NBER WORKING PAPER SERIES

IMPERFECT COMPETITION AND SANITATION:
EVIDENCE FROM RANDOMIZED AUCTIONS IN SENEGAL

Jean-François Houde
Terence R. Johnson
Molly Lipscomb
Laura A. Schechter

Working Paper 30514
http://www.nber.org/papers/w30514

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
September 2022

We thank Madelyn Bagby, Josh Deutschmann, Jared Gars, Shoshana Griffith, Ahmadou Kandji, Nitin Krishnan, Sarah Nehrling, Cheikh Samb, and Nguyen Vuong for excellent research assistance and the Bill and Melinda Gates Foundation for generous funding for the project. Thanks to Mbaye Mbeguere and the members of the Senegal Office of Sanitation (ONAS) as well as Water and Sanitation for Africa for their collaboration on this project. We thank Radu Ban, Lori Beaman, Robert Clark, Charlie Holt, Seema Jayachandran, Cynthia Kinnan, Robert Porter, Mar Reguant, and seminar participants at Dartmouth, the Econometric Society North American Winter Meetings, NBER Market Design Fall Meetings, Northwestern University, MIT, Princeton University, University of Houston, University of Michigan, University de Montréal, University of Virginia, UW Madison, UW Milwaukee, and the Washington-Area Development Conference for their comments. The views expressed herein are those of the authors and do not necessarily reflect the views of the National Bureau of Economic Research.

Imperfect Competition and Sanitation: Evidence from Randomized Auctions in Senegal
Jean-François Houde, Terence R. Johnson, Molly Lipscomb, and Laura A. Schechter
NBER Working Paper No. 30514
September 2022
JEL No. L12, L41, O55

ABSTRACT

We study the extent to which collusion can explain the under-provision of clean sanitation technologies in developing countries. Using latrine desludging services in Dakar as a case-study, we document that prices are 40% lower in competitive areas than in areas where prices are coordinated by a trade association. We then develop an experimental just-in-time auction platform with random variation in several design features to formally test for collusive conduct and estimate the welfare costs of imperfect competition. Consistent with the collusion hypothesis, we find that bidders systematically avoid competition by placing round bids and refraining from undercutting rivals. We use a K-means clustering algorithm to classify bidders as competitive, collusive, or naïve and simulate counterfactuals in which non-competitive bidders are replaced with competitive bidders. This would significantly increase take-up of the improved sanitation technology, with back-of-the-envelope calculations suggesting improvements in health of a similar magnitude to those from building improved sewerage systems.

Jean-François Houde
Department of Economics
University of Wisconsin-Madison
1180 Observatory Dr
Madison, WI 53706
and NBER
houde@wisc.edu

Terence R. Johnson
University of Virginia
School of Data Science
164C Elson Hall
Charlottesville, VA 22903
trj2j@virginia.edu

Molly Lipscomb
University of Virginia
235 McCormick Ave
Charlottesville, VA 22904
molly.lipscomb@virginia.edu

Laura A. Schechter
Department of Agricultural
and Applied Economics
University of Wisconsin, Madison
427 Lorch St.
334 Taylor Hall
Madison, WI 53706
lschechter@wisc.edu

1 Introduction

Imperfect competition resulting from coordination between firms is common in developing countries and leads to higher prices and fewer transactions. Coordination in pricing and the exertion of market power are important impediments to growth and efficient resource allocation: imperfect competition diverts business from lower-cost firms and allows high-cost firms to remain in the sector (Aghion and Griffith, 2005; Asker et al., 2019). Market power can be particularly strong in the markets of developing countries, where a few firms are highly connected and antitrust enforcement is limited, leading to high mark-ups (Barrett, 1997; Bergquist and Dinerstein, 2020; Ryan, 2020). In the context of products and services that have positive externalities, the welfare costs of market power are even more severe (for example, Barkley (2023)). State procurement can be plagued with inefficiency and high prices (Best et al., 2023). Electronic procurement and intermediation through auctions may facilitate an increase in competition and a decrease in procurement prices, although thusfar such auctions have been found to primarily increase quality rather than reduce prices (Lewis-Faupel et al., 2016).

We analyze the importance of collusion in the market for sanitation services in Dakar, Senegal where the under-provision of clean sanitation technologies increases the risk of health problems. As a consequence of rapid urbanization and under-investment in public infrastructure, most peri-urban areas of Dakar and many other large cities in developing countries are not connected to a sewerage network. Instead, households rely on individual sanitation systems such as septic tanks and unimproved pits. These systems need to be emptied periodically (in Dakar this occurs on average twice per year); a service called *desludging*. Households choose between manual and mechanized desludgings. In a manual desludging, a worker enters the pit and extracts the sludge using shovels and buckets and dumps the sludge in the street. In a mechanized desludging, a trucker pumps the sludge into a tanker truck and takes it out of the neighborhood, usually to one of three treatment centers. Survey evidence suggests that slightly more than half of desludgings in Dakar are performed using the manual option, which creates important environmental and health externalities including increased diarrhea incidence (Deutschmann et al., 2024b). The industrial organization of the market for residential sanitation services limits competition, potentially contributing to the low take-up of mechanized desludging in Dakar. Competition is limited in large part by the existence of a trade association (AAAS) controlling the prices set in the main garages where both member and non-member truckers park and meet residential clients.

Our objective in this paper is to evaluate the existence of collusion in the mechanized desludging market and the impact of collusion on the take-up of mechanized desludging and on health outcomes. To do so we use both non-experimental and experimental variation. The non-experimental variation compares outcomes in areas controlled by the Association with outcomes in the area it does not control. The experimental variation comes from an auction system for procurement of sanitation services which we developed in collaboration with the government of Senegal. The auctions were designed as a "laboratory" to facilitate the analysis of competition and bidding strategies, and to allow for counterfactual policy simulations through randomized variation in invitations and auction formats.

We start by providing non-experimental evidence in favor of the hypothesis that collusion raises prices and limits the supply of mechanized desludging. We examine the difference in prices between neighborhoods controlled by the Association and an adjacent municipality (Rufisque) dominated by unaffiliated companies. Prices in Rufisque are 40% lower than in the rest of the city, even after controlling for household and location characteristics. We document a steep price gradient outside of Rufisque as we move away from the border, suggesting that most consumers in the areas controlled by the Association do not have access to suppliers operating in Rufisque. As a result of this price difference, over 90% of the population in Rufisque uses mechanized desludging compared to 40% in the rest of the city.

We then formally test the hypothesis that providers in areas controlled by the Association behave non-competitively using data from the experimental auction platform. Specifically, our objective is to measure the fraction of suppliers behaving competitively when invited to submit anonymous bids in an auction platform. Together with the government of Senegal, we designed a just-in-time auction platform for mechanized desludging jobs. The goal was to decrease prices and increase households' access to the improved sanitation technology. Households could contact the call center to obtain a quote for a mechanized desludging. Suppliers were invited to bid by text message, and the lowest bid was presented to the client, who could accept or reject it. The design of the platform included a number of experimentally randomized features. Jobs were offered to a randomly selected group of between 8 and 20 potential bidders. The auction format was also randomized. In half of the auctions, the platform used a revisable-bid format which periodically provided invited participants with information about the current lowest standing bid and allowed bidders to revise their bid. The other half of auctions were conducted using a sealed-bid format. The bidders' identities were not revealed to the other bidders in either format.

The experimental analysis proceeds in two steps. First we analyze the bidding behavior of providers randomly invited to participate in over 5,000 procurement auctions, identifying potentially collusive strategies, and using machine learning techniques to classify bidders as competitive, collusive, or naïve. In a second step, we measure the effect of imperfect competition by simulating a counter-factual distribution of offers under the assumption that firms submit bids according to the strategy of competitive bidders.

We identify potentially collusive strategies that are inconsistent with competitive bidding, or, equivalently, strategies that are inconsistent with individual profit maximization (Chassang et al., 2022; Porter and Zona, 1993, 1999). Our identification strategy exploits the fact that competitive bidding strategies differ across revisable and sealed-bid auctions, and that bidders are randomly invited to bid in both formats. We therefore leverage the panel structure of our data to identify bidders who consistently bid competitively across auctions.

The first violation of competitive bidding which we document is the presence of frequent ties in sealed-bid auctions, associated with the use of focal prices. These are round bid values that lead to ties in a significant number of auctions, in which case the job is allocated to the participant who submitted the bid earliest. There is a simple and strictly profitable deviation. Cutting one's bid by a small amount can increase the probability of winning by up to 20%. We interpret a bidder's high propensity to submit round bids in sealed-bid auctions as evidence of his being part of a tacitly collusive agreement to soften price competition.

The second sub-optimal strategy is bidding early in the revisable-bid auction format. In this format, bidders are informed of the current lowest bid every 15 minutes, and can submit a sealed bid in the last 10 minutes of the auction. Assuming private-value costs, bidding before the closed portion of the revisable-bid auction rather than sniping in the closed portion of the auction is a sub-optimal strategy. Waiting reduces the likelihood that the bid is undercut and, to the extent that bidding is costly, firms are better off learning about rival bids before submitting their own bid in this paid-as-bid system. In contrast, submitting an early bid can be viewed as an effort to coordinate prices by sending a signal to rivals.

Our first set of results documents important differences in the distribution and timing of winning bids between the two auction formats. In particular, consistent with the presence of some competitive bids, revisable-bid auctions are significantly less likely than sealed-bid auctions to end with a round winning bid (9 pp) and less likely to end in a tie (9 pp). They are also significantly more likely to attract late winning bids (29 pp). In a significant number of revisable-bid auctions, a seller will wait until the end of the auction period and then undercut his rival in order to win. Despite these differences, not all auctions appear to be competitive. In particular, 24% of sealed-bid auctions receiving more than one bid end in a tie, largely because firms heavily rely on commonly used focal prices (57% of sealed winning bids are on a 5,000 CFA grid). Similarly, a large share of winning bids in the revisable-bid format are placed early (58%).

Next, we exploit the panel dimension of our data and use a clustering algorithm to classify bidders based on patterns in the timing and roundness of their bids (Bonhomme et al., 2022). This process yields three groups of bidders: competitive, collusive, and naïve. Types are identified based on the consistency with which bidders use non-competitive versus competitive bidding strategies across auctions and auction formats. We document a strong positive correlation between bidders' propensity to avoid focal prices (which avoids ties) in sealed-bid auctions and their propensity to submit late bids (undercutting the standing low bid) in revisable-bid auctions. Competitive bidders both avoid ties by relying less on focal prices in sealed-bid auctions and snipe at the last minute in revisable-bid auctions. Collusive bidders consistently avoid competition by submitting round bids in sealed-bid auctions and bidding early in revisable-bid auctions. We estimate that roughly 25% of active bidders are competitive types.

Giving consumers access to competitive quotes would lead to a large increase in the take-up of mechanized desludging. We estimate that changing the composition of invited bidders by swapping in one additional competitive bidder lowers the expected winning bid by 360 CFA (or 1.4% of the mean winning bid). To measure the effect of inviting *only* competitive bidders to the platform, we perform a counter-factual simulation exercise in which we predict the distribution of bids and acceptance decisions under the assumption that all invited bidders use the policy function of the competitive types. We find that doing so would lower the expected winning bid by 3000 CFA and increase the number of mechanized desludgings on the platform by 40%. This gives a lower-bound on the equilibrium effect of inviting more competitive bidders to the auctions. Very few auctions include more than one or two competitive bidders and so under the current random invitation rule competitive bidders do not need to place very low bids in order to win in a majority of auctions (this is the umbrella effect discussed in Caoui (2022)). If the platform targeted invitations to favor competitive bidders, these bidders might start bidding more aggressively.

Finally, we use the counterfactual distribution of offers to measure the potential health benefits of improving competition in the market. We rely on experimental results obtained by Deutschmann et al. (2024b) to estimate that a 1% increase in the number of mechanized desludgings among neighbors would cause a 0.6% decline in the number of diarrhea cases within a household. Assuming that the counter-factual increase in take-up on the platform would lead to an equivalent decrease in manual desludging in the market, we estimate that increasing competition would decrease negative health outcomes by 14 to 23%. A more conservative estimate uses the elasticity of demand for mechanized desludgings compared to manual desludgings from Deutschmann et al. (2024a) (rather than the elasticity of demand for mechanized desludgings on the platform compared to desludgings purchased outside the platform) and yields reduction in diarrhea cases ranging from 9 to 15%. The magnitude of these estimates is comparable to the health benefits of large-scale investments in sanitation infrastructure projects documented in the literature (Alsan and Goldin, 2019; Barreto et al., 2007; Moraes et al., 2003).

This paper builds on an extensive literature testing for collusion. Following the work of Porter and Zona (1993, 1999), we define collusive behavior as a set of actions that violate individual profit maximization. Recent papers using a similar strategy to define collusive behavior include Bernasconi et al. (2023), Chassang et al. (2022, 2023), Clark et al. (2024a), Conley and Decarolis (2016), and Kawai and Nakabayashi (2022). Our paper also relates to an extensive literature using excessive correlation in bids to detect collusive behavior (Abrantes-Metz et al., 2006; Bajari and Ye, 2003; Froeb et al., 1993). There also exists a large literature studying the inner workings of cartels in auction settings (Asker, 2010; Pesendorfer, 2000), as well as more broadly in retail and manufacturing industries (Byrne and de Roos, 2019; Clark et al., 2024b; Clark and Houde, 2013; Genesove and Mullin, 2001; Igami and Sugaya, 2022). Alé-Chilet and Atal (2020) analyze the role of a trade association in coordinating actions across a large number of players.

Our findings also relate to the literature in industrial organization identifying behavioral biases and failures to maximize profits. DellaVigna (2009) provides an early survey of empirical findings (including in auction environments), and DellaVigna and Gentzkow (2019) documents the prevalence of sub-optimal pricing strategies in retailer markets. While we interpret our findings through the lens of tacit collusion, the presence of market frictions (e.g., cash-based transactions) and information frictions (e.g., biased beliefs) can also lead bidders to make strategic 'mistakes.'

The paper additionally contributes to a growing literature studying the effect of competition (or lack thereof) on market outcomes in developing countries. Similar to our setting, Banerjee et al. (2019) and Jack (2013) analyze procurement auction experiments in developing countries. Their findings echo ours: competitive auctions can improve allocative efficiency and lower profit margins. Closely related to our context, Bergquist and Dinerstein (2020) uses a field experiment (in Kenya) to measure agricultural traders' market power, focusing on the the pass-through of cost shocks and subsidies to retail consumers. Additional papers measuring the consequence of imperfect competition in developing countries include: Alé-Chilet (2018) and Alé-Chilet and Atal (2020) (Chilean pharmacies), Barkley (2023) (Mexican pharmacies), Bernasconi et al. (2023) (Colombian electricity distribution), Brown et al. (2024) (organized crime in El Salvador), Neilson (2021) (Chilean elementary schools), Ryan (2021) (Indian electricity markets), and Walsh (2020) (Ghanaian radio broadcast markets). Less work has been done on the impact of market power on sanitation markets, particularly in developing countries, yet the welfare impacts of poor sanitation are large.

The remainder of the paper is organized as follows. In Section 2, we describe the different data sources and provide background information on the mechanized desludging market. Section 3 gives non-experimental evidence of collusion in the traditional mechanized desludging market. Section 4 describes the experimental auction design, presents evidence of collusion using auction and bid outcomes, and identifies firms behaving competitively and those behaving collusively on the platform. Section 5 uses the results from the classification exercise to measure the effect of collusion on prices, take-up of mechanized desludging, consumer surplus, and health. Section 6 concludes.

2 Residential desludging in Dakar

We start by describing households' desludging needs and the options they have for satisfying them in Senegal. We then describe the three main sources of data that we use. We proceed to use that data to describe the structure of the mechanized desludging market and provide a series of stylized facts about supply and demand.

When a household latrine pit fills, the household must empty it by having it desludged. This tends to happen once or twice a year for households in Dakar. Households then have a choice between three types of desludging services: (i) manual performed by a family member, (ii) manual performed by a hired worker (called a "baay pell" in Senegal), and (iii) mechanized using a vacuum truck. Manual desludging consists of removing the sludge using a shovel and placing it in a pit dug in the street near the house.[1] Mechanized desludgings are done by two to three workers with a vacuum truck. The truck pumps as much sludge out of the pit as possible and takes it relatively far from the household, either dumping it legally at a treatment center or illegally in a street drainage canal or the ocean.

The market for mechanized desludging is organized around three treatment centers scattered across the city and a network of garages (or parking lots) where clients meet service providers. Based on our discussions with market participants, prices in the traditional market are determined by bilateral negotiation between clients and truck drivers. Consumers meet providers at garages, and haggle over prices. Walk-in clients at a parking lot are allocated to the driver who is first in line, and drivers from the same garage do not compete over clients (similar to what often happens at taxi stands). The proximity of consumers to garages therefore affects both their ability to search and to negotiate for better prices, as well the cost of providing the service. Furthermore, repeated interactions between truckers within (and potentially across) garages limit the bargaining leverage of consumers, and firms typically initiate the negotiation by relying on "focal" prices that are common across garages.

The rising popularity of cell phones in the years leading to our sample period have somewhat eroded the bargaining power of truckers. Since households need to desludge multiple times, many clients contact truck drivers directly; either by calling their phone number or hailing them on the street. This has allowed independent truckers to compete with companies affiliated with the main garages, and offer better prices. The goal of our empirical analysis is to measure the extent to which truckers are willing to deviate from the cultural norm of using focal prices and offer competitive prices when invited to bid.

2.1 Data sources

The intervention and data collection spanned mid-2012 through mid-2015 in residential neighborhoods surrounding Dakar, Senegal. We collected (i) administrative data from just-in-time mechanized desludging auctions, (ii) a baseline and endline mechanized desludging provider survey, and (iii) four rounds of a household survey of desludging technology choice and price.

Auction platform administrative data: We use panel data from a just-in-time auction platform for the procurement of residential desludging jobs. Together with the non-governmental organization Water and Sanitation for Africa (WSA) and the government's National Office of Sanitation of Senegal (ONAS), we ran a call center which auctioned off residential desludging jobs from July 2013 through September 2015. After that, the call center was scaled up by ONAS and then given

[1] The manual option increases the risk of health-related sanitation problems (for the client, the workers, and other households in the neighborhood). Manual desludging is technically illegal, though it is rarely sanctioned. It is often a source of controversy among neighbors.

to a private sector partner, Delvic. Call-center customers are clustered around the peri-urban areas of Dakar.

As described in more detail in Section 4.1, the auctions involved randomized components which allow us to identify collusion and tease out its effects. The administrative data from these auctions includes the randomized auction format, the randomized number and identities of the desludgers invited, whether or not they bid, the time and amount of their bid if they made one, the location of the household that they were bidding on, and the winning bid.

Survey of mechanized desludging service providers: We conducted a baseline survey of 121 desludging truck operators in mid-2012, and an endline survey of 152 drivers (of which 13 were owner-operators), 75 truck owners, and 20 managers in mid-2015. Our goal was to conduct a census of trucks active in the residential desludging market. An operator is either the manager of a fleet of trucks, or a driver associated with a single license plate. From this data, we are reliably able to identify each truck's main garage.

A difficulty measuring firm attributes in this market is that it is not obvious how to define a 'firm.' Trucks are used by multiple drivers and the arrangements between owners, managers, and drivers are unclear. The different parties seem to have limited information about one another (for example their Association membership or revenue and costs). Because of these issues, there were low response rates and conflicting answers from different respondents affiliated with the same truck to key questions about the market. For example, we asked drivers, managers, and owners about both individual and company-level membership in the Association in the baseline and endline surveys. Over 60% of trucks had some form of conflicted answer about membership. This makes it difficult to use the survey to describe financial arrangements and competition. Instead, we rely on the incentivized administrative bidding data for our analysis of competition.

Household survey: We have an unbalanced panel of 16,255 observations from 9,672 households. To select households, we overlaid grid points on a map of Dakar, excluded any grid points which were in uninhabited areas or served by the city sewer network, and spiraled out from each starting point to select households. Appendix Figure B-1 displays a map with the households' locations. Our analysis focuses on seven of the nineteen arrondissements of Dakar.[2] We drop observations from households that did not receive a desludging over the past year (either mechanized or manual) or with missing responses on key variables. The final sample includes 9,805 observations from 6,121 unique households. We use the household survey to measure the distribution of prices and demand across neighborhoods. An observation corresponds to a household's most recent desludging transaction performed in the 12 months prior to the survey. Since there are no posted prices, we

[2]We chose residential arrondissements and avoided areas connected to the sanitation network and areas that frequently flooded. The arrondissements are subdivided into 43 communes d'arrondissement or CAs (admin3 and 4 on the map). The majority of households surveyed are located in five arrondissements: Pikine Dagoudane (center-west, 14%), Thiaroye (center, 30%), Guédiawaye (center-east, 14%), Niayes (north-east, 37%), and Rufisque (south-east, 5%). Note that households in Rufisque were not sampled in the second and third wave of the surveys, which explains the smaller number of observations (250 unique households).

use transaction prices reported by households for both mechanized and manual desludgings.

2.2 Description of the market for mechanized desludging

Figure 1 shows the locations of the garages and the three treatment centers. On average, consumers are located 4 km away from the closest treatment center, 1.2 km from the nearest garage, and 2 km from the second closest garage. There is wide variation in the size of garages. The largest garage hosts nearly 80 trucks, while some informal garages host only a handful. There is also a group of independent truckers who operate outside of the garage system and are typically contacted by clients on the street or by cellphone.

Much of our analysis focuses on peri-urban residential neighborhoods of Dakar, excluding the arrondissement of Rufisque. In these neighborhoods, relatively few households are connected to the sewage network and take-up of mechanized desludging is low (44%). Trade is influenced by the Association of Desludging Operators (or AAAS), which controls the operation and prices of the larger garages.[3] The official role of the Association is to help operators collaborate on the procurement of truck parts and to assign large lucrative government and commercial contracts. The influence of the Association extends beyond members and affects the provision of residential contracts by all companies throughout the city. This is because AAAS is involved in the largest garages (where non-member trucks also tend to park), and distributes contracts and services to member and non-member companies. The threat of being excluded represents a risk of reduced profits due to the loss of non-residential contracts and more difficult access to truck parts for both members and non-members. Although AAAS has never been prosecuted for price-fixing, the possibility that trade associations can facilitate collusion in large retail markets is well documented in other markets. See Alé-Chilet and Atal (2020) for an empirical analysis in Chile.

The cost of providing a mechanized desludging includes the time and fuel required to complete the job (while pumping sludge at the client's house and driving from the garage to the client to the treatment center and back to the garage again), the treatment center's dumping fee,[4] and in some cases, a referral commission paid to the garage. Mechanized desludging exhibits economies of scale due to truck maintenance and/or rental costs. The large majority of drivers, 92%, are paid a fixed salary, and about half report paying a commission for jobs the garage refers to them. A portion of these revenues is redistributed to company owners in the form of revenue-sharing agreements. For example, desludgers in the largest garages report being paid by their garage on days when they do not find work.

