

Screening Adaptive Cartels

Juan Ortner* Sylvain Chassang Jun Nakabayashi
Boston University New York University Kindai University

Kei Kawai

U.C. Berkeley

March 31, 2020

Abstract

We propose a theory of equilibrium antitrust oversight in which: *(i)* regulators launch investigations on the basis of suspicious bidding patterns; *(ii)* cartels can adapt to the statistical screens used by regulators, and may in fact use them to enforce cartel compliance. We emphasize the use of safe tests, i.e. tests that can be passed by competitive players under a broad class of environments. Such tests do not hurt competitive industries and do not help cartels support new collusive equilibria. We show that optimal collusive schemes in plausible environments fail natural safe tests, and that cartel responses to such tests explain unusual patterns in bidding data from procurement auctions held in Japan. This provides evidence that adaptive responses from cartels is a real concern that data-driven antitrust frameworks should take into account.

KEYWORDS: collusion, auctions, procurement, antitrust.

*Ortner: jortner@bu.edu, Chassang: chassang@nyu.edu, Nakabayashi: nakabayashi.1@eco.kindai.ac.jp, Kawai: kei@berkeley.edu .

1 Introduction

Competition authorities commonly rely on statistical screens to detect and investigate collusion between firms.¹ Even if formal prosecution cannot rely on data-driven evidence alone, detection tools can greatly facilitate the work of competition authorities. Such evidence can be used in court to obtain warrants or authorization for a more intrusive investigation, ultimately leading to actionable evidence and convictions (see Imhof et al., 2018, for a concrete example).² Furthermore, statistical evidence may be helpful in convincing cartel members to apply to leniency programs.³ However, this growth in the use of statistical screens raises numerous questions: Do firms adapt to the statistical tests implemented by competition authorities? If they do, what is the impact of such screens in equilibrium? Can the tests backfire and either strengthen cartels, or hurt competitive firms? Can we find tests that cannot cause harm, and yet strictly reduce the incentives to form cartels? We provide theory and evidence addressing these questions.

We model collusion in the shadow of antitrust authorities using a two-stage game. In the first stage a group of firms repeatedly participates in a first-price procurement auction. We allow firms to observe arbitrary signals about one another, and allow bidders' costs and types to be arbitrarily correlated within and across periods. In the second stage, an antitrust authority observes the entire history of firms' bidding behavior, and performs a screening test to determine whether or not they acted competitively. Firms that don't pass the test

¹Competition authorities that use statistical analysis or algorithms to screen for collusion include those in Brazil, South Korea, Switzerland and United Kingdom. A report by the OECD (2018) gives a brief description of the screening programs used in Brazil, Switzerland and the U.K. A document titled "Cartel Enforcement Regime of Korea and Its Recent Development" maintained by the Fair Trade Commission of Korea describes South Korea's bid screening program.

²Baker and Rubinfeld (1999) give an overview on the use of statistical evidence in court for antitrust litigation. In some jurisdictions, statistical evidence from screens have been used successfully to build a collusion case in court. See Mena-Labarthe (2015) for a case-study from Mexico.

³Screening of cartels can also be useful to those outside of antitrust authorities. For example, screening can help procurement offices counter suspected bidding rings by more aggressively soliciting new bidders or adopting auction mechanisms that are less susceptible to collusion. Screening may also be helpful for internal auditors and compliance officers of complicit firms to identify collusion and help contain potential legal risks arising from compliance failures.

are further investigated, and may incur penalties if found guilty of bid-rigging. Investigation may also be costly to non-cartel members. We say that the test used by the antitrust authority is a safe test if and only if it is passed with probability one by competitive bidders under arbitrary information environments. Our companion paper, Chassang et al. (2019), develops a near exhaustive class of safe tests and illustrates their relevance by applying them to Japanese procurement data.

Our main set of results suggest that antitrust oversight based on safe tests is a robust improvement over *laissez-faire*. First, we show that regulation based on firm-level safe tests does not expand the set of enforceable collusive schemes available to cartels. A similar result holds when testing is done at the industry level, provided penalties against colluding firms are large enough and we restrict attention to equilibria that are weakly renegotiation proof (Bernheim and Ray, 1989, Farrell and Maskin, 1989). This contrasts with work by Cyrenne (1999) and Harrington (2004) illustrating that arbitrary screens against collusion may backfire and enhance the cartel’s ability to collude.

Second, we show that optimal collusive schemes in plausible environments fail natural safe tests. Specifically, we show that in the maximally collusive scheme of a complete information collusion game, bidders either bid at the reserve price, or submit nearly tied bids. These configurations happen with zero-probability under competition, and naturally arise suspicion from regulators. It is plausible that cartels would adapt to the presence of such tests: in the face of a potential investigation a well working cartel would ensure that its members do not submit excessively close bids, and do not bid too close to the reserve price.

In theory, this adaptive response by a cartel should lead to noticeable patterns of “missing bids”: a missing mass of close bids, and a missing mass of bids close to the reserve price. Remarkably, we show that both patterns are present in a sample of procurement auctions from Japan: close bids, and bids close to the reserve price are unusually rare. Interestingly, as we show in Chassang et al. (2019), testing for missing bids turns out to yield safe tests. However, these are low-powered tests that only become safe asymptotically, as the sample

of bids grows large.

Our paper relates to a growing academic literature developing statistical methods to detect cartels.⁴ Collusive bidding patterns can be detected by measuring the level of correlation among bids (Bajari and Ye, 2003), by looking for price patterns predicted by the theory of repeated games (Porter, 1983, Ellison, 1994), or by exploiting changes in the auction format (Chassang and Ortner, 2019). Statistical tests of collusion have also been developed for average-price auctions (Conley and Decarolis, 2016) and multi-stage auctions with rebidding (Kawai and Nakabayashi, 2018). We complement this literature by showing that cartels do adapt to regulatory screens, and that using safe tests ensures that statistical screens do not create new collusive equilibria.

A smaller literature studies the equilibrium impact of antitrust oversight. Besanko and Spulber (1989) and LaCasse (1995) study static models of equilibrium regulation. Closer to our work, Cyrenne (1999) and Harrington (2004) study repeated oligopoly models in which colluding firms might get investigated and fined whenever prices exhibit large and rapid fluctuations. A common observation from these papers is that antitrust oversight may backfire, allowing cartels to sustain higher equilibrium profits. We provide evidence that concerns about adaptive cartels are valid, but that they can be addressed using safe tests.

Our work is also related to the literature on auction design in the presence of collusion. Abdulkadiroglu and Chung (2003), Che and Kim (2006, 2009) and Pavlov (2008) show that appropriate auction design can limit the cost of collusion when cartel members have deep pockets and can make payments upfront. Che et al. (2018) studies optimal auction design when collusive bidders are cash-constrained. Our paper complements this literature by showing how an antitrust agency can limit the effects of collusion by screening firms with safe tests.

