Detection of Collusive Networks in E-Procurement [*]

Bruno Baránek[†] Leon Musolff[‡] Vítězslav Titl[§]

April 12, 2024

Abstract

We develop a method for detecting cartels in multistage auctions. Our approach allows a firm to be collusive when facing members of its cartel yet competitive when facing others. Intuitively, as initial bids are shaded, close initial bids not only imply similar costs but also provide an incentive to undercut. We detect firm pairs that ignore this incentive when facing each other. Our algorithm predicts Ukraine's Antimonopoly Committee sanctions, yet uncovers additional collusion: 2,371 collusive firms participate in 19% of auctions, increasing costs by 2.12%. Cartels typically comprise just two members, and members often share the same ZIP code.

Keywords: Public procurement, Collusion, Online markets

[*]We gratefully acknowledge insightful comments from Sylvain Chassang, Francesco Decarolis, Adriano De Leverano, Thomas Fujiwara, Elena Fumagalli, Benny Geys, Leonardo Giuffrida, Andres Gonzalez Lira, Margaryta Klymak, Timothy Hubbard, Kate Ho, Adam Kapor, Jakub Kastl, Kei Kawai, Andres Gonzalez Lira, Deni Mazrekaj, Jonas Müller-Gastell, Dina Pomeranz and suggestions from participants of seminars at KU Leuven, Princeton, Utrecht University and IIOC 2023. We would like to also thank the Weiss Fund for financial support (award number PRN-004) and Bogdan Salagub for invaluable research assistance. Vitezslav would like to thank the Department of Economics at Princeton University for its hospitality and the Czech Science Foundation – GACR (project number 21-15480M "The Role of Institutional Factors and Information in Public Procurement Markets") for financial support. Bruno Baranek conducted this research while at Princeton.

[†]Electronic address: `bbaranek@princeton.edu`

[‡]Electronic address: `lmusolff@wharton.upenn.edu`; University of Pennsylvania.

[§]Electronic address: `v.titl@uu.nl`; Utrecht University School of Economics, Utrecht University; Department of Economics, Faculty of Law, Charles University

1 Introduction

With public procurement accounting for 12% of global GDP (Bosio et al., 2022), collusion in procurement auctions is a primary concern for academics and practitioners. While there has been significant progress in studying collusion on the procurement market, previous research has often treated firms as entirely collusive or competitive, which is likely an oversimplification. In reality, firms' behavior can vary depending on the counter-party. We propose a new method for detecting cartels in multistage auctions. This method allows a firm to be collusive when facing members of its cartel yet competitive when facing others.

To detect collusive behavior, we analyze deviations from competitive equilibrium play. In particular, as initial bids are shaded in equilibrium, close initial bids not only imply similar costs but also provide an incentive to undercut your opponent in subsequent bidding rounds. We detect firm pairs that do not update their initial bids when facing each other even though their costs would make updating profitable. By focusing on pairs of firms rather than individual firms, we allow a firm to vary its behavior depending on whether it faces a member of its own cartel or an outsider.

We apply this methodological insight to studying firm-pair level collusion in the Ukrainian e-procurement market. We detect 2,371 firms engaging in collusive behavior. These firms participate in 19% of public procurement auctions. Interestingly, the average colluding firm colludes only in 60% of the auctions it participates in, confirming our hypothesis that firm behavior varies across auctions. Furthermore, auctions with collusive pair participation result in 2.12% higher prices, suggesting collusion is an important problem in Ukraine.

We begin by presenting some stylized facts about the Ukrainian procurement market, which highlight a potential lack of competition. A striking feature of this market is the low level of firm activity during online auctions. Specifically, 52.55% of initial bids remain unchanged throughout the auction process. This pattern is unexpected, given the significant profits firms could forgo by not updating their bids. In fact, firms that choose to revise their bid after an initial losing round win the contract 25.82% of the time. Interestingly, this apparent reluctance to undercut competitors aligns with what one would expect in a cartel arrangement, where members agree on project distribution beforehand and submit initial bids primarily as cover to avoid detection. Indeed, our data shows a notable pattern of close initial bids frequently followed by a lack of activity, further supporting this hypothesis.

Our novel collusion detection algorithm confirms our suspicion of collusive behavior. We verify the algorithm's accuracy on a sample of 752 firms penalized for collusion. These firms participated in 20,340 tenders, which correspond to 4.78% of the market's total value. While our algorithm's 'collusive' label is also predictive of whether a firm pair was penalized ($t = -14.84$), we identify many more colluders than the Antimonopoly Committee of Ukraine: we find 2,371 unique firms participating in collusion. One of the significant contributions of this paper is a method that identifies whole bidding cartels and not only whether specific firms collude or whether there is collusion on the entire market. We show that the 2,371 firms are part of 1,919 distinct firm pairs and 707 cartels (a cartel is a connected component of the graph in which collusive firm pairs are linked). A majority of these cartels (496) have only two members.

The Ukrainian procurement market is a fascinating laboratory for studying collusion for three reasons. First, Ukraine has a modern procurement market, operating entirely on an electronic platform that allows unprecedented data access and transparency. Detailed information about each auction, including each bid and updates, is available online. This data can be very useful for public oversight but can also facilitate collusion (Albæk et al., 1997). Second, the market has suffered from large-scale issues with collusion. Between 11/14/16 and 9/30/19, 1,042 firms have been penalized for collusive conduct.[1] Third, the multi-round auction design allows us to propose a novel framework for the identification of collusion. Multi-round auctions are common in public procurement (Ashkenazi-Golan et al., 2023).[2] Large-scale infrastructure projects are often procured in multiple-round auctions (Sauvet-Goichon, 2007; Merrifield et al., 2002). Similar multi-round auction mechanisms are also used in real estate auctions, the U.S. government timber rights sales, and art auctions

[1]Note not all firms penalized for collusion appear in our cleaned bidding data.

[2]According to Ashkenazi-Golan et al. (2023), The Federal Transit Administration recommends a two stage auction in its manual and similar two stage auction design is also used on the procurement in the UK and France (Noumba Um and Dinghem, 2005).

(Engelbrecht-Wiggans, 1988). The exact auction design as in Ukraine (with three stages[3]) is used in other post-soviet countries, including Georgia, Moldova, and Kyrgyzstan.

The auction mechanism proceeds as follows. First, interested bidders submit bids simultaneously. Subsequently, all bidders can access an online auction where they compete for the contract by lowering their bids over several rounds. In these rounds, bidding is sequential in an order determined by the previous round's ranking of bids. The bidder with the highest initial bid starts the bidding in the online auction. This setup can be understood as a mechanism in which bidders initially bid for the order in which they will submit a final bid. There is an advantage for the last bidder as nobody can react to her bid. Consequently, bidders are incentivized to submit low initial bids. However, as bids can only be *lowered* in the updating rounds, there is no incentive to submit arbitrarily low bids in any round.

We show below that if there is a small probability that the current winner fails to update her bid,[4] then it is optimal for the bidder with the second-lowest bid to undercut the current winning bid by a small ϵ if he can do so while bidding above his costs. Hence, the (repeated) lack of undercutting within firm pairs that submit similar initial bids signals that firms do not behave competitively. Drawing from these insights and building upon our formal theoretical model, we develop a novel structural algorithm to detect collusion and identify bidding rings within public procurement. Our algorithm examines each pair of firms in turn. Using our theoretical model, we compute the probability with which a given firm should underbid a particular competitor. Subsequently, we identify collusive relationships as statistically unlikely deviations from equilibrium play. Once all collusive relationships are identified, we can reconstruct the collusive network and pinpoint the exact location of bidding rings.

Our paper contributes to the empirical literature on collusion in auctions, building upon seminal work by Porter and Zona (1993, 1999) that identify collusive behavior by highlighting that observed collusive bids are inconsistent with predictions of a competitive equilibrium, where bidders independently optimize their bids. Porter and Zona (1993) showed that bids of competitive bidders in the construction procurement market come from a distribution

[3]Our detection algorithm is not dependent on the number of stages.

[4]We observe such bid failures in our data.

that can be rationalized by a particular cost structure and competitive bidding. By contrast, cartel participants would submit phony bids that could not be rationalized like this. Porter and Zona (1999) then analyze bids in the Ohio school market and find that potential cartel members seem to differ in both their participation decisions and size of bids.

More recent papers build on this literature by providing more complex identification strategies for collusion detection (Bajari and Ye, 2003; Chassang et al., 2022; Kawai et al., 2023). Kawai and Nakabayashi (2022) study a related setting where the authors leverage a specific auction design in Japanese public procurement where rebids occur if all bids do not initially meet a secret reserve price. The authors show that the re-bidding behavior of some firms is inconsistent with a competitive equilibrium and suggests the presence of large-scale collusion. The paper proposes a firm-specific test for collusive bidding patterns. Next, Kawai et al. (2022) and Kawai et al. (2023) leverage a regression discontinuity approach that shows how small differences in bids can often be rationalized by collusion rather than a competitive alternative explanation stemming from cost differences. Kawai et al. (2023) analyze procurement in five countries and find significant evidence of collusion in one. Conley and Decarolis (2016) devise a test for collusion in average bid procurement auctions. In such auctions not the lowest bidder but the one closest to an average wins the contract. This design creates incentives to submit multiple bids to shift the distribution of the bids which the authors use to detect collusion.

Another stream of literature about collusion in auctions studies incentives and internal design within a cartel. Chassang and Ortner (2019) demonstrate that bidding constraints, such as minimum prices, can weaken cartels, highlighting the importance of institutional rules. Asker (2010) leverages unusually rich data about the inner organization of a cartel to explain the incentives and coordination of bidding among a cartel in the stamp market. Asker (2010) proposes a structural model of bidding to evaluate the inefficiency and damages caused by the cartel. Similarly, Pesendorfer (2000) investigates the internal organization of cartels that allocate contracts or use side payment internally and designate a winner that bids competitively against outsiders. Hyytinen et al. (2018) show in a sample of firms in the

Finnish manufacturing industries from 1951 to 1990 that cartels were persistent and that, when collusion was legal, almost all industries were cartelized. Finally, Clarka et al. (2020) show that investigating cartels in Montreal's asphalt industry led to higher participation in procurement auctions and decreased prices.

This paper's main contribution is an empirical test based on a structural model that can identify networks of collusive firms. We explicitly allow firms to be collusive in some settings but not in others and then identify which firm pairs collude, allowing us to reconstruct the collusive network. This detection is made possible by a microeconomic equilibrium model and unusually detailed data on collusive behavior that helps us validate our econometric test. To the best of our knowledge, this paper is the first study that presents a methodology for detecting collusive *networks*. By contrast, Porter and Zona (1993) test whether there is collusion on *the whole market* and Kawai and Nakabayashi (2022) test whether *a particular firm* behaves collusively. Neither isolates which firms are involved in a given cartel. Finally, Bajari and Ye (2003) are closer to our paper as they consider bidder pairs. However, Bajari and Ye (2003) identify collusive pairs differently (exploiting exchangeability) and, with only 25 firms, have insufficient data to speak to the network structure of collusion.

