A Field Experiment on Antitrust Compliance

Kei Kawai
University of California, Berkeley
& University of Tokyo & NBER

Jun Nakabayashi [*]
Kyoto University

September 7, 2023

Abstract

We study the effectiveness of firms' compliance programs by conducting a field experiment in which we disclose to a random subset of Japanese firms evidence of illegal bid-rigging. We find that the information that we disclose affects the bidding behavior of the treated firms: our test of bid-rigging fails to reject the null of competition when applied to the bidding data of the treated firms after the intervention. We find evidence that these changes are not the result of firms ceasing to collude, however. Our findings instead suggest that firms continue to collude even after our information disclosure intervention and that the changes in the bidding behavior we document are the result of active concealment of evidence by cartelizing firms.

Keywords: Regulatory Compliance, Scoring Auctions, Collusion.

[*]Contact information, Kawai: kawaixkei@gmail.com, Nakabayashi: nakabayashi.jun.8x@kyoto-u.ac.jp

1 Introduction

In many policy areas, there is an increasing trend towards delegating the day-to-day monitoring and enforcement of regulations to the regulated firms themselves (Sigler and Murphy, 1988, Ayres and Braithwaite, 1992).[1] This trend reflects the rapid growth in the scope, scale and complexity of government regulations without commensurate growth in regulatory resources.[2] While regulatory agencies retain the authority to investigate and intervene, in many policy areas, this authority is exercised only when firms exhibit repeated noncompliance or when firms' compliance functions are deemed ineffective. Thus, regulatory regimes often take the form of a hierarchical structure in which the firm's compliance function handles much of the routine regulatory violations, while the regulatory agencies, on the other hand, are primarily charged with overseeing the compliance function of the firms, taking direct action only under a limited set of circumstances.[3]

While the trend towards delegating monitoring and enforcement to firms have been especially salient for policy areas such as investor protection,[4] environmental protection,[5] and consumer protection,[6] this trend is also starting to affect policy areas such as antitrust which have traditionally had a strong emphasis on external enforcement.[7] For example, in some re-

[1]According to one estimate, regulatory compliance accounts for about 1.34% of the total wage bill of U.S. firms (Trebbi and Zhang, 2022).

[2]See e.g., Davis (2017) for a discussion of the rapid growth in government regulations in the U.S.

[3]The regulatory hierarchy is described in Ayres and Braithwaite (1995) as follows: "achievement of regulatory objectives is more likely when agencies display both a hierarchy of sanctions and a hierarchy of regulatory strategies of varying degrees of interventionism. The regulatory design requirement we describe is for agencies to display two enforcement pyramids with a range of interventions of ever-increasing intrusiveness (matched by ever-decreasing frequency of use). Regulators will do best by indicating a willingness to escalate intervention up those pyramids or to deregulate down the pyramids in response to the industry's performance in securing regulatory objectives."

[4]For example, the Bank Secrecy Act and the Dodd-Frank Act require financial institutions to designate a compliance officer who is responsible for implementing compliance programs within the firm. The compliance officer is required to certify to the regulators that the firm is in compliance with all regulations.

[5]See, e.g., the EPA's audit policy, formally titled "Incentives for Self-Policing: Discovery, Disclosure, Correction and Prevention of Violations". According to the EPA, the audit policy provides "several major incentives for regulated entities to voluntarily discover, self-report and correct violations of federal environmental laws and regulations", "making formal EPA investigations and enforcement actions unnecessary."

[6]For example, FTC and FDA send warning letters to firms to warn them of possible law violations, often with regards to advertising or marketing practices. The letter typically contains an explanation of why the recipient is being sent a warning letter.

[7]See, e.g., a series of papers written by Sokol, e.g., Sokol (2017, 2012)

cent antitrust violation cases, the U.S. DOJ has started "seeking court-supervised probation as a means of assuring that the company devises and implements an effective compliance program"(Assistant Attorney General Baer, 2014).[8] Implementation of antitrust compliance programs and designation of antitrust compliance officers have become key components of many consent decrees.[9] The DOJ's emphasis on effective compliance, treating it as a goal in and of itself, departs from the traditional regulatory model that focuses on ex-post punishment of violations.

Although compliance functions within firms have become an important part of antitrust enforcement, and of regulation of firms more generally, there is still relatively limited evidence on the effectiveness of within-firm compliance functions. In this paper, we provide empirical evidence on one aspect of regulatory compliance, specifically, the extent to which firms can take remedial action when confronted with evidence of illegal activity. How firms respond to evidence of regulatory violations – whether firms take steps to end wrongdoing, ignore evidence and continue to engage in wrongdoing, or actively seek to conceal incriminating evidence – can shed light on how best to incorporate within-firm compliance into the regulatory environment. It can also shed light on the effectiveness of various policy tools that seek to guide firm behavior without being legally binding, such as regulatory guidelines or warning letters sent by agencies such as the FTC and FDA. Many of these policy tools require at least some level of voluntary compliance from firms to be effective.

In order to study how firms respond to evidence of illegal activity, we conduct a field

[8]See, for example, the following remarks by Bill Baer, the Assistant Attorney General of the Antitrust Division on Sept 10, 2014.

> We also expect companies to take compliance seriously once they have pleaded guilty or have been convicted. Taking compliance seriously includes making an institutional commitment to change the culture of the company. Companies should be fostering a corporate culture that encourages ethical conduct and a commitment to compliance with the law.
>
> In such cases, the division will consider seeking court-supervised probation as a means of assuring that the company devises and implements an effective compliance program. We reserve the right to insist on probation, including the use of monitors, if doing so is necessary to ensure an effective compliance program and to prevent recidivism.

[9]See, e.g., US v. DirecTV Group Holdings and AT&T (2017), US v. Charleston Area Medical Center and St. Mary's Medical Center (2016), etc.

experiment in which we disclose to a set of construction firms in Japan evidence implicating them of illegal bid-rigging.[10] To do so, we first develop a statistical test of bid-rigging for scoring auctions in which allocation is determined by both price and quality of the proposal. We then apply the test to bidding data from auctions let by the Ministry of Land Infrastructure and Transportation in Japan. We identify 242 firms whose bidding behavior is inconsistent with competitive bidding. We then subject a random subset of these firms to a treatment in which we send out a letter that explains the statistical test we ran and the outcome of the test (i.e., competition is rejected) for the firm in question. We compare the subsequent bidding behavior of the treated and the control firms to identify how firms respond to evidence implicating them of bid-rigging.

Our first finding is that firms do respond to the information provided to them in the letter that we send out. In particular, we find that our statistical test of bid-rigging fails to reject the null of competition when applied to the bidding data of the treated firms after the intervention. For the control firms, our test continues to be effective at identifying bid-rigging behavior. Using Fisher's randomization test, we can reject (at the 1% significance level) the strong null hypothesis that the treatment induces no change in firm bidding behavior with respect to the ability of our statistical test to detect collusion.

Our second finding is that the change in firms' bidding behavior is likely to be the result of an adaptive response by the firms to evade detection without stopping collusion. First, we do not find significant changes in the level of bids or the quality of the proposals, which are often associated with a shift from collusion to competition. Moreover, we find a statistically significant decrease in the number of bidders submitting valid bids, or an increase in the proportion of bidders who submit bids above the reserve price conditional on participating. Because we study procurement auctions, bids above the reserve price have no chance of winning. These responses are consistent with continued collusive behavior. Finally, we document additional direct evidence of continued collusion using a test developed in Kawai et al. (2023). Our findings are consistent with concealment of incriminating evidence without

[10]Bid-rigging in procurement auctions is illegal in Japan and firms face fines of up to 10 % of all relevant sales in the affected market, in addition to the direct monetary gains that results from collusion.

stopping collusive behavior.

The results of our experiment suggest that, at least for the subset of construction companies participating in procurement auctions, existing levels of compliance capacity within colluding firms are unlikely to complement formal regulatory actions in achieving regulatory compliance.[11] These findings are somewhat different from previous studies that document firms stopping illegal behavior after allegations of wrongdoing are made (See, e.g., Christie et al., 1994 for the NASDAQ collusion case, and Monticini and Thornton, 2013, for the LIBOR manipulation case). Our findings suggest that, in the absence of widespread publicity, internal compliance functions may be ineffective in changing firm behavior. Our findings also suggest that regulators should be cautious in sharing evidence with firms when using informal channels to put pressure on them to change their behavior. If evidence of regulatory violations can be concealed with minimum cost and without having to stop illegal behavior, sharing detailed information with the firms may make detection and formal prosecution of illegal behavior more difficult in the future.

Related literature. There is a large literature in law and organizational behavior that analyzes how within-firm compliance functions complement the work of regulatory agencies, for example, Braithwaite (1985), Ayres and Braithwaite (1995), Parker (2002), Sokol (2012). There is also a small theoretical literature on regulatory compliance in economics that studies the trade-off between regulatory capture and efficient use of private information, e.g., Gehrig and Jost (1995) and Grajzl and Murrell (2007).[12] In these models, firms are better informed about the business environment than the regulators making firms better positioned to identify the types of regulations that are efficient. This informational asymmetry makes it potentially more efficient to delegate rule-setting to firms. However, given the obvious conflict of interest,

[11]It should be noted that in Japan, there is no equivalent of *qui tam* law suits in which private parties are compensated for assisting governments recover damages from illegal activities. For example, in the U.S., the False Claims Act allows private parties to bring suit on behalf of the government and receive a portion of the damages recovered. See, e.g., Kovacic (2001), and Engstrom (2013) for analysis on *qui tam* law suits. On the topic of private enforcement of laws more generally, see, e.g., Landes and Posner (1975) and Polinsky (1980).

[12]Relatedly, there is also a literature on self regulation, e.g., baron, Egorov, Harstad.

the regulator may not want to delegate rule-setting and enforcement entirely to the firms.[13]

Our paper is also related to papers that study firm adaptation in environments where regulatory agencies use screens to select the set of firms that go under scrutiny. For example, Wollmann (2019) and Cunningham et al. (2021) document evidence that firms adapt to the merger notification threshold set by the Hart-Scott-Rodino Act. The possibility of firm adaptation implies that the design of screens should account for the firms' equilibrium response. The importance of taking an equilibrium view of regulatory screening has been made previously by Cyrenne (1999), LaCasse (1995), Harrington (2004), Ortner et al. (2022), etc.

Previous work examining how firms react to evidence of incriminating evidence include Christie et al. (1994) and Monticini and Thornton (2013). Christie et al. (1994) document immediate changes in the quotes offered by market makers in the NASDAQ market after newspapers reported potential collusion by dealers. Monticini and Thornton (2013) find evidence of banks stopping underreporting of LIBOR rates after the Wall Street Journal reported potential manipulation of the rates. For both the NASDAQ collusion case and the LIBOR manipulation case, there was a lot of publicity created by the news reports. In contrast, the firms in our study were not subject to media exposure.

Lastly, our paper is related to the literature on detecting cartels in auctions. Early seminal work includes Hendricks and Porter (1988), Baldwin et al. (1997) and Porter and Zona (1993, 1999). More recent work includes Bajari and Ye (2003), Abrantes-Metz et al. (2006), Athey et al. (2011), Conley and Decarolis (2016), Schurter (2017), Kawai and Nakabayashi (2022), Chassang et al. (2022), Martin et al. (2022), Kawai et al. (2023) and Baránek et al. (2023).[14] The strategy for cartel detection we adopt in this paper extends ideas proposed in Kawai

[13]Other related work include Innes (1999), who studies remediation activities that are offered voluntarily by violators and Kaplow and Shavell (1994), who study self-reports of violations by perpetrating firms. There is also a literature that studies the efficacy of leniency programs and the incentives they create for firms to report their involvement in collusion, e.g., Motta and Polo (2003), Aubert et al. (2006), Spagnolo (2005), Chen and Harrington (2007), Harrington (2008a) and Miller (2009).

[14]Other related work includes Pesendorfer (2000), who studies bidding rings with and without side-payments, and Asker (2010), who studies knockout auctions among cartel members. Ohashi (2009) and Chassang and Ortner (2019) document how changes in auction design can affect the ability of bidders to sustain collusion. Clark et al. (2018) analyze the breakdown of a cartel and its implications on prices. For a survey of the literature, see Porter (2005) and Harrington (2008b).

et al. (2023) to scoring auctions.[15]

2 Institutional Background and Auction Format

Our paper analyzes the bidding behavior of firms that participate in auctions let by the Ministry of Land Infrastructure and Transportation (MLIT). This section provides a brief description of the institutional background and the auction format used by the MLIT.

MLIT is the largest procurement buyer in Japan, letting each year about 9,000 auctions for construction projects, worth a combined total of about 1.7 trillion yen (about $17 billion USD). The range of projects let by the MLIT include road paving, building and repairing bridges, installation of electrical equipment and other machinery as well as civil engineering work.

