. Alternatively, the reader can think of X as a finite set, $\mu(x)$ the probability of observing observable characteristics x , and replace Lebesgue integrals in what follows with summations over x , as in the text of Section 3. All other assumptions are the same.

Lemma 1 *A platform is incentive compatible iff (i) $p(w, x)$ is increasing in w for all x and (ii)*

$$t(w, x) = p(w, x)(w - U^0(w, x)) - U(w_x^*, w_x^*, x) - \int_{w_x^*}^w p(z, x)(1 - p_x^0 \mathbb{I}\{z \geq w_x^0\}) dz, \quad (28)$$

where w_x^* is the type that joins the platform and receives the lowest payoff among those who join.

Proof: If the prevailing search market is incentive compatible, then the envelope theorem applied to the indirect utility function yields i. Similarly, take two incentive compatibility constraints $U_0(w, w, z_i) \geq U_0(w', w, z_i)$ and $U_0(w', w', z_i) \geq U_0(w, w', z_i)$. Adding the constraints together and rearranging yields

$$(w - w')(p_0(w, z_i) - p_0(w', z_i)) \geq 0.$$

Without loss of generality, suppose $w > w'$: then for the inequality to hold, $p_0(w, z_i) \geq p_0(w', z_i)$.

■

The standard mechanism design algorithm proceeds as follows: pose a relaxed problem in which the monotonicity condition (i) is dropped; solve the relaxed problem and check whether the solution satisfies (i); if the solution satisfies (i), the relaxed solution is a solution to the original problem, and if the relaxed solution does not satisfy (i), add a pooling/ironing interval where the allocation is stochastic.

Our relaxed problem is as follows:

$$\max_{p,t} \int_x \mu(x) \int_{\underline{w}}^{\bar{w}} p(w,x) + (1-p(w,x))p_x^0 \mathbb{I}\{w \geq w_x^0\} dF[w|x] dx$$

subject to incentive compatibility,

$$t(w,x) = p(w,x)(w - U^0(w,x)) - U(w_x^*,x) - \int_{w_x^*}^w p(z,x)(1 - p_x^0 \mathbb{I}\{z \geq w_x^0\}) dz,$$

budget balance,

$$\int_x \mu(x) \int_{\underline{w}}^{\bar{w}} t(w,x) - p(w,x)c_x dF[w|x] dx + S \geq 0,$$

and individual rationality for the lowest type that trades on the platform,

$$U(w_x^*, w_x^*, x) \geq 0.$$

This means we have dropped the monotonicity condition (i) and individual rationality for all types besides the lowest that trades on the platform in defining the relaxed problem.

Note that we can set $U(w_x^*, w_x^*, x) = 0$ since if the lowest type that trades on the platform is receiving a strictly positive net payoff, we could raise the tariff at the bottom without affecting the incentive constraints for higher types, and raise more money. With a relaxed budget, we could then improve the objective. Consequently, $U(w_x^*, w_x^*, x) = 0$ at the relaxed optimum.

Now taking the expectation over $t(w,x)$ with respect to (w,x) and integrating by parts with respect to w yields expected revenue,

$$\mathbb{E}_{(w,x)}[t(w,x)] = \int_x \mu(x) \int_{\underline{w}}^{\bar{w}} \underbrace{\left\{ w - U^0(w,x) - \frac{1 - F[w|x]}{f[w|x]} (1 - \mathbb{I}\{w \geq w_x^0\} p_x^0) \right\}}_{\psi(w,x)} dF[w|x] dx.$$

This allows us to re-write the budget balance constraint as

$$\int_x \mu(x) \int_{\underline{w}}^{\bar{w}} p(w,x)(\psi(w,x) - c_x) dF[w|x] dx + S \geq 0.$$

This further simplifies the relaxed problem to

$$\max_{p,t} \int_x \mu(x) \int_{\underline{w}}^{\bar{w}} p(w,x) + (1-p(w,x))p_x^0 \mathbb{I}\{w \geq w_x^0\} dF[w|x] dx$$

subject to

$$\int_x \mu(x) \int_{\underline{w}}^{\bar{w}} p(w,x)(\psi(w,x) - c_x) dF[w|x] dx + S \geq 0.$$

We now prove existence of a solution and characterize the set of maximizers for the relaxed problem:

Lemma 2 *A solution p^* to the relaxed problem exists and the set of maximizers is convex.*

Proof: The set $[0,1]$ is compact, so by Tychonoff's theorem, $\Delta = [0,1]^{X \times [\underline{w}, \bar{w}]}$ is compact in the product topology. It is also convex, since convex combinations of numbers between 0 and 1 are also between 0 and 1. In particular, consider the space of functions $L^1(\Delta)$, with norm $\|g\| = \int_{x \in X} \mu(x) \int_{\underline{w}}^{\bar{w}} |g(w,x)| dF[w|x] dx$. This is a Banach space, so it is complete.