The extent to which firms can benefit from a thicker market or punish one another for deviating

[3]We did not survey in the historical center of the city found towards the west of the map because it is relatively well connected to the sewer network. We did survey in the arrondissement of Rufisque, which is outlined in yellow on the east of the map. Rufisque is not controlled by the Association and has a more competitive mechanized desludging market.

[4]The fee for disposal of the sludge at a treatment center is approximately 3,000 CFA.

Figure 1: Spatial distribution of average prices and demand for mechanized desludging

(a) Average transaction prices

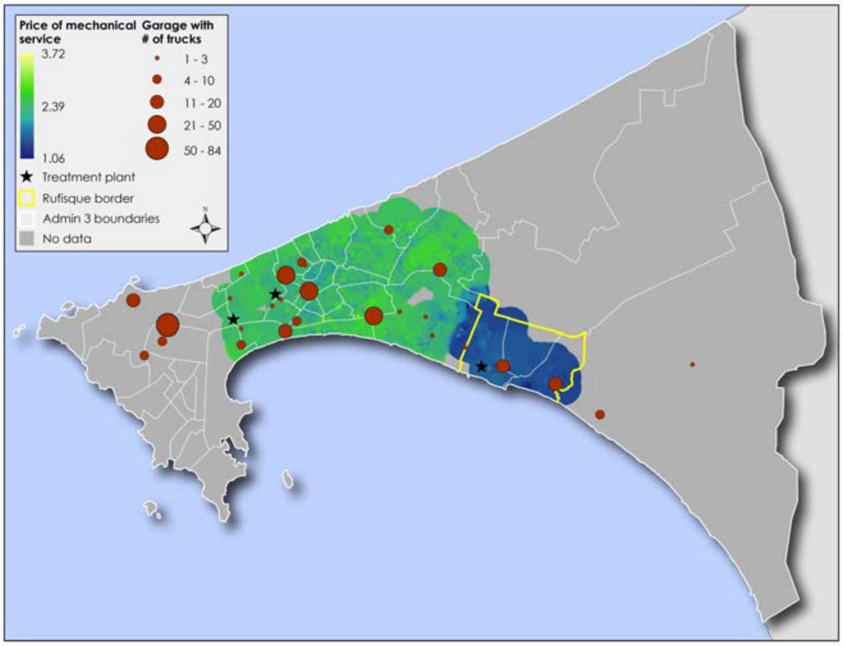

(b) Transaction probabilities

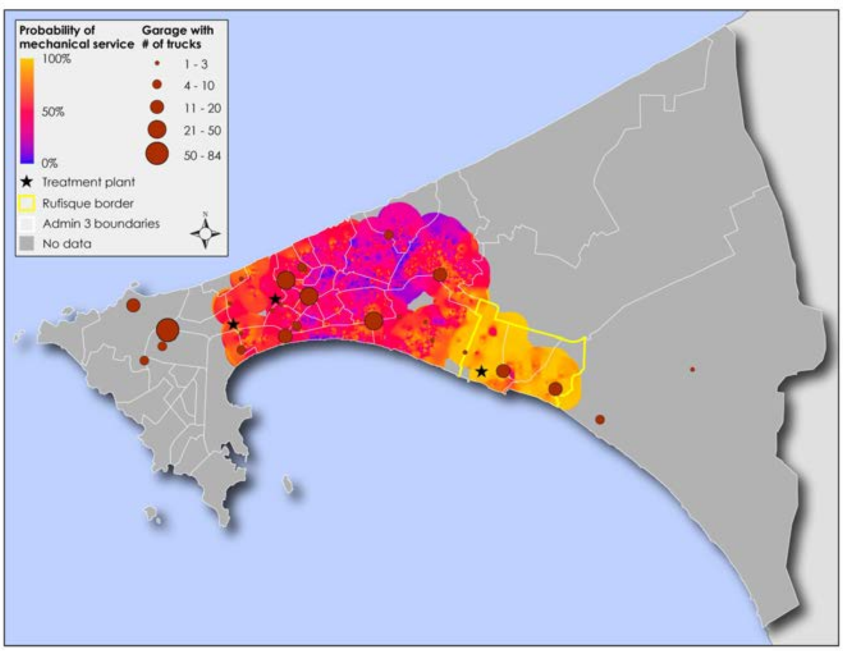

Note: Figure uses household survey data. The distance between the Camberene treatment center (the west-most star) and the Rufisque treatment center (the east-most star) is 22 km. Prices are measured in 10,000 CFA.

from collusive agreements depends on excess capacity in the industry, for which we find ample evidence. To estimate the production capacity of trucks, we ask operators how many desludging trips they performed over the last ten days. On average, trucks perform slightly more than one trip per day (residential and non-residential combined), but there is substantial heterogeneity across trucks. The top ten percent of trucks in terms of number of trips perform more than three jobs per day, and the most active truck performed over 50 trips in a ten-day period. Based on these data and discussions with providers, we estimate that a typical job takes about two hours from start to finish. Most truckers operate with substantial excess capacity, while only a small fraction operate at full capacity. Eighty-five percent of desludging operators in the baseline provider survey stated that they could find more jobs if they wanted to make more money. The evidence above suggests that the mechanized desludging market in Dakar may be collusive.

Since driving the truck and operating the vacuum pump require significant amounts of fuel, the cost of diesel and the efficiency of the truck play an important role in determining trucker costs. The fuel efficiency of the trucks varies substantially with truck size and age, and we estimate that most trucks get between three and six kilometers per liter of diesel. Since diesel prices averaged 750 CFA per liter over the period, the fuel cost per kilometer ranges between 125 CFA and 250 CFA.[5] Conversations with market participants also reveal that a single latrine pit usually fills the truck more than halfway, limiting the ability of drivers to service multiple clients without dumping the sludge at a treatment center in between. About 8% of households require more than one trip.

Table 1: Summary statistics on desludging choices and prices

Variable	**Mean**	**Std. Dev.**	**Min.**	**Max.**	**N**
Mechanized price (x1K)	22.82	7.61	10	50	4865
Baay pell price (x1K)	14.28	6.5	5	40	2339
Family price (x1K)	0.53	2.54	0	25	2601
Choice: Mechanized	0.5	0.5	0	1	9805
Choice: Baay pell	0.24	0.43	0	1	9805
Choice: Family	0.27	0.44	0	1	9805
Truck find: Garage	0.22	0.42	0	1	4865
Truck find: Phone	0.19	0.4	0	1	4865
Truck find: Referral	0.43	0.49	0	1	4865
Truck find: Street	0.1	0.3	0	1	4865
Truck find: Other	0.05	0.23	0	1	4865

Note: Table uses household survey data. Prices and search method to find the truck are conditional on the desludging method chosen. Prices are measured in 1,000 CFA. Mechanized prices are truncated to be between 10 and 50 thousand CFA while baay pell prices are truncated to be between 5 and 40 thousand CFA (the 5th and 99th percentile values).

Table 1 presents summary statistics on transaction price and desludging technology choice from

[5] In order to estimate the amount of diesel necessary per kilometer for a job, we sent an enumerator on ride-alongs with two truck drivers, filling the tank at the beginning and end of the day and recording the kilometers traveled and diesel used.

the household survey. Despite the health hazards associated with manual desludging, the reported market share of the mechanized service is only 50%. This is mostly due to the price difference between the two options. Hiring a baay pell to conduct a manual desludging tends to cost between 12,000 and 16,000 CFA ($24-$32), with an average of 14,300 CFA. Most households that have a family member conduct their manual desludging do not pay anything for the service. In contrast, households pay on average 22,800 CFA (approximately $46) for a mechanized desludging.

Figure 2: Distribution of transaction prices for mechanized and manual desludgings

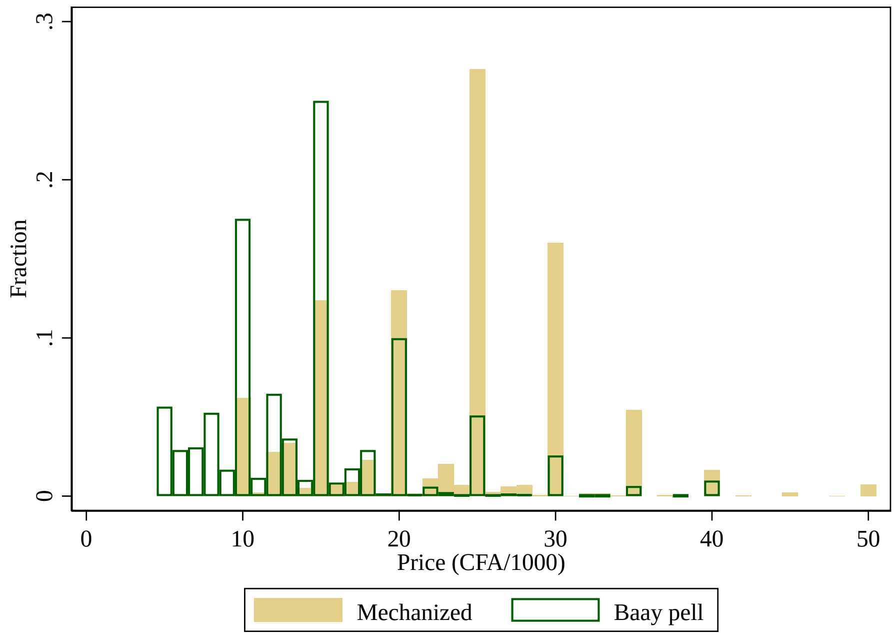

Note: Figure uses household survey data. Prices are conditional on the desludging method chosen. Prices are measured in 1,000 CFA. Mechanized prices are censored to be between 10 and 50 thousand CFA while manual baay pell prices are censored to be between 5 and 40 thousand CFA (the 5th and 99th percentile values).

Figure 2 shows the distribution of transaction prices from the household survey. Most transaction prices for mechanized desludgings are multiples of 5,000 CFA. Almost 30% of mechanized desludgings cost 25,000 CFA, with 56% of transactions costing 20,000, 25,000, or 30,000 CFA. We observe similar coarseness in the distribution of prices for manual desludging. While this may be partly due to the fact that firms and clients find it more efficient to use round numbers when performing cash-based transactions (Beaman et al., 2014), the fact that a large number of transactions have prices that are not multiples of 5,000 CFA suggests that this does not represent a hard constraint.

3 Non-experimental evidence of collusion

As described above, the mechanized desludging market shows characteristics which would be consistent with collusion. In this section, we provide more specific non-experimental evidence consistent with the presence of collusion. In particular, we characterize the distribution of prices and demand in the areas of Dakar that are controlled by the mechanized desludging association, AAAS. These neighborhoods correspond to areas in which collusion is more likely to be prevalent, as opposed to neighborhoods in Rufisque in which firms operate independently. We use this analysis to provide preliminary evidence on the importance of collusion in the market. In Section 4, we continue on to show experimental evidence of collusion.

Prices are much lower in the arrondissement of Rufisque on the eastern outskirts of Dakar which the Association does not control. In the 1990s a single company, UPAMA, provided desludging services in the independent municipality of Rufisque. UPAMA's main line of business was desludging for fish product processing companies, and it was asked by the city administration to provide affordable residential desludging. Over time, new companies entered the market to serve growing demand, but the new companies matched UPAMA's base price in order to get business. The proximity of the Rufisque treatment plant helps reduce the variable cost of the service relative to other areas of Dakar. UPAMA receives no direct subsidies, and we believe the price reflects the cost of providing a mechanized desludging in the area.

Figure 1 plots the distribution of prices and demand across the neighborhoods of Rufisque and the rest of the city. The official boundaries of Rufisque are highlighted in yellow. Panel A shows that the median price for a mechanized desludging in Rufisque is roughly 15,000 CFA, compared to 25,000 CFA in the rest of the city. This means that the price of a mechanized desludging in Rufisque is roughly equivalent to the price of a manual desludging. As a result, as shown in Panel B, nearly all households in Rufisque choose the mechanized option, compared to roughly 40% in the rest of the city.

Appendix Table A-1 formally tests for differences in observable characteristics of households in Rufisque versus the rest of Dakar. Households in Rufisque tend to be closer to both a treatment center and a garage, leading companies operating in Rufisque to have lower variable costs. Home ownership is higher in Rufisque than in the rest of Dakar, reflecting the fact that many people move out to Rufisque in the periphery of Dakar in order to be able to build their own homes.

To account for differences between households in Rufisque and other parts of Dakar, we estimate the following regression relating mechanized desludging transaction prices p for household i in month t and household characteristics (including distance to the Rufisque boundary). The household characteristics x_{it} we control for include distance to the nearest treatment center, distance to the nearest mechanized trucker garage, the number of trucks within three kilometers, household size, the number of earners living in the same household, the number of other households living in the same house and sharing the same pit, the number of rooms in the house, a house ownership indicator, a two-story house indicator, a household wealth index, and a wide road indicator.

$$p_{it} = g^k\left(\text{Distance to Rufisque}_i\right) + x_{it}\beta^k + \epsilon_{it} \quad k = \text{Rufisque, Other.} \tag{1}$$

We estimate this regression separately for the two regions to allow for differences in pricing strategies, and approximate $g^k(\cdot)$ using a step function of distance with 500 meter increments. In addition to household characteristics, the regressions also control for month-year fixed effects and standard errors are clustered at the household level. Because Rufisque is located in the southeast portion of the city, we measure the distance to Rufisque using distance to the nearest western or northern Rufisque boundary.

Figure 3: Predicted mechanized desludging prices and adoption as a function of distance to Rufisque

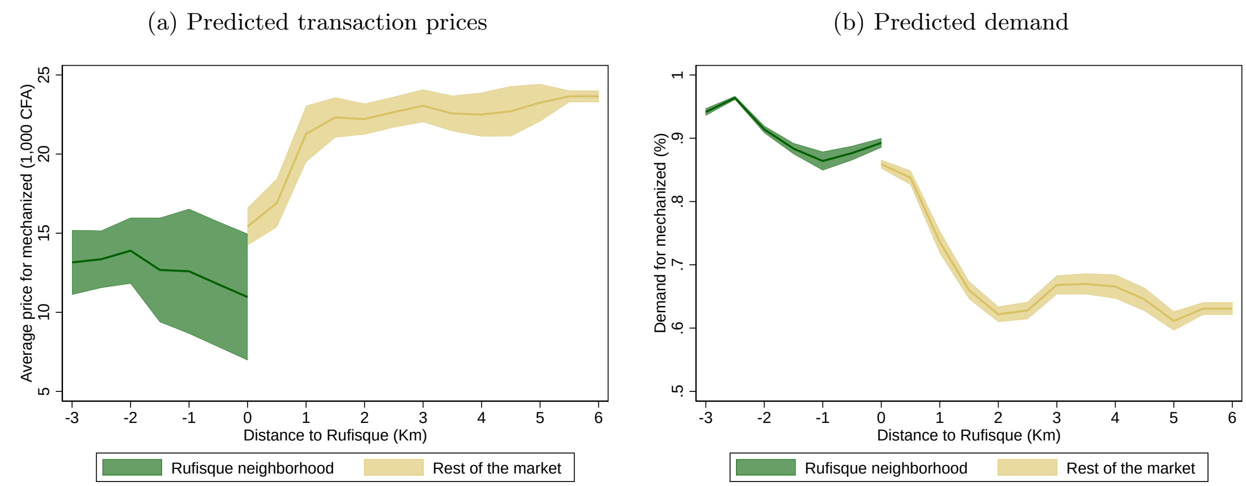

Note: Predicted prices from equation (1) are evaluated at the average characteristics of households in Rufisque. Demand is predicted using the same equation and a linear-probability model. Shading indicates 95% confidence intervals.

Figure 3a presents the predicted values from these regressions by distance to the Rufisque boundary.[6] The outcomes are predicted using the average characteristics of households living in Rufisque to eliminate any compositional differences. The green line presents the predicted prices in Rufisque, and the orange line presents the predicted prices in the rest of the market. The shaded area represents the 95% confidence interval.

This analysis confirms that lower prices in Rufisque are not due to observed differences in household or location characteristics. In the rest of the market, which is controlled by AAAS, mechanized desludging prices increase rapidly with distance from the Rufisque border, but predicted prices are flat with respect to distance from the border in Rufisque. Households in Dakar living within 500 meters of the official Rufisque boundary pay nearly the same price as their neighbors in Rufisque; roughly 15,000 CFA. The gap widens significantly as we move more than one kilometer away from

[6] Appendix Table A-2 presents the coefficients estimated by the regression.

the boundary. Low prices in Dakar near the border of Rufisque may be due to competition spilling over the boundary, or may be because the official boundary differs from the de facto boundary. Households in Dakar more than 1.5 kilometer from the Rufisque border pay close to 25,000 CFA, and the price schedule is independent of distance to Rufisque at further distances. In contrast, prices for manual desludgings (results not shown here) do not differ significantly across the region, although so few residents of Rufisque get a manual desludging that we have very few observations of manual prices within Rufisque.

Figure 3b repeats the same exercise but for demand. In particular, we re-estimate equation (1) where the outcome is an indicator for whether the most recent desludging transaction was mechanized. The figure plots predicted take-up rate for observationally equivalent households living on either side of the Rufisque boundary. Results suggest that mechanized and manual desludging are close substitutes, and as a result demand for mechanized desludging is very elastic. The mean difference in prices between the two regions is roughly 10,000 CFA, while the predicted take-up rate is 30 percentage points higher in Rufisque, leading to an elasticity estimate of roughly -3.[7]

Figure 4 illustrates another important difference between the two areas: price dispersion. Since mechanized desludging prices in Rufisque mostly reflect mechanized desludging costs, we observe very limited dispersion in transaction prices across households. In contrast, prices are very dispersed in the areas controlled by the Association in the rest of Dakar. Roughly 25% of households in the rest of Dakar pay mechanized prices comparable to those paid in Rufisque. The remaining 75% of households pay higher prices, and a sizable fraction pay more than double the average Rufisque price.

Importantly, Figure 4b shows that prices are dispersed across consumers even within narrowly defined neighborhoods. The figure plots the distribution of price residuals obtained by taking the difference between transaction prices and the corresponding neighborhood averages (27 neighborhoods). The inter-quartile range (IQR) of this residual price is 1,800 CFA in Rufisque, compared to 8,600 CFA in the rest of the market. This suggests that the main component of desludger costs, household location, only explains a small fraction of the observed dispersion in prices paid by consumers in the areas controlled by the Association. This is consistent with the presence of imperfect competition in the market, potentially due to price discrimination and search frictions.

In summary, this non-experimental analysis suggests that there are differences in competitive conduct between areas controlled by the Association and the Rufisque neighborhood. Assuming that unobserved cost differences are continuously distributed around the Rufisque boundary, the results establish that the average household in the Association-controlled neighborhoods pays a significantly higher markup for mechanized desludgings than the average household in Rufisque, and as a result is less likely to use the mechanized desludging option. This is consistent with the hypothesis that the Association is successful at restricting supply and maintaining high prices in

[7] We obtain a similar elasticity measure using the distance gradient of the predicted demand and price curves from 0 to 2 km.

Figure 4: Distribution of mechanized prices in Rufisque and in the rest of the market

(a) Price levels (b) Price residuals

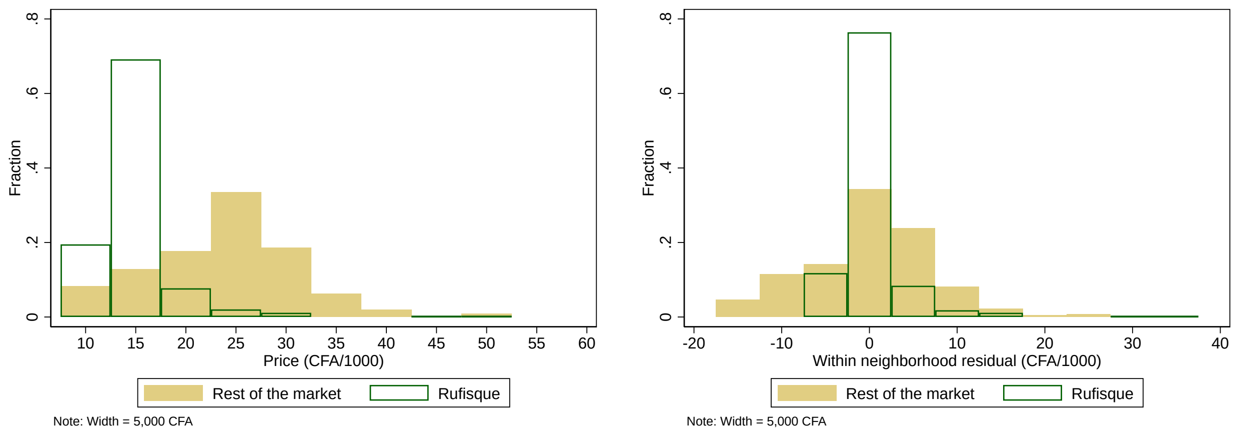

Note: Figure uses household survey data. The price residuals in panel (b) are computed as the difference between each transaction price and the average transaction price in that neighborhood (taking into account 27 neighborhoods).

most areas of Dakar. In the next section, we test for collusion using data from the experimental auction platform.

4 Experimental design and evidence of collusion

We now turn to the evidence for collusion derived from the exogenous variation in the randomized auction platform. Our objective is to measure the fraction of suppliers behaving competitively when invited to submit anonymous bids in an auction platform. In collaboration with the government of Senegal, we designed a centralized auction platform which allows us to analyze bidding behavior in a controlled environment. We randomized both the auction format as well as how many and which bidders were invited to each auction.[8] By repeatedly observing desludger participation and bidding behavior under randomly selected auction formats, we can measure firms' propensity to behave competitively.

In this section, we start by describing the auction platform and the sources of random variation. We then describe our identification strategy for auction-level and bidder-level analysis. We show results testing for collusion (or behavior inconsistent with competitive bidding) using auction-level outcomes. We then take advantage of the panel data on bidder behavior to first show that there are consistent differences in behavior across bidders and then use K-means clustering to distinguish competitive bidders from collusive bidders (and from naïve bidders). Using this categorization, we measure the effect of competitive bidders on prices. This leads to the analysis in Section 5 in which we conduct a counterfactual analysis to estimate the harm that collusion causes in decreasing

[8]The call center was advertised across Dakar in television ads, radio ads, and billboards. Consumers endogenously chose to call, making them a selected subset of the population.

uptake of mechanized desludging and on to poor health outcomes.