Since around 2006-7, almost all MLIT auctions are let using scoring auctions. In a scoring auction, each bidder, in addition to a price, submits a proposal which is converted into a scalar quality measure. Allocation is determined by each bidder's score, an index that combines both price and quality of the bid. The MLIT scores each bid according to the rule

$$s = q/p,$$

where q is the quality measure and p is the price submitted by the bidder. Bidders submit sealed bids simultaneously and the project is allocated to the bidder with the highest score, subject to the secret reserve price.

Depending on the type of project, the quality component of the bid can be more or less predictable. For simple projects, all bidders satisfying certain requirements receive the same quality measure. For these auctions, q is essentially fixed and allocation is determined by price competition. For more complicated projects, however, the quality of the proposals play a more important role in determining q. The range of quality measure a firm can obtain

[15] See e.g., Che (1993), Asker and Cantillon (2008, 2010), etc. for analysis of competitive equilibrium in scoring auctions.

is typically between $[100, 150]$. The lower bound of the quality measure is always fixed at $q = 100$, while the upper bound can be higher or lower depending on the complexity of the project.

One institutional detail that is worth mentioning is that when the price of bidder i exceeds the secret reserve price, quality of the bid is either not recorded entirely, or is assigned the lowest possible measure of 100. This fact explains certain features of the data pattern that we document in our analysis below.

3 Model and Test Statistic

This section specifies a model of scoring auctions and derives results that we use to test for collusion. Our model is similar to Che (1993), and Asker and Cantillon (2008, 2010). The results that we derive are similar to Kawai et al. (2023).

To preview the results, Corollary 1 states that, under the null of competition, if two bidders i and j almost tie for first place so that $s_i \approx s_j \geq s_k$ $(\forall k)$, the price components of the scores, p_i and p_j, must also be similar, $p_i \approx p_j$. This result should be intuitive: conditional on winning or losing by a very small margin, the winner and the loser should be similar, on average. Similarly, Corollary $1'$ states that, if i and j are almost tied for submitting the lowest price so that $p_i \approx p_j \leq p_k$ $(\forall k)$, the scores must also be similar, $s_i \approx s_j$. We use these results to test for competition: we test (1) whether or not the price component of marginal losers are the same as that of the marginal winners, on average; and (2) whether or not the scores of bidders who marginally bid the lowest are the same, on average as those of bidders who marginally bid above the lowest. If there are systematic differences, it will be evidence against competition.

Game form. A buyer procures a single item from a finite set N of potential suppliers. The procurement contract is allocated through a sealed-bid auction with a secret reserve price r, which is drawn from a distribution F_r. Each potential bidder $i \in N$ decides whether or not to participate in each auction. Conditional on participation, a bidder incurs participation cost $k > 0$, and submits a bid $\mathbf{b}$ consisting of a price-quality pair $\mathbf{b} = (p, q)$. Profit from non-participation is normalized to 0. A bid is valid if the price is below the reserve price, i.e., $p \leq r$. The bidder who submits a valid bid with the highest score is allocated the project, where the score is computed according to the formula $s = q/p$. Ties are broken with uniform probability. We denote by $\vee \mathbf{s}_t$ the highest score among participating bidders in auction t, by $\mathbf{s}_{-i,t} \equiv (s_{j,t})_{j \neq i}$ the scores of firms other than i, and by $\vee \mathbf{s}_{-i,t} \equiv \max_{j \neq i} s_{j,t}$ the highest score among i's participating competitors. Let $\vee \mathbf{s}_{-i} \prec s_i$ denote the event that bidder i wins the contract, i.e. s_i is the highest score and possible ties are broken in favor of bidder i. Bids are publicly revealed at the end of each period.

The bidder's profit conditional on winning is given by $p - C(q, \theta)$, where θ is the privately-known cost type of the bidder. The function $C(\cdot, \theta)$ represents the cost of providing higher quality for each type θ. In practice, the cost function C depends on observable auction characteristics as well, but we suppress this dependence. We assume that $C(\cdot, \theta)$ is increasing, convex and continuously differentiable for each θ.

Information. Each bidder i privately observes a signal $z_i \in Z_i$. We do not impose any assumptions on the distribution of the signal profile $\mathbf{z}_t = (z_{i,t})_{i \in N} \in Z = \prod_{i \in N} Z_i$. Signals may be arbitrarily correlated. The distribution of z_i can be asymmetric. We denote by $F_Z(\cdot)$ the joint distribution of the signals.

Cost types $\theta = (\theta_i)_{i \in N} \in \mathbb{R}^N$ are drawn independently conditional on each private signal z_i, i.e.,

$$\theta_i | z_i \sim \theta_i | \mathbf{z}_t, \theta_{-i}.$$

Bidder i's cost type does not provide information about the cost type of other bidders beyond the information already provided in the private signal z_i.[16] This class of information structures nests asymmetric independent private values, correlated values, and complete information. We denote by $F_\theta(\cdot | \mathbf{z}_t)$ the conditional distribution of the profile of cost types θ given signals $\mathbf{z}$.

[16]Because the signals are allowed to be correlated, $z_{i,t}$ helps bidder i predict the cost of other bidders.

Indirect Profit Function Because any bid (p, q) that satisfies $p/q = s$ guarantees the bidder the same winning probability, optimal bidding behavior requires a bid (p, q) with $p/q = s$ to be the solution to the following maximization problem:

$$\max_{p,q} \; p - \mathbb{E}[C(q, \theta)|z_i]$$
$$\textsf{s.t. } q/p = s,$$

where the expectation $\mathbb{E}$ is taken with respect to θ_i conditional on z_i.[17] The objective function, $p - \mathbb{E}[C(q, \theta)|z_i]$, is the profit the firm obtains if it wins. We denote the solution of this problem as $\pi_i(s, z_i)$. The function $\pi_i(s, z_i)$ corresponds to the indirect profit function of the firm when it bids a score of s and its signal is z_i. Note that $\pi_i(s, z_i)$ is continuously differentiable in s, by the Envelope theorem.[18]

Competitive Bidding We now derive implications of Bayes Nash equilibrium that we use to construct our test of competitive bidding. Our first result establishes that conditional on being a close winner or loser, any bidder believes that they win with probability greater than 50%.

Proposition 1. *For any bidder i and for any positive number $\eta > 0$, there exists ε such that*

$$\textsf{\textit{prob}}(i \textit{ wins} \mid z_i, \; |s_i - \vee \mathbf{s}_{-i}| < \varepsilon) \geq 1/2 - \eta.$$

Proof. Let $D_i(s_i|z_i)$ denote $\textsf{prob}(s_i \succ \vee \mathbf{s}_{-i}|z_i)$, the probability of winning the auction conditional on signal z_i. Incentive compatibility of the bidder implies

$$D_i(s_i|z_i)\pi_i(s_i, z_i) \geq D(s_i + \varepsilon|z_i)\pi_i(s_i + \varepsilon, z_i)$$

and

$$D_i(s_i|z_i)\pi_i(s_i, z_i) \geq D(s_i - \varepsilon|z_i)\pi_i(s_i - \varepsilon, z_i).$$

[17]If bidder i's cost, θ_i, is part of the private signal, z_i, the expectation would be unnecessary.

[18]Note that $\pi_i'(s, z_i) < 0$ for any s.

Noting that $D_i(s_i|z_i)$ must be strictly positive and continuous in equilibrium,[19]

$$
\begin{aligned}
\mathsf{prob}(i \text{ wins} \mid z_i,\ |s_i - \vee \mathbf{s}_{-i}| < \varepsilon) &= \frac{D_i(s_i|z_i) - D_i(s_i-\varepsilon|z_i)}{D_i(s_i+\varepsilon|z_i) - D_i(s_i-\varepsilon|z_i)} \\
&= \frac{1 - \frac{D_i(s_i-\varepsilon|z_i)}{D_i(s_i|z_i)}}{\frac{D_i(s_i+\varepsilon|z_i)}{D_i(s_i|z_i)} - \frac{D_i(s_i-\varepsilon|z_i)}{D_i(s_i|z_i)}} \\
&\geq \frac{1 - \frac{\pi_i(s_i,z)}{\pi_i(s_i-\varepsilon,z)}}{\frac{D_i(s_i+\varepsilon|z_i)}{D_i(s_i|z_i)} - \frac{\pi_i(s_i,z)}{\pi_i(s_i-\varepsilon,z_i)}} \\
&\geq \frac{1 - \frac{\pi_i(s_i,z_i)}{\pi_i(s_i-\varepsilon,z_i)}}{\frac{\pi_i(s_i,z_i)}{\pi_i(s_i+\varepsilon,z_i)} - \frac{\pi_i(s_i,z_i)}{\pi_i(s_i-\varepsilon,z_i)}} \\
&= \frac{\frac{1}{\pi_i(s_i,z_i)} - \frac{1}{\pi_i(s_i-\varepsilon,z_i)}}{\frac{1}{\pi_i(s_i+\varepsilon,z_i)} - \frac{1}{\pi_i(s_i-\varepsilon,z_i)}} \\
&= \frac{\frac{d}{ds}\left(\frac{1}{\pi_i(s_i,z_i)}\right) \times \varepsilon + o(\varepsilon)}{\frac{d}{ds}\left(\frac{1}{\pi_i(s_i,z_i)}\right) \times 2\varepsilon + o(\varepsilon)}
\end{aligned}
$$

As $\pi_i(s_i, z_i)$ is continuously differentiable and $\pi'(s_i, z_i) < 0$, the last term converges to 0.5 as $\varepsilon \to 0$. $\square$

Proposition 1 shows that the winning probability must not be much lower than 0.5 conditional on close bids.[20] As the next proposition shows, because at most one bidder can win, and because there are at least two close bidders conditional on the existence of close bids, it cannot be that firms' winning probability is strictly larger than 1/2. For any $\epsilon > 0$, let ϵ-close denote the event that the winning bid of the auction is within ϵ of the next highest score. For any $\varepsilon > 0$, let $\mathbb{E}[\cdot|\,\epsilon\text{-}\mathsf{close}]$ denote the expectation over $\mathbf{z}$ conditional on the event ϵ-close.

[19]Recall that participation costs k are strictly positive, which implies that $D_i(s_i)$ needs to be strictly positive. If $D_i(s_i)$ is not continuous at $s_i = s_0$, then there exists some bidder j who bids exactly s_0 with positive probability. In this case, all bidders other than j will have no incentive to bid in a small interval right below s_0, which in turn, makes bidding s_0 with positive probability suboptimal for bidder j. In other words, bidder j can gain strictly by bidding slightly below s_0. Hence, we have a contradiction.

[20]This proof is essentially the same as that of Proposition 1 of Kawai et al. (2023).

Proposition 2 (as-if random bids)**.** *For all $\eta > 0$ there exists $\epsilon > 0$ small enough such that*

$$\mathbb{E}\left[\left|\mathit{prob}(i \text{ wins} \mid z_i \text{ and } |s_i - \vee \mathbf{s}_{-i}| < \epsilon) - \frac{1}{2}\right| \,\Big|\, \epsilon\text{-}\mathit{close}\right] \leq \eta. \tag{1}$$

Proof. See Appendix. □

In words, Proposition 2 states that winning is as-if random conditional on close bids under competition. This result motivates our regression discontinuity test of competition. The following Corollary formalizes the link between equilibrium bidding and our regression discontinuity test.

Corollary 1. *For any bounded function $f : (p, q) \mapsto f(p, q) \in \mathbb{R}$,*

$$\lim_{\epsilon \searrow 0^+} |\mathbb{E}\left[f(p_i, q_i) \mid s_i - \vee \mathbf{s}_{-i} \in (0, \epsilon)\right] - \mathbb{E}\left[f(p_i, q_i) \mid s_i - \vee \mathbf{s}_{-i} \in (-\epsilon, 0)\right]| = 0.$$

In particular, if $f(p_i, q_i) = p_i$,

$$\lim_{\epsilon \searrow 0^+} |\mathbb{E}\left[p_i \mid s_i - \vee \mathbf{s}_{-i} \in (0, \epsilon)\right] - \mathbb{E}\left[p_i \mid s_i - \vee \mathbf{s}_{-i} \in (-\epsilon, 0)\right]| = 0. \tag{2}$$

.

Proof. See Appendix. □

The first part of Corollary 1 guarantees that, under the null of competition, the conditional expectation of $f(p_i, q_i)$, when bidder i is a marginal winner (i.e., $s_i - \vee \mathbf{s}_{-i} \in (0, \epsilon)$), is the same as the conditional expectation of $f(p_i, q_i)$ when bidder i is a marginal loser (i.e., $s_i - \vee \mathbf{s}_{-i} \in (-\epsilon, 0)$). In particular, setting $f(p_i, q_i) = p_i$, the second part of the Corollary 1 claims that the price of the marginal winner and the price of the marginal loser must be the same, in expectation. If, to the contrary, the conditional expectation of the marginal winner and the marginal loser is different, we can reject the null that bidder i is bidding competitively. We use expression (2) to test for competitive bidding.

