

Collusion in Auctions with Constrained Bids: Theory and Evidence from Public Procurement

Sylvain Chassang

New York University

Juan Ortner

Boston University

We study the mechanics of cartel enforcement and its interaction with bidding constraints in the context of repeated procurement auctions. Under collusion, bidding constraints weaken cartels by limiting the scope of punishment. This yields a test of collusive behavior exploiting the counterintuitive prediction that introducing minimum prices can lower the winning-bid distribution. The model's predictions are borne out in Japanese procurement data, where we find evidence that minimum prices weakened collusion. A robust design insight is that setting a minimum price at the bottom of the observed winning-bid distribution necessarily improves over a minimum price of zero.

I. Introduction

This paper studies the mechanics of cartel enforcement and its interaction with bidding constraints in the context of repeated procurement auctions. Minimum prices, which place a lower bound on the price at which procurement contracts can be awarded, are frequently used in public pro-

We are especially grateful to Rieko Ishii for pointing us toward appropriate data and sharing with us her knowledge of institutions. The paper benefited from thoughtful discussions with Andrea Prat and Joseph Harrington. We thank Susan Athey, Heski Bar-Isaac, Olivier Compte, Francesco Decarolis, Jan De Loecker, Martin Fiszbein, Chris Flinn, Michael Grubb, Bård Harstad, Philippe Jehiel, Jakub Kastl, Kevin Lang, Bart Lipman, Ludovic Renou, Hamid Sabourian, Stephane Saussier, Paulo Somaini, Giancarlo Spagnolo, and Kjetil

Electronically published June 27, 2019

[*Journal of Political Economy*, 2019, vol. 127, no. 5]

© 2019 by The University of Chicago. All rights reserved. 0022-3808/2019/12705-0007\$10.00

curement. Because minimum prices make price wars less effective, they can also make cartel enforcement more difficult. This leads to the counterintuitive prediction that the introduction of minimum prices may lead to a first-order stochastic dominance drop in the right tail of winning bids. Because this prediction does not arise in competitive environments, it provides a joint test of collusion and of the specific channel we outline: enforcement constraints are binding, and they can be affected by institution design. The model's predictions are borne out in procurement data from Japan, showing that binding enforcement constraints are an empirically relevant determinant of cartel behavior.

From a policy perspective, our findings show that in the presence of colluding bidders, attempts at surplus extraction may foster collusion and reduce the auctioneer's surplus. Inversely, providing minimum-surplus guarantees can limit collusion and improve the auctioneer's surplus. A robust takeaway from our analysis is that introducing a minimum price at the bottom of the distribution of observed bids always dominates setting no minimum price. If there is no collusion, it does not affect the distribution of bids, and if there is collusion, it can only reduce the distribution of bids.

We model firms as repeatedly playing a first-price procurement auction with independent and identically distributed (i.i.d.) production costs. We assume that costs are commonly observed among cartel members and that firms are able to make transfers. In this environment, cartel behavior is limited by self-enforcement constraints: firms must be willing to follow bidding recommendations, as well as make equilibrium transfers. We provide an explicit characterization of optimal cartel behavior: first, contract allocation is efficient, provided that price constraints are not binding; second, cartel members implement the highest possible winning bid for which the sum of deviation temptations is less than the cartel's total pledgeable surplus. This simple characterization lets us delineate distinctive predictions of the model in a transparent manner.

Our main predictions relate the introduction of minimum prices and changes in the distribution of winning bids. In our repeated-game envi-

Storesletten, as well as seminar audiences at Bocconi University, Boston College, Boston University, Brown University, the 2016 Canadian Economic Theory Conference, Collegio Carlo Alberto, Cambridge University, Einaudi Institute for Economics and Finance, London School of Economics, Northwestern University, New York University Industrial Organization Workshop, Paris School of Economics, Queen Mary University of London, the 2015 Society for the Advancement of Economic Theory Conference, Stanford Institute for Theoretical Economics, Stanford University, the 2016 Triangle Microeconomics Conference, the 2016 Sorbonne Business School Conference on Public-Private Arrangements, Universitat Autònoma de Barcelona, University College London, Universitat Pompeu Fabra, Universidad Católica de Chile, University of Essex, University of Oslo, and University of Warwick for helpful feedback and suggestions. Giulia Brancaccio, Matthew Gudgeon, and Christoph Walsh provided outstanding research assistance. Data are provided as supplementary material online.

ronment, minimum prices may weaken cartel discipline by limiting the impact of price wars. When this is the case, sustaining collusive bids above the minimum price becomes more difficult, causing a first-order stochastic dominance drop in the distribution of winning bids to the right of the minimum price. A key observation is that minimum prices have either no impact or the opposite impact in environments without collusion. Under competition, regardless of whether firms have complete or asymmetric information about costs, minimum prices lead to a weak first-order stochastic dominance increase in the right tail of winning bids. This provides a joint test of collusion and of the mechanics of cartel enforcement.

Allowing for entry lets us extend this test and generate new predictions. Under our model, minimum prices reduce the right tail of winning bids conditional on a cartel member winning. However, minimum prices have no effect on the right tail of winning bids conditional on an entrant winning. The reason for this is that cartel members seek to dissuade entry by pinning entrants' winning bids to their production costs. As a result the right tail of entrant winning bids does not depend on minimum prices. In contrast, minimum prices still affect the highest sustainable winning bid among potential cartel winners. This differential impact of minimum prices on cartel and entrant winners allows us to distinguish our model from a competitive model in which the introduction of minimum prices also increases entry. In such a model, the winning bids of both cartel and entrant winners should be affected by minimum prices.

We explore the impact of enforcement constraints on cartel behavior by using data from public procurement auctions taking place in Japanese cities of the Ibaraki prefecture between 2007 and 2016. The introduction of minimum prices in several cities lets us use difference-in-differences and change-in-changes frameworks (Athey and Imbens 2006) to recover the counterfactual distribution of winning bids after the policy change. The data consistently exhibit significant drops in the distribution of winning bids to the right of the minimum price. Using frequent participation as a proxy for cartel membership, we also show that potential cartel members are disproportionately affected by the policy change. These findings imply that (i) there is collusion, (ii) enforcement constraints limit the scope of collusion, and (iii) minimum prices successfully weaken cartel discipline.

Our paper lies at the intersection of different strands of the literature on collusion in auctions. The seminal work of Graham and Marshall (1987) and McAfee and McMillan (1992) studies static collusion in environments in which bidders are able to contract. A key takeaway from this analysis is that the optimal response from the auctioneer should involve setting more constraining reserve prices. In a procurement setting this means reducing the maximum price that the auctioneer is willing to pay. We argue, theoretically and empirically, that when bidders cannot con-

tract and must enforce collusion through repeated game play, minimum price guarantees can weaken cartel enforcement.

An important observation from McAfee and McMillan (1992) is that in the absence of cash transfers, the cartel's ability to collude is severely limited even when commitment is available. A recent strand of work takes seriously the idea that in repeated games, continuation values may successfully replace transfers. Aoyagi (2003) studies bid rotation schemes and allows for communication. Skrzypacz and Hopenhayn (2004; see also Blume and Heidhues [2008]) study collusion in environments without communication and show that while cartel members may still be able to collude, they will remain bounded away from efficient collusion. Athey, Bagwell, and Sanchirico (2004) study collusion in a model of repeated Bertrand competition and emphasize that information revelation costs will push cartel members toward rigid pricing schemes. Because we focus on obedience rather than information revelation constraints, our model simplifies away the strategic issues emphasized in this body of work: we assume complete information among cartel members and transferable utility.¹ This yields a simple characterization of optimal collusion closely related to that obtained in the relational-contracting literature (Bull 1987; Baker, Gibbons, and Murphy 1994, 2002; Levin 2003) and provides a transparent framework in which to study the effect of price constraints on winning bids.

Several recent papers study the impact of the auction format on collusion. Fabra (2003) compares the scope of tacit collusion in uniform and discriminatory auctions. Marshall and Marx (2007) study the role of bidder registration and information revelation procedures in facilitating collusion. Pavlov (2008) and Che and Kim (2009) consider settings in which cartel members can commit to mechanisms and argue that appropriate auction design can successfully limit collusion provided participants have deep pockets and can make *ex ante* payments. Abdulkadiroglu and Chung (2003) make a similar point when bidders are patient.

More closely related to our work, Lee and Sabourian (2011) as well as Mezzetti and Renou (2012) study full implementation in repeated environments using dynamic mechanisms. They show that implementation in all equilibria can be achieved by restricting the set of continuation values available to players to support repeated-game strategies. The incomplete-contracts literature (see, e.g., Bernheim and Whinston 1998; Baker et al. 2002) has suggested that the same mechanism, used in the opposite direction, provides foundations for optimally incomplete contracts. Specifically, it may be optimal to keep contracts more incomplete than needed,

¹ Note that we allow for incomplete information when we study the impact of minimum prices under competition. This ensures that our test of collusion is not driven by this stark modeling assumption.

in order to maintain the range of continuation equilibria needed to enforce efficient behavior. We provide empirical evidence that this theoretical mechanism plays a significant role in practice and can be meaningfully used to affect collusion between firms.

On the empirical side, an important set of papers develops empirical methods to detect collusion (see Harrington [2008] for a detailed survey of prominent empirical strategies and their theoretical underpinnings). Porter and Zona (1993, 1999) contrast the behavior of suspected cartel members with that of non-cartel members, controlling for observables. Bajari and Ye (2003) use excess correlation in bids as a marker of collusion. Porter (1983), along with Ellison (1994; see also Ishii [2008]), uses patterns of price wars of the sort predicted by repeated-game models of oligopoly behavior (Green and Porter 1984; Rotemberg and Saloner 1986) to identify collusion. In a multistage-auction context, Kawai and Nakabayashi (2014) argue that excess switching of second and third bidders across bidding rounds, compared to first and second bidders, is a smoking gun for collusion. We propose a test of collusion exploiting changes in the cartel's ability to implement effective punishments.

