

The Value of Privacy in Cartels: An Analysis of the Inner Workings of a Bidding Ring

Kei Kawai Jun Nakabayashi Juan Ortner*[†]
U.C. Berkeley & Kyoto University Boston University
University of Tokyo

March 26, 2025

Abstract

We study how incentive constraints can be relaxed by randomization in a repeated-game setting. Our study is motivated by the workings of a detected bidding cartel that adopted a protocol of keeping the winning bid secret from the designated losers when defection was a concern. Keeping the winning bid secret makes accurately undercutting the winning bid more difficult and makes defection less attractive as potential defectors risk not winning the auction even if they deviate. We formalize these ideas in the context of a repeated-game setting and show that a cartel can attain higher payoffs by having the preselected winner randomize its bid and keep it secret from other members. Calibration of the model to the bid data of the cartel suggests that randomization may increase firms' profits by about 56%.

KEYWORDS: procurement, collusion, bidding ring, cartel, privacy.

*Contact information, Kawai: kei@berkeley.edu, Nakabayashi: nakabayashi.jun.8x@kyoto-u.ac.jp, Ortner: jortner@bu.edu.

[†]We are especially thankful to Sylvain Chassang, for many stimulating conversations on the topic of collusion, and to Mr. Masayoshi Oura for providing us the data and answering questions about the cartel. We also thank Shunya Noda and Yuichi Yamamoto for helpful suggestions.

1 Introduction

In this paper, we explore how randomization can relax incentive compatibility (IC) constraints in a repeated-game setting when defection is a concern. Our analysis is motivated by court documents describing how an actual bidding ring, participating in procurement auctions, coped with the threat of defection by strategically hiding the bid of the designated winner from the designated losers. Specifically, the designated losers of the ring were instructed what to bid, but were deliberately kept in the dark as to what the bid of the designated winner would be. Not knowing the winning bid makes it difficult for a potential defector to correctly estimate, and undercut the winning bid. This makes defection less attractive, as potential defectors risk not winning the auction even if they deviate. We formalize these ideas by showing that randomization of the winning bid expands the set of available payoffs to a bidding ring that participates repeatedly in procurement auctions. Using non-deterministic, or random, bidding strategies makes it possible to keep the winning bid secret. We explore quantitatively the importance of randomization by calibrating our theoretical model to the bidding data of the cartel.

The bidding ring that motivates our analysis was comprised of small to mid-sized contractors located in the town of Kumatori, Japan who came under criminal investigation and were subsequently brought to trial. The court documents of the case identify features of the ring that allow us to gain an understanding of the constraints the cartel faced and the protocols adopted by the cartel to overcome them. In particular, the cartel faced a stream of projects of varying sizes, some of which were very large, and hence potentially very profitable, making defection a concern. Prior to each auction, a designated winner was selected based on a bid rotation scheme. In order to counter defection by designated losers, the cartel devised a protocol in which the losing bidders would be instructed how they should bid, but were kept in the dark as to how the designated winner would bid. Moreover, the winning bid was occasionally set substantially lower than what the losing bidders were instructed to bid, making it difficult for the designated losers to correctly predict and undercut the winning bid. Much of the communication between the bidders was mediated by an intermediary.

While it is intuitive that keeping the designated winner’s bid secret can relax IC constraints and help achieve efficiency, this idea contrasts with existing results that show that optimal cartel bidding requires the losing bids to be placed within a very small margin of the winning bid. For example, Marshall and Marx (2007, 2012) and Chassang and Ortner (2019) show that the losing bidders should bid within a very close margin of the winning bid in order to deter the winning bidder from deviating. In these papers, it is important that the designated losers know the winner’s bid.

In order to formalize the idea that randomization relaxes IC constraints, and to understand how it relates to existing work, we construct a model of a bidding ring that repeatedly participates in first-price procurement auctions. The size of the project auctioned off is drawn i.i.d. each period, and firms share the same procurement costs. Consistent with our motivating case study, our baseline model assumes that the cartel has access to a mediator, and allocates contracts through a bid rotation scheme.¹ We say that an equilibrium has *common-knowledge bids* if firms’ bids depend deterministically on the public history.

Under the cartel’s optimal bid rotation equilibrium with common-knowledge bids, winning bids are determined by either the losers’ IC constraints or the reserve price. Whenever the project is sufficiently large and IC constraints bind, the winning bid is set to satisfy the IC constraints with equality, and the losing bidders are instructed to bid within a very small margin of the winning bid. These results are similar to those in Marshall and Marx (2007, 2012) and Chassang and Ortner (2019), where bidders are constrained to use pure strategies.

Our main result shows that a cartel may strictly gain from bidding schemes without common-knowledge bids. In particular, we characterize the optimal bid rotation equilibrium, and show that this equilibrium strictly improves on the optimal bid rotation equilibrium with common knowledge bids if and only if incentive constraints bind and perfect collusion (i.e., winning bid is always equal to the reserve price) is unattainable.

¹ Both of these cartel features (i.e., a mediator and bid rotation) are common among detected cartels. Marshall and Marx (2012) describe how several uncovered cartels hired accounting firms to act as intermediaries, while Asker (2010) studies a cartel of stamp dealers who hired a taxi driver to mediate communication among its ring members. Moreover, because of its prevalence among detected cartels, bid rotation is commonly recognized as an indicator of collusion (see, for instance, the DOJ’s report on “Red Flags of Collusion.”)

Optimal bid rotation equilibria have the following properties. When the reserve price is low and IC constraints are slack, the winning bid is degenerate and set equal to the reserve price. Non-winners place marginally losing bids. When the reserve price is large and IC constraints bind, the winner's bid is drawn from a non-degenerate distribution F . Non-winners are instructed to place a losing bid just above the support of F . Distribution F has the property that the non-winner with the lowest continuation payoff is indifferent between placing a losing bid and deviating by placing a bid in the support of F . By keeping losing firms uninformed, such a bidding scheme relaxes incentive constraints and yields larger profits for the cartel. Note that, for any deviating bid by a designated loser that is in the support of F , the deviator faces the prospect of certain punishment (since all bids are made public ex-post) but a less-than-certain prospect of outbidding the designated bidder.

A key testable implication of the optimal equilibria of our model is that the cartel is more likely to use randomized bidding schemes for large, and hence potentially more profitable, projects. Moreover, optimal randomized bidding schemes tend to exhibit a sizable gap between the winning bid and the second lowest bid. These predictions are different from those of Marshall and Marx (2007, 2012) and borne out in the bidding data of the cartel in our case study.

Another implication of our results is that an auctioneer facing a cartel may strictly benefit from restricting the information that she makes public after each auction. In particular, while randomized bids can be profitable for a cartel when all bids are publicly disclosed, we show that such strategies are no longer beneficial when only the winning bid is made public. As a result, only revealing the winning bid can alleviate collusion and lead to strictly lower prices. Intuitively, if only the winning bid is public, deviations are only detected when the defector wins the auction. This contrasts with settings in which all bids are public, in which defectors can be detected and punished even when they don't win.

In the second half of the paper, we apply the theory to bidding data from a cartel that operated in Kumatori, Japan. Using detailed court documents from the cartel case, we first document features of the cartel that we model in our theoretical analysis. We then use the bidding data from the cartel to calibrate our baseline model. We find that

the calibrated model can replicate cartel bidding patterns such as the winning margin and how it varies with the reserve price, suggesting that, although parsimonious, our model captures the key constraints faced by the bidding ring. We then use the calibrated model to evaluate quantitatively the importance of randomized bidding. We find that randomizing the winning bid increases expected cartel profits by as much as 56% compared to the case with no randomization. This result suggests that the gains from randomization are of first-order importance.

We see the main contribution of the paper as highlighting the practical value of randomization in relaxing IC constraints in naturally occurring settings. While the idea that randomization can help relax IC constraints appears broadly in theoretical work on mechanism design and repeated games, empirical evidence on the use of such randomization has largely been limited to audits, such as tax audits and police traffic stops. Moreover, existing work has not attempted to quantify the value of using randomization. Our case study offers a concrete instance in which randomization is used to relax IC constraints outside of audits. Our calibration quantifies the value of its use. Our analysis suggests that randomization is broadly useful in positive theories of behavior and that there can be large gains associated with its use.

More specific to the antitrust context, our analysis sheds light on the role that transparency plays in sustaining a successful collusive scheme. It is well known from the theory of repeated games that transparency allows colluders to better coordinate and monitor each others' actions. Recently, however, Sugaya and Wolitzky (2018) identify a potential drawback of transparency: they show that transparency may hinder collusion by enabling potential defectors to devise more profitable deviations. Our case study provides an actual example of how privacy helps cartels. By keeping its bid secret from other bidders, the pre-selected winner makes it harder for potential defectors to profitably deviate.

Related literature. Several papers have highlighted how randomization may be useful in static environments. Rahman and Obara (2010) and Rahman (2012) show that randomized messages can help provide incentives by making it easier to identify deviators. Jehiel (2015)

establishes conditions under which a principal may benefit from keeping her agent in the dark about payoff relevant parameters. Ederer et al. (2018) show that randomized contracts may help prevent gaming by agents. Ortner and Chassang (2018) and Chassang and Padró I Miquel (2019) illustrate how randomized incentive schemes may help deter (static) collusion by introducing contracting frictions among collusive parties.

Within repeated-game settings, Kandori (2003), Kandori and Obara (2006), and Rahman (2014) study models with imperfect public monitoring and show how randomization may allow arbitrarily patient players to sustain larger equilibrium payoffs by improving their ability to monitor each other.² Sugaya and Wolitzky (2018) study how the monitoring structure (i.e., the signals that players observe about past prices) affects collusion. Their main observation is that, when past prices help firms predict current prices, a less precise monitoring structure can improve a cartel’s payoff. Intuitively, less precise monitoring relaxes firms’ incentive constraints by making them less informed about their competitors’ prices.

Ortner et al. (2024) also explore how mediation can help improve cartel profits. Like the current paper, they study a model of a bidding ring that has access to a mediator. Unlike the current paper, however, Ortner et al. (2024) focus on symmetric stationary equilibria, under which the allocation of projects does not depend on the history of past play. Their main result characterizes the optimal symmetric stationary equilibrium with two bidders.³ In contrast, the current paper focuses on bid rotation equilibria, under which bidders take turns winning. While bid rotation is not always optimal (indeed, Proposition 8 of Ortner et al. (2024) shows that symmetric stationary equilibria outperform bid rotation), understanding the properties of bid rotation equilibria is important, since bid rotation is a common way for cartels to allocate contracts. In the current paper, we characterize optimal bid rotation equilibria, and, moreover, show that realistic features of actual bidding, such as bid preparation costs, make it optimal for a cartel to engage in bid rotation for many periods.⁴

²Relatedly, Bernheim and Madsen (2017) show that a cartel’s optimal stationary collusive scheme may involve randomization.

³Ortner et al. (2024) show that optimal symmetric stationary equilibria have the following distinctive features: bidders place tied bids with strictly positive probability; bidders are never recommended a bid that loses for sure; and the correlation in bidders’ bids is close to zero. None of these features arises under optimal bid rotation equilibria.

⁴Another difference is that the current paper assumes random reserve prices, while the model in Ortner

Our work also relates to previous work on how cartels bid in a repeated-game setting. There is a large theoretical literature on this topic exploring issues such as monitoring (e.g., Green and Porter, 1984, Skrzypacz and Hopenhayn, 2004), efficiency and private costs (e.g., Athey and Bagwell, 2001, Athey et al., 2004, Athey and Bagwell, 2008), communication (e.g., Compte, 1998, Kandori and Matsushima, 1998, Harrington and Skrzypacz, 2011, Rahman, 2014, Awaya and Krishna, 2016), and demand shocks (e.g., Rotemberg and Saloner, 1986, Haltiwanger and Harrington, 1991). Our model adds to this strand of literature by exploring how randomized bidding schemes can help firms relax IC constraints.

Our focus on bid rotation equilibria relates our work to Kawai et al. (2023), who develop a statistical test to detect collusive bid rotation, and apply this test to bidding data from the US and Japan. Kawai et al. (2023) build a dynamic auction model in which bidders' costs each period depend on past allocations. Their main theoretical result shows that, under competition, any predetermined characteristic (e.g., backlog) of marginal winners and marginal losers must be the same, on average. Hence, differences in backlog among marginal winners and marginal losers are suggestive of collusion via bid rotation.

The empirical portion of our work is closely related to studies that test whether or not the price patterns implied by models of collusion are borne out in the data. Porter (1983) and Ellison (1994) study pricing by the Joint Executive Committee to test for the models of Green and Porter (1984) and Rotemberg and Saloner (1986). Levenstein (1997) studies the bromine cartel and finds evidence consistent with the model of Green and Porter (1984). Borenstein and Shepard (1996) study the retail gasoline market to test for the model of Haltiwanger and Harrington (1991). Wang (2009) also studies the retail gasoline market and finds evidence consistent with the equilibrium of Maskin and Tirole (1988). Wang (2009) is also one of the first papers to document evidence of mixed strategy equilibria for non zero-sum games outside of the lab. Earlier work that documents the use of mixed strategies for zero-sum games include Walker and Wooders (2001) and Chiappori et al. (2002).

