Measuring Bid Rigging at Scale

Detection and Prices in Indian Procurement

Jonathan Old*

November 1, 2026

JOB MARKET PAPER

Click here for the latest version

Abstract

Despite longstanding concerns about collusion in public procurement, its prevalence and effects on prices remain poorly measured, particularly where data are limited and enforcement is weak. This paper develops an empirical model that detects collusion among bidders in public procurement auctions when cartels use different coordination strategies, and applies it to the universe of tenders on India's government-developed e-procurement portals. Starting from twelve behavioral markers, such as bids submitted minutes apart or rotation of the winning bidder, the model identifies colluding pairs of bidders from the excess frequency of combinations of markers relative to a benchmark constructed under the assumption of no collusion. It estimates the prevalence of collusion and of each coordination strategy, and yields a ranking of bidder pairs by how likely they are to collude. Estimated on a new dataset of Indian public procurement covering 42.3 million bids on 15.7 million tenders since 2008, the model finds that collusion is widespread and most common under rural local bodies and in markets with few firms. Tenders won by the pairs most likely to collude settle 3.2 percent above comparable tenders in the same market on average. Linked rural road contracts show no difference in quality. The cost of collusion is therefore primarily a price increase, which, extrapolated across India, implies an annual transfer of approximately USD 4.0 billion from the government to colluding firms.

JEL codes: C38, D44, H57, K21, L41, O12.

Keywords: public procurement, collusion, bid rigging, cartel detection, India.

*University of California, Berkeley. Email: jonathan_old@berkeley.edu. I thank Ted Miguel, Francesco Trebbi, Fred Finan, and Cailin Slattery for excellent supervision and advice. Matthew Backus, Nano Barahona, Matilde Bombardini, Ben Faber, Benjamin Handel, Sam Kapon, Kei Kawai, Nicolas Longuet-Marx, Aprajit Mahajan, Amol Singh Raswan, Steve Tadelis, Lucio Wasserman, and Bruce Wydick as well as participants at the Berkeley Development Lunch Seminar, Berkeley IO Workshop, Berkeley Development Seminar, Robert Powell Political Economy Seminar, Berkeley Political Economy Research Lab, the Calgary Empirical Microeconomics Workshop, and the CEGA Research Retreat provided helpful feedback and comments. All errors remain my own.

1 Introduction

Public procurement – the purchase of goods, services, and infrastructure by the government – accounts for close to a third of government expenditure across countries of all income levels (OECD 2025a). In India, it amounts to roughly 8 percent of GDP.[1] This paper studies collusion among bidders in the universe of tenders conducted through India's government-developed e-procurement system. In the 2024–25 financial year alone, public bodies awarded contracts worth USD 105 billion through this system.[2]

At this scale, even small increases in procurement prices impose large costs to the government budget. Explanations for why governments overpay, particularly in developing countries, have focused on failures on the side of the government: corruption and favoritism toward firms connected to politicians and bureaucrats (F. Brugués, J. Brugués, and Giambra 2024; Lehne, Shapiro, and Vanden Eynde 2018), and variation in the effectiveness of the bureaucrats and organizations that make purchases (Best, Hjort, and Szakonyi 2023; Bandiera, Prat, and Valletti 2009). Governments have sought to address these problems by moving tendering online (Lewis-Faupel et al. 2016), centralizing the purchase of standardized goods (Bandiera, Prat, and Valletti 2009), and limiting officials' discretion through procurement regulation (Bosio et al. 2022). Yet prices can be high even when procurement officials follow the rules: bidders can coordinate and agree among themselves on who wins a contract and at what price. This practice, known as bid rigging, is a major concern for competition authorities (OECD 2025b).

Despite this regulatory attention, the prevalence of collusion in public procurement and its effects on prices remain poorly measured. In India, bid rigging is explicitly prohibited under the Competition Act, 2002, and the Competition Commission of India (CCI) investigates suspected violations and advises procuring organizations on how to prevent and detect collusion. Since 2010, the CCI has issued orders in 97 cases of bid rigging in public procurement and found a violation in 27, often years after the collusive conduct ended. Over the same period, Indian e-procurement portals recorded tens of millions of bids. These cases establish that collusion occurs, but they cannot show how widespread it is or what it costs the government: the cases that reach the CCI are not a representative sample of collusion, and they are too few to train or validate a method that detects collusion in bid records.

This paper asks two questions. First, how prevalent is collusion among bidders in public procurement in India, and what determines variation in prevalence across markets? Second, what is the effect of collusion on the prices the government pays and on the quality of the works it procures?

To answer these questions, I develop an empirical method to detect bidder collusion in repeated bidding at scale across different markets with potentially heterogeneous cartels. I then apply the method to a new dataset of every tender recorded in this system. Collusion is widespread – close to 30 percent of regular co-bidders collude – and it is most pervasive under rural local bodies and in markets with few firms. Collusive tenders are awarded at 3.2 percent higher prices than comparable tenders. For a subset of the data with available quality data from independent audits, collusive contracts do not differ in quality from other contracts. The welfare impact of collusion in this setting is therefore almost entirely a transfer from the government to colluding firms, which scaled to all

[1] Central and state budget aggregates put Indian public procurement at close to 30 percent of government expenditure. A calculation of these numbers is described in Appendix F.

[2] All monetary figures in this paper are converted from Indian rupees (INR) to USD at an exchange rate of Rs. 77/USD. The rate is the value-weighted average of the annual average exchange rate over the sample period.

Indian public procurement translates to USD 4 billion annually, more than India's largest state, Uttar Pradesh, spent on health in 2023–24.

Detecting collusion at scale requires a comprehensive dataset that records every bid, including losing ones, across repeated encounters of firms across many tenders. Measuring how prevalent collusion is, and under which conditions it is more common, further requires variation in procurement conditions, for example across levels of government, geographies, organizations, and products. I assemble such a dataset by scraping all procurement records from 76 Indian government e-procurement portals. The dataset contains 42.3 million bids on 15.7 million tenders, placed with 4,760 procuring organizations across 659 districts since 2008, for a cumulative value of USD 747 billion. The contracts included in the data cover a wide range of non-standardized goods and services, with construction and civil works being the largest categories by value. The 22 states and 8 union territories in the data account for 72% of India's population and 70% of its GDP. Existing work on Indian procurement has drawn on a single program, most often rural roads, which is only a small share of aggregate Indian public spending (Lewis-Faupel et al. 2016; Lehne, Shapiro, and Vanden Eynde 2018).

Contracts are allocated by a first-price sealed electronic auction to the lowest qualified bidder, and the portal publishes the bids of all bidders, including financial bids of winners and losers in the auction. Making the interactions between bidders analyzable requires resolving their identities: bidder names are entered as free text at registration, so a single firm may appear under many spellings, preventing it from being tracked across tenders. Entity resolution collapses 1.36 million raw name strings into 1.00 million distinct firms and contractors, which makes the tenders in which the same two firms both bid, the pair's *co-bids*, observable.

Colluding bidders enter auctions together to create the appearance of competition, and enforcing their internal agreement and dividing its gains leaves traces in their bids (Stigler 1964; McAfee and McMillan 1992). Existing methods to detect collusion rely on different traces in repeated bidding: rigid (Athey, Bagwell, and Sanchirico 2004) or round bids (Ishii 2014), bid rotation (Ishii 2009; Kawai, Nakabayashi, Ortner, et al. 2023), joint entry (Conley and Decarolis 2016), bids that are correlated across bidders after accounting for costs (Bajari and Ye 2003), or cover bids that do not respond to firms' cost shifters (Porter and Zona 1993; Porter and Zona 1999). These tests detect one trace at a time, some require a model of competitive bidding or data on costs, and they are usually validated against known cartels in one market. Across the whole procurement activity of a government, however, cartels might use different schemes, alone or in combination, and each scheme leaves different traces. The method in this paper therefore aggregates evidence across any set of traces that can be measured in bid records and models their joint distribution, so that a pair can be flagged for a combination of moderately elevated traces, even when no single trace would flag it.

Here, the traces are the "warning signs of bid rigging" listed in the CCI's *Diagnostic Toolkit* for public procurement officers, which is similar to guidance published by the OECD and the US Department of Justice (Competition Commission of India 2024; OECD 2025b; US DOJ 2015). From this list, I operationalize twelve markers of collusion that can be measured for a pair of bidders. The CCI flags, for instance, bids "submitted within short span of time" (Competition Commission of India 2024, pp. 18, 20). Two competing firms may submit close together by chance or because both bid near the deadline, but a pair that coordinates its bids should do so more often than firms with the same bidding habits. The marker is therefore the share of the pair's co-bids in which the two bids are submitted within 30 minutes of each other. Likewise, the other markers are rates or averages over

the pair's co-bids. I restrict the analysis to the 2,043,939 pairs with at least 5 co-bids in which both prices are opened. No marker uses the winning bid, the outcome in the price analysis.

To detect collusion, I estimate a mixture model in which each pair of bidders is either colluding or competing. Estimating the mixture requires knowledge of the distribution of markers among competing pairs. The approach follows large-scale hypothesis testing in statistics, where a mixture model estimates how many of a large number of null hypotheses are false and which ones (Efron 2004; Efron 2008). In economics, it has been used to estimate discrimination by individual employers in correspondence experiments, where random assignment of applicant characteristics makes the distribution under the null known (Kline and Walters 2021; Kline, Rose, and Walters 2024). Here, without a model of how firms bid, competition does not imply a particular distribution of markers, and the distribution cannot be observed directly because it is unknown which pairs compete. The main methodological contribution is a way to approximate this distribution without known cartels by generating a placebo distribution via resampling from the data.

In each of a pair's co-bids, the placebo replaces each firm's bid with a bid the same firm submitted in a randomly drawn tender its partner did not enter. The placebo thus keeps each firm's bidding habits but breaks the link between the two firms' bids. It is a valid stand-in for competition under two assumptions. First, absent collusion, a firm bids in tenders shared with its partner as it does in tenders its partner does not enter, which is plausible because bids are sealed and a competing firm does not observe who else bids. Second, absent collusion, the two firms' bids in a shared tender are independent, the analogue for markers of the conditional-independence condition of Bajari and Ye (2003). Any excess of observed over placebo markers then comes from colluding pairs. The second assumption fails if shocks specific to a tender move both firms' markers, which would overstate collusion. Two alternative placebos that preserve such shocks yield a lower but still substantial prevalence of collusion.

To pin down the share of colluding pairs, the model restricts how colluding pairs can differ from competing ones: colluding pairs are of different types, and each type is elevated on a set of up to three markers relative to the competitive distribution. Estimating the model places 33.6 percent of the scored pairs in the colluding component. The share is a lower bound: the method misses collusion that leaves no trace in bids, such as agreements not to bid, and placebo bids drawn from tenders where a firm colludes with other partners can raise the competitive benchmark. Collusion is heterogeneous: Eight types account for 84 percent of colluding pairs, the most common involving bids submitted close together in time. Most types use three markers, confirming that the joint distribution of markers matters. Models restricting collusion to use one marker at a time are rejected. Each pair's markers yield a likelihood ratio of collusion against competition, and pairs whose ratio exceeds a threshold are flagged. Given the estimated distributions, this rule is the most powerful test at a given false-positive rate (Neyman and Pearson 1933). When analyzing aggregate statistics on collusion, the estimated share of colluding pairs is used. Comparisons of colluding pairs to other firms use only the top decile of likelihood ratios, where about 7.0 percent of flagged pairs are expected to be false positives, against 37.5 percent at the estimated share.

I validate the method and its classification of flagged pairs (in the top decile of the likelihood ratio distribution) on untargeted moments, features of the data not used by any marker. For each flagged pair, I create a matched control pair from the same district that co-bids as often. First, relative to matched controls, flagged pairs' bids co-move more after accounting for bidder costs and tender

conditions, and their bid functions reject exchangeability more often (Bajari and Ye 2003). Second, flagged pairs are more often related by family: among pairs of firms named after people, 6.3 percent of flagged pairs share a surname, against 4.6 percent of other scored pairs, and the share rises with the pair's likelihood ratio. Third, flagged pairs bid consistently with cover bidding: they enter together more often than matched control pairs when few outside bidders are expected, their joint win rate falls below what competition between the two would predict, and a partner who enters after a deadline extension undercuts a "bid in hand" half as often as a late bid by another firm.

Having estimated the model, I first characterize the footprint of collusion across India. At the estimated share of colluding pairs, three quarters of scored firms belong to at least one flagged pair, and in 43.1 percent of tenders, accounting for 18.4 percent of contract value, a member of a flagged pair wins while its partner also bids. I call these tenders *collusive*: collusion determines the winner and the price. Procuring organizations account for most of the variation across markets: net of the share won by matched control pairs, organization fixed effects account for more than three quarters of the explained variance, compared to less than one fifth explained by district fixed effects. Within state-years and holding the number of firms, tenders and bidders fixed, the share of collusive tenders, compared to state-department buyers, is close to double under rural local bodies and about half under central government buyers. Within pairs, markers of collusion are also higher when the pair bids for a rural local body than when it bids for a state department. Separately, collusion is more prevalent in thin markets: holding the number of tenders in a market fixed, the collusive share is 11.6 percentage points higher in the thinnest quarter of markets than in the thickest. This relation also holds within a market over time, where the buyer is fixed.

I then estimate how much more the government pays for collusive tenders, as flagged by the top decile of likelihood ratios. The outcome is the final contract value relative to the government's engineering estimate, which makes heterogeneous purchases comparable. Within district $\times$ organization $\times$ year and product $\times$ year cells, collusive tenders settle 3.2 percent above other tenders, equivalent to 21 percent of a within-market standard deviation. The premium changes little when controlling for the number of financial bidders or, accounting for potential selection, when restricting the sample to tenders won by bidders that ever collude or to tenders in which a flagged pair bids. For matched control pairs, the estimated price premium is close to zero or negative. While there is substantial heterogeneity, the premium is positive in 90 of the 100 largest procuring organizations. Collusive tenders are 12.1 percentage points more likely to settle within one percent of the government estimate, against a base rate of 41.1 percent, consistent with cartels that face little outside competition (McAfee and McMillan 1992).

Collusion is most common in thin markets, but the premium falls with thinness. Relative to non-collusive tenders in the thickest quarter of markets, collusive tenders cost 3.1 to 3.5 log points more in every quarter, while non-collusive tenders in the thinnest quarter cost 2.3 log points more. The premium therefore falls from 3.5 log points in thick markets to 1.1 in thin ones: thinness raises the price of non-collusive tenders toward the collusive level. Cartels also supply part of the bids that minimum-bidder rules require. On two-bid portals, the late bid that completes an extended tender comes from the flagged partner in 28 percent of cases, against 17 percent for matched control pairs. On three-bid portals, three members of one cartel bid together in 24 percent of tenders with a cartel present, against 6 percent on two-bid portals.

Finally, collusion can affect welfare through the quality of the works procured, through prices, and

through quantity distortions. To measure quality, about 11 thousand tenders are linked to audit records from the national rural roads program PMGSY. Among linked tenders, the inspection grades and completion delays of collusive tenders do not differ from those of other tenders in the same markets. Firms in flagged pairs are positively selected and complete their contracts faster. The welfare cost of collusion is therefore driven by prices and quantities. Quantity distortions are small: If an organization holds its budget fixed, the efficiency loss from the smaller quantity procured is 1.6 percent of the overpayment, so the cost is almost entirely the overpayment itself. Multiplying the price premium by the value of collusive tenders at the estimated share of colluding pairs gives an overpayment of USD 428.6 million a year. If collusion is equally common in tenders that cannot be classified, the overpayment is USD 1.3 billion a year across the portals in the data and, at 8 percent of GDP, USD 4 billion a year across Indian public procurement.

This paper contributes to three strands of literature.

Procurement in low- and middle-income countries. Procurement rules and their enforcement vary widely across low- and middle-income countries (Bosio et al. 2022). Previous work on these settings has focused on the government side of procurement: the effectiveness of procurement organizations and individual officials (Best, Hjort, and Szakonyi 2023), political connections between procurement suppliers and politicians (Lehne, Shapiro, and Vanden Eynde 2018; Colonnelli and Prem 2022; F. Brugués, J. Brugués, and Giambra 2024), officials' discretion in choosing suppliers (Szucs 2024), and electronic tendering and audits as checks on corruption and discretion (Olken 2007; Lewis-Faupel et al. 2016; Gerardino, Litschig, and Pomeranz 2024). Coordination among bidders has received much less attention. Most evidence on bid rigging comes from high-income countries and from markets with a known cartel. Exceptions to this are Kawai, Nakabayashi, and Shimamoto (2025), who find evidence of bid rotation in Indonesia, but their test estimates the share of noncompetitive auctions, not which firms collude. Barkley (2023) studies a single insulin cartel in Mexican public procurement. Closest to this paper is Baránek, Musolff, and Titl (2026) who detect colluding pairs of firms on Ukraine's e-procurement platform from a bidding pattern specific to its multistage auctions. In this paper, I identify colluding pairs in all markets of one procurement system and across levels of government. I can therefore compare collusion across procuring organizations, which account for most of the variation in collusion across markets.

Detecting collusion. Existing tests for bid rigging fall into three groups. The first identifies colluding firms but uses features of one market (Porter and Zona 1999; Athey, Bagwell, and Sanchirico 2004; Ishii 2009; Ishii 2014). Bid-function tests need cost shifters such as a firm's distance to the site (Porter and Zona 1993; Bajari and Ye 2003), which many datasets do not record. Conley and Decarolis (2016) detect groups that manipulate the award threshold of Italian average-bid auctions, and Baránek, Musolff, and Titl (2026) detect pairs that do not undercut each other when bids can be revised in Ukrainian multistage auctions. The second group pools many auctions and estimates how common collusion is, but not who colludes. Kawai, Nakabayashi, Ortner, et al. (2023) reject competition among sets of Japanese bidders, and Chen (2025) estimates the share of Chinese scoring auctions where firms collude with the help of corrupt officials. The third group, supervised screens (Huber and Imhof 2019; Wallimann, Imhof, and Huber 2023; Harrington and Imhof 2022), needs adjudicated cartels as training labels, which are scarce where enforcement is weak and may not represent undetected cartels. The approach in this paper builds instead on large-scale hypothesis testing in statistics

(Efron 2004; Efron 2008) and its use in economics (Kline and Walters 2021; Kline, Rose, and Walters 2024). When applied to collusion, the distribution of markers under competition is unknown, which I overcome by building a placebo distribution from each firm's bids in tenders without its partner. This approach estimates both how many pairs of bidders collude and which ones, without cost data or known cartels. Because the markers are modeled jointly, it incorporates information across several different ways in which a regulator can detect collusion, and is naturally portable to different settings.

The cost of collusion. Overcharge estimates usually come from a prosecuted cartel with known members. In procurement, examples include school milk cartels in Ohio and Texas (Porter and Zona 1999; Caoui 2022) and paving contracts in Montreal after a police investigation (Clark, Coviello, Gauthier, et al. 2018). Outside procurement, they include a ring of stamp dealers at auctions (Asker 2010) and the US generic-drug cartel, which raised prices by about 50 percent (Starc and Wollmann 2025). Each estimate is clean for one cartel but cannot say how much of a market is collusive. Collusion can also affect outcomes other than prices. A Mexican insulin cartel raised diabetes mortality by limiting the quantity of insulin bought (Barkley 2023), and bid rigging in Brazilian federal auctions is associated with worse contract performance and more road accidents (Fazio and Zaldokas 2026). Closest to this paper, Baránek, Musolff, and Titl (2026) find that Ukrainian contracts with a cartel bidder are priced 2.0 percent of the estimated cost higher. Here, collusive tenders settle 3.2 percent above comparable tenders. Because the classification uses bids alone, I estimate both how widespread collusion is and the price premium on collusive tenders for a procurement system as a whole, and relate the two across markets. I also show that the roads built under collusive tenders are of no worse quality.

The remainder of the paper proceeds as follows. Section 2 describes the institutional setting, Section 3 the data, and Section 4 the method to detect collusion and flag pairs. Section 5 validates the classification. Section 6 describes how common collusion is, where it is more common, and what explains its variation across procuring organizations. Section 7 estimates the price premium of collusion, and Section 8 relates prevalence and prices to market thinness and minimum-bidder rules. Section 9 discusses how collusion is related to welfare through quality, prices and quantities. Section 10 concludes.

2 Institutional Setting and Data

2.1 Public procurement in India

Public procurement in India is conducted by four types of procuring organizations: central ministries (e.g., the Ministry of Defence), state governments and their departments (e.g., the Rajasthan Department of Public Works), rural local bodies, and public-sector undertakings (PSUs, government-owned firms, e.g., the Indian Oil Corporation Ltd.).

Procurement is divided into two systems according to the nature of the good or service being procured: standardized goods and services are procured via a *direct-purchase system*, everything else is procured under a *competitive tendering system*. This paper studies procurement conducted under the tendering system. Standardized "common-use" goods and services are listed on the Government e-Marketplace (GeM), an online catalog from which procuring organizations purchase listed items and small orders directly at published rates, and which they are required to use whenever the item is

listed there (Government of India 2017). GeM does not publish transaction-level data on individual purchases. All other procurement—specialized goods, non-standard services, and public works—is allocated through the competitive tendering system via auction.[3] Because its contracts are typically larger and more complex, procuring organizations record and publish detailed administrative and financial information for each tender, including the identity of every bidder and all financial bids. Of the two, the tendering system is fiscally more important: tenders covered in this paper were worth USD 105 billion in the 2024–25 financial year, against around USD 65 billion transacted on GeM nationally (Ministry of Commerce & Industry 2026).

The tendering process is governed by the regularly updated General Financial Rules (GFR), which require government purchases to be conducted in a "transparent, competitive and fair manner, to secure best value for money" (Government of India 2017, Rule 173). State organizations procure under their own financial rules, but these generally follow the same structure as the central government's (e.g. Government of Odisha 2025, Rule 204).[4].

2.2 The competitive tendering process

The tendering process is standardized across portals and organizations and the most important feature is that the contract is allocated to the lowest-price technically qualified bidder in a sealed auction. Each portal publishes the tender notice, the identity of every bidder, the financial bids, and the award; these records are the source of the data used in this paper. A tender begins when the procuring organization publishes a Notice Inviting Tender (NIT), which states the work, the government's engineering cost estimate of what the work should cost (the *government estimate*),[5] the submission deadline, and the terms of participation: an earnest money deposit (EMD), set as a round percentage of the estimate and typically between 1 and 10%; a tender fee; and other details regarding the work (such as the number of days to complete the contract). Procuring organizations may issue corrigenda before the deadline in order to modify contract details or to extend the submission window.

Bidding proceeds in two sealed stages. Bidders submit a technical proposal and a financial bid together. In the first stage, the procurement officer evaluates the technical proposals and admits or rejects bidders based on technical criteria alone. Common criteria for qualification include past experience, financial capacity, technical capacity, certifications, compliance with the tender notice, submission of all required documents, and other administrative requirements. The second stage follows the principle of a first-price sealed-bid auction: the financial bids of the technically qualified and responsive bidders are opened and the contract is awarded to the lowest bidder. This rule is refined in some tenders, most commonly when there are multiple bids at the same lowest price, which occurs in organizations that set a minimum bid as a price floor in order to discourage abnormally low bids. Other, rare, reasons for deviations from the lowest-bid-wins rule include cases in which the award is split across multiple bidders, a bidder withdraws after the technical stage, or the procurement officer rejects the lowest bid for non-compliance with the tender notice. The portal records the financial bids of all bidders that passed the technical stage, including the bids of losing bidders. Figure F.2 in Appendix F shows the notice and the published financial bids for one such tender.

[3] I refer to a single contract let under this system as a *tender*, following Indian usage, and to the procedure by which it is solicited and awarded as *tendering*.

[4] Some states have adopted separate procurement legislation (Government of Rajasthan 2012; Government of Assam 2017; Government of Punjab 2019)

[5] Portals and tender documents refer to this quantity as the tender value; throughout this paper I refer to it as the government estimate.

2.3 Structure of procurement records

Since 2008, procurement has gradually shifted from paper to electronic procurement (eProcurement). eProcurement became compulsory for central government contracts with an estimated value above INR 1 million in 2012, with the threshold successively lowered in subsequent years (Government of India, Ministry of Finance 2014). State governments mandate eProcurement under their own rules, with thresholds between INR 0.2 and 1 million (Panda and Sahu 2015, pp. 9–10). To manage the eProcurement process, the central government and most state governments and PSUs use their own procurement portal under a shared electronic infrastructure, the Government eProcurement System (GePNIC).

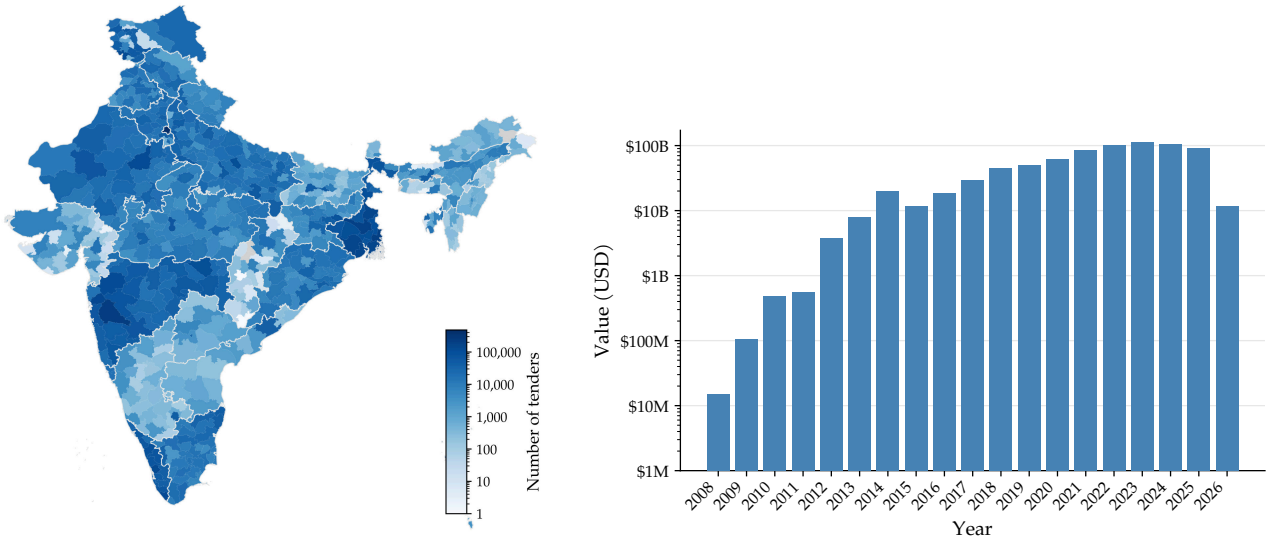
This paper introduces a new dataset assembled by web-scraping the universe of procurement records published on all 76 of these portals: those of 22 states and 8 union territories, which together account for 72% of India's population and 70% of its GDP,[6] together with the central eProcurement system, the defence portal, the rural-roads program PMGSY, and central public-sector undertakings. Thanks to this harmonization, this paper includes procurement contracts across states and organizations. Outcomes can therefore be compared across institutional settings on consistent data. Such records are rarely available elsewhere. In the United States, for example, federal contract awards are reported to a single database, the Federal Procurement Data System, whereas state and local governments publish their procurement records in separate systems and formats (Liscow, Slattery, and Nober 2025).

The records run from 2008, when eProcurement was first rolled out, to the end of the 2025–26 financial year, and cover 15.7 million tenders and 42.3 million bids placed with 4,760 procuring organizations across 659 districts, for a cumulative awarded value of INR 57.5 trillion (USD 747 billion). Rollout was gradual: 98.3% of all tenders are published from 2014 onward. Figure 1 visualizes the sample coverage: Panel (a) shades each of the 659 districts by its tender count; even states outside the GePNIC framework appear, through the central, defence, and rural-roads systems. Panel (b) traces total awarded value by year, from under a billion USD in the early years to about a hundred billion USD by the early 2020s.

The unit of observation is a *tender*, the procurement process for a single contract, and each record has three levels. The tender level gives the procuring organization, the contract type, a free-text title and work description, the location, the estimated cost, the earnest-money deposit, the dates of technical and financial evaluation and of award, and the officials who record each stage. At the level of individual bids, the data record for each submitted bid the bidder's name, time of submission, the quoted value, and the bid's status and remarks. A third layer of the data is the award, which gives the winning bidder and the awarded amount. Table 1 summarizes the tenders used in the price analysis of Section 7 at five levels: bids, tenders, bidders, organizations, and districts.

Detailed cleaning steps are described in Appendix F. In brief, making the data usable requires recovering and resolving location, product, and bidder information from the raw records. Location and product category are not published in usable form, so I recover both from the text, resolving the postal code and free-text location to a canonical district and classifying tenders into 86 sub-categories nested in 9 categories. Bidder names are entered as free text at registration, and the same firm may

[6] The remaining six states (Andhra Pradesh, Bihar, Chhattisgarh, Gujarat, Karnataka and Telangana) run their own e-procurement software, and their state-level tenders are not in the data. Tenders by central organizations and the rural-roads program PMGSY with project sites in these states are included.



(a) Tenders by district

(b) Awarded value by year

Figure 1: Geographic and temporal coverage of the sample

Table 1: Descriptive statistics for the estimation sample

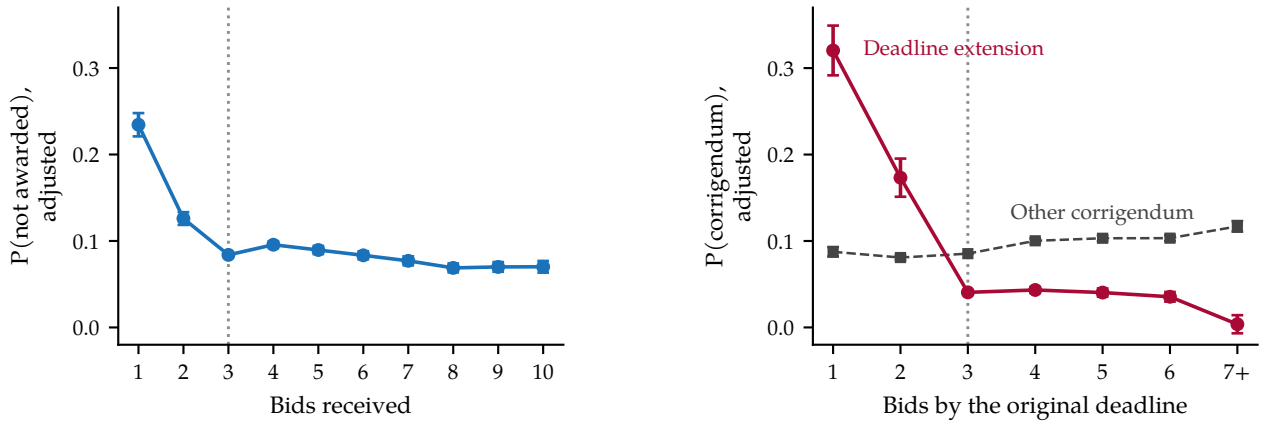
	Mean	Median	SD
Bids (N = 19,418,387)			
Bid price (bid/estimate)	0.934	0.986	0.150
Winning bid price (bid/estimate)	0.919	0.976	0.150
Tenders (N = 6,077,043)			
Tender estimate (USD, thousands)[a]	60.92	12.37	213.5
Bidders per tender[a]	3.678	3.000	2.705
Financial bidders per tender[a]	3.194	3.000	2.141
Bidders (firm × district) (N = 1,831,891)			
Financial bids submitted[a]	9.754	1.000	28.12
Financial bids won[a]	2.458	0.000	7.631
Distinct organizations bid with[a]	1.352	1.000	0.982
Organizations (N = 3,816)			
Tenders[a]	1,089.3	15.00	5,437.3
Total tender value (USD, millions)[a]	146.1	2.218	647.9
Districts (N = 656)			
Tenders[a]	9,010.2	3,028.5	18,401.7
Total tender value (USD, millions)[a]	877.6	436.6	1,470.1

Note: Mean, median, and standard deviation of selected variables at different levels, for financial bids in the sample. Bid value and tender estimate are the corrected values of Section 3.1, and the bid-to-estimate ratio is trimmed to $[0.3, 3]$. The bidder is a firm × district (Section 4). [a]Winsorized at the 99.5th percentile of its own distribution.

appear under several spellings. Because detecting collusion relies on pairs of firms that bid repeatedly across tenders, I first merge these spellings into a single firm (Section 3.1).

2.4 Rational for catel entry: minimum participation and cover bidding

Raising prices, the main goal of collusion in procurement, does not require the two firms to bid together. Instead, for a given lowest bid, having additional entrants is costly because of entry fees. However, minimum-participation rules in Indian procurement mean that a tender with too few bidders may draw additional scrutiny, have its bidding deadline extended, or be re-tendered. Figure 2 shows that tenders receiving one or two bids are more likely to be cancelled or re-tendered than those receiving three or more.[7] Tenders holding one or two bids at the original deadline are also more likely to have that deadline extended, while other adjustments to the tender are independent of the number of bidders. In addition, thin fields gain additional bidders by the final deadline compared to the original deadline (Appendix Figure F.4).



(a) Probability of non-award

(b) Probability of a deadline extension

Figure 2: Tender outcomes by number of bids received

Note: Each point is a fitted probability at the given number of bids received, from a regression of the outcome on a saturated set of bid-count indicators. Three bids are the omitted category. Both regressions absorb the following fixed effects: district-by-organization-by-year and product-by-year. Both regressions control for: indicators for tender type, form of contract, contract type, estimate decile, a cubic in the log tender estimate, an indicator for an estimate above USD 60,000, and the earnest-money rate. Missing estimates and earnest-money rates are set to zero and flagged. Panel (a) covers tenders published at least 24 months before the end of the sample period with a recorded final outcome (cancelled, re-tendered, awarded), and counts all bids the tender received. Panel (b) counts the bids received by the original submission deadline. The dashed series is corrigenda other than deadline extensions, on the same sample. Bars are 95% confidence intervals from standard errors clustered by district.

This gives cartels an incentive to submit cover bids, also called complementary bids, "designed to give the appearance of genuine competition" (Government of India, Ministry of Finance, Department of Expenditure 2017; Competition Commission of India 2024). Having agreed on which member of the cartel wins which contracts, cartels submit not only winning bids, but also deliberately losing bids to reduce the probability of a tender being cancelled, or the deadline extended to invite new competitors. Numerous case records document this practice: for example, in one CCI case, a contractor recruited four firms to file proxy bids in five separate tenders to avoid a deadline extension, and won all five (Competition Commission of India 2018). It is this practice of cover bidding that

[7] Appendix F.2 shows that the rules for cancellations and extensions differ by state. The state portals of West Bengal and Maharashtra and the Coal India portal follow a *three-bid rule*, and twelve portals, among them Uttar Pradesh, Kerala and Tamil Nadu, follow a *two-bid rule*.

makes colluding firms observable as repeated co-bidders, a feature long recognized in the literature on detecting collusion in procurement auctions (Porter and Zona 1993; Ishii 2014; Ishii 2009; Athey, Bagwell, and Sanchirico 2004; Kawai, Nakabayashi, Ortner, et al. 2023; Conley and Decarolis 2016; Chassang, Kawai, et al. 2022; Clark, Coviello, and De Leverano 2025).

2.5 Enforcement

To detect and take enforcement action against collusion, two systems operate in parallel: administrative action by the procuring organization, and antitrust enforcement by the Competition Commission of India (CCI).

Administrative action rests with the procurement officer. Once bids are opened, the officer can reject a bidder at the technical stage, cancel the tender, or re-tender it. As Section 2.4 describes, a tender that falls short of the portal's bidder threshold is subject to additional scrutiny. Procuring organizations are subject to vigilance oversight (Central Vigilance Commission 2021), and portals maintain lists of debarred firms. In total, GePNIC has debarred 179 firms from participating in any tender, with the most common reasons being fraud, non-performance, or safety violations rather than for coordination.

Antitrust enforcement rests with the CCI, which investigates and adjudicates anticompetitive conduct. Between 2010 and September 2026, it issued orders in 97 cases of bid rigging in public procurement, finding a violation in 27 and closing 69 without convictions.[8] These detected cases are a selected sample of the population of colluding bidders, so prevalence cannot be inferred from this case count. To help procurement officers detect and investigate collusion, the CCI also publishes a toolkit of warning signals and case studies (Competition Commission of India 2024). These guidelines are closely related to similar documents published by the OECD and the United States Department of Justice (OECD 2025b; US DOJ 2015). The pair-level markers (Section 3.2) follow the observable signals defined by the CCI toolkit.

3 Constructing the Analysis Data

The procurement records require cleaning to match bidders across tenders and observe pair-level attributes. I resolve bidder names and correct bid values, assemble the co-bid histories of bidder pairs, and measure patterns in those histories that may indicate coordination.

3.1 Cleaning and entity resolution

Bidder names are entered as free text, and the same firm often appears under several spellings. If those spellings are treated as separate firms, the record misses co-bids of the same two firms. Bidder entities are resolved via fuzzy string matching with additional rules to prevent common false positive matches (Appendix F.4.2 gives the procedure). This reduces 1,359,439 raw names to 999,056 firms. Resolution matters for detection because a pair's co-bids are counted over all their spellings: without it, 1.1% of the pairs the model flags would fall short of the minimum number of such tenders needed to be scored, and would leave the scored population (Appendix F.4.2).

[8] The counts cover every antitrust order published on the CCI website. Orders are first screened from their text for procurement and for collusion among bidders. The remaining cases are manually reviewed to confirm that they concern bid rigging in a public procurement and to record the verdict.

Most resolved firms are small. Among active firm-district units, 20.0% have company names, 25.5% bid under an individual's name (e.g., "Vinod Jagannath Yadav"), and 24.8% use a trading name (e.g., "Kalyani Engineering Works"). Only 4.40% of firms match Prowess, the database of Indian company financials, although those firms account for 31.02% of contract value.

Other data cleaning steps include geolocating tenders from their posted Postal Index Numbers and free-text addresses and classifying free-text tender descriptions into product categories. Financial values are sometimes mis-recorded: I convert percentage bids to rupee amounts, reconcile inconsistent value units across tender, bid, and award records, and recover missing winner flags from bid status and remarks. The data appendix documents these steps in detail (Appendix F).

3.2 Pair-level markers of collusion

Collusion is detected at the level of *pairs* of bidders. I define a bidder as a firm within a district: 77.4% of firms bid in only one district, and the mean bidder bids in 2.1 districts. Within a district, a firm with at least 5 bids meets 2.2 procuring organizations on average, so the same pair can co-bid across organizations. A tender in which both members of a pair bid is a *co-bid* of the pair, and a co-bid in which both members' financial bids are opened is a *financial co-bid*. From a pair's co-bids, I construct variables at the level of the tender-pair (such as the difference between the two firms' bids) and at the level of the pair (such as the share of co-bids in which both firms submit their bids at the same time). The full panel holds 41.0 million pairs that co-bid at least once. Because information from few co-bids is noisy, and for computational reasons, I restrict the classification to the 2.04 million pairs with at least 5 financial co-bids, excluding pairs that contain a firm with too few financial bids anywhere to be scored. The restriction is less severe at the tender level than the pair counts suggest: The pairs appear in 58.7% of the tenders with at least two bidders.

Detecting collusion relies on the assumption that cartels leave a trace in their bidding behavior. One reason for this are internal cartel constraints: When a cartel sustains high prices, each member has an incentive to undercut the others, so cartels need ways to enforce their agreements (Stigler 1964). In practice, members might prepare or file their bids together, so that no one can deviate at the last minute; the Competition Commission of India accordingly lists bids "submitted within short span of time" among the warning signals a procurement officer should look for (Competition Commission of India 2024, p. 20). They might also use round bids: a designated winner announcing its planned bid to the others has reason to pick a number that cannot be misheard, and a member bidding arbitrarily above that number has no reason to fine-tune it (Ishii 2014).

A second way to leave traces in the data can originate in the way a cartel allocates contracts among its members. For example, cartels commonly rotate the winner across contracts (Ishii 2009; Kawai, Nakabayashi, Ortner, et al. 2023). Finally, cover bidding – having a designated winner, while the other members stay out or submit complementary bids above the designated bid (McAfee and McMillan 1992; Porter and Zona 1993) – affects the distribution of bids. For example, designated losing bids are less related to firm costs than serious bids (Porter and Zona 1993), and the bids of cartel members might be more strongly correlated than those of other bidders (Bajari and Ye 2003). Cover bidding may also leave too few losing bids just above the winning bid (Chassang, Kawai, et al. 2022; Clark, Coviello, and De Leverano 2025).

The co-bidding panel makes it possible to measure patterns that could reflect such coordination within a pair. To select a set of such *markers*, I follow the warning signals in the Competition Commission of India's *Diagnostic Toolkit Towards Competitive Tenders*, a guide for procurement officers (Com-

petition Commission of India 2024). The markers are based on statistics that can be calculated from public bid records, such as bids submitted close together and a stable ordering of the pair's prices. The method in Section 4 is portable to a different set of markers.

Table 2 lists the twelve markers used in this paper. Each marker summarizes a pair's co-bids and is signed so that the higher value is the direction the Commission's guidelines treat as suspicious. Some markers use the difference or ordering of the two firms' bids. To prevent a mechanical correlation between prices and the measure of collusion, no marker used here uses the outcome in the price regressions, the winning bid or contract value. Appendix A.1 describes, for each marker, how it is constructed, which warning signal in the Commission's toolkit it operationalizes, and, if applicable, how that marker relates to existing methods to detect collusion.

Table 2: Markers

	Marker	Pair-level statistic
1	Submission proximity	Share of co-bids with both submission times recorded in which the two bids are submitted within 30 minutes of each other
2	Round bid difference	Share of the pair's price gaps that fall on a round grid
3	Margin steadiness	Standard deviation of the signed log price gap between the two bids, reversed; defined when both firms have at least ten priced tenders without the partner
4	Submission-order role	$\lvert 2p - 1 \rvert$, with p the share of co-bids in which the lower bidder submits first
5	Submission-order stability	$\lvert 2p - 1 \rvert$, with p the share of co-bids in which a given member submits first
6	Targeted disqualification	Standardized excess of co-bids in which exactly one member is disqualified, over the count implied by each firm's disqualification rate in comparable tenders without the partner; co-bids in which both are disqualified are excluded
7	Outsider disqualification	Mean share of the other bidders disqualified in the pair's co-bids
8	Lower-bidder stability	$\lvert 2p - 1 \rvert$, with p the share of co-bids in which a given member bids lower
9	Lower-bidder rotation	Distance of the lower bidder's switch rate between consecutive co-bids from the rate $2p(1-p)$ implied by the pair's lower-bidder shares
10	Backlog difference	Share of co-bids in which the relatively busier member bids higher
11	Uncompetitive fringe entrant	Share of co-bids with an outside bidder in which at least one outsider has never been the lowest bidder in its national bid record
12	Exclusive outsider	Share of co-bids that include an outside bidder who bids only when both members bid

Note: Each marker summarizes the pair's co-bids, the tenders on which both members bid, and is signed so that the higher value is the direction the Commission's guidelines treat as suspicious. The lower bidder is the member with the lower valid price. A member without a valid price counts as the higher bidder.