The paper is structured as follows. Section 2 sets up our framework to study collusion in the shadow of investigation. Section 3 introduces safe tests. Section 4 establishes that

⁴See Porter (2005) and Harrington (2008) for recent surveys of this literature.

safe tests do not create new collusive equilibria and can strictly reduce the payoffs of cartels. Section 5 provides evidence that cartels do in fact adapt to high-powered safe tests. Section 6 presents concluding remarks, and discusses deviations from safe tests. Proofs are collected in the Appendix.

2 Colluding in the Shadow of Antitrust Authorities

We model the interaction between cartel members and antitrust authorities using a two-stage game. In the first stage, cartel members participate in repeated first-price procurement auctions. In the second stage, a regulator applies tests to the data generated by the players. If a test comes out against the null hypothesis of competition, one or more firms are investigated. We begin by describing the repeated procurement game used in the first stage.

2.1 Repeated procurement

In each period $t \in \mathbb{N}$, a buyer needs to procure a single project from a finite set N of potential suppliers. The auction format is a sealed-bid first-price auction with reserve price r , which we normalize to $r = 1$. Ties are broken randomly.

At each period t , a state $\theta_t \in \Theta$ captures all the relevant past information about the environment. State θ_t is revealed to bidders at the end of period t . We assume that θ_t evolves as a Markov chain, but do not assume that there are finitely many states, that θ_t is ergodic, or that θ_t is observable to the econometrician.

Firms' realized costs are denoted by $\mathbf{c}_t = (c_{i,t})_{i \in N}$. Each firm $i \in N$ submits a bid $b_{i,t} \in [0, 1] \cup \emptyset$, where \emptyset denotes not participating. We assume that bidders incur a cost $\kappa \geq 0$ from submitting a bid in $[0, 1]$. Let $\widehat{N}_t \subset N$ denote the set of participating bidders at period t . Profiles of bids are denoted by $\mathbf{b}_t = (b_{i,t})_{i \in \widehat{N}_t}$. We let $\mathbf{b}_{-i,t} \equiv (b_{j,t})_{j \neq i}$ denote bids from firms other than firm i , and define $\wedge \mathbf{b}_{-i,t} \equiv \min_{j \neq i} b_{j,t}$ to be the lowest bid among i 's competitors. We assume that bids are publicly revealed at the end of each period. This matches standard

practices in public procurement, where legislation typically requires governments to make bids public. Our main results can be adapted if only the winning bid was made public, or if bidders could only observe the identity of the winner.

Costs. We allow for costs that are serially correlated over time, and that may be correlated across firms within each auction. In particular, we assume that period t costs \mathbf{c}_t are jointly drawn from a distribution $F(\mathbf{c}_t|\theta_t)$, where $\theta_t \in \Theta$ is the state variable.

Information. In each period t , each bidder $i \in N$ privately observes a signal $z_{i,t}$ prior to bidding.⁵ The distribution of signals $(z_{i,t})_{i \in N}$ depends only on (θ_t, \mathbf{c}_t) . Signals $(z_{i,t})_{i \in N}$ are allowed to be arbitrary, and may include information about the state θ_t , and about costs \mathbf{c}_t . This allows our model to nest many informational environments, including private and common values, correlated values, asymmetric bidders, as well as complete information.

Transfers. We allow for the possibility of transfers across firms, with $T_{i,j} \geq 0$ being a voluntary transfer from firm i to firm j . Transfer $T_{i,j}$ generates a benefit $B(T_{i,j}) \leq T_{i,j}$ to firm j . Transfers take place at the end of each bidding round t , after the project is allocated and bidders observe the state θ_t .

We denote by $\Delta T_i \equiv \sum_{j \in N \setminus i} B(T_{j,i}) - T_{i,j}$ the net transfers received by firm i . For each $t \in \mathbb{N}$ and each $i \in N$, let $T_{i,t} = (T_{i,j,t})_{j \neq i}$ denote the transfers made by firm i at period t , and let $\mathbf{T}_t = (T_{i,t})_{i \in N}$.

Overall, firm i 's profits in period t are

$$\pi_{i,t} = x_{i,t} \times (b_{i,t} - c_{i,t}) - \kappa \mathbf{1}_{i \in \hat{N}_t} + \Delta T_{i,t},$$

where $x_{i,t} \in \{0, 1\}$ denotes whether or not firm i wins the auction at time t . Firms discount future payoffs using common discount factor $\delta < 1$.

⁵Note that firm i 's bid $b_{i,t} \in [0, 1] \cup \emptyset$ determines if i will be active in the auction. Hence, this specification allows firm to choose whether or not to participate after observing their signals.

Strategies and solution concept. The public history h_t^0 in period t takes the form

$$h_t^0 = (\theta_{s-1}, \mathbf{b}_{s-1}, \mathbf{T}_{s-1})_{s \leq t}.$$

Because state θ_t is publicly revealed at the end of each period, past play conveys no information about the private types of other players. As a result we do not need to specify out-of-equilibrium beliefs. A public strategy σ_i is a mapping

$$\sigma_i : h_t^0, z_{i,t} \mapsto b_{i,t}, T_{i,t}(\mathbf{b}_t, \theta_t).$$

Since transfers take place at the end of period t , after state θ_t is realized, $T_{i,t}(\mathbf{b}_t, \theta_t)$ is allowed to depend on both realized bids \mathbf{b}_t and realized state θ_t . Our solution concept is perfect public Bayesian equilibrium (Athey and Bagwell, 2008).

2.2 Antitrust oversight

The entirety of the collusive game described above, from period $t = 0$ to $t = \infty$, takes place within the first stage of our overall game. In the second stage the antitrust authority runs a vector of tests $(\tau_i)_{i \in N}$, with $\tau_i : h_\infty \mapsto \{0, 1\}$, where $h_\infty \in H_\infty$ is the realized history of bids $h_\infty = (b_{i,t})_{i \in N, t \in \mathbb{N}}$.⁶

If test τ_i takes value 1, firm i is investigated. This yields an expected penalty K . For simplicity, we consider fixed penalties. We discuss penalties that depend on the magnitude, and impact of collusion at the end of Section 4.1. We say that the regulator uses industry-level tests if

$$\forall h_\infty \in H_\infty, \quad \tau_1(h_\infty) = \tau_2(h_\infty) = \cdots = \tau_{|N|}(h_\infty) = \tau(h_\infty).$$

⁶More generally, history h_∞ may include any data observable to the antitrust authority.

Aggregate payoffs to firm i take the form

$$(1 - \delta) \sum_{t=0}^{\infty} \delta^t \pi_{i,t} - \tau_i K.$$

Note that because penalties $\tau_i K$ in the second stage are not discounted, this game is not continuous at infinity and the one-shot deviation principle does not necessarily apply (Fudenberg and Tirole, 1991). However, the one-shot deviation principle does apply if $K = 0$, or if players use strategies under which $\tau_i(h_\infty) = 0$ with probability 1 (see Lemma A.1 in the Appendix).

3 Safe Tests

Safe tests are tests that can be passed under a broad class of environments provided firms are competitive. Following Chassang et al. (2019), we say that a firm is competitive if and only if it plays a stage-game best-reply at every history on the equilibrium path.