4.1 Description of the auction platform

The design of the platform is simple in order to encourage participation by actors on both sides of the market: households who need the service only sporadically, and desludging operators who are busy and may not have time to engage in or master complicated bidding processes. The auction platform is entirely phone based. When a household needs a desludging service, it calls the center and gives the dispatcher basic information about the location of the pit to be desludged. The call center dispatcher solicits bids via text message, and desludging operators have one hour to respond by text message. The bidder with the lowest price – after accounting for any penalties associated with past service problems – wins the auction.[9] In the case of a tie, the earliest bid at the winning price is awarded the job. The call center dispatcher makes the winning offer to the client, and if the client accepts, the client's phone number is sent to the winning bidder by text message. The client and the operator are then free to make logistical arrangements on their own, and the household pays the desludging operator directly.[10]

The final winning bid, but not the identity of the winner, is sent to all desludgers invited to participate in the auction. We reveal the winning price for two reasons. First, the Association and participating companies requested information about which bids win, which might reflect a desire for transparency and fairness, the ability to monitor and maintain cartel discipline, or feedback about which bids win in which locations. Second, we wanted the participants to learn which bids are competitive, and that would be much more difficult without clear and ongoing feedback about the winning bids.

The platform randomizes two components of the auctions. First, between 5 and 21 of the 126 registered desludgers are randomly invited to bid in each auction.[11] Invited bidders are informed of the number of bidders invited, but not their identities. Second, the auction format is randomized between *sealed-bid* and *revisable-bid* formats. In sealed-bid auctions, the bidders have one hour to submit their bid and receive no information about other bids until the winning bid is announced at the auction's conclusion. In the revisable-bid format, bidders are given updates about the standing low bid every 15 minutes and again 10 minutes before the auction closes, and are allowed to submit

[9] In our analysis we abstract away from penalties for two reasons. First, penalties were only applied in the first version of the platform and were discontinued in the later auctions. Second, the platform design made it difficult for bidders to know whether a penalty was applied to their competitors. Thus, it is reasonable for us to assume that firms believe the lowest bid wins.

[10] Some customers are surveyed by phone after the desludging takes place. The survey asks about the quality of the service and if they were charged the right amount. Desludging operators know they will be penalized in future auctions if they charge more than the agreed price or provide low quality service. The phone surveys confirm that desludgers do not try to renegotiate prices.

[11] Invitation probabilities are independent of distance to the household. In the first version of the platform, invitation probabilities differed across bidders as a piece-wise linear function of the number of valid bids submitted in the prior months. This probability was truncated at the bottom and top to ensure the invitation probability was bounded away from zero, and was less than 50%. In the later auctions, desludger invitations were unconditionally random.

revised bids at any time. In both formats, desludging operators receive reminder messages that bids are still being accepted every 15 minutes and again 10 minutes before the auction closes.

We use this experimental variation in two ways. First, as we discuss below, we exploit theoretical differences in competitive bidding strategies across formats to test for competitive conduct. Because firms and consumers are randomly assigned to auction formats, we attribute differences in bidding behavior across formats to differences in firms' strategies. We find that different bidders consistently engage in different strategies, and from this behavior we categorize them as competitive, collusive, or naïve. Second, because invitation lists are randomized and anonymous, firms cannot coordinate their behavior within auctions, and are unlikely to face the same set of rivals in future auctions. We use the auction platform as a "laboratory" to gain insights about how firms compete in the traditional market.

We ran over 5,000 auctions for mechanized desludging services through the call center in collaboration with the Senegalese Office of Sanitation (ONAS) from July 2013 through October 28th, 2016.[12] Some auctions were performed for subsidized households (671 auctions), which affected the price paid by households, though not the price received by desludgers, and so these auctions are excluded from regressions studying household acceptance.[13] The sample sizes vary across auction-level regressions as follows. There are 5,162 auctions total; 4,177 of those have at least one bid; 2,865 of those have at least two bids; 4,491 are unsubsidized auctions; and 3,621 are unsubsidized auctions with at least one bid. The research team designed and managed the platform from its inception in 2013. In January 2015, the management of the platform was transferred to ONAS. Penalties were abolished and the invitation rule was modified so that all truckers had the same probability of being invited rather than over-sampling active participants. Also, the number of truckers invited decreased on average from 14 to 11. Although the changes to the invitation rule affected the performance of the platform (by reducing competition), the assignment of bidders to formats and auctions remained random. Slightly more than half of auctions were performed prior to the design change (55%).

Table 2 presents summary statistics regarding the auctions and bidders. Around half of the auctions are randomized to be revisable format and half are sealed-bid. An average of 13 bidders are invited to each auction, and bidders were invited to bid in 524 auctions on average. The participation rate is fairly low. The probability of submitting a bid is about 10%, which leads to an average of 2.4 valid bids per auction, conditional on there being at least one bid. This low participation rate is due to the fact that a majority of desludgers rarely or ever bid, and the invitations were not targeted based on bidder availability or distance to client. As we discuss below, 52 bidders submit more than 20 bids during our sample periods and much of our bidder-

[12]Before the auctions began, the project held multiple training sessions for the truckers to help them understand the auction process and teach them to bid using the SMS message system. During the auction roll out, operators were available to take calls from the truckers at all times if they had trouble placing their bid.

[13]The subsidy intervention was conducted on a small subset of the Dakar population and is described and analyzed in Lipscomb and Schechter (2018) and Deutschmann et al. (2024b).

Table 2: Summary statistics from the auction platform

VARIABLES	(1) N	(2) mean	(3) sd	(4) min	(5) max
Revisable auction indicator	5,162	0.497	0.500	0	1
Number of invited bidders per auction	5,162	12.79	2.672	5	21
New platform indicator (2015-2016)	5,162	0.454	0.498	0	1
Zero bids indicator	5,162	0.191	0.393	0	1
Number of valid bids (at least one bid)	4,177	2.437	1.436	1	11
Number of auction invitations per bidder	126	524.0	354.5	15	1,693
Bidder participation probability	126	0.104	0.136	0	0.569
Number of auction invitations per active bidder	39	886.4	286.6	414	1,693
Active bidder participation probability	39	0.246	0.134	0.0691	0.535

Note: Table uses auction platform data at the auction level, the bidder level, and the active bidder level. Active bidders are those who submit more than 30 valid bids. Table shows random variation in auction format (revisable vs sealed-bid), and the number of invited bidders per auction. It shows an indicator for whether the auction had zero bids, and the number of valid bids conditional on there being at least one. It shows the number of auction invitations and participation probability for all bidders and then again for active bidders.

level analysis focuses on the 39 most active bidders submitting more than 30 bids. Active bidders submit bids 24% of the time.

4.2 Hypothesis and identification strategy

We construct an empirical test of competitive bidding in the auction platform, measuring the prevalence of imperfect competition by exploiting the random assignment of bidders to auctions and auction formats. We interpret a rejection of the null hypothesis of competitive bidding as evidence of tacit collusion. A bid is deemed 'competitive' if it is consistent with individual profit maximization. Conversely, bidders who systematically avoid these behaviors are deemed 'collusive,' since their actions are at odds with individual profit maximization (Chassang et al., 2023; Porter and Zona, 1993, 1999). Economic models of tacit collusion can be used to determine the factors that facilitate price coordination, but are typically silent regarding the particular collusive strategy that is selected by firms. In Appendix C, we present a model which shows that round bidding can be an optimal strategy to maintain tacit collusion in the presence of unobservable heterogeneous costs.

We assume that (a) bidders have rational expectations about the distribution of rival bids, (b) bidders observe independent and private signals of the cost of providing the service, and (c) the underlying cost distribution is continuous, smooth, and does not exhibit any mass points. These assumptions are reasonable in our context. The rational expectations assumption is justified by the fact that bidders are frequently invited to bid and receive information about the winning bids in all invited auctions, whether or not they participated. In addition, when invited to bid in a revisable-bid auction, bidders are informed about the standing low bid at minute 50, which provides useful information on the distribution of the "bid to beat" and the probability of facing sincere competition. Since bidders also compete in the traditional desludging market, they are well informed about the distribution of transaction prices in each of the neighborhoods as a benchmark for the auctions. As discussed in Section 2.2, the marginal cost of desludging is determined by the distance between the garage and the house and the treatment center, the age and size of the truck, and capacity utilization. It is therefore unlikely that the cost distribution exhibits mass points.

As a first source of evidence about collusive behavior, we investigate the presence of round bids leading to ties due to identical lowest bids in the sealed-bid auctions. Excessive correlation in bids is a common red flag used by antitrust authorities.[14] Identical bids can reflect a tacit agreement between firms to use focal prices, softening competition by allocating the job to participants who submit early bids.

Round bids at first appear to be a sub-optimal strategy, since the use of focal prices creates mass-points in the distribution of winning bids, and a bidder can do strictly better by bidding slightly below these focal prices, increasing their probability of winning the auction substantially for a very small reduction in the price they will be paid. Our model, elaborated in Appendix C, shows that round bidding is an optimal solution even in the static cartel profit-maximization problem. It also shows how round bidding aids monitoring and enforcement of a dynamic collusive agreement. The intuition comes in two steps. First, it is profit-maximizing for the cartel to restrict bidding and soften price competition to boost the expected payoffs to winning firms, even for a single auction. Pooling equilibria due to round bidding can raise profits by restricting truckers from competing away their informational rents. Second, the discreteness of round bidding makes it possible to monitor and enforce an agreement. Deviations from a round-bidding strategy are detected immediately and can be punished. An additional benefit of a round-bidding strategy is that contemporaneous communication and coordination is not necessary. This provides a rationale for why round bidding not only arises empirically, but how it plays a role in solving the complex monitoring and enforcement problems that might otherwise defeat the Association's goals of maximizing truckers' profits through the auctions. Many of the same tensions exist in the traditional (decentralized) market, explaining why we observe similar levels of price coarseness in the auction platform and in the traditional market.

Our second focus is on the timing of bidding. In particular, bidding after the final update of the standing low bid in the revisable-bid auction is profitable for competitive firms. It allows bidders to jump in at the end and "snipe" the standing low bid by slightly undercutting it (Bajari and Hortaçsu, 2003; Roth and Ockenfels, 2002). If bidding is costly (in our case the cost of sending a text message), competitive firms are better off learning about rivals' bids before bidding, rather than

[14] Antitrust laws are not well enforced in Senegal, and the fear of being detected does not play an important role. Mund (1960) and Comanor and Schankerman (1976) provide early analyses of identical bids used by cartels, and McAfee and McMillan (1992) provides a theoretical discussion of the efficiency of this type of strategy.

submitting multiple bids over the course of the auction. In contrast, when trying to collude, firms benefit from bidding early in two ways. First, in the case of collusion, bidding early increases the information provided to rivals, and therefore reduces the likelihood that the bid will be undercut. Second, for both competitive and collusive firms, in the case of a tie the bidder who bid that amount first wins the job, so that early bidding confers an advantage if a tie occurs.[15]

Consistent with the idea that it is optimal for a competitive bidder to bid in the final ten minutes of the revisable-bid auctions, desludger Cheikh Gueye explains how he learned to bid in his 2018 *Planet Money* Poop Cartel interview (Planet Money, 2018): "Generally, what I do is I wait until there are only ten minutes left. If no one takes the offer, then I propose a price. And then immediately, I go so that I have this market." The interviewer then asks: "You said earlier that the truckers were united. Did the text messages – did the auction make you less united because you were competing with each other for price?" And Cheikh Gueye responds, "Even though we used to be united – but now it's a competition. And you need to work hard in order to get something in your business."

We test for imperfect competition by identifying behaviors that fail to maximize individual expected profit: bidding of focal prices (5,000 CFA increments), early bidding in revisable-bid auctions, and bidding prior to the final update about the standing low bid in the revisable-bid auctions. Of course, deviation from profit maximization can arise for reasons other than collusion. For example, Hortaçsu et al. (2019) uses bounded rationality to explain the presence of non-serious bids. Another explanation, at least for the focal prices, is pricing frictions that restrict the ability of firms to select new prices or revise their previous bid choice to maximize profits (Beaman and Magruder, 2012). For example, cash-based transactions limit the ability of bidders to use a very fine price grid. However, we observe a non-trivial fraction of bids that deviate from the most common focal prices and reduce the probability of ties, for example using 1,000 CFA increments instead of 5,000. This is true both in the auctions and in the prices that consumers report paying in the traditional market. Finally, firms may submit non-serious bids in order to be invited more often and/or receive other government assistance. This is unlikely to be an important margin since the average participation probability is 15%, and the most active bidder on the platform participated in 60% of auctions to which he was invited. We also cannot rule out the possibility that some bidders have biased beliefs about the distribution of bids of their rivals, for example, due to inattention. However, this concern is alleviated by the fact that our bidder-level analysis focuses only on the most active bidders who had ample opportunities to learn the distribution of winning bids.

[15]One potential downside of bidding late is that it may increase the probability that a bid is rejected by the platform due to technical delays. In our platform however, bidders have ten minutes to submit a bid after the last message, and the platform gives an additional five minute grace period to ensure that all bids are received. During this period, the auction turns into a first-price sealed-bid auction with a reserve price defined by the lowest bid received prior to minute 50.

4.3 Empirical analysis

We compare auction-level outcomes across the randomized revisable versus sealed-bid auction formats in Section 4.3.1. This tells us whether one format leads to lower prices than the other. More importantly for our purposes, differences in outcomes such as round bids or timing of bids across the two formats allow us to identify trends which would be inconsistent with competitive bidding. We then leverage the panel dimension of the data to analyze heterogeneity in competitive conduct at the bidder level in Section 4.3.2. Because we see the same bidders invited to a large number of auctions of both formats, we can use a K-means clustering algorithm to categorize bidders into categories based on their behavior. This leads us to categorize each bidder as competitive, collusive, or naïve. Then we measure the impact of changing the randomized composition of invited bidders on auction outcomes such as prices, which are especially important given the sensitivity of consumers to price and the negative health effects of manual desludgings.

4.3.1 Auction-format effects

We start by analyzing the distribution of winning bids and their arrival times for revisable and sealed-bid auctions. Figure 5a plots the histogram of winning bid amounts. The distribution in the sealed-bid format exhibits clear mass points at common focal prices of 20, 25, 30, and 35 thousand CFA. The distribution of winning bids in the auctions closely resembles the distribution of negotiated prices in the traditional market in areas of Dakar outside of Rufisque displayed in Figure 4. The same mass points are present in the revisable-bid format sample, but there are clear differences. Winning bidders in the revisable-bid format are more likely to undercut those focal prices by one or two thousand CFA, which leads to a higher density at 23, 24, and 29 thousand CFA. This implies that cash-transaction frictions cannot fully explain the use of focal prices in both the traditional market and in the sealed-bid auctions.

Figure 5b shows that in sealed-bid auctions, roughly 30% of winning bids are placed in the first 5 minutes. Since the tie-breaking rule favors early bidders, bidding early is an optimal strategy for all bidders irrespective of their propensity to collude. In the revisable-bid auctions, the share of early winning bids is much smaller, and the modal winning bid is placed after the last message. Between these two extremes, the distribution of bid times reflects the nudges created by the messages. The difference in behavior across the two formats is consistent with the idea that it is optimal for competitive bidders to submit early bids in sealed-bid auctions and late bids in revisable-bid auctions. However, the fact that a significant fraction of winning bids in revisable-bid auctions arrive early suggests that not every bidder behaves competitively.

Since the format is randomly assigned to each call, we can use a simple auction-level treatment effect regression to summarize the difference in outcomes across revisable-bid and sealed-bid auctions. We estimate

$$y_t = \alpha \, \text{Revisable}_t + x_t \beta + \epsilon_t \tag{2}$$

Figure 5: Distribution of winning bids and arrival times across the two formats

(a) Distribution of winning bids (b) Distribution of winning bid time

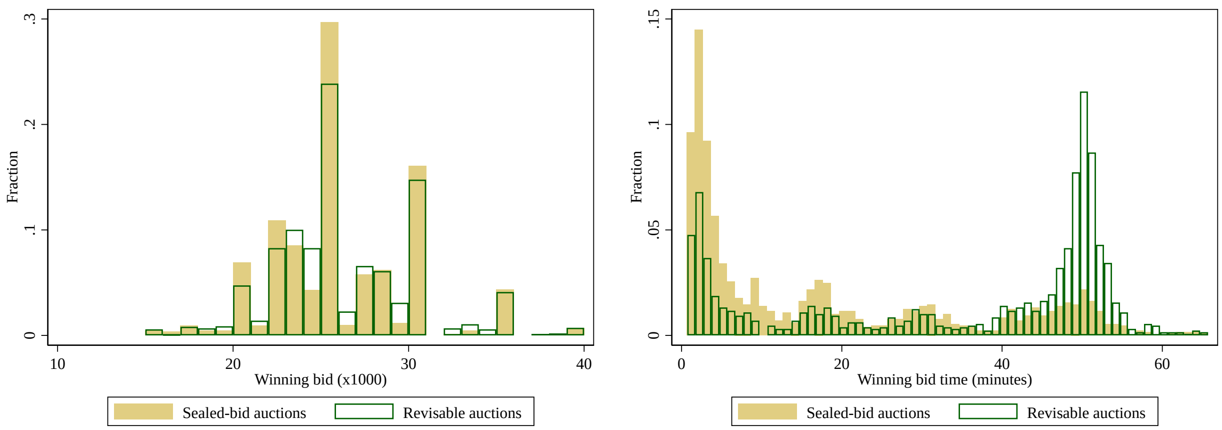

where y_t measures one of five different outcomes for auction t: (i) the winning bid amount, (ii) an indicator for the winning bid being tied, (iii) an indicator for the winning bid being divisible by 5,000, (iv) an indicator for the winning bid occurring in the last time interval, and (v) the minute of the first bid. Appendix Table A-3 presents the mean and standard deviations for the main outcome and control variables used in the analysis.

Table 3 reports the results of equation (2). Panel A is our main specification, while Panel B presents the results after controlling for auction and client characteristics. The first column shows that revisable-bid auctions lead to winning bids that are almost 300 CFA larger than sealed-bid auctions. It is interesting to learn about the effect of the format on average winning bids, although from a theoretical perspective there is no reason to believe that the two formats should be revenue equivalent under either collusion or competition. In general, collusion is thought to be easier to sustain in open auction environments (without a hard close). See Athey et al. (2011), Graham and Marshall (1987), Marshall and Marx (2007), and Robinson (1985) for theoretical and empirical analyses. Our context differs in that the revisable-bid auction format has a “hard close” and bids submitted in the last 10 minutes are not observed by rivals. The revisable-bid auction is best described as a sequential auction: open followed by closed.

The next two columns analyze the prevalence of ties and round bids. The probability that the winning bidder ties is 9 percentage points higher in the sealed-bid auction. As column (3) illustrates, this is explained by the fact that firms are significantly more likely to use bids which are divisible by 5,000 in the sealed-bid format. This is consistent with Figure 5a above. By revealing the current lowest bid at the 50th minute, the revisable-bid auction allows competitive bidders to undercut the standing low bid as of minute 50, and win the auction more often. Note, however, that the fraction of round bids and ties does not go to zero in the revisable-bid format. On average,

Table 3: Experimental treatment effect of auction format on bidding strategies

VARIABLES	(1) Winning bid	(2) 1(Ties)	(3) 1(Round)	(4) 1(Last message)	(5) First bid (min.)
Panel A: Without control variables					
1(Revisable)	0.29**	-0.090***	-0.094***	0.29***	3.29***
	(0.12)	(0.015)	(0.015)	(0.013)	(0.59)
Observations	4,177	2,865	4,177	4,177	4,177
R-squared	0.001	0.013	0.009	0.101	0.007
Panel B: With control variables					
1(Revisable)	0.24**	-0.086***	-0.092***	0.28***	3.43***
	(0.10)	(0.015)	(0.015)	(0.013)	(0.53)
Observations	4,177	2,865	4,177	4,177	4,177
R-squared	0.310	0.034	0.066	0.133	0.205
Unit of observation	Auctions	Auctions	Auctions	Auctions	Auctions
Mean dep. variable	25.8	0.19	0.53	0.28	17.3

Note: The dependent variables are (1) the value of the winning bid (in 1,000 CFA), (2) an indicator for the winning bid being tied, (3) an indicator for the winning bid being a multiple of 5,000, (4) an indicator for the winning bid coming in the last ten minutes, and (5) the minute the first bid came in. The sample includes auctions with at least one valid bid in all columns, and with at least two valid bids in column (2). Additional controls in Panel B include: log number of invited bidders (by itself and interacted with new platform dummy), log distance from the client to the nearest treatment center, log average distance from the client to all potential bidders' garages, client latitude and longitude, log population of the commune, auction number trend (linear, quadratic, and cubic), new platform indicator, indicators for morning and lunch time, and arrondissement, year-quarter, and day-of-week fixed effects. Robust standard errors in parentheses. *: $p < 0.10$, **: $p < 0.05$, ***: $p < 0.01$.

14% of revisable-bid auctions end in a tie, compared to 23% of sealed-bid auctions.

Columns (4) and (5) analyze the timing of bids. As Figure 5b suggests, the winning bid is 29 percentage points more likely to be placed in the last ten minutes of the auction in the revisable-bid format. Similarly the first bid is received three minutes later in the revisable-bid format (minute 19 versus minute 16). The first bid's early arrival in both formats explains the bimodal distribution of winning bid times. Most bidders in the sealed-bid auctions submit bids immediately after receiving the invitation, while a large fraction of winning bidders in the revisable-bid format submit late bids (roughly 42%).

This analysis confirms the presence of different forms of non-competitive behavior in the two formats; most notably the prevalence of ties in sealed-bid auctions, and the large share of winning bids that are received early instead of late in revisable-bid auctions. If all bidders were behaving competitively by maximizing expected profits, bids in the revisable-bid auction would converge to a single mass after the 50th minute, and ties would be very infrequent in both formats.

Under the tacit collusion interpretation for this result, firms avoid competing by coordinating on focal prices. In revisable-bid auctions, collusive firms can facilitate this coordination by submitting an early bid, which can serve as a reference price for other bidders. An alternative interpretation of early bids is that they are submitted by bidders trying to signal strength by submitting an aggressive first offer, hoping to discourage other bidders from competing. Under some distribution of beliefs, this could in principle rationalize early bidding as a competitive strategy.