Remarks on Corollary 1 We now make several remarks. First, there is a subset of auctions in which the quality component of the bidders are essentially fixed at some level $\overline{q}$. In other words, all qualified bidders are given the same quality level $\overline{q}$. For these auctions,

$$s_i - \vee \mathbf{s}_{-i} \to 0 \Leftrightarrow p_i - \wedge \mathbf{p}_{-i} \to 0,$$

where $\wedge \mathbf{p}_{-i}$ is the lowest price among bidder i's rivals. This implies that we can replace the conditioning set in expression (2) to be

$$p_i - \wedge \mathbf{p}_{-i},$$

when q is fixed at $\overline{q}$. Another set of auctions for which we can replace the running variable to be $p_i - \wedge \mathbf{p}_{-i}$ in Corollary 1 is when the equilibrium distribution of $\{p_i\}_{i \in N}$ is smooth. If the distribution of $\{p_i\}_{i \in N}$ is smooth, a modified version of Proposition 1 holds, i.e.,

$$\mathsf{prob}(p_i \text{ is lowest} \mid z_i,\ |p_i - \wedge \mathbf{p}_{-i}| < \varepsilon) \geq 1/2 - \eta. \tag{3}$$

Hence, a version of Corollary 1 with $p_i - \wedge \mathbf{p}_{-i}$ as the conditioning set must hold as well. The formal statement of the foregoing discussion is as follows:

Corollary 1′. *If the quality component can take only one value, $q_i = \{\overline{q}\}$ for all i, or if the distribution of $\mathbf{p}$ admits positive density on its support, then, for any bounded function $f : (p, q) \mapsto f(p, q) \in \mathbb{R}$,*

$$\lim_{\epsilon \searrow 0^+} \left| \mathbb{E}\left[f(p_i, q_i) \,|\, p_i - \wedge \vee p_{-i} \in (0, \epsilon) \right] - \mathbb{E}\left[f(p_i, q_i) \,|\, p_i - \wedge \mathbf{p}_{-i} \in (-\epsilon, 0) \right] \right| = 0.$$

In particular, if $f(p_i, q_i) = p_i/q_i$,

$$\lim_{\epsilon \searrow 0^+} \left| \mathbb{E}\left[s_i \,|\, p_i - \wedge \mathbf{p}_{-i} \in (0, \epsilon) \right] - \mathbb{E}\left[s_i \,|\, p_i - \wedge \mathbf{p}_{-i} \in (-\epsilon, 0) \right] \right| = 0. \tag{2'}$$

.

Proof. See Appendix. □

In our empirical exercise, we use both (2) and (2′) to test for competition.

In Corollary 1 and 1′, the conditional expectation is taken with respect to $f(p_i, q_i)$. It is straightforward to show that an analogous expression holds for any predetermined covariate x_i of bidder i (e.g., bidder's backlog). In Section XYZ, we exploit this fact by testing for discontinuities in the backlog between marginal winners and marginal losers. We find that marginal winners tend to have lower backlog than marginal losers, suggesting that firms in the treated groups continue to collude even after the intervention.

4 Identifying Firms That Do Not Bid Competitively

Illustration: Case of Nippo Corporation We now demonstrate how we test each firm for competitive bidding using the results obtained in the previous section. We start by focusing on bidding data from one firm, Nippo Corporation, between 2012 to 2015. Although the period 2012-2015 is not within the time frame we focus on, Nippo Corporation was involved in a bid-rigging scheme that was discovered and prosecuted in 2013 by the Japanese competition authority. By applying our test separately to the sample of auctions before and after the investigation, we illustrate the usefulness of the test in identifying collusion. We also discuss why collusive bidding is prone to fail the test.

Let T_k ($k \in \{\text{Before, After}\}$) denote the partition of the set of auctions in which Nippo Corporation participated to those taking place before and after the investigation. For each auction $t \in T_k$, let $i^*(t)$ be the winner of the auction, i.e., the bidder who had the highest score.[21] For each non-winner $i(\neq i^*(t))$ who participated in auction t, we construct $\Delta^s_{i,t} = \frac{s_{i,t}-s_{i^*(t),t}}{s_{i^*(t),t}}$ and $\Delta^p_{i,t} = \frac{p_{i,t}-p_{i^*(t),t}}{p_{i^*(t),t}}$. The variable $\Delta^s_{i,t}$ is the margin of defeat for bidder i, normalized by winner's score $s_{i^*(t),t}$. The variable $\Delta^p_{i,t}$ is the price difference between bidder i and the winner, again normalized by the winner's price, $p_{i^*(t),t}$. We then test for expression (2) by testing whether or not $\Delta^p_{i,t}$ converges to 0, as $\Delta^s_{i,t}$ converges to 0.

[21] There were no ties for the highest score in the sample.

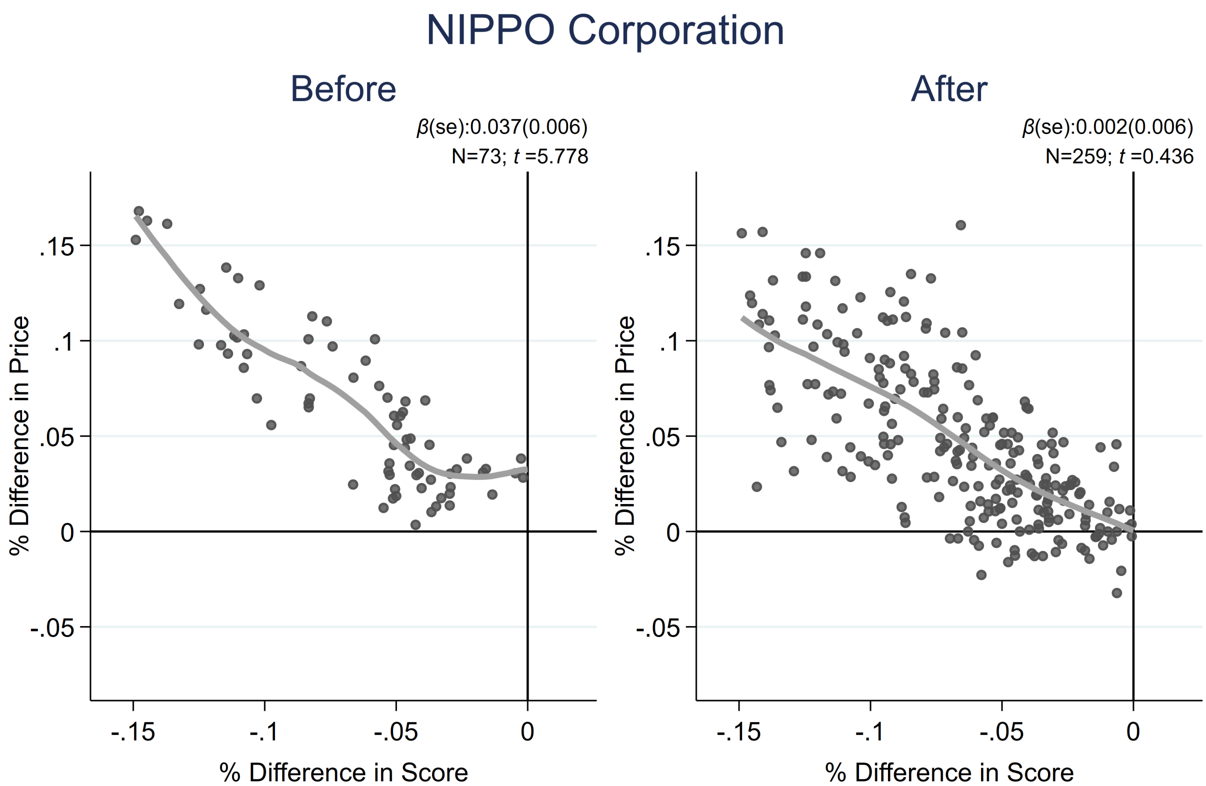

Figure 1: Binned scatter plot of $\Delta^s_{i,t}$ and $\Delta^p_{i,t}$. Left panel corresponds to the auctions in which NIPPO participated before before being investigated for collusion. Right panel corresponds to auctions after the investigation.

Figure 1 is a binned scatter plot of $(\Delta^s_{i,t}, \Delta^p_{i,t})$ for T_{Before} (left panel) and T_{After} (right panel). Figure 1 also plots the (nonparametric) mean regression of $\Delta^p_{i,t}$ on $\Delta^s_{i,t}$. Notice that, in the left panel, the mean regression does not pass through the origin. In other words, $\lim_{\epsilon \searrow 0^+} \mathbb{E}\left[\Delta^p_{i,t} | \Delta^s_{i,t} \in (-\epsilon, 0)\right] > 0$, so that $\Delta^p_{i,t}$ does not converge to 0 as $\Delta^s_{i,t}$ converges to 0. On the other hand, in the right panel of Figure 1, the mean regression goes through the origin.

Panel (A) of Table 1 reports our local linear estimates of $\mathbb{E}\left[\Delta^p_{i,t} | \Delta^s_{i,t}\right]$ at $\Delta^s_{i,t} = 0$ for T_{Before} (first column) and T_{After} (second column).[22] We find that, for the sample of auctions

[22]Our estimates are obtained using a local linear regression with a coverage error rate optimal bandwidth and a triangular kernel with a bias correction procedure as proposed in Calonico et al. (2014). We obtain standard errors by clustering at the auction level.

during which the firm was involved in bid rigging, the intercept of $\mathbb{E}\left[\Delta^p_{i,t}|\Delta^s_{i,t}\right]$ at $\Delta^s_{i,t} = 0$ is estimated to be 0.034, and statistically different from 0 at the 95% confidence level. Hence, we can reject the null hypothesis that $\lim_{\epsilon \searrow 0^+} \mathbb{E}\left[\Delta^p_{i,t}|\Delta^s_{i,t} \in (-\epsilon, 0)\right] = 0$ with 95% confidence. On the other hand, our estimate of the intercept is 0.002, and statistically indistinguishable from 0 for the sample of auctions after the investigation.

	(1)	(2)
	Before	After
Panel (A) : $\mathbb{E}\left[\Delta_i^p \mid \Delta_i^s = 0\right]$		
$\hat{\beta}$	0.037***	0.002
	(0.006)	(0.006)
Bandwidth $\hat{h}$	0.019	0.024
Obs.	73	259
Panel (B) : $\mathbb{E}\left[s_i - s_{i(1)} \mid p_i - p_{i(1)} = 0\right]$		
$\hat{\beta}$	−0.042***	−0.050
	(0.007)	(0.017)
Bandwidth $\hat{h}$	0.027	0.035
Obs.	73	211

Panel (A) corresponds to the sample of auctions in. *, **, and *** respectively denote significance at the 10%, 5%, and 1% levels.

Table 1: Intercept of Partial Linear Regression, NIPPO Corporation.

As we discussed at the end of Section 3, we also consider a second test of competition by comparing the score of the bidder who bid the lowest price with that of the bidder who bid the second lowest price. Under the null of competition, the average scores should be equal to each other when we condition only on those who bid almost the same.

In order to implement the test, let $i^\dagger(t)$ denote the bidder who submits the lowest price in auction t. For each $i \neq i^\dagger(t)$, define the running variable to be the price difference, $\tilde{\Delta}^p_{i,t} = \frac{p_{i,t} - p_{i^\dagger(t),t}}{p_{i^\dagger(t),t}}$ and the outcome variable to be $\tilde{\Delta}^s_{i,t} = \frac{s_{i,t} - s_{i^\dagger(t),t}}{s_{i^\dagger(t),t}}$. By construction, $\tilde{\Delta}^p_{i,t} > 0$. Figure 2 is a binned scatter plot of $(\tilde{\Delta}^p_{i,t}, \tilde{\Delta}^s_{i,t})$ for all bidders $i \neq i^\dagger(t)$ for T_{before} (left panel) and T_{before} (right panel). The left panel shows that bidders who bid marginally higher than the lowest price have much lower scores than $i^\dagger(t)$, the bidder who bid the lowest. This also means that these bidders have much lower quality than $i^\dagger(t)$. The right panel corresponds to the set of auctions after the investigation. We find that bidders who bid marginally higher than $i^\dagger(t)$ have, on average, the same score as $i^\dagger(t)$ in the right panel. The curves in the figure correspond to the nonparametric regression.

The bottom panel of Table 1 reports the estimate of $\mathbb{E}\left[\tilde{\Delta}^s_{i,t} | \tilde{\Delta}^p_{i,t}\right]$ at $\tilde{\Delta}^p_{i,t} = 0$. We find that the estimate is negative and statistically significant for T_{Before} while it is statistically indistinguishable from 0 for T_{After}.

Why collusive bidding may fail the test Corollary 1 and 1′ show that the regressions $\mathbb{E}\left[\Delta^p_{i,t} | \Delta^s_{i,t}\right]$ and $\mathbb{E}\left[\tilde{\Delta}^s_{i,t} | \tilde{\Delta}^p_{i,t}\right]$ should both converge to 0 as $\Delta^s_{i,t}$ and $\tilde{\Delta}^p_{i,t}$ converge to 0. The Corollaries do not, however, show that these conditions must fail under collusion. Here, we discuss why our test has power against certain types of collusive behavior.