4 Detecting Collusion

This section presents a method that detects collusion by combining evidence from several markers. It avoids the main weaknesses of two alternative ways to combine such evidence. The first is machine learning (García Rodríguez et al. 2022; Wallimann, Imhof, and Huber 2023): these methods compute a range of patterns in the bid data and train a prediction model on how those patterns relate to known cases of collusion. This requires an enforcement record to train on, which does not exist at this scale. The second is to combine the markers into an index. Such an index requires the researcher to choose which markers to include, how to combine them, and how to weight each one. The method developed here instead estimates how to combine and weight the markers by comparing the distribution of each marker in the observed data with its distribution under no collusion.

4.1 The detection problem

The data contain pairs of bidders, indexed d. Each pair has the markers defined in Section 3.2. The raw markers m_d are a vector with one entry per marker, such as the share of the pair's price gaps that fall on a round grid. They sit on different scales and are measured with different precision, so each is expressed as its position in the placebo distribution before it is modeled, as Section 4.3 sets out. I write y_d for the standardized markers. Their joint density across pairs is $f(y)$.

Each pair has a latent state $C_d \in \{0, 1\}$, equal to one if the two firms collude with each other and zero if they do not, and a share $\Lambda = \Pr(C_d = 1)$ of pairs collude. Under independent bidding the density of a pair's markers is f_0; under collusion it is f_1. The observed density in the population is thus a mixture of two distributions:

$$f(y) = (1 - \Lambda) f_0(y) + \Lambda f_1(y). \tag{1}$$

Equation (1) is the two-groups model of large-scale multiple testing, applied to pairs of bidders instead of genes or hypotheses (Efron et al. 2001; Efron 2008). Neither component is directly observed: the data provide one sample, drawn from f, while Λ, f_0 and f_1 are all unknown. It is convenient to work not with the two densities but with their ratio. Dividing through by f_0 gives

$$g(y) \equiv \frac{f(y)}{f_0(y)} = (1 - \Lambda) + \Lambda L(y), \qquad L(y) \equiv \frac{f_1(y)}{f_0(y)}, \tag{2}$$

which is increasing in the likelihood ratio L, with slope Λ. Expressing the mixture as this ratio shows that the full density under collusion does not need to be specified, only the likelihood ratio L. The ratio is also the easier object to estimate: g can be recovered by discriminating a sample drawn from f from a sample drawn from f_0, with neither density estimated at any point (DiNardo, Fortin, and Lemieux 1996; Qin and Zhang 1997).

Estimating Λ and L takes three steps. The first step is to find a valid draw from the competitive component of the distribution, f_0, so that g can be estimated from two samples: the observed data f and the competitive data f_0. Estimating a mixture from one sample of the mixture and one sample of a single component is the setting of case-control studies with contaminated controls in econometrics and of learning from positive and unlabeled data in machine learning (Lancaster and Imbens 1996; Elkan and Noto 2008; Blanchard, Lee, and Scott 2010). Kline and Walters (2021) and Kline, Rose, and Walters (2022) use related mixture methods to bound the share of employers that discriminate. To construct a draw from the competitive component, I generate placebo pairs by reshuffling observed data in a way that preserves firms' bidding habits but breaks pair-specific collusion.

The second step is to parametrize L. With f_0 known and f_1 unrestricted, the data identify only a lower bound on Λ, because any part of f_1 that resembles f_0 can be assigned to either component (Genovese and Wasserman 2004; Meinshausen and Rice 2006; Patra and Sen 2016). Imposing structure on L restores point identification (Teicher 1963; Yakowitz and Spragins 1968). The parameters of the likelihood ratio are directly interpretable and informative about which markers colluding pairs use in the Indian procurement data.

The third step is to estimate Λ and the parameters of L jointly by maximum likelihood. For pair d, the fitted $L(y_d)$ is the likelihood ratio of collusion to competition. By the Neyman–Pearson lemma, it is the most powerful statistic for separating the two hypotheses at any given test size (Neyman and Pearson 1933). Pairs can therefore be ranked by L and cut at a given threshold to classify them as colluding or competitive.

4.2 Step 1: f_0 as placebo sample

I approximate the competitive distribution f_0 by resampling the observed bidding data. For pair $d = (i, j)$ in tender t, the data record attributes (x_{it}, x_{jt}), from which the markers m_d are computed. To fix ideas, consider the marker "submission proximity". Its attribute is the firm's submission time, measured as the distance to the deadline; for every co-bid of i and j, the marker takes the gap between the two times and records whether it falls below 30 minutes. The marker is the share of co-bids of pair d in which the two bidders submit within 30 minutes of each other.

The placebo replaces every attribute of a firm in a co-bid with the same firm's attribute from a randomly drawn tender in which its partner was absent, as Figure A.1 illustrates. In tender t, $(\tilde{x}_{it}, \tilde{x}_{jt})$ replaces (x_{it}, x_{jt}), where $\tilde{x}_{it}$ is drawn at random from i's bids in tenders where j did not bid, and $\tilde{x}_{jt}$ from j's bids in tenders where i did not bid, so that the two submission times being compared are times the firms chose when they were not facing each other. Consider a pair of bidders that often submit their bids within 30 minutes of each other. If both bidders usually bid very close to the deadline, their placebo reproduces that pattern: their submission times taken from other auctions are as close together as the actual submission times in their co-bids. As a consequence, the observed distribution and the placebo distribution are similar. By contrast, consider a pair of bidders that bids at scattered times in tenders without its partner but close together in tenders with its partner. Reshuffling the data breaks this pattern, and hence the pair's observed rate of close submission times is higher than its placebo rate. Repeating this for every co-bid of every pair and recomputing the marker values m_d produces the placebo sample. The same principle is used for the other markers, with different attributes re-drawn in the same way: the bid values for markers using prices, the disqualification outcome in other tenders for the markers related to technical disqualification, and so on. Because the draws for the two bidders are independent, the placebo holds each bidder's own behavior at its observed distribution and removes the dependence between the two.

This benchmark is valid under two conditions, derived formally in Appendix A.2. The first condition is that firm i's bids in tenders without j must describe how i would have bid in their co-bids, had j not bid.[9] The main way this condition fails is if the co-bids of colluding pairs are systematically different from the tenders in which only one of them bids, in ways that are correlated with each other. Table A.3 therefore compares, within pair, characteristics of the pair's co-bids with the tenders used for its placebo draws. The two sets are similar. Given the sample size, every difference is statistically significant, but eleven of the twelve are below 0.1 standard deviations and all are below 0.2. The largest difference is in the number of bidders.

The second condition is that, absent collusion, firm i's bid and firm j's bid in the same tender are independent, conditional on their individual historical habits. This is a strong assumption, and it is violated by tender-level shocks that affect the bids of both firms.[10] Drawing bids from tenders in which i bids without j removes this tender-level dependence. Therefore, where the condition fails, the measured excess includes a competitive component.

Appendix A.2 repeats the estimation under two alternative placebo designs that do not need this condition: one reshuffles the bidders within the same tender, the other replaces each co-bid with a co-bid of one member with a different partner. Both rank pairs similarly to the headline fit (rank

[9] The pool need not be free of collusion: if i colludes with a third firm, that shifts i's own bids, and the placebo draws from those bids and inherits the shift (Appendix A.2.1).

[10] Such shocks may be irrelevant for some markers—submission times should not be related to tender-level heterogeneity—but the condition may be violated for markers based on firms' bids.

correlations 0.59 to 0.74) but place fewer pairs in the colluding component, 18.1 and 18.4 percent (Table A.4). Price regressions on either alternative fit give quantitatively similar results (Appendix Figure D.1).

4.3 A common scale for the markers

The markers are measured on different scales. To make them comparable, within marker k and bin b of the number of co-bids (5–8, 9–15, 16–30, 31–60, 61–120, or more than 120), I replace the raw marker m_{dk} by its position in the placebo distribution of the same marker and bin,

$$y_{dk} = \Phi^{-1}\left(\widehat{F}_{0,kb}(m_{dk})\right), \tag{3}$$

where $\widehat{F}_{0,kb}$ is the placebo distribution function. The inner term is the pair's percentile in the placebo and Φ^{-1} maps it to the standard normal scale, so that a pair with $y_{dk} = 1.96$ sits at the 97.5th percentile of placebo pairs with about as many co-bids.[11] By construction, the placebo scores are standard normal (Rosenblatt 1952; Angus 1994). Ties are randomized within their mass point in both arms, which preserves this property for markers with mass points (Rüschendorf 2009; Czado, Gneiting, and Held 2009). The vector $y_d = (y_{d1}, \ldots, y_{dK})$ is what enters the mixture in (1).

Because the placebo removes dependence introduced by tender-level shocks along with collusion, the bulk of a marker's observed scores can be shifted relative to the placebo even without collusion. Where that happens, I recenter and rescale the observed scores so that their median is zero and the distance between the median and the lower quartile is that of a standard normal, a quantile-based version of the empirical null in large-scale testing (Efron 2004; Jin and Cai 2007).[12] Six of the twelve markers are recentered with this method: submission proximity, both submission-order markers, both disqualification markers, and lower-bidder stability.

Finally, I drop markers that show no excess in the upper tail. For each marker, I fit the one-marker version of the model in Section 4.4, a single upward shift θ_k and the share of pairs it applies to, and keep the marker only if the estimated shift exceeds 0.05 standard deviations, which holds for every marker except submission-order stability.

Figure 3 shows the distribution of the observed and the placebo pairs for all twelve markers. This figure is where the identification can be read off directly. The placebo density is the competitive benchmark f_0 and the observed density is the mixture f, so the vertical distance between the two curves at a point y is the density ratio $g(y) = f(y)/f_0(y)$ of (2). Where the curves coincide, the observed pairs are no more concentrated than under independent bidding, and $g = 1$. For most markers, there is excess mass in the right tail of the distribution, and $g > 1$. The $\hat{\theta}_k$ and $\hat{\lambda}_k$ printed in each panel measure by how much colluding pairs are elevated compared to the placebo distribution, and what share of pairs is estimated to use the marker, respectively.

4.4 Step 2: the likelihood ratio

Equation (2) cannot be estimated without restrictions on L: f_1 is a multi-dimensional joint distribution over all markers, and with f_0 known and f_1 unrestricted, Λ is not point-identified (Genovese and Wasserman 2004; Meinshausen and Rice 2006; Patra and Sen 2016).

[11] The binning matters because the same share provides different evidence for different sample sizes: a pair that submits together three out of six times is more likely to arise by chance than a pair that submits together thirty out of sixty times.

[12] Appendix A.3 gives the rule that selects which markers are recentered, the estimator, and what the repair assumes.

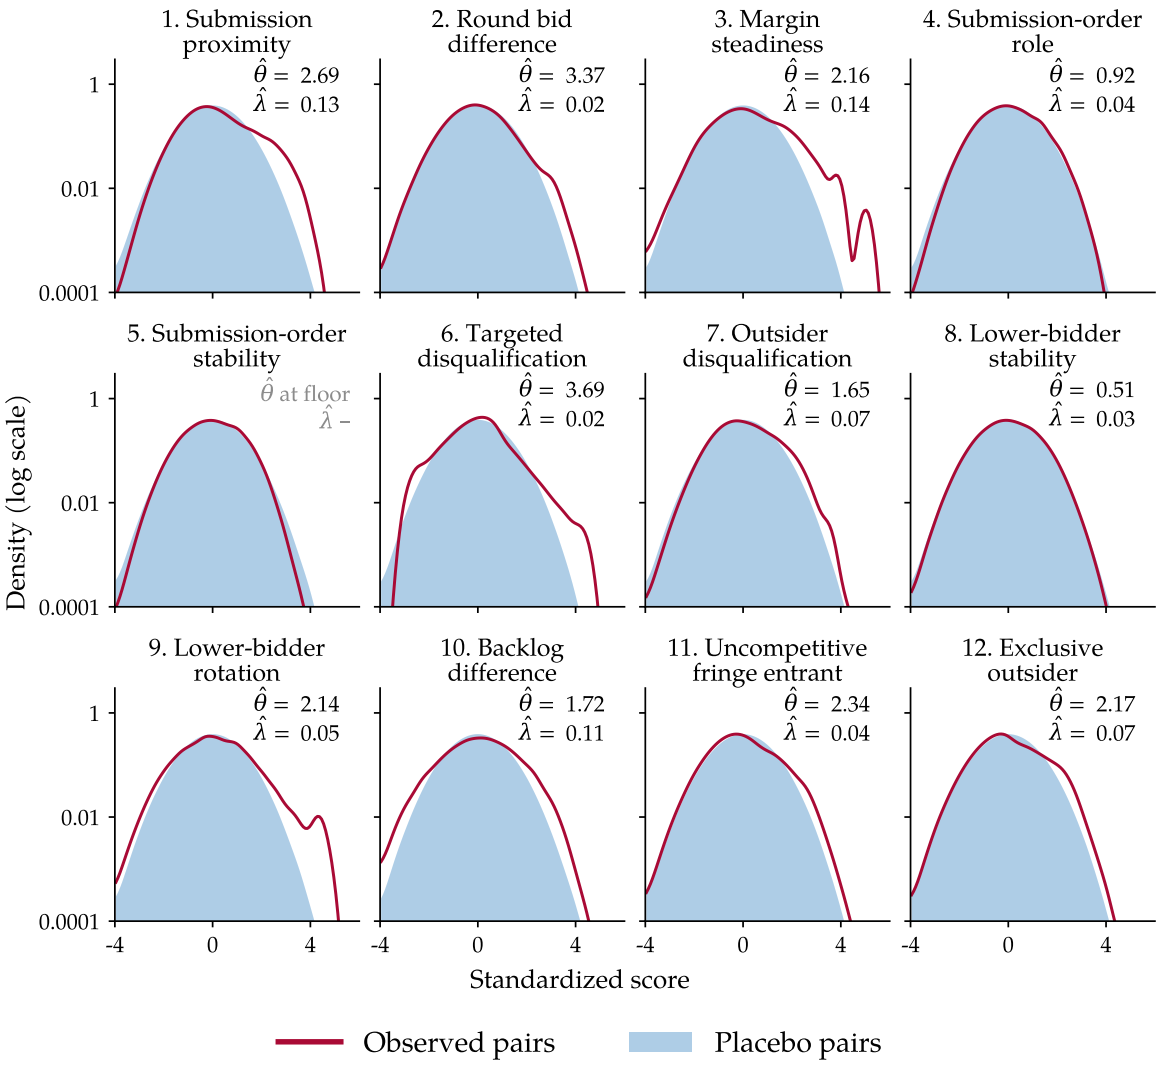


Figure 3: Observed and placebo pairs, by marker

Note: Each panel shows the log density of one marker's standardized score, comparing observed pairs (red line) to pooled placebo pairs (standard normal by construction, light blue area). Densities are smoothed with a Gaussian kernel with a bandwidth of 0.2. Each panel reports a fit from the mixture model with just one marker, using just the shift $\hat{\theta}_k$ and the share $\hat{\lambda}_k$ of pairs in the collusive component. When $\hat{\theta}_k$ is estimated to be below 0.05, the marker is left out of the joint model and $\hat{\lambda}_k$ is set to zero.

The structure imposed on L assumes that colluding pairs differ in which markers they use, and that each uses only a few markers. Let a colluding pair be of latent type S: the subset of active markers that the pair uses. Conditional on type S, markers $k \in S$ shift upward by $\theta_k > 0$ standard deviations, while markers not used by type S follow the null distribution. Given that the placebo distribution is standard normal, for a pair whose available markers are G_d, the type-specific likelihood ratio is[13]

$$L_S(y_d) = \prod_{k \in S \cap G_d} \exp\left(\theta_k y_{dk} - \tfrac{1}{2}\theta_k^2\right).$$

Letting π_S be the share of colluding pairs of type S, averaging over types gives

$$L(y_d) = \sum_{S \neq \emptyset} \pi_S L_S(y_d). \tag{4}$$

With twelve markers there are $2^{12} - 1 = 4095$ possible types. I restrict types to at most three markers and drop the one marker that does not pass the pre-fit screen of Section 4.3. This leaves $\sum_{s=1}^{3} \binom{11}{s} = 231$ possible types.

I further assume that the shares follow the log-linear model

$$\pi_S = \frac{1}{Z} \exp\left(\sum_k \alpha_k s_k + \sum_{k<j} \beta_{kj} s_k s_j\right), \qquad s_k = \mathbf{1}\{k \in S\}, \tag{5}$$

where Z normalizes over nonempty types. The coefficients α_k are log-linear weights. The interactions β_{kj} allow two markers to be more or less likely to occur in the same latent type than independent adoption would imply. In other words, α_k represents how commonly a marker is adopted in general, and β_{kj} represents whether markers k and j are complements or substitutes.

4.5 Step 3: estimation

I estimate $(\Lambda, \theta, \alpha, \beta)$ from the observed and placebo samples jointly. To do so, I stack the observed pairs, labeled $a_r = 1$, and the placebo pairs, labeled $a_r = 0$. Following DiNardo, Fortin, and Lemieux (1996) and Qin and Zhang (1997), the model can be estimated as a logit specification in terms of the density ratio g. When two samples are drawn from densities whose ratio is g, the log odds of sample membership equal $\log g$ up to a constant that absorbs the sampling proportions (Anderson 1972; Prentice and Pyke 1979). This property allows g to be estimated directly. The log odds that observation r belongs to the observed sample are:

$$\log \frac{\Pr(a_r = 1 \mid y_r)}{\Pr(a_r = 0 \mid y_r)} = \log g(y_r) = \log\left[(1 - \Lambda) + \Lambda L(y_r; \theta, \alpha, \beta)\right]. \tag{6}$$

Instead of estimating the high-dimensional distributions f or f_0 in levels, the estimator solves a classification problem: the log odds that an observation belongs to the observed data rather than to the placebo sample directly yield the parameters of the density ratio. Defining the density ratio at observation r as $g_r = (1 - \Lambda) + \Lambda L(y_r; \theta, \alpha, \beta)$, the parameters maximize the log-likelihood:

$$\ell(\Lambda, \theta, \alpha, \beta) = \sum_r \Big[a_r \log g_r - \log\left(1 + g_r\right)\Big] \tag{7}$$

[13] To see this for a single marker, note that if y_{dk} is standard normal under competition and normal with mean θ_k and unit variance under collusion, the ratio of the two densities is $\exp\{-\frac{1}{2}(y_{dk} - \theta_k)^2\} / \exp\{-\frac{1}{2}y_{dk}^2\} = \exp(\theta_k y_{dk} - \frac{1}{2}\theta_k^2)$.

Intuitively, the estimator pushes the odds g_r higher for observed data and lower for placebos, reaching its optimum where g_r equals the true density ratio f/f_0. Because the observed data are a mixture, this ratio expands to $f/f_0 = (1-\Lambda) + \Lambda(f_1/f_0)$. Therefore, solving this classification problem directly recovers the mixing weight Λ and the shape of the collusive distribution relative to the baseline (f_1/f_0), without having to estimate f_0 or f_1. The model is estimated by maximizing Equation (7) numerically. Because the likelihood of a mixture model is theoretically unbounded, I constrain the optimization routine to a bounded parameter space. Estimation is implemented using analytic gradients and multiple random starting values to ensure convergence to the consistent local maximum.

4.6 Estimation results

The estimated share of colluding pairs is $\widehat{\Lambda} = 33.6\%$ (district-clustered standard error 1.8 percentage points). A possible concern for the estimation is overfitting. To address this, I refit the model on a random set of 80% of districts and evaluate its fit on the held-out 20% of districts. The held-out districts retain 98.9% of the in-sample gain in log-likelihood over an intercept-only model.

Table 3 documents the largest estimated type shares π_S. Collusion is heterogeneous: the eight most common types account for 84% of all colluding pairs. The largest type, at 24%, uses submission proximity together with a consistent submission-order role and a steady margin. These pairs coordinate in three ways. First, they submit within half an hour of each other. Second, the cheaper bid is consistently the earlier one, or consistently the later one. Third, the gap between the two prices varies little from one co-bid to the next. The second type accounts for 20% of colluding pairs: one member is usually the cheaper of the two, the lower bidder rotates more often than under chance, and the pairs often bid with an exclusive outside bidder that never bids without the pair. Across all types, 78% of colluding pairs use three markers at once and the rest two, with no mass on single-marker types. The order role (having the lower bidder always submit first or always submit last) is used by most types.

Table 3: Most common collusion types

Pattern S	Weight	S.E.
Low-bid stab. + Low-bid rot. + Exclusive	23.5%	(4.1)
Prox. + Order role + Margin stab.	19.6%	(2.5)
Order role + Backlog	16.8%	(1.3)
Order role + Outsider disqual. + Margin stab.	5.6%	(0.9)
Outsider disqual. + Targeted disqual.	5.1%	(0.7)
Order role + Exclusive + Backlog	4.9%	(1.0)
Prox. + Outsider disqual. + Margin stab.	4.2%	(1.3)
Low-bid stab. + Order role + Margin stab.	4.1%	(0.9)

84% of the fitted pattern mass.

Note: The table reports the largest estimated shares π_S from estimating the model in Section 4.4. No mass falls on one-marker types, 22% on two-marker types, and 78% on three-marker types. Standard errors in percentage points, by the delta method from the sandwich estimator with scores clustered by district, evaluated at the estimated optimum and conditional on the standardized markers.

This result is informative about how collusion works in practice: Types combine markers, so running a separate test for each marker and pooling the results might not detect collusion in all cases. For example, a pair that is moderately elevated on three markers at once can be classified as colluding although no single marker would register it as suspicious.

To determine which markers matter most, I compute Shapley values (Shapley 1953). Because the markers enter the collusion types jointly and are correlated, the contribution of one marker cannot

simply be read off from the type shares. The Shapley value solves this by averaging: the markers are added to the fitted model one at a time, in every possible order. The importance of each marker is the average change it causes in the outcome of interest.

The Shapley values are computed for two quantities of interest. The first is fitted prevalence: how much higher the model's estimated prevalence of collusion is with a given marker than without it. The second is separation: how much the gap in the fitted mean log likelihood ratio between observed and placebo pairs grows when the marker is added to the model. Appendix A.5 gives details. The results, reported in Table 4, show that submission proximity is the most important marker: it accounts for 27.9% of separation and 15.7% of the prevalence attribution. Margin steadiness, lower-bidder rotation, and having an exclusive outsider are the next most important markers. The two attributions diverge most for the submission-order role, which carries the second-largest prevalence share of any marker, 17.5%, but almost none of the separation, 0.3%: it enters many of the fitted types, but its shift of 0.17 standard deviations is too small to move a pair's likelihood ratio. Round bid difference is the mirror case. It has the largest fitted shift, 4.68 standard deviations, but enters no type with meaningful weight and contributes nothing to either attribution.

Table 4: Shapley decomposition of fitted prevalence and likelihood-ratio separation

		Prevalence		Separation	
Marker	$\hat{\theta}_k$ (S.E.)	ϕ_k (pp)	%	ϕ_k (log LR)	%
1. Submission proximity	2.71 (0.06)	4.29	15.7	0.1971	27.9
3. Margin steadiness	2.11 (0.04)	2.06	7.6	0.1589	22.5
12. Exclusive outsider	3.02 (0.06)	6.16	22.6	0.0965	13.7
9. Lower-bidder rotation	4.33 (0.47)	1.68	6.2	0.0795	11.3
10. Backlog difference	1.55 (0.06)	1.50	5.5	0.0543	7.7
7. Outsider disqualification	1.78 (0.16)	2.75	10.1	0.0476	6.7
6. Targeted disqualification	3.20 (0.20)	0.73	2.7	0.0344	4.9
11. Uncompetitive fringe entrant	2.95 (0.12)	0.92	3.4	0.0211	3.0
8. Lower-bidder stability	1.45 (0.19)	2.40	8.8	0.0153	2.2
4. Submission-order role	0.17 (0.04)	4.77	17.5	0.0020	0.3
2. Round bid difference	4.68 (0.00)^{w}	0.00	0.0	0.0000	0.0
5. Submission-order stability	–	–	–	–	–
Total		27.26	100.0	0.7066	100.0

Note: Shapley allocations from the headline fit, ordered by separation. $\hat{\theta}_k$ is the fitted shift of marker k, in placebo standard deviations. The prevalence columns allocate v_Λ in its row-gated form, the separation columns allocate v_L. Both are defined in Appendix A.5. Submission-order stability does not pass the pre-fit screen of Section 4.3, is not in the fit, and receives no allocation. Standard errors of $\hat{\theta}_k$ in parentheses: sandwich estimator with scores clustered by district and the Hessian at the estimated optimum, conditional on the standardized markers. The Shapley allocations are point estimates. ^{w} Round bid difference loads 100.0% on directions of the Hessian that the data do not identify. Its standard error covers only the identified part.

4.7 Ranking and decision rules

For each scored pair d, the fitted model gives a likelihood ratio $L(y_d)$: how much more likely the pair's markers are under the colluding component than under the competitive one. Combined with the fitted share Λ, it gives the posterior probability that the pair colludes,

$$\Pr(C_d = 1 \mid y_d) = \frac{\Lambda L(y_d)}{(1-\Lambda) + \Lambda L(y_d)}. \tag{8}$$

The posterior is increasing in L, so both put the 2,043,939 scored pairs (9.5% of co-bidding pairs of firms with at least two financial bids in the district) in the same order. Under the fitted model, the Neyman–Pearson lemma implies that a threshold on L is the most powerful classifier. Given that the log likelihood ratios of observed and placebo pairs overlap through most of their range, any threshold misclassifies some pairs. I calculate the false discovery rate—the expected share of flagged pairs that are in fact competitive—for different thresholds.[14] For the remainder of the paper, I conservatively flag pairs in the top decile of the likelihood ratio as colluding, except for the aggregate prevalence of collusion in Section 6.1, where I flag the estimated share $\widehat{\Lambda}$ of pairs. The top-decile rule flags 204,394 pairs with a model-implied false discovery rate of 7.0%. These pairs are the flagged pairs used in the validation and price regressions that follow.

4.8 Interpretation and limitations

The method detects repeated, pair-level collusion that leaves a trace in the bids of colluding pairs. The limitations below are common to most other methods that detect collusion at the level of firm pairs.

First, several forms of collusion are not detected by this approach: firms might divide up a market (and never bid together), suppress the bids of outside bidders, or collude with procurement officials in the design of the tender. Second, the classification into colluding and competitive pairs is fixed over time. In practice, collusive arrangements can be established over time, or collusion can break down. The model scores all pairs on their full history, which weakens the evidence against colluders. Third, this method requires a sufficient history of co-bidding. In this sample, only 9.5% of co-bidding pairs of firms with at least two financial bids in the district have enough co-bids and active markers to be scored. Tenders with at least one scored pair account for 33.8% of total tender value. While repeated interaction is robustly associated with more collusion across many contexts, including in the classification here, collusion between firms that only rarely co-bid on public procurement contracts is possible.

Together, these limitations imply that the classification into flagged and non-flagged pairs is measured with error. Under the additional assumption that false positives and false negatives are unrelated to prices conditional on the regression controls, this measurement error attenuates the coefficient by a factor of approximately one minus the sum of the false-negative and false-positive rates (Aigner 1973; Bound, Brown, and Mathiowetz 2001). The estimated coefficients will therefore be lower bounds.

5 Validation

Section 4 ranks every scored pair by its likelihood-ratio evidence for collusion and classifies the top decile as flagged. This section validates the ranking using untargeted moments of the data.

Flagged pairs are systematically different from unflagged pairs. In particular, they co-bid much more often than the typical pair of bidders. To compare flagged pairs, which are likely colluding, with unflagged pairs, I build a set of *matched control pairs* by drawing, for each flagged pair, a single unflagged pair in the same district whose co-bid count is within 30 percent of its own.[15]

[14] Thresholds for common decision rules are reported in Appendix A.6.

[15] The match is drawn at random among the candidates, with probability declining in the candidate's suspicion rank within that stratum: an unweighted draw would fill the matched control pairs with pairs sitting just below the cut, which are

5.1 Excess bid correlation

Bajari and Ye (2003) show that under competitive bidding with independent private costs, bids satisfy two conditions: two firms' bids are independent conditional on firm- and auction-level cost shifters (*conditional independence*), and the relation between cost shifters and bids does not depend on the firm's identity (*exchangeability*). Both can be tested on the residuals and slopes of a bid function, here estimated as

$$\log\left(b_{it}/\hat{v}_t\right) \;=\; X_{it}'\beta \;+\; \alpha_i \;+\; \tau_t \;+\; \varepsilon_{it}, \tag{9}$$

where b_{it} is the bid firm i submits in tender t, $\hat{v}_t$ the government estimate, and α_i a firm effect, defined for the firm within the district. X_{it} is a vector of six cost shifters: three related to the bidder itself (distance, backlog, capacity utilization), together with the competing bidders' minimum distance and mean backlog, and the log estimate of the tender. τ_t is a tender fixed effect, which absorbs the estimate, every other characteristic of the tender, and any cost shock common to every bidder in it.[16]

The test for conditional independence uses the residuals $\hat{\varepsilon}_{it}$: for a pair ij, ρ_{ij} is the correlation of the two members' residuals over the pair's co-bids. To account for the mechanical negative correlation between residuals in tenders with few bidders, each pair is compared to its own field in the same tenders. The null hypothesis is therefore no *excess* correlation over the field:[17] $H_0\colon \rho_{ij} = \rho_{-ij}$, where ρ_{-ij} is the average of each member's correlation with the rest of the field in those same tenders, thus netting out shocks that might affect all bidders alike. This null is violated under coordination, where the two members' residuals co-move beyond the field, so the excess is positive. The test statistic is the Fisher z transformation of the correlations, so each pair returns one p-value against the one-sided alternative of positive excess.

The second test is for *exchangeability*. For each pair, (9) is estimated on the pair's co-bids, with separate intercepts and slopes for the two members. An F-test checks the restriction that $\beta_i = \beta_j$ for each cost shifter, first on the own shifters and then on all six.

For both tests, the question is whether the test rejects more often for flagged pairs than for matched control pairs.

Table 5 reports the results. Panel A tests for conditional independence. The residuals of flagged pairs are positively correlated and those of matched control pairs are negatively correlated. Against their other co-bidders, the excess correlation of flagged pairs is nearly four times that of matched control pairs (0.43 against 0.11), and the test rejects conditional independence at more than twice the rate (52 against 25 percent). Panel B tests for exchangeability. The test rejects the null for flagged pairs at nearly twice the rate for matched control pairs: 15 against 8 percent on the own shifters, and 28 against 16 percent on all six. Both tests reject more often for flagged pairs than for their matched control pairs, as expected if the detection method identifies coordinating pairs.

---

the pairs most likely to be coordinating without having been flagged. Flagged pairs whose stratum holds no unflagged candidate have no match and drop out of the comparison.

[16] With a tender fixed effect, the residuals of a two-bidder tender are mechanically perfectly negatively correlated, so two-bidder tenders are dropped.

[17] Formally, $\rho_{-ij} \;=\; \frac{1}{2}\left(\rho_{i,-ij} + \rho_{j,-ij}\right)$.

Table 5: Residual co-movement and exchangeability, by arm

	Flagged	Control	Δ	(s.e.)
Panel A. Conditional independence: H_0: $\rho_{ij} = \rho_{-ij}$ (71,829 matched pairs)				
Pair correlation (ρ_{ij})	0.13	−0.10	0.23	(0.003)
Member–field correlation (ρ_{-ij})	−0.30	−0.21	−0.09	(0.001)
Excess ($\rho_{ij} - \rho_{-ij}$)	**0.43**	**0.11**	**0.32**	(0.004)
Share rejecting H_0	0.52	0.25	0.26	(0.002)
Panel B. Exchangeability: H_0: $\beta_i = \beta_j$ (28,557 matched pairs)				
Share rejecting H_0, own shifters	0.15	0.08	0.07	(0.003)
Share rejecting H_0, all six shifters	0.28	0.16	0.12	(0.003)

Note: All numbers constructed by estimating (9) with bidder and tender fixed effects. Flagged pairs are the top decile of scored pairs by the estimated likelihood ratio (Section 4). The sample is the flagged pairs and their matched control pairs, kept where both the flagged pair and its matched control pair meet the sample rule: at least 10 co-bids in fields of at least three bidders (Panel A), and at least 20 co-bids (Panel B). "Flagged" and "Control" are means across the pairs of that arm. Δ is the flagged mean minus the control mean. (s.e.) is the two-sample standard error of that difference.
Panel A: Correlations ρ_{ij} are computed on the residuals of the bid function for each pair. ρ_{-ij} is each member's pooled pairwise correlation with every other bidder in the same tenders, averaged over the two members. "Share rejecting H_0" is the share of the arm's pairs rejected against $\rho_{ij} > \rho_{-ij}$, on Fisher's z transforms with standard error $\sqrt{1.5/(n-3)}$ and n the pair's co-bid count.
Panel B: The bid function is estimated for each pair on its co-bids, with separate intercepts for the two members and tender effects removed within the pair, once with the two members' slopes constrained to be equal and once unconstrained. For each pair, the p-value of the F-test on the constraint is calculated. "Own shifters": backlog, capacity utilization, distance. "All six shifters": adds the log estimate, rival minimum distance and rival mean backlog.
In both panels, the share rejecting is the share of the arm's pairs rejected at a false discovery rate q of 10 percent, applied within arm on that panel's sample. Table B.1 repeats the table on the flagged pairs and matched control pairs present in both panels.

5.2 Surname sharing

Previous work on road-paving cartels in Italy showed that several cartels included firms owned by members of the same family (Conley and Decarolis 2016). For India, an independent audit of procurement practices found bidders with a family relationship entering together, raising suspicions of collusive behavior (Comptroller and Auditor General of India 2023). Family and ownership ties are not among the twelve markers, so shared surnames, used as a proxy for those ties, provide an out-of-sample test. In 916,044 scored pairs (45 percent of all classified pairs), both bidders are named after a person.[18]

Figure 4 plots the share of pairs sharing a surname and the share sharing a rare surname, residualized with bidder fixed effects, against the percentile of the pair's likelihood ratio. The share of pairs sharing a surname rises with the likelihood ratio and is especially elevated in the top decile. Among flagged pairs, 6.3 percent share a surname and 3.1 percent a rare surname, against 4.6 and 1.2 percent among unflagged pairs. Regression results are reported in Appendix Table B.2. The out-of-sample test validates the method: Without relying on bidder names, the method flags pairs that are more likely to share a surname.

[18] I classify pairs of bidders as sharing a surname from the free-text bidder names: A bidder is named after a person if its name, apart from honorifics and initials, contains no legal, trade or place word. Two bidders share a surname if their surnames agree, and share a rare surname if, in addition, fewer than 1 percent of the state's scored bidders have the same surname.

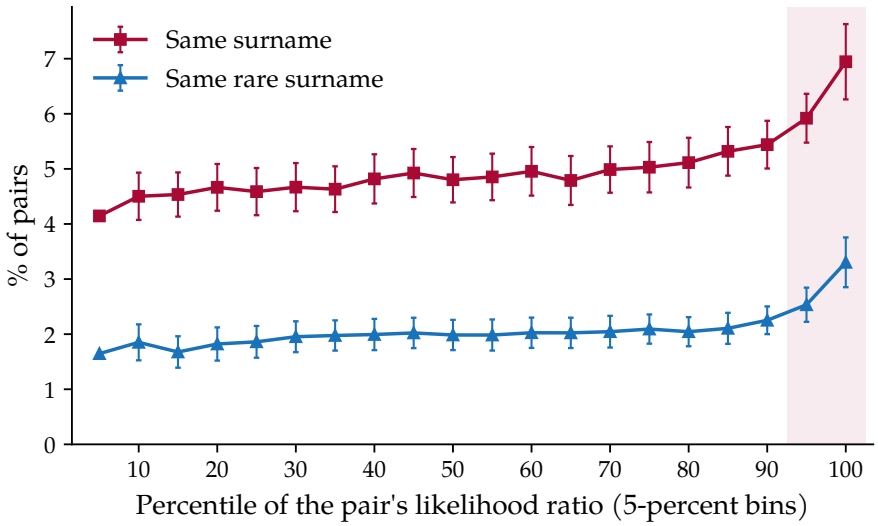


Figure 4: Shared surnames by the pair's likelihood ratio

Note: Share of pairs whose two bidders have the same surname, or the same rare surname, by 5-percent bin of the pair's likelihood ratio. The sample is scored pairs in which both bidders are named after a person (916,044 pairs). Estimates from a regression of the outcome on bin indicators and the log number of co-bids, with bidder fixed effects. The bottom bin is drawn at its raw share and each other bin at that share plus its estimated difference. Whiskers show 95 percent confidence intervals, clustered by district. The shaded area marks the flagged pairs (top decile).

5.3 Evidence for cover bidding

Coordination implies that a pair decides jointly whether its second member bids. The second bid – a designated losing bid placed to help a tender reach a minimum bidder threshold – does not respond to tender conditions the way a competitive bid does. Joint entry of flagged pairs is concentrated in thin fields. Within pair, relative to tenders with three or more outside bidders, flagged pairs' joint entry rises by 32 percentage points more than that of their matched control pairs when no outside bidder bids, and by 14 points more with one. These are the tenders in which a second bid matters most, since a tender with too few bids is extended or cancelled. The realized field is itself an outcome, so this is a correlation. When the outside field is instead predicted from tender observables, the gap keeps its sign but is much smaller (1.6 and 0.5 percentage points), in part because the prediction is noisy. Appendix B.3 describes the pair-tender panel and the regression in detail (Table B.3).

Under cover bidding one member is the designated winner and the other submits a deliberately non-competitive bid. With symmetric bidders, under competition, the win probability of each bidder is $1/n$, and the win probability of a pair is $2/n$. Under cover bidding, only one of the bidders is seriously attempting to win the auction, so the pair's win probability is then $1/(n-1)$, and the average individual win probability is $1/(2(n-1))$. In tenders in which exactly one flagged pair has both members bidding, members and non-members win at almost the same rate when pooled across field sizes (23.7 against 25.6 percent). When the field is thin, win probabilities diverge. At $n = 3$, members win 27.4 percent against 42.7 percent for non-members, and the pair wins 54.8 percent against $2/3$ under competition and $1/2$ under cover bidding. With four bidders the pair still wins less often than competition predicts, but by a smaller margin, and with at least five bidders it does not. Matched control pairs win at roughly the rates expected under independent symmetric bidding. Bidder asymmetry does not explain this: within the same firm, a flagged member wins less often when its partner also bids than when it bids without its partner, while a control member does not. Appendix B.4 plots the win rates by field size (Figure B.2), and Appendix B.5 reports the within-firm

comparison (Table B.4).

Deadline extensions provide an opportunity to directly observe the behavior of flagged pairs when it is uncertain whether a contract will be awarded. Some portals require at least two bids for a contract award, others at least three. A tender that has not received the required number of bids by its original deadline is extended, and firms that had not bid by then can still enter. On two-bid portals, among extended tenders in which one member held the only bid at the original deadline and exactly one firm bid late, the late bid is the partner's in 28 percent of tenders for flagged pairs and in 17 percent for matched control pairs. For flagged pairs, a partner's late bid beats the bid in hand in 14.0 percent of these tenders, against 27.7 percent for a late bid by another firm. Appendix B.6 shows that, for flagged pairs, the lowest bid falls less when the partner enters late than when another firm does (Figure B.3).

6 The footprint of collusion across India

With the model estimated and the classification validated, this section asks how common collusion is across India, how it varies across districts and procuring organizations, and what accounts for this variation. Because the data cover all layers of government, they show the behavior of bidders in the same place under different market conditions and under different regulatory regimes.

6.1 How common is collusion?

Collusion is common and is not concentrated among a few firms. The model estimate $\hat{\Lambda}$ is 33.6 percent, so I select the top 33.6 percent of scored pairs by likelihood ratio. Under this classification, almost two thirds (63.1 percent) of scored firms belong to at least one flagged pair, and a flagged pair bids in more than half of tenders (54.3 percent). In 43.1 percent of tenders, a member of a flagged pair wins while its partner also bids.[19] I refer to these tenders as *collusive* under this classification.

Together, these collusive tenders account for 18.4 percent of all recorded contract value in the sample. Because collusive tenders account for a larger share of tenders than of contract value, collusion is concentrated in lower-value tenders. One reason for this lies in the detection method: Because collusion is identified using repeated co-bids, the method has more power where the same two firms co-bid often, that is, on small, recurring contracts.

6.2 Where is collusion more common?

Because bidder pairs are defined at the district level, collusion can be compared across districts. Figure 5 shows the spatial variation of collusion across India. The left panel maps the district mean of each scored pair's posterior probability of belonging to the colluding component under the fitted model. The right panel shows the share of collusive tenders at the district level, with pairs flagged at the top decile.[20]

Nationally, the mean posterior is 31.4 percent, and the median district mean is close to that (29.2 percent). Across the 33 states and union territories with at least 1,000 tenders, the mean posterior ranges from 16.5 to 44.0 percent. It is highest in Arunachal Pradesh (37.2 percent), West Bengal

[19] Baránek, Musolff, and Titl (2026) find that 1,857 collusive firms take part in 15.6 percent of Ukrainian auctions. The two shares are defined differently: theirs counts auctions with a cartel bidder, mine counts tenders won by a member of a flagged pair whose partner also bids.

[20] Given that pairs with fewer than five co-bids are not classified, only 73.1 percent of tenders have a scored pair, and 65.6 percent have one linked to the winner (Figure C.1).

(44.0), Jharkhand (37.8) and Tamil Nadu (33.6), and lowest in Andaman and Nicobar (16.5 percent), Meghalaya (18.8), Goa (20.9) and Manipur (22.8).