This paper's research contributes to the literature studying inefficiencies in public procurement markets, particularly those arising from illicit activities. This literature documents significant inefficiencies attributed to corruption (for instance Decarolis et al., 2024; Kang and Miller, 2020). Mironov and Zhuravskaya (2016) document an increase in the illegal outflow of cash from firms with revenues from procurement around elections and no such increase for firms without revenue from procurement. The only plausible explanation of the outflow – offered in the paper – is bribery and corruption in public procurement. Schoenherr (2019) finds that political connections cause misallocation of procurement contracts towards connected firms, while Baranek and Titl (2024) and Titl and Geys (2019) show that contracts supplied by politically connected firms are over-priced (without compensating differences in quality). Compared to corruption losses in the Czech Republic (Baranek and Titl, 2024), collusion appears less costly per auction in Ukraine (2% vs 6% overpricing) but more prevalent (up to 20% vs 6-10% of auctions affected), though these comparisons should be interpreted cautiously due to the differing settings.

The remainder of the paper is structured as follows. Section 2 describes the Ukrainian procurement market. In Section 3, we solve for the equilibrium of the sequential auction. Section 4 shows summary statistics and suggestive evidence that the firms' behavior is potentially collusive. In Section 5, we introduce a novel method to identify collusive rings and propose a statistical test for collusion. We then validate our method using a sample of firms penalized for collusion and proceed to characterize the networks of colluding firms. Finally, section 6 concludes.

2 Description of the Market

Ukraine has faced corruption and collusion in public procurement since its independence in 1990. After the Euromaidan revolution in 2014, a group of volunteers started the ProZorro platform in February 2015 to tackle these issues. They successfully promoted this fully online tendering system, and the platform became compulsory for all public entities in 2016 after ownership transferred to the state of Ukraine. At its core, ProZorro is a unified central database of all public procurement projects conducted in Ukraine, and an API for interacting with this database. The data from this platform will serve as the primary dataset for our analysis.

2.1 Exact Tender Procedure

In Ukraine, small ("below-threshold") purchases[5] can be conducted without an online auction but are still entered on ProZorro. Larger purchases ("above-threshold") generally have to be completed as open tenders. Our analysis will focus on competitive tenders both below- and above-threshold and emphasize a critical difference between the two types of tenders: below-threshold agreements do not require an auction in the first place; they can be awarded

[5]For the exact thresholds, see Supreme Council of Ukraine (2015); the smallest threshold is $\approx$ \$7,000.

to the sole bidder should only one bidder participate in the sale. By contrast, above-threshold auctions must be repeated or canceled altogether if only one bidder exists.

We now describe the open tendering procedure for the above-threshold contracts. The tender begins with the procuring entity uploading documentation for the tender to ProZorro, at which point the period of proposal submission begins and lasts for at least 15 calendar days. After this period, the tender is automatically canceled if only one proposal has been submitted. If there are multiple proposals, the system automatically schedules and runs an online auction, which we discuss below. During the online auction, the bidders do not yet have access to each others' documents. Furthermore, while they are informed of the number of opposing bidders when the online auction starts, they remain unaware of their identity and specific proposals until it ends.

2.2 ProZorro Auction

The tendering process's critical element is the online auction, during which bidders compete for the right to complete the contract. First, initial bids are submitted together with technical proposals. Initial bids can be understood to be submitted 'simultaneously' because bidders are not aware of each others' bids at this stage. Second, bidders enter the online auction during which bids (but not bidder identities or technical proposals) are revealed:

1. Bidders are ordered in descending order according to their initial bids, and the first updating round begins:
   (a) The bidder with the highest initial bid goes first, observes all initial bids and can update her bid. However, she can only lower her bid.
   (b) After the first bidder moves, the bidder with the second-highest initial bid observes all bids, i.e., initial bids and the update by the originally highest bidder. This (second highest) bidder is again given a chance to lower his original bid.
   (c) All the bidders move sequentially until the bidder with the lowest initial bid has chosen whether to update her bid.

2. Bidders are ordered based on the size of their updated bids and again move sequentially.

3. Finally, there is a third round of bidding in which bidders are ordered based on their bids from the second round. The bidder with the lowest bid at the end of this round wins and becomes the vendor for this project.

This mechanism emulates a sequential Bertrand game. As the last-mover is advantaged, bidders are incentivized to submit low initial bids. However, as bids can only be lowered in the updating rounds, there is no incentive to submit arbitrarily low bids.

3 Model and Equilibrium

We now discuss the intuition behind the equilibrium of the ProZorro Auction, with details and proof relegated to the online appendix. The model's primary goal is to establish that – just like in a standard first-price auction – bidders have an incentive to shade their bids relative to their true cost in a ProZorro auction. As long as bids are shaded, if a firm finds itself bidding the same as its opponent, there is still room to undercut while retaining a strictly positive profit. To the extent that bid shading is a feature of many auction models, our empirical results below do not rely on firms playing the exact equilibrium we derive in this section; nevertheless, as the auction mechanism is unusual, it is important to verify that bid shading follows naturally from some simple assumptions in this mechanism.

Consider a simplified auction with two players and one updating round. The timing is as follows:

1. Bidders submit their initial bids simultaneously.
2. The initial 'loser' (the agent that submitted the higher bid) can update his bid.
3. The initial 'winner' can update her bid.

Note that bidders can only update their bids downwards, i.e., initial bids are not cheap talk.

As the timing clarifies, the equilibrium will hinge on the amount of information revealed in the initial stage of the auction. Therefore, we restrict attention to equilibria in which the initial bid is perfectly revealing (i.e., agents are not randomizing). In such equilibria, agents are fully informed about each other's cost types after the initial stage. This information generates a potential multiplicity issue: the initial loser may realize that he will lose the overall auction no matter what he bids. To resolve this multiplicity, we introduce an (arbitrarily) small probability that any given bid update is not successfully submitted. This probability captures the idea that if you know you will lose if your opponent reacts, you may as well bid in such a way as to maximize your surplus if, for some reason, your opponent fails to respond.

Our assumptions imply that when submitting the last bid in the auction, the initial winner will beat the current standing bid by the minimum amount necessary[6] if doing so is feasible. Before this, the initial loser will predict what sort of bids the initial winner can beat: if there are bids that she cannot beat and that still give the initial loser a positive surplus, he will make the highest bid satisfying these criteria. If there are none, he will update his bid to the current winning bid (hoping that she will fail to respond). More rigorously, if $b(\cdot)$ is the equilibrium bidding function in the initial round, the payoff to type c_1 from pretending to be type $\tilde{c}$ in this round is given by

$$\begin{aligned} V^{PZ}(\tilde{c}) \;=\; & \mathbb{P}\Big(b(\tilde{c}) < b(c_2) \cap c_1 < c_2\Big)\Big(b(\tilde{c}) - c_1\Big) + \\ & \mathbb{P}\Big(b(\tilde{c}) < b(c_2) \cap c_1 > c_2\Big)\Big(b(\tilde{c}) - c_1\Big) + \\ & \mathbb{P}\Big(b(\tilde{c}) > b(c_2) \cap c_1 < c_2\Big)\mathbb{E}[\min\{c_2, b(\tilde{c})\} - c_1 | c_1 < c_2, b(\tilde{c}) > b(c_2)] + \\ & \mathbb{P}\Big(b(\tilde{c}) > b(c_2) \cap c_1 > c_2\Big) \times 0 \end{aligned}$$

The four lines of this expression correspond to the four possible cases: the agent could pretend to be strong and be strong (first line), pretend to be strong and be weak (second line), pretend to be weak and be strong (third line) or pretend to be weak and be weak (fourth line). The first two lines combine to yield the payoffs from a first-price auction in which each bidder bids according to $b(\cdot)$. The last two lines can be related to the expected payoff

[6]Our formal results below assume that the minimum bid decrement is arbitrarily small.

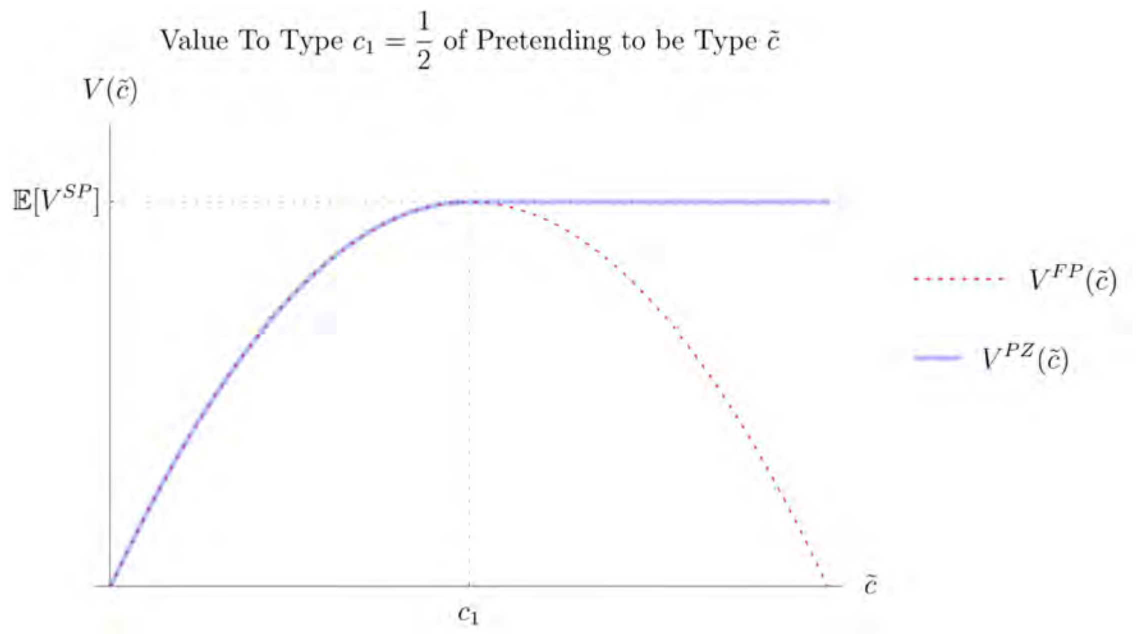

Figure 1: The Equilibrium of the ProZorro Auction

Notes: The expected utility to an agent of a given type from participating in the ProZorro auction reaches its peak at the same time as that of participating in a first-price auction, but the second-price component ensures it never drops from this level.