Table 4: Relationship between the presence of early bidding and bid amounts

VARIABLES	(1) First bid	(2) First bid	(3) Winning bid	(4) Winning bid
1(Early bid received)			-1.02***	-0.32**
			(0.17)	(0.15)
1(Revisable)	0.88***	0.85***	-0.16	0.050
	(0.13)	(0.12)	(0.16)	(0.14)
1(Revisable)x1(Early bid received)			0.89***	0.53**
			(0.26)	(0.21)
Observations	4,177	4,177	4,177	4,177
R-squared	0.011	0.238	0.009	0.325
Unit of observation	Auctions	Auctions	Auctions	Auctions
Control variables?	No	Yes	No	Yes
Mean dependent variable	27.7	27.7	25.8	25.8

Note: The dependent variable in columns (1) and (2) is the value of the first bid received (in 1,000 CFA), and the dependent variable in columns (3) and (4) is the value of the winning bid. The indicator variable 1(Early bid received) is equal to one if the first bid received was placed within the first 15 minutes of the auction. Control variables in columns (2) and (4) include: log distance from bidder to client, log number of invited bidders (by itself and interacted with new platform dummy), log distance from the client to the nearest treatment center, log average distance from the client to all potential bidders' garages, client latitude and longitude, log population of the commune, auction number trend (linear, quadratic, and cubic), new platform indicator, indicators for morning and lunch time, and arrondissement, year-quarter, and day-of-week fixed effects. Robust standard errors in parentheses. *: $p < 0.10$, **: $p < 0.05$, ***: $p < 0.01$.

Table 4 distinguishes between the competitive and collusive rationales for early bids in revisable-bid auctions. In the first column, we measure the difference in the amount of the first bid between sealed-bid and revisable-bid auctions. Column (1) shows that the first bid received in a revisable-bid auction is 880 CFA higher (roughly 3% higher) than the first bid in a sealed-bid auction. This is inconsistent with early bidding in revisable-bid auctions being a way to signal strength which would have implied a lower rather than higher bid.

Columns (3) and (4) look at the effect of early bidding on the final winning bid. We define a bid as early if it arrives within the first 15 minutes. When this happens in the revisable-bid auction format, all bidders are informed about the value of the lowest early bid, but in sealed-bid auctions no information is provided. The presence of an early bid may be correlated with the attractiveness of the job (i.e., attractive jobs receive more bids, and are more likely to receive early bids) in both auction formats. In revisable-bid auctions, it may also be a means for collusive firms to coordinate on higher prices.

To test these hypotheses, we construct a difference-in-differences estimator. In particular, we estimate the difference in the value of the winning bid associated with receiving an early bid in the revisable-bid vs sealed-bid formats. The coefficient associated with the uninteracted early bid dummy variable is negative and statistically significant (1020 CFA in column (3)), suggesting that early bids and low winning bids are correlated in sealed-bid auctions. This is because firms bid more aggressively (both lower and more quickly) for jobs they find more attractive. Controlling for characteristics of the client and the winning bidder lowers this difference to 320 CFA, consistent with the fact that observable differences across clients such as neighborhoods or distance are important determinants of cost.

The interaction term with the revisable-bid format dummy is positive and large in magnitude (890 CFA in column (3)), implying that, relative to sealed-bid auctions, consumers are offered significantly higher prices in revisable-bid auctions when we observe an early price signal. This result is consistent with the interpretation that early bidding is used by collusive firms in revisable-bid auctions to coordinate on higher bids. This difference is smaller when we add control variables in column (4). This implies that the probability of receiving a non-competitive early bid in revisable auctions is correlated with client characteristics. For instance, some neighborhoods are likely more collusive than others.

4.3.2 Categorizing bidders and their effects on outcomes

The previous section highlighted the presence of non-competitive behavior, while also giving evidence that many winning bids appear to be competitive. For example, the distribution of winning bids shows that a non-trivial fraction of bids are placed using a fine price grid. Similarly, the bimodal distribution of winning times in the revisable-bid auctions reveals a mixture of early and late bids. This heterogeneity could in principle be caused by differences in bidding strategies across auctions, with bidders submitting competitive bids for certain types of clients or time periods, and submitting insincere bids for others. Alternatively, this heterogeneity could be due to systematic differences in the propensity of some bidders to behave competitively in the auctions.

In this section, we test the hypothesis that the data are generated by two groups of bidders: competitive and non-competitive types. Recall that competitive bidders should avoid round bids in sealed-bid auctions, and submit late bids in revisable auctions. To test this hypothesis, we exploit the panel dimension of our data by measuring the importance of heterogeneity in bidders' propensity to bid competitively (as opposed to heterogeneity across bids or auctions). We then relate our measures of the competitiveness of the bidders invited to an auction with auction outcomes such as average bids and participation frequency.

We conceptualize this as a problem of unobserved heterogeneity and use the K-means clustering algorithm to recover these unobserved types, similar to Bonhomme et al. (2022).[16] This is an

[16]Because the word clustering is associated with both the computation of standard errors with correlation in

unsupervised learning algorithm that iteratively partitions the observations into K groups, seeking to minimize the total within-group deviation from the group averages. More precisely, the algorithm is:

0. Pick K observations x_j, $j = 1, 2, ..., K$, at random. Initialize the group average of group j as $\mu_j = x_j$, $j = 1, 2, ..., K$.
1. Assign each observation i to the group $G_j \subset \{1, 2, ..., N\}$ whose group average μ_j is closest to x_i in terms of Euclidean distance.
2. For each of the K current group assignments, compute the average characteristics $\mu'_j = \frac{1}{|G_j|} \sum_{i \in G_j} x_i$.
3. Repeat steps 1 and 2 until convergence of assignment or a maximum number of iterations is reached.

This is a greedy algorithm intended to minimize the within-set sum of squared errors,

$$\text{WSSE}(K, \mu) = \sum_{k=1}^{K} \sum_{i \in G_k} (x_i - \mu_k)^2.$$

Because the K-means algorithm requires credible estimates of the grouping variables that describe the bidders' participation, we only include active bidders submitting bids in at least 30 sealed-bid auctions. Since those bidders participate at a much higher rate than the average (30% compared to 10%), this sample includes the most experienced and attentive bidders. There are 39 bidders who satisfy this criterion, out of 96 bidders who submitted at least one bid. We refer to the remaining 57 bidders as "inactive." We use the following grouping variables in the algorithm:

- Round, Sealed-Bid: The proportion of bids placed by bidder i in the sealed-bid format taking the values 25,000 CFA or 30,000 CFA. (A measure of non-competitive behavior).
- Snipe, Revisable-Bid: The proportion of bids placed by i arriving after the final update in the revisable format when either no other bids have been placed or that undercut the standing low bid when there is one. (A measure of competitive behavior).
- Early, Revisable-Bid: The proportion of bids placed by i in the revisable-bid format in the first 15 minutes of the auction. (A measure of either non-competitive or naïve behavior).

Neither the timing of a bid nor whether or not it is round places any a priori restrictions on whether it is a low or high bid. Consequently, there is no mechanical relationship between the results of

outcomes across observations as well as the unsupervised learning method, we refer to the collections of firms selected by the K-means clustering algorithm as 'groups' and the variables used to select groups as 'grouping variables' to avoid confusion.

the K-means clustering algorithm and subsequent results that regress auction- or bidder-level price outcomes on the number of participants invited from each group or group membership, respectively.

Figure 6: Relationship between number of groups and within-group sum of squared errors

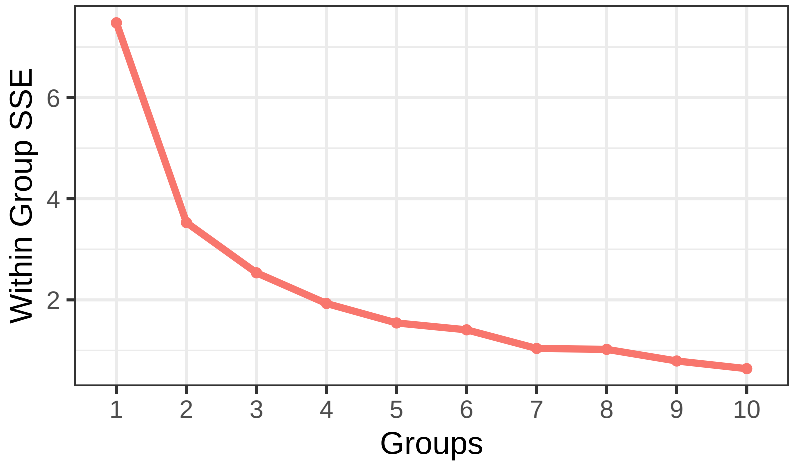

Note: Relationship between the within-group sum of squared errors from the K-means clustering algorithm and K, the number of groups.

The number of groups K is selected so that the marginal benefit of an additional group becomes "small" in terms of reducing the within-group sum of squared errors. Since a trivial global minimum in K is achieved by giving every observation its own group, the number of groups is selected so that the marginal benefit of an adding an additional group becomes roughly constant, implying that the main sources of unobserved heterogeneity have been accounted for. This is illustrated in Figure 6 which shows that increasing the number of groups past three appears to add little value in terms of decreasing the within-group sum of squared errors. We then interpret each of these three groups as a qualitatively distinct type of participant, with different strategies and competitive propensities.

Table 5: Correlation between grouping variables and other bidder characteristics

	Bid Placed	Bid Amount	Round, S-B.	Snipe, R-B.	Early, R-B.
Bid Placed	1.00				
Bid Amount	-0.04	1.00			
Round, S-B.	-0.00	0.35^{**}	1.00		
Snipe, R-B.	0.18	-0.18^{***}	-0.45^{***}	1.00	
Early, R-B.	0.12	0.08^{***}	0.38^{***}	-0.82^{**}	1.00

Note: Correlations between bidder-level characteristics. Bid placed and bid amount correspond to the proportion of invited auctions in which the bidder places a bid, and the bidder's average bid conditional on bidding. These variables are not included in the K-means clustering algorithm. The grouping variables include proportion of round bids in sealed-bid auctions, proportion of bids that arrive after the final update in the revisable format when either no other bids have been placed or that undercut the standing low bid when there is one (snipe), and proportion of bids that arrive in the first 15 minutes of revisable-bid auctions. *: $p < 0.10$, **: $p < 0.05$, ***: $p < 0.01$.

The correlations between the grouping variables given in Table 5 are useful for understanding why particular groups arise. Some of the correlations are mechanical; sniping (i.e., bidding late) in the revisable-bid format is strongly and negatively correlated with bidding early. Other correlations are not so mechanical. There is a statistically significant -0.45 correlation between tending to place round bids in sealed-bid auctions and sniping in revisable-bid auctions. This suggests that active bidders that time their bids competitively in revisable-bid auctions also avoid round bids that frequently result in ties in sealed-bid auctions. (Or, conversely, that bidders who decline to snipe in revisable-bid auctions also tend to place round bids in sealed-bid auctions.) In addition to showing correlations between the grouping variables, Table 5 shows correlations between grouping variables and two main final outcomes of interest: the share of times the firm bid and their average bid amount. The correlations are as expected. Bidders who are likely to place round and early bids are likely to place higher bids, while bidders who are more likely to snipe are more likely to place lower bids. These broad patterns illustrate that the variables exhibit relationships which can be credibly exploited through the K-means clustering algorithm.

Table 6: Group assignment

	N. Bidders	Bid Placed	Bid Amount	Round, S-B.	Snipe, R-B.	Early, R-B.
Competitive	10	0.25	26.30	0.34	0.54	0.11
Naïve	19	0.26	27.12	0.57	0.07	0.47
Collusive	10	0.22	28.41	0.68	0.25	0.21

Note: Number and average characteristics of bidders of each type as determined by the K-means clustering algorithm. The grouping variables used in the algorithm and shown in the table include proportion of round bids in sealed-bid auctions, proportion of bids that arrive after the final update in the revisable format when either no other bids have been placed or that undercut the standing low bid when there is one (snipe), and proportion of bids that arrive in the first 15 minutes of revisable-bid auctions.

The algorithm categorizes the 39 active bidders into three groups. Table 6 shows the size of each group, the average probability a member of the group places a bid, the average bid amount placed by a member of the group, and the average of the three grouping variables. Because K-means clustering is an unsupervised learning algorithm, the researcher has no direct control over the assignment of bidders to groups beyond the selection of grouping variables. Nevertheless, the groups exhibit patterns of behavior that align with our expectations of the unobserved discrete heterogeneity in the population. We label as *competitive* the group of bidders with the lowest frequency of round bids in sealed-bid auctions (0.34), the highest propensity to snipe in revisable-bid auctions (0.54), and the lowest propensity to place early bids in revisable-bid auctions (0.11). We label as *collusive* the group whose members, in contrast to those in the competitive group, place round bids in sealed-bid auctions at a high rate (0.68), snipe at almost half the rate of the competitive group (0.25), but still avoid early bidding in revisable-bid auctions (0.21). In between competitive and collusive, we label as *naïve* the group of bidders who place round bids in sealed-bid auctions at a similar rate to the collusive group (0.57), but rarely snipe in revisable bid auctions (.07), and place early bids almost half the time in revisable-bid auctions (0.47).

Although the K-means clustering algorithm did not target the average bid amount or average propensity to place a bid, these values align with the groups' labels. The bids are highest and participation lowest for the collusive group, with the competitive group bidding 2,000 CFA less on average and 3 pp more often. Finally, the naïve group participates at high rates and places relatively competitive bids, but does so in unsophisticated ways that likely undermine their overall effectiveness. This confirms that the groups recovered by the K-means clustering algorithm not only have systematic differences in timing and roundness of bids placed, but also exhibit differences in bidding behavior.

We next use regressions to look at how the groups differ in terms of behavior at the bid level and formally test whether bidding behavior is significantly different across the three groups. Appendix Table A-4 regresses bid amount and whether a valid bid was placed on group dummies in the sealed-bid format. Inactive bidders are included in the regression and are the excluded category. On average, competitive bidders place bids that are 2,300 CFA less than those placed by inactive bidders, while naïve bidders bid only 500 CFA less. Collusive bidders actually bid 500 CFA more than inactive bidders, or 2,800 CFA more than competitive bidders. This means that the bids of collusive bidders are 10% higher than the bids of competitive bidders. The differences between all three groups are statistically significant. This provides strong evidence that competitive bidders not only engage in sniping rather than bidding early and avoid placing round bids, but also place lower bids in the auctions. While the differences in the probability of placing a valid bid are not as stark or statistically significant, the competitive group is still the most likely to place a bid (27 pp more likely than inactive bidders), while the collusive group is only 17 pp more likely to place a bid than inactive bidders.

To evaluate the effect of non-competitive bidding on auction outcomes, we measure the effect of inviting bidders from the various groups on auction outcomes. We use OLS to estimate the following regression:

$$\begin{aligned} y_t =& \alpha_{\text{competitive}} N_{\text{competitive},t} + \alpha_{\text{naïve}} N_{\text{naïve},t} + \alpha_{\text{collusive}} N_{\text{collusive},t} \\ &+ \gamma_{\text{invited}} N_{\text{invited},t} + x_t \beta + \varepsilon_t, \end{aligned} \tag{3}$$

where y_t is the amount of the winning bid in auction t, the number of total valid bids, whether at least one bid was placed, and whether the winning bid was accepted conditional on at least one bid being placed or unconditionally for all auctions. The market-structure variables proxy for the competitiveness of each auction, as measured by the number and group type of invited bidders and characteristics of the job. We control for the total number of active bidders invited from each group, as well as the total number of invited bidders (which also includes the inactive bidders who are not placed in groups). Thus, the coefficients α_g represent the average change in y_t as a consequence of inviting one additional bidder from group g, controlling for the total number invited, $N_{\text{invited},t}$. Since invitations to bid in each auction were randomized, α_g measures the causal effect of swapping an inactive bidder for a bidder from group g, holding the total number invited and other auction characteristics constant.

Table 7: Effect of bidder composition on auction outcomes

VARIABLES	(1) N Bids	(2) 1(N Bids > 0)	(3) Winning bid	(4) 1(Accept)	(5) 1(Transaction)
N Invited, Competitive	0.21***	0.028***	-0.36***	0.021***	0.026***
	(0.019)	(0.0046)	(0.048)	(0.0073)	(0.0063)
N Invited, Naïve	0.22***	0.034***	-0.085**	-0.0013	0.0097**
	(0.014)	(0.0036)	(0.040)	(0.0055)	(0.0047)
N Invited, Collusive	0.16***	0.023***	-0.041	0.0017	0.011*
	(0.018)	(0.0047)	(0.048)	(0.0072)	(0.0060)
N Invited, Total	0.027***	0.0069**	-0.093***	-0.0043	-0.0018
	(0.010)	(0.0028)	(0.030)	(0.0041)	(0.0035)
Observations	5,162	5,162	4,177	3,621	4,491
R-squared	0.341	0.191	0.318	0.057	0.056
Comp.=Naïve	0.67	0.20	4.7e-07	0.0057	0.018
Comp.=Collusive	0.027	0.41	1.9e-07	0.041	0.054
Coll.=Naïve	0.0023	0.029	0.43	0.71	0.89
Mean dep. variable	1.97	0.81	25.8	0.30	0.24

Note: The dependent variables are (1) the number of bids placed, (2) an indicator for at least one bid being placed, (3) the value of the winning bid (in 1,000 CFA), (4) an indicator for whether the client accepted the desludging price offered, and (5) an indicator for whether the client received a desludging price offer and accepted it. Columns (1) and (2) include all auctions, column (3) includes all auctions with a winning bid, column (4) includes all auctions with a winning bid in which the buyer was not subsidized, and column (5) includes all auctions in which the buyer was not subsidized regardless of whether there was a winning bid. Additional controls include: log distance from the client to the nearest treatment center, log number of invited bidders (and its interaction with a new-platform dummy), log average distance from the client to all potential bidders' garages, client latitude and longitude, log population of the commune, auction number trend (linear, quadratic, and cubic), new platform indicator, revisable-bid auction indicator, indicators for morning and lunch time, and arrondissement, year-quarter, and day-of-week fixed effects. Robust standard errors in parentheses. *: $p < 0.10$, **: $p < 0.05$, ***: $p < 0.01$.

Column (3) of Table 7 shows that swapping an inactive bidder for a competitive bidder leads to a statistically and economically significant 360 CFA reduction in the winning bid. Swapping an inactive bidder for a naïve bidder would only lead to an 85 CFA reduction, and swapping an inactive bidder for a collusive bidder has no impact on the winning bid. Each additional inactive bidder invited to the auction decreases the winning bid by 93 CFA. We reject the hypothesis that the effect of inviting competitive bidders on the winning bid is the same as the effect of inviting naïve or collusive bidders. Column (1), shows that inactive bidders do not bid very often. Each additional inactive bidder increases the number of bids by only 0.03. Swapping an inactive bidder with a bidder from the competitive or naïve groups yields 0.21 or 0.22 additional bids, compared with 0.16 additional bids for each invitee from the collusive group. Some auctions end with no bids placed at all. As seen in column (2), inviting more active bidders from any group increases the probability that there is at least one bid placed, with effect sizes similar across the three groups. Competitive and naïve bidders both raise the probability by 3 pp, compared to collusive bidders who raise the probability by 2 pp.

Even if there is a winning bid, it may be so high that the client finds it unappealing and does not accept. The final two columns consider the probability that a mechanized desludging is transacted due to the auction, but for different samples. Column (4) looks at whether the client accepts the winning bid, conditional on there being a winning bid, in unsubsidized auctions. Column (5) looks at whether the auction leads to a mechanized desludging transaction and so the sample includes all unsubsidized auctions, whether or not there was a winning bid. If there are no bids in that auction, then there is no transaction.

In both cases, inviting an additional competitive bidder raises the probability the winning bid is accepted by 2-3 pp. Inviting more inactive, naïve, or collusive firms does not significantly raise the probability that the client will accept the desludging offer, conditional on there being at least one bid. Swapping inactive bidders for collusive bidders does not increase the likelihood that a winning bid is accepted, conditional on there being a winning bid, presumably because collusive bidders' bids are relatively high (seen in column (4)). But, because collusive bidders bid more often than inactive bidders, they increase the likelihood of there being a winning bid, which increases the unconditional likelihood of the auction ending in a mechanized desludging transaction (seen in column (5)). If the outcome of interest is whether the auction ends with the client purchasing a mechanized desludging in that auction as in column (5), then swapping out inactive bidders for collusive bidders would be a slight improvement. But, swapping out inactive bidders for competitive bidders would be an even larger and more significant improvement. The mean probability that a client receives an offer and accepts the price is 0.30 conditional on there being a winning bid and 0.24 for all auctions including those without bids, respectively. Thus the impact of inviting additional competitive bidders is substantial. In particular, the probability that the auction leads to a transaction increases by 3 pp for each competitive bidder, relative to a mean probability of 0.24, or a 12.5% increase.

In summary, the mix of competitive and non-competitive auction outcomes documented in Section 4.3.1 is due to bidder heterogeneity. Some bidders behave competitively, while others submit non-competitive bids either by strategically avoiding competition (collusive types) or by submitting bids using strategies inconsistent with profit maximization (naïve types). Within our auction sample, clients who were randomly matched with a larger number of competitive bidders received better prices, and as a result were more likely to purchase a mechanized desludging service from the platform. While this analysis suggests important gains from increasing competition in the market, we cannot directly extrapolate from equation (3) and Table 7 to measure the effect of collusion since only a handful of competitive bidders were invited in any given auction. In Section 5, we conduct a counter-factual simulation exercise in which we predict auction outcomes when all invited bidders are competitive.

5 Measuring economic damages from collusion

In this section, we evaluate the economic damages from the collusion that takes place on the platform. In addition to measuring the effect on prices, we also quantify the effect of collusion on take-up of mechanized desludging on the platform and on consumer surplus. From there, we make a back-of-the-envelope calculation of the impact on health. We use the observed bidding strategy of competitive types to construct a counter-factual distribution of offers that approximate environments in which only competitive type bidders are invited to the auctions. This is not meant to be a full-equilibrium prediction of how the market would operate under competition. Rather, our objective is to provide conservative estimates of the consumer's welfare gains from stimulating competition in the market under minimal behavioral assumptions.

We do so by estimating the policy function of competitive bidders using the classification obtained in Section 4.3.2. To implement this simulation exercise we estimate three functions: (i) bidders' participation decisions, (ii) bidders' bid strategies conditional on participating, and (iii) consumers' acceptance probabilities. The participation and bidding strategies are estimated in the full sample of sealed-bid auctions. We focus on sealed-bid auctions because bidders submit a single bid and do not interact with one another dynamically within the auction.[17] The demand function is estimated using the acceptance decisions of non-subsidized clients across both types of auctions.

The participation decision and bidding strategy of competitive bidders are estimated separately assuming that the decision to bid is independent of the marginal cost of providing the service. The participation decision is estimated using a Probit model, and the bids are assumed to follow a log-normal distribution. The results are summarized in Table 8. Column (1) estimates the policy function used to predict bidder participation, and column (2) estimates the log-normal distribution of bids. Both policy functions are estimated using the sample of sealed-bid auctions for competitive bidders only. The results confirm that distance is an important determinant of bidders' strategies. In column (2) we find that the bid function of competitive bidders is increasing in distance (both to the client and to the nearest treatment center), consistent with cost-based pricing. The point estimates in column (1) imply that, all else equal, competitive bidders' participation is increasing in distance. The reverse is true for other types (unreported regression results). One possible reason for this is that firms behaving collusively compete in clearly delineated territories, while competitive bidders can avoid detection by bidding on clients located far from their garage.