Lastly, our paper also relates to the literature on detecting instances of collusion. Early

et al. (2024) assumes that the size of the project is constant across periods. Having random reserve prices allows our model to deliver predictions relating the use of mediation to the distribution of project sizes.

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) and Kawai et al. (2023).⁵

2 Theory

Our baseline model is tailored to match features of the bidding ring that operated in the town of Kumatori, Japan, the bidding data from which we use to calibrate our model. Section 3 describes the inner workings of the Kumatori bidding ring in detail. Here, we highlight three key features of the bidding ring that we incorporate in the model:

- (i) the cartel faced a stream of projects of varying sizes; defection was a concern in some of the very large, and potentially very profitable, projects;
- (ii) prior to each auction, the cartel selected a predetermined winner based on a bid rotation scheme;
- (iii) much of the communication among cartel members was mediated by an intermediary.

These features are common among bidding rings. Firms participating in public procurement typically face projects of varying sizes and profitability; bid rotation is a common way in which cartels allocate projects; and several uncovered cartels have relied on intermediaries (see footnote 1 for references). We discuss at the end of this section how our results extend beyond bid rotation equilibria and to settings without mediation.

2.1 Model

Consider a repeated game in which, in each period $t \in \mathbb{N}_0$, a buyer procures a project from a set of firms $N = \{1, \dots, n\}$. We assume that the buyer uses a first-price procurement auction

⁵Other 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 surveys of the literature, see Porter (2005), Harrington (2008), Asker and Nocke (2021) and Hortaçsu and Perrigne (2021).

with a public reserve price.⁶

In each period t , firms share the same procurement cost $c(r_t) \geq 0$, which depends on the public reserve price r_t of the project auctioned off at time t .⁷ Reserve price r_t is drawn i.i.d. over time from distribution F_r with support $[\underline{r}, \bar{r}]$, with $+\infty \geq \bar{r} > \underline{r} > 0$. We assume that $c(r)$ and $r - c(r)$ are both increasing in r , with $\underline{r} - c(\underline{r}) > 0$, and that $\mathbb{E}_{F_r}[r - c(r)]$ is finite.⁸

After observing the reserve price r_t , firms submit bids $\mathbf{b}_t = (b_{i,t})_{i \in N}$. The bidder who submits the lowest bid wins, and ties are broken randomly. Firm i 's profits in period t are $x_{i,t}(b_{i,t} - c(r_t))$, where $x_{i,t} \in \{0, 1\}$ denotes whether or not bidder i won the auction at t . Firms share a common discount factor $\delta < 1$. We assume that all bids are made public at the end of each period, in line with the way the town of Kumatori runs its auctions.

Mediation. Consistent with our application, we assume that firms have access to a mediator. In each period t , prior to bidding, the mediator privately sends recommended bids $\hat{b}_{i,t}$ to each firm $i \in N$. As we explain in more detail below, bidding recommendations $\hat{\mathbf{b}}_t = (\hat{b}_{i,t})_{i \in N}$ may depend on the history of past reserve prices and bids, current reserve price r_t , and the history of past recommendations. The mediator is assumed to be a disinterested third party, with no incentives.⁹

Histories and strategies. A period- t history for the mediator,

$$h_{M,t} = \left(r_s, \hat{\mathbf{b}}_s, \mathbf{b}_s \right)_{s < t} \sqcup r_t,$$

records all previous reserve prices, bidding recommendations and realized bids, as well as the current reserve price. A strategy $\sigma_M : h_{M,t} \mapsto F_t$ for the mediator maps mediator histories

⁶In our case study, the city of Kumatori used a first-price auction with a secret reserve price. However, the reserve price was non-binding in the auctions that we study. Moreover, there is evidence suggesting that the reserve price was usually leaked to the cartel prior to each auction.

⁷Our decision to abstract away from private costs reflects the fact that, as we describe in Section 3, the designated losers in the Kumatori cartel did not even have the opportunity to estimate costs in the Obara residences auction because the plans were taken away from them.

⁸For our calibration exercise in Section 5, we assume that $c(r) = \alpha \times r$ for some constant $\alpha < 1$.

⁹In practice, the mediator could be paid a flat fee, independent of realized bids. For instance, the stamp cartel described in Asker (2010) paid their mediator a fee that was independent of the auction outcome.

to a distribution over bidding recommendations $F_t \in \Delta(\mathbb{R}_+^n)$.

A period- t history for bidder i ,

$$h_{i,t} = \left(r_s, \widehat{b}_{i,s}, \mathbf{b}_s \right)_{s < t} \sqcup (r_t, \widehat{b}_{i,t}),$$

records past reserve prices $(r_s)_{s < t}$, mediator's private recommendations $(\widehat{b}_{i,s})_{s < t}$, and realized bids $(\mathbf{b}_s)_{s < t}$, as well as current reserve price r_t and mediator's private recommendation $\widehat{b}_{i,t}$. A pure strategy $\sigma_i : h_{i,t} \mapsto b_{i,t}$ for bidder i maps bidder i histories to bids. The period- t public history prior to learning the reserve price at t is $h_t^0 = (r_s, \mathbf{b}_s)_{s < t}$, and the period- t public history after learning the reserve price at t is $h_{t+}^0 = (r_s, \mathbf{b}_s)_{s < t} \sqcup r_t$.

Solution concepts. Let $H_{i,t}$ denote the set of period- t histories of player i . Let $H_t = \times_{i \in N} H_{i,t}$ and, for each $i \in N$, let $H_{-i,t} = \times_{j \in N \setminus \{i\}} H_{j,t}$. Our solution concept is weak Perfect Bayesian Equilibrium (σ, μ) , where $\sigma = (\sigma_M, (\sigma_i)_{i \in N})$ is a strategy profile, and $\mu = (\mu_i)_{i \in N}$ are bidders' beliefs over their opponents' histories; i.e., $\mu_i : h_{i,t} \mapsto \mu_i(h_{i,t})$, where $\mu_i(h_{i,t}) \in \Delta(H_t)$ is a belief over H_t with the property that $\text{supp } \mu_i(h_{i,t}) \subseteq \{h_{i,t}\} \times H_{-i,t}$ (i.e., bidder i 's beliefs assign probability 1 to her own history $h_{i,t}$). In a weak Perfect Bayesian Equilibrium (σ, μ) , bidders' strategies $(\sigma_i)_{i \in N}$ must be sequentially rational at every history, and, for every on-path history $h_{i,t}$, bidder i 's beliefs $\mu_i(h_{i,t})$ must be consistent with the mediator's strategy σ_M . From now on, we use the word *equilibrium* to refer to a weak Perfect Bayesian Equilibrium.

Fix an equilibrium (σ, μ) . For each bidder i and each public history $h_{t+}^0 = h_t^0 \sqcup r_t$, let $\bar{b}_i(h_{t+}^0; \sigma)$ denote the largest (supremum) bid that bidder i submits under (σ, μ) at history h_{t+}^0 . We say that bidders $j \neq i$ submit complementary bids at history h_{t+}^0 if they all place a bid larger than $\bar{b}_i(h_{t+}^0; \sigma)$; i.e., if for all $j \neq i$, $\text{prob}_\sigma(b_{j,t} > \bar{b}_i(h_{t+}^0; \sigma) | h_{t+}^0) = 1$.

Definition 1. We say that equilibrium (σ, μ) is a bid rotation equilibrium if there exists a permutation $\Pi : N \rightarrow N$ such that for all $i \in N$ and all on-path public histories h_{t+}^0 with $t \bmod n = \Pi(i) - 1$, bidders $j \neq i$ submit complementary bids.

Under a bid rotation equilibrium, there is a designated winner at each auction, and bid-

ders other than the designated winner submit complementary bids that lose with probability 1. The identity of the designated winner is determined in a rotating manner: if bidder i is the winner in period t , then she becomes the winner again in period $t + n$.

Definition 2. *We say that equilibrium (σ, μ) has common-knowledge bids if, for all $i \in N$ and for all histories $h_{i,t}$, $\sigma_i(h_{i,t})$ is a pure action and depends only on h_{t+}^0 .*

We note that equilibria with common-knowledge bids correspond to pure strategy equilibria of the game without the mediator. Under an equilibrium with common-knowledge bids, firms's bids don't depend on the mediator's bidding recommendation.

The next subsection focuses on bid rotation equilibria, which correspond to the bidding scheme used by the bidding ring in our case study.¹⁰ Section 2.3 shows how our main theoretical insights extend beyond bid rotation equilibria.

2.2 Bid Rotation Equilibria

Let Σ denote the set of bid rotation equilibria, and let $\Sigma_{\text{ck}} \subset \Sigma$ denote the set of bid rotation equilibria with common-knowledge bids. For each equilibrium (σ, μ) and each $i \in N$, let $V_i(\sigma, \mu)$ denote firm i 's expected discounted payoff at the start of the game under (σ, μ) .

Define

$$\begin{aligned}\bar{V}_{\text{ck}} &\equiv \sup_{(\sigma, \mu) \in \Sigma_{\text{ck}}} \sum_{i \in N} V_i(\sigma, \mu), \text{ and} \\ \bar{V} &\equiv \sup_{(\sigma, \mu) \in \Sigma} \sum_{i \in N} V_i(\sigma, \mu),\end{aligned}$$

to be, respectively, the cartel's largest payoffs under an equilibrium in Σ_{ck} and Σ . Since $\Sigma_{\text{ck}} \subset \Sigma$, we have $\bar{V} \geq \bar{V}_{\text{ck}}$. We maintain the following Assumption throughout the paper:

Assumption 1. *There exists $W > 0$ such that*

$$W \leq \frac{1}{1 - \delta^n} \mathbb{E}_{F_r} [\min\{r - c(r), \delta^{n-1}W\}]. \quad (1)$$

¹⁰The ruling of the Osaka District Court states that the members of the Kumatori Contractors Cooperative allocated projects according to a predetermined order to even out the work of each contractor.

As we will see below, Assumption 1 guarantees that the set of bid rotation equilibria with common-knowledge bids, Σ_{ck} , is non-empty, and that $\bar{V}_{\text{ck}} > 0$. If Assumption 1 fails, Σ_{ck} is empty and our analysis is not interesting. We note that Assumption 1 holds whenever the discount factor δ satisfies $\delta^{n-1}(1 + \delta) > 1$. This condition is satisfied in our calibration.¹¹

Equilibria with common-knowledge bids. Our first result characterizes optimal equilibria in Σ_{ck} . These equilibria take the following intuitive form. The designated bidder bids the reserve price, r_t , or cost $c(r_t)$ plus the expected continuation payoff of the designated loser who won last period (who has to wait $n - 1$ periods to win again), whichever is less. Designated losers submit complementary bids. Deviations are punished by Nash reversion. Letting \bar{W}_{ck} denote the expected continuation value of the designated winner under an equilibrium in Σ_{ck} attaining \bar{V}_{ck} , we have $\bar{V}_{\text{ck}} = \sum_{i=1}^n \delta^{i-1} \bar{W}_{\text{ck}}$. As we show in the Appendix, continuation value \bar{W}_{ck} is the largest value $W > 0$ that satisfies inequality (1). Proposition 1 formalizes this discussion. All proofs are collected in the Appendix and the Online Appendix.

Proposition 1. *On the equilibrium path, any equilibrium in Σ_{ck} that attains \bar{V}_{ck} is such that, in all periods t , the winning bid is given by the minimum between r_t and $c(r_t) + \delta^{n-1} \bar{W}_{\text{ck}}$, where \bar{W}_{ck} is the largest value $W > 0$ that satisfies inequality (1).*

Proposition 1 characterizes on-path winning bids under any equilibrium attaining \bar{V}_{ck} : at any such equilibrium, the winning bid is the minimum between r_t and $c(r_t) + \delta^{n-1} \bar{W}_{\text{ck}}$. Designated losers submit complementary bids that deter the winner from increasing her bid. While there are multiple ways of achieving this, one such way is to have at least one loser bidding marginally above the winning bid.

We note that, in any equilibrium attaining \bar{V}_{ck} , the winning bid as a fraction of the reserve price, $\min_j b_{j,t}/r_t$, becomes less than 1 as r_t becomes larger. This feature of the equilibria is similar to that in Rotemberg and Saloner (1986).

¹¹In our calibrations, we set the annual discount factor between 0.88 and 0.92, and the number of bidders to 9. When adjusted for the number of auctions per year, an annual discount factor of 0.88 translates to a discount of about 0.9715 between auctions. With $n = 9$ and $\delta = 0.9715$, $\delta^{n-1}(1 + \delta) \approx 1.56$.

A previous version of this paper assumed that firms could endogenously break ties. Under endogenous tie-breaking, set Σ_{ck} is non-empty for any $\delta \in [0, 1)$.

Equilibria without common-knowledge bids. Our next result characterizes optimal equilibria in Σ , and establishes necessary and sufficient conditions for $\bar{V} > \bar{V}_{\text{ck}}$. We describe the optimal equilibrium in words before stating the results in the form of a proposition.

