

Data Driven Regulation: Theory and Application to Missing Bids

Sylvain Chassang Kei Kawai Jun Nakabayashi
New York University U.C. Berkeley Kindai University

Juan Ortner*[†]
Boston University

October 17, 2018

Abstract

We document a novel bidding pattern observed in procurement auctions from Japan: winning bids tend to be isolated. This bidding pattern is suspicious in the following sense: it is inconsistent with competitive behavior under arbitrary information structures. Building on this observation, we develop a theory of data-driven regulation based on “safe tests,” i.e. tests that are passed with probability one by competitive bidders, but need not be passed by non-competitive ones. We provide a general class of safe tests exploiting weak equilibrium conditions, and show that such tests reduce the set of equilibrium strategies that cartels can use to sustain collusion. We provide an empirical exploration of various safe tests in our data, as well as discuss collusive rationales for missing bids.

KEYWORDS: missing bids, collusion, regulation, procurement.

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

[†]We are especially indebted to Steve Tadelis for encouragement and detailed feedback. The paper benefited from discussions with Pierpaolo Battigali, Eric Budish, Yeon-Koo Che, Francesco Decarolis, Emir Kamenica, Roger Myerson, Ariel Pakes, Wolfgang Pesendorfer, Andrea Prat, Michael Riordan, Larry Samuelson, Andy Skrzypacz, Paulo Somaini, as well as comments from seminar participants at the 2018 ASSA meeting, Bocconi, the 2017 Berkeley-Sorbonne workshop on Organizational Economics, Boston University, the 2018 ESSET meeting at Gerzensee, HKUST, Johns Hopkins, NYU, the 2017 NYU CRATE conference on theory and econometrics, Princeton, Rochester, the 2018 Stanford SITE conference, the 2018 Transparency in Procurement Conference, the University of Chicago, the University of Copenhagen, and the University of Illinois at Urbana-Champaign.

1 Introduction

One of the key functions of antitrust authorities is to detect and punish collusion. Although concrete evidence is required for successful prosecution, screening devices that flag suspicious firm conduct may help regulators identify collusion, and encourage members of existing cartels to apply for leniency programs. Correspondingly, an active research agenda has sought to build data-driven methods to detect cartels using naturally occurring market data (e.g. Porter, 1983, Porter and Zona, 1993, 1999, Ellison, 1994, Bajari and Ye, 2003, Harrington, 2008). This raises a number of questions: How should regulators act on data-driven evidence of collusion? Wouldn't cartel members adapt their play to the screening tests implemented by the regulator? If so, can we build general tests that do not only target a specific pattern of behavior? Can we ensure that regulatory policies do not end up strengthening cartels?

This paper documents a suspicious bidding pattern observed in procurement auctions from Japan: The density of the bid distribution just above the winning bid is very low. These missing bids are reminiscent of patterns already noted in Hungary (Tóth et al., 2014) and Switzerland (Imhof et al., 2016), and used in regulatory screening programs. We establish that these missing bids indicate non-competitive behavior under a general class of asymmetric information models. Indeed, this missing mass of bids makes it a profitable stage-game deviation for bidders to increase their bids.

We expand on this observation and propose a theory of robust data-driven regulation based on “safe tests,” i.e. tests that are passed with probability one by competitive bidders, but not necessarily by non-competitive ones. We provide a general class of such tests exploiting weakened equilibrium conditions, and show that safe tests cannot help cartels: they necessarily constrain the set of continuation values bidders can use to support collusion. We illustrate the implications of various safe tests in our data, as well as propose several explanations for why missing bids may arise as a by-product of collusion.

Our data comes from multiple datasets of public works procurement auctions taking place

in Japan. One dataset, analyzed by Kawai and Nakabayashi (2018), reports data from 90,000 national-level auctions between 2001 and 2006. A second dataset, studied by Chassang and Ortner (forthcoming), assembles data for 1,500 city-level auctions between 2007 and 2014. We are interested in the distribution of bidders’ margins of victory/defeat. In other terms, for every (bidder, auction) pair, we are interested in the difference Δ between the bidder’s own bid and the most competitive bid among this bidder’s opponents, normalized by the reserve price. When $\Delta < 0$, the bidder won the auction. When $\Delta > 0$ the bidder lost. The finding motivating this paper is summarized by Figure 1, which plots the distribution of bid-differences Δ in the sample of national-level auctions. There is a striking missing mass around $\Delta = 0$. Our first goal is to clarify the sense in which this gap — and other patterns that could be found in data — are suspicious. Our second goal is to formulate a theory of regulatory response to such data.

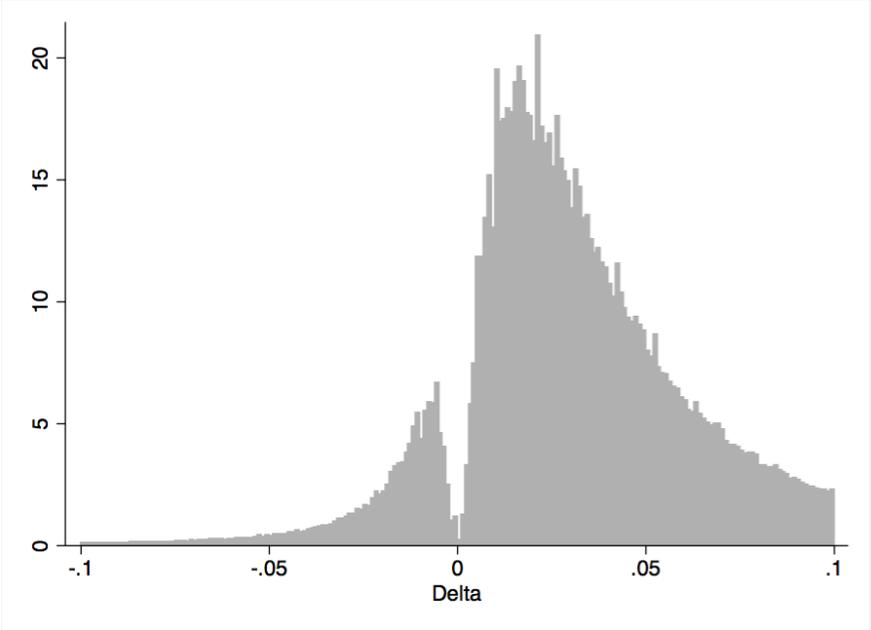


Figure 1: Distribution of bid-differences $\Delta \equiv \frac{\text{own bid} - \min(\text{other bids})}{\text{reserve}}$ over (bidder, auction) pairs in national auction data.

We analyze our data within a fairly general model of repeated play in first-price procure-

ment auctions. A group of firms repeatedly participates in first-price procurement auctions. We assume private values, and rule out intertemporal linkages between actions and payoffs. We allow players to observe arbitrary signals about one another, under the private value assumption. We allow bidders' costs and types to be arbitrarily correlated within and across periods. We say that behavior is competitive, if it is stage-game optimal under the players' information structure.

Our first set of results establishes that the pattern of missing bids illustrated in Figure 1 is not consistent with competitive behavior under any information structure. There is no stochastic process for costs and types (ergodic or not) that would rationalize observed bids in equilibrium. We exploit the fact that in any competitive equilibrium, firms must not find it profitable in expectation to increase their bids. This incentive constraint implies that with high probability the elasticity of firms' *sample counterfactual demand* (i.e., the empirical probability of winning an auction at any given bid) must be bounded above by -1. This condition is not satisfied in our data: because winning bids are isolated, the elasticity of sample counterfactual demand is close to zero. In addition we are able to derive bounds on the minimum number of histories at which non-competitive bidding must happen.

This empirical finding begs the question: what should a regulator do about it? If the regulator investigates industries on the basis of such empirical evidence, won't cartels adapt? Could the regulator make collusion worse by reducing the welfare of competitive players? Our second set of results formulates a theory of regulation based on safe tests. Like the elasticity test described above, safe tests can be passed with probability one provided firms are competitive under *some* information structure. We show how to exploit equilibrium conditions to derive a large class of safe tests. Finally, we show that regulatory policy based on safe tests is a robust improvement over laissez-faire: regulation based on safe tests cannot hurt competitive bidders, and can only reduce the set of enforceable collusive schemes available to cartels.

Our third set of results takes safe tests to the data. We delineate how different moment

conditions (i.e. different deviations) uncover different non-competitive patterns. In addition, we show that the outcomes of our tests are consistent with other proxy evidence for competitiveness and collusion. High bid-to-reserve histories are more likely to fail our tests than low bid-to-reserve histories. Bidding histories before an industry is investigated for collusion are more likely to fail our tests than bidding histories after being investigated for collusion. Altogether this suggests that although safe tests are conservative, they still have bite in practice: they seem to detect collusive industries with positive probability.

Our paper relates primarily to the literature on cartel detection.¹ Porter and Zona (1993, 1999) show that suspected cartel members and non-cartel members bid in statistically different ways. Bajari and Ye (2003) design a test of collusion based on excess correlation across bids. Porter (1983) and Ellison (1994) exploit dynamic patterns of play predicted by the theory of repeated games (Green and Porter, 1984, Rotemberg and Saloner, 1986) to detect collusion. Conley and Decarolis (2016) propose a test of collusion in average-price auctions exploiting cartel members' incentives to coordinate bids. Chassang and Ortner (forthcoming) propose a test of collusion based on changes in behavior around changes in the auction design. Kawai and Nakabayashi (2018) analyze auctions with re-bidding, and exploit correlation patterns in bids across stages to detect collusion. We propose a class of robust, systematic tests of non-competitive behavior that are guaranteed to improve over laissez-faire in equilibrium.

A small set of papers study the equilibrium impact of data driven regulation. 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.² A common observation from these papers is that data driven regulation may backfire, allowing a cartel to sustain higher equilibrium prices. We contribute to this literature by introducing the idea of safe tests, and showing that regulation based on such tests restricts the set of equilibrium

¹See Harrington (2008) for a recent survey.

²Other papers, like Besanko and Spulber (1989) and LaCasse (1995), study static models of equilibrium regulation.

values a cartel can sustain.

Our emphasis on safe tests connects our work to a branch of the microeconomic literature that seeks to identify predictions that can be made for all underlying economic environments. The work of Bergemann and Morris (2013) is particularly relevant: they study the range of behavior in games that can be sustained by some incomplete information structure. A similar exercise is at the heart of our analysis.³ Our work is also related to a branch of the mechanism design literature that considers endogenous responses to collusion (Abdulkadiroglu and Chung, 2003, Che and Kim, 2006, Che et al., 2010).

The tests that we propose, which seek to quantify violations of competitive behavior, are similar in spirit to the tests used in revealed preference theory.⁴ Afriat (1967), Varian (1990) and Echenique et al. (2011) propose tests to quantify the extent to which a given consumption data set violates GARP. More closely related, Carvajal et al. (2013) propose revealed preference tests of the Cournot model. We add to this literature by proposing tests aimed at detecting non-competitive behavior in auctions which are robust to a wide range of informational environments.

Finally, our paper makes an indirect contribution to the literature on the internal organization of cartels. Asker (2010) studies stamp auctions, and analyses the effect of a particular collusive scheme on non-cartel bidders and sellers. Pesendorfer (2000) studies the bidding patterns for school milk contracts and compares the collusive schemes used by strong cartels and weak cartels (i.e., cartels that used transfers and cartels that didn't). Clark and Houde (2013) document the collusive strategies used by the retail gasoline cartel in Quebec. We add to this literature by documenting a puzzling bidding pattern that is poorly accounted for by existing theories. We establish that this bidding pattern is non-competitive, and propose some potential explanations.

The paper is structured as follows. Section 2 describes our data and documents missing

³Also closely related is Bergemann et al. (2017) which studies properties of the first price auction under arbitrary incomplete information.