Section II sets up our benchmark model of cartels and characterizes optimal cartel behavior. Section III derives empirical predictions from this model that distinguish it from competitive behavior. Section IV briefly extends these results in a setting with entry. Section V takes the model to data. Section VI discusses endogenous participation by cartel members, nonperforming bidders, and robustness tests for our empirical analysis. Appendix OA provides empirical extensions; proofs are collected in appendix OB (apps. OA–OF available as an online supplement).²

II. Self-Enforcing Cartels

Modeling strategy.—McAfee and McMillan's (1992) classic model of cartel behavior focuses on the constraints imposed by information revelation among asymmetrically informed cartel members. Instead, we are interested in the enforcement of cartel recommendations through repeated play. Viewed from the mechanism design perspective of Myerson (1986), McAfee and McMillan (1992) focus on truthful revelation, while we focus on obedience. The implications of the two frictions turn out to be different: interpreted in a procurement context, McAfee and McMillan (1992) show that collusion makes lower maximum prices desirable; we argue that higher minimum prices may help weaken cartels.

This different emphasis is reflected in our modeling choices. We have three main goals:

² The appendixes also show how to accommodate nonperforming bidders, develop a model of endogenous participation by cartel members, tackle error in measurement, and provide a calibration exercise assessing the magnitude of our findings.

- i. we want to provide transparent intuition on how bidding constraints, here minimum prices, can affect cartel behavior and the distribution of bids;
- ii. we want to assess empirically whether enforcement constraints are a significant determinant of cartel behavior;
- iii. we want to exploit this understanding of cartel behavior to derive a test of collusion.

Given those goals, we use a tractable complete-information model of collusion when fleshing out implications of our H_1 hypothesis (“there is collusion and enforcement constraints are binding”). To ensure that our test is not dependent on this simplification, we allow for more general informational environments when we characterize behavior under our H_0 hypothesis (“there is no collusion”).

A. *The Model*

Players and payoffs.—Each period $t \in \mathbb{N}$, a buyer procures a single unit of a good through a first-price auction described below. A set $N = \{1, \dots, n\}$ of long-lived firms is present in the market. In each period t , a subset $\hat{N}_t \subset N$ of firms is able to participate in the auction. Participant set \hat{N}_t is exogenous, i.i.d. over time.

We think of this set of participating firms as those potentially able to produce in the current period.³ In period t , each participating firm $i \in \hat{N}_t$ can deliver the good at a cost $c_{i,t}$. Firm i 's cost $c_{i,t}$ is drawn i.i.d. across time periods from a cumulative distribution function F_i with support $[\underline{c}, \bar{c}]$ and density f_i with $f_i(c) > 0$ for all $c \in [\underline{c}, \bar{c}]$.⁴

Firms are able to send transfers to each other, regardless of whether or not they participate in the auction. We denote by $T_{i,t}$ the net transfer received or sent by firm i . Let $x_{i,t} \in \{0, 1\}$ denote whether firm i wins the procurement contract in period t . Let $b_{i,t}$ denote the firm's bid. We assume that firms have quasi-linear preferences, so that firm i 's overall stage game payoff is

$$\pi_{i,t} = x_{i,t}(b_{i,t} - c_{i,t}) + T_{i,t}.$$

Firms value future payoffs using a common discount factor $\delta < 1$.

The stage game.—The procurement contract is allocated according to a first-price auction with constrained bids. Specifically, each participant must submit a bid b_i in the range $[p, r]$, where r is a maximum (or reserve)

³ We consider endogenous participation by entrants in Sec. IV, and endogenous participation by cartel members in app. OD.

⁴ While we allow firms to be asymmetric, for simplicity we assume that the cost distributions $\{F_i\}_{i \in N}$ all have the same support.

price and $p < r$ is a minimum price. Bids outside this range are discarded. The winner is the lowest bidder, with ties broken randomly. The winning firm then delivers the good at the price it bid. For simplicity, we assume that $r \geq \bar{c}$.⁵

To keep the model tractable and to focus on how enforcement constraints affect bidding behavior, we assume that all firms belong to the cartel, and firms in the cartel observe one another's production costs. In addition, we assume that payoffs are transferable.⁶ The timing of information and decisions within period t is as follows.

1. The set of participating firms \hat{N}_t is drawn and observed by all cartel members.
2. The production costs $\mathbf{c}_t = (c_{i,t})_{i \in \hat{N}_t}$ of participating firms are publicly observed by cartel members.
3. Participating firms $i \in \hat{N}_t$ submit public bids $\mathbf{b}_t = (b_{i,t})_{i \in \hat{N}_t}$. This yields allocation $\mathbf{x}_t = (x_{i,t})_{i \in \hat{N}_t} \in [0, 1]^{\hat{N}_t}$ such that if $b_{j,t} > b_{i,t}$ for all $j \in \hat{N}_t \setminus \{i\}$ then $x_{i,t} = 1$; if there exists $j \in \hat{N}_t \setminus \{i\}$ with $b_{j,t} < b_{i,t}$, then $x_{i,t} = 0$. In the case of ties, we follow Athey and Bagwell (2001) and let the bidders jointly determine the allocation. This simplifies the analysis but requires some formalism (which can be skipped at moderate cost to understanding). We allow bidders to simultaneously pick numbers $\gamma_t = (\gamma_{i,t})_{i \in \hat{N}_t}$ with $\gamma_{i,t} \in [0, 1]$ for all i, t . When lowest bids are tied, the allocation to a lowest bidder i is

$$x_{i,t} = \frac{\gamma_{i,t}}{\sum_{\{j \in \hat{N}_t \text{ such that } b_{j,t} = \min_i b_{i,t}\}} \gamma_{j,t}}.$$

4. Firms make transfers $T_{i,t}$. Positive transfers are always accepted and only negative transfers will be subject to an incentive compatibility condition. We require exact budget balance within each period at the overall cartel level; that is, $\sum_{j \in \hat{N}_t} T_{j,t} = 0$.

Our model is intended to capture commonly observed features of public construction procurement (see McMillan [1991] for a reference). Governments need to procure construction services on an ongoing basis. They face a limited and stable set of firms that can potentially perform the work, a subset of which participates regularly. Legislation frequently requires participants to register, and governments make bids and out-

⁵ This assumption is largely verified in our data. Indeed, 99.02 percent of auctions in our data have a winner.

⁶ The assumption that firms can transfer money is not unrealistic. Indeed, many known cartels used monetary transfers; see, e.g., Pesendorfer (2000), Asker (2010), and Harrington and Skrzypacz (2011). In practice these transfers can be made in ways that make it difficult for authorities to detect them, such as subcontracting between cartel members or, in the case of cartels for intermediate goods, between-firms sales.

comes public after each auction is completed. The repeated and public nature of the interaction makes collusion a realistic concern.

Note that procurement auctions with minimum acceptable bids are frequently used in practice. For instance, auctions with minimum bids are used for procurement of public works in several countries in the European Union and by local governments in Japan. The common rationale for introducing minimum bids in the auction is to limit the incidence of strategic default by nonperforming contractors (Calveras, Ganuza, and Hauk 2004; Decarolis 2018). Appendix OC extends our model to allow for nonperforming contractors.⁷

The repeated game.—Interaction is repeated and firms can use the promise of continued collusion to enforce obedient bidding and transfers. Formally, bids and transfers need to be part of a subgame perfect equilibrium of the repeated game among firms.

The history among cartel members at the beginning of time t is

$$h_t = \{\mathbf{c}_s, \mathbf{b}_s, \gamma_s, \mathbf{x}_s, \mathbf{T}_s\}_{s=0}^{t-1}.$$

Let \mathcal{H}^t denote the set of period t histories and $\mathcal{H} = \cup_{t \geq 0} \mathcal{H}^t$ denote the set of all histories. Our solution concept is subgame perfect equilibrium (SPE), with strategies

$$\sigma_i : h_t \mapsto (b_{i,t}(c_t), \gamma_{i,t}(c_t), T_{i,t}(c_t, \mathbf{b}_t, \gamma_t, \mathbf{x}_t))$$

such that bids $(b_{i,t}(c_t), \gamma_{i,t}(c_t))$ and transfers $T_{i,t}(c_t, \mathbf{b}_t, \gamma_t, \mathbf{x}_t)$ can depend on all public information available at the time of decision making.

Denote by Σ the set of SPEs in the repeated stage game. Let

$$V(\sigma, h_t) = \mathbb{E}_\sigma \left[\sum_{s \geq 0} \delta^s \sum_{i \in N_{t+s}} x_{i,t+s} (b_{i,t+s} - c_{i,t+s}) | h_t \right]$$

denote the total surplus generated under equilibrium σ conditional on history h_t . We denote by

$$\bar{V}_p \equiv \max_{\sigma \in \Sigma} V(\sigma, h_0)$$

the highest surplus sustainable in equilibrium.⁸ We emphasize that this highest equilibrium value depends on minimum price p .

We say that a strategy σ_i is noncollusive whenever bids at history h_t depend only on the costs of participating bidders at history h_t , excluding

⁷ Such firms can be viewed as entrants with zero costs, producing a worthless good. Since our model and predictions focus exclusively on the bidders' side of the market, our predictions regarding bid distributions hold regardless of whether such nonperforming firms are included in the model. However, the presence of nonperforming firms would affect the welfare of the auctioneer.

⁸ The existence of surplus-maximizing and surplus-minimizing equilibria follows from proposition 2.5.2 in Mailath and Samuelson (2006).

the remainder of the public history, and the identity of other bidders: $\sigma_i(h_t) = \hat{\sigma}_i(c_{i,t}, \{c_{j,t}\}_{j \in \hat{N}_i \setminus i})$ for all histories h_t . Since there is no persistent state in this game, noncollusive strategies coincide with Markov perfect strategies.

DEFINITION 1 (Collusive and competitive environments). We say that we are in a collusive environment if firms play a Pareto-efficient SPE, that is, an SPE that attains \bar{V}_p . We say that we are in a competitive environment if firms play a weakly undominated SPE in noncollusive strategies.

Under complete information, the unique competitive equilibrium outcome is such that the winning bid is equal to the maximum between the second-lowest cost and the minimum price. The contract is allocated to the bidder with the lowest cost whenever the winning bid is above the minimum price, and is allocated randomly among all bidders with cost below the minimum price when the winning bid is equal to the minimum price.

B. Optimal Collusion

Given a history h_t and a strategy profile σ , we denote by $(\beta(\mathbf{c}_i|h_t, \sigma), \gamma(\mathbf{c}_i|h_t, \sigma))$ the bidding profile induced by strategy profile σ at history h_t as a function of realized costs \mathbf{c}_i .