Let \bar{W} be the expected discounted payoff of the designated winner under an equilibrium that attains \bar{V} . Hence, the expected continuation value of a bidder who will be the designated winner in $k \leq n - 1$ periods is $\delta^k \bar{W}$.¹² In any optimal equilibrium in Σ , along the path of play the mediator recommends bid b to the winner, with b drawn from c.d.f. $F^*(\cdot; r_t)$. The distribution $F^*(\cdot; r_t)$ is degenerate at r_t if $r_t - c(r_t) \leq \delta^{n-1} \bar{W}$ and is non-degenerate otherwise. Designated losers are recommended to bid above \bar{b} , where \bar{b} is the largest point in the support of $F^*(\cdot; r_t)$. If any bidder deviates, the mediator sends bidding recommendations $b_i = c(r)$ to all $i \in N$ from the next period onwards, and players adhere to this recommendation; i.e., they play Bertrand-Nash. Note that, while deviations by the losers are publicly observed, deviations by the winner may only be detected by the mediator (since bidding recommendations are private). The mediator's messages following a deviation provide the winner with incentives to follow the recommended bid.

Next, we show how we derive the distribution $F^*(\cdot; r)$. Recall that a bidder who will be the designated winner in $n - 1$ periods has a discounted continuation payoff of $\delta^{n-1} \bar{W}$. Moreover, all other non-winners have a continuation payoff that is larger than $\delta^{n-1} \bar{W}$. Suppose that the winning bid at time t is drawn from c.d.f. F_t . Let \bar{b} and \underline{b} denote, respectively, the largest and smallest points in the support of F_t . For the designated losers not to have an incentive to deviate and place a bid $b < \bar{b}$, F_t must satisfy:

$$\forall b < \bar{b}, \quad (1 - F_t(b))(b - c(r)) \leq \delta^{n-1} \bar{W} \iff F_t(b) \geq 1 - \frac{\delta^{n-1} \bar{W}}{b - c(r)}. \quad (2)$$

Indeed, if F_t satisfies (2), a bidder who will be designated winner in $n - 1$ periods does not gain by deviating and placing a bid below \bar{b} . Since other designated losers have larger continuation payoffs, they don't gain by deviating either whenever F_t satisfies (2). Setting $b = \underline{b}$, inequality (2) implies $\underline{b} \leq c(r) + \delta^{n-1} \bar{W}$.

¹²We show in the Appendix that optimal equilibria in Σ are stationary.

Consider now the incentives of the predetermined winner. If the mediator recommends bid $b < \bar{b}$, the winner can increase its bid to $\bar{b} - \epsilon \approx \bar{b}$ and still win the auction. For the winner to have incentives to follow the mediator's recommendation, we must have

$$\underline{b} + \delta^n \bar{W} \geq \bar{b}, \quad (3)$$

where the inequality follows since the winner's equilibrium continuation payoff is $\delta^n \bar{W}$. Since $\underline{b} \leq c(r) + \delta^{n-1} \bar{W}$ (by equation (2)), inequality (3) gives us $\bar{b} \leq c(r) + \delta^{n-1} (1 + \delta) \bar{W}$.

Distribution $F^*(\cdot; r)$ is the highest distribution (in terms of f.o.s.d.) with $\bar{b} \leq r$ satisfying (2) and (3). When $c(r) + \delta^{n-1} \bar{W} \geq r$, $F^*(\cdot; r)$ puts all its mass at r . When $c(r) + \delta^{n-1} \bar{W} < r$, $F^*(\cdot; r)$ is given by:

$$F^*(b; r) = \begin{cases} 0 & \text{if } b < c(r) + \delta^{n-1} \bar{W}, \\ 1 - \frac{\delta^{n-1} \bar{W}}{b - c(r)} & \text{if } b \in [c(r) + \delta^{n-1} \bar{W}, \min\{r, c(r) + \delta^{n-1} (1 + \delta) \bar{W}\}], \\ 1 & \text{if } b \geq \min\{r, c(r) + \delta^{n-1} (1 + \delta) \bar{W}\}. \end{cases}$$

We note two features of distribution $F^*(\cdot; r)$. First, this c.d.f. has a mass point at the highest point in its support. Second, $F^*(\cdot; r)$ has the property that a bidder who will be the designated winner in $n - 1$ periods is indifferent between submitting a losing bid or placing any bid $b < \bar{b}$ in the support of F^* .

The following proposition formalizes this discussion:

Proposition 2. *On the equilibrium path, any equilibrium in Σ that attains \bar{V} is such that, in all periods t , the winning bid is drawn from c.d.f. $F^*(\cdot; r_t)$. Distribution $F^*(\cdot; r_t)$ is degenerate at r_t if $r_t \leq c(r_t) + \delta^{n-1} \bar{W}$ and is non-degenerate otherwise, where \bar{W} is the expected continuation payoff of the designated winner under any equilibrium that attains \bar{V} .*

Moreover, $\bar{V} > \bar{V}_{ck}$ if and only if $\bar{r} > c(\bar{r}) + \frac{\delta^{n-1}}{1-\delta^n} \mathbb{E}_{F_r}[r - c(r)]$.

Proposition 2 characterizes the on-path winning bid distribution under any equilibrium attaining \bar{V} : under any such equilibrium, winning bids are drawn from c.d.f. $F^*(\cdot; r_t)$.

Designated losers submit complementary bids that deter the winner from defecting.¹³ The expression for \bar{W} , which we have not made explicit thus far, is given in the Appendix.

The first key takeaway of the results in this section is that a cartel strictly benefits from strategies without common-knowledge bids when the largest point in the support of F_r is high enough; i.e., $\bar{r} - c(\bar{r}) > \frac{\delta^{n-1}}{1-\delta^n} \mathbb{E}_{F_r}[r - c(r)]$. For our calibration exercise in Section 5, we fit a Pareto distribution to the empirical reserve price distribution by maximum likelihood. Since $\bar{r} = \infty$ for the Pareto distribution, the condition in Proposition 2 always holds. More generally, any distribution with a relatively low mean $\mathbb{E}_{F_r}[r - c(r)]$ and a relatively large upper bound \bar{r} also satisfies the condition.¹⁴

The second takeaway is that when the reserve price is high enough, the winning bid is drawn from a non-degenerate distribution and the expected winning margin (i.e., the difference between the second lowest bid and the lowest bid) is positive in any equilibrium that attains \bar{V} . The value of randomization is positive only when the reserve price is high and not when it is low.

We note that the mediator plays two roles in a bid rotation equilibrium attaining \bar{V} : she sends bidding recommendations to each firm, and acts as a whistleblower against potential defections by the designated winner.

Comparative statics. For any reserve price r , let $\underline{b}(r) \equiv \inf_b \text{supp } F^*(\cdot; r)$ and $\bar{b}(r) \equiv \sup_b \text{supp } F^*(\cdot; r)$ denote, respectively, the lowest and highest bids in the support of the optimal winning bid distribution in any equilibrium in Σ attaining \bar{V} .

Corollary 1. $\underline{b}(r)$, $\bar{b}(r)$, $\bar{b}(r) - \underline{b}(r)$, $\mathbb{E}_{F^*(\cdot; r)}[b]$ and $\bar{b}(r) - \mathbb{E}_{F^*(\cdot; r)}[b]$ are increasing in r .

Corollary 1 shows how different features of the optimal winning bid distribution $F^*(\cdot; r)$ vary with the reserve price.

In order to illustrate Propositions 1, 2 and Corollary 1, we plot the relationship between the winning bid and the reserve price in the left panel of Figure 1, and the relationship

¹³As with Proposition 1, there are multiple equilibria attaining \bar{V} , since there are multiple ways in which designated losers might bid.

¹⁴For example, we use $\delta = 0.95$ and $n = 9$ in our calibration (Section 5). The condition in Proposition 2 would also hold for the empirical distribution of the reserve prices and for any linear cost function, $c(r) = c_0 \times r$ ($c_0 \in [0, 1)$).

between the winning margin and the reserve price in the right panel of Figure 1, for parameter values that we use in our calibration (see Section 5 for details). Specifically, the left panel of Figure 1 plots five functions from bottom to top: (1) $\min\{r, c(r) + \delta^{n-1}\overline{W}_{ck}\}$; (2) $\underline{b}(r)$; (3) $\mathbb{E}_{F^*(\cdot; r)}[b]$; (4) $\overline{b}(r)$; and (5) r . Function (1) corresponds to the winning bid in the optimal bid rotation equilibria with common-knowledge bids. At the calibrated parameters, $c(r) + \delta^{n-1}\overline{W}_{ck}$ is less than r whenever r is larger than 211.2 million yen. This implies that whenever r is above this number, the winning bid is strictly less than the reserve price in the equilibrium with common-knowledge bids. This is depicted by point A in Figure 1.

Functions (2), (3), (4) correspond to features of $F^*(\cdot; r)$, the distribution of winning bids in the optimal bid rotation equilibria that obtains \overline{V} . Function (2) corresponds to the lower bound of the support of $F^*(\cdot; r)$, function (3) corresponds to the mean, and function (4) corresponds to the upper bound of the support. Whenever $r_t - c(r_t)$ is less than $\delta^{n-1}\overline{W}$, distribution $F^*(\cdot; r)$ is degenerate, and functions (2), (3), and (4) are identical. At the calibrated parameters, this is so whenever r is less than 330.4 million yen (Point B). For $r > 330.4$, functions (2) and (4) are depicted by the lower and upper contours of the shaded region. Function (3) is depicted by the solid curve. For r less than 653.2 million yen, the upper contour is equal to r , but at $r = 653.2$, we have $r = c(r) + \delta^{n-1}(1 + \delta)\overline{W}$. This implies that, for values of r above this point, function (4) is strictly less than the 45 degree line. This point is given by C in Figure 1.

The right panel of Figure 1 plots the winning margin. The solid curve corresponds to the expected winning margin, $\overline{b}(r) - \mathbb{E}_{F^*(\cdot; r)}[b]$, while the shaded region corresponds to the support of the winning margin, $[0, \overline{b}(r) - \underline{b}(r)]$. The expected winning margin and the height of the shaded region are both 0 for r less than 330.4 million yen, increase for r between 330.4 and 653.2, and stay flat for values of r greater than that.

What information should the auctioneer disclose? Our model assumes that all submitted bids are made public after each auction, consistent with how the town of Kumatori operates. We note that this is crucial for randomized bidding schemes to be profitable. Indeed, the following result shows that, if only the winning bid and the winner's identity are




Figure 1: The left panel illustrates the relationship between the winning bid and the reserve price. The shaded region corresponds to the range of winning bids under an optimal equilibrium that attains \bar{V} . The dashed line corresponds to the equilibrium winning bid under common-knowledge bids. The right panel plots the winning margin against the reserve price for an optimal equilibrium that attains \bar{V} .

made public, then bidders don't gain from using randomized bidding schemes.

Proposition 3. *Suppose that only the winning bid and the winner's identity are made public after each auction. Then, $\bar{V} = \bar{V}_{ck}$.*

On the equilibrium path, any equilibrium in Σ that attains \bar{V} is such that, in all periods t , the winning bid is given by the minimum between r_t and $c(r_t) + \delta^{n-1}\bar{W}_{ck}$, where \bar{W}_{ck} is the largest solution to (1).

Proposition 3 shows that, when only the winning bid and the winner's identity are made public, a cartel does not benefit from using randomized bidding schemes: optimal cartel profits under bid rotation can be attained with an equilibrium with common knowledge bids. In addition, cartel profits in this case coincide with profits under the equilibrium in Proposition 1.

To understand the contrast between Propositions 2 and 3, note that when all bids are made public, a defector can be punished even if she doesn't win the auction. As a result,

randomized winning bids relax the incentive constraints of designated losers. In contrast, when only winning bids are public, a deviation by a designated loser can only be detected and punished if she wins the auction. As a result, randomized winning bids are no longer profitable when only the winning bid is made public.¹⁵

2.3 Extensions

Our baseline model is tailored to match the key features of the bidding ring from Kumatori. In this section, we discuss how our results extend when we consider (1) the set of all equilibria (and not just the set of bid rotation equilibria); and (2) no mediation.

Beyond bid rotation equilibria. In our simple stylized model, bid rotation equilibria are not optimal whenever they fail to achieve perfect collusion.¹⁶ Intuitively, under a bid rotation scheme, the designated loser who won in the previous period has a strong incentive to defect, limiting the bids that can be sustained. Cartel profits can be improved by having all firms win with positive probability each period.¹⁷

In practice, however, there are features of cartels that our model abstracts from, such as capacity constraints and bid preparation costs, that may make bid rotation efficient. The Online Appendix considers a richer model in which firms incur bid preparation costs and firms' procurement costs are increasing functions of backlog. In that environment, for a range of discount factors, including those for which perfect collusion cannot be achieved, any optimal equilibrium has the property that firms rotate who wins for many periods.¹⁸ Hence,

¹⁵Ortner et al. (2024) show that the same result holds when considering the set of symmetric stationary equilibria: mediation does not improve optimal symmetric stationary equilibrium payoffs when only the winning bid and the winner's identity are made public.

¹⁶By fail to achieve perfect collusion, we mean that $\bar{r} > c(\bar{r}) + \delta^{n-1}\bar{W}$, so that the winning bid is below the reserve price for high enough realization of r under an optimal bid-rotation equilibrium.

¹⁷Proposition 8 in Ortner et al. (2024) shows that, in this model, optimal symmetric stationary equilibria outperform bid rotation equilibria.