Definition 1 (competitive histories). *Fix a common knowledge profile of play σ and a history $h_{i,t} = (h_t^0, z_{i,t})$ of player i . We say that player i is competitive at history $h_{i,t}$ if play at $h_{i,t}$ is stage-game optimal for firm i given the behavior of other firms σ_{-i} .*

We say that a firm is competitive if it plays competitively at all histories on the equilibrium path. We say that the industry is competitive if all firms play competitively at all histories on the equilibrium path.

We note that, if the industry is competitive under equilibrium σ , firms must be playing a stage-game Nash equilibrium in every period along the path of play.

Safe tests. We say that a test $\tau : H_\infty \rightarrow \{0, 1\}$ is safe if it is passed with probability one, either by competitive firms or competitive industries.⁷ Let us denote by $\lambda \equiv$

⁷As we discuss below, safe tests may have different data requirements. In practice many of the tests that can be implemented on real data will likely be slightly “unsafe”.

$\text{prob}((c_{i,t}, \theta_t, z_{i,t})_{i \in N, t \geq 0})$ the underlying economic environment, and by Λ the set of possible environments λ .

Definition 2 (safe tests). *We say that tests $(\tau_i)_{i \in N}$ are unilaterally safe if and only if for all $\lambda \in \Lambda$, and all profiles σ such that firm i is competitive, $\text{prob}_{\lambda, \sigma}(\tau_i(h) = 0) = 1$.*

We say that tests $(\tau_i)_{i \in N}$ are jointly safe if and only if for all $\lambda \in \Lambda$, and all profiles σ such that the industry is competitive, $\text{prob}_{\lambda, \sigma}(\tau_i(h) = 0) = 1$ for all $i \in N$.

In words, tests are jointly safe (resp. unilaterally safe) if they do not admit false positives. This concern over false positives coincides with concerns expressed by regulators (Imhof et al., 2016). In practice, investigation is a highly disruptive process that is only triggered if sufficiently many pieces of evidence are collected.⁸

Our main result shows that preventing harm against competitive firms serves an important strategic purpose: safe tests cannot unwittingly increase a cartel's enforcement capability. As Harrington (2004) highlights, this need not be true for tests that are not safe.

We consider two different settings:

- (i) The antitrust authority runs a unilaterally safe test τ_i on each firm $i \in N$.
Firm i incurs an undiscounted penalty of $K \geq 0$ if and only if $\tau_i(h_\infty) = 1$ (i.e., if and only if firm i fails the test).
- (ii) The antitrust authority runs a jointly safe test τ on all firms in N .
Firms in N incur a penalty of $K \geq 0$ if and only if $\tau(h_\infty) = 1$.

A notable aspect of safe tests is that they can be freely combined.

Remark 1. *If τ and τ' are safe tests, then both $\tau \wedge \tau'$ and $\tau \vee \tau'$ are safe tests.*

⁸This is not to say that a regulator would not launch investigation on the basis of somewhat imperfect evidence. Rather, that there is little cost in using safe tests.

4 Screening for Collusion in Equilibrium

4.1 Safe tests do not create new equilibria

This section provides normative foundations for safe tests. We show that safe tests can be used to place constraints on potential cartel members without creating new collusive equilibria.

For any $K \geq 0$, let $\Sigma(K)$ denote the set of perfect public Bayesian equilibria of the game with firm-specific testing and with penalty K . Similarly, let $\Sigma_{RP}(K)$ denote the set of Weakly Renegotiation Proof equilibria (Farrell and Maskin, 1989, Bernheim and Ray, 1989) of the game with penalty K , such that all players get weakly positive expected discounted payoffs in period 0. Let $\bar{K} \equiv \delta$. Note that \bar{K} serves as a rough upper bound on the difference in the continuation values a player obtains for different actions.⁹

The following result holds.

Proposition 1 (safe tests do not create new equilibria). *(i) If the regulator runs unilaterally safe tests, then $\Sigma(K) \subset \Sigma(0)$ for all $K > 0$.*

(ii) If the regulator runs jointly safe tests, then $\Sigma_{RP}(K) \subset \Sigma_{RP}(0)$ for all $K > \bar{K}$.

The proof of Proposition 1 is in the Appendix. Here, we provide an intuition for point (i), for the special case in which $K > \bar{K}$. Note first that when penalty K is large (i.e., $K > \bar{K}$), any equilibrium of the regulatory game has the property that, at all histories (both on and off path), all firms expect to pass the test with probability 1. Indeed, at every history, each firm can guarantee to pass the test by playing a stage-game best reply at all future periods.

Suppose $K > \bar{K}$ and fix $\sigma \in \Sigma(K)$. For simplicity, suppose that there are no transfers under σ . Consider a public history h_t^0 , and let $\beta = (\beta_i)_{i \in N}$ be the bidding profile that firms use at h_t^0 under σ : for all $i \in N$, $\beta_i : z_i \mapsto \mathbb{R}$ describes firm i 's bid as a function of her signal. Let $\mathbf{V} = (V_i)_{i \in N}$ be firms continuation payoffs excluding penalties after history h_t^0 under σ ,

⁹Recall that flow payoffs are normalized by $(1 - \delta)$, and that $r = 1$.

with $V_i : \mathbf{b} \mapsto \mathbb{R}$ mapping bids $\mathbf{b} = (b_j)_{j \in N}$ to a continuation value for firm i . Bidding profile β must be such that, for all $i \in N$ and all possible signal realizations z_i ,

$$\begin{aligned} \beta_i(z_i) &\in \arg \max_b \mathbb{E}_\beta[(1 - \delta)(b - c_i) \mathbf{1}_{b < \wedge \mathbf{b}_{-i}} + \delta V_i(b, \mathbf{b}_{-i}) | z_i] - \mathbb{E}_\sigma[\tau_i | h_t^0, b] K \\ \implies \beta_i(z_i) &\in \arg \max_b \mathbb{E}_\beta[(1 - \delta)(b - c_i) \mathbf{1}_{b < \wedge \mathbf{b}_{-i}} + \delta V_i(b, \mathbf{b}_{-i}) | z_i], \end{aligned}$$

where the second line follows since all firms pass the test with probability 1 after all histories under $\sigma \in \Sigma(K)$. In words, strategy profile σ is such that, at each history h_t^0 , no firm i has a profitable one shot deviation in a game without testing. The one-shot deviation principle then implies that $\sigma \in \Sigma(0)$.¹⁰

Let $\Sigma^P(0) \subset \Sigma(0)$ (resp., $\Sigma_{RP}^P(0) \subset \Sigma_{RP}(0)$) denote the set of equilibria (resp., the set of Weakly Renegotiation Proof equilibria) of the game without a regulator with the property that all firms expect to pass the test with probability 1 at every history. The following result holds:

Corollary 1. *Assume $K > \bar{K}$.*

- (i) *If the regulator runs unilaterally safe tests, then $\Sigma(K) = \Sigma^P(0)$.*
- (ii) *If the regulator runs jointly safe tests, then $\Sigma_{RP}^P(K) = \Sigma_{RP}^P(0)$.*

In words, when penalties are large, the set of equilibria of the regulatory game is equal to the set of equilibria of the game with no regulator under which all firms pass the test.