The objective function is

$$\Gamma(p) = \int_{x \in X} \mu(x) \int_{\underline{w}}^{\bar{w}} p(w,x) + (1-p(w,x)) \mathbb{I}\{w \geq w_x^0\} p_x^0 dF[w|x] dx$$

which is a linear functional of p , since $\Gamma(ap + bp') = a\Gamma(p) + b\Gamma(p')$ for all $a, b \in \mathbb{R}$, so it is also concave. Let $\varepsilon > 0$ be given. Then if $\|p - p'\| < \delta = \varepsilon$, then

$$\begin{aligned} |\Gamma(p) - \Gamma(p')| &= \left| \int_{x \in X} \mu(x) \int_{\underline{w}}^{\bar{w}} (p(w,x) - p'(w,x))(1 - \mathbb{I}\{w \geq w_x^0\} p_x^0) dF[w|x] dx \right| \\ &\leq \left| \int_{x \in X} \mu(x) \int_{\underline{w}}^{\bar{w}} |p(w,x) - p'(w,x)| (1 - \mathbb{I}\{w \geq w_x^0\} p_x^0) dx \right| \\ &\leq \left| \int_{x \in X} \mu(x) \int_{\underline{w}}^{\bar{w}} |p(w,x) - p'(w,x)| p_x^0 dF[w|x] dx \right| \\ &\leq \left| \int_{x \in X} \mu(x) \int_{\underline{w}}^{\bar{w}} |p(w,x) - p'(w,x)| dF[w|x] dx \right| \\ &\leq \|p - p'\| \\ &< \varepsilon \end{aligned}$$

Therefore the objective is concave and continuous.

The subset of Δ which is comprised of functions that satisfy the budget-balance constraint is compact and convex. Assume that the sequence $\{p_n\}_{n=1}^{\infty}$ satisfies the budget balance constraint for all n , $p_n \rightarrow p$,

$$\int_{x \in X} \mu(x) \int_{\underline{w}}^{\bar{w}} p_n(w, x) (\psi(w, x) - c_x) dF[w|x] dx + S \geq 0,$$

but

$$\int_{x \in X} \mu(x) \int_{\underline{w}}^{\bar{w}} p(w, x) (\psi(w, x) - c_x) dF[w|x] dx + S < 0,$$

which jointly imply that

$$\int_{x \in X} \mu(x) \int_{\underline{w}}^{\bar{w}} (p_n(w, x) - p(w, x)) (\psi(w, x) - c_x) dF[w|x] dx > 0.$$

But this contradicts convergence of p_n to p , since it implies there is a set of strictly positive measure such that p_n and p are sufficiently far apart that the integral is strictly positive. Therefore p must satisfy the budget balance constraint, so that the subset of Δ which is comprised of functions that satisfy the budget balance constraint is closed, and therefore compact as a subset of Δ . Now take two allocations p and p' that satisfy budget balance. Taking the convex combination of the two constraints yields

$$\int_{x \in X} \mu(x) \int_{\underline{w}}^{\bar{w}} (\lambda p(w, x) + (1 - \lambda) p'(w, x)) (\psi(w, x) - c_x) dF[w|x] dx + S \geq 0,$$

so that the convex combination satisfies budget balance. Therefore, the subset of Δ that satisfies budget balance is a compact, convex set.

Therefore, the relaxed problem considers a continuous function over a compact set, so a solution exists. Since the objective is concave and the feasible set is convex, the set of maximizers is convex, since if p_1^* and p_2^* are both solutions, $\lambda p_1^* + (1 - \lambda) p_2^*$ is in the feasible set and $\Gamma(\lambda p_1^* + (1 - \lambda) p_2^*) \geq \lambda \Gamma(p_1^*) + (1 - \lambda) \Gamma(p_2^*)$. ■

The Lagrangian for the relaxed problem then is

$$\mathcal{L}(p, \lambda) = \lambda S + \int_x \mu(x) \int_{\underline{w}}^{\bar{w}} p(w, x, \lambda) + (1 - p(w, x, \lambda)) p_x^0 \mathbb{I}\{w \geq w_x^0\} + \lambda p(w, x, \lambda) (\psi(w, x) - c_x) dF[w|x] dx. \quad (29)$$

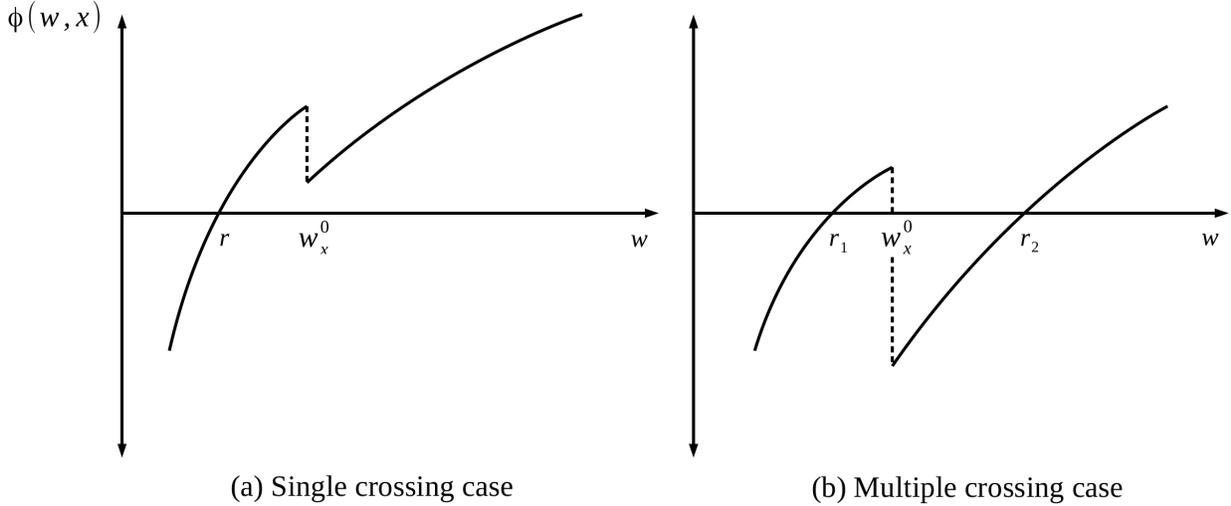
Since we are maximizing a concave function over a convex set, the critical points of the Lagrangian (including the additional constraints that $p(w, x) \geq 0$ for all (w, x) and $p(w, x) \leq 1$ for all (w, x)) are global maximizers. We will now take the Lagrange multiplier λ as given and solve the maximization problem in p , then use the budget balance constraint to solve for λ^* .

