Data Driven Regulation: Theory and Application to Missing Bids

Sylvain Chassang — New York University

Kei Kawai — U.C. Berkeley

Jun Nakabayashi — Kindai University

Juan Ortner*†

Boston University

September 18, 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 robust 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 a weakened version of equilibrium conditions, and show that safe tests make life weakly harder for cartels by constraining the set of continuation values that can be used 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, the University of Illinois at Urbana-Champaign.

1 Introduction

One of the key functions of antitrust authorities is to detect and punish collusion. While the number of cartels that are actually prosecuted is relatively small, detection and punishment of cartels are nonetheless important for deterrence. The ability of antitrust agencies to identify and prosecute cartels can also incentivize members of existing cartels to apply for leniency programs. Hence, successful detection and punshiment of cartels deter cartel formation and complement leniency programs.

In the absence of concrete leads, screening devices that flag suspicious firm conduct may be useful for regulators as a first step in identifying collusion. In this paper, we first document a bidding pattern from procurement auctions in Japan in which the density of the bid distribution just above the winning bid is very low. The pattern that we document implies that winning bids tend to be isolated, an observation that has already been made elsewhere, and even been used as part of screening programs in some countries.[1] We show 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

[1] See Imhof et al. (2016) and Tóth et al. (2014).

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 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.

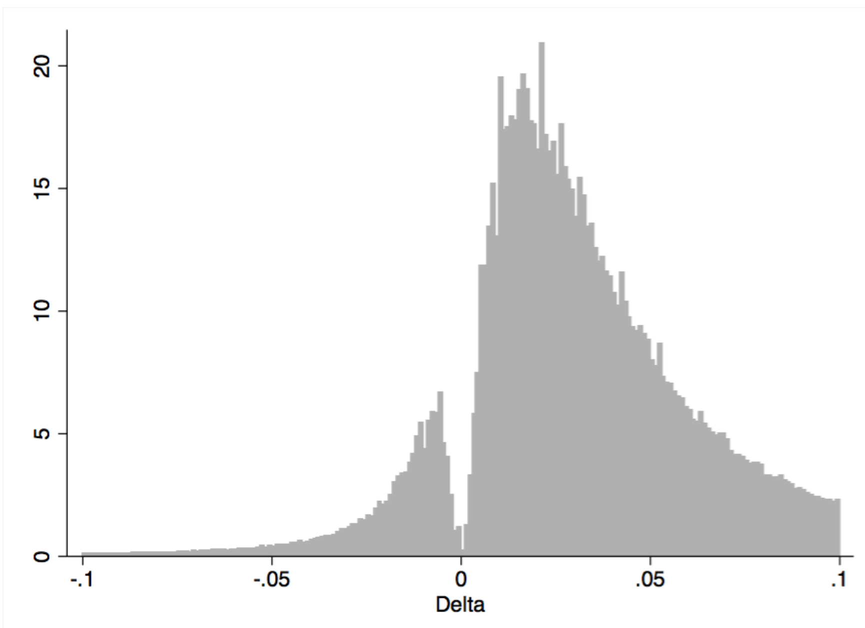

Figure 1: Distribution of bid-differences $\Delta \equiv \frac{\text{own bid}-\min(\text{other bids})}{\text{reserve}}$ over (bidder, auction) pairs.

We analyze our data within a fairly general model of repeated play in first-price procurement 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.[2] 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 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

[2]We discuss this assumption at length. [XXX sense in which our results extend]

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 robust to incomplete information, 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.[3] 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.[4] A common observation from these papers is that data driven regulation may backfire, allowing a cartel to sustain higher equilibrium prices. We add to these papers by introducing safe tests, and by showing that regulation based on such tests necessarily restricts the set of equilibrium

[3]See Harrington (2008) for a recent survey.

[4]Other 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 environment. 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.[5] 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.[6] 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

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

[6]See 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 optimality conditions from a weakened version of 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 invlove 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.[7] The mass of missing bids around a difference of 0 is

[7]Note 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.

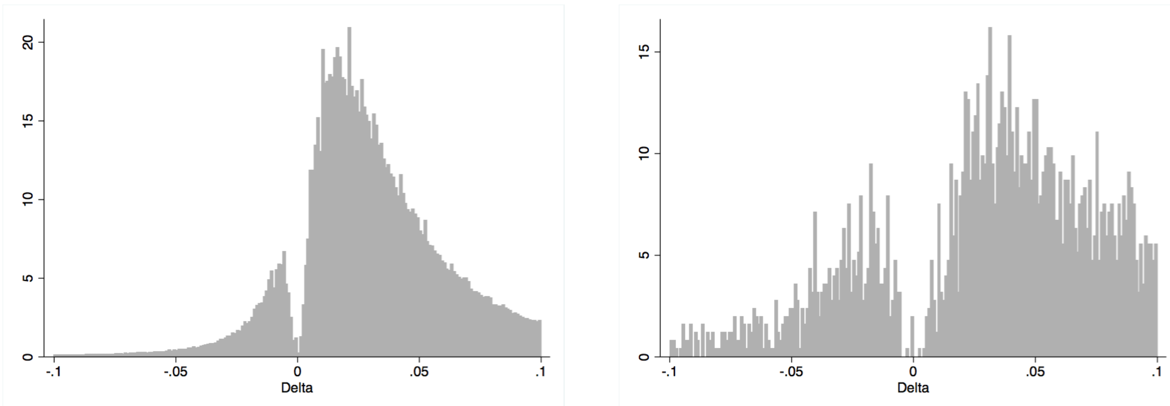

(a) all firms (b) single large firm

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

Our second dataset, analyzed in Chassang and Ortner (forthcoming), consists of roughly 1,500 auctions held between 2007 and 2014 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 presents plots the distribution of Δ for auctions held in Tsuchiura. Again, we see a significant mass of missing bids around zero.[8]

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, we now present some reduced form evidence relating missing bids to different proxies of collusion.

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

[8]Imhof 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).

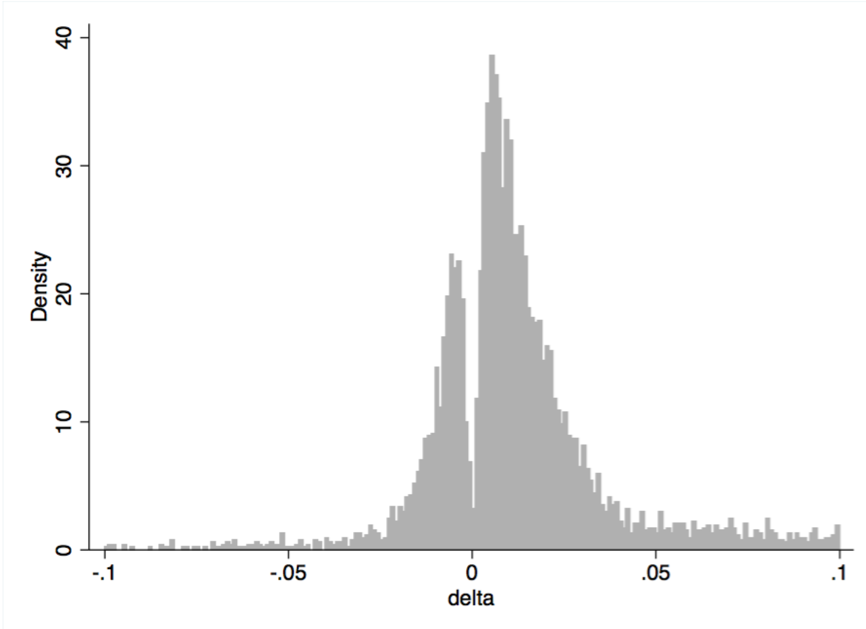

Figure 3: Distribution of bid-difference Δ – city data.

distribution of $\Delta_{i,a} = \frac{b_{i,a} - \min b_{-i,a}}{r}$ for bids 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.[9] 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.