We highlight that testing at the individual firm level is crucial for Proposition 1(i). Indeed, as Cyrenne (1999) and Harrington (2004) show, regulation based on industry level tests may backfire, allowing cartels to achieve higher equilibrium payoffs. Intuitively, when testing is at the industry level, cartel members can punish deviators by playing a continuation strategy that fails the test. This relaxes incentive constraints along the equilibrium path,

¹⁰Note that the game with $K = 0$ is continuous at infinity, and so the one-shot deviation principle holds in such game.

and may lead to more collusive outcomes. Proposition 1(ii) shows, however, that jointly safe tests don't generate new collusive equilibria that are Weakly Renegotiation Proof.

Outcome contingent penalties. For simplicity, our model assumes a fixed penalty K from failing the test. We stress, however, that Proposition 1(i) continues to hold if penalty K is allowed to depend on the outcome h_∞ of the repeated game.¹¹ This allows, for instance, for penalties that depend on the extent, or impact, of collusion.

4.2 Optimal collusive schemes fail safe tests

Optimal collusion in an idealized setting. We consider a simple special case of the repeated procurement model of Section 2 following Chassang and Ortner (2019). Each period, the profile of costs $(c_{i,t})_{i \in N}$ is drawn i.i.d. from a joint distribution with c.d.f. F . Costs are complete information among firms, and firms can engage in efficient transfers: $B(T_{i,j}) = T_{i,j}$. Note that since we use perfect public Bayesian equilibrium as our solution concept, the fact that costs are complete information means that bids will also be common knowledge among bidders. The assumption that costs are common knowledge is not unreasonable under the assumption that firms are participating in a cartel.

For simplicity we assume that participation is constant; i.e., $\widehat{N}_t = N$ for all t . This may capture minimum participation requirements by the auctioneer: the auction is only run if a minimum number of participants join. Alternatively, if we think of the cost κ of participating in auctions as the cost of formulating a quality bid, or as the cost of tying up resources in case a project is won, then this cost may fall to 0 for bidders who don't intend to win in the first place.

In this setting, Chassang and Ortner (2019) establishes that all collusive schemes that are Pareto efficient among bidders are such that:

¹¹Proposition 1(ii) also continues to hold in this case, provided penalties are always larger than \overline{K} .

- (i) In each period the project is allocated to the lowest cost bidder, i.e. the bidder i s.t. $c_i < \wedge \mathbf{c}_{-i}$.
- (ii) The winning bid is the maximum value $b_{i,t} \leq r = 1$ such that

$$\sum_{j \neq i} (b_{i,t} - c_{j,t})^+ \leq \delta \bar{V}, \quad (1)$$

where \bar{V} is the discounted expected value of bidder profits aggregated at the cartel level.

- (iii) If $b_{i,t} < 1$, one player $j \neq i$ places bid $b_{j,t} = b_{i,t}^+$, adjacent to, but above the winning bid $b_{i,t}$.

Equation (1) states that a winning bid can be sustained if and only if the sum of deviation temptations (i.e. profits from bidders tempted to undercut the intended winner) is less than the pledgeable surplus from collusion.

The rationale for having a player bid immediately above winning bidder i is the following. If the highest sustainable bid is strictly below the reserve price $r = 1$, then (1) is binding. This means that pledgeable surplus has positive value for the cartel. In addition, since $b_{i,t} < 1$, the winning bidder is potentially able to increase her bid. If the lowest bid of other bidders were strictly above the winning bid, then the winning bidder would have an incentive to bid slightly higher. This means that in order to enforce equilibrium behavior, some of the scarce pledgeable surplus would have to be used to discipline the winning bidder. By having one bidder bid $b_{i,t}^+$, the deviation temptation of the winning bidder disappears, which liberates valuable pledgeable surplus.

Two small-sample safe tests. We now describe two safe tests that the collusive behavior described above fails. Importantly these tests are safe in small samples. Test τ^0 ensures that there is no mass of close winning and losing bids. Test τ^1 checks that there is no mass of winning bids at the reserve price.

Pick $\rho \in (0, 1)$ and a finite time T . Let us denote by $b_t^{(1)}$ and $b_t^{(2)}$ the lowest and second lowest bids of the auction occurring at time t . Tests τ^0 and τ^1 are formally defined as follows:

$$\tau^0 = \mathbf{1} \left(\limsup_{\epsilon \searrow 0} \frac{|\{t \leq T, \text{ s.t. } b_t^{(2)} - b_t^{(1)} < \epsilon\}|}{T} > \rho \right),$$

$$\tau^1 = \mathbf{1} \left(\limsup_{\epsilon \searrow 0} \frac{|\{t \leq T, \text{ s.t. } r - b_t^{(1)} < \epsilon\}|}{T} > \rho \right).$$

Because these tests identify events that have zero-probability under competition, they are safe tests even using a finite amount of data. Recall that $\kappa \geq 0$ is the cost that bidders incur for participating in an auction.

Proposition 2. *If $\kappa > 0$, then τ^0 and τ^1 are jointly safe tests.*¹²

It is immediate that for $\rho > 0$ small enough, the optimal collusive behavior described above will fail these tests. Frequent high bids, as captured by test τ^1 , would naturally attract the suspicion of regulators. Interestingly, regulators frequently use variance screens (Abrantes-Metz et al., 2005, Imhof et al., 2016) that flag auctions whose bids are unusually close together. This corresponds to the pattern of bids captured by test τ^0 , with the adjustment that it focuses on the distance between the lowest and second lowest bid, rather than the variance of the bids. The match is exact when there are only two bidders.

5 Evidence of Cartel Adaptation and Large Sample Safe Tests

Tests τ^0 and τ^1 are small sample safe tests that rule out plausible collusive equilibria, and are in fact applied by regulatory agencies. If cartels adapt to regulatory screens, then it is

¹²An individually safe analogue of τ^0 , τ_i^0 , can be obtained by looking at auctions such that bidder i is the second lowest bidder. An individually safe analogue of τ^1 , τ_i^1 , can be obtained by focusing on auctions in which player i is the highest bidder.

plausible that they may adapt to these particular tests. If cartel members communicate, they may coordinate to avoid triggering such a test. In fact, if a cartel is especially careful about not triggering this test – more so than a competitive industry with nothing to fear would be – they might take a margin of safety generating noticeable patterns in large samples.

Specifically, under test τ^0 , cartel bidders will avoid bids such that $b^{(2)} - b^{(1)} < \epsilon$ for ϵ small. This means that in the sample of auction bids, the mass of close winning bids will be low. This is the pattern studied in Chassang et al. (2019) by considering the sample distribution of bid-differences $\Delta_{i,t} \equiv b_{i,t} - \wedge b_{-i,t}$. In words, bid difference $\Delta_{i,t}$ measures the margin by which bidder i wins or loses the auction at time t . A cartel that is trying hard to avoid triggering test τ^0 will generate a sample of bids such that the density of $\Delta_{i,t}$ is close to 0 around $\Delta_{i,t} = 0$.