¹⁸Formally, we first show that when the discount factor δ is above a cutoff $\delta_{\text{opt}} < 1$, perfect collusion can be sustained as a bid rotation equilibrium, where perfect collusion here implies that firms rotate who wins, with the lowest cost firm bidding the reserve price and the other firms not placing serious bids (to avoid paying duplicate bid preparation costs). We next show that there exists $\hat{\delta} < \delta_{\text{opt}}$ such that, for all $\delta \in (\hat{\delta}, \delta_{\text{opt}})$, any optimal equilibrium has the property that firms rotate who wins until (at least) some stopping time $\tau(\delta)$, with $\tau(\delta)$ diverging to $+\infty$ as $\delta \nearrow \delta_{\text{opt}}$.

in that more realistic environment, restricting attention to bid rotation equilibria is without loss of generality for a range of discount factors.¹⁹ Moreover, equilibrium cartel payoffs under common knowledge bids can be strictly improved upon by randomization.

Here, we show that even in the baseline setting (i.e., no bid preparation costs and constant marginal costs) and with no restrictions on the set of equilibria, there is still value to using randomized bidding strategies. To show this, let Σ^* denote the set of all equilibria of the repeated game of Sections 2.1-2.2, and let $\Sigma_{\text{ck}}^* \subset \Sigma^*$ denote the set of equilibria with common-knowledge bids. Recall that, for any equilibrium (σ, μ) , we denote by $V_i(\sigma, \mu)$ firm i 's expected discounted payoff under σ at the start of the game. Define

$$V_{\text{ck}}^* \equiv \sup_{(\sigma, \mu) \in \Sigma_{\text{ck}}^*} \sum_{i \in N} V_i(\sigma, \mu), \text{ and}$$

$$V^* \equiv \sup_{(\sigma, \mu) \in \Sigma^*} \sum_{i \in N} V_i(\sigma, \mu)$$

to be, respectively, the highest cartel payoff that can be attained by an equilibrium in Σ_{ck}^* and Σ^* . Since $\Sigma_{\text{ck}}^* \subset \Sigma^*$, we again have $V^* \geq V_{\text{ck}}^*$. Our next result shows that bidding schemes without common-knowledge bids are strictly beneficial whenever \bar{r} is large enough.

Proposition 4. $V^* > V_{\text{ck}}^*$ if and only if $\bar{r} > c(\bar{r}) + \frac{1}{n-1} \frac{\delta}{1-\delta} \mathbb{E}_{F_r}[r - c(r)]$.

The proof of Proposition 4 consists of three parts. In the first part, we characterize equilibria in Σ_{ck}^* that attain V_{ck}^* . Under such equilibria, on the path of play, in each period t all bidders place the same bid b_t , and win the auction with equal probability. Each bidder then earns a continuation payoff equal to $\delta V_{\text{ck}}^*/n$ after every on-path history. To deter bidders from defecting, bid b_t placed at time t must be such that

$$\frac{b_t - c(r_t)}{n} + \frac{1}{n} \delta V_{\text{ck}}^* \geq b_t - c(r_t) \iff b_t \leq c(r_t) + \frac{1}{n-1} \delta V_{\text{ck}}^*. \quad (4)$$

Indeed, if (4) did not hold, bidders would have a strict incentive to undercut the winning

¹⁹In fact, in the Kumatori bidding ring, the designated losers did not perform cost estimates, as we discuss below. Previous papers have found that estimating costs is a costly activity. Krasnokutskaya and Seim (2011), for example estimate bid preparation costs to be about 2.2% to 3.9% of the engineer's cost estimate.

bid. Hence, in any equilibrium in Σ_{ck}^* attaining V_{ck}^* , bid b_t must be equal to $c(r_t) + \frac{1}{n-1}\delta V_{\text{ck}}^*$ or r_t , whichever is less. Note that equilibria in Σ_{ck}^* attaining V_{ck}^* achieve perfect collusion whenever $\bar{r} \leq c(\bar{r}) + \frac{1}{n-1}\delta V_{\text{ck}}^*$. Hence, $V_{\text{ck}}^* = V^*$ in this case.

The second part of the proof of Proposition 4 shows that when $\bar{r} > c(\bar{r}) + \frac{1}{n-1}\delta V_{\text{ck}}^*$ (so perfect collusion can't be sustained with common-knowledge bids), cartel members can attain profits strictly larger than V_{ck}^* by having the mediator send random bid recommendations to all bidders. In particular, we use the following scheme to prove Proposition 4. If $r_0 \leq c(r_0) + \frac{1}{n-1}\delta V_{\text{ck}}^*$, the mediator recommends each bidder to bid r_0 at $t = 0$. If instead $r_0 > c(r_0) + \frac{1}{n-1}\delta V_{\text{ck}}^*$, the mediator sends i.i.d. recommendations to each firm, recommending them to bid r_0 with probability $\gamma \in (0, 1)$ and $c(r_0) + \frac{1}{n-1}\delta V_{\text{ck}}^*$ with the complement probability. Deviations are punished by Nash reversion. From $t = 1$ onwards, players play the equilibrium with common-knowledge bids attaining V_{ck}^* . Clearly, this bidding scheme generates profits strictly larger than V_{ck}^* . The proof of Proposition 4 shows that, for an appropriately chosen γ , bidders don't have an incentive to defect.

The third and last part of the proof of Proposition 4 shows that $\bar{r} > c(\bar{r}) + \frac{1}{n-1}\delta V_{\text{ck}}^*$ if and only if $\bar{r} > c(\bar{r}) + \frac{1}{n-1}\frac{\delta}{1-\delta}\mathbb{E}_{F_r}[r - c(r)]$.

No mediation. We now briefly discuss how our results extend to the case of no mediation. In our model, the mediator acts as a whistleblower against defection by the designated winner. When we dispense with the mediator, we need a way for the designated losers to verify that the designated winner actually randomizes its bid according to F^* . We show that, with at least three firms, it is possible to achieve the equilibrium that attains \bar{V} even in the absence of a mediator if bidders have access to private randomization devices, the outcome of which they can share with each other.

To see this, assume that $n \geq 3$, and consider the following protocol. (i) Two designated losers independently and privately draw a random variable from $\text{Uniform}[0, 1]$, U_{l_1} and U_{l_2} . (ii) These two bidders simultaneously and privately share U_{l_1} and U_{l_2} with the designated winner. (iii) The designated winner computes the fractional part of $U_{l_1} + U_{l_2}$; i.e., $Z = U_{l_1} + U_{l_2}$ if $U_{l_1} + U_{l_2} \leq 1$, or $Z = U_{l_1} + U_{l_2} - 1$ if $U_{l_1} + U_{l_2} > 1$. (iv) The designated winner bids

$F^{*-1}(Z; r)$, where for each $z \in [0, 1]$, $F^{*-1}(z; r) = \inf\{b \in [0, r] : F^*(b; r) \geq z\}$ is the generalized inverse of $F^*(\cdot; r)$; and the designated losers submit bids marginally above $\bar{b}(r)$, the highest point in the support of $F^*(\cdot; r)$. (v) After bids are realized, the two designated losers publicly and simultaneously share the realization of U_{l1} and U_{l2} ; if the winner's bid coincides with $F^{*-1}(Z; r)$, and all losers submit losing bids, players continue colluding next period; otherwise, they revert to Bertrand-Nash. Because Z is distributed uniformly on $[0, 1]$, the winner's bid $F^{*-1}(Z; r)$ has distribution $F^*(\cdot; r)$. Moreover, neither of the two designated losers has an incentive to deviate, since conditional on U_{li} , Z is distributed uniformly on $[0, 1]$. (If $n > 3$, the other designated losers don't have an incentive to deviate either).

3 Bid-rigging in the Town of Kumatori

This section provides a description of the internal organization of a detected bidding ring that operated in the town of Kumatori that motivates our analysis. Our description of the cartel is drawn from court documents.²⁰ We use the bidding data from this cartel to calibrate the model developed in the previous section.

3.1 Background

Auctions for construction projects in Kumatori. The town of Kumatori uses auctions to allocate construction projects that are estimated to cost above 1.3 million yen, or about 13,000 dollars. The auction format is first-price sealed bid with a reserve price. Although the reserve price is secret, it was customary for the cartel to seek out the reserve price from the town officials, as we discuss below.²¹

An important feature of the auctions is that participation is by invitation only. The town maintains a list of qualified contractors and invites a subset of firms from the list to bid. The

²⁰Online Appendix OD contains key excerpts from the original ruling, along with their English translation. The source for the court documents and the bidding data of the cartel is a booklet published by the plaintiffs who sued the mayor of Kumatori for failing to pursue liability claims against the cartel (Oura and Kashiwano, 2015). The booklet can be found at <http://www.keikawai.com/booklet.pdf>.

²¹When none of the bids meet the reserve price, the lowest bidder and the town typically engage in a bilateral negotiation. The lowest bid was always below the reserve price during our sample period, however.

town maintains separate lists for each project category as well as for different project sizes within a category. For example, for building construction, the town maintains four mutually exclusive lists of contractors (Tier A through Tier D). The town typically invites Tier A firms to bid on the largest projects, Tier B firms to bid on the next largest projects, and so on, essentially segmenting the market by project size. All of Tier A firms are headquartered outside of Kumatori and are invited to bid on exceptionally large projects. Tier B firms and below are local firms typically headquartered within the town. Most of the Tier B and C firms were members of the Kumatori Contractors Cooperative, a trade association that consisted of a little more than 20 mid to small-size contractors in Kumatori. The members of this cooperative were found to be colluding.

Bidding ring in the town of Kumatori. Reports of police investigation of the members of the Kumatori Contractors Cooperative for bid-rigging first appeared in the news on October 12, 2007. In addition to the contractors in Kumatori, more than 20 town officials, including the town mayor, were questioned by the police. In November 2007, four individuals were indicted for bid-rigging. The criminal charges focused on the defendants' involvement in a bid-rigging arrangement in a single auction that took place on August 22, 2006 that Imakatsu Construction won, an auction for rebuilding a public housing complex (Obara Residences). The defendants included Mr. Kitagawa, the owner of Imakatsu Construction and director of the Kumatori Contractors Cooperative; his son, who was an employee of Imakatsu Construction; Mr. Nishio, the vice-director of the Cooperative; and Mr. Takano, an employee of the Cooperative. While Mr. Nishio and Mr. Takano were not participants of the auction, they mediated much of the communication between Mr. Kitagawa and the other participants of the auction. For example, Mr. Nishio gave out the instructions to other bidders on how they should bid.²² All four defendants were found guilty in trial in March 2008. Mr. Kitagawa was sentenced to prison for 18 months, Mr. Nishio to 14 months, and the other defendants to 10 months.

Although the criminal case focused on the defendants' involvement in a bid-rigging ar-

²²The ruling states that Mr. Nishio played a leading role in rigging bids on previous lettings too, as part of his role as the vice-director of the Cooperative.

rangement in the Obara Residences auction alone, the court ruling of the case made it clear that the Obara Residences auction was not an isolated incident, and that there were many other individuals who actively participated in the bid-rigging scheme.²³ The court ruling called bid-rigging among members of the Kumatori Contractors Cooperative “deep-rooted”, and “habitual”, and, furthermore, stated that the members of the Cooperative allocated projects according to a predetermined order to even out the work of each contractor. While none of the town officials were formally charged, the court ruling also stated that the designated winner of the cartel would approach the town officials to seek out the engineering estimate. The designated winner would perform cost estimates based on the project plan but the designated losers did not.²⁴

In response to the ruling of the criminal case, the town of Kumatori withheld part of its payment to Imakatsu Construction for work that had been completed on the Obara residences in order to off-set liability claims. However, the mayor and town officials showed little interest in pursuing claims for damages incurred on other auctions. This inaction led some of the residents of Kumatori to file suit against the mayor asking the court to order the mayor to pursue claims against 23 firms, all members of the Kumatori Contractors Cooperative, for damages incurred on other auctions. The District Court of Osaka ruled in favor of the plaintiffs, ordering the town to pursue claims against the bidders in the amount of about 375 million yen, or about 3.75 million dollars. The mayor appealed the ruling, but the verdict was upheld by the Osaka High Court with relatively minor modifications.

3.2 Obara Residences Auction

Among the auctions that the cartel bid on, the Obara residences auction was unique because of the large value of the project. We now discuss how, despite the town’s policy of segmenting the market by project size, the members of the Kumatori Contractors Cooperative were invited to bid in the auction. We also draw on the court rulings to provide a detailed quali-

²³Our description of the cartel in this paragraph is taken from page 2 and 3 of Ruling H19 (WA) No. 6418, Osaka District Court.

²⁴See page 22 of Ruling H21 Gyo-U, No. 99.

tative description of the way in which bid-rigging was carried out for the Obara Residences auction. We include the original text of the ruling and English translations of key passages in the Online Appendix.

The town of Kumatori started planning for the rebuilding of the Obara Residence Complex, an ageing public housing project, around year 2000 according to the minutes of the meeting of the town council.²⁵ The new residence complex would consist of three separate buildings. Construction of the first and smallest of the three buildings was put to an auction in April 2004 to Tier A firms. The winning bid was 363 million yen, or about 3.6 million dollars. The winner of the auction was Asanuma Corporation, a contractor headquartered in the city of Osaka with annual sales of about 2 billion dollars.