Under most bid rigging arrangements, the ring preallocates projects so that one of its members is a designated winner and other members are designated losers. Often there is heterogeneity among the ring members in terms of their ability to submit high quality proposals. The quality of the proposals depends on the technology and machinery available to each firm. When the designated winner is a low quality firm, the designated losers will have to submit a price that is substantially higher than the designated winner since their

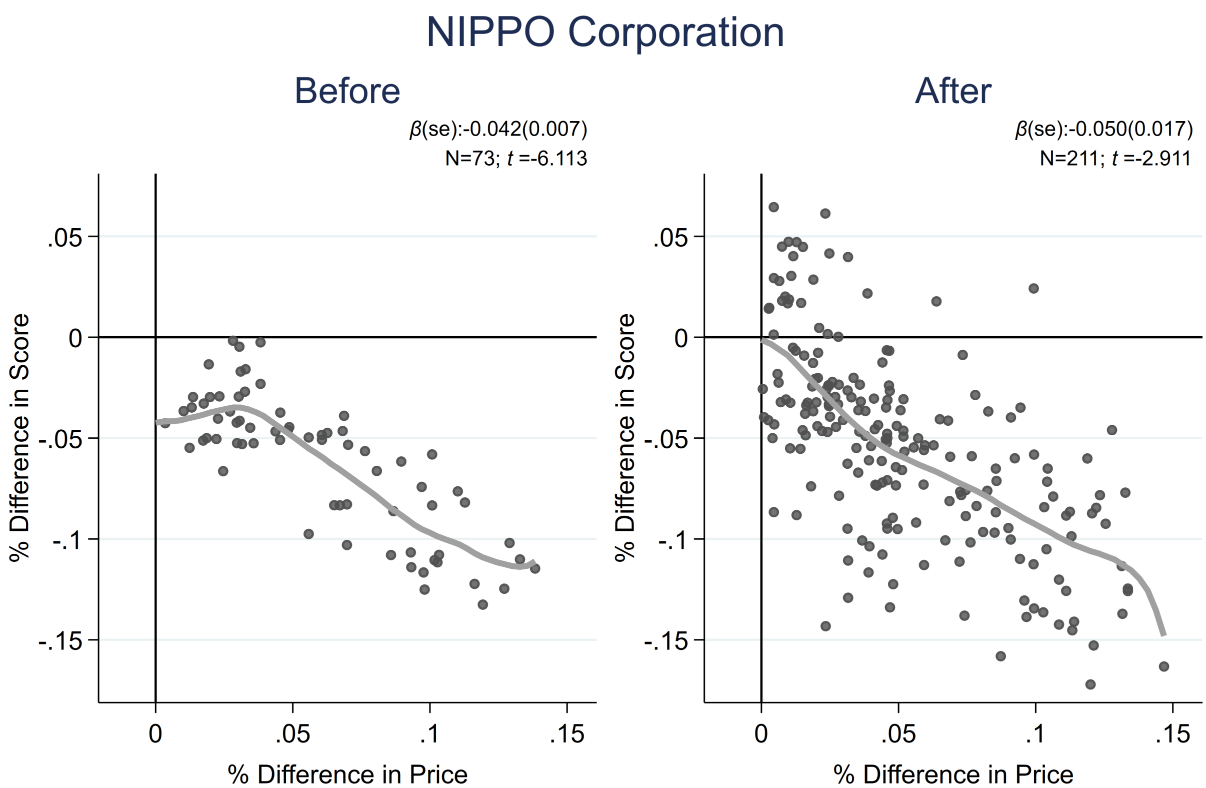

Figure 2: Nippo

proposals will likely receive a higher quality measure than that of the designated winner. For these auctions, the designated losers often come out as narrow losers, but with a much higher price and a much higher quality measure. The designated winner, on the other hand, comes out as the narrow winner, but with a much lower price and a much lower quality measure than the marginal losers. This bidding pattern is consistent with the data pattern illustrated in the left panel of Figure 1.

When the cartel designates a high quality bidder to be the winner of the auction, the low quality bidders can bid lower in terms of prices than the designated winner and still ensure that the auction is allocated to the intended bidder. In practice, however, many low quality bidders bid slightly higher than the designated winner, perhaps out of caution. When this is the case, bidders who bid marginally above the lowest bidder are often low quality designated losers whose overall scores are much lower than the designated winner. These types of bidding patterns will result in bidding patterns illustrated in the right panel of Figure 2.

Screening for non-competitive bidders For each firm that participated in at least 5 MLIT auctions between April 2015 and March 2017, we test the following two hypotheses

$$H_0 : \mathbb{E}\left[\Delta^p_{i,t} | \Delta^s_{i,t} = 0^-\right] = 0 \quad \text{v.s.} \quad H_1 : \mathbb{E}\left[\Delta^p_{i,t} | \Delta^s_{i,t} = 0^-\right] > 0$$

$$\text{and}$$

$$H_0 : \mathbb{E}\left[\tilde{\Delta}^s_{i,t} | \tilde{\Delta}^p_{i,t} = 0^+\right] = 0 \quad \text{v.s.} \quad H_1 : \mathbb{E}\left[\tilde{\Delta}^s_{i,t} | \tilde{\Delta}^p_{i,t} = 0^+\right] < 0$$

during that time period. Based on the two tests, we initially select 1,143 firms that fail either of the tests at 5% significance.[23] We then further screen these firms by visually inspecting the binned scatter plot and determine whether or not the intercept seems clearly different from zero. The reason for adopting this two-step procedure instead of mechanically selecting the firms with the highest t-statistics is because there is occasionally a disconnect between

[23]For the first test, we select firms with t-statistics of above 1.65. For the second test, we select those with t-statistics of below -1.65. Of the 1,143 firms that we initially select, 635 firms failed the first test, 601 firms failed the second test, and 93 failed both.

the magnitude of the t-statistic and how strong the evidence appears based on the binned scatter plot. For this reason, we choose a threshold of 5% confidence in terms of the t-statistic to initially preselect the set of potentially non-competitive firms. We then follow up with a visual inspection of the binned scatter plot of each firm to avoid making type I errors.[24] Ultimately, we end up with 240 firms which we feel confident classifying as noncompetitive.[25]

Because we conduct tests for all firms in the data that bid at least 5 times, a simple t-test with a significance level of 5% is going to reject about 5% of the firms even in the absence of collusion. In order to get a sense of the *false discovery rate* (FDR), or the proportion of competitive firms that we wrongly identify as uncompetitive, we plot the empirical CDF of the estimated t-statistic from our first test (i.e., test of $\mathbb{E}\left[\Delta^p_{i,t}|\Delta^s_{i,t} = 0^-\right] = 0$) along with the CDF of the standard normal in Figure 3. Note that under the null of competition, our test statistic follows the standard normal distribution. Figure 3 shows clearly, however, that the empirical CDF exhibits large excess mass at the right tail of the distribution relative to the standard normal: for example, 15% of the estimated t-statistics are higher than 1.65 as opposed to 5%. This implies that the (positive) FDR at critical value of 1.65 is at most $0.05/0.15 \approx 33\%$.[26] Using empirical Bayes shrinkage to shrink the distribution of the t-statistic only changes the fraction of t-statistic exceeding 1.65 to 14%.

As we discussed above, we use a cutoff of 1.65 in terms of the t-statistics to determine the initial set of noncompetitive firms. We then visually inspect the binned scatter plot for each of these firms to further narrow down the set to about 240. If, instead of visual inspection, we were to mechanically pick the same number of firms with the largest t-statistics, the corresponding FDR, computed using Benjamini et al. (1995), would be about 2%. To the extent that our visual inspection is no worse at screening out type I errors than mechanically choosing the firms with the highest t-value, the FDR among the chosen 240 firms is lower

[24] See Korting et al. (2023) for a discussion of the benefits of combining a formal econometric test with a visual examination to achieve lower type I errors.

[25] We initially select 242 firms, but we end up dropping 2 firms after performing the clustering procedure we discuss next section. See Online Appendix Section **??** for details.

[26] The original definition of the FDR proposed by Benjamini and Hochberg (1995) is $\mathbb{E}[V/R|R > 0]\,\Pr(R > 0)$, where R is the number of total rejections and V is the number of Type 1 errors. The positive FDR is defined to be $\mathbb{E}[V/R|R > 0]$ (See, e.g., Storey, 2002).

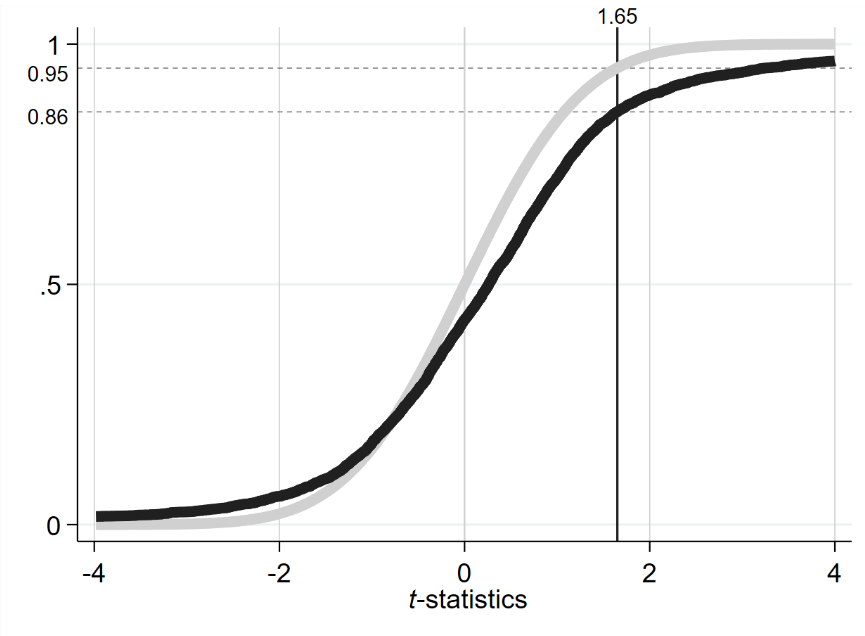

Figure 3: Histogram of t-statistic for $\mathbb{E}[\Delta^p_{i,t}|\Delta^s_{i,t} = 0^-]$.

than 2%.[27]

Finally, we note that while the selection of the 240 firms are based on somewhat subjective criteria, all of the subsequent statistical analysis are based on randomization inference which gives us exact p-values. Regardless of whether the 240 firms were selected subjectively or based on objective measures, our inference is valid because the treatment is randomized given this sample.

5 Experimental Design

Assignment of Treatment The 240 firms that we identify as noncompetitive in the previous section often bid on the same auctions and hence we expect some communication to take place among these bidders. In order to contain potential treatment spillovers across units, we partition the 240 firms by grouping those that frequently bid together using a

[27]Korting et al. (2023) conclude that visual inference achieves a type I error rate that is slightly lower than the Imbens and Kalyanaraman (2012) and Calonico et al. (2014) procedures. Moreover, they find that visual and econometric inferences are complementary.

clustering algorithm.[28] The resulting partition has the property that firms within each group bid on the same auction frequently while firms in different groups rarely do. We end up with 26 groups of firms. We then construct 13 matched pairs based on the group's geographical location, type of work (e.g., landscaping, paving, etc.), and the number of firms in the group.

Finally, we assigned treatment status with rerandomization (see, e.g., Morgan and Rubin, 2012) to achieve balance between the treatment and the control with regards to the mean winning bid, mean of the t-statistics of the two tests, and the average number of auctions firms participate.[29] We take the effect of rerandomization into account when conducting our statistical tests below.

Treatment We send (physical) letters to 13 groups, or a total of 106 firms on Feb 12, 2019. We send the letters to the addresses of the firms as recorded in the MLIT's registry. In the letter, we first explain that we have developed a screen for bid rigging and that we are exploring its usefulness and applicability. We explain the mechanics of our test by walking through Figures 1 and 2 of Nippo Corporation in the letter. We then include corresponding figures for the firm in question and discuss the similarities between the firm's bidding patterns with those of Nippo Corporation before the investigation, i.e., the regression lines do not pass through the origin. We also include a list of auctions used to analyze the firm's bidding pattern so as to emphasize that our analysis is specific to the firm. Note that there is a lag between the period of analysis contained in the letter (April 2015 to March 2017) and the date the letters are sent out (Feb 12, 2019). Finally, we ask the firm whether the firm is aware of various screening methods to detect noncompetitive bidding, and whether such screens can help the firm improve antitrust compliance. We include a return envelope for the firm to send back its reply, asking them to do so by March 15, 2019. A copy of the letter sent to the firms in the treatment groups is in the Online Appendix.

[28]In particular, we use a hierarchical agglomerative linkage procedure. We provide the details in Online Appendix Section **??**

[29]We explain in more detail the rerandomization procedure in the online appendix.