from a second price auction so that we can write the overall payoff as

$$V^{PZ}(\tilde{c}) = V^{FP}(\tilde{c}) + \mathbb{P}(b(\tilde{c}) > b(c_2))\mathbb{E}[V^{SP}|\underline{c}_{-1} < \tilde{c}],$$

where we use $\underline{c}_{-1} := \min_{j \neq 1} c_j$ to emphasize that this way of expressing the payoffs does not rely on a specific number of players. Indeed, we have the following result:

Proposition 1. *In any equilibrium in which initial bids are given by some strictly increasing $b(\cdot)$, the expected payoff from pretending to be type $\tilde{c}$ is given by $V^{PZ}(\tilde{c})$ no matter the number of updating rounds or number of players.*

We illustrate $V^{PZ}(\tilde{c})$ in Figure 1 for the case of $c_i \sim U[0, 1]$. If $\tilde{c} < c_1$, the expected value from a second price auction conditional on $c_2 < \tilde{c} < c_1$ is zero; hence, $V^{PZ}(\tilde{c}) = V^{FP}(\tilde{c})$ to the left of $\tilde{c} = c_1$. Furthermore, $V^{FP}(1) = 0$. Thus, $V^{PZ}(1) = \mathbb{E}[V^{SP}]$. However, the expected rent that bidders earn in a second-price auction is exactly the expected rent they earn in a first-price auction when pretending to be their true type. Thus, $V^{PZ}(1) = V^{PZ}(c_1)$. It turns out that the effects of decreasing rent from the first-price component of the auction and increasing rent from the second-price component of the auction exactly cancel and hence $V^{PZ}(\cdot)$ is flat to the right of $\tilde{c}$. A more formal version of this heuristic argument in the online appendix allows us to conclude that if $c_i \sim F(\cdot)$ (with associated p.d.f. $f(\cdot)$),

Proposition 2. *The ProZorro auction (with $k \geq 1$ rounds and $n \geq 1$ players) has a unique PBE in which initial bids are given by a strictly increasing $b(\cdot)$. In this equilibrium,*

$$b(c) = \frac{1}{[1 - F(c)]^{n-1}} \int_{c}^{c_{max}} s(n-1)f(s)[1 - F(s)]^{n-2} ds$$

and bids are decreased by the arbitrarily small minimum bid decrement whenever doing so is possible without bidding below one's own cost.

Thus, the initial bids in the ProZorro auction are generated by the same bidding function as in a first-price auction and hence, crucially, are shaded away from costs[7].

4 Data and Summary Statistics

We use data from three data sources, each covering the period of August 2016 to August 2019. First, we use publicly available tender-level procurement data from ProZorro. This data contains detailed information about the final price, industry code, delivering firm, procuring authority, and the reserve price for each contract procured by any public entity in Ukraine. Second, we scrape the auction platform employed by ProZorro to complement these covariates with detailed bid-level data, including bids and bidders' identities at all stages of each online auction. Third, we obtain data about firms penalized for collusive conduct in public procurement from the Antimonopoly Committee of Ukraine.

[7] To be clear, the bidding function is only the same in the limit as the minimum bid decrement goes to zero – for any positive minimum bid decrement, initial bids have to lie (slightly) above the bids in a first-price auction to account for the subsequent decreases by the minimum bid decrement.

We provide summary statistics for the tender-level data in Table 1. We draw the readers' attention to a few key stylized facts that become apparent from this table.

There is low competition in the market. Recall that above-threshold contracts are cancelled if there are less than two participants. Nevertheless, the average number of participants is just 2.47 (median: 2). While the Ukrainian data contains a large number of small contracts (frequently not reported in other countries), the low number of participants is not limited to small contracts. Indeed, competition is comparable for big contracts above 25,000,000 UAH (about 1,000,000 USD) with an average of 2.59 bidders and an (unchanged) median of 2.

Not updating leaves money on the table. Of all non-initial bids, only 32% are updates on the bidder's previous bid; indeed, there are zero bid updates in 46% of all online auctions. For big contracts, the fraction of auctions without competition during the online phase is even higher at 50%. This finding contrasts with our reasoning about the optimal behavior of participants. Indeed, the fraction of auctions with undercutting by the initial loser in which the initial loser eventually wins is 25.82% – suggestive evidence that there are strong incentives for undercutting the initial winner.

Realized updates are small. Our equilibrium discussion predicts that most updates equal the minimum bid decrement. We restrict attention to updates (i.e., situations in which a bidder lowered their bid when compared to their bid in the previous round) and define the relative step size as $\frac{b^{r+1}_{\ell(i,r)} - b^r_{w(i,r)}}{b^r_{\ell(i,r)}}$ where $b^r_{\ell(i,r)}$ is the bid of a firm currently losing an auction i in stage r and $b^r_{w(i,r)}$ is the bid of the current winner of auction i in stage r. Note that a 'positive update' means a bidder lowered their bid, but not by enough to beat the current standing winner; such 'ineffective' updates make up 15% of all updates. The mass of bid updates is very close to zero with a median update of -0.22%. The median 'effective' update is -0.43%.

Penalization for collusion is a widespread phenomenon. The Antimonopoly Committee of Ukraine penalized 1,042 firms for collusion between 11/14/16 and 9/30/19. These companies participated in 20,340 tenders accounting for 6.21% of procurement contracts and

	Panel A: All Contracts			Panel B: Big Contracts		
Variable	Mean	Median	Std. Dev.	Mean	Median	Std. Dev.
Bids						
Bid (in million UAH)	2.23	0.13	25.97	101.31	47.98	201.42
Normalized Bid	0.93	0.96	0.09	0.97	0.99	0.06
Was Bid an Update?	0.32	0.00	0.46	0.25	0.00	0.43
Update Size (x100)	-0.01	-0.22	3.38	0.30	-0.08	2.94
Is Update Effective?	0.85	1.00	0.36	0.80	1.00	0.40
Update Size (x100, Effective Only)	-1.04	-0.43	1.74	-0.65	-0.19	1.34
AMCU-Penalized Bidder?	0.03	0.00	0.17	0.06	0.00	0.23
N		3,054,189			40,467	
Tenders						
Mean Normalized Bid	0.93	0.95	0.08	0.97	0.99	0.05
Min. Normalized Bid	0.90	0.92	0.10	0.95	0.98	0.06
Number of Bidders	2.47	2.00	1.01	2.59	2.00	4.63
Any AMCU-Sanctioned Bidders?	0.07	0.00	0.25	0.11	0.00	0.31
Any Updates?	0.54	1.00	0.50	0.49	0.00	0.50
N		308,567			4,018	

Table 1: Summary of Procurement Data

Notes: This table provides summary statistics for bids (top) and tenders (bottom) for all contracts (Panel A) and large contracts (Panel B). We consider any contract with engineer's estimate above 25,000,000 UAH to be a large contract. There are N=3,054,189 bids in our data, with 40,467 of these belonging to large contracts. Note that the bid data includes multiple bids by each bidder for a given tender as the online auction proceeds over several stages. There are 308,567 tenders in our data, with 4,018 of these tenders being for large contracts. The normalized bid variable normalizes bids (in UAH) by dividing them by the engineering estimate (also in UAH). The update variables are only defined for non-initial bids.

4.78% of the total value. Firms penalized for collusion update their bids less: only 26.76% of their non-initial bids are updates. Furthermore, contracts delivered by them are more expensive: the average contract provided by a penalized firm costs 91.29% of the estimate (compare 90.22% for non-penalized firms).

Overall, these statistics point out a lack of entrants and of competitive behavior conditional on entry. These factors, together with the detailed prosecution data, hint at collusion as a crucial issue facing the market.

Figure 2: Suspiciously Close Bids and No Undercutting

Notes: This screenshot shows the initial bids in an auction for building a gymnasium. Two firms participate and submit essentially identical initial bids (the bids differ by 0.81UAH, or about $0.03.). Still, neither firm updates its bids in any of the three subsequent bidding rounds.

5 Empirical Model

Collusion appears to be a widespread phenomenon in Ukraine's market for public procurement. We now provide suggestive evidence of collusion based on bidding patterns that are hard to explain in competitive equilibrium; then, we present a formal test that isolates colluding firm pairs.

5.1 Descriptive Evidence

To begin with, recall that our discussion of the equilibrium revealed that initial bids are strictly increasing in the underlying costs. So if two bidders submit sufficiently close initial bids, their costs should also be close. Thus, we should see more undercutting in auctions where the initial bids are similar.

With this reasoning in mind, a brief review of tenders on ProZorro reveals suggestive evidence that collusion is an issue on the market. In Figure 2, we exhibit a screenshot from the online platform showing bidding on a tender in which two bidders submitted virtually identical bids: 31,864,899.19 UAH[8] and 31,864,900.00 UAH. The losing bidder did not update her bid in the following rounds even though the difference in bids was about 0.81 UAH, i.e., 3 cents. This behavior is suspicious. Below, we use the fact that many firms on the market behave in such a noncompetitive way when competing with a particular firm to identify colluding pairs of firms.

[8] Roughly 1,138,000 USD.

We study whether the data follow this suspicious pattern more generally in Figure 3a. The figure plots the fraction of auctions in which the initial loser updates his bid (vertical axis) against the difference of initial bids (horizontal axis). A competitive model predicts a declining function. This prediction partially holds on the left side of Figure 3a, in which we analyze below-threshold auctions. However, we observe the opposite on the right side: close initial bids are associated with decreased competition as measured by the likelihood of bid updating. One possible explanation for this contrast are the increased incentives for collusive bidding above the threshold. Recall above-threshold tenders are automatically canceled if only one bid exists. Thus, collusion on below-threshold contracts could take the form of market splitting; by contrast, above-threshold contracts require the submission of bids by at least two cartel members.

To further demonstrate this pattern, we plot the probability of having the lowest final bid in the auction (vertical axis) against the initial difference of bids (horizontal axis) in Figure 3b. The second-mover advantage of the initial winner explains the discontinuity at zero. Importantly, we document that for above-threshold contracts, submitting a bid just slightly above the opponent leads to a lower realized probability of winning than submitting a bid much above the opponent (i.e., the graph is increasing near the discontinuity).

Such patterns are inconsistent with a competitive equilibrium where close cost draws imply increased competition. However they are consistent with a cartel where bidders agree on the winner ex-ante. Cartel members submit bids close to each other (perhaps to make the auction seem competitive) and never update their bids. We now show that a simple collusion model in this vein can reproduce the patterns of Figure 3.