[9]See 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).

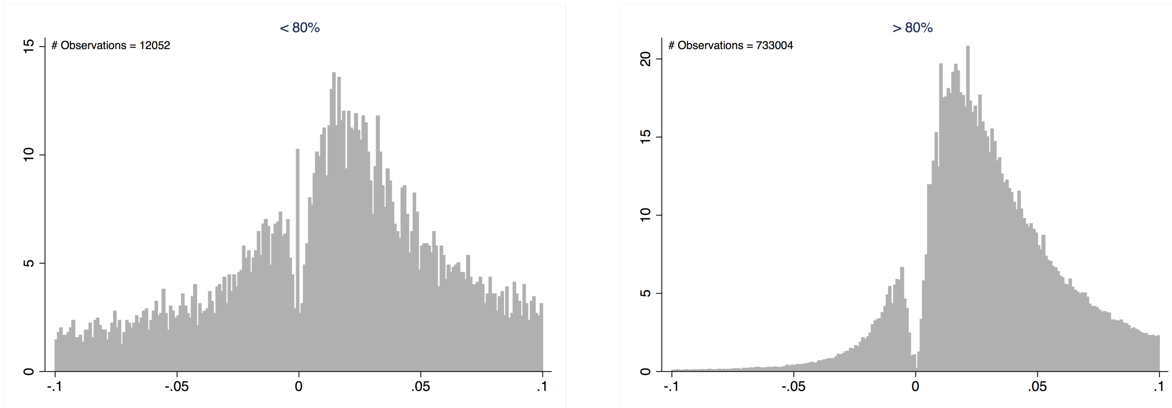

(a) Below 80%

(b) Above 80%

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 XXX shows that the bidding patterns persist even if we focus on auctions whose prices were not renegotiated up. Second, the way renegotiation works in these auctions greatly reduces firms' incentives to bid aggressively with the hope of getting a higher price: when renegotiation does occur, renegotiated prices depend on the level of the initial bid.[10]

[10]Indeed, if the project is estimated to cost more than initially thought, the renegotiated price is given by $\frac{\text{initial bid}}{\text{reserve price}} \times$ (new cost estimate).

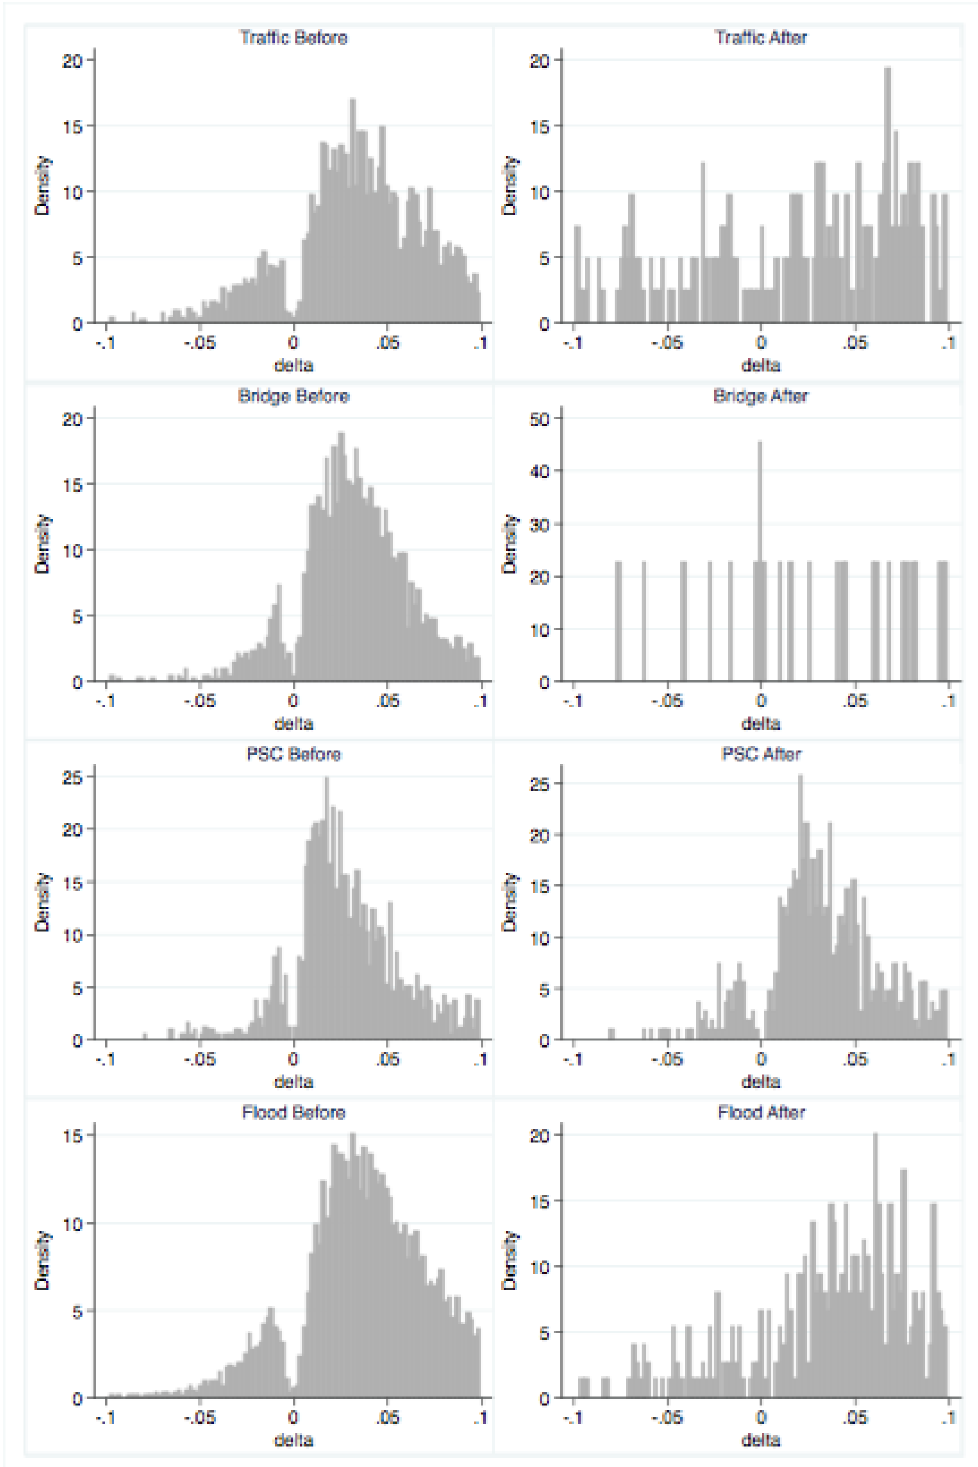

Figure 5: Distribution of bid-difference Δ – cartel cases in national data, before and after JFTC investigation.