⁴See Chambers and Echenique (2016) to a recent review of the literature on revealed preference.

bids. Section 3 introduces our theoretical framework. Section 4 shows that missing bids cannot be rationalized under any competitive model. Section 5 generalizes this analysis, and provides safe tests that systematically exploit weak optimality conditions implied by equilibrium. Section 6 proposes normative foundations for safe tests. Section 7 delineates the mechanics of safe tests in real data, and shows that their implications are consistent with other proxies of collusion. Section 8 concludes with an open ended discussion of why missing bids may arise in the context of collusion. Proofs are collected in Appendix A unless mentioned otherwise.

2 Motivating Facts

Our first dataset, described in Kawai and Nakabayashi (2018), consists of roughly 90,000 auctions held between 2001 and 2006 by the Ministry of Land, Infrastructure, Transport and Tourism in Japan (the Ministry). The auctions are first-price auctions with secret reserve price, and re-bidding in case there is no successful winner. The auctions involve construction projects, the median winning bid is USD 600K, and the median participation is 10. The bids of all participants are publicly revealed after the auctions.

For any given firm i participating in auction a with reserve price r , we denote by $b_{i,a}$ the bid of firm i in auction a , and by $b_{-i,a}$ the profile of bids by bidders other than i . We investigate the distribution of

$$\Delta_{i,a} = \frac{b_{i,a} - \min b_{-i,a}}{r}$$

aggregated over firms i , and auctions a . The value $\Delta_{i,a}$ represents the margin by which bidder i wins or loses auction a . If $\Delta_{i,a} < 0$ the bidder won, if $\Delta_{i,a} > 0$ he lost.

The left panel of Figure 2 plots the distribution of bid differences Δ aggregating over all firms and auctions in our sample.⁵ The mass of missing bids around a difference of 0 is

⁵Note that the distribution of normalized bid-differences is skewed to the right since the most competitive alternative bid is a minimum over other bidders' bids.

starkly visible. This pattern can be traced to individual firms as well. The right panel of Figure 2 reports the distribution of bid difference for a single large firm frequently active in our sample of auctions.

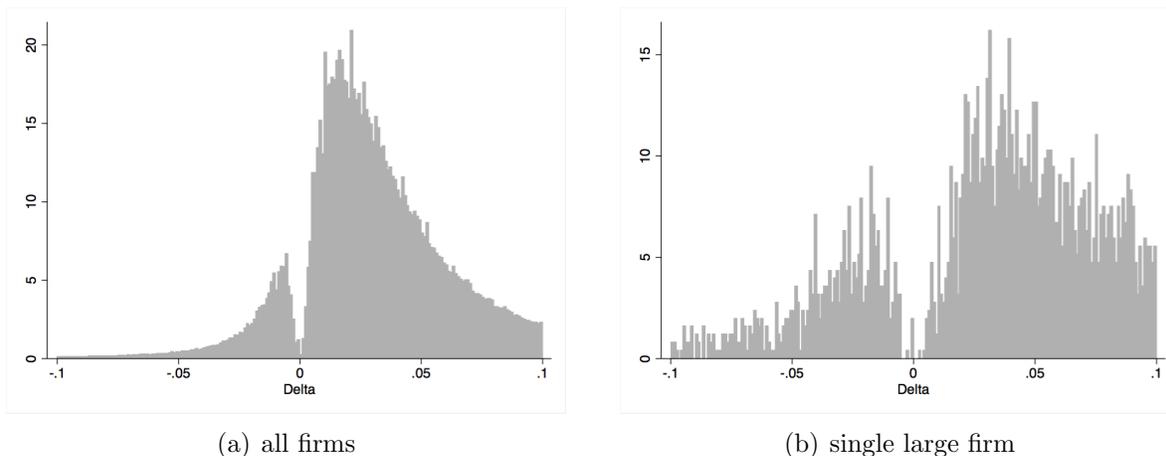


Figure 2: Distribution of bid-difference Δ – national data.

Our second dataset, analyzed in Chassang and Ortner (forthcoming), consists of roughly 1,700 auctions held between 2007 and 2016 by the city of Tsuchiura in the Ibaraki prefecture. Projects are allocated using a standard first-price auction with public reserve price. The median winning bid is USD 130K, and the median participation is 4. Figure 3 plots the distribution of Δ for auctions held in Tsuchiura. Again, we see a significant mass of missing bids around zero.⁶

One key goal of the paper is to show that the bidding patterns in Figures 2 and 3 are inconsistent with competitive behavior under any information structure. While this is different from saying that these patterns are reflective of collusion, missing bids are in fact correlated with plausible proxies of collusion.

Figure 4 breaks down the national-level data in Figure 2 by bid levels: it plots the

⁶Imhof et al. (2016) document a similar bidding pattern in procurement auctions in Switzerland: bidding patterns by several cartels uncovered by the Swiss competition authority presented large differences between the winning bid and the second lowest bid in auctions. See also Tóth et al. (2014).

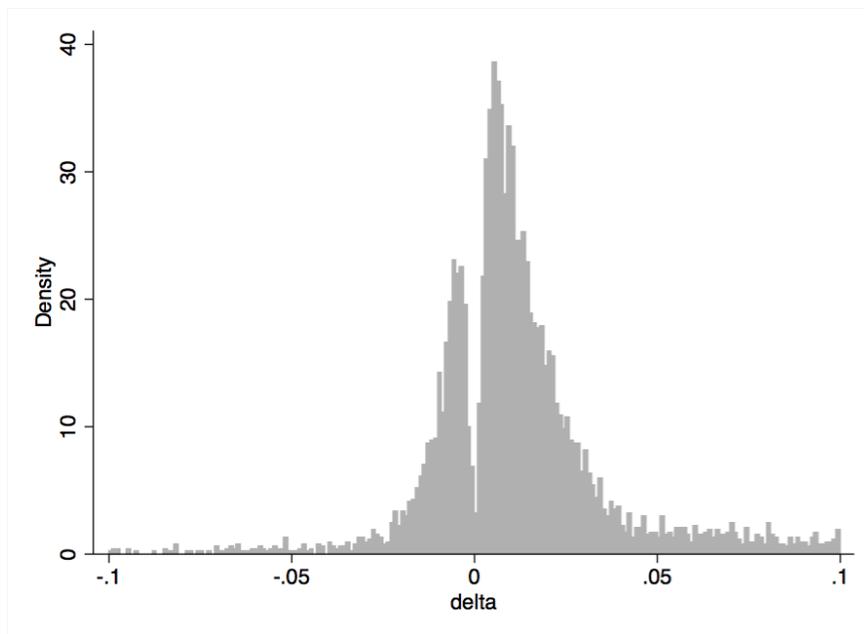


Figure 3: Distribution of bid-difference Δ – city data.

distribution of $\Delta_{i,a} = \frac{b_{i,a} - \min b_{-i,a}}{r}$ for normalized bids $\frac{b_{i,a}}{r}$ below 0.8 and above 0.8. The missing bids mass of bids in Figure 2 all but disappears when we look at bids that are low as a fraction of the reserve price.

Figure 5 presents four cases of firms participating in auctions in our national data that were implicated by the Japanese Fair Trade Commission (JFTC). The four collusion cases are: (i) firms installing traffic signs; (ii) builders of bridge upper structures; (iii) prestressed concrete providers; and (iv) floodgate builders.⁷ The left panels in Figure 5 plot the distribution of Δ before the JFTC started its investigation, and the right panels plot the distribution in the after period. In all cases except (iii), the pattern of missing bids disappears after the JFTC launched its investigation. Interestingly, court documents show that firms in case (iii) initially denied the cases against them, and continued colluding for some time during the after period.

⁷See JFTC Recommendation and Ruling #5-8 (2005) for case (i); JFTC Recommendation and Ruling #12 (2005) for case (ii); JFTC Recommendation #27-28 (2004) and Ruling #26-27 (2010) for case (iii); and JFTC Cease and Desist Order #2-5 (2007) for case (iv).

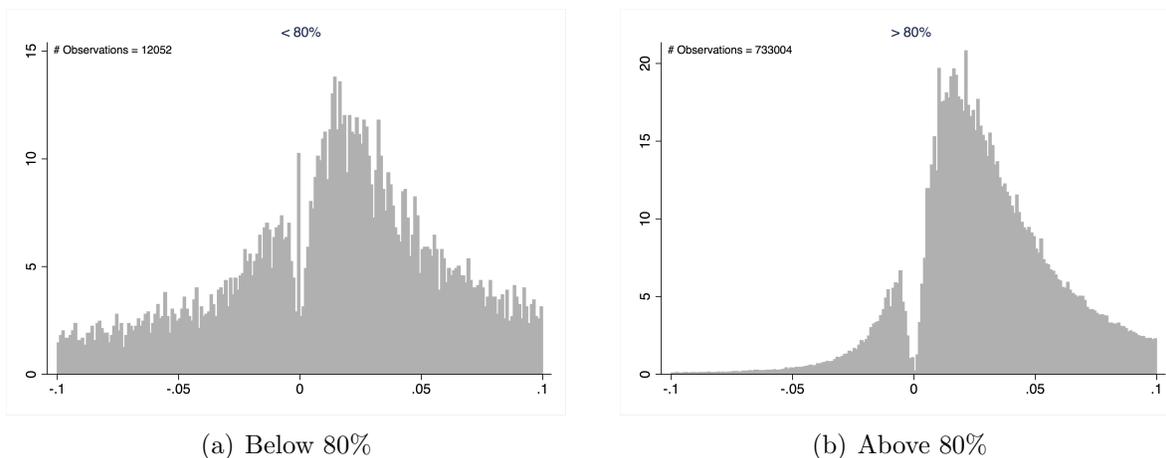


Figure 4: Distribution of bid-difference Δ by bid levels – national data. Left panel plots the distribution of Δ for bids that were below 80% of reserve price, and right panel plots the distribution of Δ for bids that were above 80% of reserve price.

This pattern is not explained by the granularity of bids. One potential explanation to the missing bids in Figures 2 and 3 is that they reflect roundness in bid increments. Figure 4 rules this explanation out: if missing bids were a consequence of the granularity of bids, we should see similar patterns across all bid levels.

This pattern is not explained by renegotiation. Another potential explanation is renegotiation. Indeed, with renegotiation, some firms might have an incentive to bid very aggressively to later renegotiate prices up.

Our national dataset contains data on renegotiated prices, and allows us to rule out this explanation. First, Figure 6 shows that the missing bids pattern persists even if we focus on auctions whose prices were not renegotiated up. Second, the auctions we study include renegotiation provisions that greatly reduce firms' incentives to bid aggressively with the hope of renegotiating to a higher price later on. Specifically renegotiated prices depend on the level of the initial bid: if the project is estimated to cost more than initially thought, the renegotiated price is given by $\frac{\text{initial bid}}{\text{reserve price}} \times (\text{new cost estimate})$. This implies that non-profitable excessively competitive bids are likely to remain unprofitable after renegotiation.