In the Fall of 2004, Mr. Kitagawa, the director of the Kumatori Contractors Cooperative and owner of Imakatsu Construction began lobbying the town's mayor and the head of the town's general affairs department to let Tier B firms bid on the second and third components of the Obara Residences Complex.²⁶ Mr. Kitagawa was an important supporter of the mayor. Despite repeated lobbying by Mr. Kitagawa, the head of the department was reluctant to let Tier B firms bid on the Obara residences project initially, according to court documents. However, the head of the department started to warm towards the idea around April 2006.²⁷

According to the Osaka District Court ruling, Mr. Kitagawa of Imakatsu Construction started to believe, around April of 2006, that Tier B firms would be invited to bid on the second part of the Obara residences project and started mentioning the project at the meetings of the Kumatori Contractors' Cooperative. In one meeting of the Cooperative in late June 2006, Mr. Kitagawa stated that he wanted others to let his firm, Imakatsu Construction, win the auction. He also told the members that he would be collecting, from each of the invited bidders, the detailed project plan that the town distributes at the on-site briefing. This was understood by the members of the Cooperative as a preventative measure to make defection more difficult by making it harder for other firms to estimate costs.²⁸

²⁵Minutes of the meeting of the town council to discuss the budget, March 2001, page 49.

²⁶Our description of the cartel in this paragraph are taken from page 17 of Ruling H21 (Gyo-U) No. 99, Osaka District Court.

²⁷Last paragraph, page 17 of the Osaka District Court ruling.

²⁸Page 18 of the Osaka District Court ruling (Ruling H21 Gyo-U, No. 99).

On August 3, 2006, the town office sent out invitations to five contractors to bid on the housing complex project, including Imakatsu Construction. All invited bidders were Tier B firms and members of the Cooperative. On or around this day, Mr. Kitagawa asked Mr. Nishio to help him with the operation and, in particular, to obtain confidential information about the project from the town, create cost breakdowns for each firm, determine what to bid, and instruct bidders the amount each bidder should bid. Mr. Nishio was the vice director of the Cooperative and a senior managing director of Nishinuki Construction at the time. Nishinuki Construction was a Tier C firm, and not one of the invited bidders.

On August 4, 2006, after the general meeting of the Cooperative, Mr. Kitagawa met with the four other bidders that were invited to bid on the Obara Residences project and repeated his intention to collect the project plans from them after the on-site briefing.

The town of Kumatori held an on-site briefing for the Obara residences project for the five invited firms on August 7. At the on-site briefing, the town distributed the detailed project plan as well as other documents required to estimate costs. These documents were collected immediately after the briefing from all bidders other than Imakatsu by an employee of the Cooperative. The plans and the documents were not returned to the bidders until August 21, the day before the auction. During this time, Imakatsu Construction was the only firm that had access to the documents required to estimate costs: Other bidders did not have access to the documents that would make cost estimates possible. After the on-site briefing, Mr. Nishio met with an official of the Buildings Division of the town to obtain information about the engineering estimate.²⁹

On August 21, one day before the day of the auction, Mr. Kitagawa and Mr. Nishio met at the office of Imakatsu Construction and decided on a bid of 630 million yen for Imakatsu Construction and bids of above 700 million yen for all other bidders. According to the court ruling, they decided to set the losing bids to be above 700 million yen with the explicit intent of making it hard for the other bidders to guess the bid of Imakatsu.³⁰

Also on the same day, Mr. Nishio contacted the four other invited bidders of the auction.

²⁹Page 19 of the Osaka District Court ruling (Ruling H21 Gyo-U, No. 99).

³⁰Page 19 of the Osaka District Court ruling (Ruling H21 Gyo-U, No. 99).

He told them that he would hand them the document containing each firm's break-down of the estimated costs on the day of the auction (a bid would need to be accompanied by this document for it to be considered valid). Mr. Nishio also told the bidders that he would give instructions on how much to bid on the day of the auction. According to the court ruling, Mr. Nishio's decision to hand the documents and give instructions on bids on the day of the auction (as opposed to two days prior to the day of the auction, as was customary for the bidding ring) was to prevent defection given the large size of the project.³¹

The auction for the public housing complex was held at the town office on August 22 2006. Mr. Nishio went to the town office and stapled together the documents containing the cost break-down of each bidder with the cover page brought by the representatives of the firms. Mr. Nishio also indicated, to each representative of the firms, the amount that each bidder should bid by showing a slip of paper with a number above 700 million yen. Importantly, the bid of Imakatsu Construction was kept secret to the other bidders. The representatives of the invited bidders bid the same or slightly above the amount shown on the slip of paper. As a result, Imakatsu Construction won the auction with a bid of 630 million yen. The other bids were 705 million yen, 707 million yen, 710 million yen and 723 million yen, respectively.

4 Data and Summary Statistics

In this section, we briefly discuss the data and key bidding patterns.

4.1 Data

Our primary data source is the bid data submitted to court by the residents of Kumatori who sued the mayor for failing to pursue damages against the bidding ring. The residents claimed that the town was owed money for a subset of the auctions won by the bidding ring from April 2003 to October 2007. We have information on all of the bids, date of the auction, reserve price, and the identity of the winner (but not the identity of the losing bidders) for

³¹Page 4 of the Osaka High Court ruling (Ruling H24 Gyo-Ko, No. 101).

those auctions. A limitation of this dataset is that it does not include auctions for which the plaintiffs did not pursue damages even if a cartel member won the auction. Moreover, lettings that were awarded to non-colluding firms are not included.

We complement our primary data source with auction data that were collected by the town of Kumatori. This dataset covers all auctions let by the town between April 2006 and December 2009. Because the second dataset includes the universe of auctions let by the town during this period, the second dataset is a superset of the first during the period the two datasets overlap, i.e., between April 2006 and October 2007. The second dataset contains information on all bids, identity of the bidders, date of the auction and the reserve price.

Table 1 reports the summary statistics of the two datasets. The top panel reports the summary statistics without conditioning on the tier of the participants and the bottom panel focuses on auctions in which Tier B firms were invited to bid. As we discussed in section 3.1, the town of Kumatori segments the market by the size of the project and invites bidders from a particular tier to bid on a given auction for almost all auctions. We are specifically interested in Tier B firms because the participants of the Obara residence auction were Tier B firms and these are the firms we know the most about from the court documents. All of the Tier B firms were members of the bidding ring.

The first two columns of Table 1 correspond to the dataset we obtained from the plaintiffs. Column (1) reports the summary statistics for all of the auctions in that sample (top panel) and those for Tier B auctions (bottom panel). Column (2) reports the summary statistics for the subset of the auctions that were let after April 2006, which is when the dataset we obtained from the town begins. We find that the average reserve price is 28.93 million yen (about \$290,000) in the top panel of column (1) while it is significantly higher, at 39.65 million yen in the top panel of column (2). The difference in the mean reserve price between column (1) and column (2) is explained by the fact that column (2) includes the Obara residence auction. The reserve price of the Obara residences auction was about 17 times larger than the other auctions. This auction raises the average reserve price and the winning bid much more in column (2) given the smaller sample size in column (2). We also find that

	Data from Plaintiffs		Data from Kumatori
	All periods (1)	2006.4 - 2007.10 (2)	2006.4 - 2007.10 (3)
All Tiers			
Reserve price (mil. Yen)	28.93 (58.48)	39.65 (105.07)	31.33 (79.97)
Lowest bid (mil. Yen)	28.06 (56.44)	38.27 (100.65)	29.21 (75.67)
Lowest bid/Reserve	0.967 (0.011)	0.964 (0.011)	0.928 (0.068)
#Bidders	10.41 (2.76)	8.00 (2.77)	9.41 (2.68)
Sample size	158	39	79
Tier B Auctions			
Reserve price (mil. Yen)	47.80 (87.97)	166.83 (244.92)	158.63 (224.63)
Lowest bid (mil. Yen)	46.46 (84.79)	161.17 (234.60)	152.69 (215.33)
Lowest bid/Reserve	0.969 (0.011)	0.971 (0.015)	0.965 (0.020)
#Bidders	10.84 (2.78)	8.33 (1.63)	8.57 (1.62)
Sample size	63	6	7

Table 1: Sample Statistics – Auctions. Columns (1) and (2) correspond to the sample of auctions in the plaintiffs dataset and column (3) corresponds to those in the dataset we obtained from the town. Column (2) corresponds to a subset of the auctions in the plaintiffs dataset that were let between April 2006 and October 2007. The bottom panel corresponds to Tier B auctions. Standard errors are reported in parenthesis.

Name	(1) Revenue	(2) #Won	(3) Obara
Imakatsu Construction	759.45	6	●
Nakabayashi Construction	488.65	12	
Nishio Gumi	454.40	8	
Tokushin construction	416.39	6	
Takada Gumi	396.25	9	○
Yamamoto Construction	195.60	7	○
Nakajima Kougyou	159.80	5	
Hannan Construction	109.70	4	○
Seiko Construction	48.23	7	○

Table 2: Sample Statistics – Tier B firms.

the reserve price and the bids in the bottom panel are larger than in the top panel, reflecting the fact that Tier B auctions are generally larger than Tier C auctions. Note that the data from the plaintiffs only include Tier B and Tier C auctions. The average winning bid is about 96-7% of the reserve price.

Column (3) reports the sample statistics for the dataset that we obtained from the town of Kumatori. There are 79 auctions that were let by the town between the beginning of the sample (April 2006) and the breakdown of the cartel in October 2007. Because the set of auctions in the town data is not a selected sample, the data include auctions in which the cartel bidders did not participate (e.g., Tier A and Tier D auctions) as well as those in which the cartel bidders participated but are not included in the plaintiffs dataset. Comparing the sample sizes between columns (2) and (3) in the top panel, we find that about half (39/79) of the auctions are included in the first dataset during the period of overlap. If we focus on the bottom panel, we find that the sample size is 6 in column (2) and 7 in column (3). The fact that 6 out of 7 auctions let during the period of overlap are contained in the first dataset suggests that sample selection issues are unlikely to be too severe for Tier B auctions. For this reason, we use the first dataset in all of the subsequent analysis.

Table 2 reports the summary statistics of the 9 contractors that were on Tier B using the

dataset from the plaintiffs.³² Column (1) reports the total amount awarded to each of the firms during the sample period. The award amount varies from a high of about 760 million yen to a low of about 50 million yen. The number of auctions awarded is reported in column (2). Column (3) reports whether or not a firm was invited to the Obara residences auction. White circles correspond to those that were invited to bid in the auction and a black circle corresponds to the winner of the auction.

4.2 Size of Auctions and Winning Margin

The Obara residences auction was unusually large compared to other auctions that Tier B firms were invited to bid on. The left panel of Figure 2 plots the reserve price of Tier B auctions in our data. The horizontal axis is the calendar date and the vertical axis is the reserve price of the auction. The figure shows that, except for the Obara residence auction that took place on August of 2006 (corresponding to the dark circle), Tier B firms were invited to bid on auctions with reserve prices below 200 million yen. The average reserve price during this period excluding the Obara residences auction is about 39.1 million yen. The reserve price of the Obara residences auction was 657 million yen, or about 17 times the average size of projects on which these bidders were invited to bid.

The right panel of Figure 2 plots the difference between the lowest and the second lowest bids as a fraction of the reserve price for these auctions.³³ The figure shows that the winning margin is always less than 4%, except for the Obara residences auction (corresponding to the dark circle). The average winning margin for lettings excluding Obara Residences is about 0.93%. The margin for Obara Residences, on the other hand, is 11.4%. The bidding patterns suggest that the cartel members kept the winning margin small for all projects except for the Obara residences project.

³²The sample of auctions used for Table 2 include all of the 63 Tier B auctions from the plaintiffs dataset.

³³We define the winning margin as the difference between the lowest bid and the second lowest bid even when the second lowest bid exceeds the reserve price (note that the lowest bid is always below the reserve price in the sample). This definition corresponds to the money left on the table (i.e., amount of surplus forgone by the winner this period) under the assumption that the bilateral negotiation between the lowest bidder and the auctioneer (that takes place when no bids meet the reserve price) results in a final price that is close to the lowest bid.

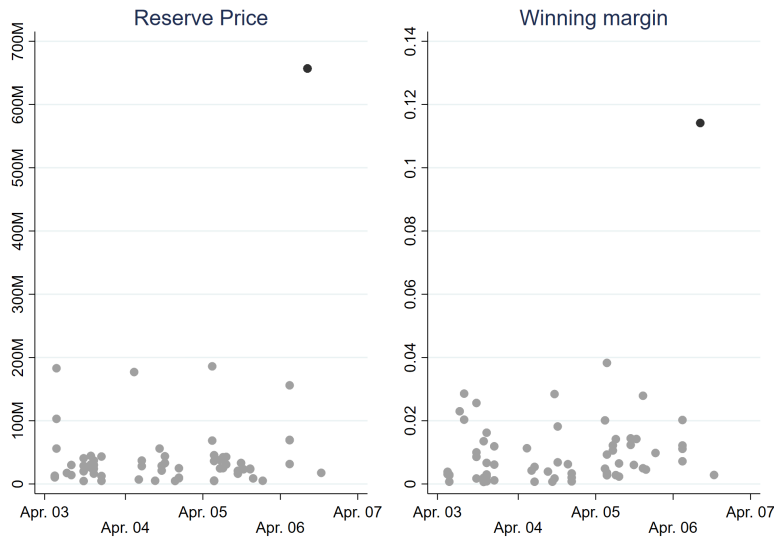


Figure 2: Reserve Price of Auctions (Left Panel) and Difference between Winning Bid and Second Lowest Bid (Right Panel). Left panel of the Figure plots the reserve price of auctions in which collusive firms in Tier B were invited to bid. The right panel plots the difference between the winning bid and the lowest losing bid, as a fraction of the reserve price.