Our objectives in this paper are: 1) formalize why this pattern is suspicious; 2) delineate what it implies about bidding behavior and the competitiveness of auctions in our sample; 3) formulate a theory of regulation based on safe tests; and 4) 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.[11] The sets 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

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

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.

Costs. We allow for costs that are serially correlated over time, and that may be correlated across firms within each auction. Denoting by $\langle ., . \rangle$ 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 bidders at the end of period t; we assume that θ_t follows a Markov chain;
- $\varepsilon_{i,t}$ 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.

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

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

3.2 Repeated Interaction and Solution Concept

Transfers. Bidders are able to make positive transfers from one to the other at the end of each period. A transfer from i to j is denoted by $T_{i \to j,t} \geq 0$. Transfers are costly, and we denote by $K\left(\sum_{j \neq i} T_{i \to j,t}\right)$ the cost to player i of the transfers she makes. We assume that K is positive, increasing and convex. Altogether, flow realized payoffs to player i in period t take the form

$$u_{i,t} = \pi_{i,t} + \sum_{j \neq i} T_{j \to i,t} - K\left(\sum_{j \neq i} T_{i \to j,t}\right).$$

Solution Concepts. The public history h_t at period t takes the form

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

where $\mathbf{T}_s$ are the transfers made in period s. 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 $i \in N$,

$$\sigma_i : h_t \mapsto (b_{i,t}(z_{i,t}), (T_{i \to j,t}(z_{i,t}, \mathbf{b}_t))_{j \neq i}),$$

where bids $b_{i,t}(z_{i,t}) \in \Delta([0,r])$ and transfers $(T_{i\to j,t}(z_{i,t}, \mathbf{b}_t))_{j\neq i} \in \Delta(\mathbb{R}^{n-1})$ 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, μ) 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, μ) 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 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 $h \in H_\infty$.

[XXX: connection with expert testing literature Foster and Vohra (1998, 1999), Olszewski and Sandroni (2008) — even though the cartel would pass the test in equilibrium, it reduces the profits from running a cartel]

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 obtain aggregates of counterfactual demand that can be estimated from data, even though the players' residual demands can vary with the history.

[XXX include elements of intuitive discussion with Jeff Ely]

[XXX Foreshadow empirical results in Section 7]

4.1 Counterfactual demand

Fix a perfect public Bayesian equilibrium (σ, μ). For all public 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

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

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

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

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

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

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.

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

[XXX discuss non asymptotic results, and practical implementation as test]

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 can be viewed as a weakening of the equilibrium requirement that beliefs be correct. It may fail in settings with sufficiently strong non-common priors.

The ability to legitimately vary the conditioning set H lets us explore the competitiveness of auctions in particular subsettings of interest.

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, μ) *be a competitive equilibrium. Then,*

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

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

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,

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

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 as well in these histories.

Proof. Consider a competitive equilibrium (σ, μ). Let u_i denote the flow payoff of player i, and let $V(h_{i,t}) \equiv \mathbb{E}_{\sigma,\mu}\left(\sum_{s\geq t} \delta^{s-t} u_{i,s} \big| h_{i,t}\right)$ 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 $b' > b$,

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

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

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

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 $b' \to b$. 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. 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,

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

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

Proposition 1 proposes a simple test of whether a 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

[XXX: emphasize the point that many histories makes it more difficult to explain the pattern through mistakes alone]

[XXX: Should we kill this section?]

Fix a perfect public Bayesian equilibrium (σ, μ) and a finite set of histories H. Let $H^{\mathsf{comp}} \subset H$ be the set of competitive histories in H, and let $H^{\mathsf{coll}} = H \backslash H^{\mathsf{comp}}$. Define $s_{\mathsf{comp}} \equiv \frac{|H^{\mathsf{comp}}|}{|H|}$ 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

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

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*

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

Proof. For any $\rho > 0$,

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

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 $\overline{R}(\rho|H^{\textsf{comp}}) - \overline{R}(0|H^{\textsf{comp}}) \leq 0$. Second,

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

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

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

Note that $\widehat{R}(\rho|H)$ is the sample analog of counterfactual revenue. A result identical to Theorem 1 establishes that $\widehat{R}(\rho|H)$ is an unbiased estimate of $\overline{R}(\rho|H)$, whenever set H is adapted. We have the following corollary to Proposition 2.

Corollary 1. *Fix a set of histories H and a scalar $\rho^* > 0$, and suppose that $s_{\textsf{comp}} \geq 1 - \kappa$ for some $\kappa > 0$. Then, there exists constants $\alpha > 0$ and $\beta > 0$ such that, with probability at* *least* $1-\beta\exp(-\alpha|H|)$,

$$\forall\rho\geq\rho^*,\qquad \frac{\widehat{R}(\rho|H)-\widehat{R}(0|H)}{\rho}\leq 2\kappa.$$

Corollary 1 gives the following statistical test with significance level $1-\beta\exp(-\alpha|H|)$. Let the null hypothesis be $H_0 = s_{\mathsf{comp}} \geq 1-\kappa$ for some $\kappa>0$, and let the alternative hypothesis be $H_1 = s_{\mathsf{comp}} < 1-\kappa$. We reject the null hypothesis if we can find $\rho\geq\rho^*$ such that $\frac{1}{\rho}[\widehat{R}(\rho|H)-\widehat{R}(0|H)] > 2\kappa$.

[TODO DATA: maybe add an example]

5 A General Class of Safe Tests

Take as given an adapted set of histories H corresponding to a set of auctions A. Take as given scalars $(\rho_n)_{n\in\mathcal{N}}$, with $\rho_n\in(-1,\infty)$ for all $n\in\mathcal{N}=\{-\underline{N},\cdots,\overline{N}\}$, $\rho_0=0$ and $\rho_n<\rho_{n'}$ for all $n'>n$. For each history $h_{i,t}\in H$, let $d_{h_{i,t},n}=D_i((1+\rho_n)b_{h_{i,t}}|h_{i,t})$. That is, $(d_{h_{i,t},n})_{n\in\mathcal{N}}$ is firm i's subjective counterfactual demand at history $h_{i,t}$. For any auction a and associated histories $h\in a$, we denote by $\omega_a=(d_{n,h},c_h,s_a)$ an environment at a, where s_a is a public signal observed by all firms; i.e, a candidate payoff and belief structure at a. We let $\omega_A=(\omega_a)_{a\in A}$.

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

For each deviation n, environment $\omega_A=(\omega_a)_{a\in A}$ and adapted set of histories $\hat{H}\subset H$ define

$$D_n(\omega_A,\hat{H})\equiv\frac{1}{|\hat{H}|}\sum_{h_{i,t}\in\hat{H}}d_{h_{i,t},n}\quad\text{and}\quad\widehat{D}_n(\hat{H})\equiv\frac{1}{|\hat{H}|}\sum_{h_{i,t}\in\hat{H}}\mathbf{1}_{(1+\rho_n)b_{h_{i,t}}<\wedge b_{-i,h_{i,t}}}.$$

We encode our inference problem as a constrained minimization problem. Specifically, given an objective function $u : \omega_a \mapsto U(\omega_a) \in \mathbb{R}$, and environments $\omega_A = (\omega_a)_{a\in A} \in \Omega$ let

$$U(\omega_A) = \sum_{a\in A} u(\omega_a).$$

For each environment ω_A, we let $H(\omega_A) \subset H$ be the adapted set of histories of interest under environment ω_A. For instance, $H(\omega_A)$ could be the set of competitive histories under environment ω_A.