Inspecting the Lagrangian, the coefficient on $p(w, x, \lambda)$ is

$$\phi(w, x, \lambda) = 1 - p_x^0 \mathbb{I}\{w \geq w_x^0\} + \lambda (\psi(w, x) - c_x).$$

Since probabilities of trade must be between 0 and 1, the relaxed solution therefore sets $p(w, x, \lambda) = 1$ if $\phi(w, x, \lambda) \geq 0$ and $p(w, x, \lambda) = 0$ otherwise. Since this is a “bang-bang” solution, the relaxed solution satisfies the monotonicity condition if $\phi(w, x, \lambda)$ satisfies the single crossing condition: there exists an $r \in [\underline{w}, \bar{w}]$ such that for all $w \leq r$ and $\lambda \geq 0$, $\phi(w, x, \lambda) \leq 0$ and for all $w \geq r$, $\phi(w, x, \lambda) \geq 0$.

Note that under the standard regularity condition that $1 - F[w|x]$ is log-concave, $\psi(w, x)$ is increasing in w everywhere except w_x^0 , where there is a discontinuous drop downward; this implies $\phi(w, x, \lambda)$ is also increasing everywhere except at w_x^0 . This implies there are at most two crossings, yielding the following graph:



Note that a sufficient condition for $\phi(w, x, \lambda)$ to satisfy the single-crossing property for all $\lambda \geq 0$ is that for all x ,

$$\psi(w_x^0, x) - c_x = w_x^0 - \frac{1 - F[w_x^0|x]}{f[w_x^0|x]}(1 - p_x^0) - c_x \geq 0.$$

The economic content of this condition is that the marginal revenue generated by the marginal type in the prevailing market w_x^0 is positive, so that a profit-maximizing platform would find it profitable to sell to w_x^0 . This implies that $\phi(w, x, \lambda)$ satisfies the single crossing property

This splits the analysis into two cases:

- (a) $\phi(w, x, \lambda)$ satisfies the single-crossing condition for all $x \in X$ and $\lambda \geq 0$, and the bang-bang solution already discussed is optimal
- (b) There is an interval $[r_1, r_2)$ for which $\phi(w, x, \lambda)$ is positive and then negative for some x and λ , and the monotonicity condition might be violated at the relaxed solution.

In case (b), the non-monotonicity occurs at the bottom, but the platform will always wish to sell to types with high hidden information. We now add the pooling interval explicitly to the analysis and show that, in fact, non-trivial randomization will never be optimal. In particular, allow for an interval $[w_x^*, a)$ at the bottom with probability of trade p_x^a and payment t_x^a . Then the individual rationality constraint for w_x^* determines the price in the pooling interval,

$$p_x^a(w_x^* - U^0(w_x^*, x)) - t_x^a = 0,$$

or

$$t_x^a = p_x^a(w_x^* - U^0(w_x^*, x)),$$

and the downwards incentive constraint at a takes the form

$$U(a, a, x) \geq p_x^a a - t_x^a.$$

This is an equality, since if it was strictly larger than zero, the platform could raise prices on a and all higher types without disrupting incentive compatibility or individual rationality, raise more money, and expand the market to more agents. This implies that for types $w > a$,

$$t(w, x) = p(w, x)(w - U^0(w, x)) - U(a, a, x) - \int_a^w p(z, x)(1 - \mathbb{I}\{z \geq w_x^0\}p_x^0)dz.$$

We can now compute the new Lagrangian incorporating the pooling interval

$$\begin{aligned} \mathcal{L}(p, p_x^a, t_x^a, w_x^*, \lambda) &= \underbrace{\int_{x \in X} \mu(x) \int_{\underline{w}}^{w_x^*} \mathbb{I}\{w \geq w_x^0\} dF[w|x] dx}_{\text{Off-platform trade}} + \underbrace{\int_{x \in X} \mu(x) \int_{w_x^*}^a p_x^a + \lambda(t_x^a - p_x^a c_x) dF[w|x] dx}_{\text{Pooled, random allocation}} \\ &+ \underbrace{\int_{x \in X} \mu(x) \int_a^{\bar{w}} p(w, x) + \lambda(p(w, x, \lambda)(\psi(w, x) - c_x) - U(a, a, x)) dF[w|x] dx}_{\text{Potential separation}} \end{aligned}$$

Substituting in the expressions for $U(a, a, x)$ and t_x^a and isolating the terms involving p_x^a yields

$$p_x^a \left\{ \int_{w_x^*}^a 1 + \lambda(w_x^* - U^0(w_x^*, x) - c_x) dF[w|x] - \lambda \int_a^{\bar{w}} (a - (w_x^* - U^0(w_x^*, x))) dF[w|x] \right\}.$$

The first term represents the gain in rent extraction from the pooled interval by raising p_x^a , and the second term represents the loss in rent extraction from all types $w \geq a$ since the allocation to the pooled interval has improved, and higher types must be incentivized not to deviate downwards. But note that the control of p_x^a is still linear/bang-bang, since it enters the objective linearly: thus

$p_x^a \in \{0, 1\}$, and randomization is not optimal. This implies that in case (b), either

$$\int_{r_1}^{r_2} 1 + \lambda(\psi(w, x, \lambda) - c_x) dF[w|x] \geq 0$$

and the platform includes all the types between the two roots, or the above term is strictly negative, and the platform excludes all the types between the two roots.