[17]In terms of stimulating use of mechanized desludgings, the sealed-bid format also yields slightly lower prices as seen in Table 3.

Table 8: Estimated bidder strategies and consumer demand

VARIABLES	(1) 1(Participate)	(2) Bid (log)	(3) 1(Accept)	(4) 1(Accept)
Winning bid (x1000)			-0.11***	-0.13***
			(0.0080)	(0.043)
Dist. (log): Garage to client	0.11***	0.061***		
	(0.031)	(0.0068)		
Dist. (log): Client to treatment	0.15**	0.092***	0.13	0.17
	(0.072)	(0.016)	(0.081)	(0.11)
No. Invited (log)	0.18	-0.045		
	(0.19)	(0.040)		
Observations	4,440	1,362	3,620	3,620
Unit of observation	Bidders	Bidders	Auctions	Auctions
Auction sample	Sealed-bid	Sealed-bid	All	All
Bidders	Competitive	Competitive	All	All
Mean dep. variable	0.31	3.24	0.30	0.30
Control function?			No	Yes
First-stage (F)				44.3
Exog. test (P-value)				0.58
Avg. demand elasticity			-3.52	

Note: Columns (1), (3), and (4) are estimated by probit while column (2) is estimated by OLS. Columns (3)-(4) exclude subsidized households. Control variables include: log distance from the client to the nearest treatment center, log average distance from the client to all potential bidders' garages, client latitude and longitude, log population of the commune, auction number trend (linear, quadratic, and cubic), new platform indicator, indicators for morning and lunch time, and arrondissement, year-quarter, and day-of-week fixed effects. Column (4) instruments for prices using a control-function estimator (Newey, 1987). The instruments are the number of invited active bidders located less than 10 km from the client and distance to the nearest invited active bidder. Robust standard errors in parentheses. *: $p < 0.10$, **: $p < 0.05$, ***: $p < 0.01$.

The acceptance probability of consumers is estimated using the full sample of auctions (excluding subsidized households). We leverage exogenous variation in offers (due to the random invitation lists) to identify the platform demand elasticity. More specifically we estimate the distribution of the consumer's willingness-to-pay on the platform using the following Probit model:

$$\Pr(\text{Accept}|b_t^{\min}, x_t) = \Pr\left(b_t^{\min} < \underbrace{x_t\beta + \epsilon_t}_{=r_t}\right) \tag{4}$$

where ϵ_t is normally distributed with variance σ_ϵ^2. The reservation price, r_t, corresponds to the value of the consumer's outside option: mechanized desludging obtained through standard channels, manual desludging, or no desludging (e.g., postpone). We conducted a small-scale informal survey of platform clients, and our conversations reveal that most clients who rejected an offer from the platform chose to purchase a mechanized desludging from the traditional market.[18] Based on this evidence, we believe that the platform provided an outside option which households used to bargain for a better price for their desludging with their local provider. Additional evidence on the impacts of the call center on local market prices is provided in Deutschmann et al. (2024a).

Our preferred specification assumes that the lowest bid is independent of the residual entering the willingness-to-pay of consumers. This assumption would be violated if, for instance, the residuals determining participation or bid amounts were correlated with unobserved market conditions when clients contact the platform. We control for this by adding quarter-year fixed-effects, flexible auction number trends, and time-of-day and day-of-week indicators.

The estimated price coefficient implies an average platform demand elasticity of -3.52 (median -3.17), confirming that consumers are very price sensitive.[19] This elasticity is slightly lower than the one obtained in Section 3 using variation in distance to the Rufisque border (-3). Because many clients use the platform to compare prices for mechanized desludgings through the platform to those they could get for a mechanized desludging in the open market, the platform elasticity is likely an upper bound (in absolute value) of the elasticity of substitution between mechanized and manual desludging.

In column (4), we relax the exogeneity restriction using the control function approach developed by Newey (1987). This specification exploits variation in the number and distance of active bidders invited to auction t to construct two instrumental variables (random by design): the number of invited active bidders located less than 10 km from the client and the distance to the nearest invited active bidder. The results validate our exogeneity assumption. We cannot reject the null hypothesis that ϵ_t is independent of the lowest bid. The estimated price coefficient is also quite similar in both specifications (i.e., -0.11 vs -0.13).

Next we use the estimated competitive policy and demand functions to simulate counter-factual auction outcomes. For each sealed-bid auction, we first simulate the participation decision (independent binomial draws) of each invited bidder, and, conditional on participating, we sample bids from the estimated log-normal distribution. In both cases we use the parameter estimates from Columns (1) and (2), and calculate the predicted strategies of each bidder using the observed characteristics of the client (primarily distance and time). We then use the estimated demand function to calculate the predicted acceptance probability evaluated at the lowest simulated bid. We repeat this process $1,000$ times, and report an average value by neighborhoods and across auctions.

[18] In the small follow-up survey with 45 clients who turned down the desludging they were offered from the platform, 64% claimed to have purchased a mechanized desludging in the traditional market, 4% purchased a manual desludging in the traditional market, 24% had not yet performed a desludging, and 7% performed a desludging but did not tell us what type it was.

[19] The elasticity for client i is given by:

$$\varepsilon_i^d = \frac{\partial \Pr(A_t = 1|X_i, b_t^{min})}{\partial b_t^{min}} \times \frac{b_t^{min}}{\Pr(A_t = 1|X_i, b_t^{min})} = -\frac{\phi(-(1/\sigma_\epsilon)b_t^{min} + x_t\beta/\sigma_\epsilon)}{\sigma_\epsilon}\frac{b_t^{min}}{\Phi(-(1/\sigma_\epsilon)b_t^{min} + x_t\beta/\sigma_\epsilon)}$$

.

Table 9: Competitive platform simulation results by neighborhoods

	(1) Num. bids Average	(2) Bid Average	(3) Winning bid Average	(4) Accepted bid Average	(5) Acceptance Probability
Panel A: Observed auctions					
Rest of market	2.04	27.55	25.77	24.87	0.30
Rufisque	0.81	26.56	24.84	17.14	0.06
Panel B: Scenario 1					
Rest of market	3.88	26.26	22.60	21.64	0.42
Rufisque	2.13	23.22	21.17	19.68	0.18
Panel C: Scenario 2					
Rest of market	2.28	26.37	23.88	22.55	0.37
Rufisque	1.22	23.34	21.95	19.99	0.16

Note: The observed and simulation samples correspond to all sealed-bid auctions for non-subsidized households (2,598 auctions). Bids are measured in 1,000 CFA. The simulations are conducted by estimating the participation (Probit) and bidding strategy (Log-normal) of competitive bidders, as well as the acceptance probability of consumers (Probit). See Table 8. The number of simulation draws for each auction is 1,000. Scenario 1: All invited bidders bid according the competitive type policy function. Scenario 2: Inactive types submit the same bid/participation decisions as in the observed sample.

We consider two counter-factual scenarios. In the first, we assume that all invited bidders decide whether to participate and how much to bid according to the policy function of the competitive type. In the second, we assume that inactive bidders (those who submitted less than 30 bids overall) do not change their behavior, and only replace the bids of collusive and naïve types. These two scenarios allow us to avoid taking a stand on whether it is collusion or high participation costs which lead inactive bidders to bid infrequently. Since the expected winning bid is decreasing in the number of bids, the second scenario provides a conservative estimate of the effect of collusion on prices and take-up.

Table 9 summarizes the simulation results. As seen in column (5) of Panel A, in the observed auctions 30% of offers are accepted in the collusive areas of Dakar, compared to only 6% in Rufisque (the low price neighborhood). In contrast, in our most competitive scenario (scenario 1 in Panel B), the average accepted offer in column (4) is CFA 21,640, and 42% of offers are accepted in the collusive area. Scenario 2 in Panel C slightly attenuates the effect of competition on take-up (37%) and accepted prices (22,550). Across all neighborhoods, the simulation results confirm that the lack of competition in the market leads to a large decline in the number of mechanized desludgings purchased through the platform.

To get a better understanding of the source of the gain from increasing competition, the table also reports the average bid amount and the number of bids received per auction. In both scenarios, column (2) shows that replacing non-competitive bidders with competitive types decreases the average bid in the collusive neighborhoods by 4.7% and 4.3% respectively. In contrast, column (3) shows that the average winning bid decreases by 12.2% in scenario 1, and by 7.3% in scenario 2. In scenario 2, this implies that more than half of the decline in winning bid is due to competitive bidders submitting lower bids, and the rest is due to the fact that collusive types are less likely to participate. This can also be seen in column (1) in the slightly higher number of bids in Dakar in scenario 2 compared to in reality. In scenario 1 which additionally replaces the inactive bidders with competitive types, more than 60% of the effect is due to increased participation, since inactive types are unlikely to participate.

Not surprisingly, the take-up rate of mechanized desludgings through the platform (the acceptance probability) remains low in Rufisque, because consumers are able to get lower prices by contracting directly with local garages. Rufisque simulations show both an increase in the average accepted price on the platform and an increase in take-up through the platform. This is because the bidders' participation rate is particularly low in Rufisque, and the counter-factual simulations decrease the number of auctions with zero bids.

Finally, we use the estimated distribution of willingness-to-pay from equation (4) to measure the effect of increasing competition on consumer surplus. To do so, we simulate demand residuals ϵ_t for each auction (simulated and observed), and compute consumer surplus: $CS_t = \max\{0, r_t - b^{\min}\}$. Averaging across simulation draws, we estimate that the average consumer surplus generated by the platform is CFA 1,570. Replacing all bidders by competitive types as in scenario 1 nearly doubles the average consumer surplus to 2,960 (or by 89%), while replacing only naïve and collusive types increases consumer surplus to 2,280 (or by 45%). Importantly, this represents a lower bound on the loss in consumer surplus due to collusion on the platform. This is because in a fully-competitive scenario, competitive bidders would likely change their bidding strategy which would lower bids even further.

5.1 Collusion's effects on health

It is generally the case that collusion leads to high prices and low consumer surplus. The welfare impacts in this setting are particularly large because mechanized desludgings are a sanitary method of desludging involving positive externalities. If consumers are discouraged from purchasing a mechanized desludging due to the high price, and instead purchase a cheaper but less sanitary manual desludging, this has ramifications for health. We make back-of-the envelope calculations of the impact that collusion has on health, namely diarrhea. Because any such calculation involves many assumptions, we construct multiple estimates involving different assumptions.

To estimate the health benefits of increasing competition, we must first calculate the effect of mechanized desludgings on health outcomes. We focus on the incidence of the mechanized desludging technology on diarrhea cases, using estimates from Deutschmann et al. (2024b). In that paper, within the same neighborhoods of Dakar, some individuals were offered high subsidies on mechanized desludgings, while others were offered low subsidies (close to the market rate). Households that were offered a high subsidy were 2.9 pp more likely to get a mechanized desludging. Each additional neighbor in a household's nearest four surveyed neighbors who was offered a high subsidy led the number of household members with diarrhea to decrease by 0.028. The implied elasticity of the number of diarrhea cases (H_i) with respect to the number of mechanized desludgings in a neighborhood (Q_i) is given by:

$$\frac{\Delta H_i}{\Delta Q_i}\frac{\bar{Q}}{\bar{H}} = \frac{\Delta H_i}{\Delta \text{Subsidy}_i}\left(\frac{\Delta Q_i}{\Delta \text{Subsidy}_i}\right)^{-1}\frac{\bar{Q}}{\bar{H}} = \frac{-0.028}{0.029}\frac{0.295}{0.49} = -0.58,$$

where "Subsidy" refers to the number of neighbors receiving a high subsidy. Therefore, this broadly suggests that a 1% increase in mechanized desludging leads to a 0.6% decrease in diarrhea.[20]

Next, we need to predict the change in mechanized desludging associated with an increase in competition. Based on the simulations above, replacing non-competitive bidders with competitive bidders would increase purchases of mechanized desludgings on the platform by 40% under scenario 1, and by 23% under scenario 2. Under the assumption that these increases are associated with an equivalent reduction in the demand for manual desludging, increasing competition on the platform would decrease the number of diarrhea cases by 14 to 23%.

This is likely an upper bound on the health benefit of the two counter-factual scenarios. This is because when the platform price goes down and more people purchase mechanized desludgings on the platform, some of the new purchasers may have switched from purchasing a mechanized desludging in the traditional market while others switched from purchasing a manual desludging in the traditional market. For comparison, Deutschmann et al. (2024a) estimate an elasticity of demand for mechanized compared to manual desludgings of -2.2 (compared to the -3.5 estimated above).[21] Using this lower elasticity, we estimate that demand for mechanized desludging would increase by 26% and 15% respectively under scenarios 1 and 2,[22] which would lead to a reduction in diarrhea cases ranging from 9 to 15%.

While the estimate of a 9 to 23% decrease in diarrhea incidence caused by an increase in competition is approximate, and involves many assumptions and caveats, the effect is reasonably large. For comparison, the roll-out of a new municipal sanitation system in Salvador Brazil decreased

[20] This estimate of the elasticity of diarrhea with respect to mechanized desludging is an underestimate since it does not take into account the health spillovers found in Deutschmann et al. (2024b) whereby more households using mechanized desludgings spills over to increase the mechanized desludging use of their neighbors.

[21] Deutschmann et al. (2024a) use survey data on the distribution of prices for mechanized and manual desludging services to estimate a price elasticity of demand of -2.2. The platform demand elasticity will naturally be larger (in absolute value) than the mechanized market demand elasticity estimated above if the outside option of consumers in the platform is mechanized service in the traditional market (as opposed to manual service). The other discrepancy is that the consumers calling the platform might be more "price savvy" than those in the survey and are therefore more price sensitive.

[22] The counterfactuals in scenarios 1 and 2 decrease the winning bids by 12% and 7% respectively.

diarrhea by 22% (Barreto et al., 2007), with the impact being twice as large in higher risk areas. Also in Salvador, using earlier data, Moraes et al. (2003) found that childhood diarrhea was 33% lower in neighborhoods with proper drainage and sanitation. Thus, decreasing collusion in mechanized desludging markets could have similar impacts on health as building improved drainage and sewerage systems - a much more expensive and disruptive undertaking.

6 Conclusion

We document the importance of imperfect competition and collusion in the market for mechanized desludging services in Dakar. Using non-experimental data on prices and transactions in the traditional desludging market, we first establish that areas in Rufisque supplied by companies not affiliated with the trade association exhibit 40% lower mechanized desludging prices and 50% higher mechanized desludging take-up than areas in central Dakar. This suggests that collusion may have a large impact on prices which in turn has a large impact on take-up which in turn has deleterious effects on sanitation and health.

We then use experimental data to test for the presence and impact of non-competitive behavior. We create an anonymous auction platform which randomly assigns firms to jobs and auction formats (sealed-bid and revisable-bid). We analyze the bidding and participation strategies of firms to detect deviations from competitive behavior. We document the presence of two main strategies inconsistent with individual profit maximization and competition: (i) the prevalence of round bids in the sealed-bid auction format, and (ii) the prevalence of early bidding in the revisable-bid auction format. We use randomized variation in auction format and bidder invitations together with a K-means clustering algorithm to establish that a large group of active participants, which we call collusive, systematically avoid competing by using both of these strategies. In contrast, a group of competitive bidders demonstrates a willingness to undercut their rivals, avoid ties, and submit late bids to snipe in the revisable-bid auctions.

We conclude that while there exists a group of suppliers willing to submit competitive bids, this group is not big enough to overcome the majority of bidders behaving collusively. This is consistent with what we observe in the traditional market where most consumers in the areas of Dakar controlled by the Association pay very high prices, while a significant fraction of consumers in this region pay prices that are as competitive as in Rufisque. Most of this dispersion in prices is due to unobserved differences across consumers, and is present even within narrowly defined neighborhoods. In light of the experimental results from the auction market, we suspect that some consumers are able to negotiate better prices by getting quotes from non-collusive truckers. As in the auctions, this group of non-collusive truckers is likely too small to serve the entire market. The fact that a majority of firms do not submit competitive quotes, even when it is secret and thus difficult to detect suggests repeated interactions between truckers and the fear of losing access to the services and contracts provided by the Association is large enough to discipline most firms, both members and non-members, which are active in the market.

Finally, using the observed bidding strategy of competitive firms, we find that eliminating collusion in this market could lead to important welfare gains by lowering prices, increasing take-up of mechanized desludging, and improving health outcomes. Our simulation results provide conservative estimates of the economic damages of collusion in the market. We find that replacing non-competitive bidders with competitive ones would increase the number of mechanized desludgings on the platform by 23-40%, increases consumer surplus by 45-89%, and decreases diarrhea by 9-23%. This illustrates the importance of market power as a first-order source of market inefficiency with significant implications for health.

References

Abrantes-Metz, R. M., Luke, F., Geweke, J., and Taylor, C. T. (2006). A variance screen for collusion. *International Journal of Industrial Organization*, 24(3):467–486.

Aghion, P. and Griffith, R. (2005). *Competition and Growth: Reconciling Theory and Evidence*. MIT Press.

Alé-Chilet, J. (2018). Gradually rebuilding a relationship: The emergence of collusion in retail pharmacies in Chile. Working paper.

Alé-Chilet, J. and Atal, J. P. (2020). Trade associations and collusion among many agents: Evidence from physicians. *Rand Journal of Economics*, 51(4):1197–1221.

Alsan, M. and Goldin, C. (2019). Watersheds in child mortality: The role of effective water and sewerage infrastructure, 1880–1920. *Journal of Political Economy*, 127(2):586–638.

Asker, J. (2010). A study of the internal organization of a bidding cartel. *American Economic Review*, 100(3):724–762.

Asker, J., Collard-Wexler, A., and De Loecker, J. (2019). (Mis)allocation, market power, and global oil extraction. *American Economic Review*, 109(4):1568–1615.

Athey, S., Bagwell, K., and Sanchirico, C. (2004). Collusion and price rigidity. *Review of Economic Studies*, 71(2):317–349.

Athey, S., Levin, J., and Seira, E. (2011). Comparing open and sealed-bid auctions: Evidence from timber auctions. *Quarterly Journal of Economics*, 126(1):207–257.

Bajari, P. and Hortaçsu, A. (2003). The winner's curse, reserve prices, and endogenous entry: Empirical insights from eBay auctions. *RAND Journal of Economics*, 34(2):329–355.

Bajari, P. and Ye, L. (2003). Deciding between competition and collusion. *Review of Economics and Statistics*, 85(4):971–989.

Banerjee, A., Hanna, R., Kyle, J., Olken, B. A., and Sumarto, S. (2019). Private outsourcing and competition: Subsidized food distribution in Indonesia. *Journal of Political Economy*, 127(1):101–137.

Barkley, A. (2023). The human cost of collusion: Health effects of a Mexican insulin cartel. *Journal of the European Economic Association*, 21(5):1865–1904.

Barreto, M. L., Genser, B., Strina, A., Teixeira, M. G., Assis, A. M. O., Rego, R. F., Teles, C. A., Prado, M. S., Matos, S. M., Santos, D. N., dos Santos, L. A., and Cairncross, S. (2007). Effect of city-wide sanitation programme on reduction in rate of childhood diarrhoea in northeast Brazil: Assessment by two cohort studies. *The Lancet*, 370(9599):1622–1628.

Barrett, C. (1997). Food marketing liberalization and trader entry: Evidence from Madagascar. *World Development*, 25(5):763–777.

Beaman, L. and Magruder, J. (2012). Who gets the job referral? Evidence from a social networks experiment. *American Economic Review*, 102(7):3574–3593.

Beaman, L., Magruder, J., and Robinson, J. (2014). Minding small change among small firms in Kenya. *Journal of Development Economics*, 108:69–86.

Bergquist, L. F. and Dinerstein, M. (2020). Competition and entry in agricultural markets: Experimental evidence from Kenya. *American Economic Review*, 110(12):3705–3747.

Bernasconi, M., Espinosa, M., Macchiavello, R., and Suarez, C. (2023). Relational collusion in the Colombian electricity market. Unpublished Manuscript.

Best, M. C., Hjort, J., and Szakonyi, D. (2023). Individuals and organizations as sources of state effectiveness. *American Economic Review*, 113(8):2121–2167.

Bonhomme, S., Lamadon, T., and Manresa, E. (2022). Discretizing unobserved heterogeneity. *Econometrica*, 90(2):69–86.

Brown, Z., Montero, E., Schmidt-Padilla, C., and Sviatschi, M. M. (2024). Market structure and extortion: Evidence from 50,000 extortion payments. *Review of Economic Studies*. Forthcoming.

Byrne, D. and de Roos, N. (2019). Learning to coordinate: A study in retail gasoline. *American Economic Review*, 109(2):591–619.

Caoui, E. H. (2022). A study of umbrella damages from bid-rigging. *Journal of Law and Economics*, 65(2):239–277.

Chassang, S., Kawai, K., Nakabayashi, J., and Ortner, J. (2022). Robust screens for non-competitive bidding in procurement auctions. *Econometrica*, 90(1):315–346.

Chassang, S., Kawai, K., Nakabayashi, J., and Ortner, J. (2023). Using bid rotation and incumbency to detect collusion: A regression discontinuity approach. *Review of Economic Studies*, 90(1):376–403.

Che, Y.-K. and Kim, J. (2007). Robustly collusion-proof implementation. *Econometrica*, 74(4):1063–1107.

Clark, R., Coviello, D., and De Leverano, A. (2024a). Complementary bidding and the collusive arrangement: Evidence from an antitrust investigation. *American Economic Journal: Microeconomics*. Forthcoming.

Clark, R., Horstmann, I., and Houde, J.-F. (2024b). Hub-and-spoke cartels: Theory and evidence from the grocery industry. *American Economic Review*, 14(3):783–814.

Clark, R. and Houde, J.-F. (2013). Collusion between asymmetric retailers: Evidence from a gasoline price-fixing case. *American Economic Journal: Microeconomics*, 5(3):97–123.

Comanor, W. S. and Schankerman, M. A. (1976). Identical bids and cartel behavior. *The Bell Journal of Economics*, 7(1):281–286.

Conley, T. and Decarolis, F. (2016). Detecting bidders groups in collusive auctions. *American Economic Journal: Microeconomics*, 8(2):1–38.

DellaVigna, S. (2009). Psychology and economics: Evidence from the field. *Journal of Economic Literature*, 47(2):315–372.

DellaVigna, S. and Gentzkow, M. (2019). Uniform pricing in U.S. retail chains. *Quarterly Journal of Economics*, 134(4):2011–2084.