For any function $T : \mathbb{N} \to \mathbb{R}^+$, let $\widehat{U}$ denote the solution to the following constrained optimization problem:

$$\widehat{U} = \max_{\omega_A} U(\omega_A) \tag{6}$$
$$\text{s.t. } \forall n, \quad D_n(\omega_A, H(\omega_A)) \in \left[\widehat{D}_n(\omega_A, H(\omega_A)) - T(|H(\omega_A)|), \widehat{D}_n(\omega_A, H(\omega_A)) + T(|H(\omega_A)|)\right].$$

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

By using different objective functions, we can solve a variety of inference objectives.

5.1 Maximum Share of Non-Competitive Histories and Auctions

We now use the results in Proposition 3 to provide estimates on the share of competitive histories in H and the share of competitive auctions in A; i.e., the set of auctions with the property that is common knowledge among bidders that play is competitive.

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 the following conditions:

Feasibility. Costs and beliefs must be feasible, satisfying

$$c_h \in [0, b_h]; \quad \forall n, \;\; d_{h,n} \in [0,1]; \quad \forall n, n' > n, \;\; d_{h,n} \geq d_{h,n'}. \tag{7}$$

Individual optimality. Bidding b_h must be optimal, given cost and subjective believes:

$$\forall n, \quad [(1+\rho_n)b_h - c_h]\, d_{h,n} \leq [(1+\rho_0)b_h - c_h] d_{h,0} \tag{8}$$

Economic plausibility. In addition to feasibility and incentive compatibility, one may be able to impose plausible ad hoc constraints on the bidder's economic environment at each history h. We focus on two intuitive constraints on the bidder's costs c_h and interim beliefs $d_h = (d_{h,n})$:

$$\frac{b_h}{c_h} \leq 1 + m \tag{9}$$

and

$$\forall n, \quad \left| \log \frac{d_{h,n}}{1-d_{h,n}} - \log \frac{D_n}{1-D_n} \right| \leq k \tag{10}$$

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.[12]

Correspondingly, given an environment ω_a at auction a, we can define the objective function

$$u(\omega_a) = \frac{1}{|H|} \sum_{h \in H} \mathbf{1}_{(d_h, c_h) \text{ satisfy (7), (8), (9), (10)}}$$

For each ω_A, let

$$H(\omega_A) = \arg \max_{H^{\mathsf{comp}} \subset H} |H^{\mathsf{comp}}|$$
$$\text{s.t. } \forall h \in H^{\mathsf{comp}} \quad \exists (d_h, c_h) \text{ satisfying (7), (8), (9), (10).}$$

[12]To see why, that that $\log \frac{d_{h,n}}{1-d_{h,n}} = \log \frac{\mathsf{prob}(Z|h)}{\mathsf{prob}(\neg Z|h)}$ for Z the event that $\wedge \mathbf{b}_{-i} > (1+\rho_n)b_h$. 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.

be the largest number of histories in H that can be rationalized as competitive. Program (6) then becomes

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

$\widehat{U}$ provides an upper bound to the share of competitive histories in H.

Alternatively, we can define the objective function to be

$$u(\omega_a) = \frac{1}{|A|} \mathbf{1}_{\forall h \in a,\ d_h, c_h \text{ satisfy (7), (8), (9), (10)}}.$$

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

$$H(\omega_A) = \arg \max_{H^{\mathsf{comp}} \subset H} |A_{H^{\mathsf{comp}}}|$$
$$\text{s.t. } \forall a \in A_{H^{\mathsf{comp}}}, \forall h \in a, \quad \exists (d_h, c_h) \text{ satisfying } (7),\ (8),\ (9),\ (10).$$

Program (6) then becomes

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

$\widehat{U}$ provides an upper bound to the fraction of competitive auctions corresponding to histories in H.

We have the following Corollary to Proposition 3.

Corollary 2. *Suppose that the true environment ω_A satisfies* (9) *and* (10)*, and that the true share of competitive histories (true share of competitive auctions) is $s_{\mathsf{comp}} \in (0, 1]$. Then,* *with probability at least* $1 - 2|\mathcal{N}|\exp(-\frac{1}{2}T(|H(\omega_A)|)^2|H(\omega_A)|)$, $\hat{U} \geq s_{\mathsf{comp}}$.

Fix $s_0 \in (0, 1]$. The null hypothesis is that the fraction s_{comp} of competitive histories (or competitive auctions) satisfies $s_{\mathsf{comp}} \geq s_0$, and the alternative hypothesis is $s_{\mathsf{comp}} < s_0$. Let τ^{safe} be a test such that $\tau^{\mathsf{safe}}(H) = 0$ if $\hat{U} \geq s_0$ and $\tau^{\mathsf{safe}}(H) = 1$ otherwise. We have:

Corollary 3. *[safe tests] Suppose function* $T(\cdot)$ *satisfies* $\exp(-\frac{1}{2}T(|H|)^2|H|) \to 0$ *as* $|H| \to \infty$. *Then,* τ^{safe} *is a safe test.*

We make two observations. First, by varying the set H of adapted histories, we can make test τ^{safe} be safe for a given firm, or for a given industry. Indeed, by taking H to be the set of histories corresponding to all the bids placed by a given firm i, test τ^{safe} is safe for firm i. Similarly, we can make τ^{safe} be safe for a given industry by taking H to be the set of histories corresponding to all the bids placed by firms in that industry.

Second, for finite data, we can choose $T(\cdot)$ to determine the significance level of test τ^{safe}. For instance, for the test to have a robust significance level of $\alpha \in (0, 1)$, we set $T(|H|)$ such that $2|\mathcal{N}|\exp(-\frac{1}{2}T(|H|)^2|H|) = \alpha$.

5.2 Maximum Lost Surplus

[XXX needs to be revisited]

Assume cartel members allocate contracts efficiently, and use reversion to competitive Nash as a threat.[13] When this is the case, any deviation temptation must be compensated by reducing the prices faced by the auctioneer. As a result, the sum of deviation temptation provides a measure of the welfare loss to the auction.

Given an environment ω_a, and constraint set $\mathcal{C}$ for environments ω_a, let

$$u(\omega_a) \equiv -\frac{1}{|A|}\sum_{h \in a}\left[\max_{n \in \mathcal{N}}[(1+\rho_n)b_h - c_h]d_{h,n} - (b_h - c_h)d_{h,0}\right] - \kappa \mathbf{1}_{\omega_a \in \mathcal{C}}$$

[13] Nash-reversion repeated-game equilibria figure prominently in the applied theory literature (e.g. Bull, 1987, Aoyagi, 2003, Baker et al., 1994, 2002).

with κ large enough.

$U(\omega_A)$ provides an estimate of surplus lost by the auctioneer. [XXX this needs more work]

6 Normative Foundations for Safe Tests

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 – for simplicity, we assume no transfers. At $t = \infty$, after firms played the game, a regulator runs a safe test on firms in N based on the realized history $h_\infty \in H_\infty$. We consider two different settings:

(i) the regulator runs a safe test τ_i on each firm $i \in N$;

(ii) the regulator runs a jointly safe test τ on all firms in N.

In case (i), a 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). In case (ii), all firms in N incur a penalty of $K \geq 0$ if and only if $\tau(h_\infty) = 1$.[14] When $K = 0$, under either form of testing the game collapses to the model in Section 3.