We now confirm that individual rationality holds for all types in either case. The net indirect utility in the mechanism for all types $w \geq w_x^*$ is

$$U(w, w, x) = \int_{w_x^*}^w 1 - \mathbb{I}\{z \geq w_x^0\} p_x^0 dz \geq (w - w_x^*)(1 - p_x^0) \geq 0,$$

so that all types have a weakly positive net benefit of participating.

The preceding analysis yields the following result:

Lemma 3 *For all λ , there exists a deterministic solution, $p^*(w, x, \lambda) \in \{0, 1\}$ for all (w, x) , which is incentive compatible and individually rational. For each $x \in X$, there is a type $w_x^*(\lambda)$ satisfying*

$$1 - p_x^0 \mathbb{I}\{w_x^*(\lambda) \geq w_x^0\} + \lambda(\psi(w_x^*(\lambda), x) - c_x) = 0 \quad (30)$$

such that if $w < w_x^(\lambda)$, w trades in the prevailing market or purchases a manual desludging, and if $w \geq w_x^*(\lambda)$, w purchases a mechanical desludging on the platform.*

This yields an allocation p^* that depends on λ , so we now solve for the λ associated with the optimal solution. To solve for λ^* , consider the function

$$\beta(\lambda) = \int_x \mu(x) \int_{w_x^*(\lambda)}^{\bar{w}} \psi(w, x) - c_x dF[w|x] dx + S$$

The budget balance constraint is satisfied if $\beta(\lambda)^* = 0$.

Due to the potential discontinuities in $w_x^*(\lambda)$, it is not obvious that $\beta(\lambda)$ is continuous in λ . To establish continuity, we will use a duality argument and Berge's Theorem of the Maximum. In

particular, consider the problem of maximizing profit,

$$\int_x \mu(x) \int_{\underline{w}}^{\bar{w}} t(w, x) - c_x dF[w|x] dx$$

subject to individual rationality, incentive compatibility, and a quota constraint.

$$\int_x \mu(x) \int_{\underline{w}}^{\bar{w}} p(w, x) + (1 - p(w, x)) p_x^0 \mathbb{I}\{w \geq w_x^0\} dF[w|x] dx \geq \bar{q}.$$

Using the same approach as for the quantity maximization problem, this problem is associated with a Lagrangian:

$$\begin{aligned} \mathcal{L}(p, \lambda) &= -\lambda \bar{q} \\ &+ \int_{x \in X} \mu(x) \int_{\underline{w}}^{\bar{w}} p(w, x) \{\psi(w, x) - c_x\} + \lambda (p(w, x, \lambda) + (1 - p(w, x, \lambda)) p_x^0 \mathbb{I}\{w \geq w_x^0\}) dF[w|x] dx \end{aligned}$$

Let $\pi^*(\lambda, \bar{q})$ be the optimized value of the objective. By Berge's Theorem, $\pi^*(\lambda, \bar{q})$ is continuous in λ and \bar{q} . Similarly, the optimized value of the objective for the quantity-maximization problem, $q^*(\lambda, S)$, is continuous in λ and S .

Lemma 4 *If p^* is optimal in the quantity-maximization problem with a subsidy of S , it is optimal in the profit-maximization problem when $q^*(\lambda, S) = \bar{q}$.*

If p^ is optimal in the profit-maximization problem with a quota of \bar{q} , it is optimal in the quantity-maximization problem when $\pi^*(\lambda, \bar{q}) + S = 0$.*

Proof: Suppose that p^* is optimal in the quantity-maximization problem with a subsidy of S , but that in the profit-maximization problem with a quota of $q^*(\lambda, S)$, there is an alternative allocation p' that yields strictly higher profits. But that implies that p' achieves the same quantity as p^* but raises strictly higher profits, implying that by extending participation in the platform to a set of types of strictly positive measure just below $w_x^*(\lambda)$, the budget constraint becomes an equality and strictly more mechanical desludgings are purchased. This is a contradiction. A similar argument establishes the second claim. ■

As a consequence of the previous Lemma, note that

$$\beta(\lambda, S) = \pi^*(\lambda, q^*(\lambda, S)) + S,$$

expresses $\beta(\lambda)$ as a composition of continuous functions in λ (by Berge's theorem applied to π^* and q^*). Therefore, β is continuous.

Note that if λ is zero, the trivial solution is for the platform to issue a desludging to everyone at a price of zero, yielding profits that are strictly negative and in excess of S , so that $\beta(0) < 0$. Conversely, if λ is taken to be sufficiently large, the platform will behave arbitrarily like a profit-maximizer. This implies it will raise a strictly positive amount of profit, plus a strictly positive subsidy. This implies there is a $\lambda' > 0$ such that $\beta(\lambda') > 0$. By the intermediate value theorem, this then implies there exists at least one solution λ^* to $\beta(\lambda) = 0$.