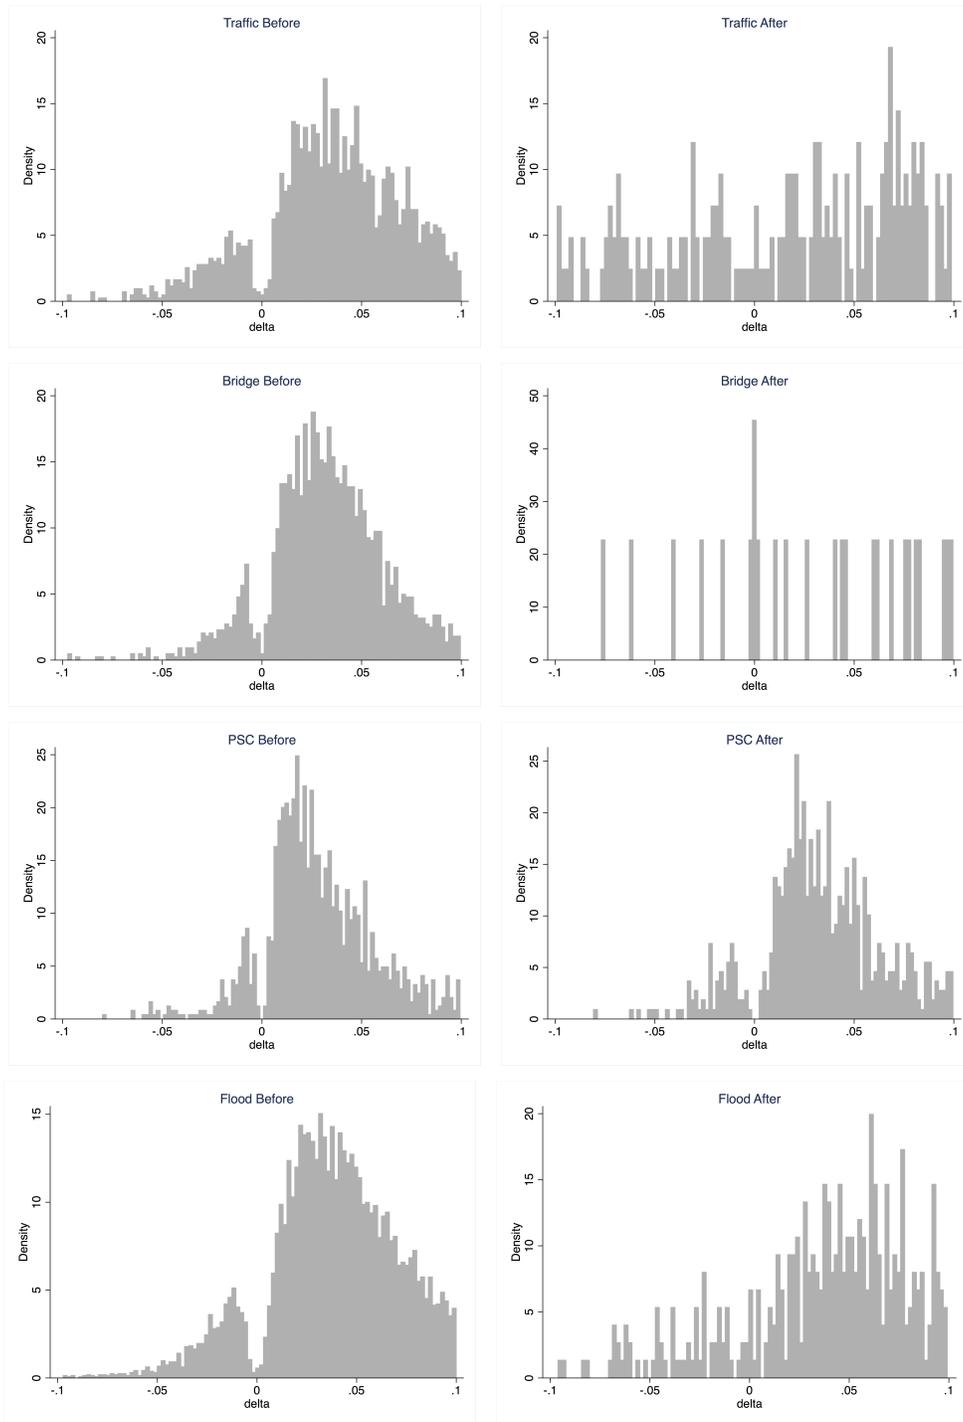


Figure 5: Distribution of bid-difference Δ – cartel cases in national data, before and after JFTC investigation.

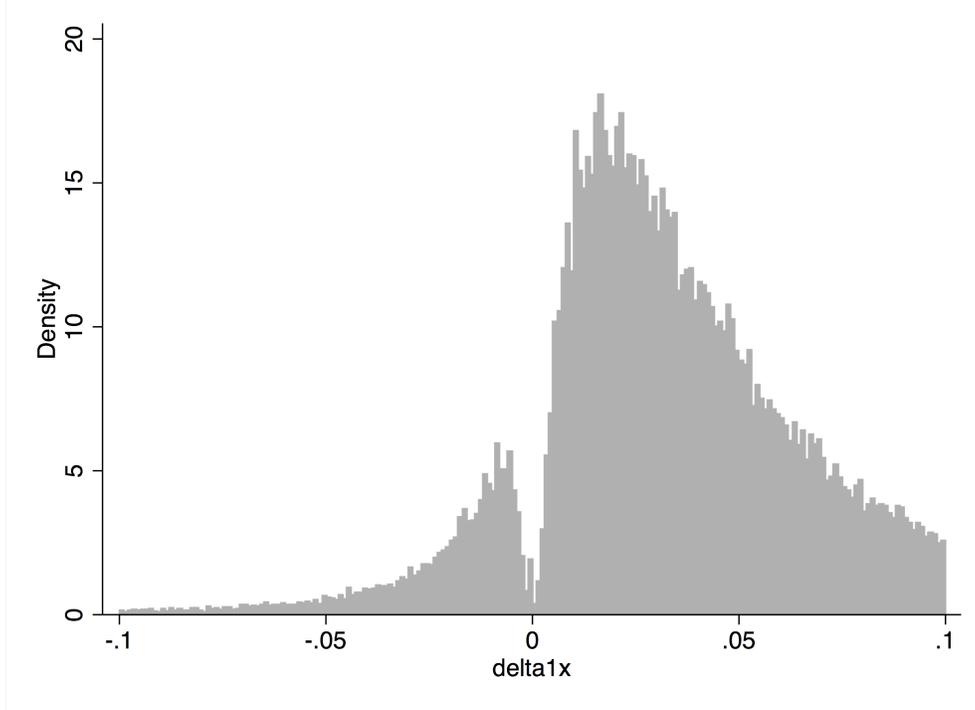


Figure 6: Distribution of bid-difference Δ – auctions whose price was not renegotiated upwards.

Our objectives in this paper are: *(i)* to formalize why this pattern is suspicious; *(ii)* to delineate what it implies about bidding behavior and the competitiveness of auctions in our sample; *(iii)* to formulate a theory of regulation based on data-driven tests; and *(iv)* to propose possible explanations for why this behavior may arise under collusive bidding. To do so we use a model of repeated auctions.

3 Framework

3.1 The Stage Game

We consider a dynamic setting in which, at each period $t \in \mathbb{N}$, a buyer needs to procure a single project. The auction format is a first-price auction with reserve price r , which we normalize to $r = 1$.

In each period $t \in \mathbb{N}$, a set $\widehat{N}_t \subset N$ of bidders is able to participate in the auction, where N is the overall set of bidders. We think of this set of participating firms as those eligible to produce in the current period.⁸ The set of eligible bidders can vary over time.

Realized costs of production for eligible bidders $i \in \widehat{N}_t$ are denoted by $\mathbf{c}_t = (c_{i,t})_{i \in \widehat{N}_t}$. Each bidder $i \in \widehat{N}_t$ submits a bid $b_{i,t}$. 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 opponents at time t . The procurement contract is allocated to the bidder submitting the lowest bid at a price equal to her bid.

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 $\gamma_t = (\gamma_{i,t})_{i \in \widehat{N}_t}$ with $\gamma_{i,t} \in [0, 1]$ for all i, t . When lowest bids are tied, the allocation to a lowest bidder i is

$$x_{i,t} = \frac{\gamma_{i,t}}{\sum_{\{j \in \widehat{N}_t \text{ s.t. } b_{j,t} = \min_k b_{k,t}\}} \gamma_{j,t}}.$$

Participants discount future payoffs using common discount factor $\delta < 1$. Bids are publicly revealed at the end of each period.⁹

Costs. We allow for costs that are serially correlated over time, and that may be correlated across firms within each auction. Denoting by $\langle \cdot, \cdot \rangle$ the usual dot-product we assume that costs take the form

$$c_{i,t} = \langle \alpha_i, \theta_t \rangle + \varepsilon_{i,t} > 0 \tag{1}$$

where

- parameters $\alpha_i \in \mathbb{R}^k$ are fixed over time;
- $\theta_t \in \mathbb{R}^k$ may be unknown to the bidders at the time of bidding, but is revealed to

⁸See Chassang and Ortner (forthcoming) for a treatment of endogenous participation by cartel members.

⁹For the sake of concision, we do not consider the possibility of transfers at this point. It does not change our analysis.

bidders at the end of period t ; we assume that θ_t follows a Markov chain but do not assume that there are finitely many states, or that the chain is irreducible;

- $\varepsilon_{i,t}$ is i.i.d. with mean zero conditional on θ_t .

In period t , bidder i obtains profits

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

Note that costs include both the direct costs of production and the opportunity cost of backlog.

The sets \widehat{N}_t of bidders are independent across time conditional on θ_t , i.e.

$$\widehat{N}_t | \theta_{t-1}, \widehat{N}_{t-1}, \widehat{N}_{t-2} \dots \sim \widehat{N}_t | \theta_{t-1}.$$

Information. In each period t , bidder i gets a signal $z_{i,t}$ that is conditionally i.i.d. given $(\theta_t, (c_{j,t})_{j \in \widehat{N}_t})$. This allows our model to nest many informational environments, including asymmetric information private value auctions, common value auctions, as well as complete information. Bids \mathbf{b}_t are observable at the end of the auction.

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

3.2 Solution Concepts

The public history h_t at period t takes the form

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

Our solution concept is perfect public Bayesian equilibrium (Athey and Bagwell, 2008). Because state θ_t is 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 perfect public Bayesian equilibrium consists only of a strategy profile σ , such that for all $i \in N$,

$$\sigma_i : h_t \mapsto b_{i,t}(z_{i,t}),$$

where bids $b_{i,t}(z_{i,t}) \in \Delta([0, r])$ depend on the public history and on the information available at the time of decision making. We let \mathcal{H} denote the set of all public histories.

We emphasize the class of competitive equilibria, or in this case Markov perfect equilibria (Maskin and Tirole, 2001). In a competitive equilibrium, players condition their play only on payoff relevant parameters.

Definition 1 (competitive strategy). *We say that σ is competitive (or Markov perfect) if and only if $\forall i \in N$ and $\forall h_t \in \mathcal{H}$, $\sigma_i(h_t, z_{i,t})$ depends only on $(\theta_{t-1}, z_{i,t})$.*

We say that a strategy profile σ is a competitive equilibrium if it is a perfect public Bayesian equilibrium in competitive strategies.

We note that in a competitive equilibrium, firms must be playing a stage-game Nash equilibrium at every period; that is, firms must play a static best-reply to the actions of their opponents.

Competitive histories. Generally, an equilibrium may include periods in which (a subset of) firms collude and periods in which firms compete. This leads us to define competitive histories.

Definition 2 (competitive histories). *Fix a common knowledge profile of play σ and a history $h_{i,t} = (h_t, z_{i,t})$ 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.

3.3 Safe Tests

Let H_∞ denote the set of coherent full public histories $(h_{i,t})_{i \in N, t \geq 0}$. A test τ is a mapping from H_∞ to $\{0, 1\}$.

Definition 3 (safe tests). *We say that τ_i is unilaterally safe for firm i if and only if for all $\lambda \in \Lambda$, and all profiles σ such that firm i is competitive, then λ -a.s. $\tau_i(h) = 0$ for all $h \in H_\infty$.*

We say that τ is jointly safe if and only if for all $\lambda \in \Lambda$, and all profiles σ such that all players $i \in N$ are competitive, then λ -a.s. $\tau(h) = 0$ for all $h \in H_\infty$.

4 Missing Bids are Inconsistent with Competition

In this section, we show how to exploit equilibrium conditions at different histories to obtain bounds on the share of competitive histories. The first step is to identify moments of counterfactual demand that can be estimated from data, even though the subjective residual demand faced by bidders can vary with the history.