This prediction holds in the data. Figure 1 illustrates the distribution of bid differences $\Delta_{i,t}$ for a sample of procurement auctions taking place in the city of Tsuchiura, in Ibaraki prefecture, Japan. The data contains approximately 400 auctions taking place between May 2007 and October 2009. The auction format is first-price sealed-bid, with a public reserve price. The median number of bidders is 5, and the median winning bids is approximately USD 98,000. In previous work, Chassang and Ortner (2019) provide evidence of bidder collusion in these auctions.

As anticipated, firms seek to avoid suspiciously close bids. Interestingly, as we show in Chassang et al. (2019), this missing mass of bids around $\Delta = 0$ is itself a suspicious pattern that can be turned into a safe test. Specifically, testing whether the elasticity of sample demand is less than -1 is a safe test which is failed by the data in Figure 1. This latter test, however, is only safe for large samples.

Consider in turn test τ^1 . A cartel seeking to avoid detection will try to ensure that there are no auctions such that $r - b^{(1)} \leq \epsilon$ for ϵ small. This means that the density of winning bids around the reserve price should be zero. As Figure 2 shows this is also the case in




Figure 1: Distribution of bid differences Δ

the sample of city level procurement auctions mentioned above. As in the case of τ^0 , this adaptation to test τ^1 creates a bidding pattern that can be exploited to generate a new safe test. For any $\epsilon > 0$ define

$$\tau^2 = \limsup_{\rho \searrow 0} \limsup_{T \rightarrow \infty} \mathbf{1} \left(\frac{|\{t \leq T \text{ s.t. } r - b_t^{(1)} \leq \epsilon\}|}{T} \leq \rho \right).$$

In words, firms fail test τ^2 whenever the density of winning bids around the reserve price is zero. The following result holds.

Proposition 3. *If $\kappa > 0$, test τ^2 is a jointly safe test.*

6 Discussion

Summary. We propose an equilibrium model of data-driven screening for cartel behavior in which cartel members can adapt to undermine regulatory oversight. We emphasize the value




Figure 2: The cumulative distribution of winning bids does not reach the reserve price

of safe tests designed to fail firms whose behavior cannot be explained by any competitive model. We show that such tests cannot help cartels sustain new collusive equilibria, and that they can be freely combined to create new safe tests.

We identify two safe tests of interest, identifying whether a mass of winning bids are very close to the second highest bid, and whether a mass of winning bids is very close to the reserve price. Such tests are safe in small samples, and rule out optimal collusive behavior in plausible settings. These tests are related to screens used by regulators in practice, and data from procurement auctions in Japan provides suggestive evidence that cartel members do in fact adapt to such tests. The bidding patterns generated by the cartels' adaptive response can be exploited to generate large-sample safe tests.

Beyond safe tests. Because they have a limited downside while still providing some bite, we view safe tests as a natural starting point for anti-competitive screens. In practice however, it seems useful for regulators to go beyond safe tests. Indeed, many safe tests are only safe for asymptotically large sample sizes. The tests that can be implemented on real

data will necessarily fail competitive bidders with positive probability in finite samples. In addition, regulators may decide to respond to patterns in the data that are unlikely to occur in a competitive equilibrium given their prior over the underlying environment.

The fact that cartels adapt to the regulatory environment suggests that it is important to ask whether a candidate screen can potentially enhance the cartel's mechanics. Answering this question requires careful thinking on a case by case basis. In the case of tests that are asymptotically safe but are implemented on finite data, it is possible to show that they do not significantly enlarge the range of possible equilibrium outcomes.

Appendix

A Proofs

Because the regulatory game we study is not continuous at infinity, the one-shot deviation principle does not hold in general. However, as a preliminary to the proof of Proposition 1, we establish a version of the one-shot deviation principle for equilibria in which the bidders pass safe tests with probability one.

We use the following notation. For any strategy profile $\hat{\sigma}$ and any history $h_{i,t} = (h_t^0, z_{i,t})$, we use $V_i(\hat{\sigma}, h_{i,t}) = \mathbb{E}_{\hat{\sigma}}[\sum_{s \geq t} \delta^s \pi_{i,s} | h_{i,t}]$ to denote firm i 's continuation payoff excluding penalties under $\hat{\sigma}$ at history $h_{i,t}$ ($\pi_{i,s}$ includes firm i 's payoff from the auction at time s , plus payoffs from transfers). Firm i 's total payoff under strategy profile $\hat{\sigma}$ given history $h_{i,t}$ is $V_i(\hat{\sigma}, h_{i,t}) - \mathbb{E}_{\hat{\sigma}}[\tau_i | h_{i,t}]K$.

Lemma A.1. *Let σ be a strategy profile with the property that all firms pass the tests with probability 1 at every history. Then, $\sigma \in \Sigma(K)$ if and only if there are no profitable one-shot deviations.*

Proof. Let σ be a strategy profile with the property that all firms pass the test with probability 1 at every history. Clearly, if $\sigma \in \Sigma(K)$, there are no profitable one-shot deviations.

Suppose next that there are no profitable one-shot deviations, but $\sigma \notin \Sigma(K)$. Then, there exists a player $i \in N$ a history $h_{i,t}$ and a strategy $\tilde{\sigma}_i$ such that

$$\begin{aligned} V_i((\tilde{\sigma}_i, \sigma_{-i}), h_{i,t}) &\geq V_i((\tilde{\sigma}_i, \sigma_{-i}), h_{i,t}) - \mathbb{E}_{(\tilde{\sigma}_i, \sigma_{-i})}[\tau_i | h_{i,t}]K \\ &> V_i(\sigma, h_{i,t}) - \mathbb{E}_\sigma[\tau_i | h_{i,t}]K = V_i(\sigma, h_{i,t}), \end{aligned}$$

where the last equality follows since σ is such that all firms pass the test with probability 1 at every history.

The proof now proceeds as in the proof of the one-shot deviation principle in games that are continuous at infinity. Let $\epsilon \equiv V_i((\tilde{\sigma}_i, \sigma_{-i}), h_{i,t}) - V_i(\sigma, h_{i,t})$. Let $T > 0$ be such that $\delta^T \times r = \delta^T < \epsilon/2$. Let $\hat{\sigma}_i$ be a strategy for firm i that coincides with $\tilde{\sigma}_i$ for all histories of length $t + T$ or less, and coincides with σ_i for all histories of length strictly longer than $t + T$. Since firms' flow payoffs are bounded above by $r(1 - \delta) = 1 - \delta$, it must be that $V_i((\hat{\sigma}_i, \sigma_{-i}), h_{i,t}) - V_i(\sigma, h_{i,t}) \geq \epsilon/2$. Moreover, since σ is such that all firms pass the test with probability 1 at all histories, and since $\hat{\sigma}_i$ differs from σ_i only at finitely many periods, all firms pass the test under $(\hat{\sigma}_i, \sigma_{-i})$ with probability 1 at every history.