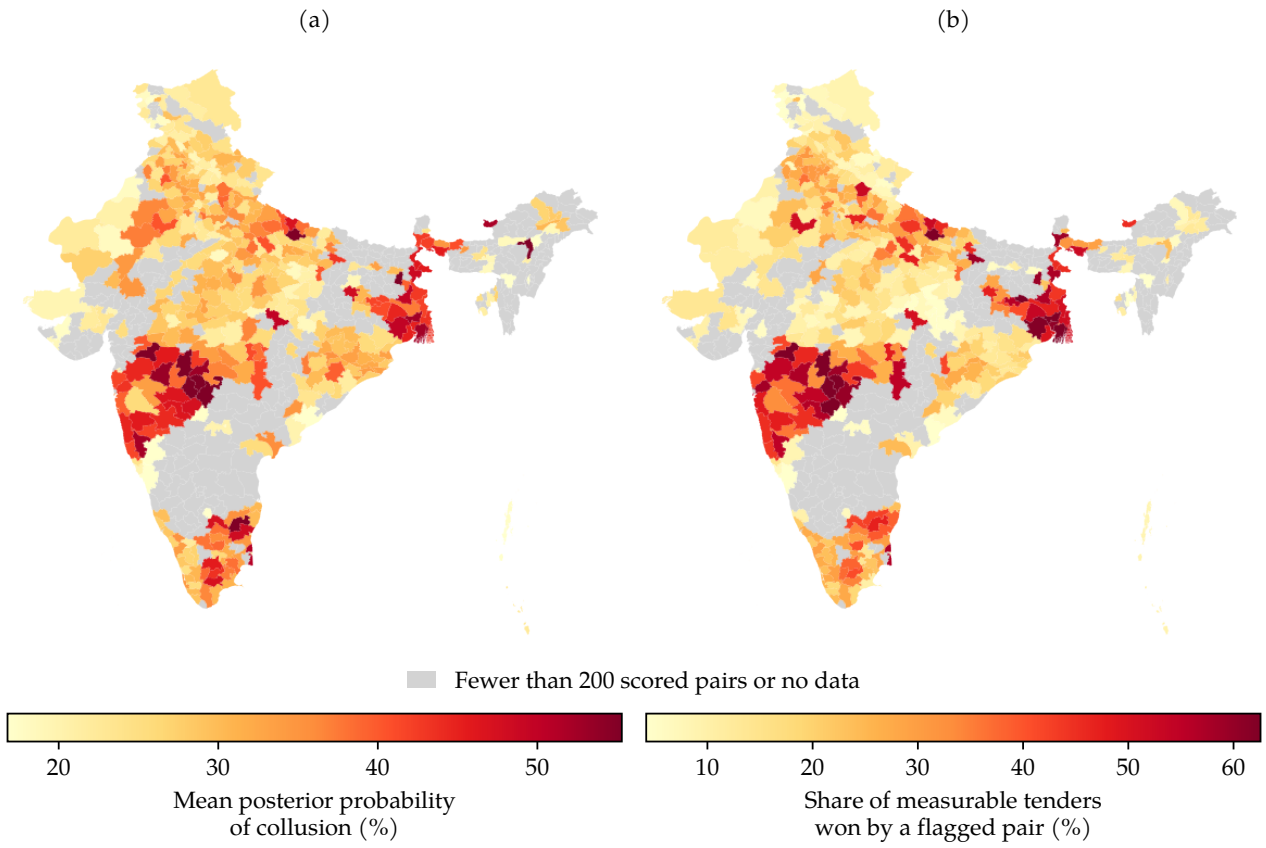


Figure 5: The geography of flagged pairs

Note: District-level maps over GADM district boundaries. Flagged pairs are the top decile of scored pairs by the estimated likelihood ratio (Section 4). Districts with fewer than 200 scored pairs are shown in gray. The respective color scales cover the 2nd to the 98th percentile of districts, values winsorized for plotting. Panel (a): district mean of each scored pair's posterior probability of belonging to the colluding component, over the scored pairs located in the district, in percent. Panel (b): share of measurable tenders (tenders with at least one scored pair that includes the winner) won by a member of a flagged (top-decile) pair while its partner also bids (*collusive* tenders).

To ask whether collusion follows places (districts) or procuring organizations, I decompose the variation in prevalence across market-years, where a market is a district × procuring organization × product cell. I regress each market-year's share of collusive tenders on district, organization, product and year fixed effects and calculate how much of the explained variance comes from each set using Shapley values: each set receives the average contribution to the R^2 over all orders in which the sets can be added (Table C.1). The four sets together have an R^2 of 0.74, of which 58 percent is explained by the organization and 37 percent by the district. Part of this variation reflects how often pairs co-bid, which determines whether a pair and therefore a tender can be scored at all. The variation in the prevalence of tenders won by matched control pairs has a similar pattern (organization 53 percent, district 46 percent). I therefore define *excess prevalence* as the flagged share minus the control share. Matched control pairs bid in the same district and co-bid about as often as flagged pairs, so the difference nets out the part of the variation that comes from how often pairs co-bid and how many tenders can be scored.

For excess prevalence, the four sets have an R^2 of 0.57, and the organization accounts for 76 percent of the explained variance against 19 percent for the district. Once co-bidding exposure is held fixed,

collusion varies mainly across organizations rather than across places.

6.3 Explaining variation across organizations

To understand what explains the variation in collusion across procuring organizations, I collect a set of variables for each procuring organization. First, each organization belongs to one of seven types: rural local bodies, urban local bodies, mixed local bodies, state departments, state boards and PSUs (public sector undertakings, i.e. state-owned enterprises), central PSUs, and the central government. Second, its rules: the shares of its tenders placed on two-bid and on three-bid portals (Section 2.4), the share allocated by lottery among tied lowest bids, and the share subject to a strict bid floor that excludes unreasonably low bids. Third, its behavior: the share of tenders it cancels when the lowest bid exceeds the government estimate. Finally, market conditions within the organization: the number of distinct firms that bid, the number of tenders, the average number of bidders per tender, the average log value of contracts, and the share of contracts that are works (as opposed to goods or services).

To relate these characteristics to the prevalence of collusion, I regress the share of collusive tenders in each district $\times$ organization $\times$ year cell (26,900 market-years) on these characteristics jointly, controlling for state $\times$ year fixed effects and the share of tenders that can be classified as collusive or non-collusive. The mean share of collusive tenders in this sample is 27.6 percent. Figure 6 summarizes the results from this regression. Together with the state $\times$ year fixed effects, these characteristics predict the share of collusive tenders. The organization-level variables explain 75 percent of the variance across organizations that remains after the fixed effects (Table C.2, column 1).

Collusive tenders are concentrated among local bodies: Within state-years and holding other characteristics fixed, the share of collusive tenders is 16.3 percentage points higher under rural local bodies than under state departments and 3.5 points higher under urban local bodies. The share is 9.0 points lower under central government procurement. Organizations whose tenders are placed on three-bid portals have a share of collusive tenders 11.5 points higher. Because the three-bid portals are the state portals of West Bengal and Maharashtra and the Coal India portal, and the regression holds the state and year fixed, this compares these portals with other portals in the same states. Organizations on two-bid portals, those that allocate more tenders by lottery, those with a strict bid floor, and those that cancel tenders more often do not differ significantly from other organizations. Market conditions also predict collusion: holding the number of tenders fixed, a one standard deviation increase in the log number of distinct firms is associated with 11.3 points fewer collusive tenders. Holding the number of firms fixed, a one standard deviation increase in the log number of tenders is associated with 11.9 points more collusive tenders. Organizations with more bidders per tender have slightly more collusive tenders and organizations that buy larger contracts slightly fewer.

This pattern reflects how the same pairs bid under different organizations. To show this, I regress, at the tender level, whether a pair has used a marker in a given tender on the organization-level variables, controlling for pair fixed effects. The results are shown in column 2 of Table C.2. The comparison is thus of the same pair bidding under different organizations.[21] In column 3, the outcome is a composite of whether a flagged pair has used a marker in a given tender. This specification uses only markers that can be observed in a given tender (for example, it excludes rotation-based markers that rely on observing multiple tenders). The results are qualitatively similar to those for the share

[21] The median flagged pair bids in tenders of a single organization, and the mean flagged pair in tenders of 1.42. In total, 30 percent of flagged pairs bid under more than one organization.

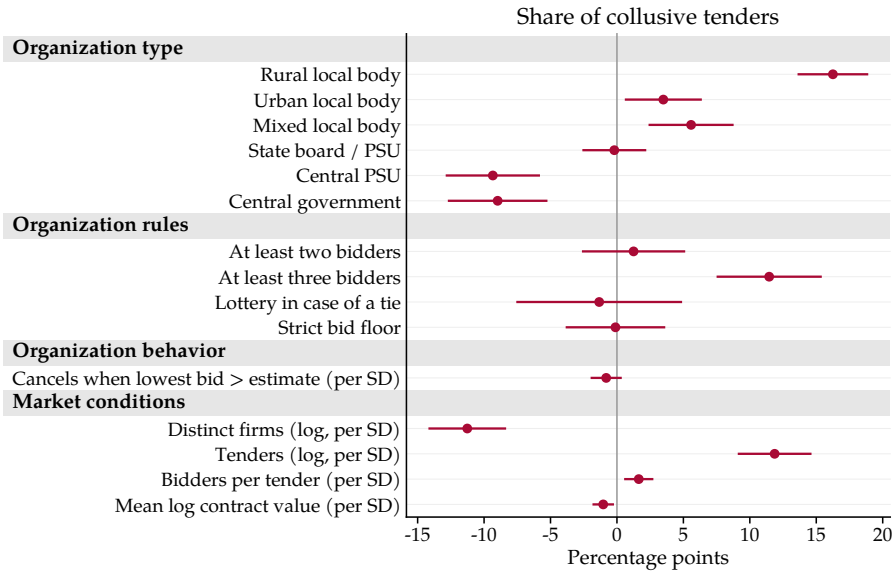


Figure 6: Organization characteristics and the share of collusive tenders

Note: Coefficients, in percentage points, from a single regression of the share of a market-year's measurable tenders won by a member of a flagged pair while its partner also bids on all characteristics jointly. Unit: district × organization × year with at least five measurable tenders in the lower two terciles of co-bid frequency. 26,900 market-years. Mean 27.6 percent. Organization types relative to state departments. Continuous regressors per standard deviation. Rule variables are shares of tenders from 0 to 1. The cancellation rate is the rate of cancellation or re-tendering when the lowest bid exceeds the estimate, for offices with at least ten such tenders (missing set to zero with an indicator). The regression also controls for the share of works (not shown), includes state × year fixed effects and the share of measurable tenders, and is weighted by the number of tenders. Bars are 95 percent confidence intervals. Standard errors clustered by district.

of collusive tenders across organizations: A given pair looks more collusive when bidding in tenders run by rural local bodies, and less collusive when a procuring organization uses lotteries to break ties.

Finally, I relate collusion to the thinness of markets. For each market-year (district × organization × product × year, at least five tenders), I count the distinct firms that bid and regress the share of collusive tenders on indicators for quartiles of that count, controlling for the log number of tenders and for state × year fixed effects. Markets with more tenders attract more firms, so the control compares market-years of similar size that differ in how many firms bid. Figure 7 plots the quartile coefficients relative to the thickest quartile.

In the thinnest quarter of markets, the share of collusive tenders is 11.6 percentage points higher than in the thickest quarter, against a mean of 26.2 percent. The collusive share decreases steadily in market thickness. This is consistent with work on incomplete rings: a ring gains less the more potential bidders it leaves out (Graham and Marshall 1987; Marshall and Marx 2012). To the extent that the other firms in thin markets are themselves members of flagged pairs, it is also the pattern Stigler (1964) predicts: collusion is easier to sustain among few firms. The same exercise for matched control pairs shows no such gradient: their share moves by at most about two points across quartiles and is slightly lower, not higher, in the thinnest markets.

Table C.3 reports a regression version with a continuous measure of market thickness. Across different specifications and controlling for the number of tenders in a market, more distinct firms in a market are robustly associated with less collusion. This also holds true within a market over time (controlling for market and year fixed effects), and when using the lagged number of firms in a market.

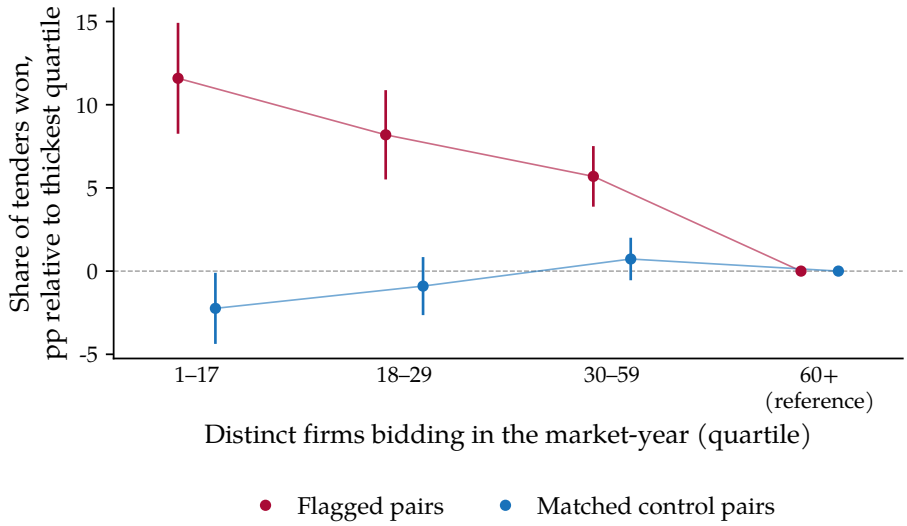


Figure 7: Thin markets and the share of tenders won by flagged pairs

Note: Coefficients from regressions of the share of a market-year's tenders won by a member of a flagged pair whose partner also bids (red), or by a member of a matched control pair whose partner also bids (blue), on indicators for quartiles of the number of distinct firms bidding in the market-year, in percentage points relative to the thickest quartile. Market: district × procuring organization × product. Market-years with at least five tenders (109,021). Quartiles hold equal numbers of market-years (cut points 17, 29 and 59 firms). The thickest quartile holds 75 percent of tenders. All regressions control for the log number of tenders, include state × year fixed effects, are weighted by the number of tenders, and cluster standard errors by district. Bars are 95 percent confidence intervals. Flagged pairs are the top decile of scored pairs by the estimated likelihood ratio (Section 4). Matched control pairs are from the same district with a co-bid count within a factor of 1.3.

A possible concern is that, controlling for the number of tenders in a market, collusion is mechanically higher in thin markets because in those markets, the same two firms co-bid more often. I therefore repeat the estimate separately for pairs that co-bid equally often, in bands of 5–9, 10–19, 20–49 and 50 or more co-bids. Within each band, markets with more firms have fewer collusive tenders relative to matched control pairs (Table C.3, columns 4–7).

In summary, collusive tenders account for a larger share of tenders than of contract value. In the median district, the mean posterior probability that a scored pair colludes is 29.2 percent. Net of co-bidding exposure, collusion varies more across organizations than across districts. It is most common under local bodies and, holding market size fixed, in markets with few firms. Section 7 turns from where flagged pairs win to what their wins cost.

7 Collusion and contract prices

This section estimates how much more the government pays in collusive tenders: tenders won by a flagged pair in which both of its members bid. Section 7.1 defines the outcome and the estimating equation, Section 7.2 reports the main estimates and the comparisons that address selection into tenders, Section 7.3 reports within-pair price premia, Section 7.4 reports placebo and robustness checks, and Section 8 relates the premium to local competition and to the procuring organization. In the preferred specification, collusive tenders settle at prices 3.19 percent (s.e. 0.23) above comparable tenders in the same market.

7.1 Measurement and estimating equation

Outcome. Studies of procurement prices compare prices either within narrowly defined product categories (see, e.g., Best, Hjort, and Szakonyi 2023; Allende et al. 2024) or relative to an engineering estimate or reserve price (see, e.g., Bajari, Houghton, and Tadelis 2014; Krasnokutskaya 2011; Chassang, Kawai, et al. 2022; Kawai, Nakabayashi, Ortner, et al. 2023). In the case of the Indian procurement data, precise product categories are not available and contracts are heterogeneous and thus difficult to compare. Therefore, the preferred outcome is the award price relative to the published government estimate, $Y_i = \log(\text{contract value}_i / \text{government estimate}_i)$.

Figure 8 plots the distribution of the award-to-estimate ratio by the government estimate. 86 percent of contracts are awarded at or below the estimate, and 13.2% exactly at it. Contracts in the middle of the size distribution are awarded at slightly lower average prices, relative to the estimate, than those at either end: the mean falls from 0.97 for tenders estimated at USD 1,000 to 0.89 at about USD 47,000 and returns to 0.99 at USD 1 million. For the tenders estimated between USD 3,000 and USD 30,000, which make up 57 percent of the sample, the interquartile range is 14 percent of the estimate; at the trough it is 20 percent.

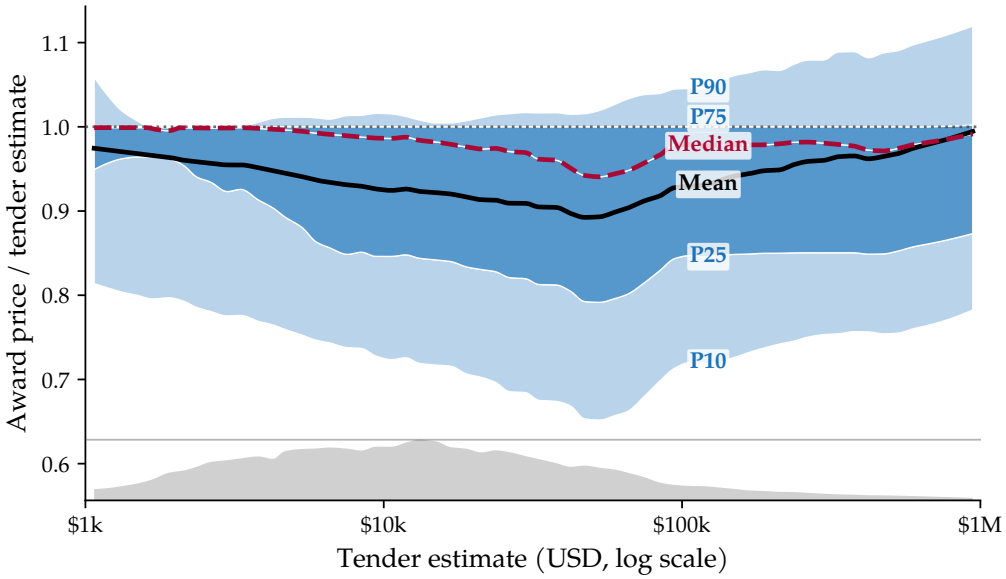


Figure 8: Award price relative to the government estimate, by tender size

Note: Distribution of the award price relative to the government estimate. Prices are the corrected values of Section 3.1. The sample is the estimation sample of Section 7.1, restricted to tenders whose price lies in $[0.30, 3.00]$ of the estimate. Estimates are converted at Rs. 77 per USD. The mean and the 10th, 25th, 50th, 75th, and 90th percentiles of the ratio are local linear fits. 13.2% of contracts are awarded at exactly the estimate. The gray shading below the plot is a kernel density estimate of the number of tenders.

Treatment. A colluding pair moves the price of a tender only if its designated winner wins and its cover bidder is in the auction. As before, a tender is therefore *collusive*, $\text{Collusive}_i = 1$, if its winner belongs to a pair flagged in Section 4 *and* that pair's other member also bid in it. The comparison group consists of all other tenders, and alternative specifications restrict it to tenders in which a flagged firm wins without its partner, or in which a flagged pair bids and loses.

Estimating equation. The baseline specification is

$$Y_i = \beta \, \text{Collusive}_i + X_i'\gamma + \alpha_{d(i)o(i)t} + \alpha_{p(i)t} + \varepsilon_i, \tag{10}$$

where the coefficient of interest is β, the price premium associated with collusion. It captures the difference, within a market and conditional on controls, between collusive tenders and other tenders. $\alpha_{d(i)o(i)t}$ are district $\times$ organization $\times$ year fixed effects, over the 659 districts and 4,760 top-level procuring organizations, and $\alpha_{p(i)t}$ are product $\times$ year fixed effects over 9 product categories. The tender controls X_i are the log estimate, the length of the submission window, the tender fee, the earnest-money percentage, and calendar-month dummies. Missing values of a control are set to zero, and X_i also includes an indicator for each control being missing. Standard errors are clustered by district. The baseline specification does not control for the number of bidders in the auction because collusion can affect it: a colluding pair may add cover bids and deter outside bidders. Subsequent specifications control for the number of bidders.

Sample. The price premium regressions use tenders with an identifiable winner and an award price within $[0.30, 3.00]$ of the estimate. The sample excludes tenders allocated by lottery among tied bids, tenders where the final contract value is not observed, and tenders with fewer than two bidders, since a single-bidder tender cannot contain a pair. Table 1 describes the sample; 26.8 percent of its tenders are collusive.

7.2 Main estimates

Table 6 reports the estimates. In the baseline specification in column 1, collusive tenders settle at prices 3.19 percent above other tenders in the same district, organization, and year. This price premium corresponds to 0.21 standard deviations of the residualized outcome (Table D.1). Controlling for the number of bidders barely moves the price premium, to 3.16 percent (column 2). Controlling instead for the number of bidders who pass the technical stage gives a similar estimate (Table D.2).

Columns 3 and 4 narrow the comparison group. First, bidders in flagged pairs may differ from other bidders in their costs, leading to different prices even in the absence of collusion. Column 3 therefore restricts the sample to tenders won by a member of a flagged pair. Among these tenders, the collusive ones, in which the winner's partner also bids, settle 3.31 percent higher (column 3). Second, flagged pairs may systematically enter jointly in tenders where they can extract a higher price premium. Section 5.3 shows that flagged pairs enter together at higher rates when there are few outside bidders. To address this, column 4 restricts the comparison group to tenders where a flagged pair bid but lost. The estimate is slightly larger, at 3.31 percent (column 4). As an alternative, Appendix Table D.4 takes the full sample and explicitly controls for selection into treatment by adding an indicator for whether a flagged pair enters the tender jointly. The selection effect is close to zero: conditional on tender controls and market fixed effects, tenders where flagged pairs jointly enter are not more expensive. The price premium is entirely driven by tenders where flagged pairs win.

Finally, column 5 replaces the indicator with a continuous score, the standardized log likelihood ratio used to flag pairs. A one-standard-deviation increase in the winner's pair score is associated with a 1.71 percent higher price, holding fixed the highest score among non-winning pairs.

Appendix D.6 shows that the price premium is driven by collusive tenders being much more likely to settle at or very close to the government estimate: the probability that a collusive tender settles

Table 6: Collusive tenders and contract prices

Dependent variable: log(contract value/government estimate)					
	(1)	(2)	(3)	(4)	(5)
Collusive tender	0.0319***	0.0316***	0.0331***	0.0331***	
	(0.0023)	(0.0020)	(0.0024)	(0.0032)	
Max winner pair score (SD)					0.0171***
					(0.0012)
N	4,594,861	4,594,861	2,377,278	1,660,170	3,217,813
Mean contract/estimate ratio	0.933	0.933	0.949	0.965	0.937
Share treated	26.8%	26.8%	51.9%	74.3%	
District × Org × Year, Product × Year FE	✓	✓	✓	✓	✓
Tender-level controls	✓	✓	✓	✓	✓
Missing-value and imputation flags	✓	✓	✓	✓	✓
N bidders control		✓	✓	✓	✓
Competitive environment control					✓
Sample	Full	Full	Any flagged winner	Flagged pair present	Scored subsample
Clustering	District	District	District	District	District

Note: Dependent variable is log(contract value/government estimate). The row "Collusive tender" is Collusive$_i$: the winner belongs to a flagged pair whose partner also bid in the tender. Column 1 is the full sample without a bidder-count control. Column 2 adds the total number of bidders (linear), which columns 3 to 5 retain. Column 3 restricts to tenders with a flagged winner. Column 4 restricts to collusive tenders and tenders in which a flagged pair that does not include the winner bid. Column 5 replaces the indicator with the standardized pair score of the winner's highest-scoring pair, on the scored subsample, controlling for the highest score among non-winning pairs. The comparison group is the complement of the treatment indicator within each column's sample. All columns include district × organization × year and product × year fixed effects and the tender controls of equation (10). Flagged pairs are the top decile of scored pairs by the estimated likelihood ratio (Section 4). Standard errors, clustered by district, in parentheses. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

within one percent of the estimate is 12.1 percentage points higher, against a base rate of 41.1 percent.

7.3 Within-pair price premia

The comparisons above hold the market fixed but compare different pairs. To hold the pair fixed, I construct a (tender, pair) panel in which tender i enters once for each flagged pair c that co-bids in that tender.

Within this sample, the regression is:

$$\text{ContractRatio}_{ic} = \beta\, \text{PairWins}_{ic} + \alpha_c + \gamma_t + \gamma_p + X_i'\gamma + \varepsilon_{ic}, \tag{11}$$

where $\text{PairWins}_{ic} = 1$ if either member wins and the pair fixed effect α_c absorbs every characteristic of the pair that does not vary across its tenders. Standard errors are clustered by tender, which enters once for each flagged pair active in it.

Table 7: Within-pair effect: winning among the pair's co-bids

	(1)	(2)	(3)
Pair member wins	0.0189***	0.0191***	0.0192***
	(0.0002)	(0.0002)	(0.0002)
N	5,397,144	5,396,273	5,396,266
Within R^2	0.0042	0.0043	0.0189
Mean dep. var.	0.968	0.968	0.968
Pair FE	✓	✓	✓
Year FE		✓	✓
Product(L1) FE		✓	✓
Tender controls			✓
Clustering	Tender	Tender	Tender

Note: Dependent variable is the contract-to-estimate ratio. Sample: (tender, pair) cells in which both members of the flagged pair bid. *Pair member wins* is an indicator for either member of the pair winning. Column 1 includes pair fixed effects only. Column 2 adds year and product (L1) fixed effects. Column 3 adds the log estimate, the tender fee, and the earnest-money percentage, with fixed effects for tender type, form of contract, and contract type. Flagged pairs are the top decile of scored pairs by the estimated likelihood ratio (Section 4). Standard errors, clustered by tender, in parentheses. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

Among a pair's co-bids, those it wins settle 1.89 points of the estimate higher than those it loses, against a mean ratio of 0.968 (column 1 of Table 7). Year and product fixed effects leave the estimate at 1.91 points (column 2), and tender characteristics leave it essentially unchanged at 1.92 (column 3); adding the bidder count gives 2.00.

7.4 Placebos and robustness

Placebo assignments. Table 8 re-estimates column 2 of Table 6 (3.16 percent under the headline flag) with different placebo flags, to show that the estimated price premium is not due to co-bidding at the pair level.

The first two rows re-estimate the price premium for matched control pairs, first matched within district and co-bid count and then also within the same organization. For both, the estimated price premium of a "placebo collusive" tender is negative and around 2 percent. This negative price premium is expected for non-colluding pairs that co-bid regularly, as joint entry by such pairs is associated with more competition, and thus lower prices, in that tender. Rows three through five assign

the flag randomly, reassign the flag within the co-bidding bins of the actual flagged pairs, and reassign it within district-years. The estimated coefficients are close to zero or negative, between -0.76 and 0.02 percent. Finally, the bottom two rows use two alternative ways to flag pairs. First, flagging the most frequent co-bidders recovers a price premium of 1.71 percent, around half of the headline result. This is expected, given that the share of flagged pairs is increasing in the number of co-bids: the score flags 20.9 percent of pairs with at least 24 co-bids, the cutoff of the co-bid screen, against 8.8 percent of pairs with fewer. Second, flagging the bottom decile of the likelihood ratio, i.e. the least suspicious ten percent of the scored pairs, yields a large negative price premium of -5.69 percent.

Table 8: Placebo tests: the price effect under alternative definitions of the suspicion flag

Treatment definition	Coef.	(SE)	N	Share treated
Headline: top decile of the likelihood ratio	0.0316^{***}	(0.0020)	4,594,861	26.8%
Matched control pairs (district × co-bid count)	$\mathbf{-0.0183}^{***}$	(0.0025)	4,594,861	19.7%
Matched control pairs, same organization	$\mathbf{-0.0174}^{***}$	(0.0022)	4,594,861	18.3%
Random 10% of scored pairs	-0.0045^{*}	(0.0025)	4,594,861	12.2%
Reassigned within co-bidding bin	-0.0076^{***}	(0.0027)	4,594,861	15.4%
Score reshuffled within district × year	0.0002	(0.0019)	4,594,861	19.9%
Most frequent co-bidders	0.0171^{***}	(0.0021)	4,594,861	34.7%
Bottom decile of the likelihood ratio	-0.0569^{***}	(0.0040)	4,594,861	14.0%

Note: Column 2 of Table 6 re-estimated with the suspicion flag replaced by each alternative. The headline flag is the top decile of scored pairs by the estimated likelihood ratio (Section 4). Every alternative is built on that fit's scored pairs and flags the same number of pairs, except the matched assignments, which leave unreplaced the flagged pairs with no unflagged counterpart in their stratum. Each is carried through the tender-level treatment unchanged, so samples are identical across rows. The first rows are matched placebos (same district and similar co-bid count; in the second row also the same modal organization). The next are mechanical placebos (a seeded random tenth; a set disjoint from the flag with the same count in each co-bid bin; a permutation of the score within district × year). The last two are comparisons that are not placebos (a co-bid-frequency screen flagging the share closest to a tenth; the bottom decile of the score). "Share treated" is the fraction of tenders in which the winner and a flagged partner both bid. It differs across rows, so the placebos are not exactly power-matched. Standard errors, clustered by district, in parentheses. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

Robustness to specification and sample. Appendix D.2 reports results from a battery of robustness checks varying the sample, the fit, and the bidder identifier. All 75 re-estimated coefficients are positive and significant at the 1 percent level, and the column-1 estimate ranges from 1.99 to 3.61 percent. In particular, the results are robust to excluding contracts priced at a round multiple of five percent of the estimate, and to scoring pairs on pre-2022 bids only and estimating the premium on later tenders, so that the flag is estimated on bids that do not enter the outcome. The results are also robust to weighting tenders by the log estimate and to dropping tenders whose winner belongs to no scored pair. Fitting the model under the two alternative placebo designs of Appendix A.2 gives similar premia. Finally, using bidder pairs defined from bidders at the national instead of the district level does not move the results.

Cutoff. Figure 9 re-estimates β at every cutoff from the 70th to the 99.5th percentile of the score, holding the comparison group fixed at the tenders that are not collusive under the 70th-percentile cutoff. The price premium is monotone in the cutoff in columns 1 to 3, and flat in column 4. In

columns 1 to 3, pairs with higher scores have higher implied price premia of collusion.

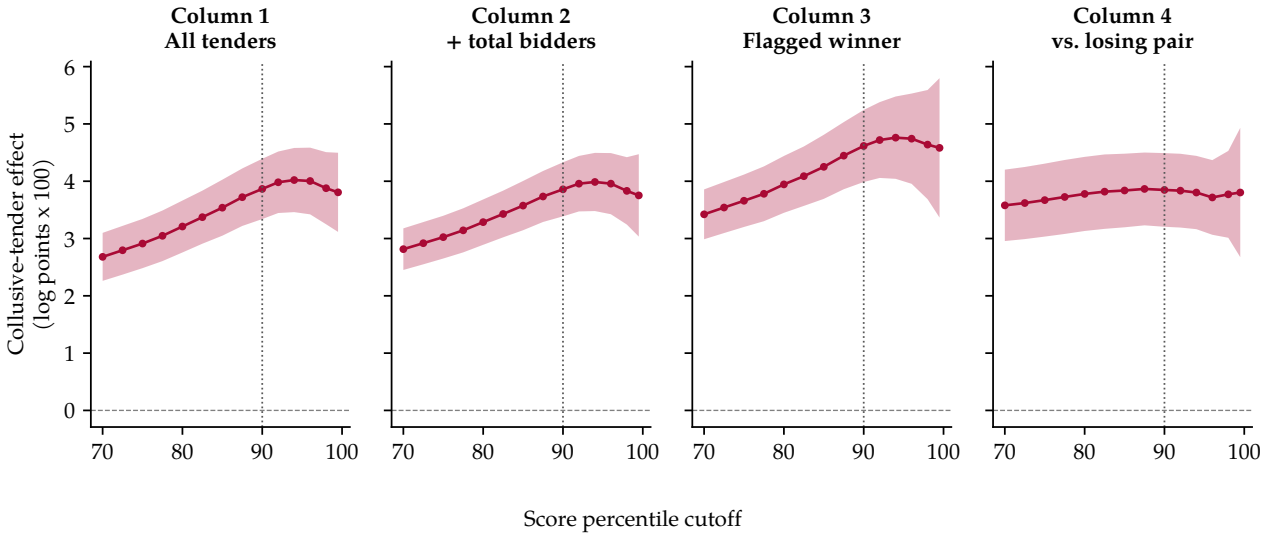


Figure 9: Price effect across suspicion-score cutoffs

Note: The coefficient β of equation (10), re-estimated at each percentile cutoff of the suspicion score from the 70th to the 99.5th, one panel for each of columns 1 to 4 of Table 6. The dotted line marks the 90th-percentile cutoff used throughout. At that cutoff the estimate differs from Table 6 because the comparison group excludes tenders that are collusive only for pairs scored between the 70th and 90th percentiles. The comparison group is held at the tenders that are not collusive under the 70th-percentile cutoff. Tenders that are collusive only for pairs scored between the 70th percentile and the cutoff are excluded. Therefore, across cutoffs, the treated group changes and becomes smaller. Flagged pairs are the top decile of scored pairs by the estimated likelihood ratio (Section 4). Shaded bands are 95% confidence intervals. Standard errors clustered by district.

8 Competition, minimum-bidder rules and collusion

Section 6 shows that flagged pairs win more of the tenders in markets where few firms bid, and Section 7 that collusive tenders cost more than non-collusive ones. This section relates both to the number of firms in the market and to the minimum-bidder rules of Indian portals. In the benchmark of McAfee and McMillan (1992), a cartel that faces no outside competition bids the reserve price. The government estimate acts as a soft reserve price, since organizations cancel or re-tender some of the tenders whose lowest bid exceeds it (Section 6.3). To the extent that a cartel faces little outside competition, the price of a collusive tender should depend less on the number of firms in the market than the competitive price does. A rule that requires a minimum number of bids adds competition only if the required bids come from firms that do not collude; colluding firms can meet it with complementary bids of their own, which give the appearance of competition (Porter and Zona 1993; Clark, Coviello, and De Leverano 2025).

8.1 Competition and prices

Prices fall with the number of firms in the market. Auctions with few bidders are common in public procurement: even in United States federal procurement, many contracts receive only one or a few bids (Kang and R. A. Miller 2022). I sort market-years (district $\times$ procuring organization $\times$ product $\times$ year, at least five tenders) into quartiles of the number of distinct firms that bid, as in Figure 7. Holding the log number of tenders fixed, non-collusive tenders cost 2.3 log points more in the thinnest quartile than in the thickest, within the same district, organization and year. The price of collusive

tenders varies much less across quartiles: about 3.5 log points above non-collusive tenders of the thickest quartile, whether the market is thin or thick (Figure 10, panel a). Collusive tenders are more common in thin markets (Section 6.3).

The premium on a collusive tender is smaller in thin markets. Figure 10 estimates the price of collusive and non-collusive tenders in each quartile in one regression with the controls and fixed effects of Table 6, so that every comparison is within the same district, organization and year. Relative to non-collusive tenders in the thickest quartile, a collusive tender costs 3.4, 3.1, 3.1 and 3.5 log points from the thinnest to the thickest quartile (panel a). Because non-collusive tenders are 2.3 log points dearer in the thinnest quartile, the premium rises from 1.1 (s.e. 0.2) in the thinnest quartile to 3.5 (s.e. 0.3) in the thickest (panel b). Market thinness moves the price of non-collusive tenders but not that of collusive tenders, so the premium is smallest in the markets where collusive tenders are most common. Collusive prices are compressed toward the estimate: the contract is within 1 percent of the estimate in 52 percent of collusive tenders in the thinnest quartile and 68 percent in the thickest, against 27 to 35 percent of non-collusive tenders. Among non-collusive tenders, matched-control tenders cost 1.0 to 2.0 log points less than the rest in the same quartile, and their price falls with that of the rest.

Prices of non-collusive tenders also fall with the number of financial bidders, by 18 log points from two to seven or more. The premium falls from 3.4 log points (s.e. 0.3) with two financial bidders to close to zero with seven or more, but about half of this decline also appears for matched control pairs, and the number of bidders depends on the pair itself (Appendix Figure D.5). Unlike the number of firms in the market, the number of bidders in the tender therefore does not separate the effect of competition from the conduct of the flagged pair.

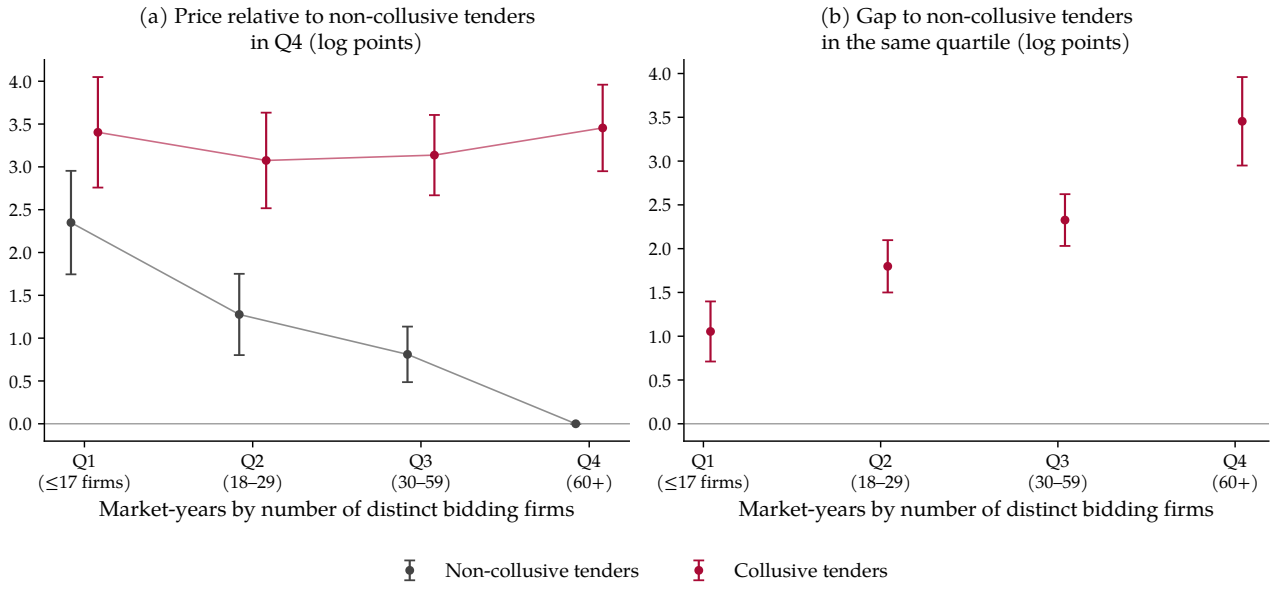


Figure 10: Market thinness and the price of collusive and non-collusive tenders

Note: Market-years (district × procuring organization × product × year, at least five tenders) are sorted into quartiles with equal numbers of market-years by the number of distinct bidding firms (cut points 17, 29 and 59). One tender-level regression of 100 times the log contract-to-estimate ratio on indicators for each group of tenders in each quartile, with the log number of tenders in the market-year and the controls and fixed effects of Table 6. Collusive tenders are won by a member of a flagged pair whose partner also bids. Non-collusive tenders are all others, including the tenders of matched control pairs. Panel (a): each group relative to non-collusive tenders in the thickest quartile. Panel (b): the difference of collusive from non-collusive tenders in the same quartile. Bars are 95 percent confidence intervals. Standard errors clustered by district.

The price of collusive tenders also responds less to the local price level than the price of tenders won by matched control pairs. I measure the local price level of a tender as the mean contract-to-estimate ratio of the other non-collusive tenders in the same district, organization, product and year. Figure 11 plots the contract-to-estimate ratio of collusive tenders and of tenders won by a member of a matched control pair against this level. The price of matched-control tenders rises with a slope of 1.05 and the price of collusive tenders with a slope of 0.81: where the other tenders settle about 20 percent below the estimate, collusive tenders settle about 15 percent below it, and at or above the estimate the two groups coincide. The difference in slopes is -0.24 (s.e. 0.04), and -0.21 (s.e. 0.02) with pair fixed effects, which compare tenders of the same pair in different markets. Collusive prices are thus again compressed toward the estimate, and the premium is largest where competition pushes the other prices furthest below it. Collusive tenders are concentrated where the premium is smallest: half of them are in markets where the other tenders settle at 95 to 100 percent of the estimate. In these markets collusive tenders cost 1.9 percent more than matched-control tenders, against 6 to 8 percent in markets below 95 percent. Across district $\times$ product markets, collusive and non-collusive tenders are both more expensive where collusive tenders make up a larger share of the market, and the gap between them narrows by 0.3 percentage points per 10 percentage points of collusive share (Appendix Figure D.6).

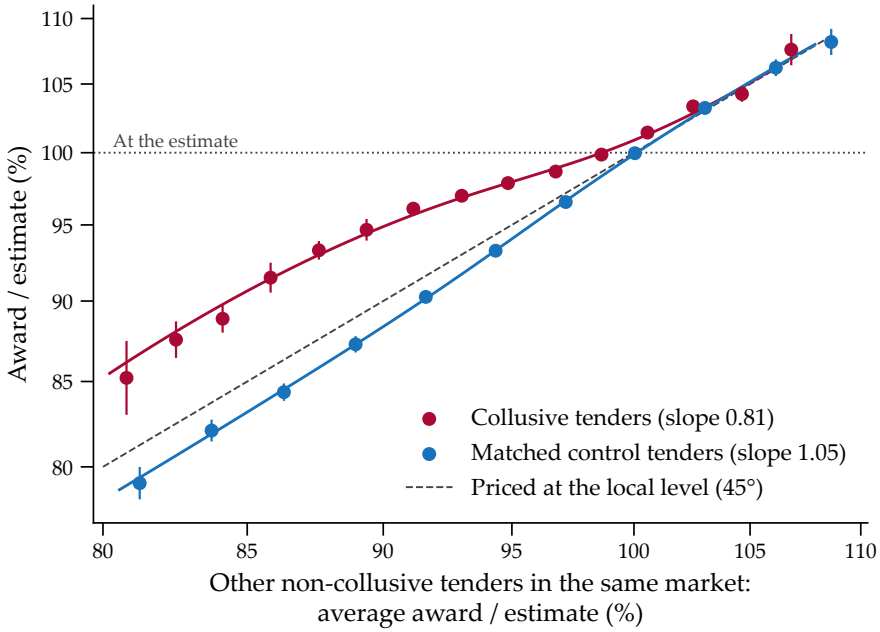


Figure 11: Collusive and matched-control tenders by local price level

Note: Local price level: mean contract-to-estimate ratio of the other non-collusive tenders in the same district $\times$ procuring organization $\times$ product $\times$ year, excluding the tender itself. Binned means of the contract-to-estimate ratio of collusive tenders (won by a member of a flagged pair whose partner also bids) and of tenders won by a member of a matched control pair whose partner also bids, with 95 percent confidence intervals; lines are fitted curves, and the slopes are from linear regressions of the ratio on the local level. The vertical distance between the two groups is the raw difference, without fixed effects. Tenders of Table 6, with its controls. Standard errors clustered by district.

The premium also varies across procuring organizations. Figure 12 estimates equation (10) separately for the 100 organizations with the largest estimation samples, with district, year and product fixed effects in place of the interacted fixed effects. The coefficient is positive in 90 organizations and positive and significant at the 5 percent level in 74; the median is 2.4 percent and the interquartile range runs from 1.0 to 4.2 percent. The estimates vary by more than sampling error alone would produce:

their standard deviation is 2.8 percentage points, against a root-mean-square standard error of 1.2. The median premium is 3.1 percent across the 43 state organizations, 2.2 percent across the 46 local bodies and 1.5 percent across the 11 central organizations. These medians weight each organization equally. In one regression over all tenders, the premium under rural local bodies is the same as under state departments (3.4 against 3.1 log points; difference 0.3, s.e. 0.4), although collusive tenders are 16.3 percentage points more common under rural local bodies (Section 6.3).