Deutschmann, J., Gars, J., Houde, J.-F., Johnson, T., Lipscomb, M., Mbeguere, M., Nehrling, S., Schechter, L., and Zhu, S. J. (2024a). Using auctions to improve sanitation at scale. Working paper.

Deutschmann, J. W., Lipscomb, M., Schechter, L., and Zhu, J. (2024b). Spillovers without social interactions in urban sanitation. *American Economic Journal: Applied Economics*, 16(3):482–515.

Doval, L. and Skreta, V. (2022). Mechanism design with limited commitment. *Econometrica*, 90(4):1463–1500.

Froeb, L. M., Koyak, R., and Werden, G. J. (1993). What is the effect of bid rigging on prices? *Economics Letters*, 42:419–423.

Genesove, D. and Mullin, W. P. (2001). Rules, communication, and collusion: Narrative evidence from The Sugar Institute case. *American Economic Review*, 91(3):379–398.

Graham, D. A. and Marshall, R. C. (1987). Collusive bidder behavior at single-object second-price and English auctions. *Journal of Political Economy*, 95(6):1217–1239.

Horner, J. and Jamison, J. (2007). Collusion with (almost) no information. *RAND Journal of Economics*, 38(3):804–822.

Hortaçsu, A., Luco, F., Puller, S. L., and Zhu, D. (2019). Does strategic ability affect efficiency? Evidence from electricity markets. *American Economic Review*, 109(12):4302–42.

Igami, M. and Sugaya, T. (2022). Measuring the incentive to collude. *Review of Economic Studies*, 89(3):1460–1494.

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

Kawai, K. and Nakabayashi, J. (2022). Detecting large-scale collusion in procurement auctions. *Journal of Political Economy*, 130(5):1585–1629.

Laffont, J. and Martimort, D. (1997). Collusion under asymmetric information. *Econometrica*, 65(4):875–911.

Lewis-Faupel, S., Neggers, Y., Olken, B. A., and Pande, R. (2016). Can electronic procurement improve infrastructure provision? Evidence from public works in India and Indonesia. *American Economic Journal: Economic Policy*, 8(3):258–83.

Lipscomb, M. and Schechter, L. (2018). Subsidies versus mental accounting nudges: Harnessing mobile payment systems to improve sanitation. *Journal of Development Economics*, 135:235–254.

Marshall, R. C. and Marx, L. (2007). Bidder collusion. *Journal of Economic Theory*, 133:374–402.

McAfee, R. and McMillan, J. (1992). Bidding rings. *American Economic Review*, 82(3):579–599.

Moraes, L., Cancio, J. A., Cairncross, S., and Huttly, S. (2003). Impact of drainage and sewerage on diarrhoea in poor urban areas in Salvador, Brazil. *Transactions of the Royal Society of Tropical Medicine and Hygiene*, 97(2):153–158.

Mund, V. A. (1960). Identical bid prices. *Journal of Political Economy*, 68(2):150–169.

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

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

Neilson, C. (2021). Targeted vouchers, competition among schools and the academic achievement of poor students. Working paper.

Newey, W. K. (1987). Efficient estimation of limited dependent variable models with endogenous explanatory variables. *Journal of Econometrics*, 36:231–250.

Pesendorfer, M. (2000). A study of collusion in first-price auctions. *Review of Economic Studies*, 67(3):381–411.

Planet Money (2018). 855: The poop cartel.

Porter, R. H. and Zona, D. (1993). Detection of bid rigging in procurement auctions. *Journal of Political Economy*, 101(3):518–538.

Porter, R. H. and Zona, D. (1999). Ohio school milk markets: An analysis of bidding. *RAND Journal of Economics*, 30(2):263–288.

Robinson, M. S. (1985). Collusion and the choice of auction. *RAND Journal of Economics*, 16(1):141–145.

Roth, A. E. and Ockenfels, A. (2002). Last-minute bidding and the rules for ending second-price auctions: Evidence from eBay and Amazon auctions on the internet. *American Economic Review*, 92(4):1093–1103.

Ryan, N. (2020). Contract enforcement and productive efficiency: Evidence from the bidding and renegotiation of power contracts in India. *Econometrica*, 88(2):383–424.

Ryan, N. (2021). The competitive effects of transmission infrastructure in the Indian electricity market. *American Economic Journal: Microeconomics*, 13(2):202–242.

Skreta, V. (2006). Sequentially optimal mechanisms. *Review of Economic Studies*, 73:1085–1111.

Walsh, C. (2020). Social impacts of new radio markets in Ghana: A dynamic structural analysis. Working paper.

A Appendix tables

Table A-1: Mean differences in characteristics across neighborhoods

	Rest of the market	Rufisque	Difference: Rest - Rufisque	
			Mean	p-value
Nearest center (km)	4.56	2.8	1.77	0
Nearest garage (km)	1.23	.99	.23	0
Num. trucks w/in 3 km	49.64	28.85	20.79	0
Household size	11.19	11.64	-.45	.139
Number of earners	3.43	3.21	.22	.055
Number of other households	1.04	.76	.27	.003
Number of rooms	7.13	7.22	-.09	.592
House ownership	.77	.87	-.09	0
Two story house	.32	.23	.08	0
Wealth index	.04	.17	-.13	.036
Wide road	.89	.95	-.06	0

Note: Table uses household survey data. Nearest center and nearest garage give the distance in kilometers from the household to the nearest treatment center and mechanized desludger garage respectively. Household size is the number of people who live in the household, the number of earners is the number of them who earn income, and number of other households is the number of other households that live in the same house and share the latrine pit. The number of rooms is the number of rooms in the house. Wide road is the enumerator's assessment of whether the road leading to the house is wide enough for a mechanized desludging truck to enter.

Table A-2: Regression of mechanized desludging prices on consumer characteristics and distance to the Rufisque boundary

VARIABLES	(1) Rest of the market	(2) Rufisque
1(.5 < Boundary dist. <1)	4.66***	2.01**
	(1.18)	(0.86)
1(1 < Boundary dist. <1.5)	6.59***	0.48
	(0.95)	(0.75)
1(1.5 < Boundary dist. <2)	7.14***	3.03***
	(0.81)	(0.73)
1(2 < Boundary dist. <2.5)	6.54***	2.79***
	(0.77)	(0.83)
1(2.5 < Boundary dist. <3)	8.17***	1.97**
	(0.86)	(0.96)
1(3 < Boundary dist. <3.5)	7.14***	2.64**
	(0.85)	(1.11)
1(3.5 < Boundary dist. <4)	7.16***	
	(1.02)	
1(4 < Boundary dist. <4.5)	7.02***	
	(0.96)	
1(4.5 < Boundary dist. <5)	7.54***	
	(1.03)	
1(5 < Boundary dist.)	8.23***	
	(0.70)	
Nearest center (km)	0.64***	-0.40**
	(0.088)	(0.20)
Nearest garage (km)	0.16	0.11
	(0.19)	(0.49)
Num. trucks w/in 3 km	0.012**	-0.071**
	(0.0047)	(0.029)
Household size	-0.038	0.030
	(0.024)	(0.039)
Number of earners	-0.17***	-0.15
	(0.062)	(0.11)
Number of other households	-0.021	0.17
	(0.076)	(0.15)
Number of rooms	0.18***	0.0094
	(0.042)	(0.076)
House ownership	-0.16	-0.060
	(0.28)	(0.55)
Two story house	1.07***	1.34***
	(0.26)	(0.47)
Wealth index	0.79***	0.14
	(0.096)	(0.17)
Wide road	1.22**	-1.76*
	(0.56)	(1.04)
Observations	4,409	394
R-squared	0.122	0.119
H0: Boundary coefficients = 0 (p-value)	0	
H0: Boundary coefficients = 0 (p-value)		0.00074
H0: Rufisque coefs. = Rest of market coefs. (p-value)		0

Note: Table uses household survey data. Nearest center and nearest garage give the distance in kilometers from the household to the nearest treatment center and mechanized desludger garage respectively. Household size is the number of people who live in the household, the number of earners is the number of them who earn income, and number of other households is the number of other households that live in the same house and share the latrine pit. The number of rooms is the number of rooms in the house. Wide road is the enumerator's assessment of whether the road leading to the house is wide enough for a mechanized desludging truck to enter. Standard errors are clustered at the household level. *: $p < 0.10$, **: $p < 0.05$, ***: $p < 0.01$.

Table A-3: Summary statistics on auction control variables for sealed-bid and revisable-bid auctions

	Sealed-bid			Revisable			
	n	mean	sd	n	mean	sd	Diff
No. invited (log)	2598	2.52	0.22	2564	2.53	0.22	0.003
Dist. from client to treatment (log)	2598	1.71	0.52	2564	1.72	0.53	0.011
Ave. dist. from potential bidders to client (log)	2598	2.42	0.21	2564	2.42	0.21	0.001
Client latitude	2598	14.76	0.02	2564	14.76	0.02	0.001
Client longitude	2598	-17.36	0.05	2564	-17.36	0.05	0.002
Population (log)	2598	11.26	0.81	2564	11.29	0.81	0.025
Auction number	2598	0.50	0.29	2564	0.50	0.29	-0.006
Auction number^2	2598	0.34	0.30	2564	0.33	0.29	-0.009
Auction number^3	2598	0.26	0.29	2564	0.24	0.28	-0.010
1(New platform)	2598	0.46	0.50	2564	0.45	0.50	-0.006
1(Morning)	2598	0.62	0.48	2564	0.61	0.49	-0.015
1(Lunch)	2598	0.21	0.41	2564	0.23	0.42	0.016
Tuesday	2598	0.19	0.40	2564	0.20	0.40	0.008
Wednesday	2598	0.19	0.39	2564	0.19	0.39	-0.000
Thursday	2598	0.17	0.37	2564	0.17	0.38	0.007
Friday	2598	0.17	0.38	2564	0.18	0.38	0.005
Saturday	2598	0.05	0.21	2564	0.05	0.21	0.001
1(Downtown)	2598	0.12	0.32	2564	0.10	0.31	-0.011
1(Guediawaye)	2598	0.13	0.34	2564	0.14	0.35	0.010
1(Pikine Dagoudane)	2598	0.09	0.29	2564	0.09	0.29	0.002
1(Rufisque)	2598	0.04	0.20	2564	0.04	0.20	-0.002
1(Thiaroye)	2598	0.32	0.47	2564	0.32	0.47	-0.008
2013Q3	2598	0.02	0.15	2564	0.02	0.15	-0.000
2013Q4	2598	0.01	0.12	2564	0.01	0.12	0.001
2014Q1	2598	0.06	0.24	2564	0.05	0.23	-0.006
2014Q2	2598	0.15	0.35	2564	0.16	0.37	0.012
2014Q4	2598	0.09	0.28	2564	0.10	0.30	0.010
2015Q1	2598	0.08	0.28	2564	0.09	0.29	0.008
2015Q2	2598	0.08	0.28	2564	0.08	0.27	-0.006
2015Q3	2598	0.12	0.32	2564	0.12	0.33	0.001
2015Q4	2598	0.06	0.24	2564	0.06	0.24	-0.001
2016Q1	2598	0.03	0.18	2564	0.03	0.18	-0.001
2016Q2	2598	0.03	0.18	2564	0.04	0.18	0.003
2016Q3	2598	0.05	0.22	2564	0.04	0.20	-0.009
2016Q4	2598	0.01	0.12	2564	0.01	0.09	-0.005*

Note: Auction number is normalized to be between 0 and 1. *: $p < 0.10$, **: $p < 0.05$, ***: $p < 0.01$.

Table A-4: Bidding behavior

VARIABLES	(1) Bid amount	(2) 1(Participation)
Type: Competitive	-2.31***	0.27***
	(0.21)	(0.0071)
Type: Naïve	-0.52***	0.23***
	(0.19)	(0.0052)
Type: Collusive	0.49**	0.17***
	(0.21)	(0.0064)
Observations	5,171	33,177
R-squared	0.265	0.119
Comp.=Naïve	0	3.9e-06
Comp.=Collusive	0	0
Coll.=Naïve	0	0
Mean dep. variable	27.1	0.15
N Bidders	89	126

Note: Bidder-level outcomes in sealed-bid auctions. The dependent variable in column (1) is every bid placed in sealed-bid auctions. The dependent variable in column (2) is the participation decision made by every bidder in every auction. Inactive bidders are included in the regression and are the excluded category. Additional control variables include: log distance from garage to client, log distance from the client to the nearest treatment center, log average distance from the client to all potential bidders' garages, client latitude and longitude, log population of the commune, auction number trend (linear, quadratic, and cubic), new platform indicator, revisable-bid auction indicator, indicators for morning and lunch time, and arrondissement, year-quarter, and day-of-week fixed effects. Robust standard errors in parentheses. *: $p < 0.10$, **: $p < 0.05$, ***: $p < 0.01$.

B Appendix figures

Figure B-1: Distribution of survey households across neighborhoods

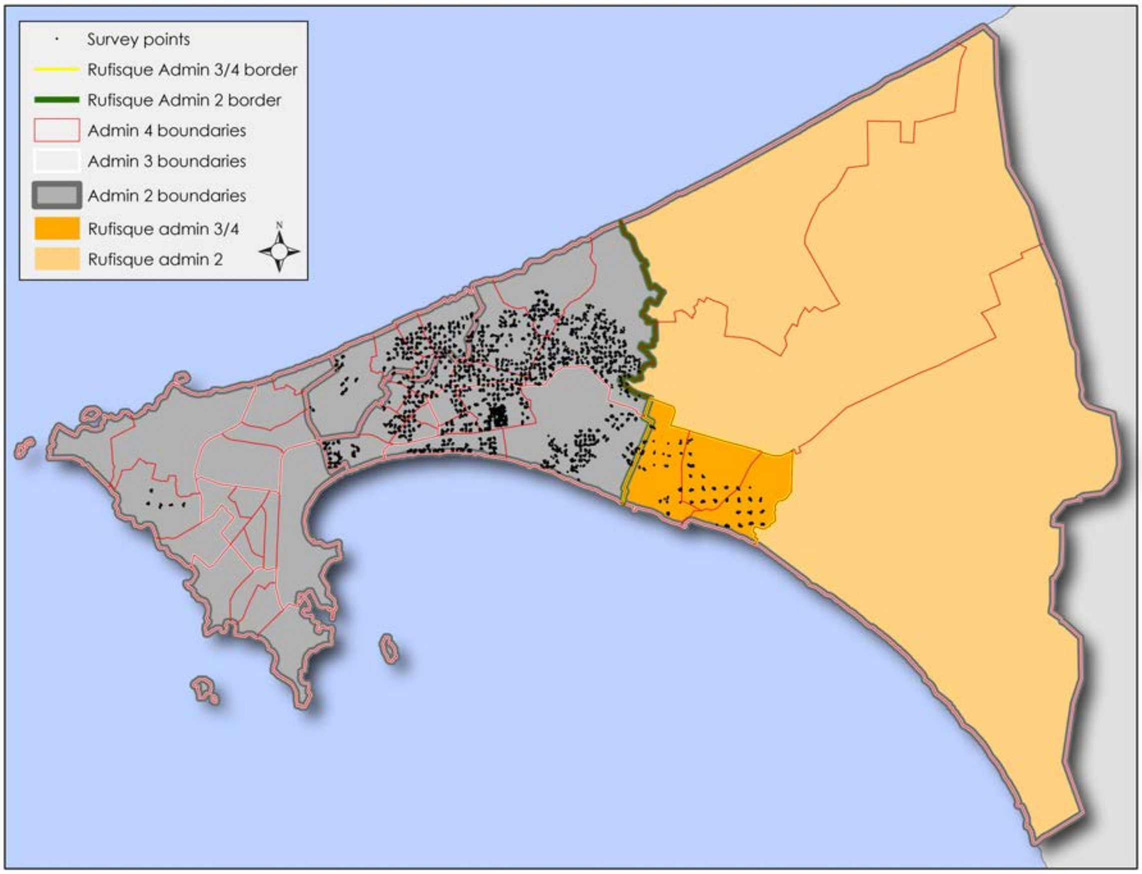

C Round bidding as optimal collusion

The analysis in the paper associates round bidding with collusive behavior. This begs the question, is round bidding an optimal or at least profitable way for a cartel to organize its collusive activities? This section uses concepts from mechanism design to answer the question in the affirmative. The intuition is that, when the cartel members have i.i.d. private information about costs, bidding on a grid of “round values” softens price competition while still allowing the cartel to use its private information to reduce total costs. This requires a punishment scheme that punishes off-grid bidding with a temporary reversion to the fully competitive equilibrium, as in standard Folk Theorem analyses of repeated games. A surprising result is that, even though the optimal direct mechanism can exploit communication between the cartel and its members, none is required at the optimum: a round bidding strategy is simply adopted, cartel members bid, and non-round deviations are punished.

Literature. The model we study is closest to Athey et al. (2004), and to a lesser extent Horner and Jamison (2007), who examine a repeated Bertrand pricing game. They come to similar conclusions about 'rigid pricing.' In the presence of incomplete information about costs, a cartel can benefit by eschewing full separation in favor of a bidding scheme in which firms collude on the monopoly price. They focus on the case in which firms are sufficiently patient to collude on the monopoly price, and consider a second case in which two prices are used. In contrast, we explicitly consider auctions rather than markets and allow for many rigid bids, rather than just one or two. The previous papers study the optimality of the agreement as posed as a problem of imperfect public monitoring, while we study the existence of collusive equilibria in the auction supergame that exhibit the empirical patterns observed in Senegal. In particular, we show that most profitable round-bidding strategies exist even for levels of patience that do not support rigid bidding on the optimal monopoly price (here, the reserve price).

In many ways, this is a commitment problem as considered in Skreta (2006) and Doval and Skreta (2022). The Cartel would like to commit its members to a certain pattern of behavior, but the members cannot themselves commit to follow it. These problems are typically studied from the perspective of a seller facing a group of buyers (to what extent can an auctioneer credibly refuse 'lowball' bids below a reserve price?). Our setting studies the perspective of a buyer facing a group of sellers (if the Cartel members can privately deviate ex post from a prescribed pattern of bidding, how can the Cartel maintain discipline ex ante?). In this sense, we find the commitment framework to be more useful than the repeated games framework.

A classic literature in auction theory following McAfee and McMillan (1992) considers how to design collusive bidding rings. This literature typically maximizes the Cartel's expected revenue, not the expected payoffs of the individual members, which avoids the issues of commitment and monotonicity that appear later. The Cartel member with the lowest cost is the one that maximizes Cartel revenue, so the issue in that literature is how to compensate the other members of the Cartel for 'taking a dive.' However, McAfee and McMillan (1992) also note that for weak cartels when side payments aren't feasible, there should arise price rigidity such that firms all collude on a single bid. Laffont and Martimort (1997) consider a version of the problem that explicitly assumes commitment to the mechanism, but generalizes other aspects of the analysis in McAfee and McMillan (1992). Che and Kim (2007) study the problem from the platform's perspective, and conclude that a simple take-it-or-leave-it offer from the platform to the Cartel is optimal. The main distinction in our case is that we study equilibrium behavior of the Cartel to verify the plausibility of our empirical observations in Senegal, as opposed to studying how the platform can best combat the presence of collusive participants.

To fix ideas and maintain verisimilitude with the auctions in Senegal, consider a paid-as-bid format in which the lowest bidder wins and is paid their bid, as long as that bid is below a reserve price R. Then in a non-collusive Bayesian Nash equilibrium, firms would maximize their private profits,

$$\pi_i(c_i) = \max_{b_i}(b_i - c_i) \times \text{pr}[\text{Submitted the lowest bid less than } R \mid b_i].$$

The reason to focus attention on the paid-as-bid mechanism in particular is that the platform must be budget-balanced, so that transfers from the household to the desludger must be equal for each job.

Incentive compatibility. In a Direct Mechanism, bidders report a type $\hat{c}_i$ rather than directly placing a bid, and then the Cartel itself maps reports to probabilities of winning $p(\hat{c}) = \{p_i(\hat{c}_1, ..., \hat{c}_N)\}_{i=1}^N$ and payments received $b(\hat{c}) = \{b_i(\hat{c}_1, ..., \hat{c}_N)\}_{i=1}^N$ that are consistent with the underlying paid-as-bid auction. So while the platform picks the winner at random from the set of bidders submitting the lowest bid below the reserve price, the Cartel can manipulate who submits which bids and their values, reflected in $\langle p, b\rangle$.

From firm i's perspective, the *interim probability* that it wins given a report of $\hat{c}_i$ is

$$p_i(\hat{c}_i) = \mathbb{E}_{\hat{c}_{-i}}\left[p_i(\hat{c}_1, ..., \hat{c}_{i-1}, \hat{c}_i, \hat{c}_{i+1}, ..., \hat{c}_N)\right]]$$

where $p_i(\hat{c}_1, ..., \hat{c}_N)$ is the probability that the Cartel selects firm i to win given the reports $\hat{c}$. Before asking what kinds of pairs $\langle b, p\rangle$ are optimal for the Cartel, the more fundamental question is, given that the bidders can lie about their types, what kinds of pairs $\langle p, b\rangle$ are actually possible, or *implementable*? For what kinds of $\langle p, b\rangle$ do bidders submit $\hat{c}_i = c_i$?

The classic characterization of Bayesian incentive compatibility is:

Proposition 1 (Myerson (1981))**.** *A Direct Mechanism $\langle p, b\rangle$ is Bayesian incentive compatible iff it satisfies the monotonicity and envelope conditions. Firm i's interim payoff is*

$$\pi_i(c_i) = \max_{\hat{c}_i} p_i(\hat{c}_i)(b_i(\hat{c}_i) - c_i) = \int_{c_i}^{R} p_i(z)dz,$$

and their equilibrium bid function is

$$b(c_i) = c_i + \frac{\int_{c_i}^{R} p_i(z)dz}{p_i(c_i)}. \tag{5}$$

A key assumption is that the Cartel can commit to the mechanism. For example, in an auction to sell a good, the auctioneer can commit to refuse to consider bids below the reserve price. This will become relevant in our environment: the extent to which the Cartel can enforce discipline is limited by its ability to detect and punish deviations.