Next, look at histories of length $t + T$. If there exists a history $h_{i,t+T}$ of length $t + T$ that is consistent with $h_{i,t}$ and such that $V_i((\hat{\sigma}_i, \sigma_{-i}), h_{i,t+T}) > V_i(\sigma, h_{i,t+T})$, then there exists a profitable one shot deviation from σ (since $\hat{\sigma}_i$ and σ_i coincide for all histories of length $t + T + 1$ or longer).

Otherwise, let $\hat{\sigma}_i^1$ be a strategy that coincides with $\tilde{\sigma}_i$ at all histories of length $t + T - 1$ or less, and that coincides with σ_i at all histories of length strictly longer than $t + T - 1$. Note that it must be that $V_i((\hat{\sigma}_i^1, \sigma_{-i}), h_{i,t}) - V_i(\sigma, h_{i,t}) \geq \epsilon/2$. We can now look at histories of length $t + T - 1$ that are consistent with $h_{i,t}$. If there exists such a history $h_{i,t+T-1}$ such that $V_i((\hat{\sigma}_i^1, \sigma_{-i}), h_{i,t+T-1}) > V_i(\sigma, h_{i,t+T-1})$, then there exists a profitable one shot deviation from σ . Otherwise, we can continue in the same way. Since $V_i((\hat{\sigma}_i, \sigma_{-i}), h_{i,t}) - V_i(\sigma, h_{i,t}) \geq \epsilon/2$, eventually we will find a profitable one shot deviation by player i , a contradiction. ■

Proof of Proposition 1. We first establish point (i). Fix $\sigma = (\sigma_j)_{j \in N} \in \Sigma(K)$. Then, for all $i \in N$, for all histories $h_{i,t}$, and for all $\tilde{\sigma}_i \neq \sigma_i$, it must be that

$$\begin{aligned} V_i(\sigma, h_{i,t}) - \mathbb{E}_\sigma[\tau_i|h_{i,t}]K &\geq V_i((\tilde{\sigma}_i, \sigma_{-i}), h_{i,t}) - \mathbb{E}_{(\tilde{\sigma}_i, \sigma_{-i})}[\tau_i|h_{i,t}]K \\ \implies V_i(\sigma, h_{i,t}) &\geq V_i((\tilde{\sigma}_i, \sigma_{-i}), h_{i,t}) - \mathbb{E}_{(\tilde{\sigma}_i, \sigma_{-i})}[\tau_i|h_{i,t}]K, \end{aligned} \quad (2)$$

where the second inequality follows since $\mathbb{E}_\sigma[\tau_i|h_{i,t}]K \geq 0$.

Pick any $\epsilon > 0$, and let $T \in \mathbb{N}$ be such that $\delta^T < \epsilon$. Let $\hat{\sigma}_i$ be a strategy that coincides with $\tilde{\sigma}_i$ at all histories of length $s \leq t + T$, and such that under $\hat{\sigma}_i$, player i plays a static best response to σ_{-i} at all histories of length $s > t + T$. Note that firm i passes the test with probability 1 by playing $\hat{\sigma}_i$ against σ_{-i} .

Since $\hat{\sigma}_i$ and $\tilde{\sigma}_i$ only differ on histories of length $s > t + T$,

$$|V_i((\tilde{\sigma}_i, \sigma_{-i}), h_{i,t}) - V_i((\hat{\sigma}_i, \sigma_{-i}), h_{i,t})| < \epsilon. \quad (3)$$

At the same time, since $\sigma \in \Sigma(K)$, it must be that

$$\begin{aligned} V_i(\sigma, h_{i,t}) - \mathbb{E}_\sigma[\tau_i|h_{i,t}]K &\geq V_i((\hat{\sigma}_i, \sigma_{-i}), h_{i,t}) - \mathbb{E}_{(\hat{\sigma}_i, \sigma_{-i})}[\tau_i|h_{i,t}]K \\ \implies V_i(\sigma, h_{i,t}) &\geq V_i((\hat{\sigma}_i, \sigma_{-i}), h_{i,t}) \geq V_i((\tilde{\sigma}_i, \sigma_{-i}), h_{i,t}) - \epsilon, \end{aligned} \quad (4)$$

where the first inequality in (4) follows since $\mathbb{E}_\sigma[\tau_i|h_{i,t}]K \geq 0$ and since firm i passes the test with probability 1 by playing $\hat{\sigma}_i$ against σ_{-i} , and the second inequality in (4) uses (3). Since $\epsilon > 0$ is arbitrary, it must be that $V_i(\sigma, h_{i,t}) \geq V_i((\tilde{\sigma}_i, \sigma_{-i}), h_{i,t})$. Since this is true for every player $i \in N$, every history $h_{i,t}$, and any deviation $\tilde{\sigma}_i \neq \sigma_i$, we have that $\sigma \in \Sigma(0)$.

We now turn to the proof of point (ii). We first show that, when penalty K is sufficiently large, any $\sigma \in \Sigma_{RP}(K)$ has the property that all firms pass the test with probability 1, both on and off the path of play.

To show this, we first note that, at any equilibrium in $\Sigma_{RP}(K)$ and at every public history h_t^0 , at least one firm's continuation payoff is larger than 0. Indeed, by definition, all firms' payoffs at the start of the game are weakly larger than 0 at any equilibrium in $\Sigma_{RP}(K)$. By weak renegotiation proofness, it must be that at every history public h_t^0 at least one firm's payoff is larger than 0.

Recall that $\bar{K} \equiv \delta$. Suppose $K > \bar{K}$, and fix $\sigma \in \Sigma_{RP}(K)$. Towards a contradiction, suppose that there exists a public history h_t^0 (on or off path) such that, at this history, firms expect to fail the test with strictly positive probability under strategy profile σ . Then, for every $\epsilon > 0$ small, there must exist a history h_s^0 with $s \geq 0$ such that, at the concatenated public history $h_t^0 \sqcup h_s^0$, firms expect to fail the test with probability at least $\frac{\bar{K}}{K} + \epsilon < 1$ under σ . Hence, under σ , at public history $h_t^0 \sqcup h_s^0$ each firm's continuation payoff at the end of period $t + s$ (i.e., after the outcome of the auction is determined) is bounded above by $\delta - \left(\frac{\bar{K}}{K} + \epsilon\right)K = -\epsilon K < 0$, a contradiction.