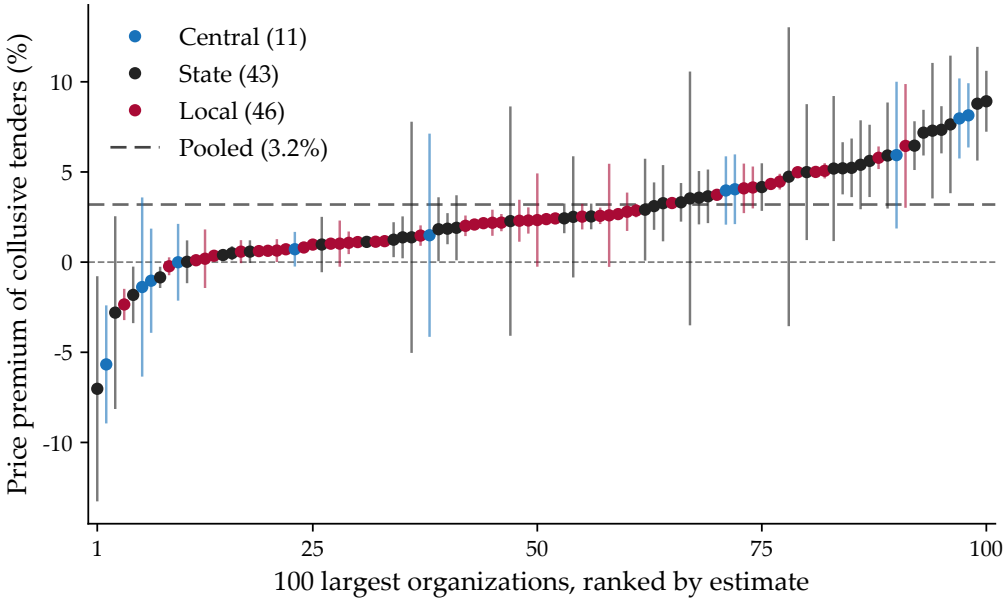


Figure 12: Price effect of collusion by procuring organization

Note: Each point is the coefficient on the collusive-tender indicator from a separate estimation of equation (10) within one of the 100 organizations with the largest estimation samples, ranked by the estimate; bars are 95 percent confidence intervals. Colors give the level of government; the number of organizations is in parentheses. Tender controls of equation (10) and district, year and product fixed effects. Standard errors clustered by district within each organization, or HC1 where the organization spans fewer than ten districts. The dashed line is the pooled estimate of column 1 of Table 6. Estimates in percent.

8.2 Minimum-bidder rules

The rules of Section 2.4 shape the number of bidders: on two-bid portals, 57 percent of tenders have exactly two financial bidders, and on three-bid portals, 70 percent have exactly three (Appendix Figure F.3).

8.3 How cartels meet the rules

Flagged pairs meet part of these requirements themselves: the partner of a flagged pair is more likely to have bid where its bid completes the count (Section 5.3), and more often supplies the late bid in an extended tender (Section 5.3). Where the rule requires three bids, a cartel that supplies the required bids itself needs a third member.

I recover larger groups from the network of flagged pairs and define a cartel as a maximal clique: a set of firms in which every pair is flagged and to which no other firm can be added (if A–B, B–C and A–C are all flagged, the three firms form one cartel; if A–C is not, they form two cartels of two). Of the 2.6 million tenders in which at least two members of one cartel bid, 80 percent have exactly two members bidding and 18 percent exactly three (Figure 13).

On three-bid portals, exactly three members of one cartel bid in 24 percent of the tenders with a cartel present, against 6 percent on two-bid portals and 8 percent on other portals (Figure 14).[22] This is consistent with a three-bid rule raising the number of bids partly by enlarging the cartels that supply them. The three-bid portals are those of a few states, which also hold the densest groups of flagged firms, so these comparisons cannot separate the rule from the portals that apply it.

Within states, collusive tenders are 11.5 percentage points more common on three-bid portals and no more common on two-bid portals (Section 6.3). The premium per collusive tender is about the same under the two rules (3.3 log points on three-bid and 3.6 on two-bid portals; difference −0.3, s.e. 0.4). It is lower on other portals (0.7), but so is the coefficient for matched control pairs, so that comparison does not isolate collusion.

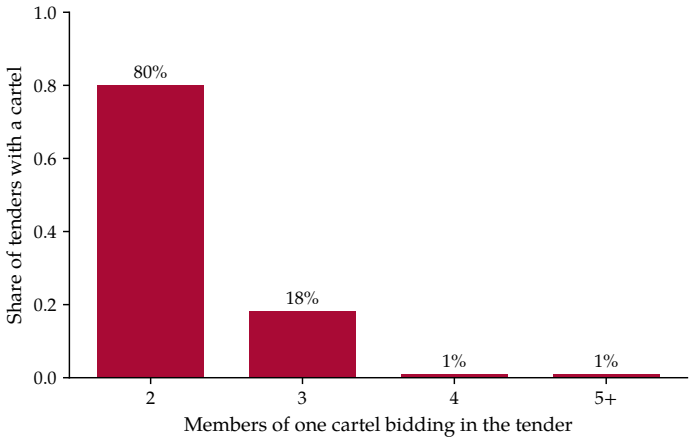


Figure 13: Size of detected cartels

Note: A cartel is a maximal clique in the network of flagged pairs: a set of firms in which every pair is flagged and to which no other firm can be added. Distribution of the 2,584,220 tenders in which at least two members of one cartel bid, by the largest number of members of one cartel bidding in the tender. All portals.

Thinness raises the price of non-collusive tenders but not that of collusive ones, so the premium per collusive tender is smallest in the thin markets where collusion is most common. Flagged pairs themselves supply part of the bids that the rules require, and on three-bid portals three members of one cartel bid together more often. A minimum price weakens cartels by limiting the scope for punishment (Chassang and Ortner 2019). A minimum number of bids, by contrast, can be met from inside the cartel, so its effect on prices depends on which firms supply the required bids.

[22] Three-bid portals also hold more tenders with exactly three bidders, so the same pairs co-bid more often, and pairs that co-bid often are more likely to be flagged; triangles of flagged pairs could be more frequent there for that reason alone. I test this with a permutation placebo that reassigns the flags at random among all scored pairs within cells of district and co-bidding frequency, holding the number of flagged pairs in each cell fixed, and rebuilds the cartels in each of 50 draws. The placebo reproduces 7.9 of the 18.2 percentage point difference between three-bid and two-bid portals, and no draw reaches the observed difference. Counting cartels instead of tenders, or measuring how often a triangle of bidders closes, does not separate the rules: closure is high on both, 71 percent on three-bid and 47 percent on two-bid portals against permutation means of 33 and 17 percent.

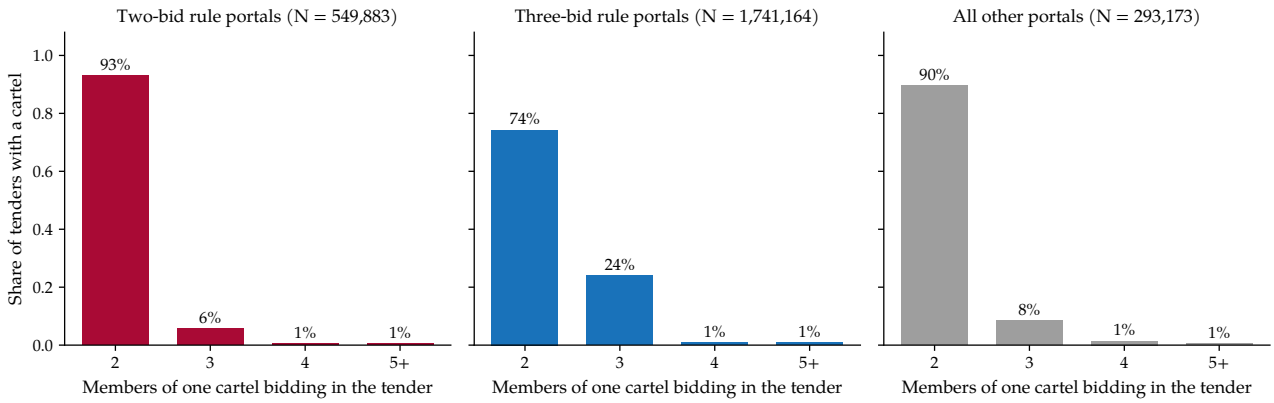


Figure 14: Cartel size by minimum-bidder rule

Note: Portals are grouped by the rule implied by their deadline extensions (Section 2.4). Tenders in which at least two members of one cartel bid, by the largest number of members of one cartel bidding. "All other portals" are the other portals of Section 2.4. Cartels are maximal cliques of flagged pairs.

9 Welfare

This section studies the welfare consequences of collusion. It first asks whether winners of collusive tenders deliver works of different quality: there is no evidence that they do, so the welfare effect can be read from prices and quantities. It then computes the transfer from procuring organizations to firms in flagged pairs and bounds the efficiency loss from the quantity response. In aggregate, the estimates imply that the Indian government overpays by about USD 13.2 billion in the sample and USD 40.1 billion on national procurement over ten years.

9.1 Quality

I first examine whether collusion is associated with worse project quality and completion. The answer matters for the interpretation of the price effect: if colluding bidders have worse project delivery outcomes, the price premium understates the cost of collusion. If flagged pairs consist of better contractors, having outside bidders enter might lower prices at the cost of worse quality. The tender records do not document project delivery. However, some of the procurement contracts are for rural roads constructed under the national rural roads program, the Pradhan Mantri Gram Sadak Yojana (PMGSY).

A separate PMGSY database records, for every road, the contractor, the award date, the original completion deadline, the actual completion date, and the grade from each inspection by the program's national and state quality monitors. 10,853 tenders are linked to PMGSY road packages on package codes, contractor and road names, award date, location, and sanction cost against the tender estimate (Appendix F.7). I construct two measures of audit quality: the share of a tender's inspections graded satisfactory and an indicator that all inspections were graded satisfactory. Two further measures capture completion: days of delay against the original deadline and an indicator for finishing more than 180 days late. In the average linked tender, a share of 0.915 of inspections is graded satisfactory, and the tender finishes 236 days after its original deadline.

Table 9 re-estimates Table 6, with the quality and completion variables as outcomes. Collusive tenders do not differ from other tenders in the same district, organization, and year on any of the four outcomes (panel A): the share of satisfactory inspections differs by 0.004 against a mean of 0.915, delay by -14.6 days against a mean of 224.0, and no coefficient is statistically different from zero. The

coefficients remain insignificant when restricting the sample to tenders in which a flagged pair bid (panel B). Tenders won by a matched control pair with its partner active finish about 53 days earlier (panel C), more than the 35-day advantage of flagged firms in panel D: firms that co-bid regularly finish their roads faster whether or not they collude.

This null result is consistent with how common flagged firms are among contractors. In the linked sample, 6.2 percent of tenders are collusive and 28.3 percent are won by a member of a flagged pair; in the full sample of Table 6 the shares are 26.8 and 51.7 percent. In panel D, tenders won by firms in flagged pairs receive the same inspection grades (a difference of 0.003) but complete 35.3 days earlier and are 4.7 percentage points less likely to finish more than 180 days late. Delivery is a persistent firm trait, and the bidder with the best record on its other roads wins 62.9 percent of road tenders.

In this subset of the data, covering one program and 3,622 tenders with usable variation, collusion raises the price without a measurable change in the roads delivered. I therefore compute the welfare effect of collusion from prices and quantities alone.

9.2 Prices and transfers

Section 7.2 shows that contract prices rise by 3.16 to 3.31 percent when a flagged pair wins with its partner active, depending on the comparison group. As a back-of-the-envelope calculation, I approximate the annual overpayment by the government to flagged pairs as a transfer, equal to the price effects of Section 7.2 times the annual volume of collusive tenders.

Transfer. The annual transfer can be calculated on the classified sample, consisting of the tenders in which at least two bidders form a scored pair so that collusion could have been detected. The volume is the average annual value of collusive tenders over the four most recent years in the data (2022 to 2025), and the transfer is:

$$\widehat{\text{Transfer}} \;=\; \hat{\beta} \cdot \overline{Q}, \qquad \overline{Q} \;=\; \tfrac{1}{4} \sum_{t=2022}^{2025} \sum_{j\,:\,\text{winner}(j) \in \mathcal{M},\; t(j)=t} V_j, \tag{12}$$

where $\mathcal{M}$ is the set of collusive tenders and $\hat{\beta}$ is the price effect in collusive tenders of Section 7.2. I report it for each specification of Table 6.

In this initial specification, the volume of collusive tenders is computed on the same sample on which $\hat{\beta}$ was estimated. Flagged pairs win USD 4.97 billion of classified-sample procurement a year on this definition, 15.1 percent of the value of classified tenders. At the headline price effect the transfer is USD 159 million a year, or USD 1.6 billion over ten years at that rate.

Panel A of Table 10 holds this volume fixed and varies the price effect across the columns of Table 6. The estimated transfer ranges from USD 157 to USD 164 million a year.

Panel B holds the column-1 price effect fixed and widens the volume it applies to. Flagging the top $\hat{\Lambda}$ (33.6 percent) of scored pairs instead of the top decile raises the collusive share of classified value from 15.1 to 40.8 percent and the transfer to USD 428.6 million a year. A large share of tenders was not classified because their bidders do not bid often enough. Extrapolating the share of collusive tenders to these tenders gives USD 0.89 billion a year. Adding the tenders in the procurement data that were dropped by the sample restrictions yields USD 1.32 billion. Finally, extrapolating the share of collusive tenders to Indian public procurement at 8 percent of GDP gives USD 4.01 billion a year, or USD 40.1 billion over ten years. The last three rows assume that collusion is as common where it cannot be detected as where it can.

Table 9: Collusive tenders and road quality

	Audit quality		Completion delay	
	Share of inspections satisfactory	All inspections satisfactory	Days late	More than 180 days late
	(1)	(2)	(3)	(4)
A. All linked tenders (Table 7, column 1)				
Collusive tender	0.004	-0.001	-14.6	-0.027
	(0.009)	(0.022)	(12.3)	(0.022)
Mean of dependent variable	0.915	0.694	224.0	0.437
Share collusive	6.5%	6.5%	6.9%	6.9%
Observations	7,749	7,749	7,497	7,497
B. Tenders in which a flagged pair bid (Table 7, column 4)				
Collusive tender	-0.025	-0.146**	-52.6*	-0.028
	(0.023)	(0.067)	(26.8)	(0.059)
Observations	597	597	612	612
C. Placebo: matched control pairs (column 1 design)				
Collusive tender (matched controls)	-0.002	0.005	-53.2***	-0.060***
	(0.010)	(0.022)	(17.8)	(0.023)
Observations	7,749	7,749	7,497	7,497
D. Firm type: winner belongs to a flagged pair (column 1 design)				
Flagged winner	0.003	-0.010	-35.3***	-0.047***
	(0.005)	(0.014)	(7.4)	(0.013)
Share with a flagged winner	28.1%	28.1%	29.2%	29.2%
Observations	7,749	7,749	7,497	7,497

Note: PMGSY-linked tenders. Columns 1–2: share of a tender's quality-monitor inspections graded satisfactory, and an indicator that all were. Column 3: days between the actual and the original stipulated completion date, averaged over the tender's completed roads. Column 4: indicator that a road finished more than 180 days after the original deadline. Panel A is the specification of column 1 of Table 6. Panel B restricts to tenders in which a flagged pair bid (collusive tenders and tenders in which a flagged pair bid and lost) and adds the total number of bidders, as in column 4. Panel C replaces the collusive-tender indicator with the same indicator for matched control pairs. Panel D replaces it with an indicator that the winner belongs to any flagged pair. All columns include district × organization × year and product × year fixed effects and the tender controls of Table 6. Standard errors, clustered by district, in parentheses. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

Table 10: Annual transfer from the procurement budget to firms in flagged pairs

	$\hat{\beta}$	Procurement ($\$$ bn/yr)	Collusive tenders: Volume ($\$$ bn/yr)	Collusive tenders: Transfer ($\$$ m/yr)	Collusive tenders: 10 years ($\$$ bn)
Panel A: Price effect from each column of Table 6					
(1) Full sample	3.19%	32.89	4.97	158.7	1.59
(2) Full sample, total bidders	3.16%	32.89	4.97	156.9	1.57
(3) Tenders with a flagged winner	3.31%	32.89	4.97	164.3	1.64
(4) Tenders with a flagged pair bidding	3.31%	32.89	4.97	164.4	1.64
Panel B: Column-1 price effect, scaled to national procurement					
Classified Table 7 tenders, top decile of pairs	3.19%	32.89	4.97	158.7	1.59
Classified Table 7 tenders, top $\hat{\Lambda}$ of pairs	3.19%	32.89	13.42	428.6	4.29
+ unclassified Table 7 tenders	3.19%	68.63	28.00	894.4	8.94
+ trimmed tenders (all portal procurement)	3.19%	101.59	41.45	1,324.0	13.24
National procurement (8% of GDP)	3.19%	307.53	125.49	4,007.9	40.08

The transfer is $\hat{\beta} \cdot \overline{Q}$. $\hat{\beta}$ is the collusive-tender coefficient of the column of Table 6 named in the row (column 5 is a per-standard-deviation effect and is omitted). $\overline{Q}$ ("Volume") is the annual value of collusive tenders, in which one member of a flagged pair won and the other bid, averaged over 2022–2025, the years with full portal coverage. The 10-year column holds the annual figure for 10 years.

Panel A uses the classified tenders of the sample on which Table 6 is estimated, in which at least two bidders belong to a scored pair so that a flag could have fired (4,012,058 of 4,594,861 tenders). Panel B holds the column-1 price effect fixed throughout. Its first row repeats panel A, column 1; its second flags the top $\hat{\Lambda}$ (34%) of scored pairs, the estimated share of colluding pairs, instead of the top decile, which raises the collusive share of classified value from 15.1% to 40.8%. The last three rows hold that share fixed and apply it to wider bases: first adding the unclassified tenders of the sample; then the tenders outside the sample (lottery awards, contract-to-estimate ratios outside $[0.3, 3.0]$, missing estimates, and tenders the sample filters drop), which gives every portal tender with a recorded contract value (5,385,398 tenders); and finally Indian public procurement at 8% of FY24 GDP. These rows assume that collusion is as common where it cannot be detected as where it can.

Flagged pairs are the top decile of scored pairs by the estimated likelihood ratio (Section 4).

The figures in this exercise are suggestive, and they are likely lower bounds. The transfer does not account for the bidders that colluding firms deter from entering. Collusive tenders attract 0.05 fewer financial bidders than comparable tenders (s.e. 0.03), and each additional bidder is associated with a 2.9 log point lower winning price ratio, which implies a further cost of about 0.2 percentage points that the reduced-form coefficient on Collusive_i does not capture. In addition, measurement error in the classification of collusion attenuates the estimated price effects.

9.3 Quantities and the efficiency loss

If the procuring organization buys less when prices rise, collusion also causes an efficiency loss. With ε the absolute elasticity of the quantity procured with respect to price, the loss is the Harberger triangle

$$\text{DWL} \;=\; \tfrac{1}{2}\,\Delta P\,|\Delta Q| \;\approx\; \tfrac{1}{2}\,\hat{\beta}^2\,\varepsilon\,\overline{Q}, \qquad \frac{\text{DWL}}{\text{Transfer}} \;=\; \tfrac{1}{2}\hat{\beta}\varepsilon, \tag{13}$$

with $\hat{\beta}$ and $\overline{Q}$ as in equation (12). Because the triangle is second order in $\hat{\beta}$ and the estimated price effect is small, the efficiency loss is small. An organization that holds its nominal budget fixed buys exactly one percent less for every percent price increase. An upper bound on the elasticity is thus $\varepsilon = 1$. At the headline $\hat{\beta} = 3.19$ percent and $\varepsilon = 1$, the efficiency loss is 1.6 percent of the transfer, about USD 3 million a year. Under these assumptions, the cost of collusion in Indian public procurement is almost entirely a transfer from the government to colluding firms.

10 Conclusion

This paper develops a pair-level likelihood-ratio screen for bidder coordination and applies it to the universe of tenders on India's government-developed e-procurement portals. The screen compares the observed distribution of behavioral markers across bidder pairs with the distribution generated by placebo pairs, which are built from the same firms' bids outside the relationship. It accommodates several latent types of coordination and requires no convictions for training.

Applying the screen to India yields three lessons about collusion in public procurement. First, collusion is heterogeneous: colluding pairs combine different markers, so screening one marker at a time can miss pairs that are moderately elevated on several markers at once. Second, collusion is widespread and follows market conditions and the structure of procurement: it affects a large share of tenders, is more common where fewer firms bid, and is more common under rural local bodies than under state departments, and under state departments than under the central government. Third, collusive tenders settle at higher prices, and linked road contracts show no difference in quality, so the cost of collusion is primarily a transfer from the government to colluding firms.

This transfer reduces welfare to the extent that the funds would otherwise have financed projects with high social returns, or that raising them requires distortionary taxation. The method uses records procuring organizations already collect, and its ranking can direct investigations to the pairs most likely to collude, although it does not establish liability. A deployed screen may also change bidder behavior. Measuring that equilibrium response is beyond the design studied here.

References

Aigner, Dennis J. (1973). "Regression with a Binary Independent Variable Subject to Errors of Observation". In: *Journal of Econometrics* 1.1, pp. 49–59. DOI: 10.1016/0304-4076(73)90005-5 ↩ 22.

Allende, Claudia, Juan Pablo Atal, Rodrigo Carril, José Ignacio Cuesta, and Andrés González-Lira (2024). "Drivers of Public Procurement Prices: Evidence from Pharmaceutical Markets". In: *International Journal of Industrial Organization* 96, p. 103086. DOI: 10.1016/j.ijindorg.2024.103086 ↩ 31.

Anderson, J. A. (1972). "Separate Sample Logistic Discrimination". In: *Biometrika* 59.1, pp. 19–35. DOI: 10.1093/biomet/59.1.19 ↩ 19.

Angus, John E. (1994). "The Probability Integral Transform and Related Results". In: *SIAM Review* 36.4, pp. 652–654. DOI: 10.1137/1036146 ↩ 17.

Asker, John (2010). "A Study of the Internal Organization of a Bidding Cartel". In: *American Economic Review* 100.3, pp. 724–762. DOI: 10.1257/aer.100.3.724 ↩ 7.

Athey, Susan, Kyle Bagwell, and Chris Sanchirico (2004). "Collusion and Price Rigidity". In: *Review of Economic Studies* 71.2, pp. 317–349. DOI: 10.1111/0034-6527.00286 ↩ 3, 6, 12, A-2.

Bajari, Patrick, Stephanie Houghton, and Steven Tadelis (2014). "Bidding for Incomplete Contracts: An Empirical Analysis of Adaptation Costs". In: *American Economic Review* 104.4, pp. 1288–1319. DOI: 10.1257/aer.104.4.1288 ↩ 31.

Bajari, Patrick and Lixin Ye (2003). "Deciding Between Competition and Collusion". In: *Review of Economics and Statistics* 85.4, pp. 971–989. DOI: 10.1162/003465303772815871 ↩ 3–6, 13, 23.

Bandiera, Oriana, Andrea Prat, and Tommaso Valletti (2009). "Active and Passive Waste in Government Spending: Evidence from a Policy Experiment". In: *American Economic Review* 99.4, pp. 1278–1308. DOI: 10.1257/aer.99.4.1278 ↩ 2.

Baránek, Bruno, Leon Musolff, and Vítězslav Titl (2026). "Detection of Collusive Networks in Multistage Auctions". In: *American Economic Journal: Applied Economics* 18.4, pp. 439–471. DOI: 10.1257/app.20240232 ↩ 6, 7, 26, A-20.

Barkley, Aaron (2023). "The Human Cost of Collusion: Health Effects of a Mexican Insulin Cartel". In: *Journal of the European Economic Association* 21.5, pp. 1865–1904. DOI: 10.1093/jeea/jvad005 ↩ 6, 7.

Benjamini, Yoav and Yosef Hochberg (1995). "Controlling the False Discovery Rate: A Practical and Powerful Approach to Multiple Testing". In: *Journal of the Royal Statistical Society. Series B (Methodological)* 57.1, pp. 289–300. URL: https://www.jstor.org/stable/2346101 (visited on 08/31/2026) ↩ A-14.

Bergman, Mats A., Johan Lundberg, Sofia Lundberg, and Johan Y. Stake (2020). "Interactions Across Firms and Bid Rigging". In: *Review of Industrial Organization* 56.1, pp. 107–130. DOI: 10.1007/s11151-018-09676-0 ↩ A-14.

Best, Michael Carlos, Jonas Hjort, and David Szakonyi (2023). "Individuals and Organizations as Sources of State Effectiveness". In: *American Economic Review* 113.8, pp. 2121–2167. DOI: 10.1257/aer.20191598 ↩ 2, 6, 31.

Blanchard, Gilles, Gyemin Lee, and Clayton Scott (2010). "Semi-Supervised Novelty Detection". In: *Journal of Machine Learning Research* 11.99, pp. 2973–3009. URL: http://jmlr.org/papers/v11/blanchard10a.html (visited on 09/30/2026) ↩ 15.

Bosio, Erica, Simeon Djankov, Edward Glaeser, and Andrei Shleifer (2022). "Public Procurement in Law and Practice". In: *American Economic Review* 112.4, pp. 1091–1117. DOI: 10.1257/aer.20200738 ↩ 2, 6.

Bound, John, Charles Brown, and Nancy Mathiowetz (2001). "Measurement Error in Survey Data". In: *Handbook of Econometrics*. Vol. 5. Elsevier, pp. 3705–3843. URL: https://EconPapers.repec.org/RePEc:eee:ecochp:5-59 (visited on 07/28/2026) ↩ 22.

Brugués, Felipe, Javier Brugués, and Samuele Giambra (2024). "Political Connections and Misallocation of Procurement Contracts: Evidence from Ecuador". In: *Journal of Development Economics* 170, p. 103296. DOI: 10.1016/j.jdeveco.2024.103296 ↩ 2, 6.

Cameron, A. Colin, Jonah B. Gelbach, and Douglas L. Miller (2011). "Robust Inference With Multiway Clustering". In: *Journal of Business & Economic Statistics* 29.2, pp. 238–249. DOI: 10.1198/jbes.2010.07136 ↩ A-16.

Campello, Ricardo J. G. B., Davoud Moulavi, and Joerg Sander (2013). "Density-Based Clustering Based on Hierarchical Density Estimates". In: *Advances in Knowledge Discovery and Data Mining*. Ed. by Jian Pei, Vincent S. Tseng, Longbing Cao, Hiroshi Motoda, and Guandong Xu. Berlin, Heidelberg: Springer, pp. 160–172. ISBN: 978-3-642-37456-2. DOI: 10.1007/978-3-642-37456-2_14 ↩ A-48.

Caoui, El Hadi (2022). "A Study of Umbrella Damages from Bid Rigging". In: *Journal of Law and Economics* 65.2, pp. 239–277. DOI: 10.1086/717755 ↩ 7.

Carbone, Carlotta, Francesco Calderoni, and Maria Jofre (2024). "Bid-Rigging in Public Procurement: Cartel Strategies and Bidding Patterns". In: *Crime, Law and Social Change* 82.2, pp. 249–281. DOI: 10.1007/s10611-024-10142-0 ↩ A-2.

CCI (2026). *Orders of the Competition Commission of India*. URL: https://www.cci.gov.in/antitrust/orders (visited on 05/31/2026) ↩ A-50.

Central Vigilance Commission (2021). *Vigilance Manual*. URL: https://cvc.gov.in/vigilancemanual (visited on 05/27/2026) ↩ 12.

Chassang, Sylvain, Kei Kawai, Jun Nakabayashi, and Juan Ortner (2022). "Robust Screens for Noncompetitive Bidding in Procurement Auctions". In: *Econometrica* 90.1, pp. 315–346. DOI: 10.3982/ECTA17155 ↩ 12, 13, 31, A-3.

Chassang, Sylvain and Juan Ortner (2019). "Collusion in Auctions with Constrained Bids: Theory and Evidence from Public Procurement". In: *Journal of Political Economy* 127.5, pp. 2269–2300. DOI: 10.1086/701812 ↩ 40.

Chen, Qianmiao (2025). *Rigging the Scores : Corruption through Scoring Rule Manipulation in Public Procurement Auctions*. World Bank Policy Research Working Paper, 11267. URL: https://documents.worldbank.org/en/publication/documents-reports/documentdetail/099651112022531100 (visited on 10/19/2026) ↩ 6.

Clark, Robert, Decio Coviello, and Adriano De Leverano (2025). "Complementary Bidding and the Collusive Arrangement: Evidence from an Antitrust Investigation". In: *American Economic Journal: Microeconomics* 17.2, pp. 66–93. DOI: 10.1257/mic.20210450 ↩ 12, 13, 36.

Clark, Robert, Decio Coviello, Jean-François Gauthier, and Art Shneyerov (2018). "Bid Rigging and Entry Deterrence in Public Procurement: Evidence from an Investigation into Collusion and Corruption in Quebec". In: *The Journal of Law, Economics, and Organization* 34.3, pp. 301–363. DOI: 10.1093/jleo/ewy011 ↩ 7.

Colonnelli, Emanuele and Mounu Prem (2022). "Corruption and Firms". In: *The Review of Economic Studies* 89.2, pp. 695–732. DOI: 10.1093/restud/rdab040 ↩ 6.

Competition Commission of India (2018). *Nagrik Chetna Manch v. Fortified Security Solutions & Ors.* URL: https://cci.gov.in/images/antitrustorder/en/5020151652433457.pdf (visited on 05/27/2026) ↩ 11.

Competition Commission of India (2024). *Diagnostic Toolkit Towards Competitive Tenders: For Public Procurement Officers*. URL: https://www.cci.gov.in/images/whatsnew/en/diagnostic-toolkit-for-public-procurement1747723682.pdf (visited on 05/27/2026) ↩ 3, 11–13, A-1, A-2, A-3.

Comptroller and Auditor General of India (2023). *Performance Audit of Implementation of eProcurement System*. Government of Tamil Nadu, Report No. 4 of 2023. URL: https://cag.gov.in/uploads/download_audit_report/2023/Report-No-4-of-2023-eProcurement-English-0668269bb4dda56.82968764.pdf (visited on 05/27/2026) ↩ 24.

Conley, Timothy G. and Francesco Decarolis (2016). "Detecting Bidders Groups in Collusive Auctions". In: *American Economic Journal: Microeconomics* 8.2, pp. 1–38. DOI: 10.1257/mic.20130254 ↩ 3, 6, 12, 24, A-14.

Czado, Claudia, Tilmann Gneiting, and Leonhard Held (2009). "Predictive Model Assessment for Count Data". In: *Biometrics* 65.4, pp. 1254–1261. DOI: 10.1111/j.1541-0420.2009.01191.x ↩ 17.

DiNardo, John, Nicole M. Fortin, and Thomas Lemieux (1996). "Labor Market Institutions and the Distribution of Wages, 1973-1992: A Semiparametric Approach". In: *Econometrica* 64.5, pp. 1001–1044. DOI: 10.2307/2171954 ↩ 15, 19.

Efron, Bradley (2004). "Large-Scale Simultaneous Hypothesis Testing: The Choice of a Null Hypothesis". In: *Journal of the American Statistical Association* 99.465, pp. 96–104. DOI: 10.1198/016214504000000089 ↩ 4, 7, 17.

– (2007). "Size, Power and False Discovery Rates". In: *The Annals of Statistics* 35.4, pp. 1351–1377. DOI: 10.1214/009053606000001460 ↩ A-11.

– (2008). "Microarrays, Empirical Bayes and the Two-Groups Model". In: *Statistical Science* 23.1, pp. 1–22. DOI: 10.1214/07-STS236 ↩ 4, 7, 15.

Efron, Bradley, Robert Tibshirani, John D Storey, and Virginia Tusher (2001). "Empirical Bayes Analysis of a Microarray Experiment". In: *Journal of the American Statistical Association* 96.456, pp. 1151–1160. DOI: 10.1198/016214501753382129 ↩ 15.

Elkan, Charles and Keith Noto (2008). "Learning Classifiers from Only Positive and Unlabeled Data". In: *Proceedings of the 14th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*. KDD08: The 14th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining. Las Vegas Nevada USA: ACM, pp. 213–220. ISBN: 978-1-60558-193-4. DOI: 10.1145/1401890.1401920 ↩ 15.

Fazio, Dimas and Alminas Zaldokas (2026). *Kamikazes in Public Procurements: BidRigging and Real Non-Market Outcomes*. egi Finance Working Paper, 1054/2025. URL: https://www.ecgi.global/sites/default/files/2026-06/kamikazes-in-public-procurements.pdf ↩ 7, A-2.

GADM (2026). *Database of Global Administrative Areas, Version 4.1*. URL: https://gadm.org/ (visited on 01/13/2026) ↩ A-47.

García Rodríguez, Manuel J., Vicente Rodríguez-Montequín, Pablo Ballesteros-Pérez, Peter E. D. Love, and Regis Signor (2022). "Collusion Detection in Public Procurement Auctions with Machine Learning Algorithms". In: *Automation in Construction* 133, p. 104047. DOI: 10.1016/j.autcon.2021.104047 ↩ 14.

Genovese, Christopher and Larry Wasserman (2004). "A Stochastic Process Approach to False Discovery Control". In: *The Annals of Statistics* 32.3, pp. 1035–1061. DOI: 10.1214/009053604000000283 ↩ 15, 17.

GeoNames (2025). *GeoNames Geographical Database*. URL: https://www.geonames.org/about.html (visited on 07/13/2025) ↩ A-47.

Gerardino, Maria Paula, Stephan Litschig, and Dina Pomeranz (2024). "Distortion by Audit: Evidence from Public Procurement". In: *American Economic Journal: Applied Economics* 16.4, pp. 71–108. DOI: 10.1257/app.20220512 ↩ 6.

Government of Assam (2017). *The Assam Public Procurement Act, 2017*. URL: https://www.indiacode.nic.in/bitstream/123456789/13886/1/the_assam_public_procurement_act%2C_2017.pdf ↩ 8.

Government of India (2017). *General Financial Rules, Updated up to 31.07.2024*. URL: https://doe.gov.in/files/circulars_document/FInal_GFR_upto_31_07_2024.pdf (visited on 09/01/2025) ↩ 8.

– (2022). *Open Government Data (OGD) Platform India*. URL: https://data.gov.in (visited on 06/23/2026) ↩ A-47.

Government of India, Ministry of Finance (2014). *Office Memorandum: Implementation of Comprehensive End-to-End e-Procurement*. URL: https://eprocure.gov.in/cppp/rulesandprocs/kbadqkdlcswfjdelrquehwuxcfmijmuixngudufgbuubgubfugbububjxcgfvsbdihbgfGhdfgFHytyhRtMTk4Nzg= (visited on 05/05/2025) ↩ 9.

Government of India, Ministry of Finance, Department of Expenditure (2017). *Manual for Procurement of Goods 2017*. URL: https://doe.gov.in/manual-procurement-goods-0 (visited on 05/27/2026) ↩ 11.

Government of Odisha (2025). *Odisha General Financial Rules 2023, Updated to 04.06.2025*. URL: https://finance.odisha.gov.in/sites/default/files/2026-02/OGFR%202023_updated%20upto%2004.06.2025.pdf ↩ 8.

Government of Punjab (2019). *The Punjab Transparency In Public Procurement Act, 2019*. URL: https://eprocure.gov.in/cppp/mmprulesandprocs/kbadqkdlcswfjdelrquehwuxcfmijmuixngudufgbuubgubfugbububjxcgfvsbdihbgfGhdfgFHytyhRtODI5ODcxMg== ↩ 8.

Government of Rajasthan (2012). *The Rajasthan Transparency in Public Procurement Act, 2012*. URL: https://sppp.rajasthan.gov.in/actrulesprocedures.php (visited on 05/27/2026) ↩ 8.

Graham, Daniel A. and Robert C. Marshall (1987). "Collusive Bidder Behavior at Single-Object Second-Price and English Auctions". In: *Journal of Political Economy* 95.6, pp. 1217–1239. DOI: 10.1086/261512 ↩ 29.

Harrington, Joseph E and David Imhof (2022). "Cartel Screening and Machine Learning". In: *Computational Antitrust* 2 ↩ 6.

Huber, Martin and David Imhof (2019). "Machine Learning with Screens for Detecting Bid-Rigging Cartels". In: *International Journal of Industrial Organization* 65, pp. 277–301. DOI: 10.1016/j.ijindorg.2019.04.002 ↩ 6, A-3.

Imhof, David (2017). "Econometric Tests to Detect Bid-Rigging Cartels: Does It Work?" In: URL: https://folia.unifr.ch/global/documents/305651 (visited on 09/17/2026) ↩ A-3.

Ishii, Rieko (2009). "Favor Exchange in Collusion: Empirical Study of Repeated Procurement Auctions in Japan". In: *International Journal of Industrial Organization* 27.2, pp. 137–144. DOI: 10.1016/j.ijindorg.2008.05.006 ↩ 3, 6, 12, 13, A-2.

– (2014). "Bid Roundness Under Collusion in Japanese Procurement Auctions". In: *Review of Industrial Organization* 44.3, pp. 241–254. DOI: 10.1007/s11151-013-9408-6 ↩ 3, 6, 12, 13, A-2.

Jaro, Matthew A. (1989). "Advances in Record-Linkage Methodology as Applied to Matching the 1985 Census of Tampa, Florida". In: *Journal of the American Statistical Association* 84.406, pp. 414–420. DOI: 10.1080/01621459.1989.10478785 ↩ A-49.

Jin, Jiashun and T. Tony Cai (2007). "Estimating the Null and the Proportion of Nonnull Effects in Large-Scale Multiple Comparisons". In: *Journal of the American Statistical Association* 102.478, pp. 495–506. URL: https://www.jstor.org/stable/27639880 (visited on 09/28/2026) ↩ 17.

Kang, Karam and Robert A Miller (2022). "Winning by Default: Why Is There So Little Competition in Government Procurement?" In: *The Review of Economic Studies* 89.3, pp. 1495–1556. DOI: 10.1093/restud/rdab051 ↩ 36.

Kawai, Kei and Jun Nakabayashi (2022). "Detecting Large-Scale Collusion in Procurement Auctions". In: *Journal of Political Economy* 130.5, pp. 1364–1411. DOI: 10.1086/718913 ↩ A-20.

Kawai, Kei, Jun Nakabayashi, Juan Ortner, and Sylvain Chassang (2023). "Using Bid Rotation and Incumbency to Detect Collusion: A Regression Discontinuity Approach". In: *Review of Economic Studies* 90.1, pp. 376–403. DOI: 10.1093/restud/rdac013 ↩ 3, 6, 12, 13, 31, A-2.

Kawai, Kei, Jun Nakabayashi, and Daichi Shimamoto (2025). "A Study of Bid-Rigging in Procurement Auctions: Evidence from Indonesia, Georgia, Mongolia, Malta, and State of California". In: *Journal of Political Economy Microeconomics*. DOI: 10.1086/737570 ↩ 6, A-2.

Kline, Patrick, Evan K. Rose, and Christopher Walters (2022). "Systemic Discrimination Among Large U.S. Employers". In: *The Quarterly Journal of Economics* 137.4, pp. 1963–2036. DOI: 10.1093/qje/qjac024 ↩ 15.

– (2024). "A Discrimination Report Card". In: *American Economic Review* 114.8, pp. 2472–2525. DOI: 10.1257/aer.20230700 ↩ 4, 7.

Kline, Patrick and Christopher Walters (2021). "Reasonable Doubt: Experimental Detection of Job-Level Employment Discrimination". In: *Econometrica* 89.2, pp. 765–792. DOI: 10.3982/ECTA17489 ↩ 4, 7, 15.

Krasnokutskaya, Elena (2011). "Identification and Estimation of Auction Models with Unobserved Heterogeneity". In: *The Review of Economic Studies* 78.1, pp. 293–327. DOI: 10.1093/restud/rdq004 ↩ 31.

Lancaster, Tony and Guido Imbens (1996). "Case-Control Studies with Contaminated Controls". In: *Journal of Econometrics* 71.1, pp. 145–160. DOI: 10.1016/0304-4076(94)01698-4 ↩ 15.

Lehne, Jonathan, Jacob N. Shapiro, and Oliver Vanden Eynde (2018). "Building Connections: Political Corruption and Road Construction in India". In: *Journal of Development Economics* 131, pp. 62–78. DOI: 10.1016/j.jdeveco.2017.10.009 ↩ 2, 3, 6.

Levenshtein, V. I. (1966). "Binary Codes Capable of Correcting Deletions, Insertions and Reversals". In: *Soviet Physics Doklady* 10, p. 707. URL: https://ui.adsabs.harvard.edu/abs/1966SPhD...10..707L (visited on 06/19/2026) ↩ A-49.

Lewis-Faupel, Sean, Yusuf Neggers, Benjamin A. Olken, and Rohini Pande (2016). "Can Electronic Procurement Improve Infrastructure Provision? Evidence from Public Works in India and Indonesia". In: *American Economic Journal: Economic Policy* 8.3, pp. 258–283. DOI: 10.1257/pol.20140258 ↩ 2, 3, 6.

Liscow, Zachary, Cailin Slattery, and Will Nober (2025). *State Capacity and Infrastructure Costs* ↩ 9.

Marshall, Robert C. and Leslie M. Marx (2012). *The Economics of Collusion: Cartels and Bidding Rings*. The MIT Press. ISBN: 978-0-262-30150-3. DOI: 10.7551/mitpress/9011.001.0001 ↩ 29, A-32.

McAfee, R Preston and John McMillan (1992). "Bidding Rings". In: *The American Economic Review*, pp. 579–599 ↩ 3, 5, 13, 36.

Meinshausen, Nicolai and John Rice (2006). "Estimating the Proportion of False Null Hypotheses among a Large Number of Independently Tested Hypotheses". In: *The Annals of Statistics* 34.1, pp. 373–393. DOI: 10.1214/009053605000000741 ↩ 15, 17.

Ministry of Commerce & Industry (2026). *GeM Completes a Decade of Transforming Public Procurement with Cumulative GMV Exceeding INR 20 Lakh Crore*. URL: https://www.pib.gov.in/www.pib.gov.in/Pressreleaseshare.aspx?PRID=2296138 (visited on 10/23/2026) ↩ 8.

Neyman, Jerzy and Egon Sharpe Pearson (1933). "IX. On the Problem of the Most Efficient Tests of Statistical Hypotheses". In: *Philosophical Transactions of the Royal Society of London, Series A: Containing Papers of a Mathematical or Physical Character* 231.694–706, pp. 289–337. DOI: 10.1098/rsta.1933.0009 ↩ 4, 15.

OECD (2025a). *Government at a Glance 2025*. Government at a Glance. OECD Publishing. ISBN: 978-92-64-31390-3 978-92-64-36321-2 978-92-64-59508-8. DOI: 10.1787/0efd0bcd-en ↩ 2.

– (2025b). *OECD Guidelines for Fighting Bid Rigging in Public Procurement (2025 Update)*. OECD Publishing. ISBN: 978-92-64-48350-7 978-92-64-54949-4. DOI: 10.1787/cbe05a56-en ↩ 2, 3, 12.

Olken, Benjamin A. (2007). "Monitoring Corruption: Evidence from a Field Experiment in Indonesia". In: *Journal of Political Economy* 115.2, pp. 200–249. DOI: 10.1086/517935 ↩ 6.

Panda, Prabir and G. P. Sahu (2015). "Electronic Government Procurement Implementation in India: A Cross Sectional Study". In: *International Journal of Business Information Systems*. DOI: 10.1504/ijbis.2015.066124 ↩ 9.

Patra, Rohit Kumar and Bodhisattva Sen (2016). "Estimation of a Two-Component Mixture Model with Applications to Multiple Testing". In: *Journal of the Royal Statistical Society Series B: Statistical Methodology* 78.4, pp. 869–893. DOI: 10.1111/rssb.12148 ↩ 15, 17.