Since $w_x^*(\lambda)$ is characterized locally by

$$1 - p_x^0 \mathbb{I}\{w_x^*(\lambda) \geq w_x^0\} + \lambda(\psi(w_x^*(\lambda), x) - c_x) = 0,$$

it is locally differentiable, and an increase in λ implies

$$\frac{dw_x^*(\lambda)}{d\lambda} = \frac{-(\psi(w_x^*(\lambda), x) - c_x)}{\lambda \psi_w(w_x^*(\lambda), x)},$$

which has the opposite sign of $\psi(w_x^*(\lambda), x) - c_x$: if a type $w_x^*(\lambda)$ is being subsidized and $\psi(w_x^*(\lambda), x) - c_x < 0$, increasing λ leads to an increase in $w_x^*(\lambda)$; conversely, if a type $w_x^*(\lambda)$ is revenue-producing and $\psi(w_x^*(\lambda), x) - c_x > 0$, increasing λ leads to a decrease in $w_x^*(\lambda)$. Since the equation characterizing $w_x^*(\lambda)$ is continuous in λ , the implicit function theorem implies that

$$\beta'(\lambda) = - \int_x \mu(x) (\psi(w_x^*, x) - c_x) f[w_x^*(\lambda)|x] \frac{dw_x^*(\lambda)}{d\lambda} dx,$$

or

$$\beta'(\lambda) = \int_x \mu(x) \frac{(\psi(w_x^*(\lambda), x) - c_x)^2}{\lambda \psi_w(w_x^*(\lambda), x)} f[w_x^*(\lambda)|x] dx,$$

which is positive. This implies $\beta(\lambda)$ is increasing, so that there is a unique solution λ^* or a convex set $[\underline{\lambda}^*, \bar{\lambda}^*]$ for which $\beta(\lambda) = 0$.

Note that if $\psi(w, x)$ is strictly increasing, there is at most one type w' where $\psi(w', x) - c_x = 0$. This implies that $\beta'(\lambda)$ is strictly positive almost everywhere, so that there is a unique root.

The preceding analysis yields the following result:

Lemma 5 *There exists a λ^* such that $\beta(\lambda^*) = 0$. If $\psi(w, x)$ is strictly increasing in w for all $x \in X$, then there exists a unique multiplier λ^* for which $\beta(\lambda^*) = 0$.*

This characterizes the optimal allocation p^* and the multiplier λ^* .

We now show that a simple posted price scheme implements the optimal outcome. If the platform uses posted prices, a household with value w purchases if

$$w - t_x \geq U^0(w, x),$$

so that the marginal type that purchases from the platform is defined by

$$t_x = w_x^* - U^0(w_x^*, x).$$

Substituting this into (30) yields

$$1 - p_x^0 \mathbb{I}\{w_x^*(\lambda) \geq w_x^0\} + \lambda \left\{ t_x^* - \frac{1 - F[w|x]}{f[w|x]} (1 - \mathbb{I}\{w \geq w_x^0\} p_x^0) - c_x \right\} = 0,$$

which can be solved for

$$t_x^* = c_x + \left(\frac{1 - F[w_x^*|x]}{f[w_x^*|x]} - \frac{1}{\lambda^*} \right) (1 - \mathbb{I}\{w_x^* \geq w_x^0\} p_x^0)$$

The preceding analysis yields the result:

Lemma 6 *The optimal mechanism can be implemented by issuing posted prices conditional on*

observables x that satisfy

$$t_x^* = c_x + \left(\frac{1 - F[w_x^*|x]}{f[w_x^*|x]} - \frac{1}{\lambda^*} \right) (1 - \mathbb{I}\{w_x^* \geq w_x^0\} p_x^0)$$

We summarize Lemmas 1 – 8 in the following theorem given in the text:

Theorem 4 *Assume the standard regularity condition that $1 - F[w|x]$ is log-concave and let λ^* be the multiplier on the expected budget balance constraint at the optimum.*

i. For all x , there is a type w_x^ that satisfies*

$$1 - \mathbb{I}\{w_x^* \geq w_x^0\} p_x^0 + \lambda^* (\psi(w_x^*, x) - c_x) = 0,$$

and in the optimal mechanism, all types $w \geq w_x^$ trade on the platform and all types $w < w_x^*$ either purchase a mechanical desludging in the prevailing search market or get a manual desludging.*

ii. The optimal direct mechanism can be implemented by making take-it-or-leave-it offers conditional on each observable type $x \in X$, where the optimal price satisfies

$$t_x^* = c_x + \left(\frac{1 - F[w_x^*|x]}{f[w_x^*|x]} - \frac{1}{\lambda^*} \right) (1 - \mathbb{I}\{w_x^* \geq w_x^0\} p_x^0). \quad (31)$$

7.1 Hazard rate dominance

From the text, an observable type x hazard-rate dominates an observable type x' if for all $w \in [\underline{w}, \bar{w}]$,

$$\frac{f[w|x]}{1 - F[w|x]} \leq \frac{f[w|x']}{1 - F[w|x']},$$

which implies that all $w \in [\underline{w}, \bar{w}]$

$$-\frac{1 - F[w|x]}{f[w|x]} \leq -\frac{1 - F[w|x']}{f[w|x']}.$$

Set $c_x = c_{x'} = c$, $p_x^0 = p_{x'}^0 = p^0$, $t_x^0 = t_{x'}^0 = t^0$, and $U^0(w, x) = U^0(w, x') = W^0(w)$. This implies that $\phi(w, x)$ and $\phi(w, x')$ both have a downward discontinuity at $w_x^0 = w_{x'}^0 = w^0$. Note also that, pointwise, $\phi(w, x) \geq \phi(w, x')$, since