Static platform and Cartel optimization. The Cartel is a collective of firms that maximize their expected interim profits,

$$\mathbb{E}_{(c_1,\ldots,c_N)}\left[\sum_{i=1}^{N}\pi_i(c_i,c_{-i})\right].$$

Before turning to the Cartel's solution, it helps to revisit the platform's problem and the standard non-collusive solution. In standard auction design without a colluding cartel, Proposition 1 is sufficient to derive the platform's expected payoff for any Bayesian incentive compatible mechanism $\langle p,b\rangle$. In particular, rearranging Equation (5) allows the expected payment for the c_i type to be expressed as

$$p_i(c_i)b_i(c_i)=p_i(c_i)c_i+\int_{c_i}^{R}p_i(z)dz.$$

Taking the expectation over c_i yields

$$\mathbb{E}_c[p_i(c)b_i(c)]=\int_{\underline{c}}^{R}\left(p_i(c_i)c_i+\int_{c_i}^{R}p_i(z)dz\right)dG(c_i)+\int_{R}^{\overline{c}}(0)dG(c_i)$$

and integration by parts with respect to c_i yields the expected payment to firm i,

$$\mathbb{E}_c[p_i(c_i)b_i(c_i)]=\mathbb{E}_c\left[p_i(c_1,\ldots,c_N)\left(c_i+\frac{G(c_i)}{g(c_i)}\right)\right].$$

The quantities

$$\phi(c_i)=c_i+\frac{G(c_i)}{g(c_i)}\quad\text{and}\quad\lambda(c_i)=\frac{G(c_i)}{g(c_i)}$$

are the *virtual cost* and *informational rent*, respectively. The informational rent captures the portion of the expected cost of the job that accrues to the firm as compensation for revealing its private information, and the virtual cost is the true cost plus the informational rent.

Let v be the value to the household of getting the job done (or the platform's estimate of that value). Then the platform deducts the expected procurement cost for each firm from the expected benefit and takes the weighted sum over firms to get expected consumer surplus:

$$\mathbb{E}_{(c_1,\ldots,c_N)}\left[\sum_{i=1}^{N}p_i(c_1,\ldots,c_N)\times(v-\phi(c_i))\right].$$

How does the platform maximize this quantity? Setting aside the monotonicity condition for the moment, the optimal choice is to pick the firm i for whom $v-\phi(c_i)$ is largest and non-negative, and refuse to award the job if the virtual cost exceeds the benefit for every firm. When is the monotonicity condition slack? As long as $\phi(c_i)$ is non-decreasing, so that the firm with the lowest virtual cost is also the firm with the lowest realized cost. Then a first-price or second-price auction with a reserve price satisfying $\phi(R)=v$ will be an optimal mechanism, as well as more exotic games like the all-pay auction or third-price auction. This is the general logic of the celebrated Revenue Equivalence result. Equation (5) determines payments as a function of the probabilities of winning, so that all Bayesian incentive compatible mechanisms that award the good the same way generate the same expected consumer surplus.

In order for the monotonicity condition be satisfied however, so that the optimal allocation rule $p(c)$ selects the firm with the lowest reported cost to do the job, it is necessary that

$$\phi(c_i) = c_i + \frac{G(c_i)}{g(c_i)}$$

be non-decreasing, and it is sufficient that

$$\frac{G(c_i)}{g(c_i)}$$

be a non-decreasing function; essentially, that $G(c_i)$ is log-concave.[23] Otherwise, the platform would want to violate the monotonicity condition's restriction that higher costs win with lower probability, and this would not be implementable. The lower-cost firm would simply pool with the higher-cost firm that is winning with higher probability, thereby increasing its payoff.

All of this is familiar auction design, but from the platform's perspective, the goal is approximately to minimize the expected cost of the job. The Cartel, however, has almost the opposite objective: to maximize expected interim profit. Indeed, if firm i's expected interim profit is

$$\pi_i(c_i) = \int_{c_i}^{R} p_i(z)dz + \pi_i(R)$$

where $\pi_i(R) = 0$, then a similar set of calculations as above yields

$$\mathbb{E}_{c_i}[\pi_i(c_i)] = \mathbb{E}_{(c_1,\dots,c_N)}\left[p_i(c_1, ..., c_N)\frac{G(c_i)}{g(c_i)}\right],$$

and summing over the firms yields the Cartel's interim profits:

$$\mathbb{E}_c\left[\sum_{i=1}^{N} \pi_i(c_i)\right] = \mathbb{E}_{(c_1,\dots,c_N)}\left[\sum_{i=1}^{N} p_i(c_1, ..., c_N)\lambda(c_i)\right].$$

The Cartel is maximizing its expected informational rents, and has preferences that are approximately opposed to those of the platform. Unconstrained, it would like to pick the winner from among the firms with the highest realized costs, since $\lambda(c_i)$ is non-decreasing under the standard regularity condition. This is not implementable for a simple reason: every firm would then report

[23] Since $D^2 \log(G) = (g'G - g^2)/G^2$ and $D(G/g) = (g^2 - g'G)/g^2$, $-D^2 \log(G)$ and $D(G/g)$ have the same sign. So G/g is non-decreasing if G is log-concave.

that it was the highest type that wins with positive probability to the Cartel in order to maximize its payment and probability of being selected to bid and win in the auction. Consequently, every Cartel-optimal solution will involve pooling when $\lambda(c_i)$ is non-decreasing.

Consider mechanisms in which the Cartel partitions $[\underline{c}, R]$ by a set of interior points $\mathcal{C} = \{c_1, c_2, ..., c_L\}$, with the notational convention that $c_0 = \underline{c}$ and $c_{L+1} = R$. To each sub-interval there corresponds a bid: $[\underline{c}, c_1)$ to b_1, $[c_1, c_2)$ to b_2, ..., and $[c_L, R]$ to R. Denote these bids by $\mathcal{B}$. We call this kind of mechanism a *round-bidding strategy.* Ultimately, focusing on round-bidding strategies is without loss of generality, because as many interior types are added to the same interval, the probability of a tie in that interval tends to zero, and strict separation occurs in the limit.

Proposition 2. *If $\lambda(c_i)$ is non-decreasing on $[\underline{c}, R]$, then the optimal mechanism for the Cartel is a round-bidding strategy. Optimal mechanisms exist.*

Any mechanism with a strictly separating interval can be improved upon: by using the ironing procedure of Mussa and Rosen (1978) and Myerson (1981) and pooling those types into one report, the expected informational rents from that interval increase. Full separation on any sub-interval is sub-optimal, since some amount of pooling can reduce competition and raise expected informational rents.

The intuition for this result is that the informational rent $\lambda(c_i)$ is a non-decreasing function, but any incentive compatible mechanism must have a non-increasing probability of winning.

Dynamic enforcement and round bidding. Can the Cartel achieve the static optimum in a repeated game when monitoring and enforcement are required? A key assumption of the static analysis is that the Cartel can *commit to the mechanism* (e.g., Doval and Skreta (2022), Skreta (2006)). For an example of commitment, in an auction to sell a good the auctioneer can threaten to refuse to consider bids below the reserve price but this might not be credible in practice. In this environment, a similar issue becomes relevant: the firms do not report to the Cartel who bids for them, they place bids themselves, and will be tempted to renege on the Cartel's agreement once they know their private information. The extent to which the Cartel can enforce discipline is therefore limited by its ability to detect and punish deviations. Bids are not submitted to the Cartel directly, but to the platform, and the individual Cartel members' bids are largely unobservable to the Cartel as a whole. If all Cartel members are supposed to adopt the round-bidding strategy or withdraw when their costs exceed the reserve price, then low-cost types will be tempted to undercut round bids and win with strictly higher probability and make arbitrarily similar payments. If this kind of deviation isn't somehow monitored and punished, Cartel discipline crumbles.

With fully separating equilibria and public monitoring of just the winning bid, every history is approximately (up to withdrawals by firms with costs above the reserve price) measure zero, so deviations can only be inferred from examining the statistical properties of long histories to detect anomalies. With round-bidding strategies, however, compliance can be directly observed from the history of the winning bid: if it is on-$\mathcal{B}$, the Cartel has followed the agreement, and if any off-$\mathcal{B}$ winning bids are observed, someone has necessarily cheated. Thus, unlike with fully separating equilibria, deviations can be detected and punished instantly with round bidding.

Let the firms' common discount factor be given by δ. Then $\mathcal{B}$ is *enforceable for* δ if there exists some set of strategies in the infinitely repeated game for which there are no profitable deviations from bidding on-$\mathcal{B}$ after any history. The static unprofitability of strict separation on any interval actually facilitates dynamic enforcement.

Theorem 3. *If the discount factor δ is sufficiently close to 1, any incentive compatible round-bidding strategy $\mathcal{B}$ is enforceable. For a fixed discount factor δ, there exists a most profitable incentive compatible round-bidding strategy $\mathcal{B}^*$ that is enforceable for δ. As δ increases, the optimal round-bidding strategy has weakly fewer bids.*

So when the Cartel cannot commit to a mechanism, the optimal grid $\mathcal{B}$ becomes a function of temptation as well as collusive motives. Relatively impatient cartels can reduce the benefits of deviating and undercutting high on-$\mathcal{B}$ bids by introducing more bid increments, compensating compliance with a higher probability of winning. This keeps competition soft relative to the fully separating equilibrium, but reduces the temptation for lower cost firms to deviate. As the discount factor goes to zero, only the fully separating equilibrium is incentive compatible and enforceable as the limit of a round-bidding strategy with infinitely many increments, and as the discount factor goes to one, any incentive compatible round-bidding strategy is enforceable. For intermediate values of δ, the optimal $\mathcal{B}$ will be finite but become denser as firms become more impatient. This rationalizes the existence of relatively rich but finite grids, when the static solution suggests an extremely coarse one that includes only the reserve price R and the non-cooperative Bayesian Nash equilibrium has full separation.

It is surprising that no coordination or communication by the Cartel is required for optimal collusion. The pooling intervals are invariant to the realization of the firms' cost types, so the optimal cartel does not need to run a pre-auction knockout or otherwise exploit idiosyncratic and transient private information to determine the winner. Likewise, there is no 'turn-taking' or 'phases of the moon' coordination that might be detected using forensic algorithms. Firms bid in what appears to be an uncoordinated and competitive fashion, hiding their collusive behavior in round bidding that can be shrugged off as boundedly rational behavior. In the context of Senegal, this is essential, because it means the Cartel does not actually have to coordinate its behavior from auction to auction on tight deadlines: truckers optimally coordinate by using the round-bidding scheme.

Secret reserve prices. The model used a fixed reserve price R and subsequent analysis assumed that the household's willingness-to-pay W was known to the platform, which presumably set the optimal reserve price in the standard way, as $\phi(R) = w$. In reality, the households' willingness-to-pay vary from job to job and are unknown to the platform in advance. One alternative way to model the household's accept-or-reject decision is as a secret reserve price: if the distribution function for the household's willingness-to-pay W is $F_W(w)$, then the probability a bid b_i is accepted is $R(b_i) = 1 - F_W(b_i)$.

This creates an additional margin that complicates the Cartel's problem. Higher bids are now more likely to be rejected, so that colluding on a high price is not necessarily profit maximizing, since there is a continuous trade-off between the probability of a bid being accepted and job remuneration, as opposed to a single reserve price where the probability the bid is accepted goes from one to zero.

The presence of a secret reserve price thereby impacts the structure of the optimal partition. When the Cartel's problem is locally concave because the secret reserve price and likelihood of rejection dominate, there will be separation in bidding as above, and when the Cartel's problem is locally convex and the non-decreasing information rent dominates, it will employ pooling. So the exact analytical properties of the optimal solution can be determined from an optimal control approach, but the qualitative properties are similar.

Proof of Proposition 2:

Proof. Consider some candidate non-increasing $p^*(c)$ function. Since $p^*(c)$ is a non-increasing and bounded function on a compact interval, it has at most a countable number of jump discontinuities. Take any interval (a, b) on which $p^*(c)$ is continuous and strictly decreasing.

Since the mechanism is fully separating on this interval and incentive compatible, the probability of winning can be decomposed into the probability of having the lowest type, multiplied by the conditional probability of winning at all, conditional on the reported type, $\alpha(c_i)$,

$$(1 - G(c_i))^{N-1}\alpha(c_i).$$

Since p^* is continuous on (a, b), $\alpha(c_i)$ must then be continuous as well.

Consider ironing this interval, so that all types in (a, b) make the same report and receive the same probability of winning. The probability of being randomly selected as the winner given conditional on reporting a type in this interval is

$$\bar{p} = \sum_{j=0}^{N-1} \frac{1}{j+1} \frac{(N-1)!}{j!(N-1-j)!}(1 - G(b))^{N-1-j}(G(b) - G(a))^j.$$

This can be rewritten as

$$\begin{aligned} N(G(b)-G(a))\bar{p} &= \sum_{j=0}^{N-1} \frac{N!}{(j+1)!(N-(j+1))!}(1-G(b))^{N-(j+1)}(G(b)-G(a))^{j+1} \\ &= \sum_{j=1}^{N} \frac{N!}{j!(N-j)!}(1-G(b))^{N-j}(G(b)-G(a))^{j} \end{aligned}$$

Note that the right-hand side is similar to the distribution of an order statistic, except for the truncation at $G(a)$. The standard induction arguments provided in Proposition 4 below, however, reduce it to the quantity

$$H(a,b) = \int_a^b N(1-G(x))^{N-1} g(x)dx,$$

so that

$$\bar{p} = \frac{\int_a^b N(1-G(x))^{N-1} g(x)dx}{N(G(b)-G(a))} = \frac{\int_a^b (1-G(x))^{N-1} g(x)dx}{G(b)-G(a)}. \tag{6}$$

Notice that this quantity can be written

$$\bar{p} = \mathbb{E}[(1-G(x))^{N-1} | a < x < b].$$

Now, by the mean value theorem, there exists an $\bar{\alpha} = \alpha(\xi), \xi \in (a,b)$, such that

$$\int_a^b (1-G(x))^{N-1}\alpha(x)\lambda(x)g(x)dx = \int_a^b (1-G(x))^{N-1}\lambda(x)g(x)dx\bar{\alpha}.$$

Imagine ironing (a,b), and giving those who bid $c_i \in (a,b)$ the probability $\bar{p}\bar{\alpha}$ of winning, which is feasible with respect to the monotonicity condition because $p^*(a)\alpha(a) \geq \bar{p}\bar{\alpha}$. Now consider the comparison

$$\int_a^b (1-G(x))^{N-1}\alpha(x)\lambda(x)g(x)dx \gtreqless \int_a^b \bar{p}\bar{\alpha}\lambda(x)g(x)dx,$$

or, by previous calculations, this also equals

$$\int_a^b (1-G(x))^{N-1}g(x)\lambda(x)dx \gtreqless \int_a^b \bar{p}\lambda(x)g(x)dx,$$

or

$$\int_a^b \frac{(1-G(x))^{N-1}}{\bar{p}}\frac{g(x)}{G(b)-G(a)}\lambda(x)dx \gtreqless \int_a^b \frac{g(x)}{G(b)-G(a)}\lambda(x)dx.$$

Considering the left-hand side of the inequality, define

$$w(x) = \frac{(1-G(x))^{N-1}}{\bar{p}}\frac{g(x)}{G(b)-G(a)}.$$

Then $\int_a^b w(x)dx = 1$, since $\bar{p} = \mathbb{E}[(1-G(x))^{N-1}|a < x < b]$ is the conditional mean of $(1-G(x))^{N-1}$ on (a, b). Likewise, $\int_a^b dG(x)/(G(b)-G(a)) = 1$ on the right-hand side. But $w(x)$ places more weight on low values of x compared to $g(x)/(G(b) - G(a))$, and less weight on high values of x. Since $\lambda(x)$ is non-decreasing,[24] that implies the left-hand side must be weakly less than the right-hand side, and ironing the interval is weakly profitable. Ironing this way is feasible since it does not violate any monotonicity constraints, and indeed relaxes them on subsequent intervals of the type space. Therefore, ironing any continuously separating interval raises the value of the objective without violating any of the constraints.

Therefore, we can restrict attention to mechanisms in which there are no strictly decreasing segments, and $p^*(c)$ is constant almost everywhere and non-increasing.

Optimal partitioning. Take any grid $\mathcal{C} \subseteq [\underline{c}, R]$. The probability of winning when reporting a type $c_i \in [c_{\ell-1}, c_\ell)$ is

$$p_\ell = p(c_{\ell-1}, c_\ell) = \sum_{j=0}^{N-1} \underbrace{\frac{1}{1+j}}_{\text{Tie-breaking}} \times \underbrace{\frac{(N-1)!}{j!(N-1-j)!} \times (G(c_\ell) - G(c_{\ell-1}))^j (1 - G(c_\ell))^{N-1-j}}_{\text{Probability of } j \text{ other bidders in the } \ell\text{-th bin, } N-1-j \text{ in strictly higher bins}} .$$

This depends only on $c_{\ell-1}$ and c_ℓ, which was previously shown to equal

$$p(c_{\ell-1}, c_\ell) = \frac{H(c_{\ell-1}, c_\ell)}{N(G(c_\ell) - G(c_{\ell-1}))}.$$

For a fixed L, a partition $\mathcal{C}$ that satisfies the monotonicity constraint exists since c_L can be selected to be very close to R, c_{L-1} very close to c_L, and so on. The probabilities that the higher elements of the partition win are arbitrarily small and the entire $\{p_\ell\}_{\ell=1}^L$ sequence is non-increasing. This is a feasible point for any finite L, so the feasible set is non-empty. Call any such $\mathcal{C}$ *admissible*.

Take any admissible $\mathcal{C}$. Then it determines $\mathcal{B}$ because we can write a set of indifference conditions for all of the boundary types that must get the same payoff from b_ℓ and $b_{\ell+1}$,

$$p_\ell(b_\ell - c_\ell) = p_{\ell+1}(b_{\ell+1} - c_\ell), \quad \ell = 1, ..., L$$

with terminal condition

$$p_L(b_L - c_L) = p_R(R - c_L).$$

The terminal condition can be solved for b_L in terms of R and c_L,

$$b_L = c_L + \frac{p_R}{p_L}(R - c_L),$$

[24] Note that if $\lambda(x)$ were decreasing, as $v - \phi(x)$, we would have the standard, opposite conclusion that full separation is optimal.

and then back-substituted through the rest of the conditions for a fixed $\mathcal{C}$ and $b_{\ell+1}$ to get b_ℓ, using the recursion

$$b_\ell = c_\ell + \frac{p_{\ell+1}}{p_\ell}(b_{\ell+1} - c_\ell).$$

The bids are therefore a continuous function of the type partition, $\mathcal{B}(\mathcal{C})$, and the mapping between them means each is entirely determined by the other.

Let

$$H(c_{\ell-1}, c_\ell) = \sum_{j=1}^{N} \frac{N!}{j!(N-j)!}(G(c_\ell) - G(c_{\ell-1}))^j (1 - G(c_\ell))^{N-j}$$

be the probability that the lowest of N draws is between $c_{\ell-1}$ and c_ℓ. Then — using the same calculations as for (6) to express $p(c_{\ell-1}, c_\ell)$ in terms of $H(c_{\ell-1}, c_\ell)$ by way of Proposition 4 — the Cartel's expected payoff in a round in which all firms participate as intended is

$$N \sum_{\ell=1}^{L+1} p(c_{\ell-1}, c_\ell) \int_{c_{\ell-1}}^{c_\ell} \lambda(x) g(x) dx = \sum_{\ell=1}^{L+1} H(c_{\ell-1}, c_\ell) \mathbb{E}[\lambda(c) | c_{\ell-1} < c < c_\ell] \tag{7}$$

where the monotonicity constraint is[25]

$$p_\ell = \frac{H(c_{\ell-1}, c_\ell)}{N(G(c_\ell) - G(c_{\ell-1}))} \geq \frac{H(c_\ell, c_{\ell+1})}{N(G(c_{\ell+1}) - G(c_\ell))} = p_{\ell+1}, \quad \ell = 1, ..., L,$$

and the bids $\mathcal{B}(\mathcal{C})$ are determined by the backwards recursion

$$b_\ell = c_\ell + \frac{p_{\ell+1}}{p_\ell}(b_{\ell+1} - c_\ell), \quad \ell = 1, ..., L$$

with boundary condition $b_{L+1} = c_{L+1} = R$.

Since the objective function is continuous in $\mathcal{C}$ and the constraints are all continuous weak inequalities, the constraint set is a closed subset of $[\underline{c}, \overline{c}]^L$ and therefore compact. Therefore, a solution to the problem exists for a fixed L, by the extreme value theorem.

Existence of a finite, optimal L. The previous work established existence of an optimal partition $\mathcal{C}_L^*$, but not that there is a finite L that corresponds to a solution. In principle, the benefit of adding additional sub-intervals to the mechanism might increase indefinitely, so that no solution exists. The next part of the proof shows this is not the case.

[25] The quantity $\dfrac{H(c_{\ell-1}, c_\ell)}{N(G(c_\ell) - G(c_{\ell-1}))}$ looks like it might be poorly behaved (discontinuous, singular, etc.) as $c_\ell \to c_{\ell-1}$, but this isn't the case, because

$$\lim_{c_\ell \downarrow c_{\ell-1}} p(c_{\ell-1}, c_\ell) = \lim_{c_\ell \downarrow c_{\ell-1}} \frac{\int_{c_{\ell-1}}^{c_\ell} N(1 - G(x))^{N-1} g(x) dx}{N(G(c_\ell) - G(c_{\ell-1}))} = \lim_{c_\ell \downarrow c_{\ell-1}} \frac{\int_{c_{\ell-1}}^{c_\ell} N(1 - G(x))^{N-1} g(x) dx / (c_\ell - c_{\ell-1})}{N(G(c_\ell) - G(c_{\ell-1}))/(c_\ell - c_{\ell-1})} = (1 - G(c_{\ell-1}))^{N-1},$$

which is correct as the partition shrinks down to a single point.

Define the *widest implementable partition* for L as the solution to

$$J_L^W = \max_{\mathcal{C}} \min_{\ell} (c_{\ell-1} - c_\ell)^2$$

subject to $p(c_{\ell-1}, c_\ell) \geq p(c_\ell, c_{\ell+1})$ for $\ell = 1, ..., L$. Just like the interim profit maximization problem, this one has a solution because the objective function is continuous and the constraints characterize a closed subset of a compact set. This partition will maximize the minimum distance between its elements in the type space, subject to being implementable. Note that this implies that any other implementable partition $\mathcal{C}$ for any finite L must have some interval on which points are closer together, by construction. The solution to this problem spaces the points of $\mathcal{C}_L^W$ as far apart as possible without violating incentive compatibility. As shown earlier, a solution exists for any fixed L, so $\mathcal{C}_L^W$ and J_L^W are well-defined for all $L \in \mathbb{N}$.