4.1 Counterfactual demand

Fix a perfect public Bayesian equilibrium σ . For all histories $h_{i,t} = (h_t, z_{i,t})$ and all bids $b' \in [0, r]$, player i 's *counterfactual demand* at $h_{i,t}$ is

$$D_i(b'|h_{i,t}) \equiv \text{prob}_\sigma(\wedge \mathbf{b}_{-i,t} > b'|h_{i,t}).$$

For any finite set of histories $H = \{(h_t, z_{i,t})\} = \{h_{i,t}\}$, and any scalar $\rho \in (-1, \infty)$, define

$$\bar{D}(\rho|H) \equiv \sum_{h_{i,t} \in H} \frac{1}{|H|} D_i((1 + \rho)b_{i,t}|h_{i,t})$$

to be the average counterfactual demand for histories in H , and

$$\widehat{D}(\rho|H) \equiv \sum_{h_{i,t} \in H} \frac{1}{|H|} \mathbf{1}_{\wedge b_{-i,t} > (1+\rho)b_{i,t}}.$$

Definition 4. *We say that set H is adapted to the players' information if and only if the event $h_{i,t} \in H$ is measurable with respect to player i 's information at time t prior to bidding.*

For instance, the set of auctions for a specific industry with reserve prices above a certain threshold is adapted. In contrast, the set of auctions in which the margin of victory is below a certain level is not. The ability to legitimately vary the conditioning set H lets us explore the competitiveness of auctions in particular settings of interest.

Lemma 1. *Consider a sequence of adapted sets $(H_n)_{n \in \mathbb{N}}$ such that $\lim_{n \rightarrow \infty} |H_n| = \infty$. Under any perfect public Bayesian equilibrium σ , with probability 1, $\widehat{D}(\rho|H_n) - \overline{D}(\rho|H_n) \rightarrow 0$.*

In other words, in equilibrium, the sample residual demand conditional on an adapted set of histories converges to the true subjective aggregate conditional demand. This result is the central equilibrium implication that we rely on. Note that it is a weaker requirement than equilibrium itself. It may fail under sufficiently non-common priors, but will hold if participants use data-driven predictors of demand satisfying no-regret.

4.2 A Test of Non-Competitive Behavior

The pattern of bids illustrated in Figures 1, 2 and 3 is striking. Our first main result shows that its more extreme forms are inconsistent with competitive behavior.

Proposition 1. *Let σ be a competitive equilibrium. Then,*

$$\forall h_i, \quad \left. \frac{\partial \log D_i(b'|h_i)}{\partial \log b'} \right|_{b'=b_i^+(h_i)} \leq -1, \quad (2)$$

$$\forall H, \quad \left. \frac{\partial \log \overline{D}(\rho|H)}{\partial \rho} \right|_{\rho=0^+} \leq -1. \quad (3)$$

In other terms, under any non-collusive equilibrium, the elasticity of counterfactual demand must be less than -1 at every history. The data presented in the left panel of Figure 2 contradicts the results in Proposition 1. Note that for every $i \in N$ and every h_i ,

$$\begin{aligned} D_i(b'|h_i) &= \text{prob}_\sigma(b' - \wedge \mathbf{b}_{-i} < 0 | h_i) \\ &= \text{prob}_\sigma(b' - b_i + \Delta_i < 0 | h_i), \end{aligned}$$

where we used $\Delta_i = \frac{b_i - \wedge \mathbf{b}_{-i}}{r} = b_i - \wedge \mathbf{b}_{-i}$ (since we normalized $r = 1$). Since the density of Δ_i at 0 is essentially 0 for some sets of histories in our data, the elasticity of demand is approximately zero at these histories.

Proof. Consider a competitive equilibrium σ . Let u_i denote the flow payoff of player i , and let $V(h_{i,t}) \equiv \mathbb{E}_\sigma \left(\sum_{s \geq t} \delta^{s-t} u_{i,s} \mid h_{i,t} \right)$ denote her discounted expected payoff at history $h_{i,t} = (h_t, z_{i,t})$.

Let $b_{i,t} = b$ be the bid that bidder i places at history $h_{i,t}$. Since $b_{i,t} = b$ is an equilibrium bid, it must be that for all bids $b' > b$,

$$\begin{aligned} \mathbb{E}_\sigma \left[(b - c_{i,t}) \mathbf{1}_{\wedge \mathbf{b}_{-i,t} > b} + \delta V(h_{i,t+1}) \mid h_{i,t}, b_{i,t} = b \right] \\ \geq \mathbb{E}_\sigma \left[(b' - c_{i,t}) \mathbf{1}_{\wedge \mathbf{b}_{-i,t} > b'} + \delta V(h_{i,t+1}) \mid h_{i,t}, b_{i,t} = b' \right] \end{aligned}$$

Since σ is competitive, $\mathbb{E}_\sigma[V(h_{i,t+1}) \mid h_{i,t}, b_{i,t} = b] = \mathbb{E}_\sigma[V(h_{i,t+1}) \mid h_{i,t}, b_{i,t} = b']$. Hence, we must have

$$\begin{aligned} bD_i(b|h_{i,t}) - b'D_i(b'|h_{i,t}) &= \mathbb{E}_\sigma \left[b \mathbf{1}_{\wedge \mathbf{b}_{-i,t} > b} - b' \mathbf{1}_{\wedge \mathbf{b}_{-i,t} > b'} \mid h_{i,t} \right] \\ &\geq \mathbb{E}_\sigma \left[c_{i,t} (\mathbf{1}_{\wedge \mathbf{b}_{-i,t} > b} - \mathbf{1}_{\wedge \mathbf{b}_{-i,t} > b'}) \mid h_{i,t} \right] \geq 0, \end{aligned} \tag{4}$$

where the last inequality follows since $c_{i,t} \geq 0$. Inequality (4) implies that, for all $b' > b$,

$$\frac{\log D_i(b'|h_i) - \log D_i(b|h_i)}{\log b' - \log b} \leq -1.$$

Inequality (2) follows from taking the limit as $b' \rightarrow b$. Inequality (3) follows from summing (4) over histories in H , and performing the same computations. ■

As the proof highlights, this result exploits the fact that in procurement auctions, zero is a natural lower bound for costs (see inequality (4)). In contrast, for auctions where bidders have a positive value for the good, there is no obvious upper bound to valuations to play that role. One would need to impose an ad hoc upper bound on values to establish similar results.

An implication of Proposition 1 is that, in our data, bidders have a short-term incentive to increase their bids. Because of the need to keep participants from bidding higher, for every $\epsilon > 0$ small, there exists $\nu > 0$ and a positive mass of histories $h_{i,t} = (h_t, z_{i,t})$ such that,

$$\delta \mathbb{E}_\sigma [V(h_{i,t+1}) | h_{i,t}, b_i(h_{i,t})] - \delta \mathbb{E}_\sigma [V(h_{i,t+1}) | h_{i,t}, b_i(h_{i,t})(1 + \epsilon)] > \nu. \quad (5)$$

In other terms, equilibrium σ must give bidders a dynamic incentive not to overcut the winning bid.

Proposition 1 proposes a simple test of whether an adapted dataset H can be generated by a competitive equilibrium or not. We now refine this test to obtain bounds on the minimum share of non-competitive histories needed to rationalize the data. We begin with a simple loose bound and then propose a more sophisticated program resulting in tighter bounds.

4.3 Estimating the share of competitive histories

It follows from Proposition 1 that missing bids cannot be explained in a model of competitive bidding. We now establish that competitive behavior must fail at a large number of histories in order to explain isolated winning bids. This implies that bidders have frequent opportunities to learn that their bids are not optimal.

Fix a perfect public Bayesian equilibrium σ and a finite set of histories H . Let $H^{\text{comp}} \subset H$ be the set of competitive histories in H , and let $H^{\text{coll}} = H \setminus H^{\text{comp}}$. Define $s_{\text{comp}} \equiv \frac{|H^{\text{comp}}|}{|H|}$ to be the fraction of competitive histories in H .

For all histories $h_{i,t} = (h_t, z_{i,t})$ and all bids $b' \geq 0$, player i 's *counterfactual revenue* at $h_{i,t}$ is

$$R_i(b'|h_{i,t}) \equiv b' D_i(b'|h_{i,t}).$$

For any finite set of histories H and scalar $\rho \in (-1, \infty)$, define

$$\bar{R}(\rho|H) \equiv \sum_{h_{i,t} \in H} \frac{1}{|H|} (1 + \rho) b_{i,t} D_i((1 + \rho) b_{i,t} | h_{i,t})$$

to be the average counterfactual revenue for histories in H . Our next result builds on Proposition 1 to derive a bound on s_{comp} .

Proposition 2. *The share s_{comp} of competitive auctions is such that*

$$s_{\text{comp}} \leq 1 - \sup_{\rho > 0} \frac{\bar{R}(\rho|H) - \bar{R}(0|H)}{\rho}.$$

Proof. For any $\rho > 0$,

$$\begin{aligned} \frac{1}{\rho} [\bar{R}(\rho|H) - \bar{R}(0|H)] &= s_{\text{comp}} \frac{1}{\rho} [\bar{R}(\rho|H^{\text{comp}}) - \bar{R}(0|H^{\text{comp}})] \\ &\quad + (1 - s_{\text{comp}}) \frac{1}{\rho} [\bar{R}(\rho|H^{\text{coll}}) - \bar{R}(0|H^{\text{coll}})] \\ &\leq 1 - s_{\text{comp}}. \end{aligned}$$

The last inequality follows from two observations. First, since the elasticity of counterfactual demand is bounded above by -1 for all competitive histories (Proposition 1), it follows that $\bar{R}(\rho|H^{\text{comp}}) - \bar{R}(0|H^{\text{comp}}) \leq 0$. Second,

$$\frac{1}{\rho}[\bar{R}(\rho|H^{\text{coll}}) - \bar{R}(0|H^{\text{coll}})] \leq \frac{1}{\rho}((1 + \rho)\bar{R}(0|H^{\text{coll}}) - R(0|H^{\text{coll}})) = \bar{R}(0|H^{\text{coll}}) \leq r = 1.$$

This concludes the proof. ■

In words, if total revenue in histories H increases by more than $\kappa \times \rho$ when bids are uniformly increased by $(1 + \rho)$, the share of competitive auctions in H is bounded above by $1 - \kappa$.

For each $\rho \in (-1, \infty)$, define

$$\hat{R}(\rho|H) \equiv \sum_{h_{i,t} \in H} \frac{1}{|H|} (1 + \rho)b_{i,t} \mathbf{1}_{\wedge b_{-i,t} > (1+\rho)b_{i,t}}.$$

Note that $\hat{R}(\rho|H)$ is the sample analog of counterfactual revenue. A result identical to Lemma 1 establishes that $\hat{R}(\rho|H)$ is a consistent estimator of $\bar{R}(\rho|H)$, whenever set H is adapted.

In the extreme case where the density of competing bids is zero just above winning bids, we have that $\bar{R}(\rho|H) - \bar{R}(0|H) \simeq \rho$ for ρ small. This implies that $s_{\text{comp}} \simeq 0$. Most histories are non-competitive.

5 A General Class of Safe Tests

We now extend the approach of Section 4 to derive a more general class of safe tests that exploits the information content of both upward and downward deviations. In addition we allow the regulator or econometrician to formulate plausible constraints on costs, and

incomplete information.