Suppose $K > \bar{K}$ and fix $\sigma \in \Sigma_{RP}(K)$. Since σ is an equilibrium, there cannot be profitable deviations; in particular, there cannot be profitable one shot deviations: for every $i \in N$, every history $h_{i,s}$, and every one-shot deviation $\tilde{\sigma}_i \neq \sigma_i$ with $\sigma_i(h_{i,t}) = \tilde{\sigma}_i(h_{i,t})$ for all $h_{i,t} \neq h_{i,s}$,

$$\begin{aligned} V_i((\tilde{\sigma}_i, \sigma_{-i}), h_{i,s}) - \mathbb{E}_{(\tilde{\sigma}_i, \sigma_{-i})}[\tau_i | h_{i,s}]K &\leq V_i(\sigma, h_{i,s}) - \mathbb{E}_\sigma[\tau_i | h_{i,s}]K \\ \iff V_i((\tilde{\sigma}_i, \sigma_{-i}), h_{i,s}) &\leq V_i(\sigma, h_{i,s}), \end{aligned} \tag{5}$$

where the second line in (5) follows since, under equilibrium σ , firms pass the test with probability 1 at every history.¹³ By the second line in (5), in the game with $K = 0$ (i.e., no regulator) no firm has a profitable one shot deviation under strategy profile σ . By the one-shot deviation principle, $\sigma \in \Sigma(0)$. Finally, since σ is weakly renegotiation proof under penalty $K > \bar{K}$, σ must also be weakly renegotiation proof under penalty $K = 0$. Therefore,

¹³The deviation $\tilde{\sigma}_i$ from σ_i at history $h_{i,t}$ can be either at the bidding stage or at the transfer stage.

$\sigma \in \Sigma_{RP}(0)$. ■

Proof of Corollary 1. We first establish part (i). We start by showing that, when $K > \bar{K}$, any $\sigma \in \Sigma(K)$ has the property that all firms pass the test with probability 1 at every history, both on and off the path of play. To see why, note first that for every $i \in N$ and every strategy profile σ_{-i} of i 's opponents, firm i can guarantee to pass the test by playing a stage-game best reply to σ_{-i} at every history. This implies that each firm's continuation payoff cannot be lower than 0 at any history.

Towards a contradiction, suppose there exist $\sigma \in \Sigma(K)$ and a public history h_t^0 (on or off path) such that, at this history, firm i expects to fail the test with strictly positive probability under σ . Then, for every $\epsilon > 0$ small, there must exist a public history h_s^0 with $s \geq 0$ such that, at the concatenated history $h_t^0 \sqcup h_s^0$, firm i expects to fail the test with probability at least $\frac{\bar{K}}{K} + \epsilon < 1$. Hence, under history $h_t^0 \sqcup h_s^0$, firm i 's continuation payoff at the end of period $t + s$ (i.e., after the outcome of the auction is determined) is bounded above by $\delta - \left(\frac{\bar{K}}{K} + \epsilon\right) K = -\epsilon K < 0$, a contradiction.

The arguments in above imply that for $K > \bar{K}$, all equilibria in $\Sigma(K)$ have the property that all firms pass the test with probability 1 at every history. Since $\Sigma(K) \subset \Sigma(0)$ (by Proposition 1(i)), it follows that, for all $K > \bar{K}$, $\Sigma(K) \subset \Sigma^P(0)$.

We now show that $\Sigma^P(0) \subset \Sigma(K)$ whenever $K > \bar{K}$. Fix $\sigma \in \Sigma^P(0)$. Since σ is an equilibrium of the game without a regulator, there cannot be profitable one shot deviations: for every $i \in N$, every history $h_{i,s}$, and every one-shot deviation $\tilde{\sigma}_i \neq \sigma_i$ with $\sigma_i(h_{i,t}) = \tilde{\sigma}_i(h_{i,t})$ for all $h_{i,t} \neq h_{i,s}$,

$$\begin{aligned} V_i((\tilde{\sigma}_i, \sigma_{-i}), h_{i,s}) &\leq V_i(\sigma, h_{i,s}) \\ \iff V_i((\tilde{\sigma}_i, \sigma_{-i}), h_{i,s}) - \mathbb{E}_{(\tilde{\sigma}_i, \sigma_{-i})}[\tau_i | h_{i,s}]K &\leq V_i(\sigma, h_{i,s}) - \mathbb{E}_{\sigma}[\tau_i | h_{i,s}]K \end{aligned}$$

where the second line follows since, under σ , all firms pass the test with probability 1 at

every history. Lemma A.1 then implies that $\sigma \in \Sigma(K)$.

We now prove part (ii). By Proposition 1(ii) and the arguments in its proof, we have that $\Sigma_{RP}(K) \subset \Sigma_{RP}^P(0)$ for all $K > \bar{K}$. To show that $\Sigma_{RP}^P(0) \subset \Sigma_{RP}(K)$, fix $\sigma \in \Sigma_{RP}^P(0)$, and note that no firm can have a profitable one-shot deviation under σ : for every $i \in N$, every history $h_{i,s}$, and every one-shot deviation $\tilde{\sigma}_i \neq \sigma_i$ with $\sigma_i(h_{i,t}) = \tilde{\sigma}_i(h_{i,t})$ for all $h_{i,t} \neq h_{i,s}$,

$$\begin{aligned} V_i((\tilde{\sigma}_i, \sigma_{-i}), h_{i,s}) &\leq V_i(\sigma, h_{i,s}) \\ \iff V_i((\tilde{\sigma}_i, \sigma_{-i}), h_{i,s}) - \mathbb{E}_{(\tilde{\sigma}_i, \sigma_{-i})}[\tau_i | h_{i,s}]K &\leq V_i(\sigma, h_{i,s}) - \mathbb{E}_\sigma[\tau_i | h_{i,s}]K \end{aligned}$$

where the second line follows since, under σ , all firms pass the test with probability 1 at every history. Lemma A.1 then implies that $\sigma \in \Sigma(K)$. Since $\sigma \in \Sigma_{RP}(0)$, it must also be that $\sigma \in \Sigma_{RP}(K)$. ■

Proof of Proposition 2. Consider a competitive equilibrium σ . Consider a bidder i at history $h_{i,t}$ who participates in the auction at time t . For each bid b , let us denote by $D_{i,t}(b) \equiv \mathbf{prob}_{h_{i,t}}(b < \wedge b_{-i,t})$ player i 's subjective probability that she will win the auction if she bids b . In a competitive equilibrium, for any $\epsilon > 0$, equilibrium bid $b_{i,t}$ must satisfy

$$D_{i,t}(b_{i,t} - \epsilon)(b_{i,t} - \epsilon - c_{i,t}) \leq D_{i,t}(b_{i,t})(b_{i,t} - c_{i,t}).$$

Since bidding is costly, it must be that $b_{i,t} - c_{i,t} \geq \kappa$. Hence for all $\epsilon < \kappa$, we obtain that

$$D_{i,t}(b_{i,t} - \epsilon) \leq D_{i,t}(b_{i,t}) \frac{b_{i,t} - c_{i,t}}{b_{i,t} - c_{i,t} - \epsilon} \leq D_{i,t}(b_{i,t}) \frac{\kappa}{\kappa - \epsilon}.$$

By construction, we have that $D_{i,t}(b_{i,t} - \epsilon) \geq D_{i,t}(b_{i,t})$. Hence, $D_{i,t}(\cdot)$ has no mass points at bid $b_{i,t}$. Altogether, this implies that $\mathbf{prob}_\sigma(\tau^0 = 1) = 0$ and $\mathbf{prob}_\sigma(\tau^1 = 1) = 0$. ■

Proof of Proposition 3. Consider a competitive equilibrium σ . Fix a bidder i , and a

history $h_{i,t}$ at which bidder i participates in the auction, bidding $b_{i,t} < r = 1$. For each bid b , let us denote by $D_{i,t}(b) \equiv \mathbf{prob}_{h_{i,t}}(b < \wedge b_{-i,t})$ player i 's subjective probability that she will win the auction if she bids b at history $h_{i,t}$. Note that $D_{i,t}(b_{i,t})(b_{i,t} - c_{i,t}) \geq \kappa > 0$, since otherwise bidder i won't participate. Hence, $b_{i,t} - c_{i,t} \geq \kappa$.