Porter, Robert H. and J. Douglas Zona (1993). "Detection of Bid Rigging in Procurement Auctions". In: *Journal of Political Economy* 101.3, pp. 518–538. DOI: 10.1086/261885 ↩ 3, 6, 12, 13, 36, A-2.

– (1999). "Ohio School Milk Markets: An Analysis of Bidding". In: *RAND Journal of Economics* 30.2, pp. 263–288. DOI: 10.2307/2556080 ↩ 3, 6, 7.

Prentice, R. L. and R. Pyke (1979). "Logistic Disease Incidence Models and Case-Control Studies". In: *Biometrika* 66.3, pp. 403–411. DOI: 10.1093/biomet/66.3.403 ↩ 19.

Qin, Jing and Biao Zhang (1997). "A Goodness-of-Fit Test for Logistic Regression Models Based on Case-Control Data". In: *Biometrika* 84.3, pp. 609–618. DOI: 10.1093/biomet/84.3.609 ↩ 15, 19.

Reimers, Nils and Iryna Gurevych (2019). "Sentence-BERT: Sentence Embeddings Using Siamese BERT-Networks". In: *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*. EMNLP-IJCNLP 2019. Ed. by Kentaro Inui, Jing Jiang, Vincent Ng, and Xiaojun Wan. Hong Kong, China: Association for Computational Linguistics, pp. 3982–3992. DOI: 10.18653/v1/D19-1410 ↩ A-47.

Rosenblatt, Murray (1952). "Remarks on a Multivariate Transformation". In: *The Annals of Mathematical Statistics* 23.3, pp. 470–472. DOI: 10.1214/aoms/1177729394 ↩ 17.

Rüschendorf, Ludger (2009). "On the Distributional Transform, Sklar's Theorem, and the Empirical Copula Process". In: *Journal of Statistical Planning and Inference*. Special Issue: The 8th Tartu Conference on Multivariate Statistics & The 6th Conference on Multivariate Distributions with Fixed Marginals 139.11, pp. 3921–3927. DOI: 10.1016/j.jspi.2009.05.030 ↩ 17.

Shapley, L. S. (1953). "A Value for N-Person Games". In: *Contributions to the Theory of Games, Volume II*. Princeton University Press, pp. 307–318. ISBN: 978-1-4008-8197-0. URL: https://www.degruyterbrill.com/document/doi/10.1515/9781400881970-018/html?lang=en (visited on 10/01/2026) ↩ 20, A-12.

Starc, Amanda and Thomas G. Wollmann (2025). "Does Entry Remedy Collusion? Evidence from the Generic Prescription Drug Cartel". In: *American Economic Review* 115.5, pp. 1400–1438. DOI: 10.1257/aer.20220276 ↩ 7.

Stigler, George J. (1964). "A Theory of Oligopoly". In: *Journal of Political Economy* 72.1, pp. 44–61. URL: https://www.jstor.org/stable/1828791 (visited on 09/21/2026) ↩ 3, 13, 29.

Szucs, Ferenc (2024). "Discretion and Favoritism in Public Procurement". In: *Journal of the European Economic Association* 22.1, pp. 117–160. DOI: 10.1093/jeea/jvad017 ↩ 6.

Teicher, Henry (1963). "Identifiability of Finite Mixtures". In: *The Annals of Mathematical Statistics* 34.4, pp. 1265–1269. DOI: 10.1214/aoms/1177703862 ↩ 15.

US DOJ (2015). *Antitrust Division | Red Flags Of Collusion*. URL: https://www.justice.gov/atr/red-flags-collusion (visited on 09/28/2026) ↩ 3, 12.

Wachs, Johannes and János Kertész (2019). "A Network Approach to Cartel Detection in Public Auction Markets". In: *Scientific Reports* 9.1, p. 10818. DOI: 10.1038/s41598-019-47198-1 ↩ A-2.

Wallimann, Hannes, David Imhof, and Martin Huber (2023). "A Machine Learning Approach for Flagging Incomplete Bid-Rigging Cartels". In: *Computational Economics* 62.4, pp. 1669–1720. DOI: 10.1007/s10614-022-10315-w ↩ 6, 14.

Winkler, William E (1990). "String Comparator Metrics and Enhanced Decision Rules in the Fellegi-Sunter Model of Record Linkage". In: *Proceedings of the Section on Survey Research Methods, American Statistical Association*, pp. 354–359 ↩ A-49.

Yakowitz, Sidney J. and John D. Spragins (1968). "On the Identifiability of Finite Mixtures". In: *The Annals of Mathematical Statistics* 39.1, pp. 209–214. DOI: 10.1214/aoms/1177698520 ↩ 15.

Appendix

A Method

A.1 The twelve markers

The markers are built from the warning signals in the Competition Commission of India's procurement toolkit (Competition Commission of India 2024). The toolkit describes conduct a procurement officer might observe in a single tender; I translate each signal into a statistic over the repeated co-bids of a bidder pair. Table A.1 defines each marker in the order of Table 2 and lists the toolkit signal behind it, the studies that use the same signal, and the share of pairs for which it is measured; Table A.2 gives examples of signals that are not used because they are not recorded in the data, not defined for a pair of bidders, or measured only through the winning bid or the price level.

Table A.1: Markers and the warning signals of the CCI toolkit they measure

	Marker	Pair-level statistic	CCI warning signal (toolkit page)	Also used in	Pairs (%)
1	Submission proximity	Share of co-bids with both submission times recorded in which the two bids are submitted no more than 30 minutes apart; co-bids in which one member stops at the technical stage are included	"Bids submitted within short span of time"; electronic documents with "edits made at one and the same time" (20)		100.0
2	Round bid difference	Share of eligible co-bids in which the price gap is round: the relative gap lies near a multiple of 0.5 percentage points or near $1/n$ ($n = 10, \ldots, 20$), or the absolute gap lies near a round rupee amount for the bids' magnitude. Eligible: both members priced, relative gap at most 20.5%, neither bid itself round, no identical bids at a focal price, no focal-price bids in a lottery tender, and bids not anchored to the government estimate	Prices "different from one another by the same margin" across procedures of the same buyer (21)	Ishii 2014	36.4
3	Margin steadiness	Standard deviation of $\log(p_i/p_j)$ over co-bids in which both members submit distinct positive prices at most a factor of three apart, multiplied by -1; defined when both firms have at least ten priced tenders without the partner	Same margin as above; "prices were the same for a long period of time" (21)	Athey, Bagwell, and Sanchirico 2004	44.5
4	Submission-order role	$\lvert 2p - 1 \rvert$, with p the share of co-bids with both submission times and a lower bidder in which the lower bidder submits first	"Bids were prepared by the same person or were prepared jointly" (19); a firm "submits both its own and a competitor's bid" (22)		56.1
5	Submission-order stability	$\lvert 2p - 1 \rvert$, with p the share of co-bids with both submission times in which member i submits first, whichever member bids lower	Same signals as marker 4 (19, 22)		61.5
6	Targeted disqualification	Count of co-bids in which exactly one member is technically disqualified, standardized by the mean and standard deviation implied by each firm's disqualification rate in tenders without the partner; co-bids in which both are disqualified are excluded	A bidder "technically disqualified although was found qualified in the past" (19); an incomplete bid from a bidder with "the potential to do the work" (20)	Fazio and Zaldokas 2026	70.6
7	Outsider disqualification	Share of the other bidders with resolved technical status that are disqualified, averaged over the pair's co-bids	"Only one bidder found qualified after number of stages of a tender" (19)		97.8
8	Lower-bidder stability	$\lvert 2p - 1 \rvert$, with p the share of co-bids with a lower bidder in which member i bids lower	"The same supplier is often the lowest bidder" (18); a bidder "always submits the lowest offers" (22)	Porter and Zona 1993	57.4
9	Lower-bidder rotation	$\lvert s - 2p(1-p) \rvert$, with s the share of consecutive co-bids in which the lower bidder switches and p as in marker 8; $2p(1-p)$ is the switch rate under independent draws, so both over-regular alternation and over-persistence score high	"Each company seems to take a turn being the winning bidder" (18)	Ishii 2009; Kawai, Nakabayashi, Ortner, et al. 2023	48.6
10	Backlog difference	Share of co-bids in which the member whose log 360-day backlog is further above its own mean over the pair's co-bids submits the higher bid	Same turn-taking signal (18)	Kawai, Nakabayashi, Ortner, et al. 2023; Kawai, Nakabayashi, and Shimamoto 2025	33.1
11	Uncompetitive fringe entrant	Share of co-bids with at least one outsider in which at least one outsider is fringe: it has submitted a qualifying bid but has never been the lowest qualifying bidder anywhere in India	"Certain companies always submit bids but never win" (18, 22)	Porter and Zona 1993; Carbone, Calderoni, and Jofre 2024	98.2
12	Exclusive outsider	Share of co-bids that include an outsider every one of whose tenders also contains both members	Bidders that "always or often participate in bidding but never win" (22)	Wachs and Kertész 2019	100.0

Note: Signals are quoted from Section II of the Competition Commission of India's *Diagnostic Toolkit Towards Competitive Tenders* (Competition Commission of India 2024). Page numbers are those printed in the toolkit. The statistics give in full the definitions shortened in Table 2. A co-bid is a tender in which both members bid. An outsider is any other bidder in it. Member i's price is p_i, and which member is i is arbitrary for markers 5 and 8. *Pairs (%)*: share of the 2,043,939 scored pairs with at least five marker-specific co-bids. This is measurement coverage, not the share classified as suspicious. Markers 1, 4 and 5 need submission times, which the procurement data in the studies cited do not contain.

Table A.2: Examples of warning signals of the CCI toolkit that are not used

Reason	CCI warning signal (toolkit page)	Why it fails	Also used in
Not recorded in the data	"Bids received from same IP address" (20)	IP addresses are not published	
	"Identical mistakes/corrections in the bid documents" (19)	Bid documents are not published; circumstantial	
Not defined for a pair	"None of the bidders have participated in the tender although there are a number of bidders available in the market" (19)	A trait of the market; the number of bidders is an outcome in Section 8	
	"Regular suppliers fail to bid in a tender they would normally be expected to bid" (18)	An absence of bids; the markers are defined on co-bids	
Uses the price outcome	"A large difference between the price of a winning bid and other bids" (21)	Uses the winning bid, the outcome of Section 7; the direction is also ambiguous, since cover bids can sit far above the winning bid or just above it, and the gap depends on cost dispersion and the number of bidders. Markers 2 and 3 use the regularity of the gap, not its size	Imhof 2017; Huber and Imhof 2019; small gaps: Chassang, Kawai, et al. 2022
	"Sudden and identical increases in price … that cannot be explained by cost increases" (20)	Uses the price level and costs	
Measured for too few pairs	"There is a geographic allocation of winning tenders" (18)	Defined for too few pairs to be estimated	

Note: Signals are quoted from Section II of the toolkit (Competition Commission of India 2024). Page numbers are those printed in the toolkit. A signal is used if it is recorded in the portal data, defined for a pair of bidders, and measurable without the winning bid or the price level. The table gives examples for each reason. Of the toolkit's 64 signals, 38 are not recorded, 8 are not defined for a pair, 3 use the price outcome, and the twelve markers of Table A.1 are built from 14 of the remaining 15.

The construction uses raw statistics, not empirical-Bayes estimates. A pair contributes to a marker only when it has at least five marker-specific observations, and the last column of Table A.1 reports how many pairs clear that gate for each marker. Precision enters through the observation gate and the placebo comparison, not through shrinkage.

A.2 Placebo designs and identifying assumptions

This subsection states what each of the three placebo designs assumes and what it delivers. Throughout, $d = (i, j)$ is a pair, T_d the set of its co-bids, and $n_d = |T_d|$. For co-bid $t \in T_d$, x_{it} is the attribute of i that the marker reads: a submission time, a price, a disqualification status. The pair's data are $X_d = (x_{it}, x_{jt})_{t \in T_d}$, and every raw pair-level marker is a deterministic functional $m_d = M(X_d)$. Write $C_d \in \{0, 1\}$ for the pair's latent state, with $C_d = 1$ if the two firms collude and $C_d = 0$ if they do not, so that $\Lambda = \Pr(C_d = 1)$. A design is valid for a marker when the distribution of M applied to the placebo draw equals the distribution of M applied to the pair's own data given $C_d = 0$.

A.2.1 Own-history placebo

For each $t \in T_d$ the draw replaces x_{it} by $\tilde{x}_{it}$, drawn from i's attributes in tenders where j does not bid, and x_{jt} by $\tilde{x}_{jt}$, drawn from j's attributes in tenders where i does not bid, independently across co-bids and across the two firms.

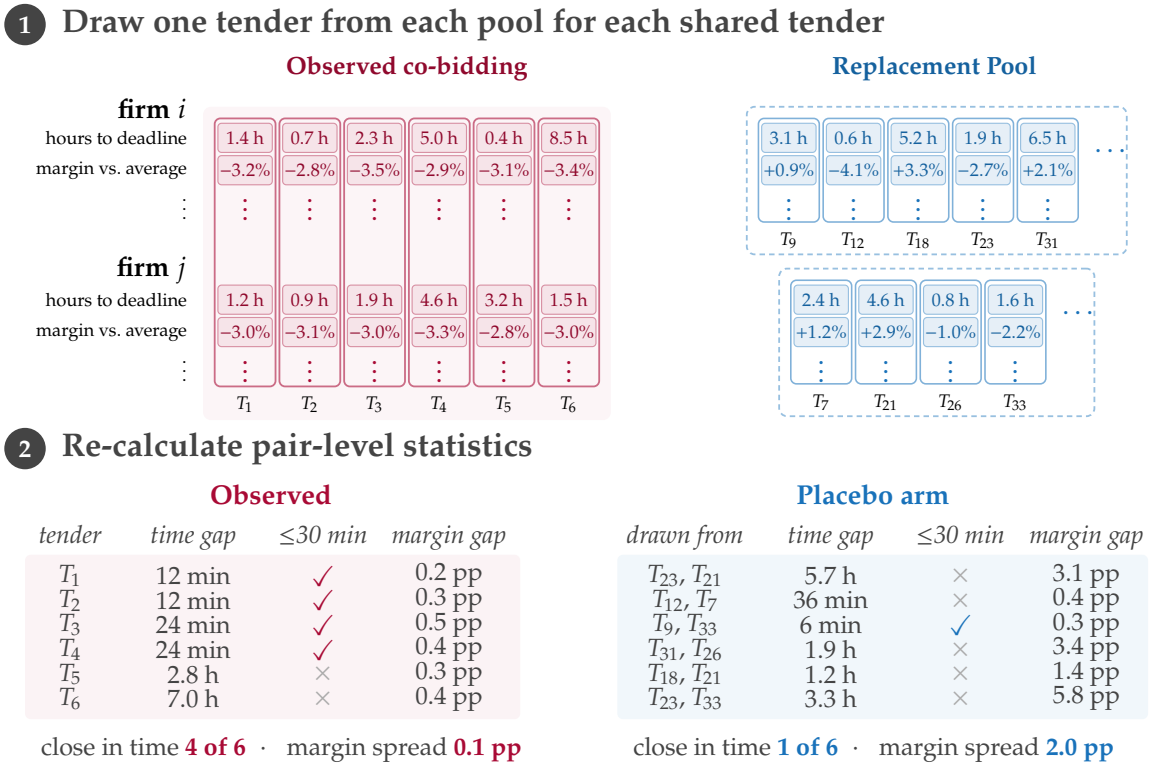


Figure A.1: The own-history redraw

Note: Each member's attribute in a co-bid is replaced by one of its own attributes from a tender the partner did not enter, drawn with replacement and independently across the two firms and across co-bids. The markers are then recomputed on the redrawn values. The two pools are disjoint by construction, so nothing links the two draws. The figure shows a donor tender supplying both of a firm's attributes at once. In the current build each marker family is drawn from its own pool.

Assumption 1 (Own-history exchangeability)**.** Conditional on $C_d = 0$, (i, j) and T_d, the attributes

$(x_{it})_{t \in T_d}$ are i.i.d. with the distribution of i's attributes in tenders where j does not bid, and likewise for j.

Assumption 2 (Pair independence)**.** Conditional on $C_d = 0$, (i, j) and T_d, the sequences $(x_{it})_{t \in T_d}$ and $(x_{jt})_{t \in T_d}$ are independent of each other.

Proposition 1. *Under Assumptions 1 and 2, X_d and $\tilde{X}_d$ have the same distribution, and therefore so do $M(X_d)$ and $M(\tilde{X}_d)$ for every marker M.*

Proof. Assumption 1 fixes the distribution of each firm's sequence as the n_d-fold product of its partner-absent distribution. Assumption 2 makes the two sequences independent of each other. The distribution of X_d is therefore the product measure over the $2n_d$ draws, which is by construction the distribution of $\tilde{X}_d$. Equality in distribution is preserved under the deterministic map M. □

Three remarks on these assumptions. First, Assumption 1 does not require i's other tenders to be free of collusion. If i colludes with a third firm k, that arrangement enters the distribution of i's partner-absent attributes and the draw inherits it; it creates no dependence between $\tilde{x}_{it}$ and $\tilde{x}_{jt}$, which come from disjoint tenders. Contamination of the reference pool therefore bears on Assumption 1 alone, and only insofar as the arrangement with k operates differently inside and outside T_d. When it operates in both, the two sides of the assumption move together and it is unaffected.

Second, Assumption 2 may be stronger than competition requires. Rival bidders respond to one another: a firm bids differently when a strong rival is present, and a common cost shock or a change of deadline moves both members at once. The draw removes this dependence together with any collusion, so wherever Assumption 2 fails the excess includes a competitive component and Λ is biased upward. The within-tender and donor-partner placebos are built to avoid this assumption.

Third, the i.i.d. requirement in Assumption 1 is not implied by equality of marginals, and the difference matters for markers computed from the sequence of co-bids and not from a pooled rate. Margin steadiness and lower-bidder rotation are sensitive to serial dependence and to drift over the sample period, both of which equality of marginals permits.

Table A.3 tests the representativeness half of Assumption 1. Because the draw is uniform over each firm's partner-absent bids, the expected characteristic of the tender supplying a drawn attribute is the leave-one-partner-out pool mean, available in closed form, so the contrast between a pair's own tenders and its draw pool is exact within pair. Across twelve tender-level characteristics the pool resembles the focal tenders closely: all but one differ by less than 0.1 standard deviations, and the largest difference, 0.156, is field size, which the draw conditions on neither directly nor through the tender. Two cautions. The table speaks to the pool's representativeness, not to the i.i.d. requirement above, which equality of means does not deliver. And it weights focal tenders equally: a tender with n bidders generates $\binom{n}{2}$ pairs, so weighting co-bids would be quadratically size-biased in field size on the focal side and only linearly so in the pool, manufacturing a field-size gap on its own.

A.2.2 Within-tender placebo

The tender is held fixed. For each co-bid, two distinct bidder slots of tender t are drawn and assigned to the pair's two roles. Each selected record travels whole, and the outsider statistics are computed relative to the selected source pair.

	Mean over tenders		Difference	
	Pair's own co-bids	Placebo draws	Std. diff.	(s.e.)
Size and competition				
Log tender estimate	13.73	13.81	0.053	(0.001)
Bidders in tender	4.30	4.95	0.156	(0.001)
Financial-stage bidders	2.88	3.07	0.060	(0.001)
Disqualification share	0.089	0.108	0.095	(0.001)
Tender terms				
Envelopes (covers)	1.86	1.86	0.002	(0.001)
Log earnest-money deposit	9.00	9.05	0.018	(0.001)
Log tender fee	5.37	5.39	0.007	(0.001)
Bidding window (days)	12.2	12.3	0.006	(0.001)
Bid validity (days)	108.1	108.2	0.002	(0.001)
Work period (days)	138.0	140.8	0.006	(0.000)
Timing				
Deadline year	2021.6	2021.5	−0.042	(0.001)
Hours to deadline	39.5	38.6	−0.005	(0.000)

Table A.3: Balance of the own-history placebo's draw pool

Note: Each row is one tender-level characteristic. *Pair's own co-bids* is the mean over the pair's actual co-bids. *Placebo draws* is the mean over the tenders that supply the attributes drawn in the own-history placebo. *Std. diff.* is the coefficient from a regression of the characteristic, standardized to mean zero and unit standard deviation on the observed tenders, on an indicator for the placebo draws (placebo minus own). The regression controls for pair fixed effects. Standard errors, clustered on the pair, are in parentheses. N=39.6M co-bids over 2.04M pairs. Every standardized difference, with the exception of the number of bidders in the tender, is below 0.1, i.e., differences between the pairs' own tenders and the placebo draws are small.

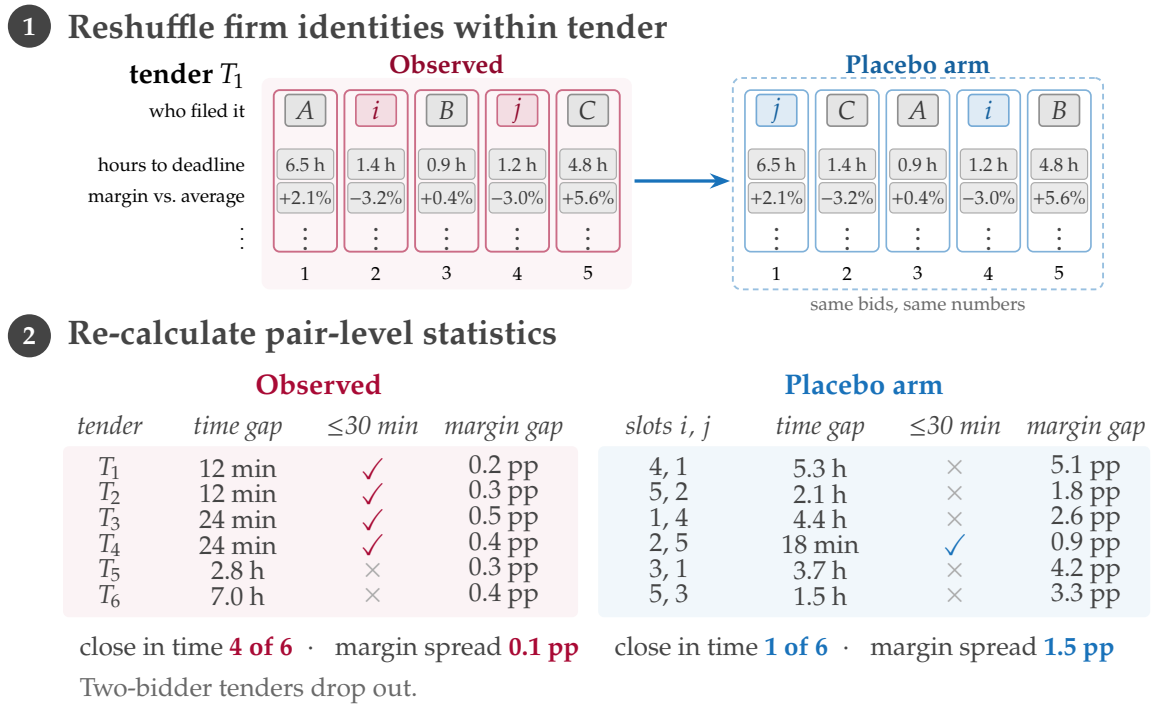


Figure A.2: The within-tender reassignment

Note: The tender is held fixed and two of its bidder slots are drawn and assigned to the focal pair's two roles. Each record travels whole, so the field, the estimate, the deadline and any common shock survive the draw. The identity of the two firms does not.

Assumption 3 (Within-tender exchangeability)**.** Conditional on $C_d = 0$ and on tender t, the joint distribution of the attributes of the focal pair's two slots equals that of any other ordered pair of distinct slots of t.

Proposition 2. *Under Assumption 3, the within-tender draw and the focal pair's record have the same distribution conditional on the tender, and therefore the same distribution unconditionally.*

Proof. Conditional on t, Assumption 3 states that the distribution of the focal pair's slot attributes is invariant to which ordered pair of distinct slots of t is selected. The placebo draws such a pair, so the drawn record and the focal record are identically distributed given t. Averaging over the distribution of tenders in T_d, which the design leaves unchanged, gives the unconditional claim. □

The design conditions on the tender, so the field, the estimate, the deadline and any common cost shock are held fixed and survive the draw. It therefore does not require Assumption 2, and does not charge ordinary strategic interdependence to collusion. The price is that firm identity does not survive: the donor slots belong to other firms, so a pair whose habits are unusual but innocent is not matched against a benchmark built from its own conduct, and Assumption 3 requires the focal pair not to differ systematically from the other bidders present. Contamination also runs the other way here. The two donor slots are taken from one real tender and kept together, so if they collude with each other their joint conduct enters f_0 directly and Λ is biased downward.

A.2.3 Donor-partner placebo

Every co-bid of (i, j) is replaced by a co-bid of i or of j with a different partner in a real tender, with a fresh anchor and a fresh donor partner drawn for each co-bid.

Assumption 4 (Partner exchangeability)**.** Conditional on $C_d = 0$, the distribution of the anchor firm's conduct in a co-bid with j equals its distribution in a co-bid with a partner drawn at random from its other partners.

Proposition 3. *Under Assumption 4, the donor co-bid and the focal co-bid have the same distribution, and the induced pseudo-history has the same distribution as X_d.*

Proof. Assumption 4 is the statement that the identity of the partner does not enter the distribution of the anchor's conduct given $C_d = 0$. Each co-bid of the pseudo-history is therefore distributed as the corresponding focal co-bid, and the draws are made afresh per co-bid, so the two histories agree in distribution. □

As Figure A.3 shows, each donor co-bid actually took place, so both firm identity and the real within-tender relation survive the draw: the deadline, the field and the estimate of the donor tender are genuine. The design therefore needs neither Assumption 2 nor Assumption 3, and breaks only the persistence specific to (i, j). Its weakness lies in the anchor's other partnerships. If the anchor also colludes with some of its donor partners, the placebo histories look more suspicious than independent bidding would, fewer observed pairs stand out against them, and Λ is understated.

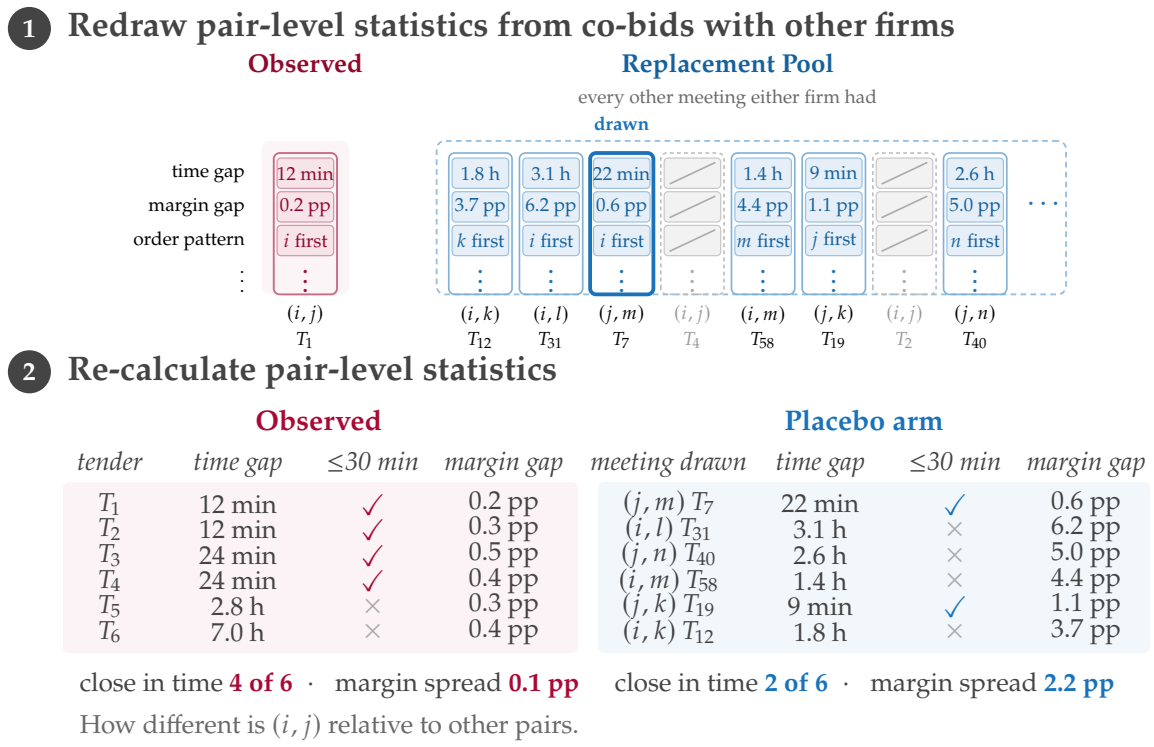


Figure A.3: The donor-partner draw

Note: Each co-bid of the focal pair is replaced by a real co-bid of one of its two members with a different partner, drawn from every other co-bid either firm had, with a fresh anchor and a fresh donor partner for each co-bid. The whole co-bid travels together, so every marker is computed off the same donor tender and the within-tender relation between the two bidders is one that actually occurred.

A.2.4 Assumed-competitive benchmark

The fourth benchmark is not a placebo. It takes each marker's no-collusion distribution from the pair's co-bids in tenders where collusion is implausible: tenders with at least three bids, other than the pair's own, below 85 percent of the estimate, or awarded by recorded lottery among at least seven bidders.

Assumption 5 (Competitive tenders)**.** In the selected tenders every pair bids as it would with $C_d = 0$, and the conduct of a pair with $C_d = 0$ is the same in the selected tenders as in its other co-bids.

The assumption is stated about real tenders rather than constructed by resampling, so it fails in different ways from the three placebos. The selected tenders are those with deep discounts or many bidders, and a non-colluding pair may bid differently there than in its other co-bids. In lottery tenders the bids tie at one price, so the markers that need a lower bidder or a price gap are not computed there.

A.2.5 Relation between the designs

The own-history placebo thus tends to overstate collusion and the other two placebos to understate it. The range of estimates across the three shows how much the results depend on the choice of benchmark.

Table A.4 puts numbers on that spread. It compares the fits under the three placebo designs, all on raw markers, with a fit against an assumed-competitive benchmark, on the pairs all four score. The

Table A.4: Placebo designs compared

	Headline	Within tender	Donor partner	Assumed competitive
Headline	**0.34**	0.62	0.70	0.41
Within tender	0.59	**0.18**	0.90	0.29
Donor partner	0.67	0.74	**0.18**	0.37
Assumed competitive	0.33	0.17	0.33	**0.39**
Fitted type mass by number of markers (%)				
One marker	0	44	39	–
Two markers	22	42	61	–
Three markers	78	14	0	–

Note: The headline, within-tender and donor-partner placebo fits use raw markers with one pooled empirical null per marker. The last column is the assumed-competitive benchmark (tenders with at least three bids below 85 percent of the estimate, or awarded by recorded lottery among at least seven bidders), fitted on atom-split markers with its own empirical null. Top panel, on the 2,027,546 pairs scored by all 4 fits: $\hat{\Lambda}$ on the diagonal, the Spearman correlation of the pair score $\log m$ below it, and above it the top-10% Jaccard index, which leaves out pairs in one fit's top 10% and the other's 10–30% band. Bottom panel: share of the fitted type mass $\sum_S \pi_S$ on types with the given number of markers (not reported for the assumed-competitive fit, whose types are built from atoms).

table shows three things. First, $\hat{\Lambda}$ on the diagonal ranges from 0.18 to 0.34 across the three placebo designs, so the fitted prevalence is not a quantity the data pin down independently of the design. Second, the ranking is far more stable than the level: among the three placebo fits the Spearman correlations are 0.59, 0.67 and 0.74, and the buffered top-decile overlaps 0.62, 0.70 and 0.90, although the diagonal they sit beside varies by a factor of two. Third, within-tender and donor agree with each other more closely than either agrees with the headline placebo, which is the pattern the bracketing argument predicts: both condition on a real tender, and the headline placebo does not.

The assumed-competitive benchmark (Appendix A.2.4) gives a $\hat{\Lambda}$ of 0.39, against 0.34 under the own-history placebo, and ranks pairs differently from all three placebo fits (Spearman correlations 0.17 to 0.33).

A.2.6 Estimates under the within-tender and donor-partner placebos

Table A.5 sets the marker estimates of Table 4 beside those under the within-tender placebo (Appendix A.2.2) and the donor-partner placebo (Appendix A.2.3). Under these designs $\widehat{\Lambda}$ is 18.1% and 18.4%, with standard errors of 0.7 and 1.8 percentage points, clustered by district.

Table A.5: Marker estimates under the three placebo designs

Marker	Headline			Within tender			Donor partner		
	$\hat{\theta}_k$	Prev. (%)	Sep. (%)	$\hat{\theta}_k$	Prev. (%)	Sep. (%)	$\hat{\theta}_k$	Prev. (%)	Sep. (%)
1. Submission proximity	2.71	15.7	27.9	2.69	22.3	38.8	2.51	22.5	45.2
2. Round bid difference	4.68[w]	0.0	0.0	2.16	0.3	0.4	3.06	0.5	1.1
3. Margin steadiness	2.11	7.6	22.5	2.41	2.1	3.7	2.74	1.9	4.2
4. Submission-order role	0.17	17.5	0.3	2.16	8.6	11.6	–	–	–
5. Submission-order stability	–	–	–	–	–	–	–	–	–
6. Targeted disqualification	3.20	2.7	4.9	1.95	6.2	3.7	2.69	3.3	3.7
7. Outsider disqualification	1.78	10.1	6.7	0.60	16.3	1.5	0.44	39.5	4.1
8. Lower-bidder stability	1.45	8.8	2.2	–	–	–	–	–	–
9. Lower-bidder rotation	4.33	6.2	11.3	2.75	4.0	9.4	2.80	9.7	22.2
10. Backlog difference	1.55	5.5	7.7	–	–	–	–	–	–
11. Uncompetitive fringe entrant	2.95	3.4	3.0	–	–	–	0.05[f]	0.0	0.0
12. Exclusive outsider	3.02	22.6	13.7	2.19	40.3	30.8	2.33	22.7	19.4
$\hat{\Lambda}$ (%)	33.6 (1.8)			18.1 (0.7)			18.4 (1.8)		
Allocated prevalence (pp)	27.26			16.12			16.08		
Allocated separation (log LR)	0.7066			0.3480			0.2253		

Note: Pairwise fits under the headline, within-tender and donor-partner placebos (Appendices A.2.1 to A.2.3). The headline columns repeat Table 4. Markers are numbered as in Table 2. The prevalence columns allocate v_Λ in its row-gated form, the separation columns allocate v_L (Appendix A.5). Dashes: the marker does not pass the pre-fit screen of Section 4.3 under that design and is not in the fit. The Shapley shares are in percent of the allocated totals at the foot of the table. [w] Round bid difference (headline) loads 100.0% on directions of the Hessian that the data do not identify. [f] Uncompetitive fringe entrant (donor-partner): $\hat{\theta}_k$ on its lower bound of 0.05, held fixed.

A.2.7 Each pair against its own placebo

The fits above standardize each pair's marker against the pooled placebo draws of pairs with the same number of trials. A pair therefore ranks as suspicious when its markers are high for a pair with its level of activity. An alternative is to compare each pair with its own placebo draws: every placebo design produces five draws of every marker for every observed pair, so each pair has a counterfactual distribution of its own. The pair-centered marker is the observed score less the mean of the pair's own draws, scaled by the within-pair dispersion of the draws pooled across pairs with the same number of trials, and the pattern model of Section 4 is then fit on these pair-centered markers in place of the pooled ones. The marker set, the placebo design and the mixture are unchanged; only the benchmark each pair is held to differs.

Under this benchmark a habit of either firm that the pair's own placebo reproduces does not count against the pair. The cost is power: each pair's null rests on five draws over its own tenders, so a pair with few co-bids cannot be far from its benchmark.

Table A.6 compares the two fits under each placebo design. The pair-centered comparison reorders a substantial share of pairs: the Spearman correlation with the pooled fit of the same design is 0.63 under the headline placebo and 0.49–0.50 under the two tender-based designs, and between 48 and 70 percent of the pair-centered fit's top decile lies in the pooled fit's top three deciles. The ranking moves with the benchmark, but the price premium of the flagged set does not change much: 3.2 against 2.4 percent under the headline placebo, 2.4 against 1.7 under the within-tender placebo, and 2.0 against 2.1 under the donor-partner placebo.

Table A.6: Pooled and pair-centered benchmarks compared

Placebo design	Agreement with the pooled-null fit		Flagged-winner premium	
	Spearman	Top-decile overlap	Pooled-null fit	Pair-centred fit
	(a)	(b)	(c)	(d)
Own history	0.63	0.70	0.0319***	0.0239***
			(0.0023)	(0.0020)
Within tender	0.49	0.55	0.0243***	0.0166***
			(0.0021)	(0.0021)
Donor partner	0.50	0.48	0.0201***	0.0209***
			(0.0020)	(0.0020)
N (pairs, columns a–b)	2,027,514			
N (tenders, columns c–d)			4,594,861	

Note: Each row compares the pattern model fit on pair-centred markers with the same model fit on the pooled markers of the same placebo design under one pooled empirical null per marker (raw markers, twelve markers, territorial exclusivity excluded). Columns (a) and (b), on the pairs scored by both fits: (a) the Spearman correlation of the two pair scores, (b) the share of the pair-centred fit's top 10% inside the pooled-null fit's top 30%. Columns (c) and (d): the coefficient on the flagged-winner indicator in column 1 of Table 6, flagging the top decile of the pair score among pairs with at least five financial co-bids under the pooled-null fit (c) and the pair-centred fit (d). Standard errors clustered by district in parentheses. ***, **, *: significant at the 1, 5 and 10 percent level.

A.3 Empirical-null normalization

When a marker's observed distribution is shifted relative to its placebo distribution as a whole, I recenter and rescale the observed scores before the fit (Section 4.3). This *central matching* assumes that the placebo misstates the null only in location and scale, and that collusion is rare enough for the center of the observed distribution to consist almost entirely of independent pairs. If collusion is common, part of the signal is absorbed by the centering, and the corrected markers understate it. I center a marker when the median of its observed scores lies more than 0.15 standard deviations from the placebo median, or when the one-marker fit described below finds a shift of less than 1.2 standard deviations on the uncorrected scores, which indicates that the whole distribution has moved and not only its upper tail. For marker k, let $\widehat{\mu}_{0,k}$ be the median of the observed scores from (3), pooled over all co-bid bins, and $\widehat{\sigma}_{0,k} = [\widehat{\mu}_{0,k} - Q_{0.25}]/0.67449$ their scale, estimated from the lower half because collusion is expected to raise the markers. Each observed y_{dk} is replaced by $(y_{dk} - \widehat{\mu}_{0,k})/\widehat{\sigma}_{0,k}$; the placebo scores are unchanged. Efron's quadratic central matching estimates the same two parameters from the curvature of the log density near the center (Efron 2007).

A.4 Restrictions on the types of colluding pairs

The headline model restricts the type shares π_S through the log-linear form in (5). Table A.7 tests three further restrictions against it, each estimated on the same observed and placebo rows. Setting every interaction $\beta_{kj} = 0$ makes colluding pairs adopt each marker independently of the others. Letting every $\beta_{kj} \to -\infty$ removes all types with more than one marker, so each colluding pair is raised on a single marker. Letting every $\alpha_k \to +\infty$, with $\beta_{kj} = 0$, puts all type mass on the type that contains every marker, so each colluding pair is raised on all markers, each by its own θ_k. The data reject all three restrictions.[23] The estimated share of colluding pairs ranges from 10.6 percent when

[23] Allowing types of up to four markers puts 4 percent of the type mass on four-marker types in the headline fit, and none under the two alternative placebo designs. The estimated share of colluding pairs is 35.4 percent, against 33.6 percent in the headline model (17.6 and 18.4 percent for the within-tender and donor-partner placebos).

every pair must use all markers to 43.8 percent when each pair may use only one, against 33.6 percent in the headline model.

Table A.7: Restrictions on the types of colluding pairs

	$\hat{\Lambda}$ (%)	Parameters	LR statistic	Restrictions	p-value
Up to three markers per type (headline)	33.6	78	–	–	–
Independent markers	27.0	23	47,956	55	<0.001
One marker per type	43.8	23	135,856	55	<0.001
All markers in every type	10.6	12	235,225	66	<0.001

Note: Headline fit, 2,043,939 scored pairs. Every model is fitted on the same rows. Independent markers: no pairwise terms. One marker per type: each colluding pair is elevated on one marker. All markers in every type: each colluding pair is elevated on all markers. LR statistic: twice the difference in log-likelihood from the headline model, with as many restrictions as parameters dropped. The p-values use the χ^2 distribution, which is approximate for the last two models: their parameters are fixed at the edge of their range.

A.5 Shapley allocation of prevalence and separation

Section 4.6 allocates two quantities across markers. Both are set up as cooperative games: a game assigns a value $v(A)$ to every subset A of the markers, the value the model would deliver with only those markers in hand, and the Shapley value ϕ_k (Shapley 1953) gives marker k the average of $v(A \cup \{k\}) - v(A)$ over all orders in which the markers could be assembled. The allocation is exact: the ϕ_k sum to the value of the full set.

Restricting the model to a subset A requires no refitting. Integrating out the markers outside A collapses each type S to $S \cap A$, because each Gaussian tilt has expectation one under the null, so a marker that is integrated out contributes a factor of one.

The first game allocates fitted prevalence. A type still flags a pair as long as at least one of its markers survives the restriction, so

$$v_\Lambda(A) = \Lambda \Pr(S \cap A \neq \emptyset), \qquad \phi_k^\Lambda = \Lambda \sum_{S \ni k} \frac{\pi_S}{|S|}. \tag{14}$$

The closed form says that each type splits its share of Λ evenly among the markers it contains.

The second game allocates separation. Writing L_A for the likelihood ratio of the model restricted to A,

$$v_L(A) = \mathbb{E}_{\text{obs}}[\log L_A(Y_A)] - \mathbb{E}_{\text{pl}}[\log L_A(Y_A)], \tag{15}$$

the gap between the observed and placebo arms in mean log likelihood ratio. The parameters are held at their full-model values throughout, so the allocation divides the fitted model's separation; it is not a comparison of refitted models.

Markers are not available for every pair, so the reported prevalence game is row-gated: a type contributes for pair d only when $S \cap A \cap G_d$ is nonempty, where G_d is the set of markers available for that pair, and the Shapley values are averaged over observed pairs.

A.6 Alternative decision rules

Section 4.7 ranks pairs by the fitted likelihood ratio and flags the top decile. The decile is a reporting choice, and Table A.8 sets it beside four alternatives: the top $\widehat{\Lambda}$ of pairs, a posterior above one half, a one-sided p-value cutoff, and a Benjamini–Hochberg cutoff at $q = 0.10$. The rules differ in how many

pairs they flag and in the false discovery rate the fitted model implies for them, and the paper keeps the two uses apart: the regressions use the fixed-share rule, which holds the treated group's size constant across specifications, while descriptive aggregates that ask how much activity is associated with the fitted colluding component use the top $\widehat{\Lambda}$.