5 Calibrating the Model to the Data

We now assess quantitatively the value of using randomized strategies relative to common-knowledge bids. In order to do so, we calibrate our model of Section 2.2 (bid rotation equilibria) and simulate equilibrium bidding behavior with and without common-knowledge bids using reserve price data from the town of Kumatori.

The key model primitives that we need to calibrate are the cost function, $c(r)$, the discount factor, δ , and the distribution of the reserve price, F_r . We explain below how we calibrate these model primitives for our baseline specification. In Online Appendix OC, we explore the robustness of our results to the choice of the discount factor and the cost parameter. We also report the results of the calibration when we estimate the reserve price distribution F_r without including the Obara Residences auction in the estimation sample.

For the cost function, $c(r)$, we assume a linear functional form, $c(r) = c_0 \times r$, so that construction costs are proportional to the reserve price. We then set c_0 to be 0.75, or equivalently, the profit margin to be 0.25. We note that 25% is the mid-point of the estimated

range of excess cartel profits for Japanese public procurement auctions reported in McMillan (1991).³⁴ Our calibrated parameter is also consistent with the fact that the average winning bid fell to 76.7% of the reserve price for procurement auctions let by the town of Kumatori in fiscal year 2007 after the cartel was prosecuted.

The second parameter that we calibrate is the discount factor between auctions. We calibrate it by setting the annual discount to be 0.9 and adjust it by the number of Tier B auctions per year.³⁵ Our calibrated parameter is approximately 0.9765.

Lastly, for F_r , we fit a Pareto distribution to the realized reserve price by maximum likelihood. Figure 3 plots the histogram of the reserve price for Tier B auctions in our sample as well as the fitted Pareto distribution.³⁶ We use the fitted Pareto distribution when computing the equilibrium outcomes.

With these calibrated parameters, for each n , we can compute \bar{V}_{ck} , \bar{V} , and the associated bid distribution $F^*(\cdot; r)$ using expressions (5), (11) and (6) in the Appendix. We set the number of bidders n to 9, which is the number of Tier B firms.

In the left panel of Figure 4, we overlay on Figure 1, a scatter plot of the realizations of the reserve price and the winning bid: Each \mathbf{x} in the figure corresponds to a realization of the reserve price and the winning bid of a Tier B auction. Similarly, in the right panel of Figure 4, we overlay the reserve price and the winning margin on Figure 1, where the winning margin is the difference between the lowest bid and the second lowest bid.

Under the calibrated parameters, we estimate \bar{W}_{ck} (the expected continuation value of the designated winner under common-knowledge bids) to be about 64 million yen, or about \$640,000. Our estimate of \bar{V}_{ck} (total cartel surplus under common-knowledge bids) is about

³⁴McMillan (1991) estimates excess cartel profits to be between 16% to 33% of the reserve price for public procurement auctions in Japan.

³⁵In particular, we use the average number of auctions in 2006 and 2007. We note that the annual discount factor of 0.9 is the product of the interest rate and the likelihood of the cartel breaking down (for instance, because of detection). Assuming a 2% annual interest rate (the interest rate on 10-year Japanese bonds was about 1.7% during that period), this implies a likelihood of cartel breakdown of about 8.2%. This rate of cartel breakdown is broadly consistent with the findings of Levenstein and Suslow (2011), who report an average cartel duration of 8.1 years.

³⁶The estimated parameters are 1.1675×10^6 for the scale parameter and 0.575 for the shape parameter.




Figure 3: The bar graph corresponds to the histogram of the reserve and the curve corresponds to the estimated Pareto distribution.

525 million yen, or about 5.25 million dollars. This implies that the continuation payoff of the loser who is last in the queue, $\delta^{n-1}\overline{W}_{ck}$, is strictly higher than the deviation gain from undercutting the reserve price, $r - c(r)$, for r less than 211.2 million yen. In the optimal bid rotation equilibria with common-knowledge bids, the winning bid is equal to the reserve price for $r < 211.2$ million, and it is strictly less than the reserve price for $r > 211.2$. The winning bid under common-knowledge bids is illustrated in the left panel of Figure 4 with a dashed line. Note that the actual realized winning bid in the Obara residences auction (630 million yen) lies significantly above the dashed line. Under common-knowledge bids, the optimal winning bid in the Obara residences auction is about 545 million yen.

In contrast, our estimate of the continuation value of the winner without common-knowledge bids, \overline{W} , is about 100 million yen, or about 1 million dollars. Our estimate of total cartel surplus, \overline{V} , is about 820 million yen, or about 8 million dollars. This implies that the continuation value of the loser who is last in the queue, $\delta^{n-1}\overline{W}$, is strictly higher than the deviation payoff of undercutting the reserve price, $r - c(r)$, for $r < 330.4$ million yen. Hence, for values of r less than 330.4 million yen, optimal bid rotation equilibria can




Figure 4: The left panel illustrates the relationship between the winning bid and the reserve price. The shaded region corresponds to the range of winning bids under an optimal equilibrium that attains \bar{V} . The dashed line corresponds to the equilibrium winning bid under common-knowledge bids. The actual realization of the winning bids are marked with **xs**. The right panel plots the winning margin (as a percent of the reserve price) against the reserve price for an optimal equilibrium that attains \bar{V} . The actual realizations are marked with **xs**.

support a winning bid that is equal to the reserve price. For r above 330.4, the optimal winning bid is a non-degenerate distribution. Note that the winning bid in the Kumatori auction is well within the range of values consistent with the winning bid in the optimal equilibria without common knowledge bids. Moreover, we find that the gain from using bidding schemes without common-knowledge bids is substantial: the gain is about 56% ($\bar{V}_{ck} = 525$ million v.s. $\bar{V} = 820$ million).³⁷

In the right panel of Figure 4, we plot the reserve price and the winning margin. For auctions excluding the Obara residences auction, the average realized winning margin is very

³⁷We note that, because we focus on a predetermined allocation schedule (i.e., bid rotation), there is no efficiency gain/loss between using common-knowledge bids and allowing for randomization in our model. Therefore, the impact of using a mechanism with randomization and mediation is limited to a transfer of money from the town of Kumatori (or the taxpayers of Kumatori) to the colluding firms. To put the gains from randomization in perspective, the budget of the town of Kumatori was about 10 billion yen, or about 100 million dollars. The town of Kumatori spends about 10% of its budget on construction.

small (although not exactly equal to zero, as theory predicts), at about 320 thousand yen, or \$3,200.³⁸ The winning margin is also small in relative terms, at less than 1% of the reserve price. On the other hand, the winning margin was large for the Obara residences auction, both in absolute terms as well as in terms of percentage of the reserve price.

Overall, the calibration results suggest that our model of cartel bidding without common-knowledge bids explains the data quite well. The calibrated model predicts the winning margins to be close to zero when the reserve price is low, which matches well the realized winning margins. The calibrated model also predicts that the winning margin can be substantial when the reserve price is as large as that of the Obara residences auction, consistent with actual margins observed in that auction. On the other hand, the equilibrium with common-knowledge bids predicts the winning bid to be much lower than the actual winning bid in the Obara residences auction.

6 Discussion

This paper highlights the practical value of keeping the winning bid private in relaxing IC constraints. We first model a bidding cartel with a mediator in a repeated-game setting, and show that the mediator randomizes the recommend bid to the winner under the optimal bid rotation equilibria.³⁹ We then use detailed court documents to build the case that the winning bid was kept secret in the Obara residences auction. Finally, we provide what we believe to be the first quantitative assessment of the value of using randomization in repeated-game settings, finding that cartel surplus can increase by as much as 50% from randomized bids.

We conclude the paper with a short discussion of: (i) our model's implications for screens

³⁸The winning margin is not exactly equal to zero, which may reflect the cartel's desire to avoid appearing suspicious and attracting the attention of the antitrust authorities.

³⁹In our model, we assumed that the reserve price is public, and the cartel keeps the winning bid private by having the mediator send randomized bid recommendations. We note, however, that a cartel participating in auctions with a secret reserve price may be able to keep the winning bid secret by only letting the designated winner find out the value of the reserve price, but not the designated losers. This mechanism may also have played a role in the Obara residence auction, where losers did not have access to project information until the very end of the bid-submission period.

of collusion; (ii) the costs and benefits of mediation for cartels; and (iii) the connection between models of mediated collusion and algorithmic collusion.

Our model predicts that when a project is relatively small (i.e., r_t is small), the winning bid measured as a fraction of the reserve price ($\frac{\min_i b_{i,t}}{r_t}$) will be close to 1 and the winning margin will be close to zero. This implies that the time-series variance in the winning bids, $\frac{\min_i b_{i,t}}{r_t}$, will be close to zero and the within-auction variance of bids will also be close to zero. These results are consistent with the premise of the price-variance screen (Abrantes-Metz et al., 2012). However, when projects are relatively large, our model predicts that there will be considerable time-series variance in the winning bids that results from randomization. Moreover, the within-auction variance of bids will not be close to zero. Hence, increases in the variance of bids should not be taken as failure of collusion when those increases are associated with increases in project sizes.

Another distinct feature of the equilibrium in Proposition 2 is that the winning bid is isolated whenever it lies below $\bar{b} = \sup_b \text{supp } F^*(\cdot; r)$. Isolated winning bids (or missing bids) are present in the bidding data analyzed in Chassang et al. (2022).⁴⁰ As Chassang et al. (2022) show, such patterns are inconsistent with competition, and hence are a marker of collusive bidding.

This suggests a trade-off for cartels when deciding how to organize themselves. On the one hand, mediation can substantially increase a cartel’s profit, as our case study and calibration exercise suggest. On the other hand, since the bidding patterns that arise from mediation may fail some screens of collusion, mediated cartels may end up facing a higher likelihood of detection and prosecution.

Lastly, we highlight a potential connection between the type of collusion facilitated by intermediaries that we study and collusion facilitated by algorithms (e.g., Harrington, 2018, Calvano et al., 2020, Assad et al., 2020, Asker et al., 2022). Indeed, price-recommendation algorithms may serve the role of mediators in helping firms coordinate their pricing decisions, as a recent investigation by the DOJ suggests (Vogell, 2022). Hence, the study of mediated

⁴⁰Tóth et al. (2014), Imhof et al. (2016) and Clark et al. (2020) document similar patterns. Clark et al. (2020) offer an explanation based on the bidders’ desire to leave some margin of error.

cartels may potentially be useful in understanding the impact of pricing AI's.

Appendix

Proof of Proposition 1. Fix an equilibrium $(\sigma, \mu) \in \Sigma_{\text{ck}}$. Rename the bidders so that under (σ, μ) bidder $i \in N$ wins the auction at period t if $(t + 1) \bmod n = i$. Hence, bidder $i = 1$ wins the auction at $t = 0$, bidder $i = 2$ wins the auction at $t = 1$, etc. For each $i \in N$, let V_i denote the expected discounted payoff that bidder i obtains at $t = 0$ under σ . Let \bar{V}^w be the highest (supremum) payoff that the designated winner at $t = 0$ can obtain under an equilibrium in Σ_{ck} . Note that for each i , it must be that $V_i \leq \delta^{i-1} \bar{V}^w$ (since bidder i is the winner at period $t = i - 1$). Hence, $\sum_i V_i \leq \sum_i \delta^{i-1} \bar{V}^w$. Since this inequality holds for all $(\sigma, \mu) \in \Sigma_{\text{ck}}$, $\bar{V}_{\text{ck}} \leq \sum_i \delta^{i-1} \bar{V}^w$.