For each public strategy profile σ and each public history h_t, let $V_i(\sigma, h_t)$ denote firm i's expected continuation payoff under σ at h_t. 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. Define $\mathcal{V}(K)$ to be the set of perfect public Bayesian equilibrium values.

Proposition 4. *[safe tests do not create new repeated game equilibria] Suppose the regulator runs firm specific safe tests. Then, there exists $\overline{K} > 0$ such that, for all $K > \overline{K}$, $\Sigma(K) \subset \Sigma(0)$, and hence $\mathcal{V}(K) \subset \mathcal{V}(0)$.*

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

We now give an intuition as to why Proposition 4 holds. Note first that when the penalty K is large enough (i.e., $K > \overline{K}$), 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 $K > \overline{K}$ 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,

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

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 $\sigma \in \Sigma(0)$.

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 $K > \overline{K}$. In fact, the following stronger result holds:

Corollary 4. *For all $K > \overline{K}$, $\Sigma(K) = \Sigma^P(0)$.*

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 now show that, in some settings, data-driven regulation based on safe-tests may not have such unintended consequences. In particular, we show that running safe-tests at the industry level does not generate new collusive equilibria if we restrict attention to equilibria in which deviations are punished by Nash reversion.

Consider the game with joint testing. For each $i \in N$ and each value θ_{t-1} of the Markov state at the previous period, let $V_i^{NE}(\theta_{t-1})$ denote the expected discounted payoff firm i obtains from playing stage-game Nash at every future period.[15] For each K, let $\Sigma_J(K)$ denote the set of public perfect equilibria Bayesian equilibria of the game with joint testing and penalty K. Define

$$\Sigma_J^{RP}(K) \equiv \left\{\sigma \in \Sigma_J(K) : \text{ for all histories } h_t \text{ and all } i \in N, V_i(\sigma, h_t) \geq V_i^{NE}(\theta_{t-1})\right\},$$

to be the set of public perfect equilibria with the property that firms' continuation values are always above competitive payoffs. One reason to focus on equilibria in $\Sigma_J^{RP}(K)$ is that such equilibria satisfy a mild form of renegotiation proofness.

Let $\mathcal{V}_J^{RP}(K) \equiv \{(V_i(\sigma, h_0))_{i\in N} : \sigma \in \Sigma_J^{RP}(K)\}$ denote the set of values supported by equilibria in $\Sigma_J^{RP}(K)$. We then have:

Proposition 5. *Suppose the regulator runs a joint safe test. Then, there exists $\hat{K} > 0$ such that, for all $K > \hat{K}$, $\mathcal{V}_J^{RP}(K) \subset \mathcal{V}_J^{RP}(0)$.*

Beyond safe tests.

[15] If the stage game admits multiple stage-game Nash for some value θ, we take the one that minimizes firm i's payoff.

7 Empirical Evaluation

7.1 A Case Study

We start by illustrating our method with the data from the city of Tsuchiura.[16] In particular, we show how different deviations $\rho \in (-1, \infty)$ affect our estimates of the share of competitive histories. Throughout the Section, we fix an adapted set of histories H. We focus here on estimating the share of competitive histories. In Appendix [XXX] we present our estimates on the share of competitive auctions.

An upward deviation. Fix an adapted set of histories H, and assume for now that all histories in H are competitive. Consider a first a setting with only one single upward deviation $\rho_1 > 0$. By incentive compatibility, beliefs $(d_{h,0}, d_{h,1})$ and costs c_h at each history $h \in H$ satisfy:

$$[(1+\rho_1)b_h - c_h]\, d_{h,n} \leq [b_h - c_h] d_{h,0} \iff \frac{b_h}{c_h}[d_{h,0} - (1+\rho_1)d_{h,1}] \geq d_{h,0} - d_{h,1}. \tag{11}$$

Using $\frac{b_h}{c_h} \leq 1 + m$, we get that at each history, beliefs $(d_{h,0}, d_{h,1})$ satisfy

$$d_{h,0} \geq \left(\frac{\rho_1}{m} + 1 + \rho_1\right) d_{h,1}.$$

Summing this inequality across all histories and using $D_0(\omega_A, H) = \frac{1}{H}\sum_{h \in H} d_{h,n}$ yields

$$D_0(\omega_A, H) \geq \left(\frac{\rho_1}{m} + 1 + \rho_1\right) D_1(\omega_A, H)$$

[16] The city of Tsuchiura changed its auction format on October 29th 2009, introducing minimum prices; see Chassang and Ortner (forthcoming) for details. Here we focus on auctions that happened before that date.

Since $D_n(\omega_A, H) \in [\widehat{D}_n(H) - T(|H|), \widehat{D}_n(H) + T(|H|)]$ for all n, it must be that

$$\widehat{D}_0(H) + T(|H|) \geq \left(\frac{\rho_1}{m} + 1 + \rho_1\right)(\widehat{D}_1(H) - T(|H|)). \tag{12}$$

That is, if all histories in H are competitive, sample counterfactual demand $(\widehat{D}_0(H), \widehat{D}_1(H))$ must fall sufficiently fast when one moves from $\rho = 0$ to $\rho_1 > 0$.

When (12) is not satisfied, our estimate of competitive histories will be strictly less than 1. Let $H_1(H) = \{h \in H : (1+\rho_1)b_{h_{i,t}} < \wedge b_{-i,h_{i,t}}\}$ be the set of histories in H such that (i) the bidder associated with that history won the auction; and (ii) this bidder would have still won the auction if she had placed bid $(1+\rho_1)b_{h_{i,t}}$ instead of bid $b_{h_{i,t}}$. Note that the missing bids in Figure 1 imply that $H_1(H)$ is large whenever ρ_1 is small. Note further that, for any $\hat{H}$, $\widehat{D}_1(\hat{H}) = \frac{|H_1(\hat{H})|}{|\hat{H}|}$.

Let n_1 be the minimum number of histories in $H_1(H)$ that need to be dropped so that (12) after these histories are dropped; if (12) holds, set $n_1 = 0$. Then, our estimated number of competitive histories using a single upward deviation $\rho > 0$ is $\widehat{U}_{\mathsf{up}} = 1 - \frac{n_1}{|H|}$.

Adding a small downward deviation. Consider next adding a small downward deviation, $\rho_{-1} = -\epsilon \approx 0$. If all the histories in H are competitive, then for all $h \in H$

$$[(1+\rho_{-1})b_h - c_h]\, d_{h,-1} \leq [b_h - c_h] d_{h,0}.$$

Let $d_{h,0-}$ be the limit of $d_{h,-1}$ as $\rho_{-1} \to 0$. Incentive compatibility implies $d_{h,0} \geq d_{h,0-}$. Summing across all histories in H, we get $D_0(\omega_A, H) \geq D_{0^-}(\omega_A, H) = \lim_{\rho_{-1}\to 0} D_{-1}(\omega_A, H)$. Since $D_n(\omega_A, H) \in [\widehat{D}_n(H) - T(|H|), \widehat{D}_n(H) + T(|H|)]$ for all n, we get that, for $\rho_{-1} = -\epsilon \approx 0$,

$$\widehat{D}_0(H) + T(|H|) \geq \widehat{D}_{-1}(H) - T(|H|) \iff \widehat{D}_{-1}(H) - \widehat{D}_0(H) \leq 2T(|H|). \tag{13}$$

Note that, when $\rho_{-1} = -\epsilon \approx 0$, $\widehat{D}_{-1}(H) - \widehat{D}_0(H)$ is equal to the fraction of histories at which at least two firms tied in the first place.