Consider what happens as L becomes large: the optimal partitions $\mathcal{C}_L^W$ include more points that are partitioning the same bounded set, so that J_L^W must be a non-increasing sequence bounded below by a limit of zero. The probability of a tie then converges to zero almost everywhere, since the odds of another firm being in the same sub-interval $[c, c+\Delta)$ goes to zero:

$$\lim_{\Delta \downarrow 0} p_\Delta(c) =$$

$$\lim_{\Delta \downarrow 0} \left\{ (1 - G(c+\Delta))^N + \sum_{j=1}^{N-1} \frac{1}{1+j} \times \frac{(N-1)!}{j!(N-1-j)!} \times (G(c+\Delta) - G(c))^j (1 - G(c+\Delta))^{N-1-j} \right\}$$

$$= (1 - G(c))^N.$$

But this implies that $\mathcal{C}_L^W$ gets arbitrarily close to implementing the assignment rule

$$\overline{p}_i(c) = \begin{cases} \dfrac{1}{|\{k : c_k = \min_j c_j|\}|}, & i \in \text{argmin}_j c_j \\ 0, & \text{otherwise} \end{cases}$$

which is the fully-separating solution. That is the decision rule that minimizes the Cartel's objective function by minimizing informational rents, so this limit mechanism cannot be optimal. More formally, for any $\varepsilon > 0$, there exists a $\tilde{L}$ and a sub-interval of $[\underline{c}, R]$ of strictly positive measure where the difference in the sup-norm between the assignment rule based on $\mathcal{C}_L^W$ and $\overline{p}_i(c)$ is less than ε, and on that particular sub-interval there is a strictly profitable deviation for the Cartel that involves coarsening the partition to induce more pooling and less separation.

But $\mathcal{C}_L^W$ is the partition that spaces the points apart as widely as possible. Any other sequence — particularly the one that corresponds to the optimal partition with L points for each L — is going to exhibit the same phenomenon because it is not constructed to space the points as widely as it can. For any $\varepsilon > 0$, there will be some sub-interval of $[\underline{c}, R]$ where the points are clustered together at least as tightly as $\mathcal{C}_L^W$, and are even closer to the fully separating limit. Therefore, the optimal sequence of partitions $\mathcal{C}_L^*$ for L sufficiently large must have sub-intervals that include even more points than $\mathcal{C}_L^W$, and are even closer to full separation. For the same reason, this will result in sub-intervals that are arbitrarily close to fully-separating as L becomes arbitrarily large, which cannot be optimal if the Cartel is free to reduce L and adopt a coarser partition.

Define

$$J(L) = \max_{\mathcal{C}} \sum_{\ell=1}^{L+1} H(c_{\ell-1}, c_\ell)\mathbb{E}[\lambda(c)|c_{\ell-1} < c < c_\ell].$$

Then there must be some finite $\overline{L}$ such that for $L' > \overline{L}$, $J(\overline{L}) > J(L')$. Then since the set $\{1, ..., \overline{L}\}$ is compact, a solution to the problem $\max_{L \in 1,...,\overline{L}} J(L)$ must exist, which corresponds to the optimal mechanism.

$\square$

Proof of Theorem 3:

Proof. From Proposition 2, we know that strict separation on any interval is not profit-maximizing for the Cartel, so attention can be restricted to round-bidding strategies on some grid $\mathcal{B}$.

The Cartel's problem with dynamic enforcement. Fix an incentive compatible round-bidding strategy $(\mathcal{C}, \mathcal{B}(\mathcal{C}))$ (see Proposition 2), and denote the expected payoff to a single compliant Cartel member,

$$\Lambda(\mathcal{C}) = \frac{1}{N} \sum_{\ell=1}^{L+1} H(c_{\ell-1}, c_\ell)\mathbb{E}[\lambda(c)|c_{\ell-1} \leq c < c_\ell]$$

and the Bayesian Nash payoffs in the strictly separating equilibrium,

$$\pi^* = \int_{\underline{c}}^{R} (1 - G(c_i))^{N-1}\lambda(c_i)dG(c_i) + \int_{R}^{\overline{c}} 0dG(c_i).$$

Consider a perfect public monitoring strategy of the type:

– Phase 1: Bid according to $\mathcal{B}$. If the winning bid is on-$\mathcal{B}$, stay in Phase 1, otherwise move to Phase 2.

– Phase 2: Bid the non-collusive Bayesian Nash strategies

$$b^*(c_i) = c_i + \frac{\int_{c_i}^{R}(1 - G(z))^{N-1}dz}{(1 - G(c_i))^{N-1}}.$$

With probability ψ, stay in Phase 2, and with probability $1 - \psi$, return to Phase 1.

The Poisson punishment is useful because the system of equations that describes the firms' payoffs is continuous in $\psi \in [0, 1]$, while a discrete punishment period over $\mathbb{N}$ is not. Likewise, since the

Cartel is programming its round-bidding strategy, it is advantageous to maximize over compact, convex sets like $[0,1]$ that have convex images rather than a countably infinite set like $\mathbb{N}$.

Semi-continuity and continuity of deviation payoffs. Because round-bidding creates discontinuities in the probability of winning as a function of the type report $\hat{c}$, maximizers of the optimal deviation problem will typically not exist. Instead, we require that compliance delivers a better payoff than any available by deviating, and consequently, the supremum of the payoffs achievable by deviating.

Notice that since only the winning bid is revealed to the Cartel, deviations are only detected if the deviation wins the auction. It is logically possible that some type c_i has a strictly profitable deviation, but the deviation in bidding is not detected in a given period. As a consequence of the One Shot Deviation principle, however, it suffices to consider a deviation from on-$\mathcal{B}$ bidding — which might be detected and trigger a deviation, or not — followed by a return to compliance. The intuition is that because we seek compliance, a deviation that is profitable and recurrent would already constitute a contradiction. The goal is to deter even a single deviation from compliance with a suitable punishment.

Let $\Delta(c_i, \mathcal{C})$ correspond to this supremum, the highest payoff a type c_i can get when faced with the grid $\mathcal{C}$ from an optimal one shot deviation:

$$\begin{aligned}\Delta(c_i, \mathcal{C}) = \sup_{b'} \Bigg\{ &\text{pr[Lowest bid less than } R|b'] \left(b' - c_i + \frac{1}{1-\psi\delta} \left(\pi^* + \frac{(1-\psi)\delta}{1-\delta} \Lambda(\mathcal{C}) \right) \right) \\ &+ (1 - \text{pr[Lowest bid less than } R|b']) \left(0 + \frac{\delta}{1-\delta} \Lambda(\mathcal{C}) \right) \Bigg\}, \end{aligned} \tag{8}$$

where the discounted value of the Poisson punishment is already calculated and included on the first line. The consideration of the optimal deviation b' is over the trade-off between winning and triggering the punishment, versus getting the chance to strategically deviate again.

Since $\mathcal{B}(\mathcal{C})$ is a finite subset of $[\underline{c}, R]$, there are a finite number of jump discontinuities in $\Delta(c_i, \mathcal{C})$. Define

$$p^d(c_i) = \limsup_{c \uparrow c_i} p(c) \quad \text{and} \quad b^d(c_i) = \limsup_{c \uparrow c_i} b(c),$$

where $\langle p, b \rangle$ correspond to the round-bidding strategy. These are everywhere continuous except at points on $\mathcal{C}$ and $\mathcal{B}$, which correspond to the supremum of the values the function takes in that neighborhood. Since the Revelation Principle guarantees that every game of incomplete information has an alternative representation as a Direct Mechanism, we can substitute the functions $\mu^d = \langle p^d, b^d \rangle$ into Equation (8) as a reporting game in which the deviator picks a type to report rather than a bid.

Substituting these definitions yields

$$\sup_{c'}\left\{p^d(c')\left(b^d(c')-c_i+\frac{1}{1-\psi\delta}\left(\pi^*+\frac{(1-\psi)\delta}{1-\delta}\Lambda(\mathcal{C})\right)\right)+(1-p^d(c'))\left(0+\frac{\delta}{1-\delta}\Lambda(\mathcal{C})\right)\right\}$$

and $\Delta(c_i,\mathcal{C})$ is actually an upper semi-continuous function, since the payoffs at the discontinuities are now being adjusted up to the supremum of values that the function takes in an arbitrarily small neighborhood of the points of discontinuity. A maximizer is then guaranteed to exist by an extension of the Extreme Value Theorem to allow for upper semi-continuous objective functions, allowing us to write

$$\max_{c'}\left\{p^d(c')\left(b^d(c')-c_i+\frac{1}{1-\psi\delta}\left(\pi^*+\frac{(1-\psi)\delta}{1-\delta}\Lambda(\mathcal{C})\right)\right)+(1-p^d(c'))\left(0+\frac{\delta}{1-\delta}\Lambda(\mathcal{C})\right))\right\}.$$

With this adjustment, all type reports in the deviation mechanism μ^d except those on $\mathcal{C}$ are dominated. These are the only points at which the probability of winning changes, and it jumps up, so adopting an off-$\mathcal{C}$ report just reduces the payment the deviator might receive by reducing their bid but holding the probability of winning constant since their opponents' bids are only on $\mathcal{B}$: why further undercut yourself if you've already exploited the gains from undercutting the atoms? More formally, if a firm is already undercutting some b_ℓ, the probability $p^d(c')$ of winning does not vary on the interval $[c_{\ell-2},c_{\ell-1})$, so there is no reason to further reduce one's bid below $b^d(c')$.

That means the maximization can be re-written as a simple linear programming problem in which the c_i type selects the probability of strategically undercutting each atom so that the payoff approaches the supremum:

$$\max_{x_1,\ldots,x_L}\sum_{\ell=1}^{L}x_\ell\left\{p^d_\ell\left(b^d_\ell-c_i+\frac{1}{1-\psi\delta}\left(\pi^*+\frac{(1-\psi)\delta}{1-\delta}\Lambda(\mathcal{C})\right)\right)+(1-p^d_\ell)\left(0+\frac{\delta}{1-\delta}\Lambda(\mathcal{C})\right)\right\}\tag{9}$$

subject to $0\le x_\ell\le 1$ and $\sum_{\ell=1}^{L}x_\ell=1$.

Now, the solution correspondence x^* will be an extreme, bang-bang, solution. It will correspond to a unique ℓ, or the convex face of a fixed set of ℓ's which deliver the same expected payoff. Make a deterministic selection from the correspondence $\ell^*(c_i)$ that is right-continuous in c_i. Then the maximum payoff to deviating is

$$\begin{aligned}\Delta^*(c_i,\mathcal{C}) &= p^d_{\ell^*(c_i)}\left(b^d_{\ell^*(c_i)}-c_i+\frac{1}{1-\psi\delta}\left(\pi^*+\frac{(1-\psi)\delta}{1-\delta}\Lambda(\mathcal{C})\right)\right)+(1-p^d_{\ell^*}(c_i))\left(0+\frac{\delta}{1-\delta}\Lambda(\mathcal{C})\right)\\ &= p^d_{\ell^*(c_i)}\left(b^d_{\ell^*(c_i)}-c_i+\frac{\pi^*-\psi\delta\Lambda(\mathcal{C})}{1-\psi\delta}\right)+\frac{\delta}{1-\delta}\Lambda(\mathcal{C}).\end{aligned}$$

The partition $\mathcal{C}$ picks out L bids $\mathcal{B}(\mathcal{C})$ along with the exogenous reserve price R, and only these types and bids are relevant to deciding the optimal deviation. Since those bids and probabilities vary continuously in $\mathcal{C}$, the payoff from the best undercutting bid likewise varies continuously in $\mathcal{C}$, and $\Delta(c_i, \mathcal{C})$ is a continuous function in the partition and type. Then $\Delta^*(c_i, \mathcal{C})$ varies continuously in $\mathcal{C}$, and the Theorem of the Maximum implies the value function also varies continuously in type c_i, as well as δ.

Enforcement with variable δ. A particular round-bidding strategy $(\mathcal{C}, \mathcal{B}(\mathcal{C}))$ is enforceable if, for all c_i,

$$p_\ell(b_\ell - c_i) + \frac{\delta}{1-\delta}\Lambda(\mathcal{C}) \geq \Delta^*(c_i, \mathcal{C}) = p^d_{\ell^*(c_i)}\left(b^d_{\ell^*(c_i)} - c_i + \frac{\pi^* - \psi\delta\Lambda(\mathcal{C})}{1-\psi\delta}\right) + \frac{\delta}{1-\delta}\Lambda(\mathcal{C}).$$

Re-arranging the above inequality yields

$$p^d_{\ell^*(c_i)}\frac{\psi\delta\Lambda(\mathcal{C}) - \pi^*}{1-\psi\delta} \geq p^d_{\ell^*(c_i)}(b^d_{\ell^*(c_i)} - c_i) - p_\ell(b_\ell - c_i). \tag{10}$$

The right-hand side is non-negative, since the flow payoff of the best deviation must be weakly greater than the flow payoff of honest reporting, or no deviation would be tempting. A simple upper bound on the right-hand side is $R - c_i$. On the left-hand side it must be the case that $\Lambda(\mathcal{C}) > \pi^*$, or the proposed collusive agreement would not be profitable. Further simplification yields

$$\frac{\psi\delta\Lambda(\mathcal{C}) - \pi^*}{1-\psi\delta} \geq \max_{c_i} \frac{p^d_{\ell^*(c_i)}(b^d_{\ell^*(c_i)} - c_i) - p_\ell(b_\ell - c_i)}{p^d_{\ell^*(c_i)}} \equiv D, \tag{11}$$

or

$$\psi\delta > \frac{\pi^* + D}{\Lambda(\mathcal{C}) + D}.$$

Since the right-hand side is strictly less than one, there is a locus of pairs $(\Psi^*, \Delta^*) \subset [0,1]^2$ for which any pair $(\psi, \delta) > (\Psi^*, \Delta^*)$ will satisfy the inequality. So for δ sufficiently close to 1, there is a punishment probability $\psi(\delta)$ for which the proposed strategies constitute a sub-game perfect Nash equilibrium and any incentive compatible $(\mathcal{C}, \mathcal{B}(\mathcal{C}))$ is enforceable for δ.

Enforcement with fixed δ. For a fixed $\mathcal{C}$ and δ, enforceability is characterized by

$$\Delta^*(c_i, \mathcal{C}, \psi) =$$
$$\max_{x_1,\ldots,x_L} \sum_{\ell=1}^{L} x_\ell \left\{ p^d_\ell \left(b^d_\ell - c_i + \frac{1}{1-\psi\delta}\left(\pi^* + \frac{(1-\psi)\delta}{1-\delta}\Lambda(\mathcal{C})\right)\right) + (1-p^d_\ell)\left(0 + \frac{\delta}{1-\delta}\Lambda(\mathcal{C})\right)\right\} \tag{12}$$

and the compliance constraint

$$0 \geq \Delta^*(c_i, \mathcal{C}, \psi) - \left(p_\ell(b_\ell - c_i) + \frac{\delta}{1-\delta}\Lambda(\mathcal{C})\right) \tag{13}$$

for all $c_i \in [\underline{c}, R]$.

Since $\Delta^*(c_i, \mathcal{C}, \psi)$ is continuous in its arguments and $c_i \in [\underline{c}, R]$, the compliance constraints can be subsumed into a single non-linear inequality,

$$0 \geq \rho(\mathcal{C}, \psi) \equiv \max_{c_i \in [\underline{c}, R]} \left\{ \Delta^*(c_i, \mathcal{C}) - \left(p_\ell(b_\ell - c_i) + \frac{\delta}{1-\delta}\Lambda(\mathcal{C}) \right) \right\}.$$

Since $\rho(\mathcal{C}, \psi)$ is continuous in $\mathcal{C}$ and ψ and the inequality is weak, the set of grids $\mathcal{C} = \{c_1, ..., c_L\}$ that satisfy the constraint is a closed subset of the set of all grids, $\mathcal{G}$, and therefore compact.

Then the Cartel can solve the constrained maximization problem:

$$\max_{\psi, \mathcal{C}} \sum_{\ell=1}^{L+1} H(c_{\ell-1}, c_\ell) \mathbb{E}[\lambda(c) | c_{\ell-1} \leq c < c_\ell]$$

subject to

$$\begin{aligned} 0 &\geq \rho(\mathcal{C}, \psi) \\ \frac{H(c_{\ell-1}, c_\ell)}{N(G(c_\ell) - G(c_{\ell-1}))} &\geq \frac{H(c_\ell, c_{\ell+1})}{N(G(c_{\ell+1}) - G(c_\ell))}, \quad \ell = 1, ..., L. \end{aligned}$$

For a fixed δ and small L, there might be no feasible points, even if $\psi \to 1$ to maximize the severity of the punishment for deviating. If the Cartel members are too impatient, it must add additional bids to the grid to deter deviations, as in the case with a single reserve price bid R. However, the constraint set is never empty, since as L goes to infinity, the fully separating equilibrium is feasible and trivially enforceable. Therefore, a solution exists for the problem of maximizing Cartel profits for a fixed δ that generically involves multiple bid increments, since the above maximization problem entails maximization of a continuous function over a compact set.

Comparative statics with δ. Consider the correspondence Γ that maps discount factors $\delta \in (0, 1)$ into a set of incentive compatible and enforceable partitions of $\mathcal{C}$ that is a subset of the set of all grids, $\mathcal{G}$. The monotonicity constraints are a set of L non-negativity constraints that correspond to a non-empty subset of the set of all grids, $\mathcal{G}$, for all L. The subset of $\mathcal{G}$ on which the enforceability constraint is satisfied, however, might be empty for some L, as noted, but as $L \to \infty$, the separating equilibrium becomes a feasible point in $\mathcal{G}$ for which the enforceability constraint is satisfied, so there is always a solution to the Cartel's problem, and $\Gamma(\delta) \neq \varnothing$ for all $\delta \in (0, 1)$.

Recall Equation (11),

$$\frac{\psi\delta\Lambda(\mathcal{C}) - \pi^*}{1 - \psi\delta} \geq \max_{c_i} \frac{p^d_{\ell^*(c_i)}(b^d_{\ell^*(c_i)} - c_i) - p_\ell(b_\ell - c_i)}{p^d_{\ell^*(c_i)}}.$$

The right-hand side does not depend on δ or ψ, and the left-hand side is non-decreasing in δ and

ψ. Then if $\delta' > \delta$ and the constraint is satisfied for $\mathcal{C}$ at δ, the constraint is also satisfied at $\delta' > \delta$. That implies that any $\mathcal{C}$ that is enforceable for δ is enforceable for $\delta' > \delta$, and $\Gamma(\delta') \supseteq \Gamma(\delta)$. Since the Cartel prefers to minimize separation, it will only add additional bids to $\mathcal{B}$ if the enforcement constraint binds, and if $\delta' > \delta$, $\mathcal{C}^*(\delta')$ will have weakly fewer elements than $\mathcal{C}^*(\delta)$. $\square$

Proposition 4. *The probability of winning by making a report in* $[c_{\ell-1}, c_\ell)$ *or bidding* b_ℓ *is*

$$p(c_{\ell-1}, c_\ell) = \frac{H(c_{\ell-1}, c_\ell)}{N(G(c_\ell) - G(c_\ell))} = p_\ell$$

and

$$\frac{\partial H(c_{\ell-1}, c_\ell)}{\partial c_\ell} = N(1 - G(c_\ell))^{N-1} g(c_\ell).$$

Proof. Define

$$H(c_{\ell-1}, c_\ell) = \sum_{j=1}^{N} \frac{N!}{j!(N-j)!} (G(c_\ell) - G(c_{\ell-1}))^j (1 - G(c_\ell))^{N-j},$$

as appears in the proof of Proposition 2, as well as

$$H_k(c_{\ell-1}, c_\ell) = \sum_{j=k}^{N} \frac{N!}{j!(N-j)!} (G(c_\ell) - G(c_{\ell-1}))^j (1 - G(c_\ell))^{N-j},$$

so that $H_1(c_{\ell-1}, c_\ell) = H(c_{\ell-1}, c_\ell)$. These are similar to the distributions of order statistics, so some of the same recursion relations are satisfied, namely,

$$H_k(c_{\ell-1}, c_\ell) = H_{k-1}(c_{\ell-1}, c_\ell) - \frac{N!}{(k-1)!(N-k+1)!} (G(c_\ell) - G(c_{\ell-1}))^{k-1} (1 - G(c_\ell))^{N-k+1}. \quad (14)$$

Define $h_k(c_{\ell-1}, c_\ell) = \partial H_k(c_{\ell-1}, c_\ell)/\partial c_\ell$. We prove by induction that

$$h_k(c_{\ell-1}, c_\ell) = \frac{N!}{(k-1)!(N-k)!} (G(c_\ell) - G(c_{\ell-1}))^{k-1} (1 - G(c_\ell))^{N-k} g(c_\ell). \quad (15)$$

The base case is

$$H_N(c_{\ell-1}, c_\ell) = (G(c_\ell) - G(c_{\ell-1}))^N (1 - G(c_\ell))^0$$

with

$$h_N(c_{\ell-1}, c_\ell) = N(G(c_\ell) - G(c_{\ell-1}))^{N-1} g(c_\ell),$$

so the result is true for the $k = N$ case.

To complete the induction step, suppose that Equation (15) holds at the k^{th} step:

$$h_k(c_{\ell-1}, c_\ell) = \frac{N!}{(k-1)!(N-k)!} (G(c_\ell) - G(c_{\ell-1}))^{k-1} (1 - G(c_\ell))^{N-k} g(c_\ell). \quad (16)$$

Differentiating the recurrence relation, Equation (14), yields:

$$\begin{aligned}
\frac{\partial H_k(c_{\ell-1}, c_\ell)}{\partial c_\ell} &= \frac{\partial H_{k-1}(c_{\ell-1}, c_\ell)}{\partial c_\ell} \\
&\quad - \left[\frac{N!}{(k-1)!(N-k+1)!}(k-1)(G(c_\ell) - G(c_{\ell-1}))^{k-2}(1-G(c_\ell))^{N-k+1} g(c_\ell)\right] \\
&\quad + \left[\frac{N!}{(k-1)!(N-k+1)!}(N-k)(G(c_\ell) - G(c_{\ell-1}))^{k-1}(1-G(c_\ell))^{N-k} g(c_\ell)\right].
\end{aligned}$$

The left-hand side and third line cancel out by the induction hypothesis, leaving

$$h_{k-1}(c_{\ell-1}, c_\ell) = \frac{N!}{(((k-1)-1)!(N-(k-1))!}(G(c_\ell) - G(c_{\ell-1}))^{(k-1)-1}(1-G(c_\ell))^{N-(k-1)} g(c_\ell)$$

which is the desired expression for h_{k-1}.

Finally, note that $H(c_{\ell-1}, c_\ell) = H_1(c_{\ell-1}, c_\ell) = H_1(c_{\ell-1}, c_{\ell-1}) + \int_{c_{\ell-1}}^{c_\ell} h_1(c_{\ell-1}, z)dz$. Then, by the Fundamental Theorem of Calculus,

$$\begin{aligned}
H(c_{\ell-1}, c_\ell) &= 0 + \int_{c_{\ell-1}}^{c_\ell} \frac{N!}{0!(N-1)!}(G(z) - G(c_{\ell-1}))^0 (1-G(z))^{N-1} g(z)dz \\
&= \int_{c_{\ell-1}}^{c_\ell} N(1-G(z))^{N-1} g(z)dz,
\end{aligned}$$

as in Proposition 2. $\square$