Baseline Summary Statistics Table 2 and 3 report the summary statistics of the auctions and the bidders for the sample of 240 firms that we select as noncompetitive. The summary statistics correspond to the sample of auctions we used to detect noncompetitive bidding, i.e., auctions let between April 2015 to March 2017. The average reserve price of the auctions is about 128 million yen, or about $1.3 million. The winning bid is about 94% of the reserve price and the average number of bidders is 4.81. The fact that there are only very minor differences between the treatment and the control is by design. We assigned treatment by a matched pairs design and, moreover, we rerandomized the assignment to achieve balance for the winning bid.

Table 3 reports firm characteristics. We report the mean annual sales, profits, the number of engineers employed by the firm, and the t-statistic corresponding to our intercept estimate. The first three variables are obtained from the registry maintained by the MLIT. Although the total number of employees is not recorded in the registry, the ratio of engineers to the total number of employees is usually about 1:2 to 1:3, based on the subset of firms that report the number of employees on their web pages. Given that the average number of engineers is about 27, this implies that the total number of employees is likely to be about 55 to 80, on average. The average t-statistics associated with the intercept estimates of each firm before the intervention is about 3.5 for the first test (i.e., test of $\mathbb{E}\left[\Delta^p_{i,t}|\Delta^s_{i,t}=0^-\right]=0$), and it is about -3.0 and -3.5 for the second test (i.e., test of $\mathbb{E}\left[\tilde{\Delta}^s_{i,t}|\tilde{\Delta}^p_{i,t}=0^+\right]=0$). By construction, for each firm, the t-statistics are higher than 1.65 for the first test or lower than -1.65 for the second test.

6 Results

Changes in the bidding behavior We first document changes in the firms' bidding behavior in relation to the ability of our test to detect collusion. Figures 4 and 5 plot

	Treatment	Control
Reserve	128.358	128.787
	(72.350)	(72.432)
WinBid / Reserve	0.939	0.939
	(0.034)	(0.034)
Bid / Reserve	0.942	0.942
	(0.039)	(0.039)
Quality	155.722	155.746
	(6.836)	(6.839)
Number of Bidders	4.794	4.794
	(2.792)	(2.799)
N	1,300	1,289

Note: The table shows the summary statistics of auctions let during fiscal years 2015 through 2017. Reserve is reported in millions of yen. Standard errors are in parenthesis.

Table 2: Summary Statistics (Auctions)

	Treatment	Control
Annual Sales	2,087.12	2,122.89
	(2,278.15)	(3,623.24)
Annual Profits	146.52	137.95
	(192.54)	(297.86)
# Engineers	26.75	27.63
	(18.97)	(34.29)
t-stat ($\mathbb{E}\left[\Delta^p_{i,t})\right]$)	3.45	3.75
	(4.09)	(8.69)
t-stat ($\mathbb{E}\left[\Delta^s_{i,t})\right]$)	-3.00	-3.54
	(4.39)	(4.97)
N	107	133

Note: Sales and profits are reported in million of yen. There are 240 firms in our sample. We could not get the data on the annual sales and profits, and the number of engineers for one firm in the control group. The reported numbers for these variables for the control group are based on the averages of 132 firms.

Table 3: Summary Statistics (Firms)

(Δ^p, Δ^s) for non-winners before and after the intervention.[30] Figure 4 corresponds to the

[30] The sample used for the left panel (before period) consist of those let between April 1, 2015 and February 15, 2019, the date we sent out the letters. The sample for the right panel (after period) consist of those let

treatment group and Figure 5 corresponds to the control group. Comparing the left and right panels of Figure 4, we find that the intercept of the regression line (depicted in gray) at $\Delta^s = 0$ is positive in the left panel while it is close to 0 in the right panel. This implies that marginal losers stop bidding substantially higher than the marginal winner in terms of prices after the intervention, suggesting that bidder behavior changed after the intervention for the treated firms. For the control group, the regression lines illustrated in Figure 5 seem to intersect the y-axis above the origin in both panels.

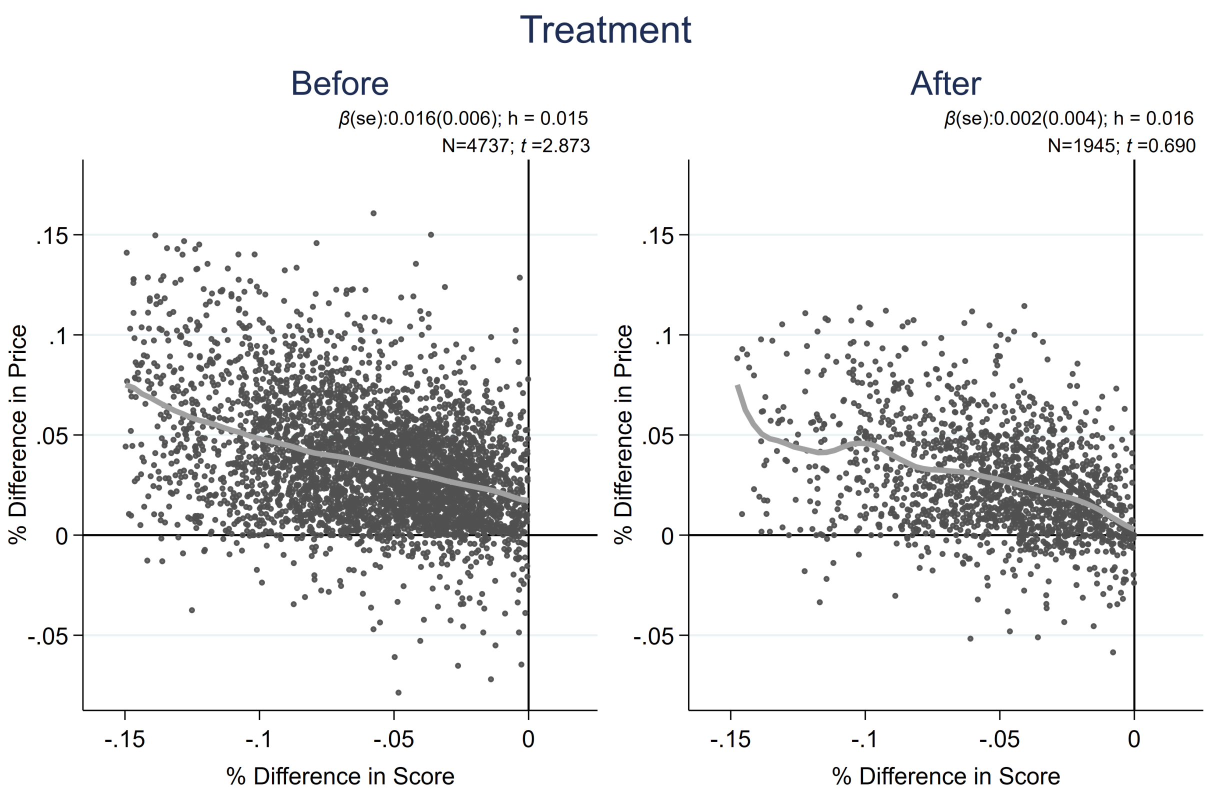

Figure 4: $\mathbb{E}[\Delta^p \mid \Delta^s]$ for Treatment Firms

Figures 6 and 7 plot the binned scatter plot of $(\tilde{\Delta}^p, \tilde{\Delta}^s)$. For the treatment sample (Figure 6), we find that the estimated intercept of the the regression line changes from -0.017 before the intervention to -0.006 after the intervention, while for the control sample (Figure 7),

between March 15, 2019, the date by which we asked the firms to respond to our survey and March 31, 2021.

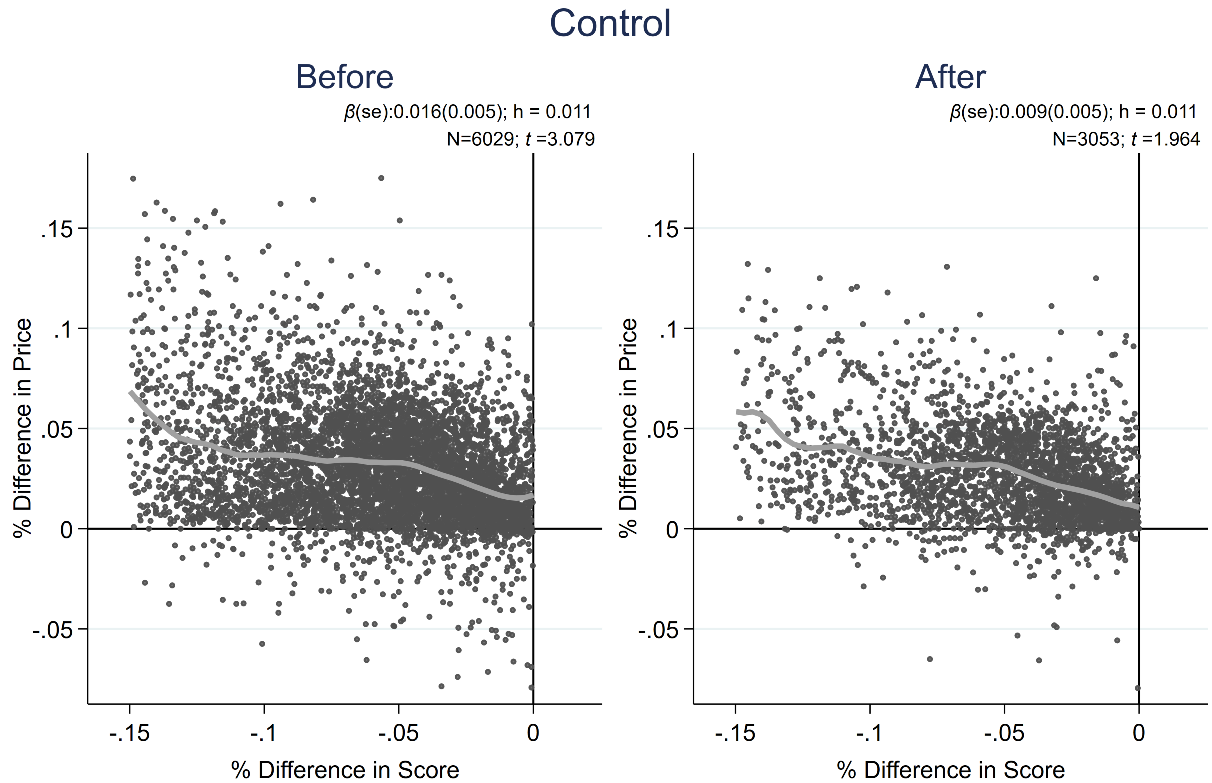

Figure 5: $\mathbb{E}[\Delta^p \mid \Delta^s]$ for Control Firms

the estimate does not change much: The estimate is -0.016 before the intervention, and it is -0.014 after.

While the figures are suggestive that the intervention induced changes in the bidding behavior of at least some firms, we need to ascertain that these changes are unlikely to be the result of chance. In order to do so, we now run a Fisher's test of the strong null hypothesis (See, e.g., Imbens and Rubin (2015) for a textbook treatment).

Specifically, for each group $g \in \{1, \cdots, 26\}$, let Y_g denote the change in the t-statistic before and after the intervention,

$$Y_g = t_g^{\mathsf{After}} - t_g^{\mathsf{Before}},$$

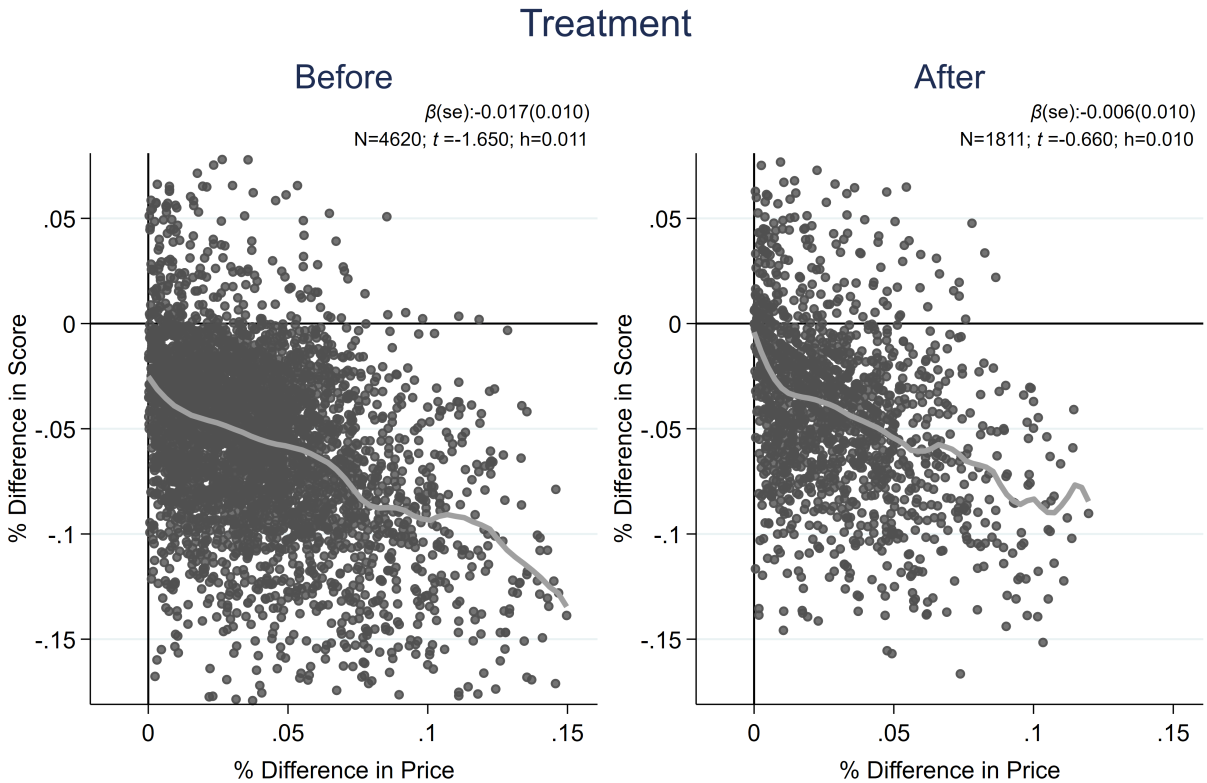

Figure 6: $\mathbb{E}[\tilde{\Delta}^s \mid \tilde{\Delta}^p]$ for Treatment Firms

where t_g^{Before} and t_g^{After} are the t-statistic of the estimated intercept, computed using the sample before the intervention and after the intervention. Now, for any partition of the groups G and G', consider the difference in the average Y_g between groups in G and groups in G':

$$\overline{Y}_{G-G'} = \frac{1}{\mid G \mid} \sum_{g \in G} Y_g - \frac{1}{\mid G' \mid} \sum_{g \in G'} Y_g.$$