Let n_{ties} be the minimum number of histories with tied winning bids that need to be dropped so that (13) holds; if (13) holds, set $n_{\mathsf{ties}} = 0$. Then, our estimated share of competitive histories using one single upward deviation $\rho > 0$ and a small downward deviation $\rho_{-1} = \epsilon \approx 0$ is $\widehat{U}_{\mathsf{up/ties}} = 1 - \frac{n_1}{|H|} - \frac{n_{\mathsf{ties}}}{|H|}$.

Adding a medium-sized downward deviation. We now show that, under certain conditions, adding a medium sized downward deviation $\rho_{-2} < 0$ can give us a tighter bound on the share of competitive histories. By incentive compatibility, at all competitive histories there must exist beliefs and cost satisfying

$$[(1 + \rho_{-2})b_h - c_h]\, d_{h,-2} \leq [b_h - c_h]d_{h,0} \iff d_{h,-2} - d_{h,0} \geq \frac{b_h}{c_h}[(1 + \rho_{-2})d_{h,-2} - d_{h,0}]. \quad (14)$$

Our estimate on the share of non-competitive histories can be computed by finding the largest subset $H^{\mathsf{comp}} \subset H$ such that, for all $h \in H^{\mathsf{comp}}$, there exists beliefs $(d_{h,n})_n$ and costs c_h satisfying all the relevant constraints (i.e., (7), (9), (10), (14), (11), (13), and the constraints in Program (6)).

It seems intuitive that adding downward deviation $\rho_{-2} < 0$ would lead to a tighter bound on the share of competitive histories. If estimated counterfactual demand $(\widehat{D}_n(H))_n$ increases fast as we move from $\rho_0 = 0$ to $\rho_{-2} < 0$, bidders might find it attractive to decrease their bids to raise their chances of winning. For such a deviation to not be profitable, firms' costs must be sufficiently large. But firms' incentives to raise their bids increase when costs are large; see equation (11). Hence, adding constraints (14) and (11) should lead to a lower estimate of the share of competitive histories.

However, in the absence of additional restrictions on beliefs, IC constraint (14) is not binding. Indeed, for all histories h_i at which bidder i won the auction, we can set $c_h \leq b_h$

and $d_{h,n} = \mathbf{1}_{(1+\rho_n)b_{h_{i,t}} < \wedge b_{-i,h_{i,t}}} = 1$ for $n = 0, -2$. And for all histories h_i at which bidder i lost the auction, we can again set $d_{h,n} = \mathbf{1}_{(1+\rho_n)b_{h_{i,t}} < \wedge b_{-i,h_{i,t}}}$ for $n = 0, -2$, and $c_h = b_h$. Note that, with these beliefs and costs, constraint (14) is satisfied at every history, and $D_n(H, \omega_A) = \hat{D}_n(H)$ for $n = -2, 0$.

Information constraints (10) rule out such extreme beliefs. As our estimates below show, in the presence of constraint (10), adding a mid-sized downward deviation ρ_{-2} leads to tighter estimates on the share of competitive histories.

Figure (6) presents our estimates on the share of competitive histories for our dataset from the city of Tsuchiura, as a function of parameter k in constraint (10). For these estimates and all the estimates that we presetnbelow, we set function $T(|H|)$ to satisfy $2|\mathcal{N}| \exp(-\frac{1}{2}T(|H|)^2|H|) = 5\%$, so that our tests have a robust confidence level of 5%.

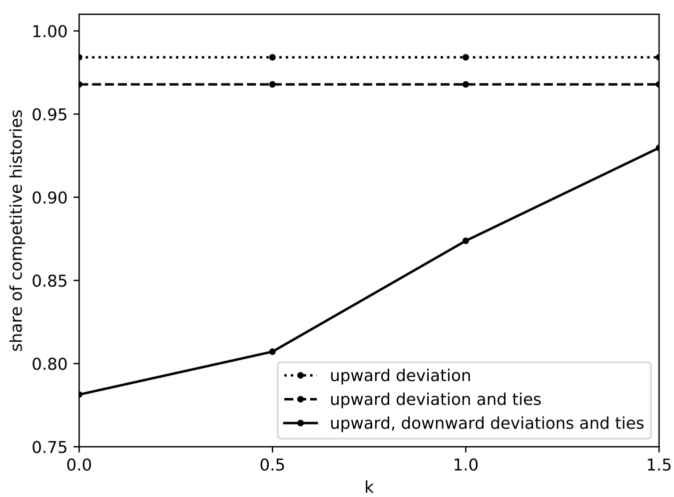

Figure 6: Estimated share of competitive auctions, city-level data.

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, in the analysis that follows 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.

Figure 7 plots our estimates for the share of competitive histories for the different sets of histories in Figure 4.[17] 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.

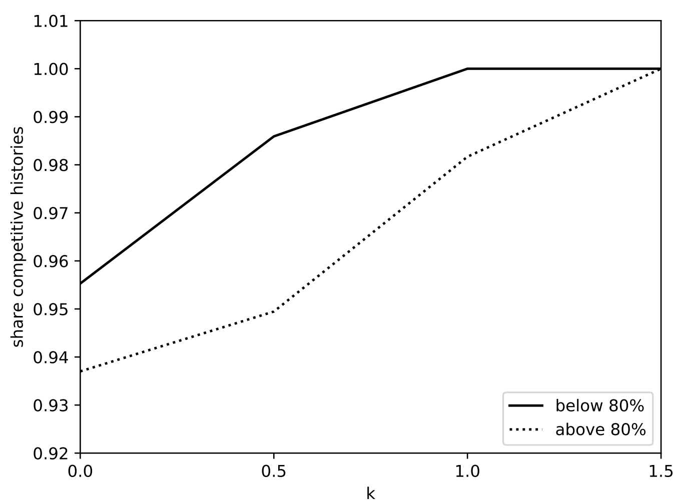

Figure 7: Estimated share of competitive auctions by bid level, national-level data.

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

Before and after prosecution. Figure 8 shows our estimates 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.