Table A.8: Alternative decision rules

Rule	Share flagged	Model-implied FDR
Top $\widehat{\Lambda}$ of pairs	33.6%	37.5%
Posterior $> 1/2$	20.3%	21.6%
$p \leq 0.05$	19.1%	20.1%
Top decile (main regression rule)	10.0%	7.0%
Benjamini–Hochberg, $q = 0.10$	9.8%	6.8%

Note: Shares are among the 2,043,939 scored pairs. The model-implied false discovery rate of a rule is the mean over the pairs it flags of the posterior probability that the pair does not collude, $1 - \hat{\Lambda}\hat{L}_d/(\hat{\Lambda}\hat{L}_d + 1 - \hat{\Lambda})$, with $\hat{L}_d$ the pair's fitted likelihood ratio.

B Validation

B.1 Excess bid correlation

Figure B.1 shows the distributions behind Table 5: excess bid correlation by arm in Panel A, and the share of pairs rejecting each null by score ventile in Panel B. Table B.1 re-estimates Table 5 on the flagged pairs whose matched control pair, like the flagged pair itself, enters every statistic.

This is in the spirit of Conley and Decarolis (2016), who compare suspect bidder groups with comparable groups, and Bergman et al. (2020), who compare a cartel's bids with the same firms' bids after the cartel ended.

Because each test is run on tens of thousands of pairs, we report the share of pairs rejected at a false discovery rate q of 10 percent, using the Benjamini-Hochberg procedure to account for multiple hypothesis testing (Benjamini and Hochberg 1995).

Margin steadiness and round bid difference are markers that are themselves regularities of the two members' bids, and the stability of the gap between the two bids is close to what the independence test measures. Controlling for both at the pair level, on the 54 percent of Panel A pairs for which both are observed, lowers the flagged-minus-control difference in excess correlation from 0.38 to 0.22 and in the rate of rejecting conditional independence from 32 to 19 percentage points; the exchangeability differences change by less than one percentage point.

Additional evidence from the tests is provided in Figure B.1: Panel A shows that the distribution of excess correlation of flagged pairs first-order stochastically dominates that of matched control pairs. Panel B plots the share rejecting each null by the pair's percentile of the likelihood ratio of Section 4.7. For both tests, the share rejecting the two conditions of independent bidding rises monotonically with the likelihood ratio, from 8 to 55 percent for conditional independence and from 8 to 29 percent for exchangeability.

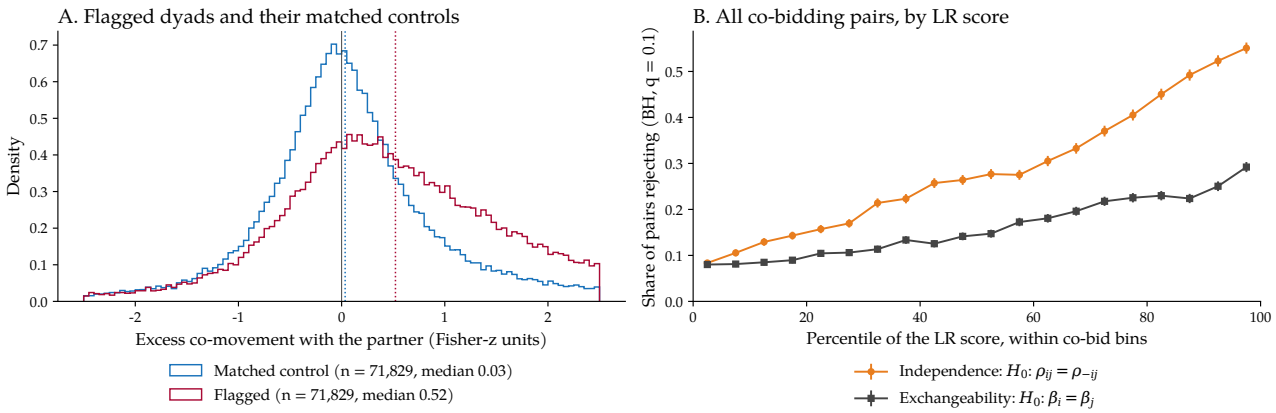


Figure B.1: Excess bid correlation by arm, and rejection shares by score

Note: Panel A: density of excess$_{ij}$ in Fisher-z units, by arm, with arm medians dotted. The arms are the flagged pairs (the top decile of scored pairs by the estimated likelihood ratio) and their matched control pairs.
Panel B: share of pairs rejecting each null, in ventiles of the score of Section 4.7 (the log likelihood ratio), taken within co-bid-count bins (5–8, 9–15, 16–30, 31–60, 61–120 and above 120 financial co-bids). Sample: 130,595 co-bidding pairs covered by the score with at least twenty co-bids and at least ten co-bids with at least three bidders. Orange: share rejected by the conditional independence test of Table 5, Panel A (H_0: $\rho_{ij} = \rho_{-ij}$, one-sided). Gray: share rejected by the exchangeability F-test of Table 5, Panel B, on all six shifters (H_0: $\beta_i = \beta_j$). Rejection is by Benjamini–Hochberg at $q = 0.10$ over all pairs, separately for each test.

Table B.1: Residual co-movement and exchangeability, constant sample of flagged and matched control pairs

	Flagged	Control	Δ	(s.e.)
Panel A. Conditional independence: H_0: $\rho_{ij} = \rho_{-ij}$ (26,947 matched pairs)				
Pair correlation (ρ_{ij})	0.03	−0.19	0.21	(0.005)
Member–field correlation (ρ_{-ij})	−0.36	−0.26	−0.10	(0.002)
Excess ($\rho_{ij} - \rho_{-ij}$)	**0.39**	**0.08**	**0.31**	(0.006)
Share rejecting H_0	0.52	0.25	0.27	(0.004)
Panel B. Exchangeability: H_0: $\beta_i = \beta_j$ (26,947 matched pairs)				
Share rejecting H_0, own shifters	0.15	0.08	0.07	(0.003)
Share rejecting H_0, all six shifters	0.28	0.16	0.12	(0.004)

Note: As Table 5, restricted to the flagged pairs for which both the flagged pair and its matched control pair enter every statistic in the table: the independence test (at least ten co-bids in fields of three or more bidders) and both exchangeability tests (at least twenty co-bids and a defined F-statistic). Keeping both members of each match preserves the one-to-one comparison of the arms.

B.2 Shared surnames

Table B.2: Shared surnames and the pair's likelihood ratio

	Dependent variable (percentage points)					
	Same surname			Same rare surname		
	(1)	(2)	(3)	(4)	(5)	(6)
Likelihood ratio of the pair (s.d.)	0.40***	0.50***	0.69***	0.21***	0.28***	0.70***
	(0.05)	(0.05)	(0.06)	(0.04)	(0.03)	(0.08)
Log co-bids control	✓	✓	✓	✓	✓	✓
District FE	✓			✓		
Firm FE		✓	✓		✓	✓
Firms with a namesake partner			✓			✓
Mean of dependent variable	4.73	4.73	6.14	1.36	1.36	3.72
Pairs	916,044	916,044	705,751	916,044	916,044	334,730
Clustering	District	District	District	District	District	District

Note: OLS, one observation per scored pair in which both firms are named after a person (Section 5.2). The outcome is 100 if the two firms have the same surname (columns 1–3) or the same rare surname, one carried by fewer than 1 percent of the state's scored firms (columns 4–6). The likelihood ratio of the pair is standardized within each sample. Every column controls for the log number of joint bids. Firm FE: a fixed effect for each of the pair's two firms. Firms with a namesake partner: pairs in which at least one firm has a scored partner with the same surname (column 3) or the same rare surname (column 6). Standard errors, clustered by district, in parentheses. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

B.3 Joint entry

I build a pair-tender panel of every tender on which at least one member of the pair bid, within the organization × product cells in which the pair has co-bid at least once. With tenders appearing across pairs, this yields about 68 million pair-tender observations. On this panel, Table B.3 reports unadjusted joint-entry rates by arm. Panel B replaces the realized outside field with the outside field that can be predicted from tender observables (for more details, see Appendix F.6).

To account for pair-level heterogeneity, I estimate a regression of joint entry of pair ij in tender t on

the number of outside bidders:

$$\text{JointEntry}_{ij,t} = \sum_{k=0}^{2} \beta_k B_{t,k} + \sum_{k=0}^{2} \gamma_k B_{t,k} \cdot D_{ij} + X'_{ij,t}\delta + X'_{ij,t}\eta \cdot D_{ij} + \alpha_{ij} + \varepsilon_{ij,t}$$

where $B_{t,k}$, for $k \in \{0, 1, 2\}$, is an indicator for the tender having k outside bidders not belonging to pair ij, with three or more outside bidders the omitted category; D_{ij} is an indicator for flagged pairs so that γ_k measures how differently flagged pairs respond to the outside field than matched control pairs; $X_{ij,t}$ is a vector of additional tender and pair characteristics, and α_{ij} is a pair fixed effect. Each pair is weighted by the inverse of the number of tenders it is at risk in, so that pairs exposed to many tenders do not dominate the fit. Because the same tender recurs across many different pairs, standard errors are clustered by pair and by tender (Cameron, Gelbach, and D. L. Miller 2011).

Unadjusted, joint entry is similar across arms with two outside bidders, but with no or only one outside bidder flagged pairs enter together 14 and 21 percentage points more often, respectively – consistent with cover bidding. The regression shows this gap survives controlling for the rest of the tender: relative to three or more outside bidders, matched control pairs' joint entry rises by 26 percentage points when the field holds no outside bidders, and flagged pairs' rises by a further 32 points on top of that (14 with one outside bidder).

The coefficients are small because of measurement error due to the fact that the model only predicts the number of bidders slightly better than a model that predicts the modal number of bidders in the district $\times$ organization cell.

Finally, flagged pairs respond differently to other tender characteristics than matched control pairs do. Flagged pairs' joint entry is less than half as responsive to the tender value as joint entry of matched control pairs ($+0.57$ compared to $+0.78$): a genuine competitor has more to gain from a larger prize, while a colluding pair's second bid depends less on what is at stake. A higher earnest money deposit lowers joint entry in both arms, more for flagged pairs (-0.44 against -0.30 percentage points; the difference, -0.14, is significant at the 5 percent level). Flagged pairs react more strongly to their backlog than matched control pairs do ($+0.28$ against $+0.15$).

Table B.3: Joint entry and the competitive field

	Joint-entry rate			Within-pair regression		
Outside bidders	Flagged	Control	Difference	Flagged	Control	Difference
Panel A. Realized outside field						
0	42.22	27.97	+14.24	+58.57***	+26.14***	+32.43***
				(0.36)	(0.35)	(0.49)
1	66.65	45.91	+20.74	+45.14***	+31.36***	+13.79***
				(0.15)	(0.14)	(0.20)
2	10.25	10.75	-0.50	-8.42***	-7.19***	-1.23***
				(0.08)	(0.06)	(0.09)
3 or more	20.63	16.53	+4.11	*omitted category*		
Panel B. Predicted outside field						
0	20.64	14.23	+6.40	-2.33***	-3.94***	+1.61***
				(0.31)	(0.16)	(0.32)
1	24.78	19.87	+4.91	-1.28***	-1.75***	+0.47*
				(0.28)	(0.12)	(0.27)
2	20.64	16.85	+3.79	-1.74***	-0.92***	-0.82***
				(0.29)	(0.11)	(0.28)
3 or more	28.15	17.15	+10.99	*omitted category*		
Control coefficients, Panel B regression, estimated separately by arm						
Log tender value				+0.57***	+0.78***	-0.21**
				(0.08)	(0.03)	(0.09)
Cover envelopes				+5.65***	+0.50***	+5.15***
				(0.43)	(0.17)	(0.44)
Log EMD amount				-0.44***	-0.30***	-0.14**
				(0.07)	(0.02)	(0.07)
Backlog, mean				+0.28***	+0.15***	+0.13***
				(0.01)	(0.01)	(0.02)
Observations	68,343,568					
Pairs, flagged / control	190,271 / 192,462					

Note: Pair-tender panel. The sample collects, for every pair, every tender on which at least one member of the pair bid, within the organization × product cells where the pair has co-bid. The outcome is that both members bid. Flagged pairs are the top decile of scored pairs by the estimated likelihood ratio (Section 4). Each is matched to an unflagged pair in the same district with a co-bid count within 30 percent of its own. Left columns: joint-entry rates by arm, conditional on at least one pair member entering. Right columns: Linear probability model, with joint entry as the outcome. Regressions control for pair fixed effects. The Panel B regression includes the controls whose coefficients are reported in the lower block of the table. Cover envelopes is the number of technical documents needed to submit for technical bidding. EMD amount is the earnest money deposit, a security deposit for the bid. Backlog is the pair's mean log value of contracts won in the past 360 days, z-scored within firm. Panel A counts outside bidders from the bid records. It depends on the outcome and is shown for comparison only. Panel B uses the predicted field size net of the pair's two members, $\lfloor E[N \mid X] - 2 \rfloor$. Three or more outside bidders is the omitted category. Pairs weighted equally. Standard errors clustered by pair and tender. ***, **, and * denote significance at the 1, 5, and 10 percent levels.

B.4 Win rates by field size

The test compares win rates in tenders in which exactly one flagged pair has both members bidding (56.5 percent of all tenders with any pair active), with a disclosed winner, at least three bidders, and a single awarded bidder. Figure B.2 plots win rates by the number of bidders, per bidder in Panel (a) and for the pair in Panel (b).

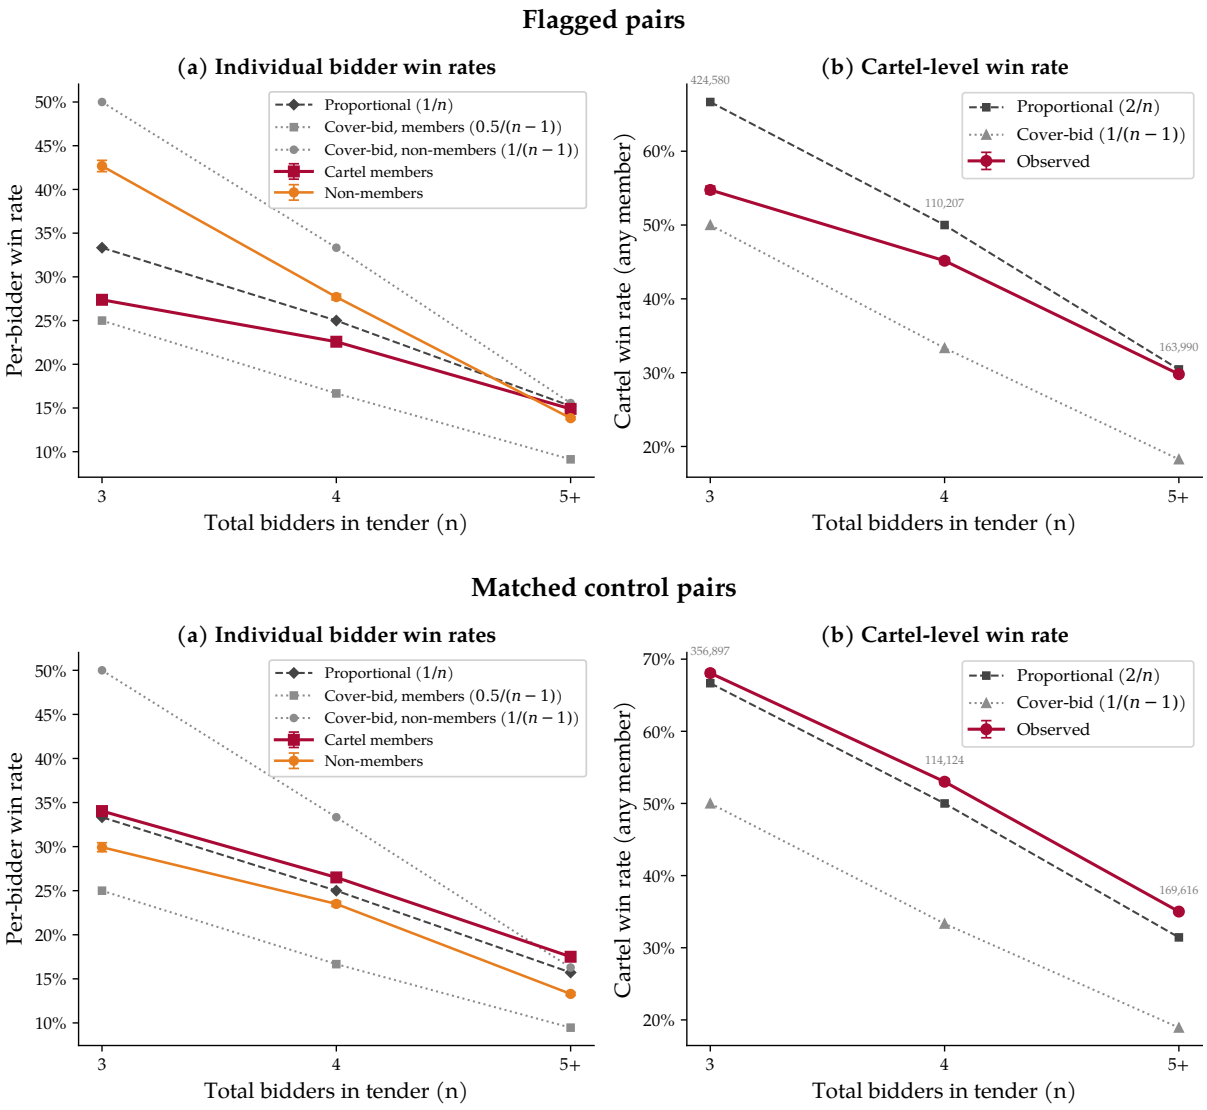


Figure B.2: Win rates in tenders with exactly one flagged pair present

Note: Tenders with at least three bidders, one disclosed winner, and exactly one flagged pair with both members bidding (698,777 tenders). Top row: flagged pairs, the top decile of scored pairs by the estimated likelihood ratio (Section 4). Bottom row: the same test on their matched control pairs. Members are the two firms of that pair. Non-members are firms of other flagged pairs whose partner is absent. Panel (a): mean win rate per bidder by field size. The dashed line is the proportional benchmark $1/n$, and the dotted lines are the cover-bid predictions, under which each member wins with probability $0.5/(n-1)$ and every other bidder with $1/(n-1)$. Panel (b): the pair's win rate against $2/n$ (proportional) and $1/(n-1)$ (cover bid). Tender counts shown above the Panel (b) markers.

B.5 Cover bidding within firm

Flagged pairs could win less often in thin fields because their members are weak bidders rather than because one member covers for the other. The check compares each firm with itself. Let $y_{it} = 100\,(\text{win}_{it} - 1/n_t)$, let F_i indicate that firm i belongs to a flagged rather than a matched control pair, and let P_{it} indicate that i's partner also bid in tender t:

$$y_{it} \;=\; \delta^F P_{it} F_i \;+\; \delta^C P_{it}(1 - F_i) \;+\; \mu_i \;+\; \varepsilon_{it}, \tag{16}$$

where μ_i is a firm $\times$ arm fixed effect, so δ^F compares a flagged firm in tenders where its partner bid with the same firm in tenders where the partner did not. When the partner bids, a flagged firm's win rate changes by -7.75 percentage points at $n = 3$ (s.e. 0.26), -3.72 at $n = 4$, and -0.97 at $n \geq 5$ (Table B.4). A control firm's win rate rises by 3.63, 3.02 and 2.01 points. The decline for flagged firms is largest where the field is thinnest, which is where a cover bid changes the allocation most. Field size is not assigned, since it partly reflects the pair's own entry decision, so these are conditional correlations.

Table B.4: Cover bidding at the bidder level, within firm

	$n = 3$	$n = 4$	$n \geq 5$
	(1)	(2)	(3)
Flagged pair member	-7.75***	-3.72***	-0.97***
	(0.26)	(0.26)	(0.13)
Matched control pair member	3.63***	3.02***	2.01***
	(0.27)	(0.23)	(0.12)
p(flagged = control)	<0.001	<0.001	<0.001
Firms identified (%)	48.0	40.0	52.6
Firm $\times$ arm fixed effects	✓	✓	✓
Observations	2,115,084	707,804	1,524,409

Note: Estimates of Equation (16), one column per field size. The unit of observation is a (bidder, tender) cell in the sample of Figure B.2, and the outcome is the bidder's win indicator minus $1/n$, in percentage points. Each coefficient is the effect of the partner also bidding, within firm. "Firms identified" is the percent of flagged firms observed with the partner both present and absent inside the bin. The matched control pairs are those of Table B.3. 25.85 percent of control-arm tenders also host a flagged pair. Standard errors, clustered two-way on the tender and the bidder, in parentheses. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

B.6 Behavior of flagged pairs in deadline extensions

I take extended tenders in which one member of a flagged pair or of a matched control pair had bid by the original deadline, and measure how much the lowest bid falls between the original deadline and the final closing, depending on who enters late (Figure B.3). A late bid above the lowest bid by the original deadline leaves the price unchanged.

When an outsider enters late, the lowest bid falls by about the same amount in both arms: 2.5 percent for flagged pairs and 2.3 percent for matched control pairs. When the absent partner enters instead, the arms diverge. For matched control pairs the lowest bid falls by 2.2 percent, as much as when an outsider enters; for flagged pairs it falls by only 0.9 percent. This is consistent with partners in flagged pairs entering to supply additional bids to meet the minimum bid requirements, and not to win the contract. Rebidding gives a similar test in other settings. When all bids in a Japanese auction miss the secret reserve, the same firms bid again, and the second-lowest bidder rarely undercuts the

lowest one (Kawai and Nakabayashi 2022); in Ukrainian multistage auctions, colluding pairs do not revise their bids against each other (Baránek, Musolff, and Titl 2026).

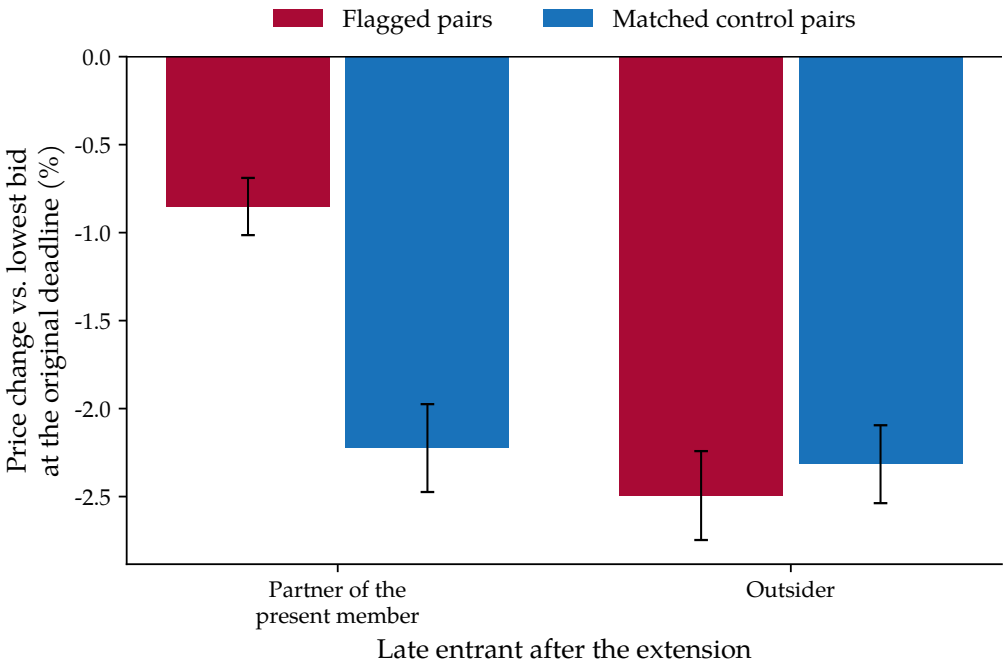


Figure B.3: Change in the lowest bid after a deadline extension, by who entered late

Note: Extended tenders in which a member of a flagged pair or of a matched control pair had bid by the original deadline. Each bar is the mean change in the lowest bid between the original deadline and the close, $100 \times \log$(lowest bid at the close / lowest bid at the original deadline). Negative values mean the price fell. A late bid above the lowest bid by the original deadline is a zero price change. Partner: one member of the pair had bid by the original deadline and only its partner entered after it. Outsider: one or both members had bid by the original deadline and only firms belonging to no pair present in the tender entered after it. Tenders: 14,031 and 78,566 (flagged, partner and outsider), 17,172 and 72,969 (control). 95 percent confidence intervals, clustered by district. Flagged pairs are the top decile of scored pairs by the estimated likelihood ratio (Section 4).

C The footprint of collusion

This appendix collects the exhibits behind Section 6: the share of each district's tenders on which any pair can be scored (Figure C.1), the decomposition of the variation in collusive tenders across places and organizations (Table C.1), the organization characteristics and the conduct of flagged pairs across organizations (Table C.2), and the relation of collusive tenders to the number of firms in a market (Table C.3).

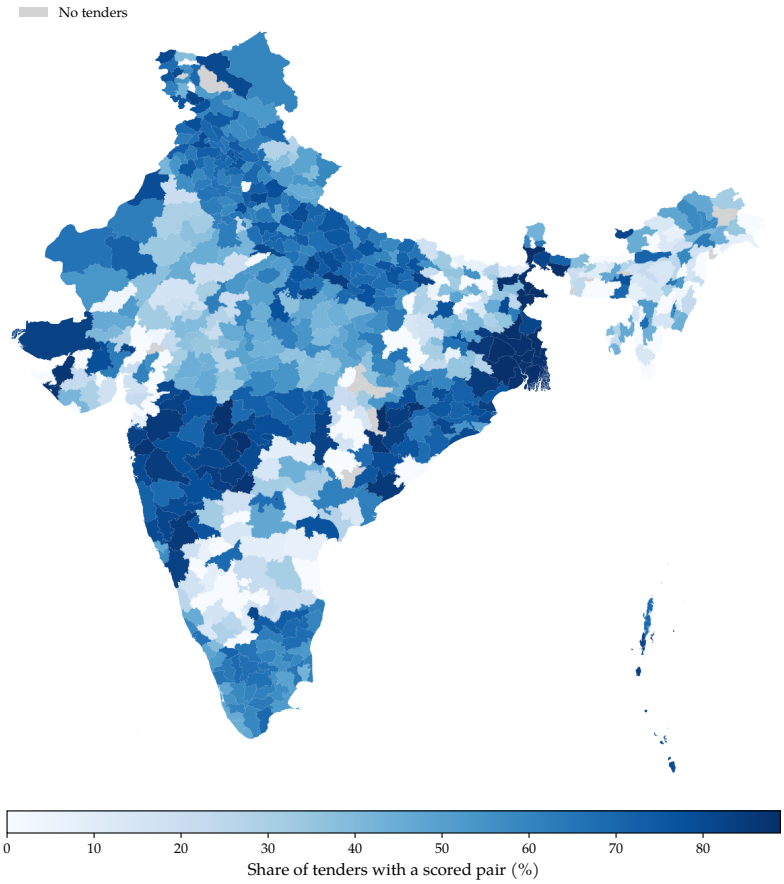


Figure C.1: Measurability: share of tenders with at least one scored pair, by district

Note: Share of a district's tenders with at least one scored pair, the denominator of the footprint in Figure 5(b). Districts in gray have no tenders. Scored pairs are defined in Section 4.

Table C.1: Variation in the prevalence of collusion across places and organizations

	Flagged	Control	Excess
A. Shapley decomposition of R^2			
District	0.27	0.30	0.11
Organization	0.43	0.35	0.43
Product	0.01	0.01	0.00
Year	0.03	0.00	0.02
All four (R^2)	0.74	0.67	0.57
Mean (pp)	26.2	18.4	7.8
Market-years		109,039	
Tenders		4,735,678	

Note: Unit of observation: market × year, where a market is a district × procuring organization × product, with at least five tenders in the year. Outcomes, in percentage points: the share of the market's tenders won by a member of a flagged pair whose partner also bids (Flagged), the same for matched control pairs (Control), and their difference (Excess). Shapley decomposition of the tender-weighted R^2 of district, organization, product and year fixed effects. Each entry is the group's marginal R^2 averaged over all subsets of the other groups, and the four entries sum to the R^2 of all four. Flagged pairs are the top decile of scored pairs by the estimated likelihood ratio (Section 4). Each flagged pair's matched control is a pair from the same district with a co-bid count within a factor of 1.3.

Table C.2: Organization characteristics, the share of collusive tenders, and the conduct of flagged pairs

	Share of collusive tenders (pp) (1)	Bids within 30 minutes (pp) (2)	Marker composite (SD) (3)
Organization type (vs. state department)			
Rural local body	16.3***	5.41***	0.020
	(1.4)	(1.03)	(0.026)
Urban local body	3.5**	2.02**	−0.050**
	(1.5)	(0.83)	(0.022)
Mixed local body	5.6***	−3.13***	−0.018
	(1.6)	(1.16)	(0.042)
State board / PSU	−0.2	0.43	0.029
	(1.2)	(1.27)	(0.043)
Central PSU	−9.3***	0.06	0.084
	(1.8)	(1.66)	(0.054)
Central government	−9.0***	3.37	0.164*
	(1.9)	(2.75)	(0.096)
Organization rules (share of tenders)			
Two-bid rule	1.3	4.82	0.201*
	(2.0)	(3.58)	(0.120)
Three-bid rule	11.5***	2.79	0.041
	(2.0)	(2.78)	(0.107)
Lottery in case of a tie	−1.3	−5.94***	−0.266***
	(3.2)	(1.75)	(0.100)
Strict bid floor	−0.1	3.52**	−0.055
	(1.9)	(1.39)	(0.072)
Organization behavior			
Cancels when lowest bid > estimate (per SD)	−0.8	0.91***	0.081***
	(0.6)	(0.32)	(0.020)
Market conditions			
Distinct firms (log, per SD)	−11.3***	2.84***	0.178***
	(1.5)	(0.74)	(0.051)
Tenders (log, per SD)	11.9***	−2.35***	−0.101**
	(1.4)	(0.62)	(0.050)
Bidders per tender (per SD)	1.6***	0.08***	0.013***
	(0.6)	(0.03)	(0.002)
Mean log contract value (per SD)	−1.0**	−0.02	0.054**
	(0.4)	(0.32)	(0.023)
Share of works (0 to 1)	0.6	−4.11	−0.143*
	(1.6)	(2.65)	(0.077)
Unit	Market-year	Co-bid	Co-bid
Fixed effects	State × year	Pair, year	Pair, year
Mean of outcome	27.6	15.1	0
Observations	26,900	7,378,631	7,395,088
Within R^2	0.314	0.001	0.002
Organization-level R^2	0.75		
10-fold cross-validated R^2	0.55		

Note: Column 1 is the regression of Figure 6: the share of a market-year's measurable tenders that are collusive, in percentage points, for district × organization × year cells with at least five measurable tenders, with state × year fixed effects and the share of measurable tenders as a control, weighted by the number of tenders. Columns 2 and 3 use every co-bid of a flagged pair (a tender in which both members bid) in the same market-years, with the characteristics of the co-bid's district × organization × year cell, and pair and year fixed effects. The outcome in column 2 is 100 if the two bids were submitted within 30 minutes of each other. The outcome in column 3 is a composite of the tender-level markers: bids within 30 minutes, exactly one member technically disqualified (net of the probability implied by each firm's own disqualification rate and the tender's strictness), a fringe outsider present, a captive outsider present, a round bid difference, and the share of outsiders disqualified. Each marker is standardized and weighted by its Shapley share of separation in Table 4, renormalized to sum to one over these markers and over the markers observed in the co-bid. The composite is standardized to mean zero and standard deviation one. Organization types relative to state departments. Continuous regressors per standard deviation. Rule variables are shares of tenders from 0 to 1. Missing values set to zero with an indicator. The organization-level R^2 is the share of the variance across organizations, in the outcome net of the fixed effects and the control, that the characteristics explain. The cross-validated R^2 predicts each market-year from a regression estimated on the other nine of ten random folds. Standard errors clustered by district in parentheses. ***, **, and * denote significance at the 1%, 5%, and 10% levels, respectively.

Table C.3: The number of firms and the share of tenders won by flagged pairs

	All pairs			Pairs by number of co-bids			
	Main	Within market	Lagged firms	5–9	10–19	20–49	50+
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Panel A. Flagged pairs							
Log distinct firms	−9.31***	−7.45***	−2.04***	0.67***	−0.28	−2.56***	−7.71***
	(0.58)	(0.66)	(0.29)	(0.20)	(0.22)	(0.31)	(0.68)
Mean of outcome (percent)	26.2	26.2	27.0	4.4	6.0	8.7	10.9
Panel B. Matched control pairs							
Log distinct firms	−2.44***	−1.54***	−0.19	1.32***	0.41*	−1.06***	−2.94***
	(0.71)	(0.41)	(0.26)	(0.15)	(0.23)	(0.36)	(0.49)
Mean of outcome (percent)	18.4	18.4	18.9	3.9	5.0	6.6	5.7
Panel C. Difference (flagged minus control)							
Log distinct firms	−6.87***	−5.90***		−0.66***	−0.68***	−1.51***	−4.77***
	(0.86)	(0.83)		(0.20)	(0.26)	(0.42)	(0.74)
Mean of outcome (percent)	7.8	7.8		0.5	1.0	2.1	5.3
Observations (market-years)	106,196	99,794	88,819	106,196	106,196	106,196	106,196
District, organization, product × year	✓		✓	✓	✓	✓	✓
Market, year		✓					

Note: Unit of observation: market × year, where a market is a district × procuring organization × product with at least five tenders in the year. Outcome, in percentage points: the share of the market's tenders won by a member of a flagged pair whose partner also bids (Panel A), the same for matched control pairs (Panel B), and their difference (Panel C). Each column regresses the outcome on the log number of distinct firms that bid in the market-year (column 3: in the previous year) and the log number of tenders, weighted by the number of tenders. Column 1 has district × year, organization × year and product × year fixed effects, column 2 market and year fixed effects. Columns 4–7 repeat column 1 for pairs with 5–9, 10–19, 20–49 and 50 or more co-bids. Standard errors clustered by district in parentheses. ***, **, and * denote significance at the 1%, 5%, and 10% levels.

D Price effects

D.1 Size of the premium relative to the dispersion of prices

Table D.1 divides the coefficients of Table 6 by the standard deviation of log(contract value/ government estimate) in each column's estimation sample. The main rows use the standard deviation of the outcome after residualizing it on the column's fixed effects and controls, which is the variation that identifies the coefficient. The bracketed rows use the raw outcome. The column-1 premium equals 0.21 standard deviations.

Table D.1: Price effects of collusion, rescaled by the dispersion of contract prices

	(1)	(2)	(3)	(4)	(5)
Collusive tender	0.207***	0.208***	0.253***	0.288***	
	(0.015)	(0.013)	(0.019)	(0.028)	
Max winner pair score (SD)					0.120***
					(0.009)
Relative to the raw SD	[0.166]	[0.165]	[0.196]	[0.219]	[0.092]
SD, within market	0.154	0.152	0.131	0.115	0.143
SD, raw	0.192	0.192	0.168	0.151	0.186
Coefficient, log points	0.0319	0.0316	0.0331	0.0331	0.0171
N	4,594,861	4,594,861	2,377,278	1,660,170	3,217,813
Share treated	26.8%	26.8%	51.9%	74.3%	
District × Org × Year, Product × Year FE	✓	✓	✓	✓	✓
Tender-level controls	✓	✓	✓	✓	✓
N bidders control		✓	✓	✓	✓
Competitive environment control					✓
Sample	Full	Full	Any flagged winner	Flagged pair present	Scored subsample
Clustering	District	District	District	District	District

Note: As Table 6, with each coefficient and standard error divided by the standard deviation of the dependent variable in that column's estimation sample. Unbracketed rows use the residual of the dependent variable on the column's fixed effects and controls, excluding the treatment variable, with the standard deviation adjusted for degrees of freedom. Bracketed rows use the dependent variable itself. Stars are those of Table 6, since rescaling leaves the *t*-statistics unchanged. The coefficient in log points is repeated from Table 6.

D.2 Alternative samples, fits, and bidder identifiers

Figure D.1 re-estimates the five columns of Table 6 under fifteen variations in three groups. The first group changes the sample or the flag. Pairs are scored on data through March 2022 and the premium is estimated on later tenders, so the flag cannot reflect the prices it is used to explain; the cutoff rises from the 90th to the 96th percentile; tenders are weighted by the log of the government estimate; contracts that settle at a round multiple of five percent of the estimate are dropped; tenders whose winner belongs to no scored pair are dropped; the flag is restricted to pairs with fewer than nine, or with nine or more, of the twelve markers observed; and the top decile is taken separately within these two coverage groups. The second group replaces the headline fit with another fit of the same model: on the within-tender placebo, on the donor-partner placebo, on the headline placebo with each pair benchmarked against its own draws (Appendix A.2.7), with the markers entering independently (the

product fit), and with each marker's no-collusion distribution taken from tenders in which at least three other bids fall below 85 percent of the estimate rather than from a placebo. The third group defines bidders at the national rather than the district level, once keeping generic firm names split by district and once pooling them. All 75 coefficients are positive and significant at the 5 percent level. The column-1 estimate ranges from 1.99 percent (pairs with fewer than 9 of 12 markers observed) to 3.61 percent (product fit), against 3.19 percent in Table 6; scoring on pre-2022 data alone gives 2.24 percent.

One row deserves comment. Markers are defined only for pairs that co-bid often enough to be scored, so a tender whose winner never formed a scored pair is coded as not collusive because it could not be assessed, not because it was assessed and found clean. These unclassified tenders are about 30 percent of the estimation sample and 41 percent of the tenders that are not collusive. Dropping them leaves the coefficient at 3.03 percent in column 1, against 3.19 percent, and 3.50 percent in column 4, against 3.31 percent. Column 5 is unchanged by construction, since the continuous score is already estimated on the scored subsample.

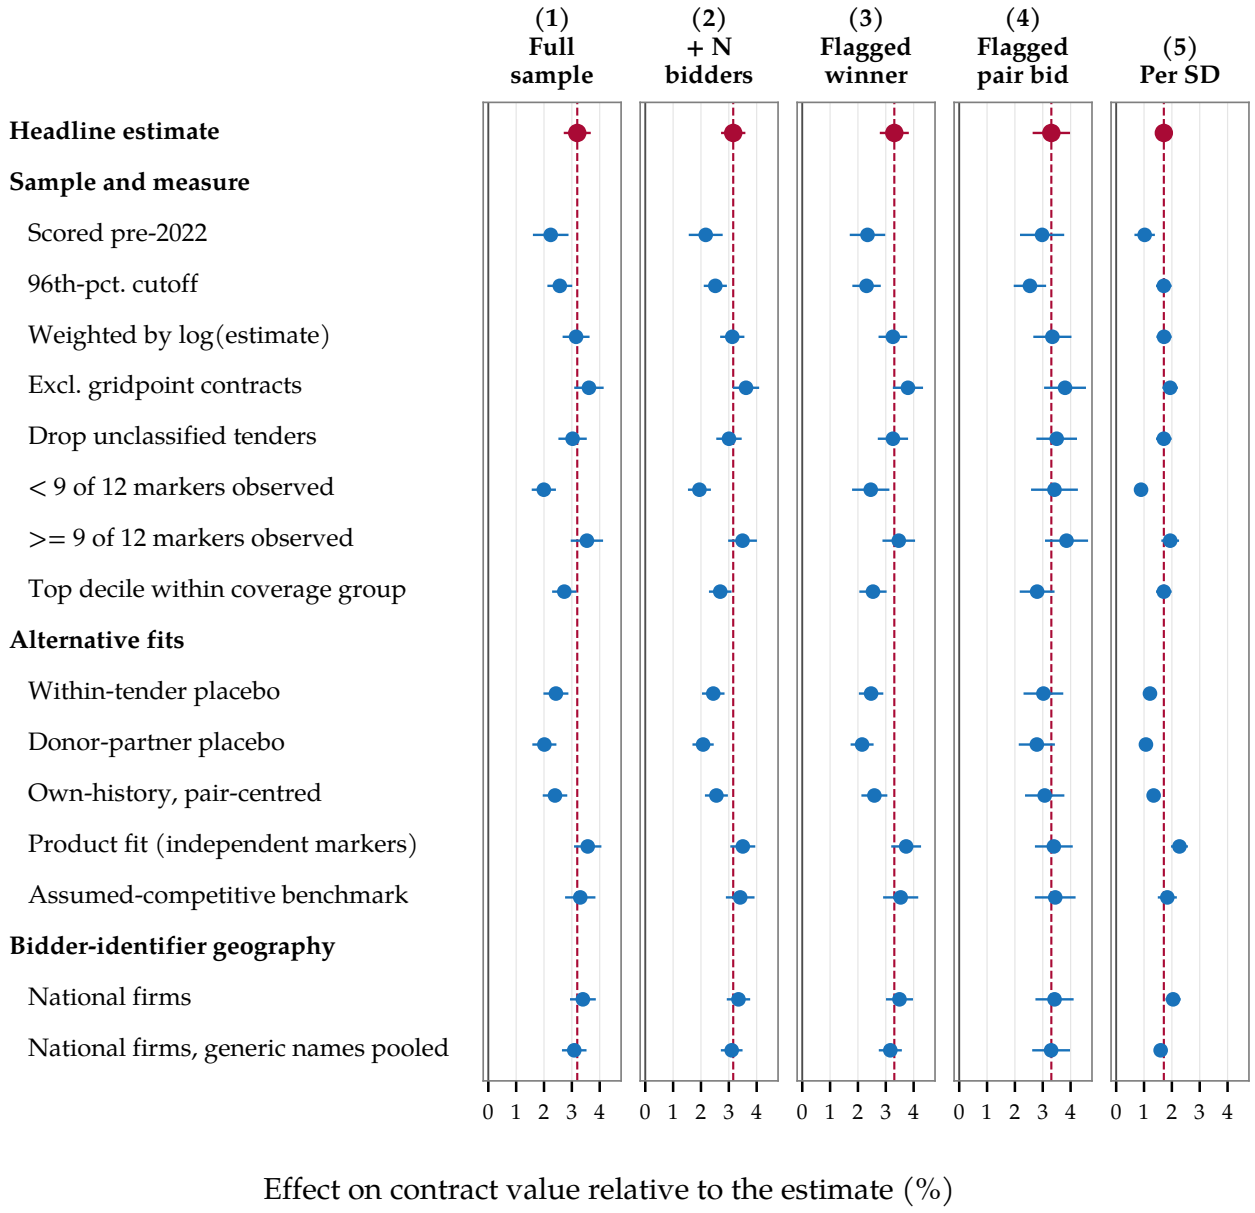


Figure D.1: Robustness: Table 6 under alternative samples, fits, and identifiers

Note: Each panel is one column of Table 6. Each row is one variant. Points are coefficients and horizontal bars 95% confidence intervals from standard errors clustered by district. The top row of each panel, in red, is the headline estimate of Table 6. Panels 1–4 estimate the effect of the binary collusive-tender indicator. Panel 5 estimates the effect of the continuous standardized pair score, in standard deviations. Estimates are in percent of the government estimate.

D.3 Competition controls and clustering

Columns 2 to 5 of Table 6 control linearly for the total number of bidders. Table D.2 instead adds one competition control at a time to column 1, on a common sample: the total number of bidders or the number of financial bidders, each entered linearly, in logs, and in bins. The coefficient is 3.19 percent with no competition control, 2.72 to 3.16 percent under the total count, and 2.69 to 3.04 percent under the financial count.

The paper clusters standard errors by district, the level at which bidders are identified and at which tenders share local shocks. Table D.3 varies the clustering. Clustering by state, by organization, or two-way by district and organization or by state and year raises the column-1 standard error from 0.0023 to at most 0.0036; clustering by product, winning bidder, fixed-effect cell, or tender package lowers it. Every coefficient remains significant at the 1 percent level; the smallest t-statistic across columns and schemes is 7.5.