LEMMA 1 (Stationarity). Consider a subgame perfect equilibrium σ that attains \bar{V}_p . Equilibrium σ delivers surplus $V(\sigma, h_t) = \bar{V}_p$ after all on-path histories h_t . There exists a fixed bidding profile (β^*, γ^*) such that, in a Pareto-efficient equilibrium, firms bid $(\beta(\mathbf{c}_i|h_t, \sigma), \gamma(\mathbf{c}_i|h_t, \sigma)) = (\beta^*(\mathbf{c}_i), \gamma^*(\mathbf{c}_i))$ after all on-path histories h_t .

For any $i \in N$ and any $\sigma \in \Sigma$, let

$$V_i(\sigma, h_t) = \mathbb{E}_\sigma \left[\sum_{s \geq 0} \delta^s (x_{i,t+s}(b_{i,t+s} - c_{i,t+s}) + T_{i,t+s}) | h_t \right]$$

denote the expected discounted payoff that firm i gets in equilibrium σ conditional on history h_t . For each $i \in N$, let

$$\underline{V}_{i,p} \equiv \min_{\sigma \in \Sigma} V_i(\sigma, h_0)$$

denote the lowest possible equilibrium payoff for firm i .

Given a bidding profile (β, γ) , let us denote by $\beta^W(\mathbf{c})$ and $\mathbf{x}(\mathbf{c})$ the induced winning bid and allocation profile for realized costs \mathbf{c} . For each firm i , we define

$$\rho_i(\beta^W, \gamma, \mathbf{x})(\mathbf{c}) \equiv \mathbf{1}_{\beta^W(\mathbf{c}) > p} + \frac{\mathbf{1}_{\beta^W(\mathbf{c}) = p}}{1 + \sum_{j \in \hat{N} \setminus \{i\} : x_j(\mathbf{c}) > 0} \gamma_j(\mathbf{c})}.$$

The term $\rho_i(\beta^W, \gamma, \mathbf{x})(\mathbf{c})$ corresponds to a deviator's highest possible chance of winning the contract by attempting to undercut the equilibrium winning bid.

LEMMA 2 (Enforceable bidding). A winning bid profile $\beta^W(\mathbf{c})$ and an allocation $\mathbf{x}(\mathbf{c})$ are sustainable in SPE if and only if, for all \mathbf{c} ,⁹

$$\sum_{i \in \hat{N}} (\rho_i(\beta^W, \gamma, \mathbf{x})(\mathbf{c}) - x_i(\mathbf{c})) [\beta^W(\mathbf{c}) - c_i]^+ + x_i(\mathbf{c}) [\beta^W(\mathbf{c}) - c_i]^- \leq \delta \left(\bar{V}_p - \sum_{i \in \hat{N}} \underline{V}_{i,p} \right). \quad (1)$$

As in Levin (2003), a bidding profile can be implemented in SPE if and only if the sum of deviation temptations (from both bidders abstaining from bidding above their cost and bidders having to bid below their cost) is less than or equal to the total pledgeable surplus $\delta(\bar{V}_p - \sum_{i \in \hat{N}} \underline{V}_{i,p})$, that is, the difference between the highest possible aggregate continuation surplus and the sum of minimal individual continuation surpluses guaranteed to each player in equilibrium. When this condition is satisfied, we can always find feasible transfers that provide bidders with incentives not to deviate.

For each cost realization \mathbf{c} , let $\mathbf{x}^*(\mathbf{c})$ denote the efficient allocation. It allocates the procurement contract to the participating firm with the lowest cost (ties are broken randomly). We define

$$b_p^*(\mathbf{c}) \equiv \sup \left\{ b \leq r : \sum_{i \in \hat{N}} [1 - x_i^*(\mathbf{c})] [b - c_i]^+ \leq \delta \left(\bar{V}_p - \sum_{i \in \hat{N}} \underline{V}_{i,p} \right) \right\}.$$

For cost realizations \mathbf{c} with $b_p^*(\mathbf{c}) > p$, this value is the highest enforceable winning bid when the cartel allocates the good efficiently. Indeed, by equation (1), when the allocation is efficient, a winning bid $b > p$ is sustainable if and only if $\sum_{i \in \hat{N}} [1 - x_i^*(\mathbf{c})] [b - c_i]^+ \leq \delta(\bar{V}_p - \sum_{i \in \hat{N}} \underline{V}_{i,p})$. Note that $b_p^*(\mathbf{c})$ is always weakly greater than the second-lowest cost.

PROPOSITION 1. On the equilibrium path, the bidding strategy in any SPE that attains \bar{V}_p sets winning bid $\beta_p^*(\mathbf{c}) = \max\{b_p^*(\mathbf{c}), p\}$ in every period. Moreover, the allocation is conditionally efficient: whenever $\beta_p^*(\mathbf{c}) > p$, the contract is allocated to the bidder with the lowest procurement cost.

This result follows from obedience constraint (1). Bid $\beta_p^*(\mathbf{c})$ is the highest enforceable bid. Furthermore, allocating the good efficiently increases the surplus accruing to the cartel while also relaxing obedience constraint (1). Indeed, the lowest-cost bidder has the largest incentives to undercut other bidders.

The following bidding profile implements the optimal collusive scheme when $\beta_p^*(\mathbf{c}) > p$. The firm with the lowest cost bids $\beta_p^*(\mathbf{c})$ and wins the contract at this price. At least one other firm bids immediately above $\beta_p^*(\mathbf{c})$.¹⁰

⁹ We note that cost vector $\mathbf{c} = (c_i)_{i \in \hat{N}}$ uniquely determines the set \hat{N} of participating bidders.

¹⁰ Tie-breaking profile γ is needed to make this statement precise.

COROLLARY 1. The following comparative statics hold:

- i. winning bid $\beta_p^*(\mathbf{c})$ is decreasing in the procurement cost of each participating firm $i \in \hat{N}$;
- ii. for all $N_0 \subsetneq N$ and all $i \in N \setminus N_0$, $\mathbb{E}[\beta_p^*(\mathbf{c}) | \hat{N} = N_0] \geq \mathbb{E}[\beta_p^*(\mathbf{c}) | \hat{N} = N_0 \cup \{i\}]$.

Corollary 1 shows that the winning bid is decreasing in bidders' procurement costs and in the set of participating bidders. Indeed, obedience constraint (1) gets tightened with a decrease in participating firms' costs or an increase in the number of participating firms.

The firm's behavior in a competitive environment with complete information is an immediate corollary to proposition 1: it coincides with collusive behavior in a game with discount factor $\delta = 0$. For any profile of cost realizations \mathbf{c} , let $c_{(2)}$ denote the second-lowest cost.

COROLLARY 2 (Behavior under competition). In a competitive environment, the winning bid is $\beta_p^{\text{comp}}(\mathbf{c}) = \max\{p, c_{(2)}\}$.

We now clarify how minimum prices affect the set of payoffs that firms can sustain in SPE. We denote by $\beta_0^*(\underline{c})$ the lowest bid in a Pareto-efficient SPE when there is no minimum price. We note that $\beta_0^*(\underline{c})$ is observable from data: it is the lowest equilibrium winning bid.

PROPOSITION 2 (Worst-case punishment).

- i. For all $i \in N$, $\underline{V}_{i,0} = 0$ and $\underline{V}_{i,p} > 0$ whenever $p > \underline{c}$;
- ii. there exists $\eta > 0$ such that for all $p \in [\beta_0^*(\underline{c}), \beta_0^*(\underline{c}) + \eta]$, $\bar{V}_p - \sum_{i \in N} \underline{V}_{i,p} < \bar{V}_0 - \sum_{i \in N} \underline{V}_{i,0}$.

Proposition 2i shows that with no minimum price, the cartel can force a firm's payoff down to a minmax value of 0, but that minmax values are bounded away from zero when the minimum price is within the support of procurement costs. Proposition 2ii establishes that the pledgeable surplus $\bar{V}_p - \sum \underline{V}_{i,p}$ that the cartel can use to provide incentives decreases after introducing a low minimum price. The reason for this is that a minimum price p in the neighborhood of $\beta_0^*(\underline{c})$ increases the firms' lowest equilibrium value $\underline{V}_{i,p}$ by an amount bounded away from 0, even for $\eta > 0$ small. This tightens enforcement constraint (1) and reduces the bids that the cartel can sustain in equilibrium.¹¹

III. Empirical Implications

The effect of minimum prices on the distribution of bids.—We now delineate several empirical implications of our model. Specifically, we contrast the

¹¹ In contrast, a minimum price significantly larger than $\beta_0^*(\underline{c})$ may increase the cartel's pledgeable surplus.

effect that a minimum price has on the distribution of winning bids under competition and under collusion.

PROPOSITION 3 (The effect of minimum prices on bids). Under collusion, minimum prices can induce a first-order stochastic dominance drop in the right tail of winning bids. Under competition, minimum prices do not affect the right tail of winning bids. Formally,

- i. there exists $\eta > 0$ such that, for all $p \in [\beta_0^*(\underline{c}), \beta_0^*(\underline{c}) + \eta]$ and all $q > p$,

$$\text{prob}(\beta_p^* \geq q | \beta_p^* \geq p) \leq \text{prob}(\beta_0^* \geq q | \beta_0^* \geq p),$$

the inequality being strict for some $q > p$ whenever $\text{prob}(\beta_0^* < r) > 0$;

- ii. for all $p > 0$ and all $q > p$,

$$\text{prob}(\beta_p^{\text{comp}} \geq q | \beta_p^{\text{comp}} > p) = \text{prob}(\beta_0^{\text{comp}} \geq q | \beta_0^{\text{comp}} > p).^{12}$$

Consider now equilibrium bidding data from auctions without a minimum price. Bidders may be either collusive or competitive. Let β_0^{obs} denote the lowest observed winning bid. Since competitive bids are not affected when the minimum price is below the observed distribution of winning bids, we obtain the following corollary.

COROLLARY 3 (Robust policy takeaway). Regardless of whether there is collusion or not, setting a minimum price $p \leq \beta_0^{\text{obs}}$ causes a weak first-order dominance drop in procurement costs.

This corollary is a robust policy takeaway. Setting a minimum price at the bottom of the distribution of observed winning bids weakly dominates setting no minimum price. Setting a minimum price strictly within the distribution of observed winning bids may increase procurement costs if there was little or no collusion.¹³

One design subtlety worth emphasizing is that the minimum prices studied in this paper are not indexed on bids. In some settings (e.g., Italy), minimum prices are set as an increasing function of submitted bids, for example, a quantile of submitted bids (Conley and Decarolis 2016; Deca-

¹² Conditioning on a strict inequality is meaningful because the distribution of winning bids may have mass points at the minimum price, which we need to correctly take care of. When the mass of bids at the minimum price is small, the conditioning events in parts i and ii of proposition 3 coincide. In data from our lead example city, Tsuchiura, 1.2 percent of auctions with a minimum price have a winning bid equal to the minimum price.