We will conduct simulations for auctions with precisely two players. There are two possibilities: either the auction is competitive, or there is a collusive pair participating. In the competitive situation, the bidders act in line with their equilibrium strategies. Collusive pairs, on the other hand, are modeled in the following way. The pair designates a winner ex-ante; the designated loser just exists to submit a phantom bid, thereby causing the auction to go ahead (and perhaps ensuring regulators do not investigate the relevant market). Thus, the

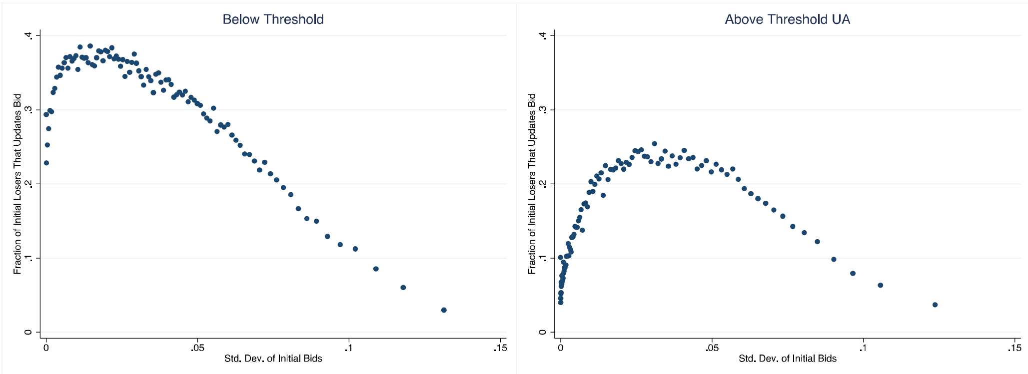

(a) Fraction of Initial Losers That Updates Bid (Binscatter)

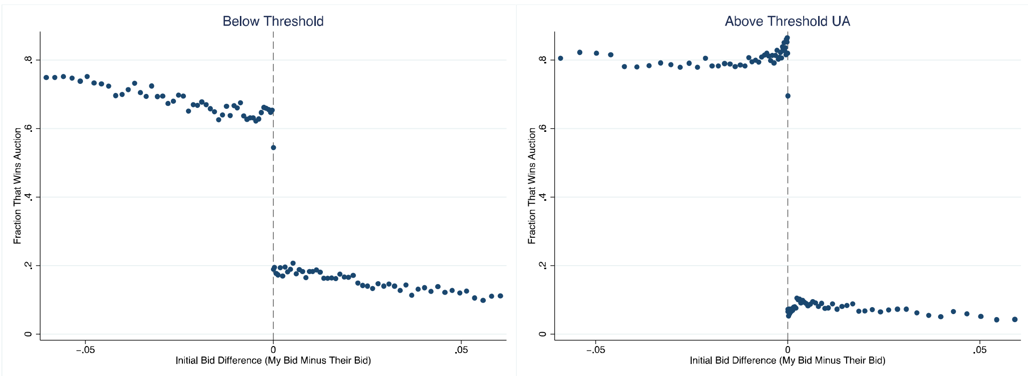

(b) Probability Own Bid Should Win Given Initial Bid Difference (Binscatter)

Figure 3: Concerning Patterns in Bidding Data

Notes: Panel (a) is a binscatter of the fraction of initial losers that updates their bid against the standard deviation of initial bids. Panel (b) is a binscatter of the probability of having the lowest standing bid at the end of the auction against the initial bid difference. For simplicity, we restrict the sample of auctions for Panel (b) to those with precisely two bidders. Furthermore, for easier visibility, we have restricted the sample to an absolute initial bid difference below the 75th percentile.

designated loser submits a bid $(1 + a)b_{c,w}$, where a is a small constant and $b_{c,w}$ is the bid of the designated winner. The selected winner submits a bid above the equilibrium competitive bid. Thus, collusive bidders do not undercut each other as they have no incentive to do so.

We present results of a simulation where bidders have uniform costs, i.e., $c_i \sim U[0, 1]$. Furthermore, we assume phantom bid multipliers are also uniformly drawn, i.e., we assume $a \sim U[0, 1/20]$. While competitive bids follow $b(c) = \frac{1}{2} + \frac{c}{2}$ in equilibrium, we assume that the designated winner's collusive bids instead are generated by collusive bidding function $\tilde{b}(c) = \frac{2}{3} + \frac{c}{3}$. Collusive bidders never undercut each other, but competitive bidders behave according to a smoothed version of the equilibrium we derive above: they undercut with probability $\Phi(0.5)$ if there is a profitable bid weakly below their current bid minus the minimum bid decrement of 0.01, and they undercut with probability $\Phi(-0.5)$ if there is no such bid. We simulate 600 bidders, each of which interacts with all other bidders in exactly two auctions. Of all bidding pairs, 30% are in a cartel.

Using data from this model we reproduce Figure 3b and 3a. We see that patterns in these figures can be explained with this simple model of collusion. The above-threshold auctions have a higher incentive for collusion and are comparable to the simulation of collusive behavior. By contrast, below-threshold auctions are comparable to the competitive ones.

As collusive agreements can take on many different forms and our model assumes one particular (if natural) form, this evidence should be interpreted as suggestive only. However, it is striking that we can explain all anomalies in our data with such a simple model.

5.2 Testing for Collusion

We now *(i)* develop a market-wide test for collusion, *(ii)* verify our detection on a sample of firm pairs fined by the Antimonopoly Committee of Ukraine, *(iii)* present a statistical test of collusive behavior for individual pairs of firms, and *(iv)* discuss characteristics of the detected cartels.

To begin with, define an indicator u_i that equals one if the initial loser $\ell(i)$ updates his bid against initial winner $w(i)$ in any round of auction i and equals zero otherwise. From the

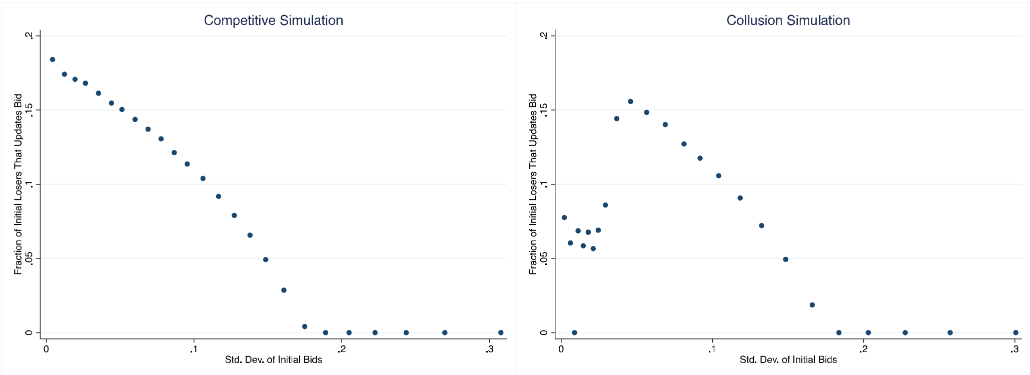

(a) Simulation: Fraction of Initial Losers That Updates Bid (Binscatter)

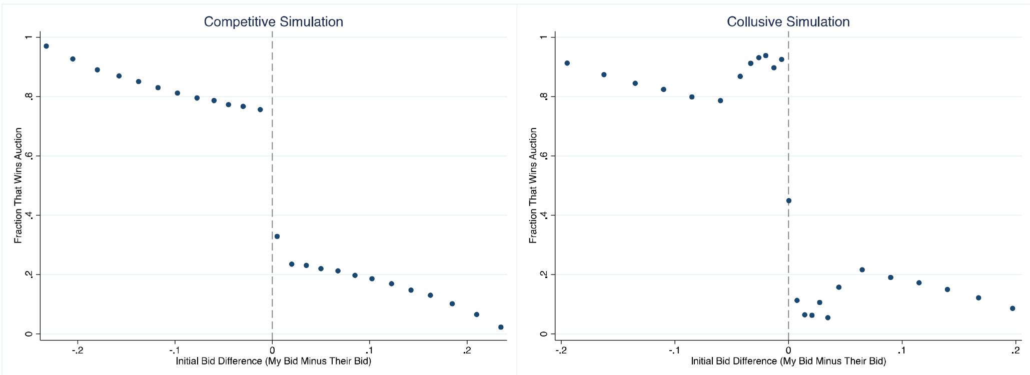

(b) Simulation: Probability Own Bid Should Win Given Initial Bid Difference (Binscatter)

Figure 4: Results of Bidding Simulation.

Notes: Panel (a) is a binscatter of the fraction of initial losers that updates their bid against the standard deviation of initial bids. For easier visibility, we have restricted the sample to an absolute initial bid difference below the 75^{th} percentile. Panel (b) is a binscatter of the probability of having the lowest standing bid at the end of the auction against the initial bid difference. For easier visibility, we have restricted the sample to an absolute initial bid difference below the 75th percentile.

discussion of the equilibrium, we know that bidders should update as long as the standing lowest bid is above their costs.[9] However, as we argue in the equilibrium section, under the assumption of fully revealing initial period bids, there is a unique and strictly increasing function $b(\cdot)$ that maps costs to initial bids. Thus, we can use the initial bid $b_{\ell(i)}$ by the initial loser $\ell(i)$ to control for his costs in auction i. In particular, the model implies

$$u_i = 1\{b^{-1}(b_{\ell(i)}) < b_{w(i)}\}.$$

While our empirical specification builds on this model prediction, it extends it in several ways to account for the fact that $b^{-1}(\cdot)$ is unknown, bidding mistakes, bid submission failures, and potential collusion:

$$u_i = 1\{b_{w(i)} - \phi(b_{\ell(i)}) + \delta_{\ell(i),w(i)} + \epsilon_i \geq 0\}. \tag{1}$$

We now discuss the differences between the model prediction and (1) in turn.

Firstly, the equilibrium bidding function $b(\cdot)$ is not known to the econometrician; hence we proxy $b^{-1}(b_{\ell(i)})$ with a non-parametric function of the initial loser's bid $\phi(b_{\ell(i)})$. In practice, $\phi(\cdot)$ is a third-degree polynomial.[10]

Secondly, we are interested in the possibility of collusion and hence allow the likelihood of undercutting to depend on a pairwise[11] fixed effect $\delta_{\ell(i),w(i)}$ that reflects the tendency of the initial loser $\ell(i)$ to undercut against the initial winner $w(i)$. In a competitive model, the fixed effects are irrelevant as only the costs of bidders matter for their undercutting behavior. Hence, under competition, the true value of any pairwise fixed effect is $\delta_{a,b} \equiv 0$. However, this might not be the case in a collusive model. For instance, our analysis above revealed that penalized cartels are much less likely to undercut each others' bids in the online auction.

Finally, we introduce an idiosyncratic exogenous auction-specific shock ϵ_i. This shock accounts for the previously referenced idiosyncratic chance of bid submission failure, but

[9] As in the discussion of the equilibrium, we will assume an arbitrarily small minimum bid decrement.