5.1 A General Result

Take as given an adapted set of histories H , corresponding to a set A of auctions. Take as given scalars $\rho_n \in (-1, \infty)$ for $n \in \mathcal{M} = \{-\underline{n}, \dots, \bar{n}\}$, such that $\rho_0 = 0$ and $\rho_n < \rho_{n'}$ for all $n' > n$. For each history $h_{i,t} \in H$, let $d_{h_{i,t},n} \equiv D_i((1 + \rho_n)b_{h_{i,t}}|h_{i,t})$. That is, $(d_{h_{i,t},n})_{n \in \mathcal{M}}$ is firm i 's subjective counterfactual demand at history $h_{i,t}$, when applying a coefficient $1 + \rho_n$ to its original bid. For any auction a and associated histories $h \in a$, an environment at a is a tuple $\omega_a = (d_{n,h}, c_h)_{h \in a}$. We let $\omega_A = (\omega_a)_{a \in A}$ denote the profile of environments across auctions $a \in A$.

Definition 5. *A set of histories $H' \subset H$ is adapted conditional on ω_A if and only if for all firms i and history h_i , the event $h_i \in H'$ is measurable with respect to the information of firm i at h_i implied by environment ω_A .*

This allows us to adjust relevant sets of histories to the environment used to best rationalize the data. For each deviation n , environment $\omega_A = (\omega_a)_{a \in A}$ and adapted set of histories $H' \subset H$ define

$$D_n(\omega_A, H') \equiv \frac{1}{|H'|} \sum_{h_{i,t} \in H'} d_{h_{i,t},n} \quad \text{and} \quad \widehat{D}_n(H') \equiv \frac{1}{|H'|} \sum_{h_{i,t} \in H'} \mathbf{1}_{(1+\rho_n)b_{h_{i,t}} < \wedge b_{-i,h_{i,t}}}.$$

We formulate the problem of inference about the environment ω_A as a constrained maximization problem defined by three objects to be chosen by the econometrician:

- (i) $u(\omega_a)$ the objective function to be maximized. For instance, the number of histories $h \in a$ that are rationalized as competitive under ω_a .
- (ii) Ω the set of plausible economic environments ω_A .

(iii) An adapted set of histories $H(\omega_A) \subset H$ which may depend on environment ω_A .

For instance, the set of competitive histories under environment ω_A .

Let $U(\omega_A) = \sum_{a \in A} u(\omega_a)$. For any tolerance function $T : \mathbb{N} \rightarrow \mathbb{R}^+$, we define inference problem (P) as

$$\begin{aligned} \widehat{U} &= \max_{\omega_A \in \Omega} U(\omega_A) & (P) \\ \text{s.t. } \forall n, \quad D_n(\omega_A, H(\omega_A)) &\in \left[\widehat{D}_n(\omega_A, H(\omega_A)) - T(|H(\omega_A)|), \right. \\ &\quad \left. \widehat{D}_n(\omega_A, H(\omega_A)) + T(|H(\omega_A)|) \right]. \end{aligned}$$

In words, estimator \widehat{U} is a robust upper-bound of moment $U(\omega_A)$. It exploits the assumption that ω_A is within the set of plausible environments Ω , as well as the fact that the sample counterfactual demand aggregated over any set of adapted histories must converge to the bidders' expected counterfactual demand. The following result holds.

Proposition 3. *Suppose the true environment is $\omega_A \in \Omega$. Then, with probability at least $1 - 2|\mathcal{M}| \exp(-\frac{1}{2}T(|H(\omega_A)|)^2|H(\omega_A)|)$, $\widehat{U} \geq U(\omega_A)$.*

We now show that by applying Proposition 3 to different tuples $(u, \Omega, H(\cdot))$, we can derive robust bounds on different measures of non-competitive behavior. These bounds imply natural safe tests.

5.2 Maximum Share of Competitive Histories

We first use Proposition 3 to extend Proposition 1, and provide a tighter upper bound on the share of competitive histories in H . Note that the set of competitive histories H_{comp} is adapted to the bidders' private information. At every competitive history $h \in H$, there must

exist costs c_h and subjective demands $d_h = (d_{h,n})_{n \in \mathcal{N}}$ satisfying

$$\text{feasibility: } c_h \in [0, b_h]; \quad \forall n, d_{h,n} \in [0, 1]; \quad \forall n, n' > n, d_{h,n} \geq d_{h,n'}. \quad (\text{F})$$

$$\text{incentive compatibility: } \quad \forall n, \quad [(1 + \rho_n)b_h - c_h] d_{h,n} \leq [(1 + \rho_0)b_h - c_h] d_{h,0} \quad (\text{IC})$$

The objective function u simply counts the number of histories satisfying these conditions:

$$u(\omega_a) \equiv \frac{1}{|H|} \sum_{h \in a \cap H} \mathbf{1}_{(d_h, c_h) \text{ satisfy (F) \& (IC)}}. \quad (6)$$

Demand consistency is checked for competitive histories: $H(\omega_A) = H_{\text{comp}}(\omega_A) = \{h \in H \text{ s.t. } (d_h, c_h) \text{ satisfy (F), (IC)}\}$.

We illustrate the use of plausible economic environments Ω by considering markup constraints, as well as constraints on the informativeness of signals bidders get:

$$\forall h \in H_{\text{comp}}, \quad \frac{b_h}{c_h} \leq 1 + m \quad \text{and} \quad \forall n, \quad \left| \log \frac{d_{h,n}}{1 - d_{h,n}} - \log \frac{D_n}{1 - D_n} \right| \leq k \quad (\text{EP})$$

where $m \in [0, +\infty]$ is a maximum markup, and $k \in [0, +\infty)$ provides an upper bound to the information contained in any signal. Note that $\log \frac{d_{h,n}}{1 - d_{h,n}} = \log \frac{\text{prob}(Z|h)}{\text{prob}(\neg Z|h)}$ for Z the event that $\mathbf{b}_{-i} > (1 + \rho_n)b_h$. Hence, k is a bound on the log-likelihood ratio of signals that bidders get. One focal case in which $k = 0$ is that of i.i.d. types.

Program (P) then becomes

$$\begin{aligned} & \max_{\omega_A \in \Omega} \frac{|H(\omega_A)|}{|H|} \\ \text{s.t. } & \forall n, \quad D_n(\omega_A, H(\omega_A)) \in \left[\widehat{D}_n(\omega_A, H(\omega_A)) - T(|H(\omega_A)|), \right. \\ & \left. \widehat{D}_n(\omega_A, H(\omega_A)) + T(|H(\omega_A)|) \right]. \end{aligned}$$

\hat{U} provides an upper bound to the share of competitive histories in H , yielding the following corollary.

Corollary 1. *Suppose that the true environment ω_A satisfies (EP), and that the true share of competitive histories under environment ω_A is $s_{comp} \in (0, 1]$. Then, with probability at least $1 - 2|\mathcal{M}| \exp(-\frac{1}{2}T(|H(\omega_A)|)^2|H(\omega_A)|)$, $\hat{U} \geq s_{comp}$.*

This robust bound lets us define safe tests. For any threshold fraction $s_0 \in (0, 1]$ of competitive histories, we define our candidate safe test by $\tau \equiv \mathbf{1}_{\hat{U} \geq s_0}$.

Corollary 2. *Whenever $T(\cdot)$ satisfies $\lim_{|H| \rightarrow \infty} \exp(-\frac{1}{2}T(|H|)^2|H|) = 0$, τ is a safe test.*

By varying the initial set H of adapted histories, we can make test τ safe for a given firm, or for a given industry. Specifically, if H is the set of histories of firm i , then test τ is unilaterally safe for firm i .

For finite data, we can choose $T(\cdot)$ to determine the significance level of test τ . For instance, for the test to have a robust significance level of $\alpha \in (0, 1)$, we set $T(|H|)$ such that $2|\mathcal{M}| \exp(-\frac{1}{2}T(|H|)^2|H|) = \alpha$. Note that this is a very conservative estimate of significance. One may obtain less conservative significance estimates by using asymptotic concentration bounds tighter than the ones used to establish Lemma 1.

5.3 Bounds on Other Moments

Corollary 1 uses Proposition 3 to obtain bounds on the share of competitive histories. We now show how to use Proposition 3 to bound other moments of interest: the share of competitive auctions, and the expected profits left on the table by non-optimizing bidders.

Maximum share of competitive auctions. The bound on the share of competitive histories provided by Corollary 1 allows some histories in the same auctions to have different competitive/non-competitive status. This may underestimate the prevalence of non-competition. In particular, if one player is non-competitive, she must expect other players to

be non-competitive in the future. Otherwise her stage-game best reply would be a profitable dynamic deviation.

For this reason we are interested in providing an upper bound on the share of competitive *auctions*, where an auction is considered to be competitive if and only if every player is competitive at their respective histories, and that is common knowledge. This common knowledge requirement ensures that the set of auctions that are competitive in our sense is adapted.

We define the objective function to be

$$u(\omega_a) = \frac{1}{|A|} \mathbf{1}_{\forall h \in a, (d_h, c_h) \text{ satisfy (F) \& (IC)}}$$

For each $H' \subset H$, let $A_{H'} \subset A$ denote the set of auctions corresponding to histories in H' . For every environment ω_A , let

$$\begin{aligned} H(\omega_A) &\in \arg \max_{H' \subset H} |A_{H'}| \\ \text{s.t. } \forall a \in A_{H'}, \forall h \in a, & \quad (d_h, c_h) \text{ satisfy (F) \& (IC)}. \end{aligned}$$

Program (P) then becomes

$$\begin{aligned} \widehat{U} &= \max_{\omega_A \in \Omega} \frac{|A_{H(\omega_A)}|}{|A|} \\ \text{s.t. } \forall n, \quad D_n(\omega_A, H(\omega_A)) &\in \left[\widehat{D}_n(\omega_A, H(\omega_A)) - T(|H(\omega_A)|), \right. \\ &\quad \left. \widehat{D}_n(\omega_A, H(\omega_A)) + T(|H(\omega_A)|) \right]. \end{aligned}$$

\widehat{U} provides an upper bound to the fraction of competitive auctions.

Total deviation temptation. Regulators may want to investigate an industry only if firms fail to optimize in a significant way. We show how to use Proposition 3 to derive a lower bound on the bidders' deviation temptation. Given an environment ω_a , we define

objective

$$u(\omega_a) \equiv \frac{1}{|A|} \sum_{h \in a} \left[(b_h - c_h) d_{h,0} - \max_{n \in \{-\bar{n}, \dots, \bar{n}\}} [(1 + \rho_n) b_h - c_h] d_{h,n} \right]$$

and estimate demand using histories $H(\omega_A) = H$.

In this case, with large probability, $-\widehat{U}$ is a lower bound for the average total deviation-temptation per auction. This lets a regulator assess the extent of firms' failure to optimize before launching a costly audit. In addition, since the sum of deviation temptations must be compensated by a share of the cartel's future excess profits (along the lines of Levin (2003)), \widehat{U} provides an indirect measure of the excess profits generated by the cartel.

6 Normative Foundations for Safe Tests

In this section we provide normative foundations for safe tests. We show that when punishments are severe enough, they can be used to place constraints on potential cartel members without creating new collusive equilibria.

A game of regulatory oversight. We study the equilibrium impact of data driven regulation within the following framework. From $t = 0$ to $t = \infty$, firms in N play the infinitely repeated game in Section 3. At $t = \infty$, after firms played the game, a regulator runs a safe test on firms in \mathcal{N} based on the realized history $h_\infty \in H_\infty$. We consider two different settings:

- (i) The regulator runs a unilaterally safe test τ_i on each firm $i \in N$.

Firm i incurs an un-discounted 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 regulator 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 $\tau(h_\infty) = 1$.¹⁰

When $K = 0$, under either form of testing the game collapses to the model in Section 3.

¹⁰Since the penalty K is undiscounted, the game is not continuous at infinity whenever $K > 0$.

Individually safe tests. 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 . Let $\bar{K} \equiv \frac{\delta}{1-\delta}$. \bar{K} serves as a rough upper bound on the difference in the continuation values a player obtains for different actions.