By incentive compatibility, for any $\varepsilon > 0$, bid $b_{i,t} < 1$ must satisfy:

$$\begin{aligned} D_{i,t}(b_{i,t})(b_{i,t} - c_{i,t}) &\geq D_{i,t}(b_{i,t} + \varepsilon)(b_{i,t} + \varepsilon - c_{i,t}) \\ \iff D_{i,t}(b_{i,t}) &\geq D_{i,t}(b_{i,t} + \varepsilon) \frac{b_{i,t} + \varepsilon - c_{i,t}}{b_{i,t} - c_{i,t}} \geq D_{i,t}(b_{i,t} + \varepsilon) \frac{\kappa + \varepsilon}{\kappa}, \end{aligned} \quad (6)$$

and so $D_{i,t}(b_{i,t}) > D_{i,t}(b_{i,t} + \varepsilon)$ for all $\varepsilon > 0$ small. This implies that, for all $\epsilon > 0$, the distribution of winning bids induced by σ must place positive mass on $[r - \epsilon, r]$. Indeed, suppose not, and let $\bar{b} < r = 1$ be the supremum of the support of the distribution of winning bids induced by σ . Note then that, for every history $h_{j,s}$, $D_{j,s}(b) = D_{j,s}(\bar{b})$ for all $b \in (\bar{b}, r]$. By definition of \bar{b} , there exists a sequence of histories $(h_{i_k, t_k})_{k \in \mathbb{N}}$ with $b_{i_k, t_k} \rightarrow \bar{b}$ as $k \rightarrow \infty$. Note that, for all $\epsilon > 0$, and using the inequality in (6),

$$\lim_{k \rightarrow \infty} D_{i_k, t_k}(b_{i_k, t_k} + \epsilon) = \lim_{k \rightarrow \infty} D_{i_k, t_k}(b_{i_k, t_k}) \geq \lim_{k \rightarrow \infty} D_{i_k, t_k}(b_{i_k, t_k} + \epsilon) \frac{\kappa + \epsilon}{\kappa},$$

a contradiction. ■

References

- ABDULKADIROGLU, A. AND K.-S. CHUNG (2003): "Auction design with tacit collusion," *Columbia and Northwestern University Working Paper*.
- ABRANTES-METZ, R., L. FROEB, J. GEWEKE, AND C. TAYLOR (2005): "A variance screen for collusion. US Federal Trade Commission," Tech. rep., Bureau of Economics Working Paper.
- ATHEY, S. AND K. BAGWELL (2008): "Collusion with persistent cost shocks," *Econometrica*, 76, 493–540.

- BAJARI, P. AND L. YE (2003): “Deciding between competition and collusion,” *Review of Economics and Statistics*, 85, 971–989.
- BAKER, J. AND D. RUBINFELD (1999): “Empirical Methods in Antitrust Litigation: Review and Critique,” *American Law and Economics Review*, 1, 386–435.
- BERNHEIM, B. D. AND D. RAY (1989): “Collective dynamic consistency in repeated games,” *Games and Economic Behavior*, 1, 295–326.
- BESANKO, D. AND D. F. SPULBER (1989): “Antitrust enforcement under asymmetric information,” *The Economic Journal*, 99, 408–425.
- CHASSANG, S., K. KAWAI, J. NAKABAYASHI, AND J. M. ORTNER (2019): “Data Driven Regulation: Theory and Application to Missing Bids,” Tech. rep., National Bureau of Economic Research.
- CHASSANG, S. AND J. ORTNER (2019): “Collusion in auctions with constrained bids: Theory and evidence from public procurement,” *Journal of Political Economy*, 127, 2269–2300.
- CHE, Y.-K., D. CONDORELLI, AND J. KIM (2018): “Weak cartels and collusion-proof auctions,” *Journal of Economic Theory*, 178, 398–435.
- CHE, Y.-K. AND J. KIM (2006): “Robustly Collusion-Proof Implementation,” *Econometrica*, 74, 1063–1107.
- (2009): “Optimal collusion-proof auctions,” *Journal of Economic Theory*, 144, 565–603.
- CONLEY, T. G. AND F. DECAROLIS (2016): “Detecting bidders groups in collusive auctions,” *American Economic Journal: Microeconomics*, 8, 1–38.
- CYRENNE, P. (1999): “On antitrust enforcement and the deterrence of collusive behaviour,” *Review of Industrial Organization*, 14, 257–272.
- ELLISON, G. (1994): “Theories of cartel stability and the joint executive committee,” *The Rand journal of economics*, 37–57.
- FARRELL, J. AND E. MASKIN (1989): “Renegotiation in repeated games,” *Games and economic behavior*, 1, 327–360.
- FUDENBERG, D. AND J. TIROLE (1991): “Game theory,” .
- HARRINGTON, J. E. (2004): “Cartel pricing dynamics in the presence of an antitrust authority,” *RAND Journal of Economics*, 651–673.
- (2008): “Detecting cartels,” *Handbook of antitrust economics*, 213, 245.

- IMHOF, D., Y. KARAGÖK, AND S. RUTZ (2016): “Screening for bid-rigging-does it work?” Tech. rep., Université de Fribourg.
- IMHOF, D., Y. KARAGÖK, AND S. RUTZ (2018): “Screening for Bid Rigging — Does It Work?” *Journal of Competition Law & Economics*, 14, 235–261.
- KAWAI, K. AND J. NAKABAYASHI (2018): “Detecting Large-Scale Collusion in Procurement Auctions,” *Available at SSRN 2467175*.
- LACASSE, C. (1995): “Bid rigging and the threat of government prosecution,” *The RAND Journal of Economics*, 398–417.
- MENA-LABARTHE, C. (2015): “Mexican Experience in Screens for Bid-Rigging,” *Competition Policy International*.
- OECD (2018): “Summary of the workshop on cartel screening in the digital era,” .
- PAVLOV, G. (2008): “Auction design in the presence of collusion,” *Theoretical Economics*, 3, 383–429.
- PORTER, R. H. (1983): “A study of cartel stability: the Joint Executive Committee, 1880-1886,” *The Bell Journal of Economics*, 301–314.
- (2005): “Detecting collusion,” *Review of Industrial Organization*, 26, 147–167.