[10] We present the coefficients from this regression in Table 4 in the online appendix.

[11] We impose $\delta_{a,b} = \delta_{b,a}$ for maximum power.

also further extends the model by allowing a bidder to undercut by mistake.

To complete the model, it remains to specify a distribution for ϵ_i. To keep estimation feasible and side-step a potential incidental parameter problem, we specify ϵ_i in such a way as to yield a linear probability model[12]. This implies (1) can be consistently estimated via OLS, which allows us to handle high-dimensional fixed effects without encountering computational constraints. Furthermore, we note that if our data had been generated by competitive equilibrium play, this would imply $\delta \equiv 0$ and hence we would have $\hat{\delta} \sim N(0, \sigma^2)$ for some σ^2 (Kwon, 2023).

We estimate the linear probability corresponding to (1) and plot the empirical distribution of the estimated links (fixed effects) in Figure 5a. We estimate fixed effects only for pairs that we repeatedly observe in the data. When compared to a Normal distribution, the plotted data has an excess mass below the mean. Using Kolmogorov-Smirnov tests, we reject the hypothesis that the distribution is normal (p < 0.001) and the (weaker) hypothesis that the distribution is symmetric (p < 0.001). The additional mass on the left side of the distribution shows a significant number of pairs that are less likely to undercut each other's bid compared to the competitive baseline – an observation that would be expected on a market with many collusive firms.

Verification of the Algorithm

We now show that our algorithm successfully detects penalized pairs of firms. In particular, we utilize the data on which firms were penalized by the Antimonopoly Committee of Ukraine as part of the same collusion case. In Figure 5b, we compare the estimated pairwise fixed effects for penalized firm pairs to those of non-penalized firm pairs. We observe that the penalized pairs are concentrated in the left part of the distribution with the suspicious additional mass. However, even among firms that were not penalized, there is still excess mass on the left, suggesting that the courts did not identify all colluding firms. The figure shows that our algorithm works well[13] for the subset of sanctioned colluders and could be

[12]This implies ϵ_i is uniformly distributed.

[13]The relevant t-statistic from regressing the collusive pair dummy on the pairwise FE is -14.84.

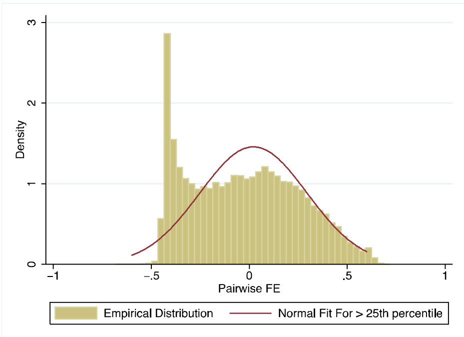

(a) Distribution of Pairwise Fixed Effects

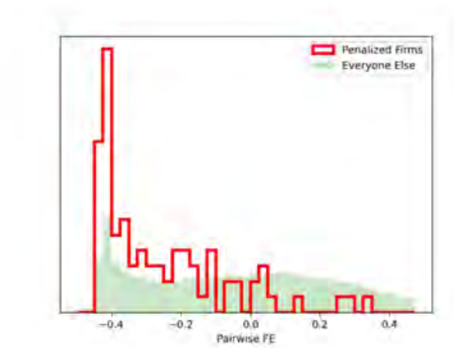

(b) Histogram by Penalization Status

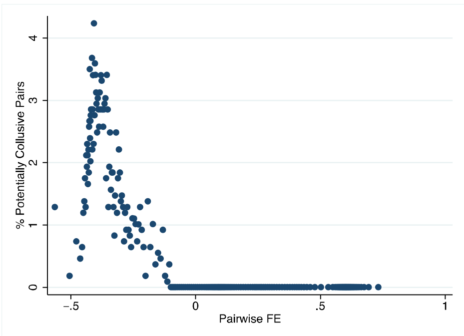

(c) Rejection of Null of Competition

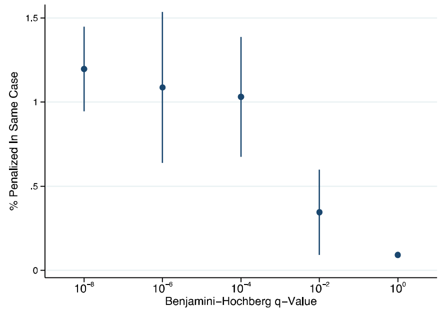

(d) Fraction of Pairs Penalized

Figure 5: Pairwise Fixed Effects & Statistical Test For Collusion

Notes: In 5a, we show the distribution of the pairwise fixed effects of (1). As a visual guide to emphasize departure from normality, we also exhibit a normal fit for the data above the 25th percentile. There is excess mass on low fixed effects, indicating that certain bidder pairs never undercut each other. In 5b, we break up the distribution of pairwise fixed effects by whether a pair was penalized in a collusion case by the Antimonopoly Committee of Ukraine and plot the distributions of pairwise FE for penalized and non-penalized firms separately. In 5c, we see that for the intermediate negative fixed effects, we can frequently reject the null of competition. Note this is not true for the most extreme negative fixed effects: these are mostly driven by sampling error. In 5d, we show that pairs with lower q-values (i.e., those our test detects as colluding) are also more likely to have been penalized in the same case for collusive conduct (the t-stat from a binary linear probability model regressing an indicator for penalization on the q-values is -9.68).

used for further identification of other likely colluding firms.

Detecting Collusive Pairs

Pairwise fixed effects, our measure of competitive behavior within pairs, are simply coefficients from a regression, which means that we can apply standard statistical tests to them. First, under the null of competitive behavior, we would have $\delta_{ab} = 0$ for all firms a and b. Second, suppose that cartels collude by not undercutting each other with some positive probability. Under this assumption, $\delta_{ab} < 0$ if a and b are members of the same cartel. This observation means that by testing whether δ_{ab} is not positive, one also examines whether the firms in the pair behave non-competitively against each other.

Accordingly, we form a one-sided test to assess for each pairwise fixed effect δ_{ab} whether it lies significantly below zero. Formally, let n be the number of observations and k the number of regressors (including the fixed effects) in the linear regression corresponding to (1). Define

$$t_{ab} := \frac{\hat{\delta}_{ab}}{se(\hat{\delta}_{ab})} \sim t_{n-k}.$$

Then, if F is the CDF of a t_{n-k} distribution, we can find the p-value associated with t_{ab}:

$$p_{ab} = 1 - F(-t_{ab}).$$

This gives us a large set of p-values. To avoid the classic multiple hypothesis testing problem, we will not control the size of each individual test, but rather the false discovery rate, i.e., the expected fraction of discovered collusive links that turn out to be false positives. To do so, we transform the p-values into an associated set of q-values (Benjamini and Hochberg, 1995), and finally reject the null hypothesis only for links for which $q_{ab} < 0.05$. These links are the ones declared collusive by our procedure. To further verify the soundness of our detection algorithm, we show that pairs with low q-values – i.e., those our test detects as colluding – are more likely to have been penalized in the same case for collusive conduct by the Antimonopoly Committee of Ukraine. We present these findings in Figure 5d. The relevant t-statistic from a binary linear probability model regressing an indicator for penalization on the q-values is -9.68.

There are three main caveats to our procedure. Firstly, our procedure jointly tests the null of our model being correctly specified and firms acting competitively; hence, our tests could reject for no other reason than a misspecified model. Secondly, our procedure is unable to detect sophisticated collusion including, for instance, a cartel that imitates competitive play under inflated costs. This is in line with classical impossibility results (Bajari and Ye, 2003). Finally, our procedure aims to be conservative: thus, it may miss collusive links for some pairs of players that have not faced each other often (as the null hypothesis is that firms are not colluding.) This issue is amplified by using the Benjamini-Hochberg procedure, which corrects for multiple hypothesis testing by requiring smaller p-values to reject the null of competition.

Another potential issue with our approach could be caused firms that are generally inactive in auctions for reasons other than collusion. To test whether our procedure is robust to such behaviour, we rerun our analysis while employing an alternative specification of Equation 1 that includes bidder fixed effects. This specification accounts for firm-specific undercutting patterns. Note, however, that in doing so, this specification makes it impossible to capture collusion for bidders who tend to collude in all of the auctions they participate in, which is why we do not use this as our main specification. Our algorithm nevertheless remains a good predictor of the collusion detected by the Antimonopoly Committee of Ukraine. We report the results in Appendix E.

Characteristics of Network Links

In this section, we present a descriptive analysis of the network characteristics exhibited by colluding firms. Our investigation begins by examining the geographical distribution of these firms based on their registered office locations. It is reasonable to assume that cartel-involved firms are more likely to be situated within the same geographical region, as this may facilitate the development of trust and coordination of illicit activities (Price, 2008).

To approximate firm locations, we utilize ZIP codes.

Our findings reveal that approximately 5.9% of non-collusive firm pairs are situated within the same ZIP code area, whereas this proportion increases to 11% for collusive firm pairs. This pattern persists when evaluating larger geographical areas, such as those sharing the same first 2, 3, or 4 digits of ZIP codes. For instance, in the largest geographical areas defined by the first two digits of ZIP codes, we find that 18.5% of non-collusive firm pairs and 31.2% of collusive firm pairs are located within the same region. The observed differences in means are statistically significant across all examined ZIP code variations, with corresponding t-statistics of -9.5, -10.3, -13.7, and -14.3 for the full code and each truncated version, respectively.

Subsequently, we investigate the dimensions and structures of collusive networks. Our analysis identifies 2,371 unique firms participating in collusion, forming 1,919 distinct firm pairs. These firms form 707 cartels. A majority of these cartels (496) have two members only. The remaining 211 cartels have an average size of 6.5 firms and the median size of 4 firms. In these 211 cartels, 61% of members maintain a single link to another firm, and the maximum number of links across cartel members constitutes 62.6% of the cartel's size on average. The magnitude of this fraction suggests that large cartels typically rely on a central cartel member for cooperation — evidence for a hub-and-spoke cartel structure.

5.3 Cost Analysis

To provide evidence of the economic impact of cartels on procurement auctions, we investigate whether auctions involving cartels result in higher prices. In particular, we regress normalized prices[14] on a dummy variable indicating whether a firm pair identified as collusive by our algorithm participated in a given auction.

The findings from this analysis are shown in Table 2. We present four specifications, in which, from Column (1) to (4), we gradually add fixed effects for industry, time, procurer, and winner identities. In all specifications, we hold the number of bidders constant. Thus,

[14]The normalized price is defined as the ratio of the final price of a contract over the estimated cost.