Proposition 4 (safe tests do not create new equilibria). *Assume the regulator runs unilaterally safe tests. For all $K > \bar{K}$, $\Sigma(K) \subset \Sigma(0)$.*

We now give an intuition as to why Proposition 4 holds. Note first that when the penalty K is large enough (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 $\sigma \in \Sigma(K)$. Consider a public history h_t , and let $\beta = (\beta_i)_{i \in N}$ be the bidding profile that firms use at h_t 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 under σ , with $V_i : \mathbf{b} \mapsto \mathbb{R}^{|N|}$ 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 ,

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

where the second line follows since all firms pass the test with probability 1 after all histories. In words, strategy profile σ is such that, at each history h_t , no firm i has a profitable one shot deviation in a game without testing. The one-shot deviation principle then implies that $\sigma \in \Sigma(0)$.¹¹

¹¹Note that the game with $K = 0$ is continuous at infinity, and so the one-shot deviation principle holds in such game.

Let $\Sigma^P(0) \subset \Sigma(0)$ denote the set of equilibria of the game without a regulator with the property that, for all $\sigma \in \Sigma^P(0)$, all firms expect to pass the test with probability 1 at every history. The arguments above imply that $\Sigma(K) \subset \Sigma^P(0)$ for all $K > \bar{K}$. In fact, the following stronger result holds:

Corollary 3. *For all $K > \bar{K}$, $\Sigma(K) = \Sigma^P(0)$.*

We highlight that testing at the individual firm level is crucial for Proposition 4. 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, and may lead to more collusive outcomes.

We note that although the set inclusion of Proposition 4 is weak, there are models of bidder-optimal collusion such that the safe tests described in Section 5.2 strictly reduce the surplus attainable by collusive bidders. We provide an explicit example Section 8. The same observation applies to Proposition 5 below.

Jointly safe tests. An analogue of Proposition 4 holds for jointly safe test, provided we impose Weak Renegotiation Proofness (see Farrell and Maskin, 1989).

Let $\Sigma_{RP}^+(K)$ denote the set of Weakly Renegotiation Proof equilibria of the game with joint testing and penalty K , such that all players get weakly positive expected discounted payoffs in period 0. The following result holds.

Proposition 5. *Assume the regulator runs a jointly safe test. For all $K > \bar{K}$, $\Sigma_{RP}^+(K) \subset \Sigma_{RP}^+(0)$.*

Beyond safe tests. Sufficiently punitive safe tests does not make collusion more profitable, and does not affect competitive industries. As a result, they should discourage the formation

of cartels. In contrast non-safe tests may increase cartel formation, either by reducing the payoffs of competition, or enabling the cartel to sustain new collusive equilibria.

However, a regulator with a strong prior may be willing to take some risks and implement tests that are not safe, but very likely to lead to good outcomes under his prior. In this sense, safe tests should be viewed as a starting point for regulation. Note that in our existing framework, regulators may express some of their views on the underlying environment when choosing the set of plausible economic environments Ω .

7 Empirical Evaluation

In this section, we explore the implications of our safe tests in real data. We argue that safe tests tend not to fail when applied to data from likely competitive industries, and tend to fail when applied to data from likely non-competitive industries.

7.1 A Case Study

We first illustrate the mechanics of inference using data from the city of Tsuchiura, represented in Figure 7.¹² Specifically, we show how three different deviations $\rho \in \{-.015, -.0005, .001\}$ affect our estimates of the share of competitive histories.

Figure 8 presents our estimates on the share of competitive histories as a function of correlation parameter k in constraint (EP). For these estimates and all the estimates that we present below, we set function $T(|H|)$ so that $2|\mathcal{M}|\exp(-\frac{1}{2}T(|H|)^2|H|) = 5\%$, so that our tests have a robust confidence level of 5% (recall that $\mathcal{M} = \{-\underline{n}, \dots, \bar{n}\}$). Again we emphasize that this leads to very conservative bounds which could be improved by using tighter asymptotic concentration result for sample demand in Lemma 1.

¹²Chassang and Ortner (forthcoming) studies the impact of a change in the auction format used by Tsuchiura that took place on October 29th 2009. We use data from auctions that took place before that date.

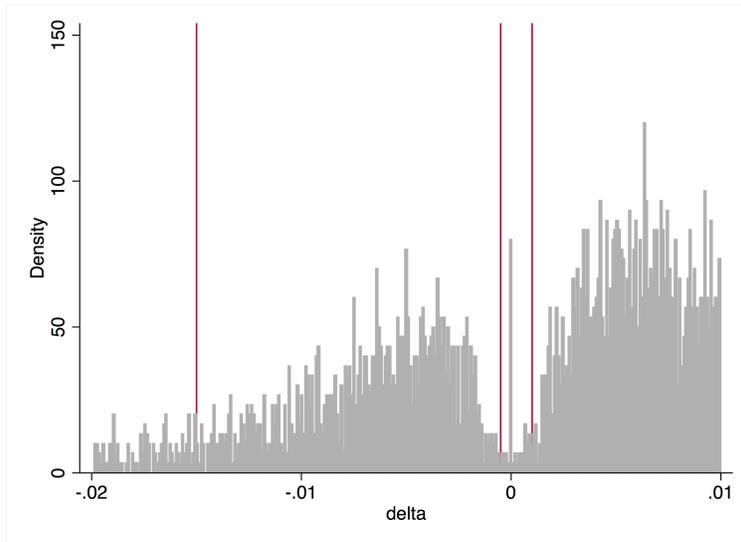


Figure 7: Distribution of Δ for the city of Tsuchiura, 2007–2009. Red vertical lines indicate deviation coefficients $\rho \in \{-.015, -.0005, .001\}$.

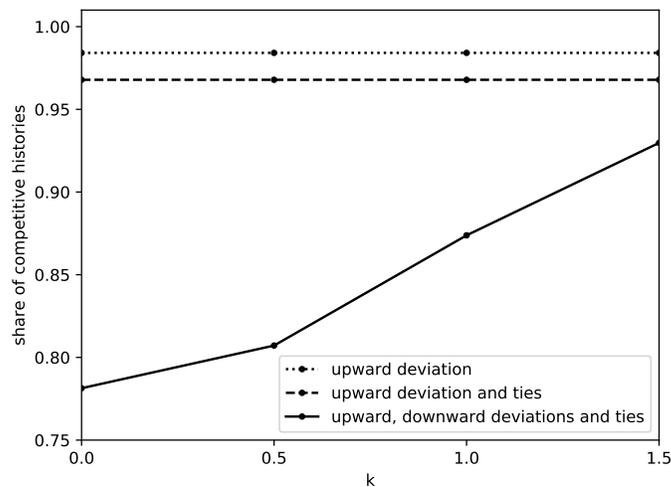


Figure 8: Estimated share of competitive auctions, city-level data.

We now delineate the mechanics of inference.

An upward deviation. We first consider a small upward deviation $\rho = .001$, corresponding to the analysis of Section 4. Because the mass of values of Δ falling between 0 and .001 is very small, this deviation hardly changes a bidder’s likelihood of winning an auction. As

a result this is a profitable deviation inconsistent with competition. If this low elasticity was estimated with great precision, Proposition 2 implies that a large number of histories would have to be deleted for the remainder to be competitive. Because the estimate is noisy, it suffices to delete a small share of histories for the data to be rationalizable within a robust 95% confidence interval. In addition, the upward deviation is least profitable (and so the data is best explained) when costs are low.

Adding a small downward deviation. Consider next adding a small downward deviation, $\rho = -.0005$. Because there is a surprisingly large mass of auctions such that the lowest and second lowest bids are tied or extremely close, this increases a bidders likelihood of winning by a non-vanishing amount, while reducing profits by a negligible amount, provided margins are not zero. The corresponding histories cannot be competitive.

Adding a medium-sized downward deviation. We now show that, under certain conditions, adding a medium sized downward deviation $\rho = -.015$ yields a tighter bound on the share of competitive histories. Intuitively this is because the aggregate counterfactual demand $\widehat{D}_\rho(H)$ increases by a large amount for relatively small deviations in bids: a 1.5% drop in prices leads to a 34% increase in the probability of winning the auction. If costs of production sufficiently below bids, this seems likely to be an attractive deviation.

We use d_ρ and D_ρ to denote subjective demand following a deviation $\rho = -.015$. Incentive compatibility (IC) implies that there must exist beliefs and costs satisfying

$$[(1 + \rho)b_h - c_h] d_{h,\rho} \leq [b_h - c_h] d_{h,0} \iff d_{h,\rho} - d_{h,0} \geq \frac{b_h}{c_h} [(1 + \rho)d_{h,\rho} - d_{h,0}]. \quad (7)$$

It turns out that in the absence of restrictions on beliefs or costs, IC constraint (7) is not binding. For all histories h_i at which bidder i won the auction, we set $c_{h_i} \leq b_{h_i}$ and $d_{h_i,\rho} = \mathbf{1}_{(1+\rho)b_{h_i,t} < \wedge b_{-i,h_i,t}} = 1$. For all histories h_i at which bidder i lost the auction, we set $d_{h_i,n} = \mathbf{1}_{(1+\rho)b_{h_i,t} < \wedge b_{-i,h_i,t}}$ and $c_{h_i} = b_{h_i}$. Under these beliefs and costs, constraint (7) is

satisfied at every history and $D_\rho(\omega_A, H(\omega_A)) = \widehat{D}_\rho(\omega_A, H(\omega_A))$.

Economic plausibility constraint (EP) rules out such extreme beliefs and costs. As Figure 8 shows, adding a mid-sized downward deviation $\rho = -.015$ leads to tighter estimates of the share of competitive histories, provided we also impose economic plausibility constraints. Alternatively, Figure 8 clarifies what assumptions about the environment are needed to claim with confidence that a significant share of histories are non-competitive.

7.2 Consistency between safe tests and proxies for collusion

We now show that our estimates on the share of competitive histories are consistent with different proxies of collusive behavior. For computational tractability, we focus on estimating the share of competitive histories only using the three deviations described above.

High vs. low bids. In Figure 4, we divide the histories in our national sample according to the bid level relative to the reserve price, and plot the distribution of Δ for the different subsamples. As the figure shows, the pattern of missing bids is more prevalent when we focus on histories at which bidders placed high bids. To the extent that missing bids are a marker of non-competitive behavior, Figure 4 suggests that histories at which firms placed lower bids are more likely to be competitive.

The left panel of Figure 9 plots our estimates for the share of competitive histories for the different sets of histories in Figure 4.¹³ The fraction of competitive histories is lower at histories at which bids are high relative to the reserve price, a finding that is consistent with the idea that collusion is more likely at periods at which bidders place higher bids. The right panel of Figure 9 plots our estimates for the share of competitive histories for two different subsamples of our city-level data: histories with bids below 0.9% of the reserve price, and histories with bids above 0.95% of reserve price. Again, the fraction of competitive histories

¹³Note that the set of histories with bids in some range $[x\%, y\%]$ of the reserve price is adapted to the bidders' information.

is lower at histories at which bids are high relative to the reserve price.

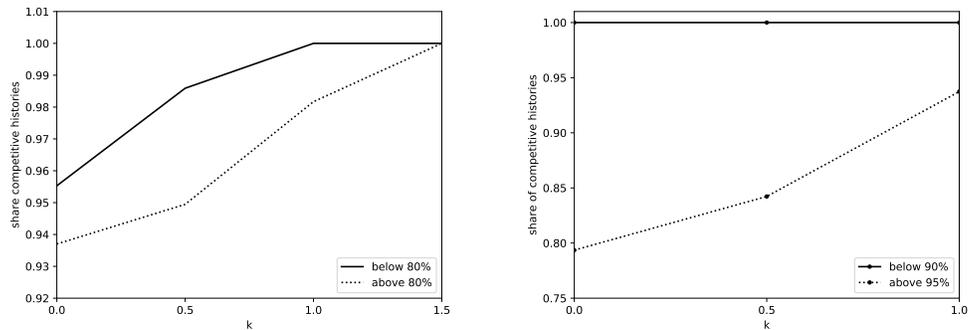


Figure 9: Estimated share of competitive auctions by bid level. Left-panel: national data; right-panel: city-level data.