$$\begin{aligned}\psi(w, x) &= w - U^0(w) - \frac{1 - F[w|x]}{f[w|x]}(1 - \mathbb{I}\{w \geq w^0\}p^0) \\ &\geq w - U^0(w) - \frac{1 - F[w|x']}{f[w|x']}(1 - \mathbb{I}\{w \geq w^0\}p^0) = \psi(w, x'),\end{aligned}$$

so that

$$1 - p^0 \mathbb{I}\{w \geq w^0\} + \lambda(\psi(w, x) - c) \geq 1 - p^0 \mathbb{I}\{w \geq w^0\} + \lambda(\psi(w, x') - c),$$

implying that $\phi(w, x) \geq \phi(w, x')$. It is still possible, however, that $w_{x'}^* > w_x^*$ if case (b) obtains, however. To ensure that $w_x^* > w_{x'}^*$, we add the condition that $\psi(w_x^*, x)$ and $\psi(w_{x'}^*, x')$ are both non-negative, ensuring a single crossing that is ordered by the hazard rate. This also implies that $w_x^* < w^0$, since all types w^0 and above generate positive marginal revenue and will be included regardless of the size of the subsidy, so that $\mathbb{I}\{w_x^* \geq w_x^0\} = 0$.

Then $\phi(w_x^*, x) = 0$ implies

$$t_x^* = c + \left(\frac{1 - F[w_x^*|x]}{f[w_x^*|x]} - \frac{1}{\lambda^*} \right) \leq c + \left(\frac{1 - F[w_{x'}^*|x']}{f[w_{x'}^*|x']} - \frac{1}{\lambda^*} \right) = t_{x'}^*,$$

so that within this class of observable types, the prices are ordered the same way as the hazard rates.

8 Appendix B: 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?

9 Appendix C: 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).

10 Tables

Table 1: Demand model: Tobit

	Selection	Mechanical Price	Manual Price
Desludging frequency	-0.0064*** (0)	0.0052*** (0)	0.0171*** (0.0001)
Water greater than 5,000 CFA	0.054*** (0.0128)	0.4525*** (0.1694)	-1.3411 (0.9818)
Precarious House	-1.936*** (0.2928)	0.3149 (0.9712)	6.1943 (41.1283)
Concrete House, 1 story	-1.5212*** (0.2779)	0.8 (0.5397)	5.7112 (38.9457)
Rooming House	-1.332*** (0.3505)	0.9198 (1.3274)	6.5951 (46.2919)
Other households in compound	0.0507*** (0.0011)	-0.0201 (0.0129)	0.4064*** (0.0932)
Own House	-0.3507*** (0.0233)	-0.9104*** (0.249)	-0.0772 (2.1791)
Pit distance to road	-0.0051*** (0.0001)	0.0383*** (0.0017)	-0.0156 (0.01)
Last trips greater than one	0.4515*** (0.1288)	5.9097*** (0.7752)	-4.6435 (14.832)
Electricity bill	0.0385*** (0)		
Number persons in household	0.0058*** (0.0002)		
Number women in household	0.0448*** (0.0012)		
Household head educated	0.4241*** (0.0192)		
Constant	1.7874*** (0.3032)	16.601*** (0.8404)	3.5416 (46.7036)
$\ln(\sigma_{mech})$	1.4526*** (0.0015)		
$\ln(\sigma_{man})$	2.0821*** (0.0083)		
$atanh(\rho_{mech,0})$	-0.5460*** (0.0149)		
$atanh(\rho_{man,0})$	-1.0937*** (0.0611)		
<i>N</i>	773	530	243

Notes: High education, electricity, household size, number of women excluded from second stage pricing equations.

Table 2: 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 3: Balance Tests: Household level

	(1) Control(SD)	(2) Diff Treat - Control(SE)	(3) Observations
Number of Members of Household	6.829 (4.17)	-0.122 (0.23)	2944
Number of Women in Household	2.435 (1.70)	-0.0293 (0.08)	2937
Indicator: Respondent completed Secondary	0.306 (0.46)	-0.0314 (0.03)	2944
Number of years lived at compound	18.82 (13.63)	-0.250 (1.09)	2863
Indicator: Precarious House	0.126 (0.33)	-0.00629 (0.02)	2944
Indicator: Concrete living structure, 1 story	0.766 (0.42)	0.0168 (0.02)	2944
Indicator: Rooming House	0.0452 (0.21)	0.0127 (0.01)	2944
House Owned by Inhabitants	0.770 (0.42)	0.000835 (0.02)	2944
Water bill more than 5,000 CFA	0.484 (0.50)	-0.0633* (0.03)	2944
Electricity Bill (in thousands of CFA)	13.86 (15.43)	-1.457 (0.93)	2944
Latrine Pit Distance to Road	5.517 (4.42)	-0.810** (0.31)	2944
Two tanks used last Desludging	0.0234 (0.15)	-0.0149** (0.01)	2944
N Months Between Desludgings	27.18 (27.19)	2.984 (1.54)	2944
Last Desludging was Manual	0.263 (0.44)	0.0197 (0.03)	2305
Has ever used Manual	0.222 (0.42)	0.000327 (0.03)	2944
Has never Desludged Here	0.312 (0.46)	0.0288 (0.03)	2944
Compound has 1 pit only	0.336 (0.47)	-0.0630* (0.03)	2944
Needed Extension Hose Last Time	0.243 (0.43)	-0.0462 (0.03)	1510
Number of desludgings done at Hhd	10.76 (27.46)	-1.136 (1.85)	1793
Of 5, N closest neighbors using Manual	0.538 (1.03)	-0.00854 (0.08)	2497

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. There are 92 clusters: 40 control clusters with 1284 households and 52 treatment clusters, with 1660 treatment households.