	(1) Norm. Price	(2) Norm. Price	(3) Norm. Price	(4) Norm. Price
Cartel Participation	0.0314	0.0309	0.0234	0.0212
	(0.0013)	(0.0013)	(0.0008)	(0.0010)
No of Bids	-0.0237	-0.0237	-0.0218	-0.0186
	(0.0004)	(0.0004)	(0.0003)	(0.0003)
Industry FE	Yes	Yes	Yes	Yes
Year FE	No	Yes	Yes	Yes
Procurer FE	No	No	Yes	Yes
Firm FE	No	No	No	Yes
N	291,368	291,368	291,368	291,368

Table 2: Cartel Participation and Normalized Price

Notes: Standard errors clustered at the procuring entity level in parentheses. The dependent variable is the ratio of the final price of a contract over the estimated cost. The main independent variable $Cartel\ Participation$ is a dummy variable equal to one if a pair of firms identified as (statistically significantly) collusive by our algorithm, otherwise 0.

the Cartel Participation coefficient measures the difference in the normalized price between the situation when there is a cartel and the counterfactual where the cartel breaks down. Procurement contracts with two bidders are the most common in the data. In this case, we draw a comparison between contracts involving two bidders in active competition and contracts where the two bidders are part of a cartel. The first three sets of fixed effects ensure that we compare contracts of the same type (4 digits CPV industry codes), procured in the same year and by the same buyer. In Column (4), winner fixed effects are added. Then, the coefficient measures the difference in the normalized price for the situation when a particular firm wins a contract while being in a cartel with the other bidder versus when it is not in a cartel with the other bidder.

We consistently find a positive price effect of collusion. In the most saturated specification, the estimated difference in the prices of contracts with and without a cartel is equal to 2.12% of the estimated cost.[15] Interestingly, our analysis reveals that the impact of cartel participation is comparable, in absolute value, to the influence of the number of bidders—albeit with opposite effects. This result appears plausible as one can think of a cartel of two firms

[15] A limitation of our approach is that cartels are not assigned to auctions randomly.

	(1) Norm. Price	(2) Norm. Price	(3) Norm. Price	(4) Norm. Price
Cartel Participation	0.0206	0.0204	0.0140	0.0141
	(0.0017)	(0.0017)	(0.0014)	(0.0014)
Cartel Alone	0.0131	0.0129	0.0116	0.00922
	(0.0018)	(0.0018)	(0.0015)	(0.0014)
No of Bids	-0.0230	-0.0230	-0.0212	-0.0181
	(0.0004)	(0.0004)	(0.0003)	(0.0003)
Industry FE	Yes	Yes	Yes	Yes
Year FE	No	Yes	Yes	Yes
Procurer FE	No	No	Yes	Yes
Firm FE	No	No	No	Yes
N	291,368	291,368	291,368	291,368

Table 3: Cartel Participation with and without Competitors

Notes: Standard errors clustered at the procuring entity level in parentheses. The dependent variable is the ratio of the final price of a contract over the estimated cost. The main independent variables are $Cartel\ Participation$, which is a dummy variable equal to one if a pair of firms identified as (statistically significantly) collusive by our algorithm, otherwise 0, and $Cartel\ Alone$ is a term indicating whether a cartel is alone in the procurement auction.

participating in an auction as de facto just one player. So the effect of a cartel's presence is similar to a decline in competition by one firm.

Cartels participate in 18.88% of public procurement contracts in the sample. According to our estimates, such contracts are overpriced by 2.12% in comparison to non-collusive contracts. A simple back-of-the envelope calculation yields an estimate of the total cost of collusion on this market equal to $0.1888 \cdot 0.0212 \approx 0.004$, i.e. approximately 0.4% of the value of public procurement market in Ukraine. Thus, the cartels we detect are responsible for significant losses to the Ukrainian government.

The participation of competitive firms might hinder the ability of cartels to inflate the price. We examine this in Table 3, in which we run the same regression as in Table 2 but we include an interaction term indicating whether the cartel members were alone in the auction. The findings indicate that the effect of cartel participation on the normalized price is about 0.9pp higher when the cartel is alone in the auction. When a cartel participates in an auction in which non-cartel bidders are also present, however, contracts are still overpriced by 1.4%.

6 Conclusion

We developed a novel algorithm to detect collusion in multistage auctions. Intuitively, we identify collusive firm pairs from their repeated lack of willingness to undercut each other's bids. In doing so, we explicitly allow firms to collude or compete depending on who else is participating on the same auction with them. Our findings show that collusion is widespread in the Ukrainian public procurement market: 2,371 collusive firms participate in 19% of auctions, increasing costs by 2.12%. Cartels typically comprise just two members, and members often share the same ZIP code.

While our algorithm relies on a particular auction mechanism, our results are nevertheless relevant outside the Ukrainian context. Firstly, the algorithm is immediately applicable in all countries that implemented similar multi-round e-procurement systems – a group that includes Georgia, Kyrgyzstan, and Moldova. Secondly, our analysis speaks to the extent of collusion in procurement auctions in developing countries, which our results suggest to be considerable.

References

Albæk, S., Møllgaard, P., and Overgaard, P. B. (1997). Government-assisted oligopoly coordination? a concrete case. *The Journal of Industrial Economics*, 45(4):429–443.

Ashkenazi-Golan, G., Tsodikovich, Y., and Viossat, Y. (2023). I want to tell you? maximizing revenue in first-price two-stage auctions. *Econ Theory*, 76:1329–1362.

Asker, J. (2010). A study of the internal organization of a bidding cartel. *American Economic Review*, 100(3):724–62.

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

Baranek, B. and Titl, V. (2024). The cost of favoritism in public procurement. *Journal of Law and Economics*.

Benjamini, Y. and Hochberg, Y. (1995). Controlling the false discovery rate: a practical and powerful approach to multiple testing. *J. Roy. Statist. Soc. Ser. B*, 57(1):289–300.

Bosio, E., Djankov, S., Glaeser, E., and Shleifer, A. (2022). Public procurement in law and practice. *American Economic Review*, 112(4):1091–1117.

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

Chassang, S. and Ortner, J. (2019). Collusion in auctions with constrained bids: Theory and evidence from public procurement. *Journal of Political Economy*, 127(5):2269–2300.

Clarka, R., Coviellob, D., and De Leveranoc, A. (2020). Complementary bidding and the collusive arrangement: Evidence from an antitrust investigation. *SSRN*.

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

Decarolis, F., Fisman, R., Pinotti, P., and Vannutelli, S. (2024). Rules, discretion, and corruption in procurement: Evidence from italian government contracting. *Journal of Political Economy: Microeconomics*.

Demokratizatsiya, S. S. (2017). Public procurement reform in Ukraine: the implications of neopatrimonialism for external actors. *demokratizatsiya.pub*.

Engelbrecht-Wiggans, R. (1988). On a possible benefit to bid takers from using multi-stage auctions. *Management Science*, 34(9):1109–1120.

Hyytinen, A., Steen, F., and Toivanen, O. (2018). Cartels uncovered. *American Economic Journal: Microeconomics*, 10(4):190–222.

Kang, K. and Miller, R. A. (2020). Winning by default: Why is there so little competition in government procurement?

Kawai, K. and Nakabayashi, J. (2022). Detecting large-scale collusion in procurement auctions. *Journal of Political Economy*, 130(5):1364–1411.

Kawai, K., Nakabayashi, J., Ortner, J., and Chassang, S. (2023). Using Bid Rotation and Incumbency to Detect Collusion: A Regression Discontinuity Approach. *The Review of Economic Studies*, 90(1):376–403.

Kawai, K., Nakabayashi, J., and Shimamoto, D. (2022). A study of bid-rigging in procurement auctions: Evidence from indonesia, georgia, mongolia, malta, and state of california. Technical report.

Kwon, S. (2023). Optimal shrinkage estimation of fixed effects in linear panel data models. Working Paper.

Merrifield, A., Manchidi, T. E., and Allen, S. (2002). The asset procurement and operating partnership system (apops) for prisons in south africa. *International Journal of Project Management*, 20(8):575–582.

Mironov, M. and Zhuravskaya, E. (2016). Corruption in procurement and the political cycle in tunneling: Evidence from financial transactions data. *American Economic Journal: Economic Policy*, 8(2):287–321.

Noumba Um, P. and Dinghem, S. (2005). Private participation in infrastructure projects in the Republic of Korea. Policy Research Working Paper Series 3689, The World Bank.

Partnership, O. C. (2020). Ukraine S&L: MEL Framework. pages 1–20.

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

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

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

Price, M. K. (2008). Using the spatial distribution of bidders to detect collusion in the marketplace: evidence from timber auctions. *Journal of Regional Science*, 48(2):399–417.

Sauvet-Goichon, B. (2007). Ashkelon desalination plant — a successful challenge. *Desalination*, 203(1):75–81.

Schoenherr, D. (2019). Political connections and allocative distortions. *The Journal of Finance*, 74(2):543–586.

Supreme Council of Ukraine (2015). Ukrainian Public Procurement Law. pages 1–41.

Titl, V. and Geys, B. (2019). Political donations and the allocation of public procurement contracts. *European Economic Review*, 111:443–458.

Transparency International Ukraine (2017). Co-Creation of ProZorro: An Account of the Process And Actors. Technical report.

Transparency International Ukraine (2019). Behind the Scenes of ProZorro. Does Ukrainian business trust public procurement? Technical report.

World Bank (2020). 2019 Public Expenditure and Financial Accountability (PEFA) Performance Assessment Report.

Online Appendix

A Institutional Background

The story of public procurement in Ukraine is long and complicated (for a summary see Transparency International Ukraine, 2017). While a first real effort to develop procurement legislation in 1997 was motivated by the need to harmonize regulations with WTO standards, the resulting law introduced in 2000 was lacking in detail and clarity (Transparency International Ukraine, 2017). The situation deteriorated substantially when the newly established 'Tender Chamber of Ukraine' was put in charge of all public procurement in 2005 and promptly began exercising its power to unduly influence bidder selection (Demokratizatsiya, 2017). An interim period followed in which there were several unsuccessful attempts to fix the system.

In 2013, the suspension of negotiations with the European Union by Ukrainian president Viktor Yanukovych sparked demonstrations. It marked the beginning of a period of political turmoil, the 'Euromaidan'. As protests spread, Yanukovych fled the country, and parliament relieved him of his duty. While an interim government led the country, the head of the Ministry of Economic Development and Trade (MoE) asked volunteers to organize themselves and research possibilities for reforming various governmental institutions. Public procurement was one of them. After meetings with Georgian and EU procurement experts, the volunteers agreed to model their system on the Georgian example.