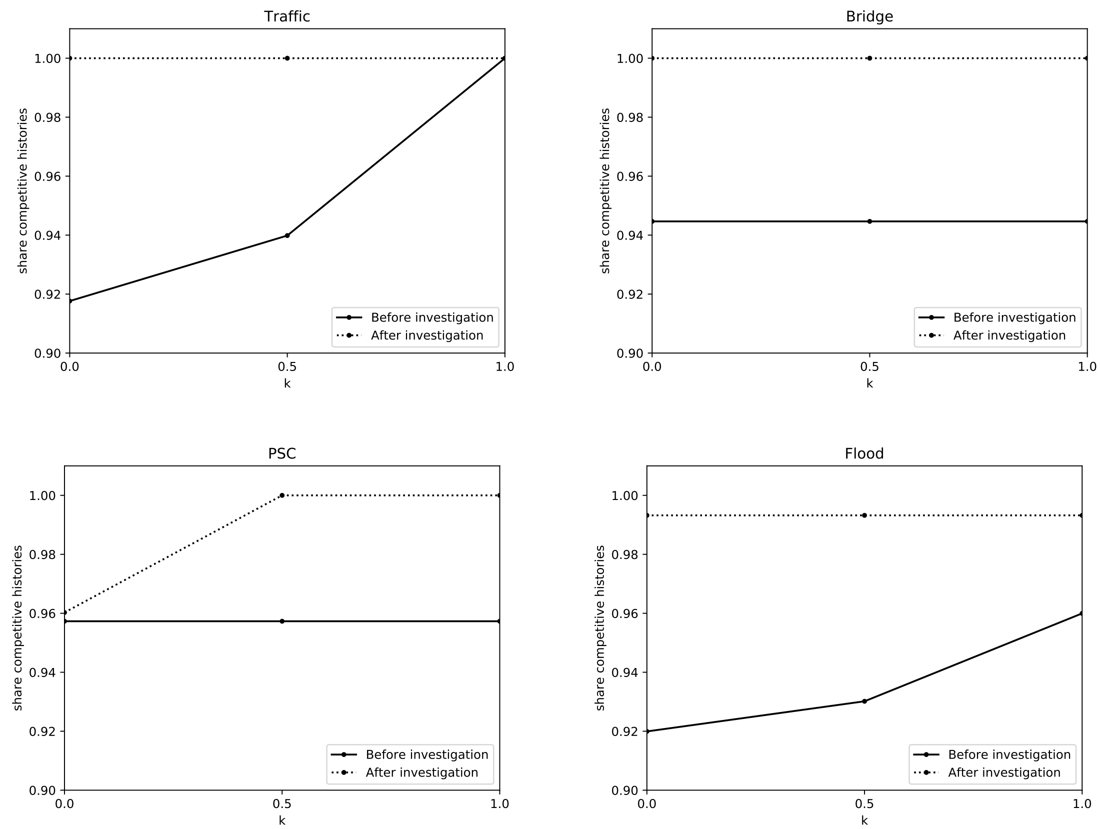

Figure 8: Estimated share of competitive auctions, before and after FTC investigation, national-level data.

7.3 Zeroing-in on specific firms

[TODO DATA: test outcomes for 3 largest firms in a given industry/in known cartelized industries]

8 Why Do Cartels Exhibit Missing Bids?

We conclude with an open ended discussion of why missing bids may be occurring in the first place.

This section has two objectives. First, we want to highlight that the bidding behavior we observe in our data is not easily explained by standard models of collusion. Second, we put forward an explanation for the bidding patterns we observe in these two datasets.

[XXX update text below] Finally, we propose a tentative explanation for missing bids, and why they could plausibly arise as an implication of collusive behavior. This is not entirely obvious because missing bids are not rationalized by standard models of tacit collusion (i.e., Rotemberg and Saloner (1986), Athey and Bagwell (2001, 2008)). In these models, the cartel's main concern is to incentivize losers not to undercut the winning bid. The behavior of designated winners is stage game optimal. We show that missing bids arise as an optimal repsonse to noise. Keeping the designated winner's bid isolated ensures that small trembles in play do not cause severe misallocations.

Missing bids is not a natural prediction of standard models.

Missing bids as coordination challenges.

Missing bids as a side effect of regulatory oversight.

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

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

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 $s \geq 0$, $\mathbb{E}_{\sigma,\mu}[\varepsilon_{i,t} | h_{j,t-s}] = \mathbb{E}_{\sigma,\mu}[\mathbb{E}_{\sigma,\mu}[\mathbf{1}_{\wedge b_{-i,t} > b_{i,t}(1+\rho)} | h_t, z_{i,t}] - \mathbf{1}_{\wedge b_{-i,t} > b_{i,t}(1+\rho)} | h_{t-s}, z_{j,t-s}] = 0$.[18]

Number the histories in H as $1, ..., |H|$ 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

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

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

[18] This 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.

A.2 Proofs of Section 4

Proof of Corollary 1. Suppose $s_{\mathsf{comp}} \geq 1-\kappa$ for some $\kappa > 0$. Then, for all $\rho > 0$,

$$\begin{aligned}
\frac{1}{\rho}[\widehat{R}(\rho|H) - \widehat{R}(0|H)] &= \frac{1}{\rho}[\overline{R}(\rho|H) - \overline{R}(0|H) + \widehat{R}(\rho|H) - \overline{R}(\rho|H) + \widehat{R}(0|H) - \overline{R}(0|H)] \\
&\leq 1 - s_{\mathsf{comp}} + \frac{1}{\rho}[\widehat{R}(\rho|H) - \overline{R}(\rho|H) - \widehat{R}(0|H) + \overline{R}(0|H)] \\
&\leq \kappa + \frac{1}{\rho}[\widehat{R}(\rho|H) - \overline{R}(\rho|H) - \widehat{R}(0|H) + \overline{R}(0|H)], \qquad (15)
\end{aligned}$$

where the first inequality follows from Proposition 2 and the second follows since $s_{\mathsf{comp}} \geq 1-\kappa$.

Next, note that for any scalar $\rho' \in (-1, \infty)$,

$$\overline{R}(\rho'|H) - \widehat{R}(\rho'|H) = \sum_{h_{i,t} \in H} \varepsilon_{i,t},$$

where

$$\varepsilon_{i,t} = \mathbb{E}_{\sigma,\mu}[(1+\rho')b_{i,t}\mathbf{1}_{\wedge b_{-i,t} > b_{i,t}(1+\rho')}|h_{i,t}] - (1+\rho')b_{i,t}\mathbf{1}_{\wedge b_{-i,t} > b_{i,t}(1+\rho')}].$$

By the law of iterated expectations, for all $h_{j,t-s} \in H$ with $s \geq 0$,

$$\mathbb{E}_{\sigma,\mu}[\varepsilon_{i,t}|h_{j,t-s}] = \mathbb{E}_{\sigma,\mu}[\mathbb{E}_{\sigma,\mu}[(1+\rho')b_{i,t}\mathbf{1}_{\wedge b_{-i,t} > b_{i,t}(1+\rho')}|h_t, z_{i,t}] - (1+\rho')b_{i,t}\mathbf{1}_{\wedge b_{-i,t} > b_{i,t}(1+\rho')}|h_{t-s}, z_{j,t-s}] = 0.$$

As in the proof of Theorem 1, number the histories in H as $1, ..., |H|$ 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

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

Note that, for all $\hat{k} \leq |H|$, $S_{\hat{k}} \equiv \sum_{k=1}^{\hat{k}} \varepsilon_k$ is a Martingale, with increments $\varepsilon_{\hat{k}}$ whose absolute value is bounded above by 1.[19] By the Azuma-Hoeffding Inequality, for all $\alpha > 0$, $\mathsf{prob}(|S_{|H|}| \geq |H|\alpha) = \mathsf{prob}(|\overline{R}(\rho'|H) - \widehat{R}(\rho'|H)| \geq \alpha) \leq 2\exp(-\alpha^2|H|/2)$.