Before and after prosecution. Figure 10 shows our bounds on the share of competitive histories for the four groups of firms that were investigated by the JFTC in Figure 5. Our estimates suggests non-competitive behavior in the before period across the four groups of firms. Moreover, with the exception of firms producing prestressed concrete, our estimates show essentially no collusion in the after period.

7.3 Zeroing-in on specific firms

We now apply our safe test to specific firms. We focus on the three firms that won most auctions in our national level data. Figure 11 plots our estimate on the share of competitive histories as a function of correlation parameter k in constraint (EP). As the figure shows, our tests have enough power that can potentially be used to test individual firms.

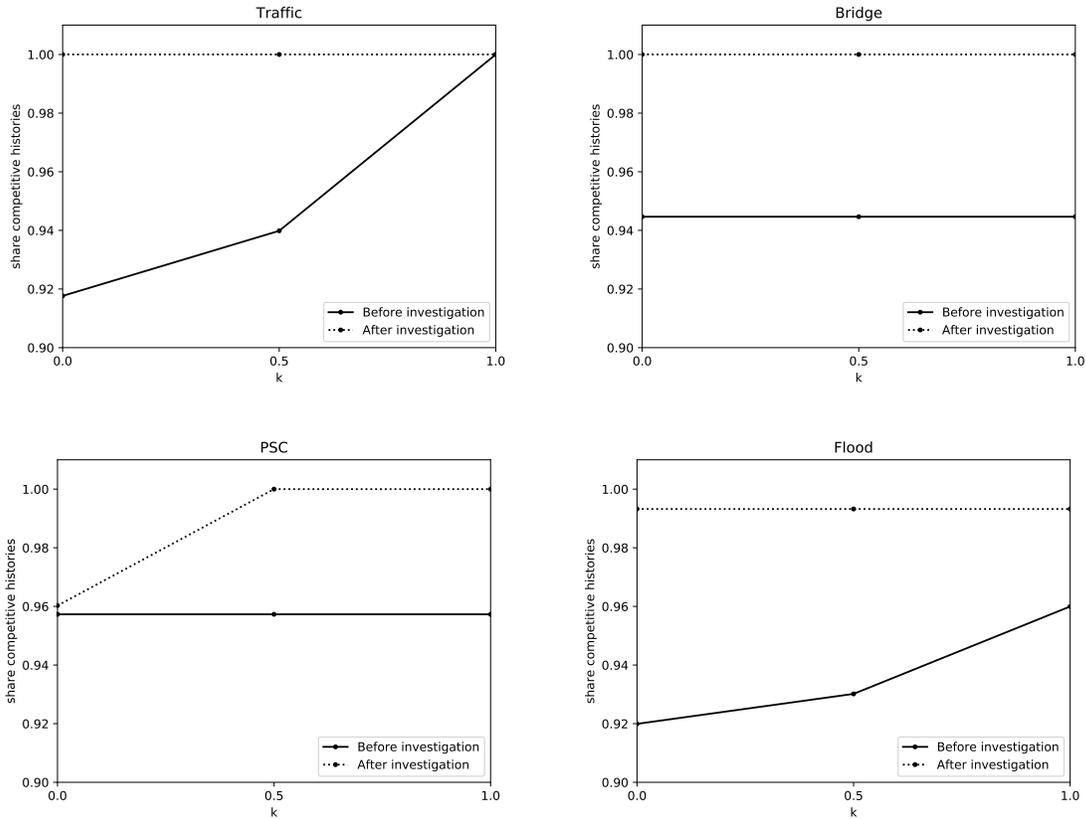


Figure 10: Estimated share of competitive auctions, before and after FTC investigation, national-level data.

8 Why Would Cartels Exhibit Missing Bids?

This paper focuses on the detection of non-competitive behavior in procurement auctions. As we argue in Section 6 the robust detection of non-competitive behavior allows for data-driven regulatory intervention that can only reduce the value of establishing a cartel. As Section 7 highlights, the corresponding tests exploit the missing bid pattern, as well as other aspects of the data: the large number of approximately tied bids, and the surprisingly large elasticity of counterfactual demand to the left of winning bids.

We conclude with an open ended discussion of why missing bids may be occurring in the first place. We argue that missing bids are not easily explained by standard models of collusion, and put forward two potential explanations. Along the way we establish that safe

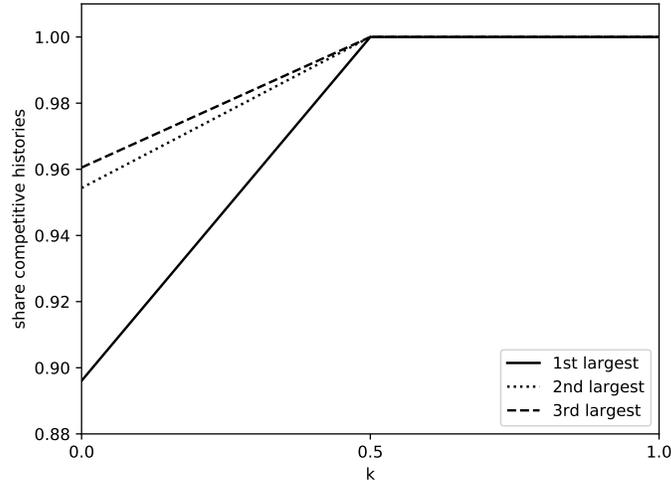


Figure 11: Estimated share of competitive auctions by bid level, four largest firms, national-level data.

tests can *strictly* reduce the surplus of cartels.

Standard models do not account for missing bids. In standard models of tacit collusion (see for instance Rotemberg and Saloner (1986), Athey and Bagwell (2001, 2008)), winning bids are typically common knowledge among bidders, and the cartel’s main concern is to incentivize losers not to undercut the winning bid. In contrast, the behavior of the designated winner is stage game optimal. This is achieved by having a losing bidder bid just above the designated winner. As a consequence, standard models of collusion would result in a point mass at $\Delta = 0$, rather than missing bids.

This is true in the framework of Chassang and Ortner (forthcoming), which builds on Levin (2003): costs are i.i.d. and common-knowledge; players have transferable utility and can make transfers; bidders play a pure strategy equilibrium. Chassang and Ortner (forthcoming) establishes that

- (i) All Pareto efficient equilibria generate the same discounted surplus \bar{V} on the equilibrium path and bidding behavior is stationary.

- (ii) Given costs $\mathbf{c} = (c_i)_{i \in N}$, the lowest cost bidder wins at the largest bid $b \leq r$ such that $\sum_i (1 - x_i)(b - c_i)^+ \leq \delta \bar{V}$
- (iii) Whenever $b < r$, another bidder bids arbitrarily close, and to the right of winning bid b .

This behavior is optimal for the cartel because otherwise some of the cartel's pledgeable surplus $\delta \bar{V}$ would have to be spent on keeping the winner from increasing her bid. It is more efficient to increase winning bids and use the pledgeable surplus to keep losing bidders from undercutting.

One implication of this is that there exist safe tests that strictly reduce the surplus of cartels. Specifically consider tests that rule out frequent approximately-tied winning bids:

$$\tau_{\kappa, \epsilon}(H) \equiv \begin{cases} 1 & \text{if } |\{a \in A_H, \text{ s.t. } b_{(2)} - b_{(1)} < \kappa\}| / |A_H| > \epsilon \\ 0 & \text{otherwise} \end{cases}$$

where $\kappa, \epsilon > 0$. Given $\epsilon > 0$, conditional on equilibrium markups being bounded away from 0, there exists κ small enough that $\tau_{\kappa, \epsilon}$ is a safe test under the assumption that competitive firms are privately informed about their costs. This test fails the optimal collusive equilibria described above.

Missing bids as robust coordination. One possible role for missing bids consists in facilitating coordination on a specific designated-winner. Being able to guarantee the identity of the winner may be important for a cartel for two reasons: (i) allocative efficiency; (ii) reducing the costs of dynamic incentive provision when utility is not transferable.

In this respect keeping winning bids isolated ensures that the designated winner does win the contract, even if bidders cannot precisely agree on exact bids ex ante, or if bids can be perturbed by small trembles (say a fat finger problem).

Missing bids as a side-effect of existing regulatory oversight. Interestingly, regulatory oversight itself may be at the origin of the missing bid pattern. If, along the lines of test $\tau_{\kappa, \epsilon}$, the regulator scrutinizes auctions in which tied bids occur, then collusive bidders sharing information about their intended bids may naturally wish to avoid this scrutiny, and ensure that there is a minimal distance between bids. Over many auctions, this leads to the detectable missing bids pattern highlighted in this paper.

Appendix

A Proofs

A.1 Proofs of Section 3

Proof of Lemma 1. Let H be a set of histories, and fix $\rho \in (-1, \infty)$. For each history $h_{i,t} = (h_t, z_{i,t}) \in H$, define

$$\begin{aligned} \varepsilon_{i,t} &\equiv \mathbb{E}_\sigma[\mathbf{1}_{\wedge b_{-i,t} > b_{i,t}(1+\rho)} | h_{i,t}] - \mathbf{1}_{\wedge b_{-i,t} > b_{i,t}(1+\rho)} \\ &= \text{prob}_\sigma(\wedge \mathbf{b}_{-i,t} > b_{i,t}(1+\rho) | h_{i,t}) - \mathbf{1}_{\wedge b_{-i,t} > b_{i,t}(1+\rho)}. \end{aligned}$$

Note that $\widehat{D}(\rho|H) - \overline{D}(\rho|H) = \frac{1}{|H|} \sum_{h_{i,t} \in H} \varepsilon_{i,t}$.

Note further that, by the law of iterated expectations, for all histories $h_{j,t-s} \in H$ with $s \geq 0$, $\mathbb{E}_\sigma[\varepsilon_{i,t} | h_{j,t-s}] = \mathbb{E}_\sigma[\mathbb{E}_\sigma[\mathbf{1}_{\wedge b_{-i,t} > b_{i,t}(1+\rho)} | h_t, z_{i,t}] - \mathbf{1}_{\wedge b_{-i,t} > b_{i,t}(1+\rho)} | h_{t-s}, z_{j,t-s}] = 0$.¹⁴

Number the histories in H as $1, \dots, |H|$ such that, for any pair of histories $k = (h_s, z_{i,s}) \in H$ and $k' = (h_{s'}, z_{j,s'}) \in H$ with $k' > k$, $s' \geq s$. For each history $k = (h_t, z_{i,t})$, let $\varepsilon_k = \varepsilon_{i,t}$, so that

$$\widehat{D}(\rho|H) - \overline{D}(\rho|H) = \frac{1}{|H|} \sum_{k=1}^{|H|} \varepsilon_k.$$

¹⁴This holds since, in a perfect public Bayesian equilibrium, bidders' strategies at any time t depend solely on the public history and on their private information at time t .