However, two issues remained. There was a worry that a centrally administrated system would not provide sufficient incentives for ease of use. Furthermore, there was no apparent source of funding for the project: perhaps surprisingly, the official procurement department did not yet support the reform. Ukraine adopted a 'hybrid' system in which access to a central database of procurement contracts is mediated by various marketplaces that are allowed to charge a fee for this access but, in turn, provided initial funding for the development of the system. Transparency International Ukraine agreed to manage the budget during the pilot phase of the project, collected the contributions from the marketplaces, and selected a company for the necessary software development.

With initial funding secured, a pilot of what would eventually become the ProZorro e-procurement system went live in February 2015. However, at this stage, the project was still entirely a volunteer-led reform initiative: things only changed when a volunteer representative became the head of the Department of Public Procurement Regulation in March 2015. Thus, the status of the project was elevated, parliament passed new legislation in November 2015, and new funding from multiple international organizations allowed various refinements of the pilot necessary for full deployment. Finally, in April (August) 2016, the use of ProZorro became compulsory for many (all) public entities.

At its core, ProZorro is *(i)* a unified central database of all public procurement projects conducted in Ukraine and *(ii)* an API for interacting with this database. Appropriate legislation ensures that procurers post all public tenders to this database, and (crucially!) read-only access (e.g., for monitoring or research) is always free. Procuring entities and tenderers interact with the database via one of several profit-oriented marketplaces that allow the (free) posting of and (fee-incurring) participation in tenders via their unique interfaces. However, the 'auctions' themselves are run by the central database so that marketplaces cannot unduly influence their result.

The marketplaces (or the whole system) are often referred to as 'eBay for public procurement' in the media. Such simplification, however, falsely suggests that the main innovation of the system is the easy access to new tenderers through the use of information technology. While this plays a part in the success of ProZorro, the platform's primary purpose is better described by its name: 'transparency.' By design, all the information that exists about a tender is readily available publicly. All interested parties can, therefore, easily monitor procurement contracts.

The fact that transparency was the primary purpose of the development of ProZorro becomes even more salient when we examine several initiatives built to complement and support the platform. Firstly, the 'analytics module' allows quick access to summary statistics;

the module is sufficiently interactive to allow for productive exploration of the data at a journalistic level. Furthermore, the MoE and ProZorro have introduced several procurement qualifications. While the university courses mainly aim at teaching potential future civil servants how to *run* successful tenders, there are also online courses with a more explicit focus on monitoring, for example, the aptly named 'Monitoring of Public Procurement; Or: How To Look for Betrayal.'

The introduction of ProZorro has been widely lauded as a highly positive step for public procurement in Ukraine. Indeed, ProZorro has received several awards (such as being rated #1 by the World Procurement awards 2016 in the Public Sector nomination). The World Bank in 2020 assigned the Ukraine letter-grades of A in nearly all scored dimensions of public procurement. The sole exception was the 'procurement methods' score since only 78.1% of the total cost of all public procurement covered by the relevant law in 2018 was tendered in competitive procedures (World Bank, 2020).

In the main text, we argue that while the formal institutions in Ukraine have greatly improved, a closer examination of the bidding suggests that collusion and shill-bidding have become costly problems. Indeed, only 13.3% of respondents in a 2017 survey agreed that 'the system helps increase competition and achieves value for money' (Partnership, 2020). When asked a similar question in 2019, this number improved, and 46.3% of respondents said that the level of corruption in public procurement had slightly or significantly decreased after the launch of ProZorro (though 12.2% said it had increased) (Transparency International Ukraine, 2019). However, 24.2% still stated that they had personally encountered situations in which they were 'forced to pay a bribe or resort to nepotism after ProZorro was launched, and 34.2% say that corruption is the most severe problem facing the platform. Our analysis supports the public perception of widespread collusion.

B Equilibrium of ProZorro Auction

Proposition 1. *In any equilibrium in which initial bids are given by some strictly increasing $b(\cdot)$, the expected payoff from pretending to be type $\tilde{c}$ is given by $V^{PZ}(\tilde{c})$ no matter the number of updating rounds or number of players.*

Proof. We consider the PZA auction with $k+1$ rounds (i.e., k updating rounds) and n players; we will index rounds by r and players by i. We will refer to the bid by player i in round r as b_i^r and use $\underline{b}_i^r$ to notate the standing lowest bid *before* i moves in round r. Note that bidding in updating rounds is not (necessarily) in order of player indices as updating priority is based on the ranking of the bids from the previous round; hence, we also introduce $\sigma(r,t)$ as notation for the index of the player that moves in position $t = 1, \ldots, n$ in round r. Thus, for example, $b^1_{\sigma(2,3)}$ refers to the first round bid by the player who moves third in the second round.

We assume that initial bids are fully revealing, and hence can let $\hat{c}_i := b^{-1}(b_i^1)$ be the shared (point-)belief of $j \neq i$ about the cost type of player i. As we are considering only deviations by P1 (wlog), we have $\hat{c}_i = c_i$ for all $i \neq 1$. This also implies that all bidders but P1 move in order of their costs in the first updating round; the position of P1 is determined by the cost-type he chooses to imitate.

We will regularly need to refer to the optimal bid of a player i who anticipates that no firm moving after her is capable of beating a standing bid of x but at least one is capable of beating all higher bids. Such a player would like to bid x, but may be constrained by her own cost. If x is below her cost, the player – anticipating that she will be beaten – would be indifferent between all other bids were it not for the possibility of bid submission failure. As it is, however, there is a small but positive probability $p > 0$ that any given subsequent bid submission attempt will fail. This gives her a chance to nevertheless win the auction: for instance, if there is just one player to move after her that could beat $y > c_i$, she could submit y and hope that this player will fail to submit his bid. Even if she believes that all players to move after her can beat her own cost c_i (as all players believe in equilibrium), there is still a chance that they all fail (repeatedly) at submitting their bids, in which case she can win by undercutting the standing winning bid by Δ. More generally, we will refer to the optimal undercut as $\Delta^*(i,r)$ without characterizing it further and introduce the following notation for the optimal bid of a player i in round r who anticipates that he would not be beaten if she bid x:

$$g_i^r(x) = \begin{cases} \max\{b : b \leq x, b \leq \underline{b}_i^r - \Delta, b \geq c_i\} & \text{if this set is nonempty,} \\ \underline{b}_i^r - \Delta^*(i,r) & \text{o/w and if } \underline{b}_i^r - \Delta^*(i,r) \geq c_i \\ b_i^{r-1} & \text{o/w.} \end{cases}$$

It of course remains to characterize the value of x after each history, which we will now do by proceeding with backward induction. Firstly, noting that $\sigma(k+1,n)$ is the last player to move in the last round, we claim that in any SPE,

$$b_{\sigma(k+1,n)}^{k+1} = \begin{cases} b_{\sigma(k+1,n)}^{k} & \text{if } b_{\sigma(k+1,n)}^{k} = \underline{b}_{\sigma(k+1,n)}^{k+1} \\ g_{\sigma(k+1,n)}^{k+1}\big(\underline{b}_{\sigma(k+1,n)}^{k+1}\big) & \text{o/w} \end{cases}$$

Thus, the last agent to move will simply undercut by as much as necessary in order to win the contract (assuming this yields positive profit). Anticipating this, all other agents in the last round would like to scoop, i.e., ensure that their bid cannot be undercut by anyone moving after them. Hence, they will anticipate that they can win if and only if they bid no more than the bid decrement Δ above the cost of whoever they believe to be the lowest cost agent moving after them. They thus bid

$$\forall t < n : b_{\sigma(k+1,t)}^{k+1} = g_{\sigma(k+1,t)}^{k+1}\Big(\min\{\hat{c}_{\sigma(k+1,s)} | s > t\} + \Delta\Big).$$

It should be noted that if $k=1$, this implies that all players but P1 and $\sigma(k+1,n)$ will simply undercut the current standing bid by Δ (if possible without going under their cost). This is because the order in which players are moving is exactly the order of player strength given their beliefs: hence they anticipate never being able to win the auction if no bid submission failure occurs. The same is true for P1 as long as he is pretending to be either a stronger type than he actually is or his true type. If he is pretending to be a weaker type, then and only then can he successfully ‘scoop’.

If $k > 1$, the argument in the preceding paragraph still applies as long as the order of players hasn’t changed between updating rounds. However, it may change due to the behavior of P1. Nevertheless, the strategies stated above are still optimal.

Moving backward, given the situation in the last updating round, all agents anticipate that the agent with the lowest cost will win. Thus, all agents (including P1) are in the same situation in round k as in $k+1$, and hence they will play essentially the same strategies: all players but P1 undercut in the hope of a bid submission failure, and P1 scoops if he is actually the lowest type but was initially pretending not to be. Why does P1 scoop ‘early’ rather than ‘late’? By scooping early, he guards against the fact that his own late scooping bid may not go through. Thus, strategies in earlier updating rounds are mostly unchanged from later updating rounds:

$$\forall\, 1 < r < k+1 : \forall t : b^r_{\sigma(r,t)} = g^r_{\sigma(r,t)}\Big(\min\{\hat{c}_{\sigma(r,s)} | s = 1, \ldots, n\} + \Delta \Big).$$

Finally, note that if $\tilde{c} \leq c_1$, then $\Delta^*(i,r) \equiv \Delta$ as all agents (including P1) anticipate that all agents moving after them can beat their own costs. If $\tilde{c} > c_1$, this is not true anymore: in particular, P1 may anticipate that some firms that will get to update their bid after him cannot beat his costs. However, either c_1 is the lowest cost draw or not. If it is, then P1 will never be forced to contemplate the case in which he relies on bid submission failure to win, and as $p \to 0$, his payoff from pretending to be $\tilde{c}$ will converge towards that he would get if there was no bid submission failure chance. If it is not the lowest cost, then with probability approaching one, P1 will not win the auction. Hence, his payoff will be zero, no matter what complicated undercutting strategies Δ^* he employs in the meantime.

Thus, as we take the limits $p \to 0$, $\Delta \to 0$, the strategies derived in this proof imply the following payoff from pretending to be type $\tilde{c}$ in the initial round (when your true type is c_1):

$$\begin{aligned} V^{PZ}(\tilde{c}) &= \mathbb{P}\Big(b(\tilde{c}) < \min_{j\neq 1} b(c_j)\Big)\Big(b(\tilde{c}) - c_1\Big) + \\ &\quad \mathbb{P}\Big(b(\tilde{c}) > \min_{j\neq 1} b(c_j) \cap c_1 < \min_{j\neq 1} c_j\Big) \times \\ &\quad \mathbb{E}[\min\{c_j : c_j < \tilde{c}, j \neq 1\} - c_1 | c_1 < \min_{j\neq 1} c_j, b(\tilde{c}) > \min_{j\neq 1} b(c_j)] \end{aligned}$$

$\square$

Proposition 2. *The ProZorro auction (with $k \geq 1$ rounds and $n \geq 1$ players) has a unique PBE in which initial bids are given by a strictly increasing $b(\cdot)$. In this equilibrium,*

$$b(c) = \frac{1}{[1 - F(c)]^{n-1}} \int_c^{c_{max}} s(n-1)f(s)[1 - F(s)]^{n-2} ds$$

and bids are decreased by the arbitrarily small minimum bid decrement whenever doing so is possible without bidding below one's own cost.