¹³ Note that the rule of thumb described in corollary 3 can be extended if the distribution of costs changes over time. Minimum prices should be adjusted to be as high as possible without being binding. If a mass of bids is concentrated at the minimum price, the minimum price should be lowered.

rolis 2018). We expect such minimum-price policies to be less effective than fixed minimum prices in deterring collusion: by coordinating on low bids, cartel bidders can still bring minimum prices down, limiting the effect that the policy has on punishments.

Proposition 3 provides a joint test of collusion and of the fact that cartel enforcement constraints are binding. Consider the introduction of a minimum price close to the minimum observed winning bid. Under collusion, the introduction of such a minimum price will lead to a first-order stochastic dominance drop in the distribution of winning bids to the right of the minimum price. Under competition, the introduction of minimum prices will lead to a (weak) first-order stochastic dominance increase in the distribution of winning bids.

We emphasize that the predictions under competition in proposition 3ii do not rely on the assumption that firms can make monetary transfers: indeed, no transfers are used in competitive equilibrium. Moreover, the results under collusion in proposition 3i also continue to hold in the absence of monetary transfers: minimum prices still reduce the cartel’s ability to punish deviators, thereby lowering the highest sustainable bid.

We now strengthen this test by showing that proposition 3ii extends to asymmetric-information settings.

Competitive comparative statics under asymmetric information.—We assume now that firms are privately informed about their own procurement cost. For simplicity, we assume that firms are symmetric, with $F_i = F$ for all $i \in N$. Let $b_0^{\text{AI}} : [\underline{c}, \bar{c}] \rightarrow \mathbb{R}_+$ denote the equilibrium bidding function in the unique symmetric equilibrium of the first-price procurement auction with reserve price r and no minimum price.

PROPOSITION 4. Under private information, a first-price auction with reserve price r and minimum price $p < \min\{r, \bar{c}\}$ has a unique symmetric equilibrium with bidding function b_p^{AI} .

- i. If $b_0^{\text{AI}}(\underline{c}) \geq p$, then $b_p^{\text{AI}}(c) = b_0^{\text{AI}}(c)$ for all $c \in [\underline{c}, \bar{c}]$.
- ii. If $b_0^{\text{AI}}(\underline{c}) < p$, there exists a cutoff $\hat{c} \in (\underline{c}, \bar{c})$ with $b_0^{\text{AI}}(\hat{c}) > p$ such that

$$b_p^{\text{AI}}(c) = \begin{cases} b_0^{\text{AI}}(c) & \text{if } c \geq \hat{c}, \\ p & \text{if } c < \hat{c}. \end{cases}$$

An immediate corollary of proposition 4 is that minimum prices can only yield a first-order stochastic dominance increase in the right tail of winning bids. Let $\beta_p^{\text{AI}}(c) \equiv \min_i b_p^{\text{AI}}(c_i)$ denote the winning bid.

COROLLARY 4. For all $p > 0$ and all $q > p$,

$$\text{prob}(\beta_p^{\text{AI}} \geq q | \beta_p^{\text{AI}} > p) \geq \text{prob}(\beta_0^{\text{AI}} \geq q | \beta_0^{\text{AI}} > p).$$

This strengthens the test of collusion provided in proposition 3. A first-order stochastic dominance drop in the right tail of winning bids cannot be explained away by a competitive model with incomplete information.

Similar results continue to hold when bidders are asymmetric and face interdependent costs. Under competition, setting a binding minimum price creates a mass of bids at the minimum price, and a gap in the support of the winning-bid distribution just above the minimum price. As a result, in these competitive environments, a minimum price cannot generate a first-order stochastic dominance drop in the right tail of winning bids.

IV. Entry

We now extend the model of Section II to allow for endogenous entry. The goal of this extension is twofold. First, we want to show that the testable predictions in proposition 3 continue to hold when non-cartel members can participate. Second, this extension allows us to derive additional predictions on the differential effect of minimum prices on cartel members and entrants. These additional predictions are important since they let us distinguish our model from a competitive one in which minimum prices somehow increase entry.¹⁴

We assume that in each period t , a short-lived firm may bid in the auction along with participating cartel members \hat{N}_t . To participate, the short-lived firm has to pay an entry cost k_t drawn i.i.d. over time from a cumulative distribution F_k with support $[0, \bar{k}]$. The distribution of entry costs may have a point mass at 0. We let $E_t \in \{0, 1\}$ denote the entry decision of the short-lived firm in period t , with $E_t = 1$ denoting entry.

Upon paying the entry cost, the short-lived firm learns its cost $c_{e,t}$ for delivering the good, which is drawn i.i.d. from a cumulative distribution function F_c with support $[\underline{c}, \bar{c}]$ and density f_c . We assume that the short-lived firm's entry decision and its procurement cost upon entry $c_{e,t}$ are publicly observed.

The timing of information and decisions within each period t is as follows:

1. The short-lived firm's entry cost k_t is drawn and privately observed. The short-lived firm makes entry decision E_t , which is observed by cartel members.
2. The set of participating cartel members \hat{N}_t is drawn and observed by both cartel members and the short-lived firm.
3. The production costs \mathbf{c}_t of participating firms are drawn and publicly observed by all firms.

¹⁴ See app. OD for a model of endogenous participation by cartel members.

4. Participating firms submit public bids $\mathbf{b}_t = (b_{i,t})$ and numbers $\gamma = (\gamma_{i,t})$ with $\gamma_{i,t} \in [0, 1]$, resulting in allocation $\mathbf{x}_t = (x_{i,t})$.¹⁵
5. Cartel members make transfers $T_{i,t}$ to one another.

The history at the beginning of time t is now $h_t = \{E_s, \mathbf{c}_s, \mathbf{b}_s, \gamma_s, \mathbf{x}_s, \mathbf{T}_s\}_{s=0}^{t-1}$, and is observed by both cartel members and entrants. Let \mathcal{H}^t denote the set of period t public histories and $\mathcal{H} = \cup_{t \geq 0} \mathcal{H}^t$ denote the set of all histories. Our solution concept is subgame perfect public equilibrium, with strategies

$$\sigma_i : h_t \mapsto (b_{i,t}(E_t, \mathbf{c}_t), \gamma_{i,t}(E_t, \mathbf{c}_t), T_{i,t}(E_t, \mathbf{c}_t, \mathbf{b}_t, \gamma_t, \mathbf{x}_t))$$

for cartel members and strategies

$$\sigma_e : h_t \mapsto (E_t(k_t), b_{e,t}(k_t, \mathbf{c}_t), \gamma_{e,t}(k_t, \mathbf{c}_t))$$

for the short-lived firms.

We note that the cartel in this model is not all-inclusive. In each period participating cartel members compete against short-lived entrants.¹⁶

The analysis of this model is essentially identical to that of the model of Section II except that now the cartel will deter entry in addition to enforcing collusive bidding. Given that procurement costs are observed after entry, entry depends only on cost k_t and takes a threshold form. Entrants with entry costs above a certain level are deterred from entering, while entrants with an entry cost below this threshold participate in the auction.

For concision, we focus on extending the main empirical predictions of our model. Appendix OB provides further details on optimal cartel behavior.

PROPOSITION 5 (The effect of minimum prices on bids).

- i. Under collusion, there exists $\eta > 0$ such that for all $p \in [\beta_0^*(\mathcal{L}), \beta_0^*(\mathcal{L}) + \eta]$, $q > p$, and $E \in \{0, 1\}$,

$$\text{prob}(\beta_p^* \geq q | \beta_p^* \geq p, E) \leq \text{prob}(\beta_0^* \geq q | \beta_0^* \geq p, E).$$

- ii. Under competition, for all $p > 0$, $q > p$, and $E \in \{0, 1\}$,

$$\text{prob}(\beta_p^{\text{comp}} \geq q | \beta_p^{\text{comp}} > p, E) = \text{prob}(\beta_0^{\text{comp}} \geq q | \beta_0^{\text{comp}} > p, E).$$

In other words, the contrasting comparative statics of proposition 3 continue to hold conditional on the entrant's entry decision.¹⁷

¹⁵ The allocation is determined in the same way as in Sec. II.

¹⁶ See Hendricks, Porter, and Tan (2008) or Decarolis, Goldmanis, and Penta (2016) for recent analyses of cartels that are not all-inclusive.

¹⁷ Note that similarly to proposition 4, a competitive model with incomplete information and entry cannot explain the predictions of proposition 5i. Indeed, in such a model

Differential impact.—A notable new prediction is that under collusion, minimum prices have a different impact on cartel and entrant winners.

PROPOSITION 6 (Differential effect of minimum prices on bids). Under collusion, there exists $\eta > 0$ such that, for all $p \in [\beta_0^*(\mathbf{c}), \beta_0^*(\mathbf{c}) + \eta]$ and all $q > p$,

- i. $\text{prob}(\beta_p^* \geq q | \beta_p^* \geq p, \text{ cartel wins}) \leq \text{prob}(\beta_0^* \geq q | \beta_0^* \geq p, \text{ cartel wins});$
- ii. $\text{prob}(\beta_p^* \geq q | \beta_p^* > p, \text{ entrant wins}) = \text{prob}(\beta_0^* \geq q | \beta_0^* > p, \text{ entrant wins}).$

In words, minimum prices should only affect the right tail of winning bids when the winners are cartel members. Importantly, this result holds without conditioning on the set of participants. This is practically valuable since in many settings, only winning bidders are observed. The intuition behind this stark prediction is straightforward. Since costs are complete information, under optimal entry deterrence, entrants win either at the minimum price or at their production cost. As a result, the right tail of winning bids conditional on an entrant being the winner is independent of the cartel's pledgeable surplus and independent of the minimum price. Appendix OE shows how to extend proposition 6 to settings in which cartel membership is measured with error.

Qualitative implications of proposition 6 continue to hold if cartel members only get a noisy signal of the entrants' production costs. Indeed, the winning bid of entrants may even increase after the introduction of a minimum price. The reason for this is that when the entrant's cost is noisily observed, optimal entry prevention may require cartel members to bid below their cost in the event of entry. When the cartel's enforcement power is weakened by minimum prices, it becomes more difficult for the cartel to sustain such low bids following entry.