Let $b(r_0)$ be the winning bid at time $t = 0$ under (σ, μ) . Note that $b(r_0) \leq c(r_0) + \delta^{n-1} \bar{V}^w$. Indeed, the continuation payoff of bidder n at $t = 0$ can't be larger than $\delta^{n-1} \bar{V}^w$ (since bidder n is the winner at $t = n - 1$). If $b(r_0) > c(r_0) + \delta^{n-1} \bar{V}^w$, bidder n would have a strict incentive to undercut $b(r_0)$. Since $b(r_0)$ must also be lower than r_0 , $b(r_0) \leq \min\{r_0, c(r_0) + \delta^{n-1} \bar{V}^w\}$. Hence, $V_1 \leq \mathbb{E}_{F_r}[\min\{r - c(r), \delta^{n-1} \bar{V}^w\}] + \delta^n \bar{V}^w$. Since the inequality holds for all $(\sigma, \mu) \in \Sigma_{\text{ck}}$,

$$\bar{V}^w \leq \mathbb{E}_{F_r}[\min\{r - c(r), \delta^{n-1} \bar{V}^w\}] + \delta^n \bar{V}^w \iff \bar{V}^w \leq \frac{1}{1 - \delta^n} \mathbb{E}_{F_r}[\min\{r - c(r), \delta^{n-1} \bar{V}^w\}].$$

Let \bar{W}_{ck} be the largest $W \geq 0$ solving

$$W = \frac{1}{1 - \delta^n} \mathbb{E}_{F_r}[\min\{r - c(r), \delta^{n-1} W\}]. \quad (5)$$

We now show that Assumption 1 implies that $\bar{W}_{\text{ck}} > 0$. By Assumption 1, set $\{W > 0 : W \leq \frac{1}{1 - \delta^n} \mathbb{E}_{F_r}[\min\{r - c(r), \delta^{n-1} W\}]\}$ is non-empty. Since both sides of (5) are continuous in W , and since the RHS of (5) is bounded, we have that $\bar{W}_{\text{ck}} = \sup\{W > 0 : W \leq \frac{1}{1 - \delta^n} \mathbb{E}_{F_r}[\min\{r - c(r), \delta^{n-1} W\}]\} > 0$.

Note next that $\bar{V}^w \leq \bar{W}_{\text{ck}}$.⁴¹ We now show that $\bar{V}^w = \bar{W}_{\text{ck}}$. Consider the following strat-

⁴¹Since the right-hand side of (5) is bounded by $\frac{1}{1 - \delta^n} \mathbb{E}_{F_r}[\min\{r - c(r)\}]$, we have that $W >$

egy profile. Along the equilibrium path, at each period t bidder i with $(t+1) \bmod n = i$ bids $b_{i,t} = \min\{r_t, c(r_t) + \delta^{n-1}\bar{W}_{\text{ck}}\}$. All other bidders bid $\min\{r_t, c(r_t) + \delta^{n-1}\bar{W}_{\text{ck}}\} + \epsilon$ for some $\epsilon \in (0, \delta^n \bar{W}_{\text{ck}})$. Defections are punished by Nash reversion.⁴² Note that no player gains by deviating. Indeed, a designated loser who will be the winner in $k \leq n-1$ periods obtains a payoff of $\delta^k \bar{W}_{\text{ck}}$ by placing a losing bid, and would earn $\min\{r_t - c(r_t), \delta^{n-1}\bar{W}_{\text{ck}}\} \leq \delta^k \bar{W}_{\text{ck}}$ by undercutting the winning bid. Moreover, the designated winner does not gain by increasing her bid either: if $c(r_t) + \delta^{n-1}\bar{W}_{\text{ck}} < r_t$, her payoff from bidding $b' \in (c(r_t) + \delta^{n-1}\bar{W}_{\text{ck}}, c(r_t) + \delta^{n-1}\bar{W}_{\text{ck}} + \epsilon)$ is at most $\delta^{n-1}\bar{W}_{\text{ck}} + \epsilon + \delta \times 0 \leq \delta^{n-1}\bar{W}_{\text{ck}} + \delta^n \bar{W}_{\text{ck}}$, while her payoff from following her strategy is $\delta^{n-1}\bar{W}_{\text{ck}} + \delta^n \bar{W}_{\text{ck}}$. Hence, this strategy profile is an equilibrium in Σ_{ck} , giving the winner at $t = 0$ an expected discounted payoff of \bar{W}_{ck} . Therefore, we have that $\bar{V}^w = \bar{W}_{\text{ck}}$. Since $\bar{V}_{\text{ck}} \leq \sum_i \delta^{i-1} \bar{W}_{\text{ck}}$, this equilibrium attains \bar{V}_{ck} . ■

For each reserve price $r \in [\underline{r}, \bar{r}]$ and each $W \geq 0$, let $F(\cdot; r, W)$ be the c.d.f. given as follows. If $c(r) + \delta^{n-1}W \geq r$, $F(\cdot; r, W)$ puts all its mass at r . If $c(r) + \delta^{n-1}W < r$, $F(\cdot; r, W)$ is given by:

$$F(b; r, W) = \begin{cases} 0 & \text{if } b < c(r) + \delta^{n-1}W, \\ 1 - \frac{\delta^{n-1}W}{b - c(r)} & \text{if } b \in [c(r) + \delta^{n-1}W, \min\{r, c(r) + \delta^{n-1}(1 + \delta)W\}], \\ 1 & \text{if } b \geq \min\{r, c(r) + \delta^{n-1}(1 + \delta)W\}. \end{cases} \quad (6)$$

Proof of Proposition 2. Fix an equilibrium $(\sigma, \mu) \in \Sigma$. Rename the bidders so that under (σ, μ) bidder $i \in N$ wins the auction at period t if $(t+1) \bmod n = i$. Hence, bidder $i = 1$ wins the auction at $t = 0$, bidder $i = 2$ wins at $t = 1$, etc. For each $i \in N$, let V_i denote the expected discounted payoff that that bidder i obtains at $t = 0$ under (σ, μ) . Let \hat{V}^w be the highest (supremum) payoff that the winner at time $t = 0$ can obtain under an equilibrium in Σ . By the same arguments as in the proof of Proposition 1, $\sum_i V_i \leq \sum_i \delta^{i-1} \hat{V}^w$. Since this inequality holds for all $(\sigma, \mu) \in \Sigma$, we have that $\bar{V} \leq \sum_i \delta^{i-1} \hat{V}^w$.

⁴² $\frac{1}{1-\delta^n} \mathbb{E}_{F_r} [\min\{r - c(r), \delta^{n-1}W\}]$ for all $W > \bar{W}_{\text{ck}}$.

⁴²We don't specify the mediator's strategy or bidders' beliefs, since in an equilibrium with common-knowledge bids firms' bidding behavior depends solely on the public history.

Let $F(b; r_0)$ be the c.d.f. from which the winning bid at $t = 0$ is drawn under (σ, μ) . Let $\bar{b} = \sup_b \text{supp } F(\cdot; r_0)$ and $\underline{b} = \inf_b \text{supp } F(\cdot; r_0)$. Note that all designated losers at $t = 0$ must place a bid weakly higher than \bar{b} under σ . Moreover, it must be that

$$\forall b < \bar{b}, \quad (1 - F(b; r_0))(b - c(r_0)) \leq \delta^{n-1} \hat{V}^w, \quad (7)$$

$$\forall b \in \text{supp } F(\cdot; r_0), \quad \bar{b} - c(r_0) \leq b - c(r_0) + \delta^n \hat{V}^w. \quad (8)$$

If inequality (7) didn't hold for some $b < \bar{b}$, then bidder n (who wins in $n - 1$ periods) would have a strict incentive to bid b and win the auction with probability $1 - F(b; r_0)$. If inequality (8) didn't hold for $b \in \text{supp } F(\cdot; r_0)$, the winner would have an incentive to bid $\bar{b} - \epsilon$ with $\epsilon \approx 0$ instead of $b < \bar{b}$. Condition (7) implies

$$\underline{b} \leq c(r_0) + \delta^{n-1} \hat{V}^w \text{ and } \forall b < \bar{b}, F(b; r_0) \geq 1 - \frac{\delta^{n-1} \hat{V}^w}{b - c(r_0)}. \quad (9)$$

Condition (8), together with $\bar{b} \leq r_0$, implies

$$\bar{b} \leq \min\{r_0, \underline{b} + \delta^n \hat{V}^w\}. \quad (10)$$

Consider the problem of finding the c.d.f. F that maximizes expected winning bid $\int b dF$, subject to (9) and (10). When $c(r) + \delta^{n-1} \hat{V}^w \geq r_0$, the c.d.f. F that solves this problem puts all its mass at r_0 . When $c(r) + \delta^{n-1} \hat{V}^w < r_0$, the c.d.f. that solves this problem is given by $F(\cdot; r, W)$ in (6), with $r = r_0$ and $W = \hat{V}^w$: indeed, since $F(\cdot; r_0, \hat{V}^w)$ satisfies the inequalities in (9) and (10) with equality, it first-order stochastically dominates any other distribution satisfying (9)-(10).

The arguments above imply that $V^w \leq \mathbb{E}_{F_r} \left[\int b dF(b; r, \hat{V}^w) - c(r) \right] + \delta^n \hat{V}^w$. Since the inequality holds for all $(\sigma, \mu) \in \Sigma$, $\hat{V}^w \leq \frac{1}{1-\delta^n} \mathbb{E}_{F_r} \left[\int b dF(b; r, \hat{V}^w) - c(r) \right]$.

Let $\bar{W} \geq 0$ be the largest solution to

$$W = \frac{1}{1 - \delta^n} \mathbb{E}_{F_r} \left[\int b dF(b; r, W) - c(r) \right], \quad (11)$$

and note that $\hat{V}^w \leq \bar{W}$.⁴³ Note further that $\int bdF(b; r, W) \geq \min\{r, c(r) + \delta^{n-1}W\}$ for each r, W , with strict inequality whenever $c(r) + \delta^{n-1}W < r$.⁴⁴ Hence, $\bar{W} \geq \bar{W}_{ck}$, with strict inequality whenever $c(\bar{r}) + \delta^{n-1}\bar{W}_{ck} < \bar{r}$. For each r , let $F^*(\cdot; r) = F(\cdot; r, \bar{W})$.

We now show that $\hat{V}^w = \bar{W}$. Consider the following strategy profile. Along the equilibrium path, in each period t with $(t + 1) \bmod n = i$, the mediator sends bidding recommendation $\hat{b}_{i,t} = b_t$ to the designated winner, with b_t drawn from c.d.f. $F^*(\cdot; r_t) = F(\cdot; r_t, \bar{W})$. If $\bar{b}_t = \sup_b \text{supp } F^*(\cdot; r_t) = r_t$, all designated losers $j \neq i$ get a bidding recommendation $\hat{b}_{j,t} > \bar{b}_t$. If $\bar{b}_t < r_t$, the mediator recommends a bid $\hat{b}_{j,t} > \bar{b}_t$ to all losers $j \neq i$ who will be winners in $k < n - 1$ periods; and recommends a random bid $\hat{b}_{j',t} = b'$ to the designated loser j' who will win in $n - 1$ periods, with b' drawn uniformly over $[\bar{b}_t, \bar{b}_t + \epsilon]$ for some $\epsilon > 0$ small. Deviations from the mediator's recommendations are punished with Nash reversion. In particular, following a deviation, the mediator sends bidding recommendations $\hat{b}_{i,t} = c(r_t)$ to each bidder i , which all bidders follow. (Firms' beliefs are derived from mediator's strategy.)

Note that designated losers don't gain by deviating. Indeed, a designated loser who will be a winner in $k \leq n - 1$ periods obtains a payoff of $\delta^k \bar{W} \geq \delta^{n-1} \bar{W}$ from following her recommendation, and would get at most $\delta^{n-1} \bar{W}$ from submitting a bid below \bar{b}_t .⁴⁵

Since $\bar{b}_t \leq c(r_t) + \delta^{n-1}(1 + \delta)\bar{W}$, and since $\underline{b}_t = \inf_b \text{supp } F^*(\cdot; r_t) = \min\{r_t, c(r_t) + \delta^{n-1}\bar{W}\}$, the designated winner doesn't gain from deviating to a bid $b' \in (\hat{b}_{i,t}, \bar{b}_t]$ whenever she is recommended $\hat{b}_{i,t} < \bar{b}_t$. Moreover, when $\bar{b}_t = c(r_t) + \delta^{n-1}(1 + \delta)\bar{W} < r_t$, deviations by the designated winner to a bid $\tilde{b} \in (\bar{b}_t, \bar{b}_t + \epsilon]$ are not profitable either whenever $\epsilon > 0$ is small enough: her payoff from such a deviation is $\frac{c(r_t) + \delta^{n-1}(1 + \delta)\bar{W} + \epsilon - \tilde{b}}{\epsilon}(\tilde{b} - c(r_t)) + \delta \times 0$, which, for $\epsilon > 0$ small, is strictly smaller than $\delta^{n-1}(1 + \delta)\bar{W}$ (which, in turn, is equal $\underline{b}_t - c(r) + \delta^n \bar{W}$, bidder i 's payoff when recommended bid \underline{b}_t). Therefore, this strategy profile is an equilibrium in Σ , and gives the designated winner at $t = 0$ a payoff equal to \bar{W} . Hence, $\hat{V}^w = \bar{W}$. Since

⁴³Since the right-hand side of (11) is bounded, $W > \frac{1}{1 - \delta^2} \mathbb{E}_{F_r} [\int bdF(b; r, W) - c(r)]$ for all $W > \bar{W}$.

⁴⁴When $c(r) + \delta^{n-1}W \geq r$, $r = \int bdF(b; r, W)$. When $c(r) + \delta^{n-1}W < r$, $\int bdF(b; r, W) > c(r) + \delta^{n-1}W$, since $c(r) + \delta^{n-1}W$ is the lowest point in the support of $F(b; r, W)$.

⁴⁵The designated loser who will be a winner in $n - 1$ periods gets an expected payoff strictly larger than $\delta^{n-1} \bar{W}$ from following her recommendation when she is recommended to bid $b' = \bar{b}_t$. Hence, she has an even smaller incentive to defect in this case. Note, however, that this bidder receives a recommendation equal to $b' = \bar{b}_t$ with probability zero.