Proof. We have

$$\begin{aligned} V^{PZ}(\tilde{c}) &= \mathbb{P}\Big(b(\tilde{c}) < \min_{j\neq 1} b(c_j)\Big)\Big(b(\tilde{c}) - c_1\Big) + \\ &\quad \mathbb{P}\Big(b(\tilde{c}) > \min_{j\neq 1} b(c_j) \cap c_1 < \min_{j\neq 1} c_j\Big) \times \\ &\quad \mathbb{E}[\min\{c_j : c_j < \tilde{c}, j \neq 1\} - c_1 | c_1 < \min_{j\neq 1} c_j, b(\tilde{c}) > \min_{j\neq 1} b(c_j)] \end{aligned}$$

Say $c_i \sim F(\cdot)$ with $\max \operatorname{supp} c_i = c_{max}$. We will use

$$G(\tilde{c}) = 1 - [1 - F(\tilde{c})]^{n-1}$$

as a short-hand to refer to the distribution of the minimum of the $n-1$ other costs. Then

$$\begin{aligned} V^{PZ}(\tilde{c}) &= [1 - G(\tilde{c})]\big(b(\tilde{c}) - c_1\big) + [1 - G(c_1)] \times \\ &\quad \max\Big\{\frac{G(\tilde{c}) - G(c_1)}{1 - G(c_1)}, 0\Big\}\Big(\frac{1}{G(\tilde{c}) - G(c_1)} \int_{c_1}^{\tilde{c}} c dG(c) - c_1\Big), \end{aligned}$$

where we used the fact that $\min\{c_j : c_j < \tilde{c}\} = \min_{j\neq 1} c_j$ given that $\tilde{c} > \min_{j\neq 1} c_j$.

Although on first glance it may seem[16] like $V^{PZ}(\tilde{c})$ is not differentiable at $\tilde{c} = c_1$, this is in fact wrong because the potentially non-differentiable part of $V^{PZ}(\tilde{v})$ is multiplied by the expected rent from a second price auction conditional on your strongest opponent having a cost draw below $\tilde{c}$, which tends to zero as $\tilde{c} \to c_1$. After recognizing this, it is easy to see that

$$V'(c_1) = (1 - G(c_1))b'(c_1) - (b(c) - c)g(c),$$

where $g(c) = G'(c)$. Together with the boundary condition $b(c_{max}) = 0$, this differential equation is uniquely solved by

$$b(c) = \frac{1}{1 - G(c)} \int_c^{c_{max}} sdG(s),$$

which is just the classic first-price auction equilibrium bidding strategy. □

C Data Manipulation

We note that a large share of bids is 'too good to be true'. Such bids are likely to be provided without showing that the company is reliably able to deliver the demanded project which leads to the subsequent disqualification of the bids. As other firms can see such low bids at the start of the auction and anticipate that the suspiciously low bidder will be disqualified. In such cases, the optimal behavior would change and the bidders would only compete against other bidders and not the low bidder. To alleviate this problem we conduct our analysis only on the sample of auctions without very low bids, which we define as containing any auction where the lowest bid is below a conservative threshold of 80% of the highest bid of other participants. This leads to omitting around 35% of all auctions. Our results are robust to both using the specified sub-sample or the whole sample of all auctions. There are also other reasons why a firm might get disqualified but as these are not easily predicted both from the

[16]As $\max\left\{\frac{G(\tilde{c})-G(c_1)}{1-G(c_1)}, 0\right\}$ is not differentiable at this point.

data available to companies before the auction starts or from the ex-post data available to researchers we choose to not explicitly model them.

D Functional Form

In the main analysis, we run a linear probability regression model corresponding to Equation 1 to obtain pairwise fixed effects. There is potential sector, time, and procuring entity heterogeneity in the procurement tenders in our dataset. To test for robustness of our findings, we estimate the same regressions with sector and time fixed effects as well as procuring entity fixed effects. We present the coefficient from this regression in Table 4 in the online appendix. Note that these coefficients are not to be interpreted causally, but are shown for the sake of transparency in our detection algorithm.

Furthermore, we discuss the correlation between pairwise fixed effects and the penalization by the Antimonopoly Committee of Ukraine. In Table 5, we present the findings using the baseline specification (*Baseline*), the specification with sector[17] and time fixed effects (*Sector and time FEs*), the specification with procuring entity fixed effects (*Entity FEs*), and the specification including all three sets of fixed effects (*Sector, time, and entity FEs*). In all specifications, the coefficient is negative and strongly significant. The negative sign shows that with smaller pairwise fixed effects, the chances that a company was penalized increases. This confirms that the pairwise fixed effects are good predictors of actual penalization by the Antimonopoly Committee of Ukraine.

E Bidder Fixed Effects

To further check robustness of our findings, we rerun the main analysis while including bidder fixed effects to Equation 1. In Figure 6, we present the main findings. First, we show the distributions of pairwise fixed effects separately for the pairs penalized by the

[17] Sectors are implemented as four digits CPV codes.

Table 4: Pairwise fixed effects estimation

	(1) Update	(2) Update	(3) Update	(4) Update
Initial loser's bid	-5.250	-5.168	-5.137	-5.093
	(0.5370)	(0.5359)	(0.5124)	(0.5119)
(Initial loser's bid)^2	6.777	6.633	6.523	6.442
	(0.7096)	(0.7085)	(0.6771)	(0.6767)
(Initial loser's bid)^3	-3.315	-3.239	-3.154	-3.108
	(0.3061)	(0.3057)	(0.2921)	(0.2919)
log(Estimated Cost)	0.00590	0.00762	0.00664	0.00871
	(0.0009)	(0.0009)	(0.0009)	(0.0009)
Industry FE	No	Yes	No	Yes
Year FE	No	Yes	No	Yes
Procurer FE	No	No	Yes	Yes
N	1,773,740	1,773,740	1,773,740	1,773,740

Standard errors in parentheses

Notes: This table summarizes estimates from the regression model corresponding to Equation 1. The dependent variable is equal to $u_i - b_{w(i)}$. Initial loser's bid ($b_{\ell(i)}$) is the bid of the initial loser and the log(Estimated Cost) is the estimated costs of the public procurement contract.

Antimonopoly Committee of Ukraine for collusion and those that were not penalized (Figure 6a). And second, we show the relationship between q-values (i.e., the measure of being detected as a colluder by our algorithm) and the penalization status (Figure 6b). In this estimation, we filter out any firm-specific average inactivity. Both figures are substantively similar to the findings in the main analysis.

Finally, we use the new sample of auctions detected as collusive in the estimation including bidder fixed effects and rerun the cost analysis as in Section 5.3. The findings are reported in Table 6. The point estimates are slightly smaller than in the main analysis in most specifications; however, the point estimate is virtually identical in our preferred specification in Column (4) and all the estimates are statistically significant.

Table 5: Pairwise fixed effects and penalization by the Antimonopoly Committee of Ukraine

	results		
	Coefficient	t-stats	p-vals
Baseline	-.184	-14.838	0
Sector and time FEs	-.195	-15.228	0
Entity FEs	-.178	-13.764	0
Sector, time, and entity FEs	-.190	-14.200	0

Notes: This table shows the coefficients, t-statistics, and p-values from the regressions of the dummy of being a collusive pair according to the Antimonopoly Committee of Ukraine on the size of the pairwise fixed effects. Each row in the table represents the relevant statistics from the regressions using different specifications to obtain the pairwise fixed effects. The row labelled "Baseline" is the specification used in the main analysis (the linear regression corresponding to Equation 1). In the following three rows, we use the specifications with sector and time fixed effects, the specification with procuring entity fixed effects, and the specification including all three sets of fixed effects, respectively.

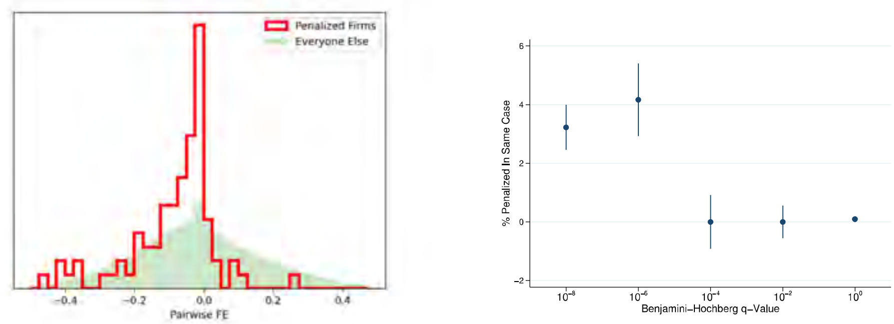

(a) Histogram by Penalization Status

(b) Fraction of Pairs Penalized

Figure 6: Pairwise Fixed Effects & Statistical Test For Collusion with Bidder FEs

Notes: In 6a, we show the distribution of the pairwise fixed effects of (1) broken up by whether a pair was penalized in a collusion case by the Antimonopoly Committee of Ukraine or not and plot the distributions of pairwise FE for penalized and non-penalized firms separately. In 6b, we show that pairs with lower q-values (i.e., those our test detects as colluding) are also more likely to have been penalized in the same case for collusive conduct.

	(1) Norm. Price	(2) Norm. Price	(3) Norm. Price	(4) Norm. Price
Cartel Participation	0.0268	0.0265	0.0198	0.0189
	(0.0015)	(0.0014)	(0.0009)	(0.0010)
No of Bids	-0.0235	-0.0235	-0.0216	-0.0184
	(0.0004)	(0.0004)	(0.0003)	(0.0003)
Industry FE	Yes	Yes	Yes	Yes
Year FE	No	Yes	Yes	Yes
Procurer FE	No	No	Yes	Yes
Firm FE	No	No	No	Yes
N	291,368	291,368	291,368	291,368

Table 6: Cartel Participation and Normalized Price - Bidder Fixed Effects Included

Notes: Standard errors clustered at the procuring entity level in parentheses. The dependent variable is the ratio of the final price of a contract over the estimated cost. The main independent variable is a dummy variable equal to one if a pair of firms identified as (statistically significantly) collusive by our algorithm, otherwise 0. The analysis is identical to the analysis in Table 2; however, the estimation used to detect collusion includes bidder fixed effects.