Alternative models of entry.—Proposition 6 is important because it lets us distinguish our model from models in which minimum prices are associated with greater (potentially unobserved) entry, but for reasons unrelated to collusion—for instance, because of media coverage of the policy change. Under such a model, minimum prices could reduce the distribution of winning bids even in a competitive environment. However, greater entry would decrease the winning bids of both entrants and long-run players. Under this alternative model, unlike ours, minimum prices should have a similar qualitative impact on long-run players and entrant winners.

the introduction of a binding minimum price generates a mass point of bids at the minimum price and a gap in the support of the winning-bid distribution just above the minimum price. As a result, such a model cannot generate a first-order stochastic dominance drop in the right tail of the winning-bid distribution.

V. Empirical Analysis

Sections II, III, and IV lay out a theoretical mechanism through which minimum prices can affect the distribution of winning bids, and clarify its implications for data. This empirical section aims to assess the relevance of this mechanism in a real-life context and answer the following questions: are enforcement constraints binding? are they affected by minimum prices? what is the impact on cartel members? what is the impact on entrants?

We provide empirical answers to these questions using auction data from Japanese cities located in the Ibaraki prefecture.

A. Data and Empirical Strategy

Context.—Local procurement in Japan is an appropriate context for us to test the model developed in Sections II, III, and IV. McMillan's (1991) account of collusive practices in Japan's construction industry vindicates many of our assumptions. It confirms the role of transfers in sustaining collusion, as well as the importance of selective tendering and observed participation in limiting entry, especially at the local level.¹⁸ In addition, recently Ishii (2008) and Kawai and Nakabayashi (2014) provide evidence of widespread collusion in Japanese procurement auctions. This suggests that local procurement in Japan is an environment in which minimum-price constraints could plausibly have an effect.

Sample selection.—We collected our data as follows. In a study of paving auctions, Ishii (2008) notes the use of minimum prices in Japanese procurement auctions. The author was able to point us to one of our treatment cities. We then proceeded to search for all publicly available data for the 30 most populous cities in the prefecture. We kept all cities that had public data available covering the relevant period. This left us with the fourteen cities included in the study. We treat these fourteen cities as distinct markets.¹⁹ The data cover public work projects auctioned off between May 2007 and March 2016, corresponding to 10,533 auctions.

Throughout the period, all cities use first-price auctions. Six cities—Hitachiomiya, Inashiki, Toride, Tsuchiura, Tsukuba, and Tsukubamirai—experience at least one policy change going from a zero minimum price to a positive minimum price. Within this set, Tsuchiura provides us with the richest data, including bidder names, nonwinning bids, and minimum prices.²⁰ Cities other than the six mentioned above use first-price

¹⁸ We further refer to McMillan (1991) for details on real-world collusion, including organizational steps taken to ensure that high-level managers could deny any knowledge of collusion.

¹⁹ We discuss this assumption in Sec. VI.

²⁰ Notable trivia: Tsuchiura is a sister city of Palo Alto, California.

auctions with no minimum price throughout the period covered in our data.

Policy documents available from municipal websites clarify that minimum prices are chosen by a formal rule and contain no more information than reserve prices. Reserve prices are computed by adding up engineering estimates of material, labor, administrative, and financing costs. Minimum prices are obtained by multiplying each expense category by a predetermined coefficient.

Publicly available policy documents, as well as exchanges with city officials, confirm that minimum prices were introduced to avoid excessively low bids that could only be executed at the expense of quality.²¹ We found no evidence that policy changes were triggered by city-specific factors also affecting the distribution of bids.

Descriptive statistics.—Some facts about our sample of auctions are worth noting. The first is that although all auctions include a reserve price, these reserve prices are not set to extract greater surplus for the city along the lines of Myerson (1981) or Riley and Samuelson (1981). Rather, consistent with recorded practice, reserve prices are engineering estimates (Ohashi 2009; Tanno and Hirai 2012; Kawai and Nakabayashi 2014) that provide an upper bound to the range of possible costs for the project. This is largely verified in our data, since 99.02 percent of auctions have a winner. This lets us treat reserve prices as an exogenous scaling parameter and use it to normalize the distribution of bids to $[0, 1]$. Normalized winning bids are defined as follows:

$$\text{norm_winning_bid} = \frac{\text{winning_bid}}{\text{reserve_price}}.$$

This normalization lets us take the comparative statics of propositions 3, 4, and 5 to the data, even though there is heterogeneity in minimum prices. As a robustness test, we also study the distribution of log-winning-bids using reserve prices as a control variable (see table OA.7). Our findings are unchanged.

The distribution of winning bids is closely concentrated near reserve prices. Throughout all our data, the aggregate cost savings from running an auction rather than using reserve prices as a take-it-or-leave-it offer are equal to 7.5 percent. This could be because reserve prices are obtained through very precise engineering estimates, but this provides justifiable concern that collusion may be present.

In the city of Tsuchiura, for which we observe minimum prices, the median minimum price is in the first decile of the distribution of normalized winning bids. This matches the theoretical requirement that mini-

²¹ Our model can capture nonperforming bidders by treating them as entrants with zero costs. We elaborate on this point in Sec. VI and app. OC.

minimum prices should be in the lower quantiles of the observed distribution of winning bids (propositions 3, 5, and 6). Minimum prices in Tsuchiura range from .75 to .85 of the reserve.

Empirical strategy.—The data let us evaluate the prediction of propositions 3 and 6 directly. If there is no collusion the introduction of a low minimum price should not change the right tail of winning bids. In fact, in a competitive environment, introducing such a low minimum price should have a very limited effect on bidding behavior. In contrast, if there is collusion, we anticipate a drop in the right tail of winning bids.

We measure the impact of a policy change on the distribution of winning bids at the city level by forming either change-in-changes (Athey and Imbens 2006) or difference-in-differences estimates for the sample of normalized winning bids above a threshold of .8.²² Given the heterogeneity in city characteristics reported in table 1, we match each treatment city to two cities that are most suitable as controls according to the following criteria:

- the control city has data before and after the treatment city's policy change; during that period the control city does not itself experience a policy change;
- the control city minimizes the distance between the treatment city t and potential control city c according to distance

$$d_{t,c} = \left| \frac{\text{population}_t - \text{population}_c}{\text{population}_t} \right| + \left| \frac{\text{density}_t - \text{density}_c}{\text{density}_t} \right|.$$

When two minimum-price increases occur in the treatment city we keep only data corresponding to the first policy change. This corresponds to going from no minimum price to a positive minimum price and matches the premise of our theoretical results. We let cities experiencing a policy change serve as a control city when they do not experience a policy change.²³ Table 2 shows how treatment and control cities are matched.²⁴

We report our findings in three steps. We first provide a detailed description of our approach using Tsuchiura as a treatment city. We use Tsuchiura as a benchmark because it is the city for which we have the richest and most abundant data. We observe minimum prices, nonwinning bids, and the identity of all bidders. We then present aggregated results cluster-

²² This is the median minimum-to-reserve price ratio in Tsuchiura. The results are unchanged if we consider the distribution of normalized winning bids conditional on prices being above .78 or .82 of the reserve price. See app. OA for details.

²³ For instance, the city of Tsuchiura changed its policy in October 2009, and serves as a control for the city of Toride, which experienced a policy change in March 2012.

²⁴ We note that both this match and the resulting findings are largely unchanged if cities are matched according to an estimate of their likelihood of introducing minimum prices.

TABLE 1
CITY CHARACTERISTICS

City	Population	Density	Minimum Price Increases	Number of Auctions	Data Time Range
Hitachiomiya	42,438	122	2011-04, 2012-12	1,371	2007-07–2016-03
Hokota	51,519	253		597	2011-04–2016-04
Inashiki	43,178	210	2010-07, 2012-04	1,003	2007-04–2016-03
Kamisu	94,551	642		671	2012-05–2016-02
Kasumigaura	45,373	382		487	2013-07–2014-10
Moriya	64,644	1,810		312	2012-04–2016-03
Sakuragawa	49,387	275		472	2010-05–2016-01
Shimotsuma	45,289	560		335	2007-07–2016-02
Toride	109,926	1,571	2012-03, 2013-04	423	2010-05–2016-09
Tsuchiura	142,931	1,162	2009-10, 2014-04	1,748	2007-05–2016-03
Tsukuba	215,306	758	2014-04	1,855	2009-04–2016-03
Tsukubamirai	48,807	617	2014-04	290	2008-05–2015-12
Ushiku	81,532	1,385		672	2008-04–2016-03
Yuki	51,429	782		242	2014-04–2016-01

ing standard errors at the (city, year) level, performing wild bootstrap to obtain p -values (Cameron, Gelbach, and Miller 2008). For completeness we report individual treatment-city regressions in appendix OA.

B. Findings for Tsuchiura

1. The Impact of Minimum Prices on the Distribution of Winning Bids

Propositions 3 and 5 suggests a test of the mechanism we model. Under collusion, the introduction of a small minimum price should lead to a first-order stochastic dominance drop in the right-tail distribution of winning bids. Under competition, we should not expect to see such a change. We begin our analysis applying this test to Tsuchiura.

Figure 1 plots distributions of normalized winning bids above .8 for Tsuchiura and the corresponding control cities Tsukuba and Ushiku, using data 2 years before and 2 years after the policy change (October 28,

TABLE 2
TREATMENT AND CONTROL CITIES, MATCHED
ACCORDING TO POPULATION AND DENSITY

Treatment	Control 1	Control 2
Hitachiomiya	Inashiki	Sakuragawa
Inashiki	Hitachiomiya	Sakuragawa
Toride	Ushiku	Tsuchiura
Tsuchiura	Ushiku	Tsukuba
Tsukuba	Kamisu	Shimotsuma
Tsukubamirai	Shimotsuma	Kasumigaura

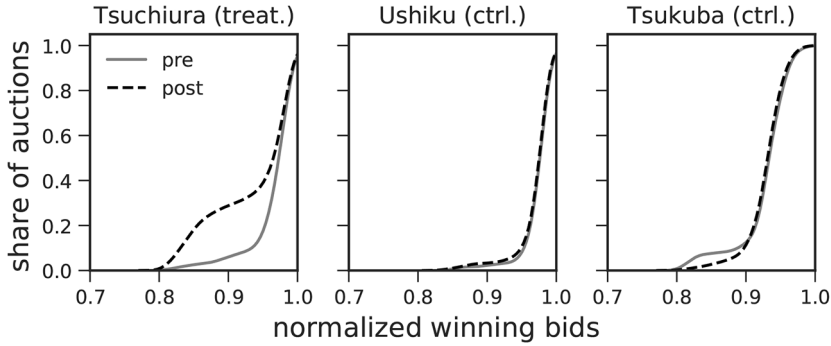


FIG. 1.—Distribution of winning bids above .8, before and after treatment: 2007–9, 2009–11.