Note that, for all $\hat{k} \leq |H|$, $S_{\hat{k}} \equiv \sum_{k=1}^{\hat{k}} \varepsilon_k$ is a Martingale, with increments ε_k whose absolute value is bounded above by 1. By the Azuma-Hoeffding Inequality, for every $\alpha > 0$, $\text{prob}(|S_{|H|}| \geq |H|\alpha) \leq 2 \exp\{-\alpha^2|H|/2\}$. Therefore, with probability 1, $\frac{1}{|H|}S_{|H|} = \hat{D}(\rho|H) - \bar{D}(\rho|H)$ converges to zero as $|H| \rightarrow \infty$. ■

A.2 Proofs of Section 4

Proof of Proposition 3. By Lemma 1, under the true environment $\omega_A \in \Omega$,

$$\text{prob}(|\hat{D}_n(H(\omega_A)) - D_n(\omega_A, H(\omega_A))| \geq T(|H(\omega_A)|)) \leq 2 \exp(-T(|H(\omega_A)|)^2|H(\omega_A)|/2)$$

for each deviation n . It then follows that

$$\text{prob}(\forall n, |\hat{D}_n(H(\omega_A)) - D_n(\omega_A, H(\omega_A))| \geq T(|H(\omega_A)|)) \leq 2(\bar{n} + \underline{n} + 1) \exp(-T(|H(\omega_A)|)^2|H(\omega_A)|/2).$$

This implies that, with probability at least $1 - 2(\bar{n} + \underline{n} + 1) \exp(-T(|H(\omega_A)|)^2|H(\omega_A)|/2)$, the constraints in Program (P) are satisfied when we set the environment equal to ω_A . Hence, with probability at least $1 - 2(\bar{n} + \underline{n} + 1) \exp(-T(|H(\omega_A)|)^2|H(\omega_A)|/2)$, $\hat{U} \geq U(\omega_A)$. ■

Proof of Corollary 2. Suppose the true environment $\omega_A \in \Omega$ is such that the industry (or the firms who placed bids in histories H) is competitive. Then $H(\omega_A) = H$, and so the true share of competitive histories under ω_A is $s_{\text{comp}} = 1 \geq s_0$. By Corollary 1, with probability at least $1 - 2(\bar{n} + \underline{n} + 1) \exp(-T(|H|)^2|H|/2)$, $\hat{U} \geq s_{\text{comp}} = 1 \geq s_0$. Since $2(\bar{n} + \underline{n} + 1) \exp(-T(|H|)^2|H|/2) \rightarrow 0$ as $|H| \rightarrow \infty$, firms in this industry pass test τ with probability approaching 1 as $|H| \rightarrow \infty$. ■

A.3 Proofs of Section 6

Proof of Proposition 4. We start by showing that, when penalty K is sufficiently large, any $\sigma \in \Sigma(K)$ has the property that all firms pass the test with probability 1, 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 equilibrium payoff cannot be lower than 0 at any history.

Let $\bar{K} = \frac{1}{1-\delta}r = \frac{1}{1-\delta}$ (recall that the reserve price r is normalized to 1), and suppose that $K > \bar{K}$. Towards a contradiction, suppose there exist $\sigma \in \Sigma(K)$ and a public history h_t (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 history h_s with $s \geq 0$ such that, at the concatenated history $h_t \sqcup h_s$, firm i expects to fail the test with probability at least $\frac{\bar{K}}{K} + \epsilon < 1$. At history $h_t \sqcup h_s$, firm i 's continuation payoff is bounded above by $\frac{1}{1-\delta} - \left(\frac{\bar{K}}{K} + \epsilon\right)K = -\epsilon K < 0$, a contradiction.

For any strategy profile $\hat{\sigma}$ and any history $h_{i,t} = (h_t, z_{i,t})$, let $V_i(\hat{\sigma}, h_{i,t}) = \mathbb{E}_{\hat{\sigma}}[\sum_{s \geq t} \delta^s u_{i,s} | h_{i,t}]$ denote firm i 's continuation payoff excluding penalties under σ at history $h_{i,t}$. Firm i 's total payoff from under strategy profile $\hat{\sigma}$ given history $h_{i,t}$ is $V_i(\hat{\sigma}, h_{i,t}) - \mathbb{E}_{\hat{\sigma}}[\tau_i | h_{i,t}]K$.

Suppose $K > \bar{K}$ and fix $\sigma \in \Sigma(K)$. Since σ is an equilibrium, there cannot be profitable deviations; in particular, there cannot be profitable one shot deviations:¹⁵ for every $i \in N$, every history $h_{i,\tau}$, 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 $h_{i,t} \neq h_{i,\tau}$,

$$\begin{aligned} V_i((\tilde{\sigma}_i, \sigma_{-i}), h_{i,\tau}) - \mathbb{E}_{(\tilde{\sigma}_i, \sigma_{-i})}[\tau_i | h_{i,\tau}]K &\leq V_i(\sigma, h_{i,\tau}) - \mathbb{E}_{\sigma}[\tau_i | h_{i,\tau}]K \\ \iff V_i((\tilde{\sigma}_i, \sigma_{-i}), h_{i,\tau}) &\leq V_i(\sigma, h_{i,\tau}), \end{aligned} \tag{8}$$

¹⁵Note that we are not using the one-shot deviation principle here (which may not hold since the game is not continuous at infinity when $K > 0$); we are only using the fact that, in any equilibrium, no player can have a profitable deviation.

where the second line in (8) follows since, under equilibrium $\sigma \in \Sigma(K)$, all firms pass the test with probability 1 at every history. By the second line in (8), in the game with $K = 0$ (i.e., no regulator) no firm has a profitable one shot deviation under strategy profile σ . Hence, by the one-shot deviation principle $\sigma \in \Sigma(0)$.¹⁶ ■

The following Lemma establishes a weaker version of the one-shot revelation principle for the game with a regulator.

Lemma A.1. *Let σ be a strategy profile with the property that all firms pass the test 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 $\frac{\delta^T}{1-\delta} \times r = \frac{\delta^T}{1-\delta} < \epsilon/2$. 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 $r = 1$, 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

¹⁶Note that the game with $K = 0$ is continuous at infinity, and so the one-shot deviation principle holds.

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 Corollary 3. Fix $K > \bar{K}$. The proof of Proposition 4 shows that, in all equilibria in $\Sigma(K)$, all firms pass the test with probability 1 at every history. Since $\Sigma(K) \subset \Sigma(0)$, it follows that $\Sigma(K) \subset \Sigma^P(0)$.

We now show that $\Sigma^P(0) \subset \Sigma(K)$. 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 $h_{i,\tau}$, 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 $h_{i,t} \neq h_{i,\tau}$,

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

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 $\sigma \in \Sigma(K)$. ■

Proof of Proposition 5. The proof is similar to the proof of Proposition 4. 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 history h_t , 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 h_t at least one firm's payoff is larger than 0.

Recall that $\bar{K} \equiv \frac{1}{1-\delta}$. Suppose $K > \bar{K}$, and fix $\sigma \in \Sigma_{RP}^+(K)$. Towards a contradiction, suppose that there exists a history h_t (on or off path) such that, at this history, firms expect to fail the test with strictly positive probability. Then, for every $\epsilon > 0$ small, there must exist a history h_s with $s \geq 0$ such that, at the concatenated history $h_t \sqcup h_s$, firms expect to fail the test with probability at least $\frac{\bar{K}}{K} + \epsilon < 1$. At history $h_t \sqcup h_s$, the continuation payoff of each firm is bounded above by $\frac{1}{1-\delta} - \left(\frac{\bar{K}}{K} + \epsilon\right)K < -\epsilon K < 0$, 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 $h_{i,\tau}$, 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 $h_{i,t} \neq h_{i,\tau}$,

$$\begin{aligned} V_i((\tilde{\sigma}_i, \sigma_{-i}), h_{i,\tau}) - \mathbb{E}_{(\tilde{\sigma}_i, \sigma_{-i})}[\tau_i | h_{i,\tau}]K &\leq V_i(\sigma, h_{i,\tau}) - \mathbb{E}_\sigma[\tau_i | h_{i,\tau}]K \\ \iff V_i((\tilde{\sigma}_i, \sigma_{-i}), h_{i,\tau}) &\leq V_i(\sigma, h_{i,\tau}), \end{aligned} \tag{9}$$

where the second line in (9) follows since, under equilibrium σ , firms pass the test with probability 1 at every history. By the second line in (9), 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,

$\sigma \in \Sigma_{RP}^+(K)$. ■

References

- ABDULKADIROGLU, A. AND K.-S. CHUNG (2003): “Auction design with tacit collusion,” *Columbia and Northwestern University Working Paper*.
- AFRIAT, S. N. (1967): “The construction of utility functions from expenditure data,” *International economic review*, 8, 67–77.
- ASKER, J. (2010): “A study of the internal organization of a bidding cartel,” *The American Economic Review*, 724–762.
- ATHEY, S. AND K. BAGWELL (2001): “Optimal collusion with private information,” *RAND Journal of Economics*, 32, 428–465.
- (2008): “Collusion with persistent cost shocks,” *Econometrica*, 76, 493–540.
- BAJARI, P. AND L. YE (2003): “Deciding between competition and collusion,” *Review of Economics and Statistics*, 85, 971–989.
- BERGEMANN, D., B. BROOKS, AND S. MORRIS (2017): “First-Price Auctions With General Information Structures: Implications for Bidding and Revenue,” *Econometrica*, 85, 107–143.
- BERGEMANN, D. AND S. MORRIS (2013): “Robust predictions in games with incomplete information,” *Econometrica*, 81, 1251–1308.
- BESANKO, D. AND D. F. SPULBER (1989): “Antitrust enforcement under asymmetric information,” *The Economic Journal*, 99, 408–425.

- CARVAJAL, A., R. DEB, J. FENSKE, AND J. K.-H. QUAH (2013): “Revealed preference tests of the Cournot model,” *Econometrica*, 81, 2351–2379.
- CHAMBERS, C. P. AND F. ECHENIQUE (2016): *Revealed preference theory*, vol. 56, Cambridge University Press.
- CHASSANG, S. AND J. ORTNER (forthcoming): “Collusion in Auctions with Constrained Bids: Theory and Evidence from Public Procurement,” *Journal of Political Economy*.
- CHE, Y.-K., D. CONDORELLI, AND J. KIM (2010): “Weak Cartels and Collusion-Proof Auctions,” Tech. rep., Working paper, Columbia University.
- CHE, Y.-K. AND J. KIM (2006): “Robustly Collusion-Proof Implementation,” *Econometrica*, 74, 1063–1107.
- CLARK, R. AND J.-F. HOUDE (2013): “Collusion with asymmetric retailers: Evidence from a gasoline price-fixing case,” *American Economic Journal: Microeconomics*, 5, 97–123.
- 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.
- ECHENIQUE, F., S. LEE, AND M. SHUM (2011): “The money pump as a measure of revealed preference violations,” *Journal of Political Economy*, 119, 1201–1223.
- 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.

- GREEN, E. J. AND R. H. PORTER (1984): “Noncooperative collusion under imperfect price information,” *Econometrica: Journal of the Econometric Society*, 87–100.
- 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.
- 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.
- LEVIN, J. (2003): “Relational incentive contracts,” *The American Economic Review*, 93, 835–857.
- MASKIN, E. AND J. TIROLE (2001): “Markov perfect equilibrium: I. Observable actions,” *Journal of Economic Theory*, 100, 191–219.
- PESENDORFER, M. (2000): “A study of collusion in first-price auctions,” *The Review of Economic Studies*, 67, 381–411.
- PORTER, R. H. (1983): “A study of cartel stability: the Joint Executive Committee, 1880–1886,” *The Bell Journal of Economics*, 301–314.
- PORTER, R. H. AND J. D. ZONA (1993): “Detection of Bid Rigging in Procurement Auctions,” *Journal of Political Economy*, 101, 518–38.
- (1999): “Ohio School Milk Markets: An Analysis of Bidding,” *RAND Journal of Economics*, 30, 263–288.

ROTEMBERG, J. J. AND G. SALONER (1986): “A supergame-theoretic model of price wars during booms,” *The American Economic Review*, 390–407.

TÓTH, B., M. FAZEKAS, Á. CZIBIK, AND I. J. TÓTH (2014): “Toolkit for detecting collusive bidding in public procurement,” *With examples from Hungary*.

VARIAN, H. R. (1990): “Goodness-of-fit in optimizing models,” *Journal of Econometrics*, 46, 125–140.