Table 4: Balance Tests: Cluster level

	(1) Control(SD)	(2) Diff Treat - Control(SE)	(3) Observations
Number of Members of Household	6.808 (1.12)	-0.126 (0.23)	92
Number of Women in Household	2.439 (0.37)	-0.0418 (0.08)	92
Indicator: Respondent completed Secondary	0.316 (0.15)	-0.0371 (0.03)	92
Number of years lived at compound	18.75 (5.58)	-0.119 (1.09)	92
Indicator: Precarious House	0.119 (0.11)	-0.000878 (0.02)	92
Indicator: Concrete living structure, 1 story	0.769 (0.12)	0.0139 (0.02)	92
Indicator: Rooming House	0.0458 (0.05)	0.0138 (0.01)	92
House Owned by Inhabitants	0.769 (0.09)	0.000489 (0.02)	92
Water bill more than 5,000 CFA	0.490 (0.13)	-0.0644** (0.03)	92
Electricity Bill (in thousands of CFA)	14.27 (5.58)	-1.744* (1.03)	92
Latrine Pit Distance to Road	5.581 (1.60)	-0.802** (0.34)	92
Two tanks used last Desludging	0.0241 (0.03)	-0.0158*** (0.01)	92
N Months Between Desludgings	27.30 (7.98)	2.799* (1.59)	92
Last Desludging was Manual	0.260 (0.14)	0.0285 (0.03)	92
Has ever used Manual	0.218 (0.12)	0.00545 (0.03)	92
Has never Desludged Here	0.313 (0.16)	0.0258 (0.03)	92
Compound has 1 pit only	0.335 (0.13)	-0.0619* (0.03)	92
Needed Extension Hose Last Time	0.244 (0.14)	-0.0498 * (0.03)	91
Number of desludgings done at Hhd	9.913 (8.93)	-1.089 (1.69)	92
Of 5, N closest neighbors using Manual	0.532 (0.38)	0.0116 (0.07)	92

Note: The first column provides the mean of the variables averaged at the cluster level 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. There are 92 neighborhoods: 40 control neighborhoods with 1284 households and 52 treatment neighborhoods, with 1660 treatment households.

Table 5: 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 6: 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 7: Market Share Effects of Treatment

	Mkt share Mechanical	Pooled Effect	By Price Group
Overall	71.2	0.052* (0.028)	
Price of 10k group	45.5		0.091* (0.052)
Price of 15k group	77.7		0.016 (0.049)
Price of 17.5k group	85.7		0.029 (0.051)
Price of 20k group	94.2		-0.037 (0.074)
<i>N</i>		92	300

Notes: column (1) gives the average neighborhood market share for each price group of mechanical desludging at baseline where market share is defined as $\frac{n_{mechanical\ desludgings}}{N_{mechanical+manual\ desludgings}}$. Column (2) provides the OLS estimate of the pooled effect with observations at the neighborhood cluster level. Column (3) 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). Controls are included 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, electricity bill, latrine pit distance to road, two tanks used last desludging, months between desludgings, compound only has 1 pit, household requires an extension hose, and the stratification variable—number of households in cluster with low walls). Results are very similar if controls for controls only for variables unbalanced at the household level (available on request).

Table 8: 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
Targeted Price Group	0.017** (0.008)		0.013* (0.007)		-0.011 (0.008)		-0.013* (0.007)	
Int Treatment*N Desludg.			0.005 (0.007)	0.004 (0.007)			0.004 (0.009)	0.004 (0.009)
Int 10k group*Treatment		0.030* (0.016)		0.028* (0.014)		-0.029* (0.017)		-0.032** (0.016)
Int 15k group*Treatment		0.010 (0.013)		0.007 (0.012)		0.001 (0.012)		-0.002 (0.012)
Int 17k group*Treatment		0.015 (0.015)		0.011 (0.017)		-0.025 (0.018)		-0.029 (0.019)
Int 20k group*Treatment		0.045* (0.024)		0.041 (0.025)		0.008 (0.032)		0.005 (0.031)
10k price group		-0.115*** (0.020)		-0.114*** (0.020)		0.119*** (0.020)		0.119*** (0.019)
15k price group		-0.089*** (0.022)		-0.088*** (0.023)		0.083*** (0.022)		0.085*** (0.022)
17k price group		-0.064** (0.024)		-0.062** (0.025)		0.066** (0.026)		0.067** (0.026)
20k price group		-0.063** (0.030)		-0.062** (0.030)		0.044 (0.030)		0.045 (0.030)
Last Desludging Mech	0.154*** (0.017)	0.145*** (0.017)	0.154*** (0.017)	0.145*** (0.017)	-0.164*** (0.017)	-0.155*** (0.018)	-0.164*** (0.017)	-0.155*** (0.018)
Never Desludged	0.075*** (0.017)	0.066*** (0.018)	0.075*** (0.017)	0.066*** (0.018)	-0.078*** (0.017)	-0.068*** (0.018)	-0.078*** (0.017)	-0.068*** (0.018)
Constant	-0.098*** (0.019)		-0.097*** (0.020)		0.096*** (0.019)		0.097*** (0.019)	
<i>N</i>	2944	2944	2944	2944	2944	2944	2944	2944
<i>R</i> ²	0.775	0.851	0.775	0.851	0.205	0.273	0.205	0.273
<i>mean</i>	0.334	0.334	0.334	0.334	0.080	0.080	0.080	0.080

Notes: Observations are at the household level. The constant is dropped in specifications (2),(4),(6), and (8). Specifications (3) and (4) and specifications (7) and (8) include a control for the interaction between the treatment variable and the number of desludgings done by the household. Included, but not shown for brevity are controls for the variables not balanced at baseline: distance from latrine pit to the road, more than one trip necessary on last desludging, water bill greater than 5,000 CFA, and an indicator for the compound has a single latrine pit. A control for the stratification variable: less than half of compound walls in the neighborhood are high, is also included but not shown. All specifications include dummies for the number of desludgings done by the household, with the dummy for no desludgings dropped. Standard errors, clustered by neighborhood, are in parentheses.