2009). The three cities are broadly comparable: their populations range from 82,000 to 215,000, with Tsuchiura at 143,000. They are located within 15 km of one another and within 75 km of Tokyo. The data appear well suited to a difference-in-differences approach. Remarkably, the distribution of normalized winning bids in the control cities seems essentially unchanged. Figure 1 also suggests a first-order stochastic dominance drop in normalized winning bids of the treatment city above .8.

Change-in-changes.—The framework of Athey and Imbens (2006) allows us to formalize this observation by estimating the counterfactual distribution of normalized winning bids in our treatment city, absent minimum prices. Following Athey and Imbens (2006), we assume that the normalized winning bid $\text{norm_winning_bid}_a$ in auction a in period $\in \{\text{pre}, \text{post}\}$ under minimum-price status $\text{min_price} \in \{0, 1\}$ satisfies the relationship

$$\text{norm_winning_bid}_a = h_{\text{min_price}}(U_a, \text{period}),$$

where $U_a \in [0, 1]$ summarizes unobservable auction-level characteristics and where $h_{\text{min_price}}(u, \text{period})$ is increasing in u .

By recovering the respective distributions of U_a in the treatment and control cities, the method of Athey and Imbens (2006) allows us to estimate the counterfactual distribution of winning bids in the post period of the treatment city, if no minimum price had been introduced.

The actual and counterfactual quantiles of normalized winning bids, conditional on prices being above 80 percent of the reserve price, are given in table 3. We use both Tsukuba and Ushiku as a controls.²⁵

²⁵ We do not merge the control data. This would bias results since the relative sample size of the pre and post period is different across control cities. Instead, we separately run the algorithm of Athey and Imbens (2006) for each control city, and then average the corresponding counterfactual estimates. We report bootstrapped standard errors for our aggregated estimates.

TABLE 3
CHANGE-IN-CHANGES ESTIMATES: QUANTILES OF THE ACTUAL AND COUNTERFACTUAL
CONDITIONAL DISTRIBUTIONS OF NORMALIZED WINNING BIDS (> .8)

	QUANTILE OF CONDITIONAL DISTRIBUTION				
	.1	.25	.5	.75	.9
Actual – counterfactual	-.051***	-.054***	.003	.004	.004*
Standard error	(.018)	(.02)	(.005)	(.003)	(.003)
Actual	.833	.881	.959	.977	.984
Counterfactual	.884	.935	.956	.973	.98

* Significant at the .1 level.

** Significant at the .05 level.

*** Significant at the .01 level.

Table 3 shows that the right-tail distribution of winning bids fell in terms of first-order stochastic dominance after the policy change, consistent with our model's predictions under collusion.

Difference-in-differences.—The findings displayed in table 3 are confirmed by a difference-in-differences approach including additional controls.²⁶ We define variable

$$\text{policy_change} = \mathbf{1}_{\text{date} \geq \text{October 28, 2009} \ \& \ \text{city} = \text{Tsuchiura}}$$

and perform both ordinary least squares (OLS) and quantile regressions of the linear model with quarterly national GDP controls, city fixed effects, month fixed effects, year fixed effects, and city-specific trends:

$$\begin{aligned} \text{norm_winning_bid}_a = & \beta_0 + \beta_1 \text{policy_change} + \beta_2 \log \text{GDP} \\ & + \text{city_fe} + \text{month_fe} + \text{year_fe} \quad (2) \\ & + \text{city_trends} + \varepsilon_a. \end{aligned}$$

We continue to use the cities of Ushiku and Tsukuba as controls. To match the theoretical predictions of proposition 3, we perform regressions on the subsample of auctions whose normalized winning bid is above .8, corresponding to the sample of auctions whose winning bids are (or would have been) above the minimum price. For completeness, we also report mean effects for the unconditional sample of auctions. Throughout this section we refer to the sample as *conditional* when normalized winning bids are constrained to be above .8, and as *unconditional* when normalized winning bids are unconstrained.

For now, we present standard errors for our estimates assuming that shocks are independent at the auction level. In Section V.C we deal with

²⁶ Throughout the empirical analysis, we winsorize the normalized winning bids at 1 percent and 99 percent.

possible city-level shocks by aggregating the data from all cities in our sample and clustering errors at the (city, year) level.

The outcome of regression (2), summarized in table 4, vindicates the mechanism we explore in Sections II, III, and IV. The introduction of a minimum price leads to lower average winning bids in the conditional sample. Consistent with propositions 3 and 5, the estimates of the quantile regressions show that the policy change is associated with a first-order stochastic dominance drop in the right tail of winning bids. The implication is not only that there is collusion, but that cartel enforcement constraints are binding and that the sustainability of collusion is limited by price constraints.

2. Whom Does the Policy Change Affect?

Proposition 6 offers another test of the mechanism analyzed in Sections II, III, and IV. Under collusion, our theory predicts that the price paid by winning long-run bidders should go down, but not the price paid by winning entrants. Under competition, winning long-run bidders and winning entrants should be similarly affected. Importantly, these predictions hold even when we do not control for the set of participating bidders.

Defining long-run players.—The key step consists of deciding which firms are long-run players and which firms are likely entrants. As we show in appendix OE, proposition 6 remains true when using a proxy for long-run firms that is a superset of the actual group of long-run firms. For this reason we err on the side of inclusiveness when proxying for long-run players. If a proxy misclassifies entrants as long-run firms, the measured effect of the policy change on entrants remains equal to zero. If long-run

TABLE 4
DIFFERENCE-IN-DIFFERENCES ANALYSIS OF THE EFFECT OF MINIMUM
PRICES ON NORMALIZED WINNING BIDS

	UNCONDITIONAL SAMPLE MEAN EFFECT	SAMPLE SUCH THAT $\text{norm_winning_bid} > .8$				
		Mean Effect	$q = .2$	$q = .4$	$q = .6$	$q = .8$
policy_change	-.008 (.007)	-.026*** (.005)	-.084*** (.010)	-.021*** (.005)	-.006* (.003)	.003 (.002)
log GDP	.065 (.103)	.027 (.070)	.010 (.148)	-.047 (.068)	-.022 (.049)	-.065* (.036)
R^2	.201	.248				
Observations	3,705	3,459	3,459	3,459	3,459	3,459

NOTE.—OLS estimates for unconditional sample and quantile regression estimates for conditional sample; regressions include city fixed effects, month fixed effects, year fixed effects, and city-specific time trends.

* Effects significant at the .1 level.

** Effects significant at the .05 level.

*** Effects significant at the .01 level.

firms are misclassified as entrants, the measured effect of the policy on entrants may become nonzero.

For Tsuchiura, participation data allow us to form a measure directly consistent with the theory. We can classify firms as long-run bidders and entrants according to the frequency with which they participate in auctions. Tsuchiura exhibits considerable heterogeneity in the degree of bidder activity. The median number of auctions a bidder participates in is 4, whereas the mean is at 20. The 25 percent most active bidders make up 83 percent of the $\widehat{\text{auction}} \times \text{bidder}$ data. Accordingly, we define long-run bidder measure $\widehat{\text{long_run}}_i$, which takes a value of 1 if bidder i belongs to the 25 percent most active bidders (82 out of 330 bidders), and is equal to zero otherwise.

This measure cannot be computed for cities other than Tsuchiura: we observe winners but not participants. For this reason, we use winning an auction as a proxy for participation. Accordingly, we define long-run bidder measure $\widetilde{\text{long_run}}_i$, which takes a value of 1 whenever bidder i belongs in the set of bidders who belong to the 35 percent of bidders who *win* auctions most often out of those who win at least once (71 firms; 77 percent of the $\text{auction} \times \text{bidder}$ data in Tsuchiura), and is equal to zero otherwise.

The threshold 35 percent is the round-number threshold that generates the best overlap between $\widehat{\text{long_run}}$ and $\widetilde{\text{long_run}}$. All but five of the firms in our data that belong to long-run measure $\widehat{\text{long_run}}_i$ also belong to long-run measure $\widetilde{\text{long_run}}_i$. It is plausible that $\widetilde{\text{long_run}}$ may be less precise than $\widehat{\text{long_run}}$ because winning events are approximately 4 times rarer than participation events.²⁷

Findings.—We estimate the differential impact of minimum prices on long-run bidders and entrants using both $\widehat{\text{long_run}}$ and $\widetilde{\text{long_run}}$ measures. Since $\widetilde{\text{long_run}}$ is available for all cities, it can be used in a difference-in-differences approach estimating the linear model

$$\begin{aligned} \text{norm_winning bid}_a = & \beta_0 + \beta_1 \text{policy_change} + \beta_2 \widetilde{\text{long_run}} \\ & + \beta_3 \widetilde{\text{long_run}} \times \text{policy_change} \\ & + \beta_2 \log\text{GDP} + \text{city_fe} + \left(\text{city, } \widetilde{\text{long_run}} \right)_fe \\ & + \text{month_fe} + \text{year_fe} + \text{city_trends} + \varepsilon_a \end{aligned} \quad (3)$$

on the sample of auctions with normalized winning bids above 80 percent. Fixed effects include city-specific time trends, city fixed effects, (city, long-run status) fixed effects, as well as month and year fixed effects.²⁸

²⁷ On average, 3.8 bidders participate in each auction in the city of Tsuchiura.

²⁸ Our specification, with (city, long-run status) fixed effects, allows long-run bidders to behave differently in the treatment and control cities.

To confirm the findings obtained for long-run measure $\widetilde{\text{long_run}}$, we replicate them with measure $\widehat{\text{long_run}}$, which is potentially more accurate but is only available for Tsuchiura. We take a before-after approach and estimate linear model

$$\begin{aligned} \text{norm_winning bid}_a &= \beta_0 + \beta_1 \text{policy_change} + \beta_2 \widehat{\text{long_run}} \\ &+ \beta_3 \widehat{\text{long_run}} \times \text{policy_change} \quad (4) \\ &+ \beta \text{controls} + \varepsilon_a \end{aligned}$$

on conditional and unconditional auction data. Table 5 reports estimates for (3) and (4).