The statistic $\overline{Y}_{G-G'}$ corresponds to a measure of the extra decline in the t-statistics exhibited by groups in G relative to G'. If we set $G = G_T$ and $G' = G_C$, where G_T and G_C are the set of treated groups and control groups, $\overline{Y}_{G_T - G_C}$ corresponds to the extra decrease in the t-statistic for the treatment groups relative to the control groups. Note, however, that we can compute $\overline{Y}_{G-G'}$ for arbitrary partition G and G', not just for G_T and G_C.

Under the strong null that our intervention had no effect on the bidding behavior of any

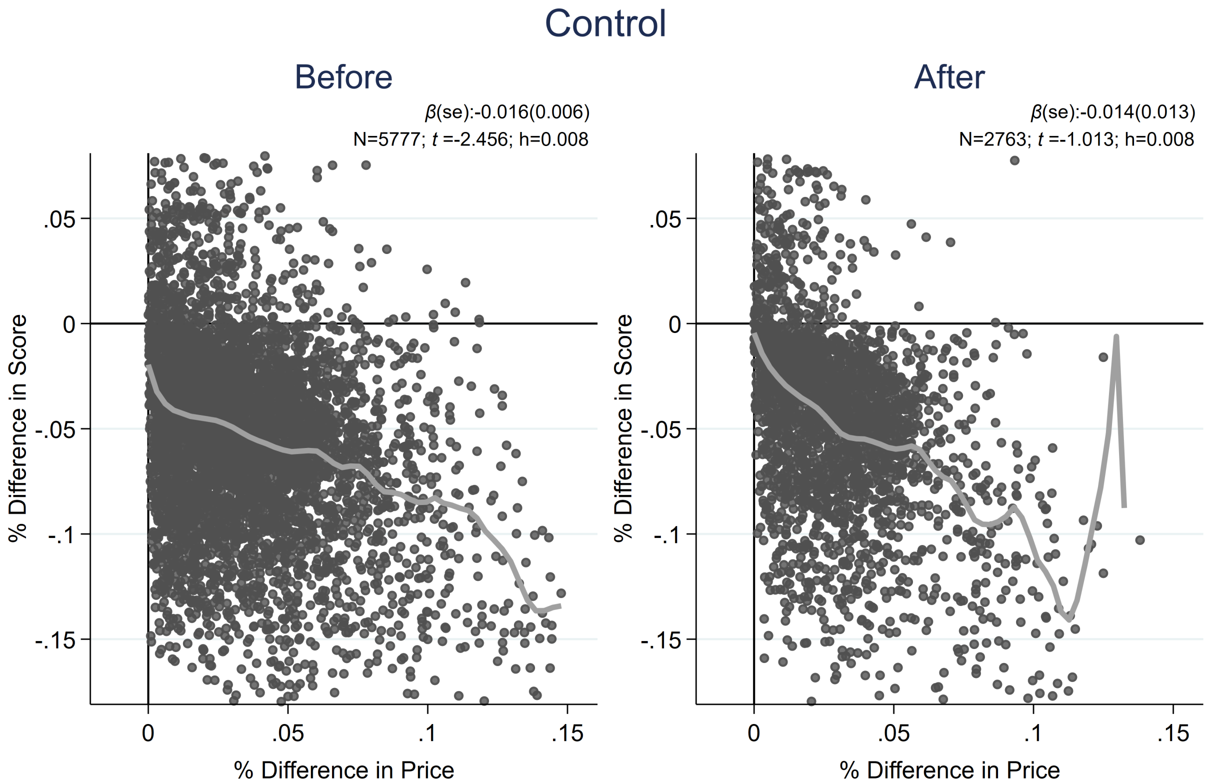

Figure 7: $\mathbb{E}[\tilde{\Delta}^s \mid \tilde{\Delta}^p]$ for Control Firms

of the treated groups whatsoever, the partition of the groups into the set of treated groups, G_T, and the control groups, G_C, is of no particular significance. In other words, under the strong null of no effect, the random variable $\overline{Y}_{G_T - G_C}$, should have the same distribution as $\overline{Y}_{G-G'}$ for arbitrary partition G and G'. Fisher's randomization test compares the realized value of $\overline{Y}_{G_T - G_C}$ to the distribution of $\overline{Y}_{G-G'}$ for all possible partitions G and G'. If the value of $\overline{Y}_{G_T - G_C}$ is extreme relative the distribution of $\overline{Y}_{G-G'}$, we can reject the null that the treatment has no effect.

Figure 8 reports the distribution of $\overline{Y}_{G-G'}$ for all possible partitions of G and G' as well as the value of $\overline{Y}_{G_T - G_C}$. The left panel corresponds to the t-statistic of $\mathbb{E}\left[\Delta^p_{i,t} | \Delta^s_{i,t} = 0^-\right]$ and the right panel corresponds to $\mathbb{E}\left[\tilde{\Delta}^s | \tilde{\Delta}^p = 0^+\right]$. The number of possible partitions are a subset of $2^{13} (= 8,192)$ that satisfy the rerandomization criteria.[31] The left panel corresponds

[31]Recall that we have a matched pair design in which one of the pair is treated and the other is not. The

to the distribution of $\overline{Y}_{G-G}$. The realization of $\overline{Y}_{G_T-G_C}$ is marked as a vertical line in the figure. Since the realization of $\overline{Y}_{G_T-G_C}$ is far off from realizations of $\tau_{G-G'}$ for other possible partitions (p-value of 0.031), we can reject the null that the intervention had no effect on the treated groups. We plot the corresponding figure for the case in which the underlying t-statistic is computed using $\tilde{\Delta}^p$ as the running variable and $\tilde{\Delta}^s$ as the outcome variable. For reasons noted in the previous section, we do not find that the value of $\overline{Y}_{G_T-G_C}$ is different from realizations of $\overline{Y}_{G-G'}$ for other partitions. Hence, we cannot reject the null that firms changed the bidding behavior relative to the power of the second test to detect collusion.

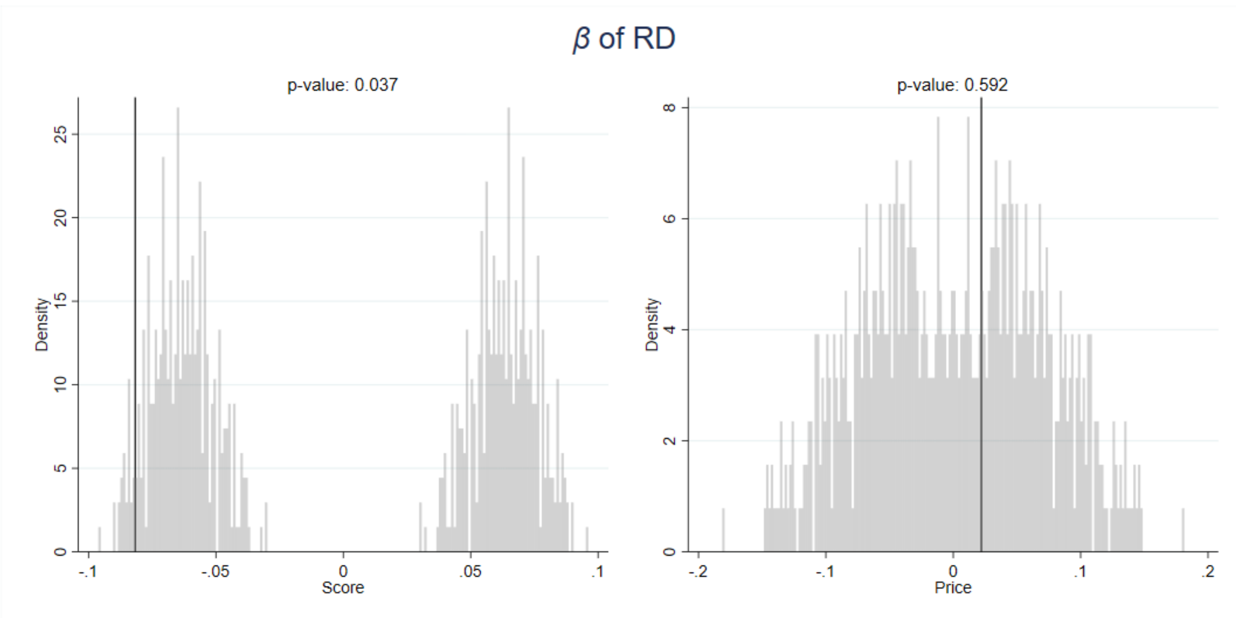

Figure 8: β of RD

Changes to other outcome variables In this section, we explore the effect of the intervention on other auction outcomes to understand whether the changes in the bidding pattern that we document in the previous section are the result of cartel firms stopping to collude or it is the result of concealment of incriminating evidence with continued collusion. Figure 9 plots the time series average of the winning bid (Left panel) and the average losing bids

number of all possible treatment assignments is 2^{13}. Because we rerandomize to maintain balance between the treated and the control groups, we have a smaller number () of possible assignments.

(Right panel) separately by treatment and control groups. We normalize the bids by the reserve price to make them comparable across auctions. The light colored line corresponds to the control group, the dark line corresponds to the treatment group, and the solid vertical line corresponds to the date of the intervention. We find that both for the treatment and the control groups, the bids are quite stable across time, at around 95% (of the reserve price) for the winning bid and at around 97% for the losing bids. In particular, we do not see any breaks in the bids for the treatment group around the time of the intervention.

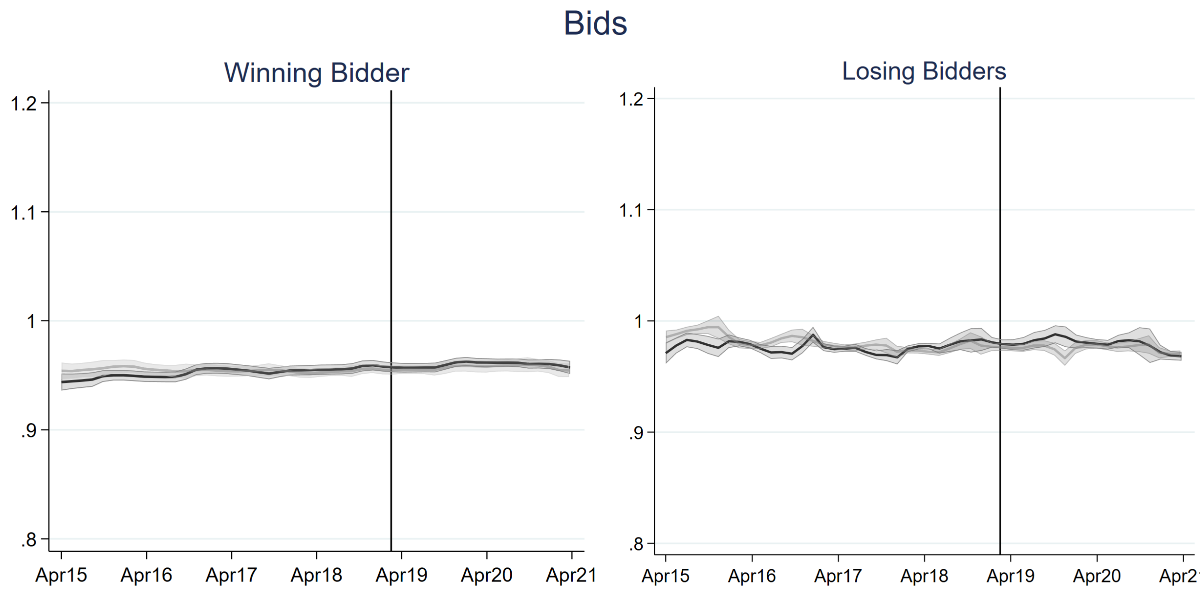

Figure 9: Time Series Plot of Bids, Normalized by the Reserve Price. Left panel corresponds to the winning bid (divided by the reserve price) and the right panel corresponds to the losing bids (divided by the reserve price). Light colored lines correspond to the control group and the dark lines correspond to the treatment group.