$\bar{V} \leq \sum_i \delta^{i-1} \hat{V}^w$, this equilibrium attains \bar{V} .

Finally, we show that $\bar{V} > \bar{V}_{\text{ck}}$ if and only if $c(\bar{r}) + \delta^{n-1} \frac{\mathbb{E}_{F_r}[r-c(r)]}{1-\delta^n} < \bar{r}$. Our argument above shows that $\bar{W} \geq \bar{W}_{\text{ck}}$, with strict inequality whenever $c(\bar{r}) + \delta^{n-1} \bar{W}_{\text{ck}} < \bar{r}$. Since $\bar{V} = \sum_i \delta^{i-1} \bar{W}$ and $\bar{V}_{\text{ck}} = \sum_i \delta^{i-1} \bar{W}_{\text{ck}}$, we have that $\bar{V} > \bar{V}_{\text{ck}} \iff c(\bar{r}) + \delta^{n-1} \bar{W}_{\text{ck}} < \bar{r}$.

To complete the proof, we show that $c(\bar{r}) + \delta^{n-1} \bar{W}_{\text{ck}} < \bar{r} \iff c(\bar{r}) + \delta^{n-1} \frac{\mathbb{E}_{F_r}[r-c(r)]}{1-\delta^n} < \bar{r}$. Since $\bar{W}_{\text{ck}} \leq \frac{\mathbb{E}[r-c(r)]}{1-\delta^n}$ by the definition of \bar{W}_{ck} (see equation (5)), $c(\bar{r}) + \delta^{n-1} \frac{\mathbb{E}[r-c(r)]}{1-\delta^n} < \bar{r} \implies c(\bar{r}) + \delta^{n-1} \bar{W}_{\text{ck}} < \bar{r}$. Next, we show that $c(\bar{r}) + \delta^{n-1} \frac{\mathbb{E}[r-c(r)]}{1-\delta^n} \geq \bar{r} \implies c(\bar{r}) + \delta^{n-1} \bar{W}_{\text{ck}} \geq \bar{r}$ (the contrapositive of $c(\bar{r}) + \delta^{n-1} \bar{W}_{\text{ck}} < \bar{r} \implies c(\bar{r}) + \delta^{n-1} \frac{\mathbb{E}[r-c(r)]}{1-\delta^n} < \bar{r}$). Suppose $c(\bar{r}) + \delta^{n-1} \frac{\mathbb{E}[r-c(r)]}{1-\delta^n} \geq \bar{r}$ and consider the following strategy profile. At each period t , bidder i such that $(t+1) \bmod n = i$ bids $b_{i,t} = r_t$, and all bidders $j \neq i$ bid $b_{j,t} = r_t + \epsilon$ for some $\epsilon > 0$. Defections are punished by Nash reversion.⁴⁶ The winner's payoff under this strategy profile is $\frac{\mathbb{E}[r-c(r)]}{1-\delta^n}$. At each time t , a loser who will win in $k \leq n-1$ periods obtains a continuation payoff of $\delta^k \frac{\mathbb{E}[r-c(r)]}{1-\delta^n}$ from playing according to her strategy, and gets at most $r_t - c(r_t) + 0$ from defecting. Since $\delta^{n-1} \frac{\mathbb{E}[r-c(r)]}{1-\delta^n} \geq \bar{r} - c(\bar{r})$, and since $r - c(r)$ is increasing in r , it follows that no bidder has an incentive to defect. Hence, this strategy profile is an equilibrium, and so $\bar{W}_{\text{ck}} = \frac{\mathbb{E}[r-c(r)]}{1-\delta^n}$. Therefore, $c(\bar{r}) + \delta^{n-1} \frac{\mathbb{E}[r-c(r)]}{1-\delta^n} \geq \bar{r} \implies c(\bar{r}) + \delta^{n-1} \bar{W}_{\text{ck}} \geq \bar{r}$. ■

Data Availability Statement. The data and code underlying this research are available on Zenodo at <https://doi.org/10.5281/zenodo.11412053>.

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.
- ABRANTES-METZ, R., M. KRATEN, A. METZ, AND G. SEOW (2012): “Libor manipulation?” *Journal of Banking & Finance*, 36, 136–150.
- ASKER, J. (2010): “A study of the internal organization of a bidding cartel,” *The American Economic Review*, 724–762.

⁴⁶We don't specify the mediator's strategy or bidders' beliefs, since under this proposed strategy profile bidders' bids at each period depend solely on the public history.

- ASKER, J., C. FERSHTMAN, AND A. PAKES (2022): “The Impact of AI Design on Pricing,” Tech. rep., Working Paper.
- ASKER, J. AND V. NOCKE (2021): “Collusion, mergers, and related antitrust issues,” in *Handbook of industrial organization*, Elsevier, vol. 5, 177–279.
- ASSAD, S., R. CLARK, D. ERSHOV, AND L. XU (2020): “Algorithmic pricing and competition: Empirical evidence from the German retail gasoline market,” .
- ATHEY, S. AND K. BAGWELL (2001): “Optimal collusion with private information,” *RAND Journal of Economics*, 32, 428–465.
- (2008): “Collusion with persistent cost shocks,” *Econometrica*, 76, 493–540.
- ATHEY, S., K. BAGWELL, AND C. SANCHIRICO (2004): “Collusion and price rigidity,” *The Review of Economic Studies*, 71, 317–349.
- 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.
- AWAYA, Y. AND V. KRISHNA (2016): “On Communication and Collusion,” *American Economic Review*, 106, 285–315.
- 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.
- BERNHEIM, D. AND E. MADSEN (2017): “Price Cutting and Business Stealing in Imperfect Cartels,” *American Economic Review*, 107, 387–424.
- BORENSTEIN, S. AND A. SHEPARD (1996): “Dynamic Pricing in Retail Gasoline Markets,” *The RAND Journal of Economics*, 27, 429–451.
- CALVANO, E., G. CALZOLARI, V. DENICOLO, AND S. PASTORELLO (2020): “Artificial intelligence, algorithmic pricing, and collusion,” *American Economic Review*, 110, 3267–3297.
- CHASSANG, S., K. KAWAI, J. NAKABAYASHI, AND J. ORTNER (2022): “Robust Screens for Non-Competitive 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.

- CHASSANG, S. AND G. PADRÓ I MIQUEL (2019): “Crime, intimidation, and whistleblowing: A theory of inference from unverifiable reports,” *The Review of Economic Studies*, 86, 2530–2553.
- CHIAPPORI, P.-A., S. LEVITT, AND T. GROSECLOSE (2002): “Testing mixed-strategy equilibria when players are heterogeneous: The case of penalty kicks in soccer,” *American Economic Review*, 92, 1138–1151.
- 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.
- CLARK, R., D. COVIELLO, AND A. D. LEVERANO (2020): “Complementary Bidding and the Collusive Arrangement: Evidence from an Antitrust Investigation,” .
- COMPTE, O. (1998): “Communication in repeated games with imperfect private monitoring,” *Econometrica*, 597–626.
- CONLEY, T. G. AND F. DECAROLIS (2016): “Detecting bidders groups in collusive auctions,” *American Economic Journal: Microeconomics*, 8, 1–38.
- EDERER, F., R. HOLDEN, AND M. MEYER (2018): “Gaming and strategic opacity in incentive provision,” *The RAND Journal of Economics*, 49, 819–854.
- ELLISON, G. (1994): “Theories of cartel stability and the joint executive committee,” *The Rand journal of economics*, 37–57.
- GREEN, E. J. AND R. H. PORTER (1984): “Noncooperative collusion under imperfect price information,” *Econometrica: Journal of the Econometric Society*, 87–100.
- HALTIWANGER, J. AND J. HARRINGTON (1991): “The Impact of Cyclical Demand Movements on Collusive Behavior,” *The RAND Journal of Economics*, 22, 89–106.
- HARRINGTON, J. E. (2008): “Detecting cartels,” *Handbook of antitrust economics*, 213, 245.
- (2018): “Developing competition law for collusion by autonomous artificial agents,” *Journal of Competition Law & Economics*, 14, 331–363.
- HARRINGTON, J. E. AND A. SKRZYPACZ (2011): “Private monitoring and communication in cartels: Explaining recent collusive practices,” *The American Economic Review*, 101, 2425–2449.
- HENDRICKS, K. AND R. PORTER (1988): “An Empirical Study of an Auction with Asymmetric Information,” *American Economic Review*, 78, 865–883.

- HORTAÇSU, A. AND I. PERRIGNE (2021): “Empirical perspectives on auctions,” in *Handbook of Industrial Organization*, Elsevier, vol. 5, 81–175.
- IMHOF, D., Y. KARAGÖK, AND S. RUTZ (2016): “Screening for bid-rigging-does it work?” Tech. rep., Université de Fribourg.
- JEHIEL, P. (2015): “On transparency in organizations,” *The Review of Economic Studies*, 82, 736–761.
- KANDORI, M. (2003): “Randomization, communication, and efficiency in repeated games with imperfect public monitoring,” *Econometrica*, 71, 345–353.
- KANDORI, M. AND H. MATSUSHIMA (1998): “Private observation, communication and collusion,” *Econometrica*, 627–652.
- KANDORI, M. AND I. OBARA (2006): “Efficiency in repeated games revisited: The role of private strategies,” *Econometrica*, 74, 499–519.
- 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,” *The Review of Economic Studies*, 90, 376–403.
- KRASNOKUTSKAYA, E. AND K. SEIM (2011): “Bid preference programs and participation in highway procurement auctions,” *The American Economic Review*, 101, 2653–2686.
- LEVENSTEIN, M. (1997): “Price Wars and the Stability of Collusion: A Study of the Pre-World War I Bromine Industry,” *The Journal of Industrial Economics*, 45, 117–137.
- LEVENSTEIN, M. C. AND V. Y. SUSLOW (2011): “Breaking up is hard to do: Determinants of cartel duration,” *The Journal of Law and Economics*, 54, 455–492.
- MARSHALL, R. C. AND L. M. MARX (2007): “Bidder collusion,” *Journal of Economic Theory*, 133, 374–402.
- (2012): *The economics of collusion: Cartels and bidding rings*, Mit Press.
- MASKIN, E. AND J. TIROLE (1988): “A theory of dynamic oligopoly, II: Price competition, kinked demand curves, and Edgeworth cycles,” *Econometrica: Journal of the Econometric Society*, 571–599.
- MCMILLAN, J. (1991): “DANGO: JAPAN’S PRICE-FIXING CONSPIRACIES,” *Economics & Politics*, 3, 201–218.

- 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. AND S. CHASSANG (2018): “Making corruption harder: Asymmetric information, collusion, and crime,” *Journal of Political Economy*, 126, 2108–2133.
- ORTNER, J., T. SUGAYA, AND A. WOLITZKY (2024): “Mediated Collusion,” *Journal of Political Economy*, 132, 1247–1289.
- OURA, M. AND T. KASHIWANO (2015): *Report on Kumatori’s Collusion and Resident Lawsuit*, Nakaishoten.
- PESENDORFER, M. (2000): “A study of collusion in first-price auctions,” *The Review of Economic Studies*, 67, 381–411.
- PORTER, R. (2005): “Detecting Collusion,” *Review of Industrial Organization*, 26, 147–167.
- PORTER, R. H. (1983): “A study of cartel stability: the Joint Executive Committee, 1880–1886,” *The Bell Journal of Economics*, 301–314.
- 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.
- RAHMAN, D. (2012): “But who will monitor the monitor?” *American Economic Review*, 102, 2767–97.
- (2014): “The power of communication,” *American Economic Review*, 104, 3737–51.
- RAHMAN, D. AND I. OBARA (2010): “Mediated partnerships,” *Econometrica*, 78, 285–308.
- ROTEMBERG, J. J. AND G. SALONER (1986): “A supergame-theoretic model of price wars during booms,” *The American Economic Review*, 390–407.
- SCHURTER, K. (2017): “Identification and inference in first-price auctions with collusion,” Tech. rep., working Paper, University of Chicago.
- SKRZYPACZ, A. AND H. HOPENHAYN (2004): “Tacit collusion in repeated auctions,” *Journal of Economic Theory*, 114, 153–169.
- SUGAYA, T. AND A. WOLITZKY (2018): “Maintaining privacy in cartels,” *Journal of Political Economy*, 126, 2569–2607.
- TÓTH, B., M. FAZEKAS, Á. CZIBIK, AND I. J. TÓTH (2014): “Toolkit for detecting collusive bidding in public procurement,” *With examples from Hungary*.

- VOGELL, H. (2022): “Department of Justice Opens Investigation Into Real Estate Tech Company Accused of Collusion with Landlords,” ProPublica <https://www.propublica.org/article/yieldstar-realpage-rent-doj-investigation-antitrust>.
- WALKER, M. AND J. WOODERS (2001): “Minimax Play at Wimbledon,” *American Economic Review*, 91, 1521–1538.
- WANG, Z. (2009): “(Mixed) Strategy in Oligopoly Pricing: Evidence from Gasoline Price Cycles Before and Under a Timing Regulation,” *Journal of Political Economy*, 117, 987–1030.