The findings are consistent with the predictions of our model under collusion. Absent minimum prices, long-run firms obtain contracts at higher prices; the introduction of minimum prices has a disproportionately larger effect on long-run winners than entrant winners.

C. Findings for All Cities

To deal with city-level shocks, we extend the analysis to all treatment and control cities listed in table 2 and aggregate the results. Figure 2 contrasts the before and after distributions of normalized winning bids above .8 for treatment and control cities. It exhibits patterns consistent with those of Tsuchiura: control cities experience little change; treatment cities experience either little change or a first-order stochastic dominance drop.

TABLE 5
EFFECT OF MINIMUM PRICES ON LONG-RUN BIDDERS AND ENTRANTS (Tsuchiura)

	long run $\equiv \widetilde{\text{long_run}}$		long run $\equiv \widehat{\text{long_run}}$	
	Unconditional Sample	Conditional Sample	Unconditional Sample	Conditional Sample
long_run	.053*** (.008)	.026*** (.006)	.068*** (.009)	.032*** (.008)
long_run \times policy_change	-.036*** (.009)	-.021*** (.007)	-.048*** (.010)	-.029*** (.008)
policy_change	.024** (.010)	-.007 (.007)	.047*** (.014)	.009 (.011)
R ²	.224	.255	.164	.117
Observations	3,705	3,459	1,748	1,660

* Significant at the .1 level.
 ** Significant at the .05 level.
 *** Significant at the .01 level.

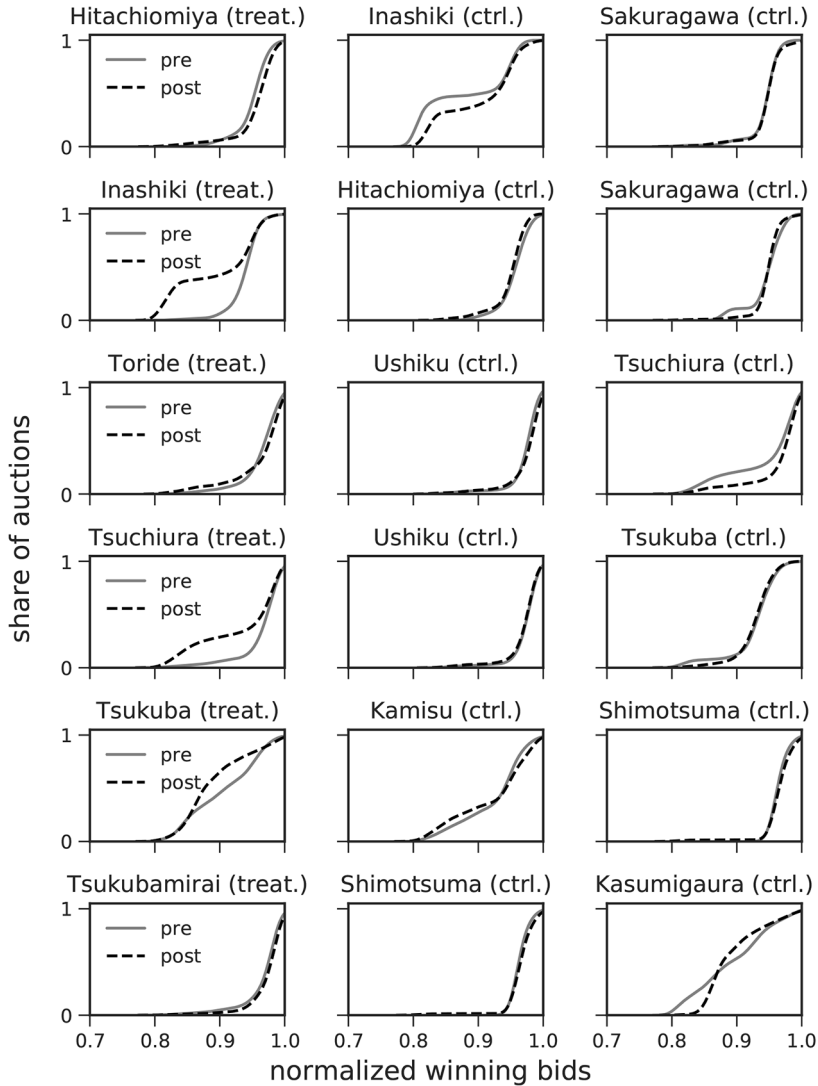


FIG. 2.—Distribution of winning bids above .8, before and after treatment

For each individual policy group g consisting of one treatment and two control cities, we assume that linear models (2) and (3) extend as follows:

$$\begin{aligned} \text{norm_winning bid}_a = & \beta_0 + \beta_1 \text{ policy_change} + \beta_g \text{ controls} \\ & + \text{fixed_effects}_g + \varepsilon_a, \end{aligned} \tag{5}$$

$$\begin{aligned} \text{norm_winning bid}_a &= \beta_0 + \beta_1 \text{policy_change} + \beta_2 \widetilde{\text{long_run}} \\ &+ \beta_3 \widetilde{\text{long_run}} \times \text{policy_change} \\ &+ \beta_g \text{controls} + \text{fixed_effects}_g + \varepsilon_a, \end{aligned} \tag{6}$$

where the g subscript in fixed_effect_g indicates that fixed-effect coefficients can vary with the treatment group.

Models (5) and (6) are naturally aggregated assuming that the impact of the policy change is the same across cities.²⁹ For an auction a and a policy group g , we denote by $a \in g$ the event that auction a is included in the relevant data for policy group g . We define $N_a \equiv \text{card}\{g, \text{ such that } a \in g\}$ as the number of policy groups in which auction a is included. For an auction a , averaging over the treatment groups g in which auction a appears yields

$$\begin{aligned} \text{norm_winning_bid}_a &= \beta_0 + \beta_1 \text{policy_change} \\ &+ \frac{1}{N_a} \sum_{g, \text{ such that } a \in g} (\beta_g \text{controls} + \text{fixed_effects}_g) \\ &+ \varepsilon_a, \end{aligned} \tag{7}$$

$$\begin{aligned} \text{norm_winning bid}_a &= \beta_0 + \beta_1 \text{policy_change} + \beta_2 \widetilde{\text{long_run}} \\ &+ \beta_3 \widetilde{\text{long_run}} \times \text{policy_change} \\ &+ \frac{1}{N_a} \sum_{g, \text{ such that } a \in g} (\beta_g \text{controls} + \text{fixed_effects}_g) \\ &+ \varepsilon_a. \end{aligned} \tag{8}$$

Estimates from (7) and (8), reported in table 6, corroborate findings from Tsuchiura. The introduction of minimum prices lowers winning bids above the minimum price, and the effect is disproportionately borne by potential long-run bidders.

VI. Robustness

This section briefly discusses the robustness of our theoretical and empirical findings. A deeper treatment of these robustness checks is provided in the appendixes.

A. Theory

Collusion with many firms.—The cartel measures we propose, $\widetilde{\text{long_run}}$ and $\widehat{\text{long_run}}$, classify a sizeable proportion of firms as long-run bidders.

²⁹ This is true under the H_0 assumption that the policy change has no impact.

TABLE 6
DIFFERENCE-IN-DIFFERENCES ANALYSIS OF THE EFFECT OF MINIMUM
PRICES ON NORMALIZED WINNING BIDS

	UNCONDITIONAL MEAN EFFECT		CONDITIONAL MEAN EFFECT	
	(1)	(2)	(3)	(4)
policy_change	-.015	.006	-.021**	-.008
<i>p</i> -value	.378	.713	.016	.238
policy_change × long_run		-.023***		-.016***
<i>p</i> -value		.006		.006
<i>R</i> ²	.284	.307	.308	.326
Observations	8,958	8,958	8,236	8,236

NOTE.—OLS estimates from eqq. (7) and (8), for unconditional sample (cols. 1 and 2) and conditional sample (cols. 3 and 4); regressions include city fixed effects, year fixed effects, month fixed effects, and city-specific time trends. Standard errors are clustered at the (city, year) level, and *p*-values are calculated by wild bootstrap.

* Effects significant at the .1 level.

** Effects significant at the .05 level.

*** Effects significant at the .01 level.

In Tsuchiura, the quartile of the most active firms (which represents 80 percent of auction × bidder data) includes 82 firms. This is partly because we want our proxy to be a superset of cartel members. In addition, although this is a large number, it is consistent with known cases of collusion among construction firms. In 2008 the United Kingdom's Office of Fair Trading filed a case against 112 firms in the construction sector. At least 80 of these firms have admitted to bid rigging and reported the use of monetary transfers. Another example is the Dutch construction cartel, which included on the order of 650 firms (Eftychidou and Maiorano 2015).

It is not implausible that comparable levels of collusion could exist in Japan's construction industry. However, rationalizing collusion within such a large group could be a challenge for the model of Section II. Pledgeable surplus is bounded, and many bidders must be compensated for their deviation temptation. We show in the appendixes that provided the cartel can endogenously control participation in auctions, then, a cartel can continue to collude even as the number of cartel members grows large.

The intuition for this result is the following. When participation is endogenous, the cartel faces two enforcement constraints: (i) bidders participating in an auction must accept bidding according to plan; (ii) bidders instructed not to participate in an auction must comply. While enforcement constraint i fails when the number of participants in an auction is large, enforcement constraint ii can be satisfied for arbitrary cartel sizes.

Indeed, imagine a cartel member that participates in an auction in which it was not supposed to bid. This unauthorized bidder is no different from an entrant. By bidding sufficiently low, other bidders can ensure that the unauthorized bidder makes zero flow profits. This implies

that the cartel can control participation very effectively and therefore keeps incentive constraint enforceable. An additional prediction is that introducing minimum prices will make it more difficult to keep cartel members from participating in auctions.

The data are consistent with this theory. In Tsuchiura (where we observe all bidders), participation is limited: the mean and median numbers of bidders per auction are both between 3 and 4. Table OA.2 confirms that introducing a minimum price leads to greater participation by both entrants and likely cartel members.