Since breakdowns in cartels typically result in significant drops in the price, Figure 9 suggests that the firms in the treatment groups are unlikely to have stopped colluding after the intervention.

Figure 10 plots the time series average of the quality measure of the winning bidders (left panel) and the losing bidders (right panel). The dark lines correspond to the treatment and the light colored lines correspond to the control. Although there do not seem to be significant changes in the quality measure for the winning bidders, there is a substantial drop in the quality for the treatment group around the time of the intervention for the losing bidders.

As we mentioned in Section 2, firms that submit bids exceeding the secret reserve price are assigned a quality of 100 points, which is the lowest possible points attainable. In order to study the extent to which the changes in the right panel of Table 10 are driven by bids that exceed the reserve price, the left panel of 11 plots the time series average of the probability that a losing bidder submits a bid that exceeds the reserve price.[32] The left panel of the figure suggests that there is a large increase in the probability of submitting a bid above the reserve price for the treatment group after the intervention.

The right panel of 11 plots the time series average of the number of bidders that submit a bid below the reserve price. The right panel of the figure suggests that the number of bidders that submit bids below the reserve price falls for the treatment group after the intervention, consistent with the results of the left panel.

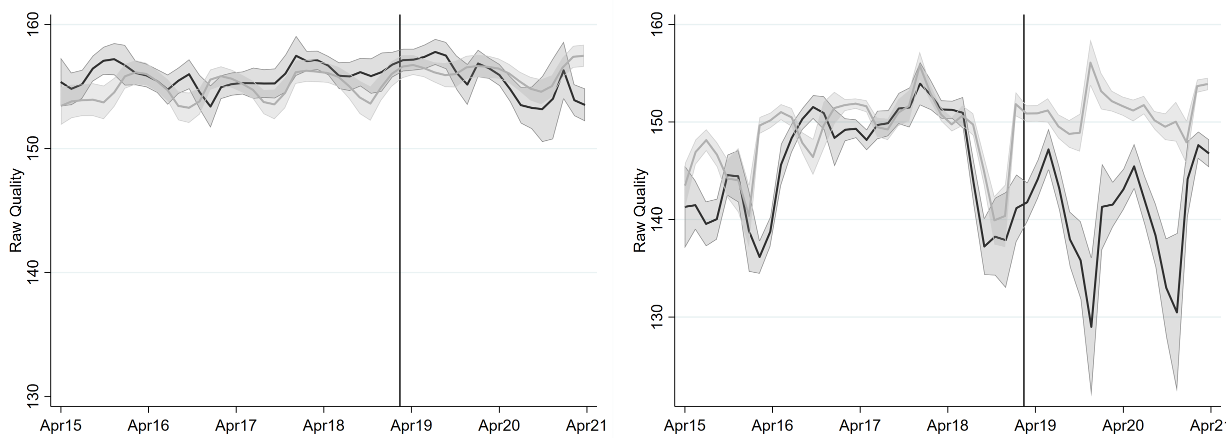

Figure 10: Winner's Raw Quality (Left) and Losers Raw Quality (Right)

The dark line on the left panel represents the raw quality of the winning bid in the auction in which one or more treatment firms participate. The light colored line represents the raw quality of the winning bid for the auction in which one more more control firms participate. The dark line in the right panel represents the raw quality of the losing bid in the auction in which one or more treatment firms participate. The light colored line represents the raw quality of the losing bid in the auction in which one more more control firms participate.

[32]The probability that the winner's bid exceeds the secret reserve price is very close to zero, so we do not

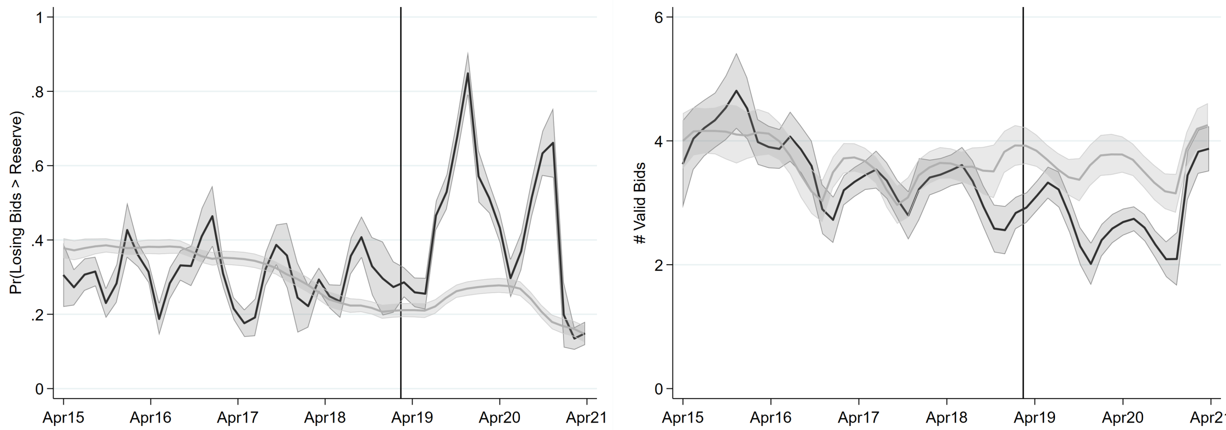

Figure 11: Losing Bids Above Reserve (Left) and Number of Valid Bids (Right)

The dark line on the left panel represents the probability that the losing bid is above the reserve price in the auction in which one or more treatment firms participate. The light colored line represents the probability for the auction in which one more more control firms participate. The dark line in the right panel represents the probability that the losing bid is above the reserve price in the auction in which one or more treatment firms participate. The light colored line represents the probability for the auction in which one more more control firms participate.

In order to assess the statistical significance of these findings, we again conduct Fisher's randomization test for each outcome. In particular, for each outcome variable Y_g, construct $\tau_{G-G'}$ as before for all possible partitions G and G', and compare the value of $\tau_{G_T-G_C}$ against the distribution. The left panel of Figure 12 corresponds to the winning bid and the right panel corresponds to the losing bids. We find that we cannot reject the null that the intervention had an impact of the winning bid.

Figure 13 plots the histogram for quality. The left panel corresponds to the winner and the right panel corresponds to the losers. While we cannot reject the null that the winner's quality is not affected by the treatment, we can reject the null that the treatment has no effect on the losers' quality at 10% significance.

We can conduct analogous exercises for the probability that one of the losers submits a bid above the reserve price, and the number of bidders that submit bids below the reserve price.

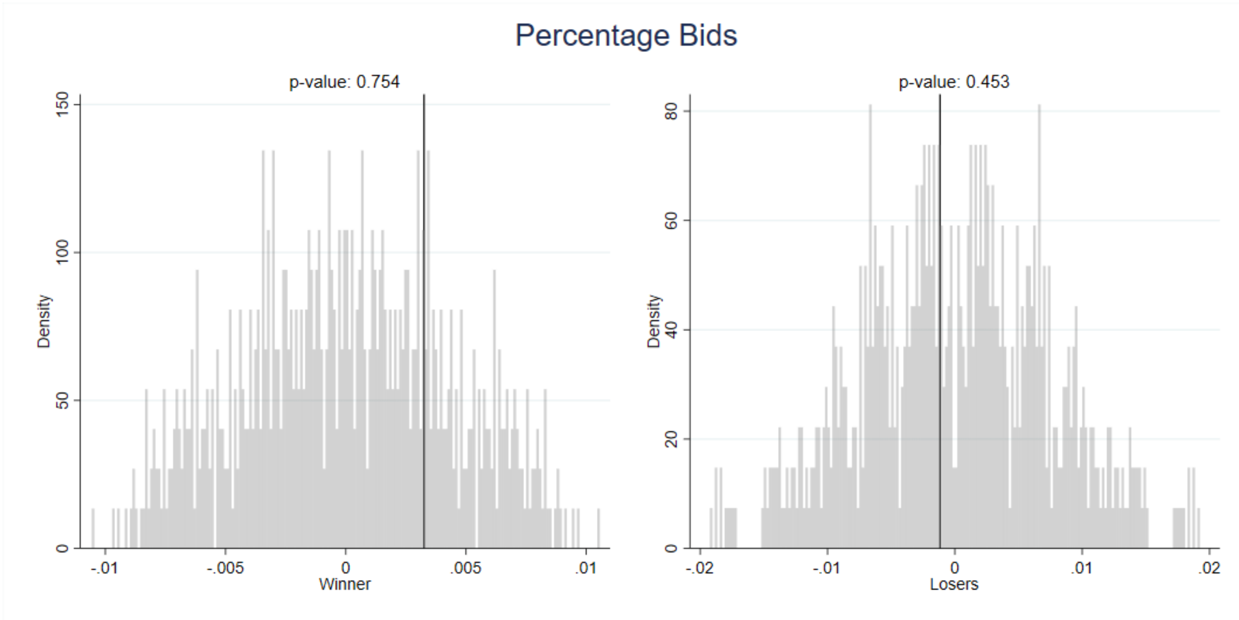

Figure 12: Percentage Bids

Figure 13: Quality

Figures 14 and **??** correspond to the Fisher's test. We reject the null at 10% significance.

It is difficult to rule out completely the possibility that the changes in the bidding behavior

report the corresponding figure for the winner.

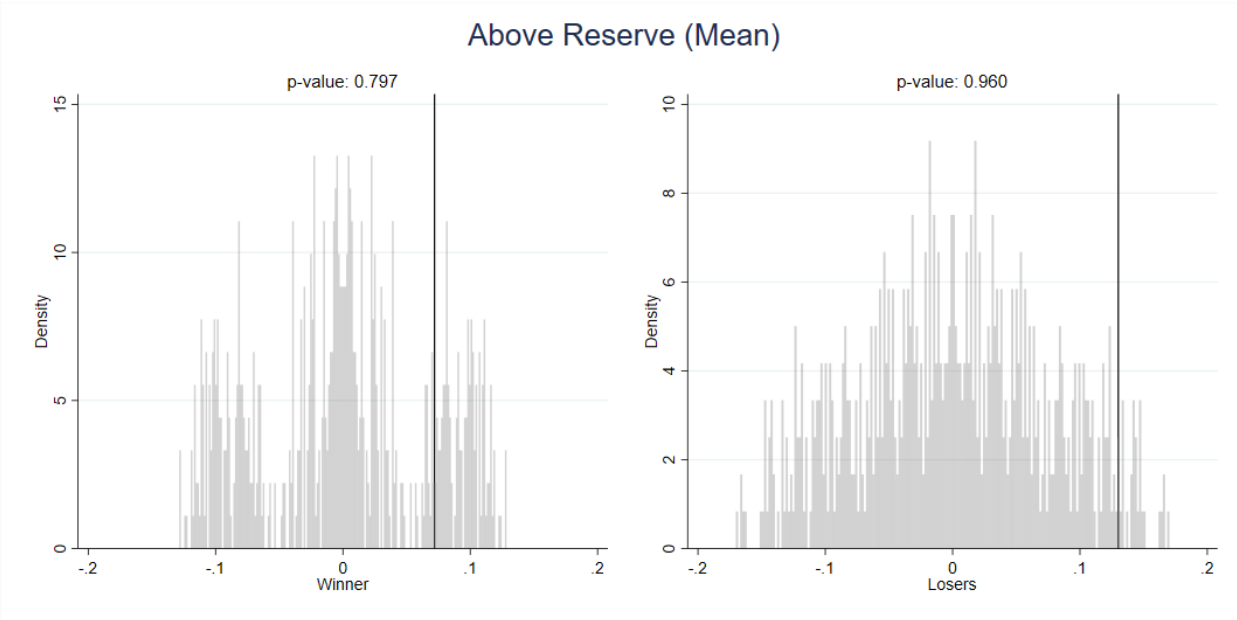

Figure 14: Bids Above Reserve

of the treatment firms captured by Figures 4 and 8 result from breakdown of collusion. However, the fact that the bids do not decrease and the quality of the losers *decrease* for treated firms suggest that this is unlikely. Typically, increased competition is associated with decreases in bids and increases in quality and number of bids. Our results are most consistent with the hypothesis that bidders continued to collude, but adapted to our screen by changing how they bid.