Table D.2: Robustness: alternative competition controls

Dependent variable: log(contract value/government estimate)							
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
		N bidders			N fin. bidders		
	Baseline	Linear	Log	Bins	Linear	Log	Bins
Collusive tender	0.0319***	0.0316***	0.0288***	0.0272***	0.0304***	0.0291***	0.0269***
	(0.0023)	(0.0020)	(0.0018)	(0.0018)	(0.0018)	(0.0018)	(0.0017)
N	4,594,861	4,594,861	4,594,861	4,594,861	4,594,861	4,594,861	4,594,861
Mean contract/estimate ratio	0.933	0.933	0.933	0.933	0.933	0.933	0.933
Share treated	26.8%	26.8%	26.8%	26.8%	26.8%	26.8%	26.8%
District × Org × Year, Product × Year FE	✓	✓	✓	✓	✓	✓	✓
Tender-level controls	✓	✓	✓	✓	✓	✓	✓
N bidders control		✓	✓	✓			
N fin. bidders control					✓	✓	✓
Sample	Full	Full	Full	Full	Full	Full	Full
Clustering	District	District	District	District	District	District	District

Note: Same specification as Table 6 column 1, adding one competition control at a time. Column 1 repeats the baseline with no competition control. Columns 2–4 control for the total number of bidders and columns 5–7 for the number of financial bidders (bidders with a disclosed bid value), each entered linearly, in logs, and in bins (2, 3, 4, 5–9, 10 or more, with 2 omitted). All columns are estimated on the same sample. Standard errors, clustered by district, in parentheses. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

Table D.3: Robustness: alternative clustering of standard errors

	(1)	(2)	(3)	(4)	(5)
Collusive tender	**0.0319*** **	**0.0316*** **	**0.0331*** **	**0.0331*** **	
Max winner pair score (SD)					**0.0171*** **
District (baseline)	(0.0023)	(0.0020)	(0.0024)	(0.0032)	(0.0012)
State	(0.0035)	(0.0033)	(0.0039)	(0.0040)	(0.0014)
Organization	(0.0034)	(0.0030)	(0.0039)	(0.0043)	(0.0014)
Product	(0.0007)	(0.0005)	(0.0007)	(0.0012)	(0.0004)
District and organization (two-way)	(0.0036)	(0.0031)	(0.0040)	(0.0044)	(0.0015)
State and year (two-way)	(0.0034)	(0.0032)	(0.0038)	(0.0038)	(0.0014)
Winning bidder	(0.0004)	(0.0004)	(0.0005)	(0.0006)	(0.0003)
District $\times$ Org $\times$ Product $\times$ Year	(0.0005)	(0.0005)	(0.0006)	(0.0009)	(0.0003)
Tender package	(0.0002)	(0.0002)	(0.0003)	(0.0003)	(0.0001)

Note: The specification, sample, and coefficients are those of Table 6. Only the clustering of the standard errors varies. Columns 1–4 report the coefficient on the collusive-tender indicator and column 5 the coefficient on the standardized maximum likelihood-ratio score among the scored pairs that include the winner. The first row reports the point estimates. Each subsequent row reports the standard errors under one clustering scheme, starting with the district clustering used throughout the paper. The two-way schemes cluster on both dimensions at once rather than on their interaction. Under the winning-bidder scheme, each tender whose winner carries no bidder-district identifier (5,777 in the panel) forms its own cluster. Significance stars refer to the district clustering. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

D.4 Individual markers

Figure D.2 flags a pair if it is in the top decile of a single marker instead of the fitted model. In column 1 the premium ranges from −4.04 to 3.28 percent across the twelve markers. Margin steadiness (3.28 percent) and submission proximity (3.17 percent) give the largest premia, and several markers give negative ones. In column 4, which compares tenders in which a flagged pair bid, every marker gives a positive premium, from 1.92 to 5.65 percent. Many pairs are not observed on a given marker, and these rows count such a pair as unflagged. Dropping instead the untreated tenders in which the winner bid with a partner that the marker does not observe leaves these ranges almost unchanged: the column-1 range becomes −4.02 to 3.71 percent and the column-4 range 1.65 to 5.29 percent.

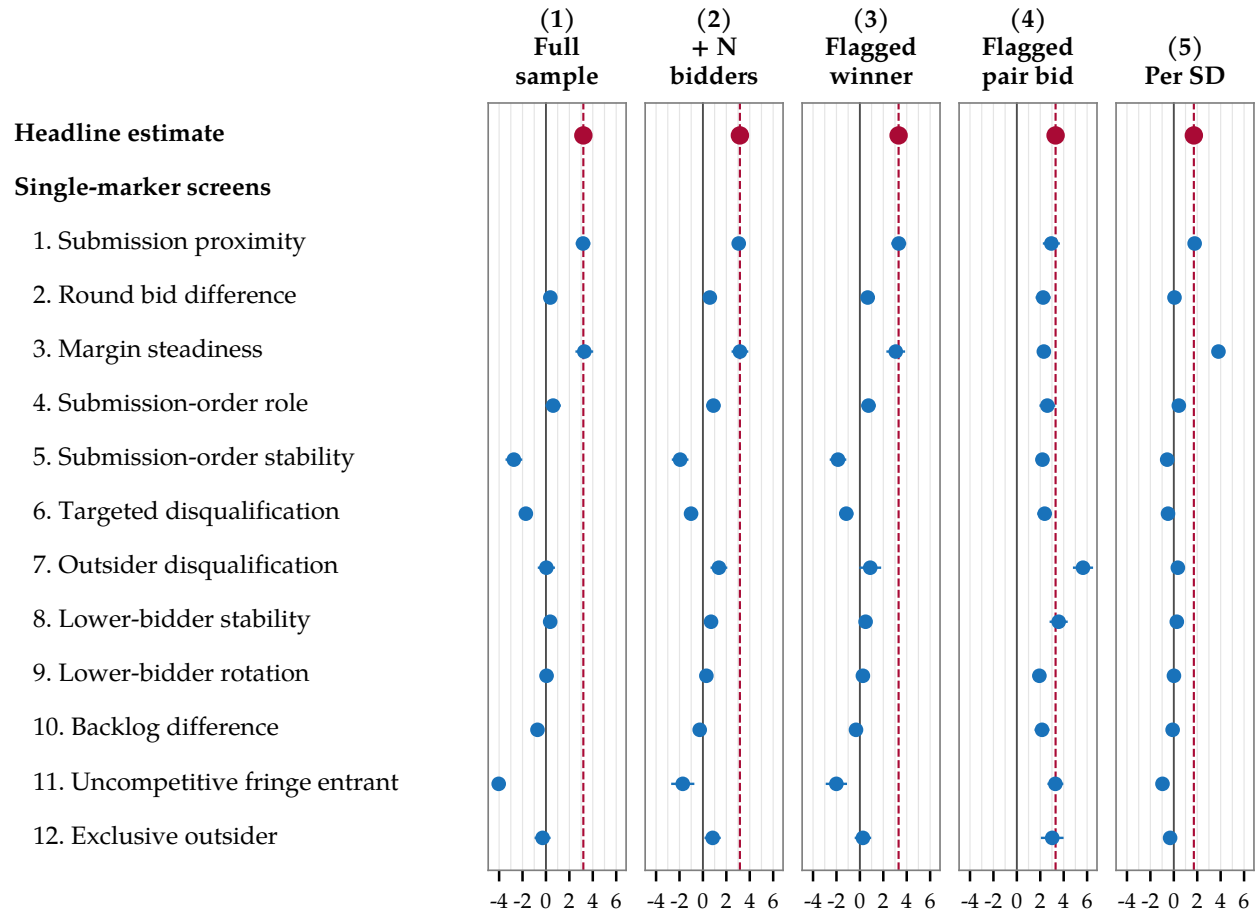


Figure D.2: Replacing the index with a single marker

Note: Layout as in Figure D.1. Each row redefines the treatment indicator using the top-decile classification of a single marker in place of the fitted model, holding the specification of Table 6 fixed. Markers are defined in Appendix A.1.

D.5 Placebo and alternative screens

Figure D.3 re-estimates the five columns of Table 6 with the headline flag replaced by another assignment on the same scored pairs; Table 8 reports column 2.

The matched designs come closest to the flagged pairs. Each flagged pair is replaced by an unflagged pair from the same district with a similar number of financial co-bids, drawn with probability declining in its own likelihood-ratio rank so that the replacements do not sit just below the cutoff; the second design also requires the same modal procuring organization and a tighter co-bid match. The replacements share the flagged pairs' geography and co-bidding intensity but not their markers. Under both designs the estimate is negative in columns 1 to 3 and 5: -2.25 and -2.11 percent in column 1, against 3.19 percent for the flagged pairs. In column 4, which compares tenders in which a flagged pair bid with tenders in which an unflagged pair bid and lost, it is 1.11 and 1.18 percent, against 3.31 percent.

Three further placebos flag a random tenth of the scored pairs, reassign the flag to unflagged pairs within bins of the financial co-bid count, or permute the likelihood-ratio score within district $\times$ year before reapplying the cutoff. Their column-2 estimates are -0.45, -0.76, and 0.02 percent, against 3.16 percent for the headline flag. Flagging the most frequent co-bidders gives 1.71 percent, and flagging the bottom decile of the score gives -5.69 percent.

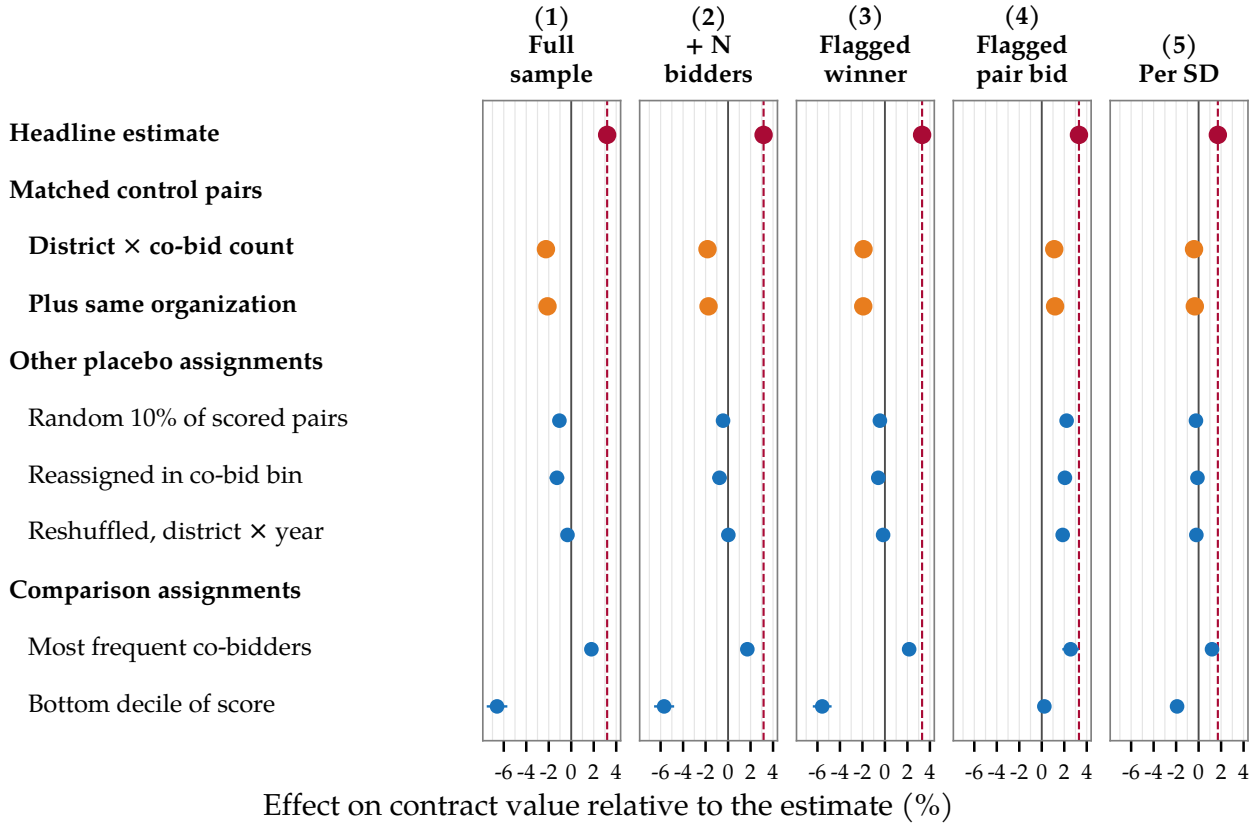


Figure D.3: Placebo and alternative screens

Note: Layout as in Figure D.1. The matched control pairs are shown first, in orange. The random, within-bin, and within-district-by-year assignments are the other placebos. The frequent-co-bidder and bottom-decile rows are comparisons.

D.6 Selection, and prices at and above the estimate

Table D.4 estimates the two margins jointly, $Y_i = \beta_1 D_i + \beta_2 \text{Collusive}_i + X_i'\gamma + \alpha_{d(i)o(i)t} + \alpha_{p(i)t} + \varepsilon_i$, where D_i indicates that both members of a flagged pair bid. Column 1 is the estimate summarized in Section 7.2: the presence term is -0.92 percent and the collusive-tender term 3.90 percent. Adding the total number of bidders (column 2) moves them to 0.33 and 2.90 percent: a flagged pair's presence is associated with fewer bidders, so conditioning on the count compares a tender in which a flagged pair co-bids to one without that drew equally few bidders.

The price increase in Table 6 comes from contracts moving onto the estimate. Overall, 41.1 percent of tenders settle within one percent of the estimate; the probability is 12.1 percentage points higher in collusive tenders (s.e. 0.7, column 1 of Table 6 with this outcome), and 11.9 points higher against the losing-pair comparison group of column 4. Settling at the estimate is consistent with colluding firms raising the price as far as they can without provoking buyer resistance (Marshall and Marx 2012, pp. 90–93), if the officer evaluating the lowest bid treats the published estimate as the benchmark for a reasonable price.

Table D.5 reports column 2 of Table 6 on four such outcomes, and Figure D.4 traces the effect window by window across the contract-to-estimate ratio. Only 13.4 percent of tenders settle above the estimate. In collusive tenders that probability is 2.5 percentage points higher in column 2, an 18.4 percent relative increase, by 2.5 points under the flagged-winner restriction (column 3) and by 3.1 points against the losing-pair comparison group (column 4); the continuous score raises it by 1.3 points per standard deviation (column 5).

Table D.4: Decomposing the price effect: presence versus winning

Dependent variable: log(contract value/government estimate)	(1)	(2)
Co-bidding flagged pair present (β_1)	-0.0092**	0.0033
	(0.0042)	(0.0052)
Collusive tender (β_2)	0.0390***	0.0290***
	(0.0029)	(0.0040)
Combined effect ($\beta_1 + \beta_2$)	0.0298***	0.0323***
	(0.0029)	(0.0026)
Observations	4,566,669	4,566,669
Mean contract/estimate ratio	0.933	0.933
Share with a co-bidding flagged pair	36.2%	36.2%
Share collusive	26.8%	26.8%
District × Org × Year, Product × Year FE	✓	✓
Tender-level controls	✓	✓
Missing-value and imputation flags	✓	✓
Total number of bidders		✓

Note: Dependent variable is log(contract value/government estimate). Each column regresses it jointly on D_i, an indicator that both members of a flagged pair bid in the tender, and on the collusive-tender indicator of Table 6. β_1 is the price difference between tenders in which a flagged pair co-bid but neither member won and tenders without such a pair. β_2 is the additional difference in collusive tenders. "Combined effect" is $\beta_1 + \beta_2$, with delta-method standard error. Column 2 adds the total number of bidders (linear), an outcome of the entry margin. All columns include district × organization × year and product × year fixed effects and the tender controls of equation (10). Flagged pairs are the top decile of scored pairs by the estimated likelihood ratio (Section 4). Standard errors, clustered by district, in parentheses. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

Table D.5: Where collusive tenders settle relative to the government estimate

	Exactly at estimate	Within ±1%	Above estimate	At least 10% below
Collusive tender	0.064^{***}	0.121^{***}	0.025^{***}	-0.110^{***}
	(0.004)	(0.007)	(0.004)	(0.006)
Mean of outcome	0.235	0.411	0.134	0.315
Tenders	4,594,861	4,594,861	4,594,861	4,594,861

Note: Each column is column 2 of Table 6 (full sample, number-of-bidders control) with a different outcome, an indicator of the contract-to-estimate ratio r. Exactly at estimate: $|r-1| \leq 0.0005$. Within ±1%: $|r-1| \leq 0.01$. Above: $r > 1$. At least 10% below: $r \leq 0.90$. Standard errors, clustered by district, in parentheses. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

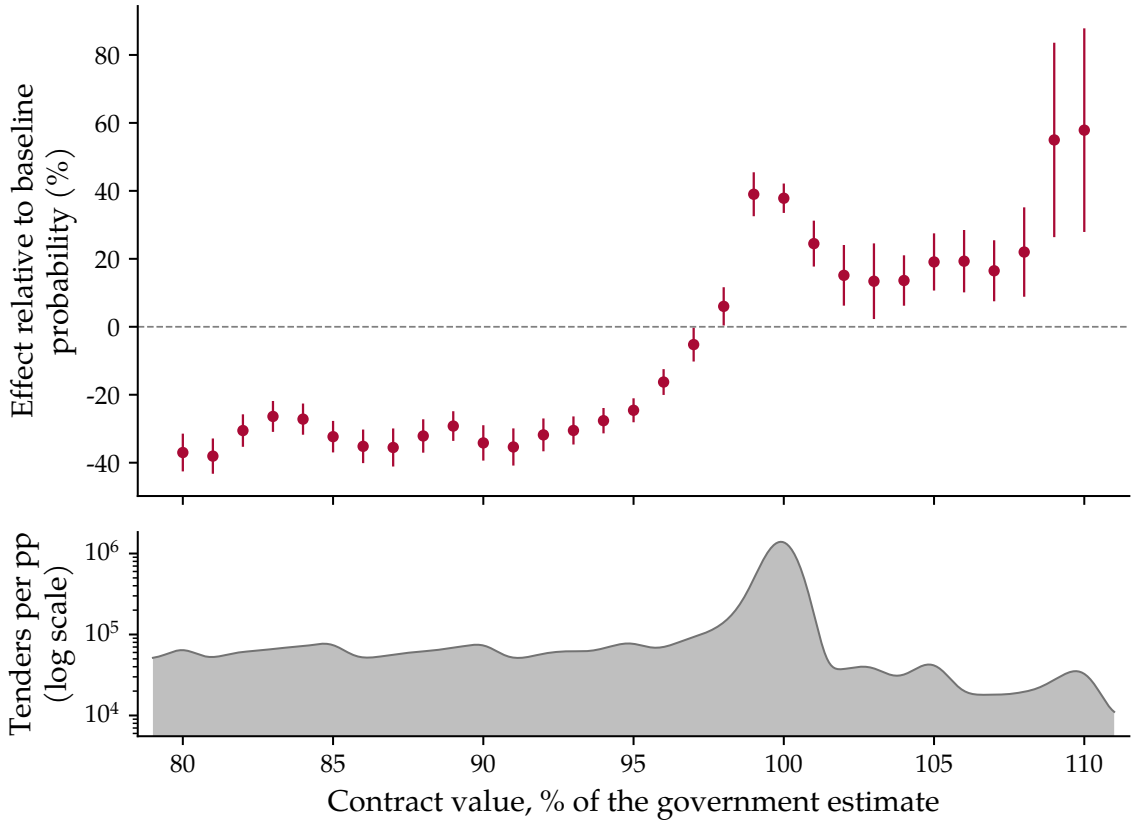


Figure D.4: Where collusive tenders settle: effect by window of the contract-to-estimate ratio

Note: Top: one regression per window centre c, from 80 to 110 percent of the estimate, each column 2 of Table 6 with outcome equal to one if the contract-to-estimate ratio lies within $c \pm 1$ percentage point. Coefficients are divided by the share of other tenders in the window. Bars are 95% confidence intervals, standard errors clustered by district. Bottom: density of the ratio, log scale. The windows cover 79.2 percent of tenders.

D.7 Price effects by field size

Figure D.5 reports the price of non-collusive tenders and the premium by the number of financial bidders, for flagged and matched control pairs (Section 8.1).

Figure D.7 breaks the comparison of Table 6 down by the number of financial bidders, for flagged pairs and for matched control pairs, on three outcomes: the contract price, the probability of a contract within 1 percent of the estimate, and the dispersion of bids.

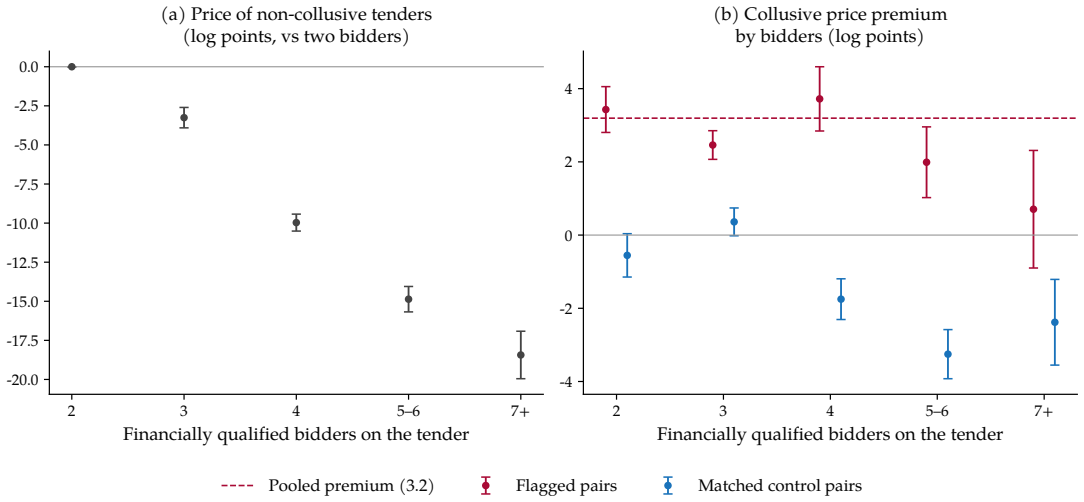


Figure D.5: Bidders on the tender, prices and the premium

Note: Panel (a): coefficients on bins of financial bidders, relative to two, in a regression of 100 times the log contract-to-estimate ratio on tenders that are not collusive. Panel (b): the collusive-tender coefficient interacted with the same bins, all tenders. The dashed line is the pooled coefficient of Table 6. Blue: the same coefficient for tenders won by a member of a matched control pair whose partner also bids, estimated after dropping the tenders that are collusive under the flag. Controls and fixed effects as in Table 6. Bars are 95 percent confidence intervals. Standard errors clustered by district.

Figure D.6 relates the mean price of collusive and of non-collusive tenders in a district $\times$ product market to the share of the market's classified tenders that are collusive (Section 8.1).

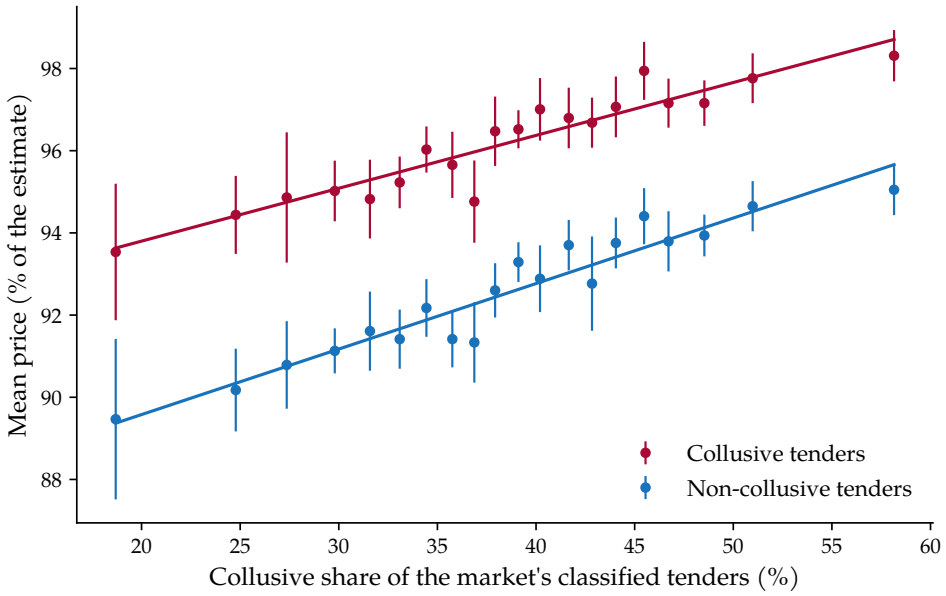


Figure D.6: Prices and the collusive share of a market

Note: 3,838 district × product (L1) markets with both collusive and non-collusive tenders. Collusive share: collusive tenders over classified tenders in the market. Twenty binned means of the mean contract-to-estimate ratio of each group of tenders, with 95 percent confidence intervals, after controls for rule groups, lotteries, the log numbers of firms and tenders, the classified share, the tender mix, and state fixed effects; markets weighted by classified tenders. Slopes per 10 percentage points of collusive share: 1.29 (s.e. 0.23) for collusive and 1.59 (0.25) for non-collusive tenders; mean gap 3.7 percentage points, falling by 0.31 (0.09) per 10 percentage points. Standard errors clustered by state.

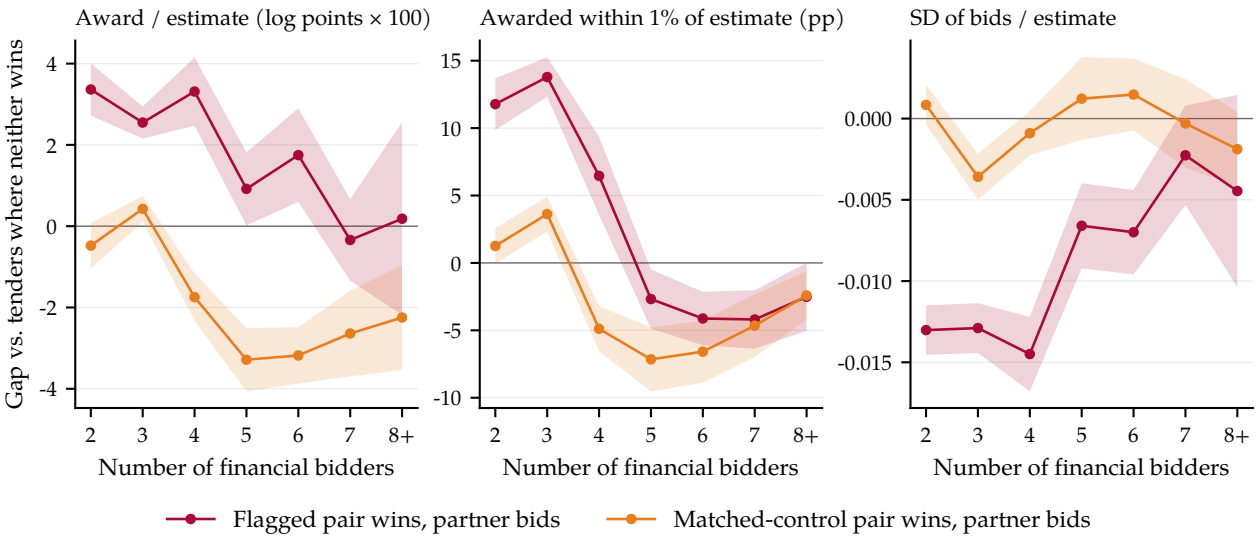


Figure D.7: Prices and bid dispersion when a flagged pair wins, by number of financial bidders

Note: Each panel reports, by number of financial bidders, the gap between tenders that a flagged pair wins with its partner bidding and tenders in which neither a flagged pair nor a matched control pair wins with its partner bidding (red), and the same gap for matched control pairs (orange). Gaps are coefficients on bin-specific indicators in a single regression per outcome, with the fixed effects and tender-level controls of Table 6, on its sample. Outcomes are log(contract value/ government estimate) in log points ×100, an indicator for a contract within 1 percent of the estimate in percentage points, and the standard deviation of the bid-to-estimate ratio across the tender's bids. Shaded bands are 95 percent confidence intervals from standard errors clustered by district.

E Transfer by pair score

In Section 9.2, the transfer is calculated as the estimated price premium multiplied by the value of collusive tenders. Both quantities were estimated using pairs in the top decile of the pair score. In this section, I estimate the transfer more precisely by estimating price premia separately for pairs in different bands of the score up to the estimated share of colluding pairs $\hat{\Lambda}$ (33.6 percent): 0–5, 5–10, . . . , 30–33.6 percent. I estimate a separate price premium $\hat{\beta}_b$ for every band b, comparing collusive tenders (in which the winner and one partner form a pair in that band of the score distribution) against non-collusive tenders (in which the highest-scoring pair is in the bottom $1 - \hat{\Lambda}$ of the score distribution) and unclassified tenders. The transfer is the band analogue of (12), the sum over bands of each band's price premium times its annual volume of collusive tenders: $\widehat{\text{Transfer}} = \sum_b \hat{\beta}_b \overline{Q}_b$, where $\overline{Q}_b$ is $\overline{Q}$ of (12) restricted to the collusive tenders of band b.

Panel A of Figure E.1 shows the estimated price premia, which are highest for the highest-scoring pairs. Panel B plots the implied value of the transfer. At the top-decile cutoff it is USD 187.1 million a year, close to the headline transfer of USD 159 million in Table 10. Extending the cutoff to the top $\hat{\Lambda}$ share adds bands with smaller premia and raises it to USD 329.8 million.

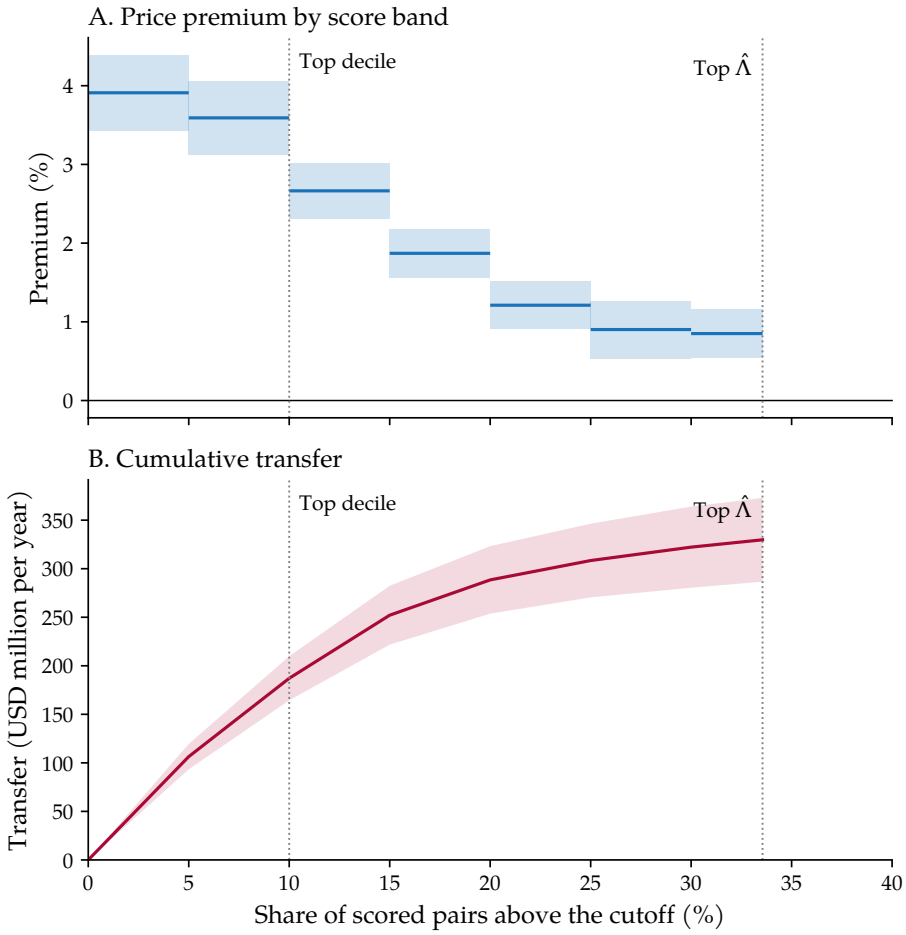


Figure E.1: Price premium and transfer by pair-score band

Note: Panel A: price premium of each 5-percentage-point band of the pair score (horizontal lines), relative to non-collusive tenders (tenders whose highest-scoring pair is in the bottom $1 - \hat{\Lambda}$ and unclassified tenders). Estimated using specification of column 2 of Table 6. Standard errors clustered by district, 95% confidence interval shaded. Panel B: annual transfer to pairs above the cutoff: the sum over bands of the price premium times the average annual volume of that band, with 95% confidence interval. Dotted lines: the top decile (the flag) and the top $\hat{\Lambda}$ share.

F Data

F.1 Data collection: web scraping

I scrape the universe of public procurement records published on India's government e-procurement portals. Most Indian states (e.g., Uttar Pradesh, Maharashtra, Tamil Nadu), union territories (e.g., Delhi, Chandigarh), major central agencies (e.g., Ministry of Defense, Ministry of Road Transport and Highways), and public-sector undertakings (e.g., Indian Railways, Coal India Limited) publish procurement notices and results through instances of the National Informatics Centre's Government e-Procurement System (GePNIC), deployed under the Central Public Procurement framework. Some states (e.g., Gujarat, Bihar) use a different procurement software. For these states, only central government and PMGSY rural roads tenders with project locations in those states are available on the GePNIC portals. The unit of observation is a *tender*, the procurement process for a single contract. Each tender has a unique identifier, the tender ID. The tender record includes information on the procuring organization, details of the contract, the estimated value, the location, and the timeline of the tender process. The tender records include technical and financial bids submitted by potential suppliers.

I scrape 76 portals. Table F.1 lists every portal, its address, and the number of tenders and bids I collect from it, grouped into state and union-territory portals (Panel A), central-government and public-sector-undertaking portals (Panel B), and the portals of the rural-roads program PMGSY (Panel C). The state-level portals cover 22 states and 8 union territories, together representing 72% of the population and 70% of the GDP of India. Figure F.1 maps the portals available in each state and the number of tenders by district.

In total I collect 15.7 million distinct tenders and 43.3 million bids. These include cancelled tenders, tenders that received no bids, and tenders whose status was never updated, so the regression samples are smaller.

F.1.1 Extracted variables

Collection proceeds in two stages. First, for each portal I obtain the list of identifiers of all tenders in the portal's public tender listing by iterating over all calendar dates and tender status combinations.[24] Second, for each identifier I retrieve the tender's complete published record, which includes the following variables:

- *Tender attributes*: title, organization chain, type and category, form of contract, location and pincode, the estimated value, fee and earnest-money-deposit (EMD) amounts, the period of work, corrigenda, and a set of stage indicators marking whether the tender passed technical-bid opening, technical evaluation, financial-bid opening, financial evaluation, and award of contract (AOC), plus the associated stage dates.
- *Bids*: one row per submitted bid, recording the bidder name, the quoted value, the bid status and remarks, and an indicator of whether the bid was the winning bid.
- *Awards*: one row per award, recording the awarded bidder, the award value, and the award currency.

[24] The listing pages are protected by a character CAPTCHA, solved using optical character recognition (OCR). This collection adhered to each portal's terms of use and respected any limitations specified in its robots.txt file.

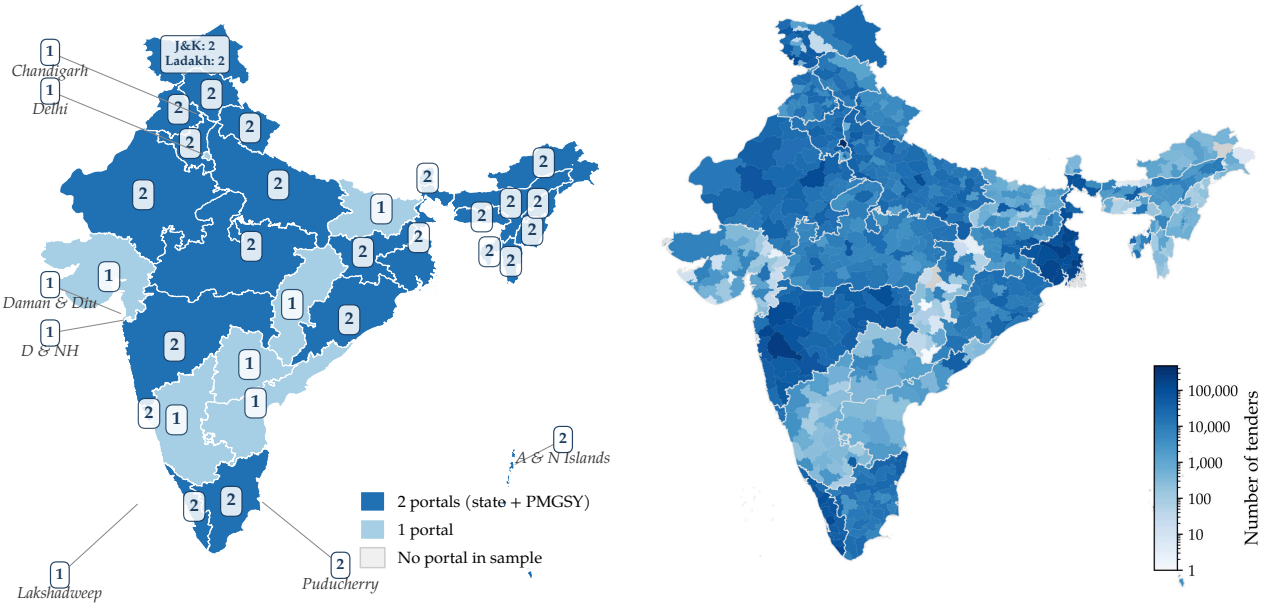
Table F.1: Government e-procurement portals scraped, with collection counts

Jurisdiction / entity	e-Procurement portal	Tenders	Bids	Orgs
Panel A. State and union-territory portals				
Andaman & Nicobar Islands	eprocure.andamannicobar.gov.in	21,613	56,831	38
Arunachal Pradesh	arunachaltenders.gov.in	2,589	5,830	26
Assam	assamtenders.gov.in	74,200	238,843	120
Chandigarh	etenders.chd.nic.in	69,826	183,398	12
Dadra & Nagar Haveli	dnhtenders.gov.in	7,608	16,454	2
Daman & Diu	ddtenders.gov.in	4,434	7,646	2
Delhi	govtprocurement.delhi.gov.in	359,849	1,480,538	218
Goa	eprocure.goa.gov.in	25,393	49,474	105
Haryana	etenders.hry.nic.in	571,344	1,110,228	4
Himachal Pradesh	hptenders.gov.in	179,509	450,900	100
Jammu & Kashmir	jktenders.gov.in	881,544	1,911,868	89
Jharkhand	jharkhandtenders.gov.in	134,595	538,594	134
Kerala	etenders.kerala.gov.in	1,285,543	2,073,827	473
Ladakh	tenders.ladakh.gov.in	37,963	93,062	27
Lakshadweep	tendersutl.gov.in	3,375	9,462	3
Madhya Pradesh	mptenders.gov.in	549,541	1,163,406	246
Maharashtra	mahatenders.gov.in	2,018,623	6,104,889	383
Manipur	manipurtenders.gov.in	3,080	7,086	57
Meghalaya	meghalayatenders.gov.in	1,748	2,741	36
Mizoram	mizoramtenders.gov.in	416	980	25
Nagaland	nagalandtenders.gov.in	696	2,118	31
Odisha	tendersodisha.gov.in	501,698	2,455,711	133
Puducherry	pudutenders.gov.in	20,905	53,345	63
Punjab	eproc.punjab.gov.in	314,810	567,539	57
Rajasthan	eproc.rajasthan.gov.in	885,205	2,874,710	320
Sikkim	sikkimtender.gov.in	754	2,372	40
Tamil Nadu	tntenders.gov.in	727,418	1,216,479	192
Tripura	tripuratenders.gov.in	67,824	171,593	120
Uttar Pradesh	etender.up.nic.in	1,925,093	4,610,042	372
Uttarakhand	uktenders.gov.in	97,138	304,919	319
West Bengal	wbtenders.gov.in	2,264,299	6,657,979	346
Subtotal		*13,038,633*	*34,422,864*	*4,093*
Panel B. Central-government and public-sector-undertaking portals				
Bharat Electronics (BEL)	eprocurebel.co.in	3,604	11,130	1
Bharat Heavy Electricals (BHEL)	eprocurebhel.co.in	53,181	148,345	1
Central Public Procurement (eProcure)	eprocure.gov.in	824,237	2,582,339	691
Central Public Procurement (eTenders)	etenders.gov.in	260,080	673,850	117
Chennai Petroleum (CPCL)	cpcletenders.nic.in	9,145	18,139	2
Coal India	coalindiatenders.nic.in	344,773	1,039,550	13
Defence Procurement	defproc.gov.in	668,362	2,714,234	16
Goa Shipyard (GSL)	eprocuregsl.nic.in	1,820	4,918	1

continued on next page

Table F.1 continued

Jurisdiction / entity	e-Procurement portal	Tenders	Bids	Orgs
Hindustan Shipyard (HSL)	eprocurehsl.nic.in	5,150	21,495	1
Indian Oil (IOCL)	iocletenders.nic.in	188,473	914,451	1
Mazagon Dock (MDL)	eprocuremdl.nic.in	6,442	18,990	1
NTPC	eprocurentpc.nic.in	95,939	214,030	8
Subtotal		*2,461,206*	*8,361,471*	*853*
	Panel C. PMGSY rural-roads programme portals			
Andaman & Nicobar Islands	pmgsytendersani.nic.in	143	146	—
Andhra Pradesh	pmgsytendersap.gov.in	2,106	2,857	—
Arunachal Pradesh	pmgsytendersarn.gov.in	1,243	3,535	—
Assam	pmgsytendersasm.gov.in	12,376	14,559	—
Bihar	pmgsytendersbih.gov.in	81,060	281,230	—
Chhattisgarh	pmgsytenderscg.nic.in	0	0	—
Goa	pmgsytendersgoa.gov.in	16	16	—
Gujarat	pmgsytendersguj.gov.in	838	509	—
Haryana	pmgsytendershry.nic.in	375	850	—
Himachal Pradesh	pmgsytendershp.gov.in	4,367	6,804	—
Jammu & Kashmir	pmgsytendersjk.gov.in	17,411	6,466	—
Jharkhand	pmgsytendersjhr.gov.in	6,380	13,018	—
Karnataka	pmgsytenderskar.gov.in	0	0	—
Kerala	pmgsytendersker.gov.in	4,381	2,761	—
Ladakh	pmgsytendersla.nic.in	121	229	—
Madhya Pradesh	pmgsytendersmp.gov.in	2,592	7,608	—
Maharashtra	pmgsytendersmah.gov.in	9,580	19,020	—
Manipur	pmgsytendersman.gov.in	2,097	6,503	—
Meghalaya	pmgsytendersmeg.gov.in	1,719	2,665	—
Mizoram	pmgsytendersmiz.gov.in	747	1,798	—
Nagaland	pmgsytendersngl.gov.in	210	533	—
National	pmgsytenders.gov.in	11,041	40,893	—
Odisha	pmgsytendersori.gov.in	14,186	23,963	—
Puducherry	pmgsytenderspy.nic.in	185	208	—
Punjab	pmgsytenderspb.gov.in	1,257	2,182	—
Rajasthan	pmgsytendersraj.gov.in	8,213	16,830	—
Sikkim	pmgsytendersskm.gov.in	400	525	—
Tamil Nadu	pmgsytenderstn.gov.in	8,168	14,399	—
Telangana	pmgsytendersts.nic.in	0	0	—
Tripura	pmgsytenderstrp.gov.in	966	1,063	—
Uttar Pradesh	pmgsytendersup.gov.in	5,661	20,362	—
Uttarakhand	pmgsytendersuk.gov.in	4,774	14,212	—
West Bengal	pmgsytenderswb.gov.in	13,002	19,166	—
Subtotal		*215,615*	*524,910*	—
Total		**15,715,454**	**43,309,245**	**4,946**



(a) Portal coverage

(b) Number of tenders by district

Figure F.1: Geographic coverage of e-procurement portals and tender counts

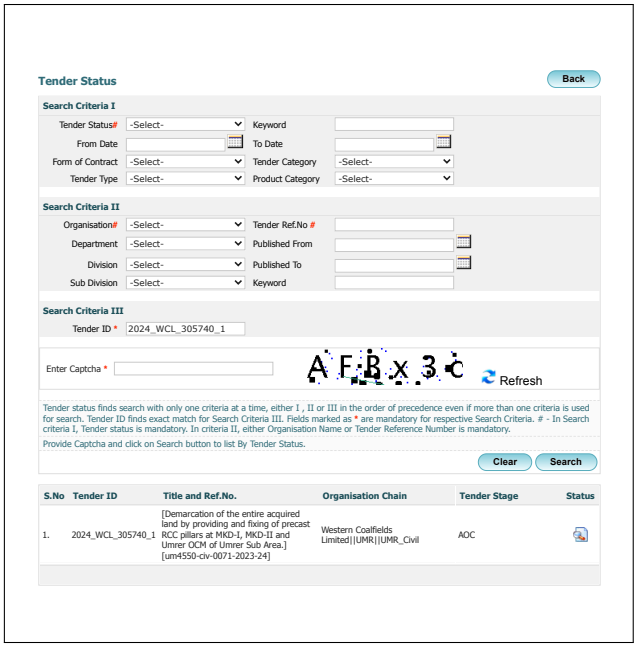
Note: Panel (a): Number of distinct state-level portals (state portal and/or PMGSY rural-roads portal) available by state. Central-government and public-sector-undertaking portals (Panel B of Table F.1) are excluded. Panel (b): Number of tenders by district. Districts with no tenders shown in gray.