Table 9: Mean Baseline Characteristics by Price Group

	10000	15000	17500	20000	Pooled
Phone Credit use over past week	1107 (1929)	1754 (2930)	4078 (5660)	4882 (10195)	2157 (4152)
Number of Refrigerators	0.168 (0.430)	0.507 (0.665)	0.927 (0.771)	1.371 (0.726)	0.530 (0.706)
Number of Cars	0.061 (0.248)	0.298 (0.577)	0.671 (0.838)	1.529 (1.073)	0.357 (0.683)
Number of Air Conditioners	0.016 (0.157)	0.081 (0.359)	0.477 (1.004)	1.486 (1.909)	0.199 (0.722)
Ever Desludged Mech	0.357 (0.479)	0.571 (0.495)	0.621 (0.486)	0.686 (0.468)	0.524 (0.499)
Expected Price Mechanical (CFA)	12792 (4717)	14103 (4743)	15847 (5550)	16716 (7120)	14243 (5173)
Last used Manual	0.510 (0.500)	0.219 (0.414)	0.153 (0.361)	0.030 (0.171)	0.263 (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 10: Impact of Treatment on Children's Diarrhea

	(1)	(2)
	Any Child Diarrhea	Any Child Diarrhea
Targeted Price Group	-0.009 (0.017)	
Int 10k group*Treatment		-0.071** (0.034)
Int 15k group*Treatment		0.016 (0.022)
Int 17k group*Treatment		0.015 (0.036)
Int 20k group*Treatment		0.015 (0.068)
10k price group		0.079* (0.042)
15k price group		0.005 (0.038)
17k price group		0.010 (0.044)
20k price group		-0.007 (0.053)
Number Household Children	0.011** (0.004)	0.010** (0.004)
Constant	0.028 (0.036)	
<i>N</i>	1772	1772
<i>R</i> ²	0.013	0.145
<i>mean</i>	0.129	0.129

Notes: observations are at the household level, standard errors are clustered at the neighborhood cluster level. The dependent variable is whether a child in the household has had diarrhea in the past 7 days. Specification (1) is the pooled effect across all price groups. Specification (2) provides the estimates of the effect for each price group. The constant has been suppressed in regression (2) in order to retain coefficients on all price group variables. 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. Included, but not shown for brevity are controls for the variables not balanced at baseline: distance from latrine pit to the road, more than one trip necessary on last desludging, water bill greater than 5,000 CFA, and an indicator for the compound has a single latrine pit. A control for the stratification variable: less than half of compound walls in the neighborhood are high, is also included but not shown. All specifications include dummies for the number of desludgings done by the household, with the dummy for no desludgings dropped. Standard errors, clustered by neighborhood, are in parentheses.

Table 11: Market Share Effects Comparison: Treatment versus Counterfactual Subsidies

	Mkt share Mechanical	Pooled Effect		By Price Group	
		Target Prices	CF Subsidies	Target Prices	CF Subsidies
Overall	71.2	0.052*	0.000		
		(0.028)	(0.029)		
Price of 10k group	45.5			0.091*	0.015
				(0.052)	(0.072)
Price of 15k group	77.7			0.016	-0.008
				(0.049)	(0.031)
Price of 17k group	85.7			0.029	0.008
				(0.051)	(0.038)
Price of 20k group	94.2			-0.037	-0.020
				(0.074)	(0.0153)
<i>N</i>		92	92	300	300

Notes: Column (1) gives the market share for each price group of mechanical desludging at baseline. Column (2) provides the OLS estimate of the targeted price pooled effect with observations at the neighborhood cluster level. Column (3) provides the OLS estimate of the counterfactual subsidies pooled effect with observations at the neighborhood cluster level. Column (4) gives the OLS estimate for the market share effect for the targeted price treatment each price group within a neighborhood cluster (not all neighborhood clusters include households from each price group). Column (5) gives the OLS estimate for the market share effect for the counterfactual subsidies for each price group within a neighborhood cluster. Controls are included for the neighborhood means (in columns 2 and 3) or neighborhood-price group means (in columns 4 and 5) of the variables not balanced at baseline at the household level (water bill more than 5,000 CFA, electricity bill, latrine pit distance to road, two tanks used last desludging, months between desludgings, compound only has 1 pit, household requires an extension hose, and the stratification variable—number of households in cluster with low walls). Results are very similar if controls for controls only for variables unbalanced at the household level (available on request).

Table 12: Ex post platform budget estimates

Cost Scenario:	Budget (CFA)	Budget (USD)
Round-by-neighborhood auction price	-98.81	-0.18
Last observed auction price	-98.85	-0.18
Mean negotiation price	1361.19	2.47
Last observed negotiation price	2901.19	5.27