Nonperforming bidders.—The official rationale for introducing minimum prices is that it reduces the incidence of nonperforming bidders, that is, bidders unable to execute the tasks described in the procurement contract. Introducing such bidders explicitly does not change the findings from our analysis.

Nonperforming bidders can be modeled within the framework of Section IV as short-term entrants whose cost of production is set to 0. As Calveras et al. (2004) suggest, this may be because a bidder near bankruptcy is protected by limited liability. Propositions 5 and 6 continue to hold in the presence of such nonperforming bidders, since they rely only on the bidder side of the market. See the appendixes for details.

B. Empirics

Our model and our interpretation of the data rely on several assumptions that can be motivated from the data. We provide a summary below and present details in appendix OA.

Smooth equilibrium adjustments.—Propositions 3 and 6 provide a test of collusion by contrasting the comparative statics of the distribution of winning bids following the introduction of minimum prices, depending on whether we are in a collusive or competitive environment. These comparative statics presume that bidders are in equilibrium given the existing policy, which is necessarily an approximation. Indeed, although communication with city officials suggests that the move to a minimum price format was unexpected, it is still possible that the anticipation of the change may have affected behavior before the change, or that behavior after the change did not immediately move to the equilibrium corresponding to the new policy.

A priori, smooth equilibrium adjustment would bias estimates against our findings. Replicating our analysis excluding auctions occurring in the 6 month period before and after the policy change does not affect our results (see table OA.6).

Separate markets.—Our difference-in-differences analysis presumes that control cities are not affected by the policy change. One potential concern is that some of the long-run bidders active in a treatment city

may also be active in control cities. If that is the case, the introduction of minimum bids in a treatment city may also cause a shift in the distribution of bids in control cities.

This possible effect does not change the interpretation of our findings. Indeed, it should lead to an attenuation bias: part of the treatment effect would be interpreted as a common shock. We also argue in appendix OA that the assumption of separate markets is plausible: in Tsuchiura the bulk of active long-run bidders are geographically much closer to Tsuchiura than its control cities.

Observable participation.—Our model assumes that bidders observe participation at the bidding stage. This assumption can be motivated from data. We estimate the effect of entrant participation and cartel participation on realized bids (winning or not). Table OA.9 shows that even controlling for auction size through reserve prices, both entrant and cartel participations have a significant effect on nonwinning bids. This suggests that participants do have information about the set of bidders.

Are the theory and empirics consistent?—In appendix OF, we gauge the potential effect that minimum prices can have on bidding behavior in our model by conducting a back-of-the-envelope calibration exercise. We calibrate the model's parameters to match key statistics of bidding data from the city of Tsuchiura.

Our calibration exercise produces three main results. First, the introduction of minimum prices at the levels implemented by Tsuchiura has a negative effect on conditional winning bids (i.e., winning bids above the minimum price), ranging from -28 percent to -0.03 percent. Second, the effect of such minimum prices on average winning bids may be negative or positive, ranging from -11 percent to $+11$ percent. Third, consistent with corollary 3, a key factor explaining whether the unconditional treatment effect is negative or positive is the level at which the minimum price is introduced: average winning bids fall when the minimum price is relatively low, while they increase when the minimum price is high.

References

- Abdulkadiroglu, A., and K.-S. Chung. 2003. "Auction Design with Tacit Collusion." Working paper, Columbia Univ. and Northwestern Univ.
- Aoyagi, M. 2003. "Bid Rotation and Collusion in Repeated Auctions." *J. Econ. Theory* 112:79–105.
- Asker, J. 2010. "A Study of the Internal Organization of a Bidding Cartel." *A.E.R.* 100:724–62.
- Athey, S., and K. Bagwell. 2001. "Optimal Collusion with Private Information." *RAND J. Econ.* 32:428–65.
- Athey, S., K. Bagwell, and C. Sanchirico. 2004. "Collusion and Price Rigidity." *Rev. Econ. Studies* 71:317–49.
- Athey, S., and G. W. Imbens. 2006. "Identification and Inference in Nonlinear Difference-in-Differences Models." *Econometrica* 74:431–97.

- Bajari, P., and L. Ye. 2003. "Deciding between Competition and Collusion." *Rev. Econ. and Statis.* 85:971–89.
- Baker, G., R. Gibbons, and K. J. Murphy. 1994. "Subjective Performance Measures in Optimal Incentive Contracts." *Q.J.E.* 109:1125–56.
- . 2002. "Relational Contracts and the Theory of the Firm." *Q.J.E.* 117:39–84.
- Bernheim, B. D., and M. D. Whinston. 1998. "Incomplete Contracts and Strategic Ambiguity." *A.E.R.* 88:902–32.
- Blume, A., and P. Heidhues. 2008. "Modeling Tacit Collusion in Auctions." *J. Institutional and Theoretical Econ.* 164:163–84.
- Bull, C. 1987. "The Existence of Self-Enforcing Implicit Contracts." *Q.J.E.* 102:147–59.
- Calveras, A., J.-J. Ganuza, and E. Hauk. 2004. "Wild Bids: Gambling for Resurrection in Procurement Contracts." *J. Regulatory Econ.* 26:41–68.
- Cameron, A. C., J. B. Gelbach, and D. L. Miller. 2008. "Bootstrap-Based Improvements for Inference with Clustered Errors." *Rev. Econ. and Statis.* 90:414–27.
- Che, Y.-K., and J. Kim. 2009. "Optimal Collusion-Proof Auctions." *J. Econ. Theory* 144:565–603.
- Conley, T. G., and F. Decarolis. 2016. "Detecting Bidders Groups in Collusive Auctions." *American Econ. J.: Microeconomics* 8:1–38.
- Decarolis, F. 2018. "Comparing Public Procurement Auctions." *Internat. Econ. Rev.* 59:391–419.
- Decarolis, F., M. Goldmanis, and A. Penta. 2016. "Marketing Agencies and Collusive Bidding in Sponsored Search Auctions." Manuscript, Dept. Econ., Univ. Bocconi.
- Eftychidou, T., and F. Maiorano. 2015. "Serial Offenders: Why Some Industries Seem Prone to Endemic Collusion." [http://www.oecd.org/officialdocuments/publicdisplaydocumentpdf/?cote=DAF/COMP/GF\(2015\)4&docLanguage=En](http://www.oecd.org/officialdocuments/publicdisplaydocumentpdf/?cote=DAF/COMP/GF(2015)4&docLanguage=En).
- Ellison, G. 1994. "Theories of Cartel Stability and the Joint Executive Committee." *RAND J. Econ.* 25:37–57.
- Fabra, N. 2003. "Tacit Collusion in Repeated Auctions: Uniform versus Discriminatory." *J. Industrial Econ.* 51:271–93.
- Graham, D. A., and R. C. Marshall. 1987. "Collusive Bidder Behavior at Single-Object Second-Price and English Auctions." *J.P.E.* 95:1217–39.
- Green, E. J., and R. H. Porter. 1984. "Noncooperative Collusion under Imperfect Price Information." *Econometrica* 52:87–100.
- Harrington, J. E. 2008. "Detecting Cartels." In *Handbook of Antitrust Economics*, edited by P. Buccirossi, 213–58. Cambridge, MA: MIT Press.
- Harrington, J. E., and A. Skrzypacz. 2011. "Private Monitoring and Communication in Cartels: Explaining Recent Collusive Practices." *A.E.R.* 101:2425–49.
- Hendricks, K., R. Porter, and G. Tan. 2008. "Bidding Rings and the Winner's Curse." *RAND J. Econ.* 39:1018–41.
- Ishii, R. 2008. "Collusion in Repeated Procurement Auction: A Study of a Paving Market in Japan." ISER discussion paper, Inst. Soc. and Econ. Res., Osaka Univ.
- Kawai, K., and J. Nakabayashi. 2014. "Detecting Large-Scale Collusion in Procurement Auctions." Available at SSRN: <https://ssrn.com/abstract=2467175>.
- Lee, J., and H. Sabourian. 2011. "Efficient Repeated Implementation." *Econometrica* 79:1967–94.
- Levin, J. 2003. "Relational Incentive Contracts." *A.E.R.* 93:835–57.
- Mailath, G. J., and L. Samuelson. 2006. *Repeated Games and Reputations*, vol. 2. Oxford: Oxford Univ. Press.
- Marshall, R. C., and L. M. Marx. 2007. "Bidder Collusion." *J. Econ. Theory* 133:374–402.

- McAfee, R. P., and J. McMillan. 1992. "Bidding Rings." *A.E.R.* 82:579–99.
- McMillan, J. 1991. "Dango: Japan's Price-Fixing Conspiracies." *Econ. and Politics* 3:201–18.
- Mezzetti, C., and L. Renou. 2012. "Repeated Nash Implementation." Available at SSRN: <https://ssrn.com/abstract=2096184>.
- Myerson, R. B. 1981. "Optimal Auction Design." *Math. Operations Res.* 6:58–73.
- . 1986. "Multistage Games with Communication." *Econometrica* 54:323–58.
- Ohashi, H. 2009. "Effects of Transparency in Procurement Practices on Government Expenditure: A Case Study of Municipal Public Works." *Rev. Indus. Org.* 34:267–85.
- Pavlov, G. 2008. "Auction Design in the Presence of Collusion." *Theoretical Econ.* 3:383–429.
- Pesendorfer, M. 2000. "A Study of Collusion in First-Price Auctions." *Rev. Econ. Studies* 67:381–411.
- Porter, R. H. 1983. "A Study of Cartel Stability: The Joint Executive Committee, 1880–1886." *Bell J. Econ.* 14:301–14.
- Porter, R. H., and J. D. Zona. 1993. "Detection of Bid Rigging in Procurement Auctions." *J.P.E.* 101:518–38.
- . 1999. "Ohio School Milk Markets: An Analysis of Bidding." *RAND J. Econ.* 30:263–88.
- Riley, J. G., and W. F. Samuelson. 1981. "Optimal Auctions." *A.E.R.* 71:381–92.
- Rotemberg, J. J., and G. Saloner. 1986. "A Supergame-Theoretic Model of Price Wars during Booms." *A.E.R.* 76:390–407.
- Skrzypacz, A., and H. Hopenhayn. 2004. "Tacit Collusion in Repeated Auctions." *J. Econ. Theory* 114:153–69.
- Tanno, T., and T. Hirai. 2012. "Collusion, Breakdown, and Competition in Procurement Auctions." Manuscript, Dept. Econ., Atomi Univ.