F.1.2 Example

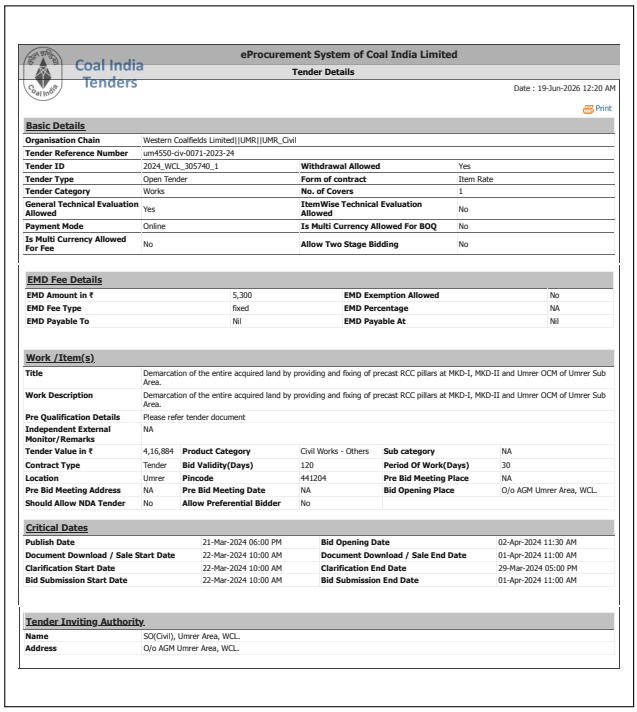
Figure F.2 illustrates the records collected for a single tender: the portal landing page (Panel a), the tender-status search and listing (Panel b), the tender's detail page (Panel c), and the combined bidder list, opened financial bids, and award of contract (Panel d). The final panel is the source of the bid-level data and identifies the winner and contract value.



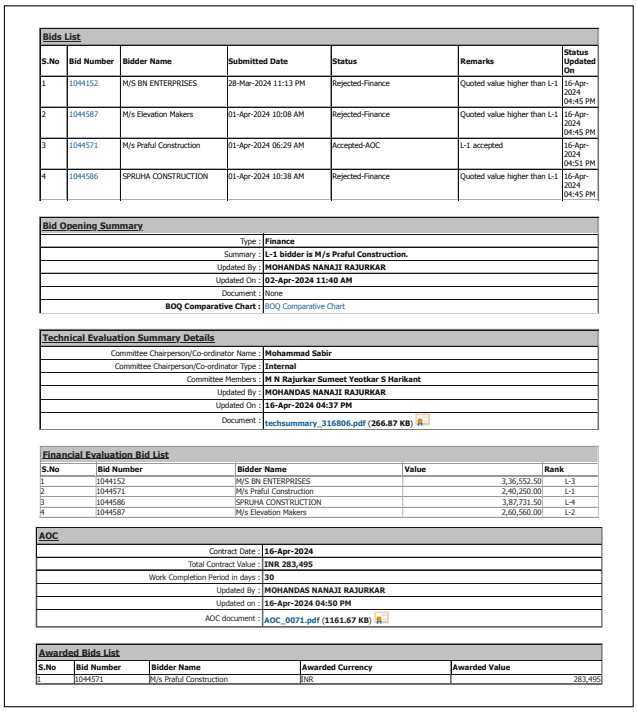
(a) Portal landing page



(b) Tender-status search and listing



(c) Tender detail page



(d) Bidders, financial bids, and award

Figure F.2: Example of the records collected for one tender.

Note: Screenshots of one tender on a government e-procurement portal (Coal India Limited), in the order in which the pages appear: the portal landing page, the tender-status search and listing, the tender detail page, and the combined bidder list and award-of-contract page. In this example, bids exclude 18% Goods and Services Tax (GST), and the award value includes it.

F.2 Minimum-bidder rules

Different procuring organizations use different rules regarding the minimum number of bidders required for a tender to be awarded. I infer each portal's threshold from its deadline extensions (Figure 2, panel b): a portal follows a two-bid rule if at least 20 percent of the tenders holding one bid at the original deadline are extended, at most 15 percent of those holding two, and the share of tenders with only one bid that are extended is at least double the share of tenders with only two bids that are extended. The three-bid rule is defined analogously on two and three bids. Each portal is assigned the rule that covers at least 60 percent of its tenders.

Figure F.3 shows that the minimum-bidder rules bind: more than half of tenders in portals with two-bidder rules have exactly two financial bidders, and two-thirds of tenders in portals with three-bidder rules have exactly three financial bids. For other portals, around 30% of tenders have two bids, and around 25% have three bids.

Tender extensions are associated with more financial bidders: tenders with one bid at the original deadline end with 1.6 bids on average, and tenders with two end with 2.2; tenders with three or four bids gain almost none (Figure F.4).

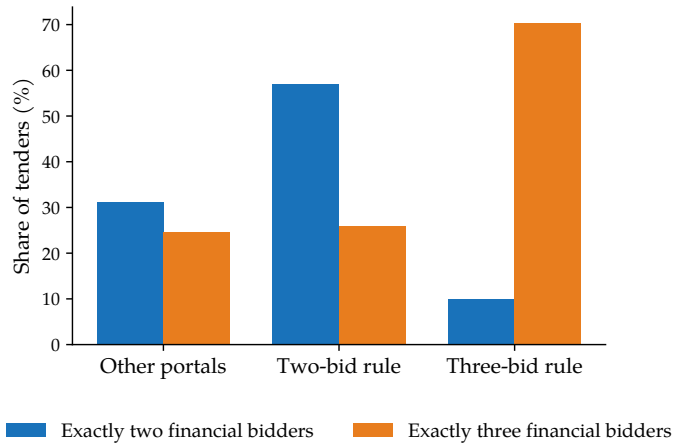


Figure F.3: Number of financial bidders by minimum-bidder rule

Note: Share of tenders with exactly two and exactly three financial bidders. Portals are grouped by the rule implied by their deadline extensions (Section 2.4).

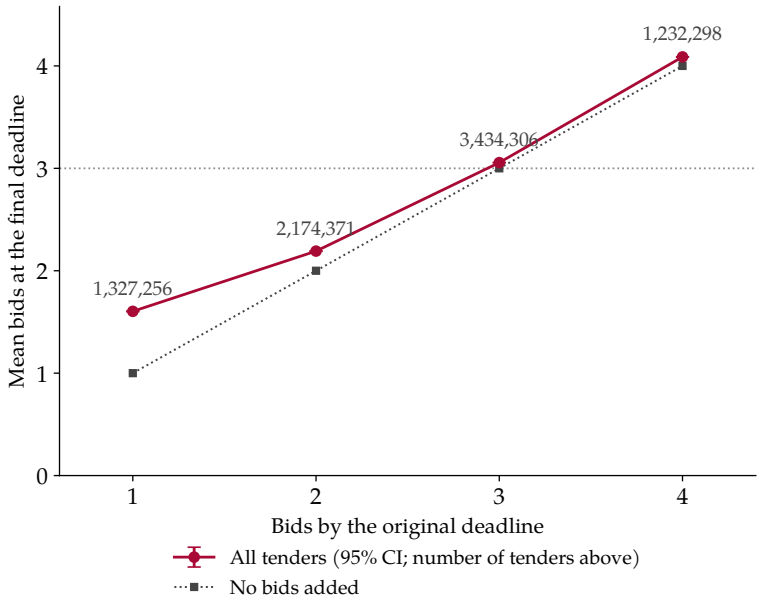


Figure F.4: Final bids received by the final deadline, by bids in hand at the original deadline

Note: Mean bids received by the final deadline (vertical axis), by bids in hand at the original deadline (horizontal axis). The dotted line marks no bids added. Sample includes all tenders, including tenders whose deadline was not extended.

F.3 Tender-level information

This section describes how I build the tender-level dataset from the raw scraped records. Each tender is identified by its automatically assigned tender ID. I list the tender-level variables, describe how the raw variables are cleaned, and construct additional variables (geographic location, product classification). The analysis samples are defined in Section F.5.6.

F.3.1 Tender-level variables

Each tender record carries identification and administrative variables, a free-text description, a location, a schedule of dates, financial and contract terms, bidding-rule variables, and the officials who record each procurement stage. Table F.2 lists these variables, grouped by type, with example values and the share of non-missing observations for each variable. Given that some tenders do not reach the award stage, award variables have lower coverage. The portals also record status updates, document attachments, and various internal timestamps. These are not used and not reported here.

Table F.2: Tender-level variables

Variable	Description	Example values	% non-missing
Identification & administration			
Tender ID	Portal-assigned unique identifier	2021_MPHID_134197_1	100
Reference number	Agency's internal reference number	MPHB SDL-18	100
Procuring organization	Department or body issuing the tender	MP Housing and Infrastructure Dev. Board	100
Tender type	Open, limited, EOI, etc.	Open Tender; Limited; Single	100
Tender category	Works, goods, or services	Works; Goods; Services	100
Contract type	Tender, rate contract, or empanelment	Tender; Rate Contract; Empanelment	100
Current status	Lifecycle status label	AOC; Financial Bid; Financial Evaluation	100
Number of corrigenda	Count of amendments issued	1; 2; 3	26.6
Corrigenda types	Types of amendments issued	Date; Cancellation of Tender; Other	26.6
Description & products			
Title	Short tender title	Construction of sub health centre	100
Work description	Free-text scope of work	Construction of sub health centre	100
Product category	Agency-entered product or work line	Civil Works; Civil Works - Roads; Civil Works - Others	100
Sub-category	Finer agency sub-category (sparse)	Civil Work; Civil Works; Road Work	12.1

continued on next page

Table F.2 continued

Variable	Description	Example values	% non-missing
Location			
PIN code	Six-digit postal code	180001; 226001; 110010	98.5
Location text	Free-text work location	Lucknow	94.4
Dates & schedule			
Publication date	Date the tender was published	2021-03-17	100
Bid start	Date bid submission opens	2021-03-17	100
Bid deadline	Date bid submission closes	2021-03-31	100
Bid opening	Date technical bids are opened	2021-04-01	100
Financial opening	Date financial bids are opened	2021-04-06	77.4
Award date	Date the contract is awarded	2021-06-30	47.1
Value & contract terms			
Estimated value	Agency pre-tender cost estimate	0; 1,000,000; 500,000	86.0
EMD amount	Earnest-money deposit	0; 50,000; 10,000	100
EMD type	EMD set as a fixed sum or a percentage	Fixed; Percentage	100
EMD percentage	EMD rate when set as a percentage	2; 1; 10	23.9
Tender fee	Document or processing fee	0; 500; 1,000	93.5
Award value	Contract value at award	500,000; 1,000,000; 1	47.5
Period of work (days)	Work period stated in tender	30; 180; 90	95.8
Work completion (days)	Completion period at award	180; 90; 30	47.5
Bidding rule & mechanism			
Form of contract	Item-rate, percentage, lump-sum, etc.	Percentage; Item Rate; Works	100
Number of covers	Envelope/cover stages required	2; 1; 3	100
Bid validity (days)	Days a submitted bid stays valid	90; 120; 180	100
Two-stage bidding	Two-stage process permitted	False; True	100
Preferential bidder	Preference/reservation permitted	False; True	100
EMD exemption	EMD waiver permitted	False; True	100
Fee exemption	Fee waiver permitted	False; True	98.9
Recording officials & committee			
Recorded by: tech. bid opening	Official who recorded technical-bid opening	Mangilal Dhakse	82.8
Recorded by: tech. evaluation	Official who recorded technical evaluation	Mangilal Dhakse	89.3
Recorded by: fin. bid opening	Official who recorded financial-bid opening	Ram Pratap Verma	68.5

continued on next page

Table F.2 continued

Variable	Description	Example values	% non-missing
Recorded by: fin. evaluation	Official who recorded financial evaluation	Mangilal Dhakse	59.6
Recorded by: award	Official who recorded the award	Sanjay Agrawal	47.5
Tech. evaluation committee chair	Chair of the technical evaluation committee	K. K. Sharma	87.2
Fin. evaluation committee chair	Chair of the financial evaluation committee	K. K. Sharma	58.7

F.3.2 Tender status

For every record, the portal displays a label for the tender status. The label advances as the tender moves through technical and financial evaluation to award, and it can also record that a tender was cancelled or re-tendered. Table F.3 reports the distribution of this status across all scraped tenders.

The status is updated manually by a procurement officer and is not always current, so whenever more recent information is available I override the status label with a more accurate outcome. For example, if a tender is recorded as being in the evaluation stage but a bid has been awarded and a contract is recorded, I treat the tender as awarded.

Of all tenders, 34 percent are recorded as awarded or concluded, and 17 percent as cancelled or retendered. The remaining tenders carry a status that records no outcome: 13 percent record no bid, stage, award, or cancellation, and 36 percent are recorded in a submission or evaluation stage. The last two columns of Table F.3 report what share of tenders in each category have received a bid and what share has been awarded. Tenders in the intermediate statuses have almost always received bids, and a winner can be identified for 16 percent of them.

Table F.3: Tender status

				Share of status (%)	
Status	Meaning	Tenders	%	Any bid	Award
To be Opened	No bid or stage recorded	2,088,686	13.3	0.0	0.0
Technical bid	Technical-bid submission stage	1,115,292	7.1	81.0	0.0
Technical evaluation	Technical bids under evaluation	1,127,578	7.2	99.1	0.2
Financial bid	Financial-bid stage reached	2,022,976	12.9	98.4	0.0
Financial evaluation	Financial bids under evaluation	1,341,118	8.5	99.2	65.8
AOC	Award of contract published	5,410,908	34.4	99.5	99.5
Concluded	Procurement process concluded	170	0.0	100.0	0.0
Cancelled	Tender cancelled	1,782,139	11.3	27.3	2.7
Retender	Cancelled and re-tendered	826,587	5.3	16.0	0.7
Total		15,715,454	100.0	72.2	40.2

Note: Lifecycle status over all 15,715,454 scraped tenders, ordered by procurement stage. The last two columns report the share of tenders in that status that record at least one bid, and for which an awarded bid can be identified, respectively.

F.3.3 Exclusion of test and placeholder records

Some records are portal artifacts rather than genuine procurements. I drop test and dummy tenders, identified by the substrings "_TEST_" or "_NIC_" in the tender ID or by explicit labels in the title or work description (e.g. "test tender", "dummy tender", "do not bid", "for training purpose").[25] This removes 852 tenders. Some tender values contain a placeholder, usually a repeated-digit value such as 999,999,999, 1,111,111, or 555,555. I set the estimate to missing when it contains a single non-zero digit repeated at least five times. This affects 1,218 tenders.

F.3.4 Geographic location

I identify the geographic location of each tender using two pieces of information. First, all tenders contain a six-digit Postal Index Number (PIN), the unit of the India Post addressing system (Government of India 2022). I use the GeoNames postal-code file to match each PIN to a state and district (GeoNames 2025). The PIN is precise, but it sometimes records the address of the procuring government office instead of the location of the project site.

Therefore, second, I use the free-text location description drawn from the title, work-description, and location fields in the tender. These fields are often missing or uninformative, but when they mention the project site, they provide information beyond the PIN. I extract place names (districts, blocks or taluks, gram panchayats, and villages) with regular expressions and look up each place name in the GADM 4.1 boundaries for India (GADM 2026), which cover districts and sub-districts. Villages and towns are looked up in the GeoNames gazetteer (GeoNames 2025), which lists about 300,000 Indian settlements with coordinates. When a name has multiple candidates, I use the location closest to the PIN.

When the PIN and the free text agree on the district, I treat the PIN as reliable and use its coordinates. When they disagree, I use the free-text place. Taken together, each tender is mapped to a state, a district, and the coordinates inferred from either the PIN code or the free-text location description.

F.3.5 Contract duration

Two variables record contract length: the planned period of work stated in the tender notice and the planned period reported at award. For 98.6% of tenders, the two variables either agree or only one is populated. Where they disagree, I use the duration reported at award, falling back to the duration at notice when the value at award lies outside a plausible window of one week to three years and the duration at notice does not. The reconciled duration is then winsorized to that same window.

F.3.6 Product classification

The raw data contain no standardized product classification, so I classify tenders by text clustering. For each tender I concatenate the title, work description, and product-category variables into one description string. Next, common abbreviations are expanded and stopwords removed. I map each description to a dense vector with a pre-trained sentence-embedding model (`all-mpnet-base-v2`) (Reimers and Gurevych 2019) and normalize the vectors. I then reduce the estimated embeddings

[25] Genuine procurement contracts such as "testing and commissioning" are retained.

into five dimensions (using UMAP) and cluster the reduced vectors (using HDBSCAN) (Campello, Moulavi, and Sander 2013). I fit the model on a random subsample of 500,000 tenders. The classification of the remaining tenders is then obtained by mapping their normalized embedding vectors onto the fitted topic model.

I impose a minimum topic size of 1,000 tenders. The classifier recovers 87 topics, one of which is a residual unclassified category. These topics are taken as sub-categories (level 2), and named by hand by looking at the most distinctive terms and a sample of the tenders classified in each sub-category. I then group the topics by subject into 9 broad level-1 (L1) categories: roads, buildings, water supply, drainage and sewerage, irrigation and flood control, electrical works, industrial and mining, supply and procurement, and services and maintenance. Table F.4 displays the L1 categories with a representative tender and tender counts.

Table F.4: Product categories (level 1)

Category	Example tender	Tenders	%
Roads	Construction of Cement Concrete Road	4,347,526	27.7
Buildings	Construction of Anganwadi Building	3,621,021	23.0
Water Supply	Retrofitting of Existing Piped Water Supply Scheme	1,611,737	10.3
Services & Maintenance	Outsourcing of Manpower for Sub-Station Operation	1,374,767	8.7
Drainage & Sewerage	Improvement of Passage and Drainage	1,231,452	7.8
Supply & Procurement	Supply of Non-Expendable Medical Stores	1,041,338	6.6
Irrigation & Flood Control	Repair of Sluice Gate in Minor Irrigation Tank	972,708	6.2
Electrical Works	Repairing and Maintenance of Electrical Overhead Line Works	804,504	5.1
Industrial & Mining	Supply of Pump Spares at Refinery	673,914	4.3
Unclassified	Civil Work	36,487	0.2
Total		15,715,454	100.0

F.4 Bidder-level information: entity resolution

Bidder names are entered as free text when firms and contractors register for bidding. Hence, the same firm can appear under several spellings. This section describes how I standardize the raw bidder strings, group them into entities, and link them to external datasets.

F.4.1 Name standardization

I first standardize punctuation, case, common abbreviations, and misspellings:

- **Case and punctuation:** lowercase, replace punctuation with spaces, and collapse whitespace.
- **Honorifics and prefixes:** drop leading non-identifying words (`m/s`, `mr`, `dr`, `the`, ...) and standardize `shri/sri` to `shree` and the Hindu-Undivided-Family suffix to `huf`.
- **Company type:** map abbreviations and legal descriptors to a harmonized form (`private limited`, `(p) ltd` → `pvt ltd`; `limited` → `ltd`).
- **Misspellings:** correct a manual list of common typos (`consturction` → `construction`; `lndia` → `india`).
- **Empty values:** treat `nil`, `na`, `none`, `n/a`, and "`-`" as missing.

F.4.2 Entity resolution

Cleaned names are grouped into entities by fuzzy self-matching. Each pair of names is scored by averaging two different string-similarity measures: (i) a Jaro–Winkler similarity on the full cleaned strings (Jaro 1989; Winkler 1990) and (ii) a Levenshtein similarity (Levenshtein 1966) computed after deleting generic stopwords (legal-entity types and industry words such as `construction` or `traders`) and spaces. The entity resolution algorithm has three steps:

1. *Block on first letter.* For computational efficiency, only names sharing a first letter are compared.
2. *Accept high-similarity pairs.* To minimize false-positive matches, a pair is accepted only when the averaged score is at least 0.972 (cross-validated against a manual subsample).
3. *Combine into connected components.* Accepted pairs are linked transitively into connected components. Each connected component is a single entity. To reduce the number of false-positive matches, merges are blocked in two cases. First, when the two components carry differing sector keywords ("X Construction" and "X Consulting" are never merged, even if their strings are very similar). Second, any merge that would grow a component beyond 30 distinct original names is blocked.

For the analysis in the paper, bidders are identified at the bidder×district level. Blank entries receive a unique identifier and are never merged with one another or with a named firm.

The 1,359,439 distinct raw bidder strings reduce to 1,065,912 after name standardization and to 999,056 resolved entities after fuzzy matching. Table F.5 illustrates the name resolution with four firms that appear under several reported spellings and are recovered as a single entity.

Table F.5: Firms recovered from variant spellings: examples

Reported bidder strings	Resolved entity
LARSEN & TOUBRO LIMITED Larsen & Toubro Ltd Larsen and Toubro Limited M/S LARSEN and TOUBRO LIMITED	Larsen Toubro Ltd
A V ENGINEERS PRIVATE LIMITED A V ENGINEERS PVT LTD A.V.ENGINEERS PVT LTD A.V.Engineers Pvt Ltd	A V Engineers Pvt Ltd
JAY BALA JI CONSTRUCTION JAY BALAJI CONSTRUCTION M/S JAY BALAJI CONSTRUCTION COMPANY M/s Jay Balaji Construction	Jay Bala Ji Construction
D P CHATTERJEE AND CO D.P. Chatterjee & Co. D.P.Chatterjee&Co DP CHATTERJEE AND CO	D P Chatterjee and Co

Note: Each row shows several spellings of one firm as reported by procuring entities, all mapped by the resolution procedure to a single entity. The reported strings differ in case, punctuation, abbreviation of the legal-entity type, honorific prefixes (`M/s`), and word spacing.

F.4.3 Linking to external datasets

Resolved entities are linked to two external sources: the Prowess database of firm financials and the firms named in antitrust orders of the Competition Commission of India (CCI) (CCI 2026). The matching between datasets uses the same name cleaning as in Section F.4.1, and fuzzy string matching based on the same averaged string-similarity score of Section F.4.2.

I first generate candidates for the matching: for each external entity, I block on the first letter of the cleaned name and compute the averaged similarity score against every bidder entity. I retain the top two (Prowess) or three (CCI) bidder candidates above a similarity cutoff of 0.8 for review. Perfect matches are accepted automatically. Candidate pairs scoring between 0.8 and 0.99 are reviewed with an LLM (Gemini 3.1 Pro). A manually maintained blocklist of known false-positive firm names is used to prevent spurious matches. Since CCI orders name groups of firms that colluded, I keep the pair identifier to link co-bidding dyads to pairs that colluded in the past.

Table F.6 reports merge statistics with the two external datasets.

Table F.6: Linkage of resolved bidder entities to external datasets

	External firms			Resolved bidder entities	
External source	In source	Matched	%	Linked	%
Prowess (firm financials)	54,817	12,727	23.2	20,215	2.02
CCI antitrust defendants	544	291	53.5	334	0.03

Note: Each external source is linked to the resolved bidder entities by the procedure of Section F.4.3. *In source* is the number of distinct firms in the external dataset (each source's native firm identifier); *Matched* is how many are linked to at least one bidder entity, with the adjacent % that share. *Linked* is the number of the 999,056 resolved entities carrying the link, and the final % its share of all entities. The two shares have different denominators: coverage of the external list versus prevalence of the attribute among bidders. A single external firm may map to several bidder entities and vice versa, so *Matched* and *Linked* need not coincide.

F.5 Cleaning of financial values

The portals record financial values inconsistently. Amounts appear in different units (rupees, lakhs of 100,000 rupees, or crores of 10,000,000 rupees). In some contracts, bids are reported as percentages or as unit rates (instead of final amounts). Furthermore, the winning bid, the award amount, and the contract value can disagree. I clean three variables used in the analysis: the tender estimate, each bid, and the final contract value.

F.5.1 Tender estimate

For each tender, the procuring entity publishes an estimate of its value. Of the 11,345,235 tenders with at least one bid, 2,242,039 lack a usable estimate: either none is reported, or the reported value is implausible. I treat a value as implausible if it is non-positive, below 10,000 rupees, or implies an earnest-money deposit (EMD) above 20 percent of the estimate (652,610 tenders).

For these tenders I impute the estimate from the EMD, which is set as a fixed share of the estimate, typically 1, 2, 2.5, 3, 5, or 10 percent. I predict the applicable rate from tender characteristics observed independently of the estimate and divide the EMD by the predicted rate. The classifier is trained on tenders that report both quantities and recovers the correct rate in 98.2% of a held-out sample. This imputes 1,661,601 estimates; the remaining 580,438 tenders report no EMD and keep a missing estimate. In total, 10,764,797 tenders (94.9%) have an estimate, 15.4% of them imputed.

F.5.2 Contract value

Awarded tenders can report three measures of the contract value: the award-of-contract (AOC) amount, the sum of the award lines, and the winning bid. For 3,827,921 tenders the AOC amount and the award total agree exactly, and I use that value. Otherwise I apply the following rules in order:

First, if the AOC amount and the award total differ by more than 10 percent, I take the one within a factor of two of the tender estimate when the other is more than a factor of four away. Second, if the award total exceeds the AOC amount by exactly a Goods and Services Tax rate (5, 12, or 18 percent), the two differ only by tax, and I use the tax-inclusive award total. Third, if the winning bid equals one of the two, I use that value. Fourth, remaining disagreements are assigned the value closer to the estimate. If none of these applies, I use the award total, or the AOC amount if no award total is reported.

Overall, 6,293,928 tenders have a contract value. Of these, 3,827,921 require no reconciliation, 1,470,141 are resolved by the rules above, 34,362 rest on the fallback, and 908,530 take the winning bid of a reconstructed winner (Section F.5.4). A further 52,974 tenders without a recorded award receive a value after having their bids converted, mostly because they were rate contracts whose bids are rescaled to the estimate-implied total (Section F.5.3). The remaining 5,051,307 tenders have no contract value, mostly because they were never awarded or were cancelled.

F.5.3 Nonstandard bid formats

Percentage and formula bids. For some auctions, bidders bid percentage discounts or markups instead of the final monetary amount. In those cases, I rescale all bids in the tender to reflect monetary amounts. This affects 13,972 bids in 4,295 tenders.

Rate contracts and scale errors. A bid far below the contract value may be a genuine unit rate or an amount recorded at the wrong scale. Whenever the final contract value is a multiple of a power of ten of the bid value, the bid is treated as a unit error and rescaled. Otherwise, if the contract form specifies a unit-rate contract, bids are converted to totals using the best available contract, award, or estimate anchor. The procedure rescales 544,221 rate-contract bids and corrects 23,442 scale errors.

Financing tenders. Some tenders procure loans from private and government-owned banks. For these contracts, the quoted bid is an interest rate. I convert the rate to the present value of the implied interest stream using the reported principal and contract duration. This affects 959 bids.

F.5.4 Winner identification

Some awarded tenders do not record the winner. I assign the win to the bid with, in order of precedence, an accepted award status, a unique lowest-bidder marker in the bid remarks, or the only accepted financial bid. The winner is reconstructed this way in 776,295 tenders retained in the main sample. Because these tenders generally lack a recorded contract value, the contract value is set equal to the winning bid.

F.5.5 Bid-level usability

Some bids remain unusable after cleaning. A bid is marked unusable if:

- its value is non-positive
- after correction, it is more than 100 times above or below the final contract value (bids rescaled as rate, percentage, or unit-error corrections are exempt from this check)
- it is the winning bid and differs from the recorded award amount by more than a factor of ten
- it belongs to a sale of government property (e.g., scrap disposal), in which the highest bid wins
- it is a repeated-digit placeholder (at least five identical digits) and lies above three times or below 0.3 times the median of the other bids in the tender
- it is a rate-contract bid rescaled using the government estimate or the EMD, which would make its ratio to the estimate mechanical.

Unusable bids and bids without a recorded value do not enter price outcomes but are counted where financial bids are not needed.

Among the 4,594,861 tenders in the final price-analysis sample (Section F.5.6), 586,638 (12.8%) use an EMD-imputed tender estimate, 657,072 (14.3%) use the reconstructed winner's bid as the contract value, and 33,907 (0.7%) contain at least one usable bid imputed from the winning bid (when there are tied lottery participants), the contract value, or the award value.

F.5.6 Analysis samples

I construct the analysis samples from the 15,715,454 scraped tenders in several steps.

First, I keep tenders with usable bid records:

(i) the tender carries at least one recorded bid. This removes 4,370,219 tenders, around half of which never opened to bidding and the remainder of which were cancelled or retendered;

(ii) it is not an *empanelment* tender, in which bidders bid to be admitted to a panel of suppliers rather than to win a single contract, and in which multiple bidders are admitted at a common rate (30,159);

(iii) it is not a placement of surplus funds, in which a buyer distributes cash across banks that quote deposit rates and awards all of them (319);

(iv) it records at most 100 distinct bidders. Tenders with more bidders are rare outliers and make pair-level statistics expensive to compute (1,866).

This gives 11,312,891 usable tenders. The descriptive statistics in this appendix are computed on these tenders.

Panel A of Table F.7 constructs the *dyad scoring sample*, on which I compute the bidder-pair statistics used to estimate collusion. Because a pair must meet in a tender, I drop tenders with a single distinct bidder. This leaves 10,130,103 tenders. In these, I observe 7,045,102 distinct pairs of firms bidding on a common tender within a district. A pair is scored only if the two firms share at least 5 financial bids, below which the components of the index are shares of too few observations to rank, which leaves 2,043,939 scored dyads.

Panel B restricts Panel A's tenders to those recording a completed procurement in a standard format, which removes 4,385,047 tenders:

(a) 4,306,384 tenders are dropped because they lack an identified winner: 1,635,131 are not awarded (and have a technical-bid or technical-evaluation status), 2,214,593 are not awarded (and have financial-bid or financial-evaluation status), 359,237 were cancelled and 96,032 retendered after receiving bids, and 1,391 record an award whose winning bid cannot be identified (i.e., the tender was awarded without bidding)

(b) 31,272 tenders are dropped because they are from a highest-price auction in which the government sells a product

(c) 47,391 tenders are dropped because their status is Cancelled or Retender.

The 5,745,056 surviving tenders form the main sample. The price regression sample imposes further restrictions:

(d) it is not a *multi-award* tender, in which the award is split among several bidders (dropping 144,724 tenders);

(e) if its winner is reconstructed (Section F.5.4), the reconstructed winner's bid is the lowest bid; otherwise the record likely contains a portal-side entry error (23,333);

(f) its winner-bid-to-estimate ratio, where it can be formed, lies in $[0.30, 3.00]$ (116,414);

(g) its contract-to-estimate ratio is in $[0.30, 3.00]$ (13,866);

(h) at least one of its bids is usable and positive (Section F.5.5) (154,552);

(i) its estimate and district are recorded (195,714);

(j) it was not awarded by lottery (167,370);

(k) it has a positive contract value (6,967);

(l) at least two of its bids have a disclosed value (325,799);

(m) the district, organization, product, year, and month used for the fixed effects are recorded (1,456).

The price regression sample contains 4,594,861 tenders, or 45.4% of the dyad scoring sample and 40.5% of the tenders with at least one recorded bid.

Table F.7: Sample selection

Restriction	Observations	Δ	% of raw	% of scoring sample
Raw normalized tenders	15,715,454		100.0	
$\geq$ 1 recorded bid	11,345,235	-4,370,219	72.2	
Empanelment	11,315,076	-30,159	72.0	
Surplus-fund placement	11,314,757	-319	72.0	
More than 100 distinct bidders	11,312,891	-1,866	72.0	
Tenders with usable bid records	**11,312,891**		**72.0**	
Panel A: dyad scoring sample				
Single bidder (no pair observed)	10,130,103	-1,182,788	64.5	
Dyad scoring sample	**10,130,103**		**64.5**	**100.0**
at the level of firm pairs				
Co-bidding pairs observed	7,045,102			
Fewer than 5 shared financial bids	2,075,034	-4,970,068		
Firm with fewer than 5 financial bids overall	2,043,939	-31,095		
Scored dyads	**2,043,939**			
Panel B: price regression sample				
No identified winner	5,823,719	-4,306,384	37.1	57.5
Auction (government selling)	5,792,447	-31,272	36.9	57.2
Cancelled or Retender	5,745,056	-47,391	36.6	56.7
In main sample	**5,745,056**		**36.6**	**56.7**
Multi-award	5,600,332	-144,724	35.6	55.3
Uncorroborated reconstructed winner	5,576,999	-23,333	35.5	55.1
Bid/estimate $\in [0.30, 3.00]$	5,460,585	-116,414	34.7	53.9
Contract/estimate $\in [0.30, 3.00]$	5,446,719	-13,866	34.7	53.8
No usable positive bid	5,292,167	-154,552	33.7	52.2
Missing estimate or district	5,096,453	-195,714	32.4	50.3
Awarded by lottery	4,929,083	-167,370	31.4	48.7
No positive contract value	4,922,116	-6,967	31.3	48.6
Fewer than two financial bidders	4,596,317	-325,799	29.2	45.4
Missing fixed-effect variables	4,594,861	-1,456	29.2	45.4
Price regression sample	**4,594,861**		**29.2**	**45.4**

Notes: Each row applies one additional sample restriction. Δ is the number of observations removed at that step.

F.6 Expected number of bidders

Panel B of Table B.3 replaces the realized number of outside bidders with a prediction from tender observables. I fit a gradient-boosted classifier (scikit-learn's `HistGradientBoostingClassifier`) on the 5,982,467 main-sample tenders with at least one bid record, separately for all bidders and for financial bidders. All predictors are observed when the tender is published: location, organization, and product identifiers; the estimate, fee, and EMD; the submission window, corrigenda, bid covers, work duration, and bid validity; exemption, empanelment, auction, and multi-award indicators; the month; and the number of tenders in the same district, product category, and year with a near-

identical title.

I treat field size as categorical, with classes one to seven and eight or more (few tenders have more than eight bidders). The prediction of the model is the expected count $E[N \mid X] = \sum_k k\, P(k \mid X)$. The model is trained on the earliest 75 percent of tenders by publication date and evaluated on the most recent 25 percent.

Table F.8 reports how well the model predicts the realized number of bidders, compared to using the most common category (three bidders), or the most common category within each district × organization as a benchmark prediction. For all bidders, the model clearly beats the most common class (60.5 against 43.9 percent) but only slightly beats the most common class of the district × organization cell (59.1 percent). Hence, most of what is predictable about the number of bidders is a property of the market: beyond the usual field size of the district × organization cell, knowing tender characteristics does not supply additional predictive power.

Table F.8: Predicting the number of bidders

	All bidders		Financial bidders	
Forecast	Exactly	Within one	Exactly	Within one
Most common class	43.9%	78.6%	44.6%	77.8%
Most common class in district × organization	59.1%	79.7%	62.5%	83.4%
Gradient-boosted classifier	60.5%	80.9%	63.9%	84.5%

Notes: Share of held-out tenders whose bidder class is predicted exactly or within one class. Model evaluated for held-out tenders, the most recent 25 percent by publication date (1,495,167 tenders). Bidder counts are grouped into classes from one through seven and eight or more. Financial bidders also have a class of zero.

F.7 Linkage to PMGSY road records

The public tender data do not include project completion, quality, or final cost. However, these data are available for a subset of tenders related to the national rural roads program (PMGSY). The quality and completion outcomes in Table 9 are separately scraped from the PMGSY online management system (`omms.nic.in`). For each road, the system records an identifier (the package) under which it was awarded, the contractor, the sanction cost, the award date, the original and the actual completion date, and the grade from each inspection by the program's national and state quality monitors. The scraped records cover 203,466 road segments in 121,116 packages. A package is the unit of award and can contain several roads.

The PMGSY records do not carry the tender identifier, so I link packages to tenders in three steps. First, I generate candidate pairs within each state. Candidate pairs are package-tender pairs where (1) the package code equals the tender's reference number or title, (2) the package code appears in the tender's title or work description, or (3) the contractor's name equals the name of a bidder in the tender and the two records name the same district. Candidates whose award dates lie more than 180 days apart are dropped.

Second, I score each candidate on six measures of agreement: the gap between the two award dates, the ratio of the sanction cost to the tender estimate, the similarity of the contractor's name to the bidder's name, the text similarity of the road name to the tender's title and to its work description, and the distance between the road and the tender's location. I compute an agreement score as the weighted mean of the measures of agreement, with weights from the first principal component of

the measures among candidates linked by an exact package code. Third, among the remaining candidates, I keep those within 0.05 of the best score.

This procedure links 12,791 tenders to at least one package. Of these, 10,853 remain after the sample restrictions of Table F.7 and form the road sample of Table 9. For the linked records, the median gap between the two award dates is 14 days, the median distance between road and tender is 28.8 km, and the correlation between the log sanction cost and the log contract value is 0.72, with a median ratio of contract value to sanction cost of 0.87. Figure F.5 maps the linked roads.

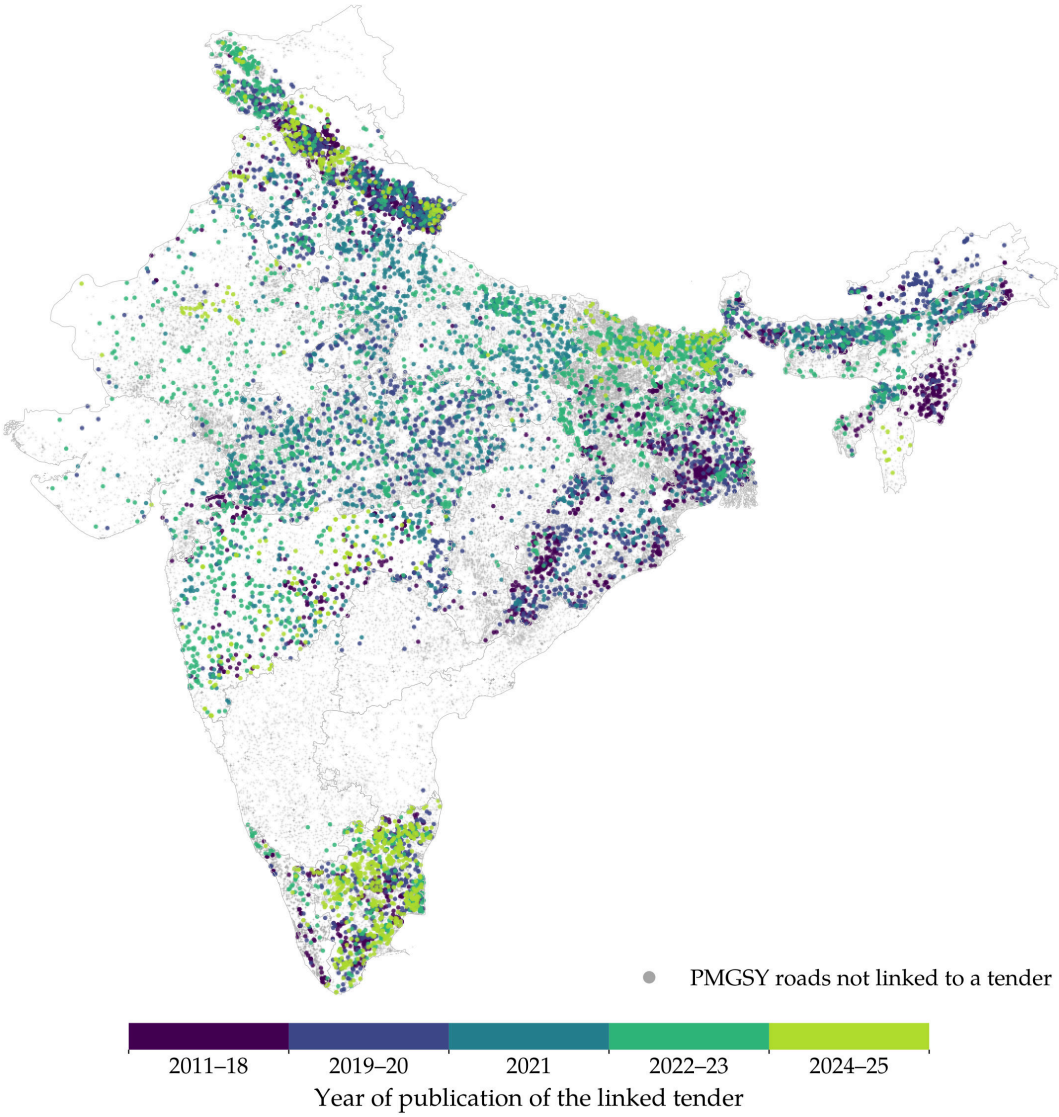


Figure F.5: Tenders linked to PMGSY road records

Note: Each dot is a geocoded PMGSY road. Gray: roads not linked to a tender. Colored: roads procured by a linked tender, colored by the tender's year of publication, grouped into five bins with roughly equal numbers of linked tenders.