Additional evidence of collusion Lastly, we provide additional direct evidence that bidders in the treated group continued to collude even after the intervention. In order to do so, we apply the test developed in Kawai et al. (2023) directly to the sample of treated firms. In Kawai et al. (2023), we construct a test that compares the backlog (i.e., amount of recently awarded projects) of marginal winners and marginal losers. Under the null of competition, any bidder is just as likely to be the marginal winner as the marginal loser. This implies that marginal winners and losers should, on average, have the same amount of backlog under the null.

Figure 15 corresponds to the bin scatter plot of backlog measures. We take Δ^s on the horizontal axis. Backlog is constructed by summing up the value of auctions won by each firm in a T-day window before the auction. In the figure, we use T equal to 30, 45, 60, 90, 120 and 150. For each T, there are two bin scatter plots, one corresponding to the period before the intervention and the period after. We find that there is a visible discontinuity for $T = 45, 60$ and 90 both for the period before and after the intervention. These results suggest that the treated firms continued to collude even after the intervention.

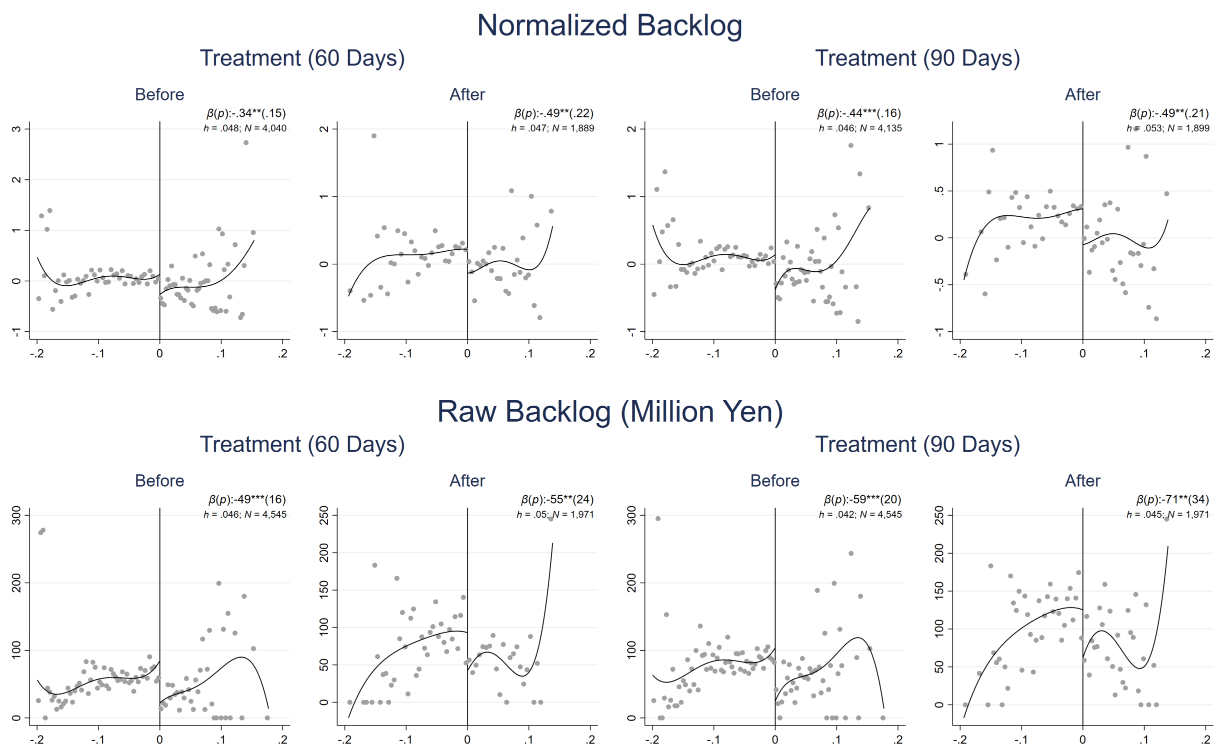

Figure 15: Binned scatter plot of backlog

References

Abrantes-Metz, R., L. Froeb, J. Geweke, and C. Taylor (2006): "A variance screen for collusion," *International Journal of Industrial Organization*, 24, 467–486.

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

Asker, J. and E. Cantillon (2008): "Properties of scoring auctions," *The RAND Journal of Economics*, 39, 69–85.

——— (2010): "Procurement when price and quality matter," *The RAND Journal of Economics*, 41, 1–34.

Athey, S., J. Levin, and E. Seira (2011): "Comparing open and Sealed Bid Auctions: Evidence from Timber Auctions*," *The Quarterly Journal of Economics*, 126, 207–257.

Aubert, C., P. Rey, and W. Kovacic (2006): "The impact of leniency and whistle-blowing programs on cartels," *International Journal of Industrial Organization*, 24, 1241–1266, the Economics of Cartels, Cartel Policy, and Collusion.

Ayres, I. and J. Braithwaite (1992): *Responsive regulation: Transcending the deregulation debate*, Oxford University Press, USA.

——— (1995): *Responsive regulation: Transcending the deregulation debate*, Oxford University Press, USA.

Baer, B. (2014): "Prosecuting Antitrust Crimes," Remarks as Prepared for the Georgetown University Law Center Global Antitrust Enforcement Symposium.

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

Baldwin, L. H., R. C. Marshall, and J.-F. Richard (1997): "Bidder collusion at forest service timber sales," *Journal of Political Economy*, 105, 657–699.

Baránek, B., L. Musolff, and V. Titl (2023): "Detection of Collusive Networks in E-Procurement," *Working Paper.*

Benjamini, Y. and Y. Hochberg (1995): "Controlling the False Discovery Rate: A Practical and Powerful Approach to Multiple Testing," *Journal of the Royal Statistical Society. Series B (Methodological)*, 57, 289–300.

Benjamini, Y., A. Krieger, and D. Yekurieli (1995): "Adaptive linear step-up procedures that control the false discovery rate," *Biometrika*, 93, 491–507.

Braithwaite, J. (1985): *To Punish or Persuade*, State University of New York Press.

Calonico, S., M. Cattaneo, and R. Titiunik (2014): "Robust Nonparametric Confidence Intervals for Regression-Discontinuity Designs," *Econometrica*, 82, 2295–2326.

Chassang, S., K. Kawai, J. Nakabayashi, and J. Ortner (2022): "Robust Screens for Noncompetitive Bidding in Procurement Auctions," *Econometrica*, 90, 315–346.

Chassang, S. and J. Ortner (2019): "Collusion in Auctions with Constrained Bids: Theory and Evidence from Public Procurement," *Journal of Political Economy*, 127, 2269–2300.

Che, Y.-K. (1993): "Design competition through multidimensional auctions," *The RAND Journal of Economics*, 668–680.

Chen, J. and J. E. Harrington (2007): "Chapter 3 The Impact of the Corporate Leniency Program on Cartel Formation and the Cartel Price Path," in *The Political Economy of Antitrust*, ed. by V. Ghosal and J. Stennek, Elsevier, vol. 282 of *Contributions to Economic Analysis*, 59–80.

Christie, W., J. Harris, and P. Schultz (1994): "Why Did NASDAQ Market Makers Stop Avoiding Odd-Eighth Quotes?" *The Journal of Finance*, 49, 1841–1860.

Clark, R., D. Coviello, J.-F. Gauthier, and A. Shneyerov (2018): “Bid rigging and entry deterrence in public procurement: Evidence from an investigation into collusion and corruption in Quebec,” *The Journal of Law, Economics, and Organization*, 34, 301–363.

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

Cunningham, C., F. Ederer, and S. Ma (2021): “Killer acquisitions,” *Journal of Political Economy*, 129, 649–702.

Cyrenne, P. (1999): “On antitrust enforcement and the deterrence of collusive behaviour,” *Review of Industrial Organization*, 14, 257–272.

Davis, S. (2017): “Regulatory complexity and policy uncertainty: headwinds of our own making,” *Working Paper*.

Engstrom, D. (2013): “Public Regulation of Private Enforcement: Empirical Analysis of DOJ Oversight of Qui Tam Litigation under the False Claims Act,” *Northwestern University Law Review*, 107, 1689–1756.

Gehrig, T. and P. Jost (1995): “Quacks, lemons, and self regulation: A welfare analysis,” *Journal of Regulatory Economics*, 7, 309–325.

Grajzl, P. and P. Murrell (2007): “Allocating lawmaking powers: Self-regulation vs government regulation,” *Journal of Comparative Economics*, 35, 520–545.

Harrington, J. (2008a): “Optimal Corporate Leniency Programs,” *The Journal of Industrial Economics*, 56, 215–246.

Harrington, J. E. (2004): “Cartel pricing dynamics in the presence of an antitrust authority,” *RAND Journal of Economics*, 651–673.

——— (2008b): “Detecting cartels,” *Handbook of antitrust economics*, 213, 245.

Hendricks, K. and R. Porter (1988): "An Empirical Study of an Auction with Asymmetric Information," *American Economic Review*, 78, 865–883.

Imbens, G. and K. Kalyanaraman (2012): "Optimal Bandwidth Choice for the Regression Discontinuity Estimator," *The Review of Economic Studies*, 79, 933–959.

Imbens, G. and D. Rubin (2015): *Causal Inference for Statistics, Social, and Biomedical Sciences: An Introduction*, Cambridge University Press.

Innes, R. (1999): "Self-Policing and Optimal Law Enforcement When Violator Remediation is Valuable," *Journal of Political Economy*, 107, 1305–1325.

Kaplow, L. and S. Shavell (1994): "Optimal Law Enforcement with Self-Reporting of Behavior," *Journal of Political Economy*, 102, 583–606.

Kawai, K. and J. Nakabayashi (2022): "Detecting Large-Scale Collusion in Procurement Auctions," *Journal of Political Economy*, 130, 1364–1411.

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

Korting, C., C. Lieberman, J. Matsudaira, Z. Pei, and Y. Shen (2023): "Visual Inference and Graphical Representation in Regression Discontinuity Designs*," *The Quarterly Journal of Economics*, 138, 1977–2019.

Kovacic, W. (2001): "Private Monitoring and Antitrust Enforcement: Paying Informants to Reveal Cartels," *George Washington Law Review*, 69.

LaCasse, C. (1995): "Bid rigging and the threat of government prosecution," *The RAND Journal of Economics*, 398–417.

Landes, W. and R. Posner (1975): "The Private Enforcement of Law," *The Journal of Legal Studies*, 4, 1–46.

Martin, J., M. Pesendorfer, and J. Shannon (2022): "Underbidding for Oil and Gas Tracts," *Working Paper*.

Miller, N. (2009): "Strategic Leniency and Cartel Enforcement," *The American Economic Review*, 99, 750–768.

Monticini, A. and D. Thornton (2013): "The effect of underreporting on LIBOR rates," *Journal of Macroeconomics*, 37, 345–348.

Morgan, K. and D. Rubin (2012): "Rerandomization to improve covariate balance in experiments," *The Annals of Statistics*, 40, 1263 – 1282.

Motta, M. and M. Polo (2003): "Leniency programs and cartel prosecution," *International Journal of Industrial Organization*, 21, 347–379.

Ohashi, H. (2009): "Effects of transparency in procurement practices on government expenditure: A case study of municipal public works," *Review of Industrial Organization*, 34, 267–285.

Ortner, J., S. Chassang, K. Kawai, and J. Nakabayashi (2022): "Screening Adaptive Cartels," .

Parker, C. (2002): *The open corporation: Effective self-regulation and democracy*, Cambridge University Press.

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

Polinsky, M. (1980): "Private versus Public Enforcement of Fines," *The Journal of Legal Studies*, 9, 105–127.

Porter, R. (2005): "Detecting Collusion," *Review of Industrial Organization*, 26, 147–167.

Porter, R. H. and J. D. Zona (1993): "Detection of Bid Rigging in Procurement Auctions," *Journal of Political Economy*, 101, 518–38.

——— (1999): "Ohio School Milk Markets: An Analysis of Bidding," *RAND Journal of Economics*, 30, 263–288.

Schurter, K. (2017): "Identification and inference in first-price auctions with collusion," Tech. rep., working Paper, University of Chicago.

Sigler, J. and J. Murphy (1988): *Interactive Corporate Compliance: An Alternative to Regulatory Compultion*, Quorum Books.

Sokol, D. (2012): "Cartels, Corporate Compliance, and What Practitioners Really Think About Enforcement," *Antitrust Law Journal*, 78, 201–240.

Sokol, D. D. (2017): "Antitrust Compliance," in *The Oxford Handbook of Strategy Implementation*, Oxford University Press.

Spagnolo, G. (2005): "Divide Et Impera: Optimal Leniency Programs," *SSRN Working Paper*.

Storey, J. (2002): "A Direct Approach to False Discovery Rates," *Journal of the Royal Statistical Society. Series B*, 64, 479–498.

Trebbi, F. and M. B. Zhang (2022): "The Cost of Regulatory Compliance in the United States," *Working Paper*.

Wollmann, T. G. (2019): "Stealth consolidation: Evidence from an amendment to the Hart-Scott-Rodino Act," *American Economic Review: Insights*, 1, 77–94.