Fix any $\rho \geq \rho^*$. Since this bound holds for all $\rho' \in (-1, \infty)$, it follows that

$$\mathsf{prob}(|\overline{R}(\rho|H) - \widehat{R}(\rho|H)| \geq \frac{\rho\kappa}{2} \text{ and } |\overline{R}(0|H) - \widehat{R}(0|H)| \geq \frac{\rho\kappa}{2}) \leq 4\exp(-(\rho\kappa)^2|H|/2) \leq 4\exp(-(\rho^*\kappa)^2|H|/2).$$

Combining this with equation (15), it follows that with probability at least $1-4\exp(-(\rho^*\kappa)^2|H|/2)$, $\frac{1}{\rho}[\widehat{R}(\rho|H) - \widehat{R}(0|H)] \leq 2\kappa$. ■

Proof of Proposition 3. By Lemma 1, under the true environment ω_A, $\mathsf{prob}(|\hat{D}_n(H(\omega_A)) - D_n(\omega_A, H(\omega_A))| \geq T(|H(\omega_A)|) \leq 2\exp(-T(|H(\omega_A)|)^2|H(\omega_A)|/2)$ for each deviation n. It then follows that

$$\mathsf{prob}(\forall n, |\hat{D}_n(H(\omega_A)) - D_n(\omega_A, H(\omega_A))| \geq T(|H(\omega_A)|) \leq 2|\mathcal{N}|\exp(-T(|H(\omega_A)|)^2|H(\omega_A)|/2).$$

This implies that, with probability at least $1 - 2|\mathcal{N}|\exp(-T(|H(\omega_A)|)^2|H(\omega_A)|/2)$, the constraints in Program (6) are satisfied when we set the environment equal to ω_A. Hence, with probability at least $1 - 2|\mathcal{N}|\exp(-T(|H(\omega_A)|)^2|H(\omega_A)|/2)$, $\hat{U} \geq U(\omega_A)$. ■

Proof of Corollary 3. Suppose ω_A 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 (or competitive auctions) is $s_{\mathsf{comp}} = 1$. By Corollary 2, with probability at least $1 - 2|\mathcal{N}|\exp(-T(|H|)^2|H|/2)$, $\hat{U} \geq s_{\mathsf{comp}} = 1$. Since $2|\mathcal{N}|\exp(-T(|H|)^2|H|/2) \to 0$ as $|H| \to \infty$, firms in this industry pass test τ^{safe} with probability approaching 1 as $|H| \to \infty$. ■

[19] This follows since we normalized reserve price to 1.

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

For any strategy profile σ and any history $h_{i,t} = (h_t, z_{i,t})$, let $V_i(\sigma, h_{i,t}) = \mathbb{E}_\sigma[\sum_{s \geq t} u_{i,s} | h_t]$ denote firm i's continuation payoff excluding penalties under σ at history $h_{i,t}$. Firm i's total payoff from under σ given history $h_{i,t}$ is $V_i(\sigma, h_{i,t}) - \mathbb{E}_\sigma[\tau_i | h_{i,t}]K$. Recall that a one-shot deviation by player i from a strategy σ_i is a strategy $\tilde{\sigma}_i \neq \sigma_i$ such that there exists a unique history $h_{i,\tau}$ such that $\sigma_i(h_{i,t}) = \tilde{\sigma}_i(h_{i,t})$ for all $h_{i,t} \neq h_{i,\tau}$.

Suppose $K > \overline{K}$ 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:[20] 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 $h_{i,t} \neq$

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

$h_{i,\tau}$,

$$V_i((\tilde{\sigma}_i, \sigma_{-i}), h_{i,\tau}) - \mathbb{E}_{(\tilde{\sigma}_i,\sigma_{-i})}[\tau_i|h_{i,\tau}]K \leq V_i(\sigma, h_{i,\tau}) - \mathbb{E}_{\sigma}[\tau_i|h_{i,\tau}]K$$
$$\iff V_i((\tilde{\sigma}_i, \sigma_{-i}), h_{i,\tau}) \leq V_i(\sigma, h_{i,\tau}), \tag{16}$$

where the second line in (17) follows since, under equilibrium σ, all firms pass the test with probability 1 at every history. By the second line in (17), 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, $\sigma \in \Sigma(0)$. ■

The following Lemma establishes 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. Clearly, if $\sigma \in \Sigma(K)$, then 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 $\tilde{\sigma}$ such that

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

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 σ is such that all firms pass the test with probability 1 at all histories, and since $\hat{\sigma}_i$ differs from σ_i at finitely many periods, all firms also pass the test under $(\hat{\sigma}_i, \sigma_{-i})$. Then, it must be that $V_i((\hat{\sigma}_i, \sigma_{-i}), h_{i,t}) - V_i(\sigma, h_{i,t}) \geq \epsilon/2$, where the strict inequality follows since firms' flow payoffs are bounded above by $r = 1$.

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 $t + T + 1$).

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 4. Fix $K > \overline{K}$. 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 $\Sigma(K) \subset \Sigma^P(K)$.

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

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

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. Recall that, at any equilibrium in $\Sigma^{RP}(K)$, firms' payoffs at every history h_t are bounded below by their Nash equilibrium profits $V_i^{NE}(\theta_{t-1})$.

Let $\hat{K} \equiv \sup_{i,\theta} \frac{1}{1-\delta} - V_i^{NE}(\theta)$. Suppose $K > \hat{K}$, and let $\sigma \in \Sigma^{RP}(K)$. Towards a contradiction, suppose that there exists a history h_t (on or off path) such that, at this history, firm i expects 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 $h_t \sqcup h_s$, firm i expects to fail the test with probability at least $\frac{\hat{K}}{K} + \epsilon < 1$. At history $h_t \sqcup h_s$, firm i's continuation payoff is bounded above by $\frac{1}{1-\delta} - \left(\frac{\hat{K}}{K} + \epsilon\right) K = \inf_{j,\theta} V_j^{NE}(\theta) - \epsilon K < V_i^{NE}(\theta_{t+s-1})$, a contradiction.

Suppose $K > \hat{K}$ 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 $h_{i,t} \neq h_{i,\tau}$,

$$\begin{aligned} V_i((\tilde{\sigma}_i, \sigma_{-i}), h_{i,\tau}) - \mathbb{E}_{(\tilde{\sigma}_i, \sigma_{-i})}[\tau_i | h_{i,\tau}] K &\leq V_i(\sigma, h_{i,\tau}) - \mathbb{E}_{\sigma}[\tau_i | h_{i,\tau}] K \\ \Longleftrightarrow V_i((\tilde{\sigma}_i, \sigma_{-i}), h_{i,\tau}) &\leq V_i(\sigma, h_{i,\tau}), \end{aligned} \tag{17}$$

where the second line in (17) follows since, under equilibrium σ, all firms pass the test with probability 1 at every history. By the second line in (17), 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, $\sigma \in \Sigma(0)$. ■

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.

Aoyagi, M. (2003): "Bid rotation and collusion in repeated auctions," *Journal of economic Theory*, 112, 79–105.

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.

Baker, G., R. Gibbons, and K. J. Murphy (1994): "Subjective Performance Measures in Optimal Incentive Contracts," *The Quarterly Journal of Economics*, 1125–1156.

——— (2002): "Relational Contracts and the Theory of the Firm," *The Quarterly Journal of Economics*, 117, 39–84.

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.

Bull, C. (1987): "The existence of self-enforcing implicit contracts," *The Quarterly journal of economics*, 102, 147–159.

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.

Foster, D. P. and R. Vohra (1999): "Regret in the on-line decision problem," *Games and Economic Behavior*, 29, 7.

Foster, D. P. and R. V. Vohra (1998): "Asymptotic calibration," *Biometrika*, 85, 379–390.

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.

Maskin, E. and J. Tirole (2001): "Markov perfect equilibrium: I. Observable actions," *Journal of Economic Theory*, 100, 191–219.

Olszewski, W. and A. Sandroni (2008): "Manipulability of Future-Independent Tests," *Econometrica*, 76, 1437–1466.

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.