

NBER WORKING PAPER SERIES

DATA DRIVEN REGULATION:
THEORY AND APPLICATION TO MISSING BIDS

Sylvain Chassang
Kei Kawai
Jun Nakabayashi
Juan M. Ortner

Working Paper 25654
<http://www.nber.org/papers/w25654>

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
March 2019

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, Jozsef Sakovics, Larry Samuelson, Andy Skrzypacz, Paulo Somaini. The views expressed herein are those of the authors and do not necessarily reflect the views of the National Bureau of Economic Research.

NBER working papers are circulated for discussion and comment purposes. They have not been peer-reviewed or been subject to the review by the NBER Board of Directors that accompanies official NBER publications.

© 2019 by Sylvain Chassang, Kei Kawai, Jun Nakabayashi, and Juan M. Ortner. All rights reserved. Short sections of text, not to exceed two paragraphs, may be quoted without explicit permission provided that full credit, including © notice, is given to the source.

Data Driven Regulation: Theory and Application to Missing Bids
Sylvain Chassang, Kei Kawai, Jun Nakabayashi, and Juan M. Ortner
NBER Working Paper No. 25654
March 2019
JEL No. C72,D44,H57,L51

ABSTRACT

We document a novel bidding pattern observed in procurement auctions from Japan: winning bids tend to be isolated. There is a missing mass of close losing bids. This 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.

Sylvain Chassang
Department of Economics
New York University
19 West 4th Street
New York, NY 10012
and NBER
sylvain.chassang@gmail.com

Jun Nakabayashi
Kindai University
3-4-1 Kowakae
Higashiosaka City
Osaka 577-8502
Japan
nakabayashi.1@eco.kindai.ac.jp

Kei Kawai
Department of Economics
University of California, Berkeley
530 Evans Hall #3880
Berkeley, CA 94720
and NBER
kei@berkeley.edu

Juan M. Ortner
Boston University
270 Bay State Rd,
Boston, MA 02215
jortner@bu.edu

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 paper seeks to make progress on several questions relevant to this literature. 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 target only a specific pattern of behavior? Can we ensure that regulatory policies do not end up strengthening cartels?

We begin by documenting a suspicious bidding pattern observed in procurement auctions in Japan: the density of the bid distribution just above the winning bid is very low. There is a missing mass of close losing bids. These missing bids are related to bidding patterns of collusive firms in Hungary (Tóth et al., 2014) and Switzerland (Imhof et al., 2016). We establish that these missing bids indicate non-competitive behavior under a general class of asymmetric information models, and under the presence of arbitrary unobserved heterogeneity. 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 two datasets of public works procurement auctions in Japan. The first dataset, analyzed by Kawai and Nakabayashi (2018), contains data on approximately 78,000 national-level auctions held between 2001 and 2006 by the Ministry of Land, Infrastructure and Transportation. The second dataset, studied by Chassang and Ortner (forthcoming), contains information on approximately 1,500 city-level auctions held between 2007 and 2014. We are interested in the distribution of bidders' margins of victory (or defeat). For every (bidder, auction) pair, we compute Δ , the difference between the bidder's own bid and the most competitive bid among this bidder's opponents, divided 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 the data — are suspicious. Our second goal is to formulate a theory of regulatory response to such data.

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 allow players to observe arbitrary signals about one another, and rule out intertemporal linkages between actions and payoffs. 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




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

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 such 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, provided penalties against colluding firms are large enough,

it 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. The sample of histories such that bids are close to the reserve price is more likely to fail our tests than histories where bids are low relative to the reserve price. 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.

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.² We propose a class of robust, systematic tests of non-competitive behavior guaranteed to make cartel formation less attractive 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.³ A common

¹See Harrington (2008) for a recent survey.

²Also related is Schurter (2017), who proposes a test of collusion based on exogenous shifts in the number of bidders.

³Other papers, like Besanko and Spulber (1989) and LaCasse (1995), study static models of equilibrium regulation.

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 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: for a given finite game, they study the range of behavior that can be sustained by some incomplete information structure.⁴ A similar exercise is at the heart of our analysis, though we choose to consider optimality conditions that are weaker and more easily satisfied than equilibrium. 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., 2018).

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.⁵ 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

⁴Also closely related is Bergemann et al. (2017), which characterizes bounds on equilibrium revenue in first-price auctions under arbitrary incomplete information. Instead, we focus on implications for bidding behavior under arbitrary incomplete information, and impose conditions that are weaker (and hence more plausibly satisfied) than equilibrium. See also Doval and Ely (2019) for an extension allowing for arbitrary extensive form in addition to arbitrary incomplete information.

⁵See Chambers and Echenique (2016) for a recent review of the literature on revealed preference.

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 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 indicators of collusion. Section 8 concludes with an open ended discussion of why missing bids may arise in the context of collusion. Appendix A reports descriptive statistics of our data, as well as additional empirical results. Appendix B clarifies the connection between our approach and Bayes correlated equilibrium (Bergemann and Morris, 2016). Appendix C shows how our results extend under common values. Appendix D describes our computational strategy. Proofs are collected in Appendix E unless mentioned otherwise.

2 Motivating Facts

Our first dataset, described in Kawai and Nakabayashi (2018), consists of roughly 78,000 auctions held between 2001 and 2006 by the Ministry of Land, Infrastructure and Transportation in Japan (the Ministry). The auctions are sealed-bid first-price auctions with a secret reserve price, and re-bidding in case there is no successful winner. The auctions involve construction projects. The median winning bid is 1 million USD, and the median number of participants is 10. The bids of all participants are publicly revealed after each auction.

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 $\mathbf{b}_{-i,a}$ the profile of bids by bidders other than i . We investigate the distribution of

$$\Delta_{i,a} = \frac{b_{i,a} - \wedge \mathbf{b}_{-i,a}}{r}$$

aggregated over firms i , and auctions a , where $\wedge$ denotes the minimum operator. 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 $\Delta_{i,a} > 0$ she lost.

The left panel of Figure 2 plots the distribution of bid differences Δ aggregating over all firms and auctions in our sample.⁶ The mass of missing bids around a difference of 0 is 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 normalized bid-differences Δ – national-level 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 Ibaraki prefecture, also in Japan. Projects are allocated using a sealed-bid first-price auction with a public reserve

⁶Note 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.

price. The median winning bid is 130,000 USD, and the median number of participants 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.⁷




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

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

Figure 4 breaks down the auctions in Figure 2 by bid levels: the figure plots the distribution of $\Delta_{i,a} = \frac{b_{i,a} - \wedge \mathbf{b}_{-i,a}}{r}$ for normalized bids $\frac{b_{i,a}}{r}$ below 0.8 and above 0.8. The mass of missing bids in Figure 2 is considerably reduced when we look at bids that are low compared to the reserve price.

Figure 5 plots the distribution of $\Delta_{i,a}$ for participants of auctions held by the Ministry

⁷Imhof 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).

that were implicated by the Japanese Fair Trade Commission (JFTC). The JFTC implicated a total of four bidding rings that were participating in the auctions in our data: (i) firms installing traffic signs; (ii) builders of bridge upper structures; (iii) pre-stressed concrete providers; and (iv) floodgate builders.⁸ The left panels in Figure 5 plot the distribution of Δ for auctions that were run 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, firms in case (iii) initially denied the charges against them, unlike in the other three cases, and seem to have continued colluding for some time (see Kawai and Nakabayashi (2018) for a more detailed account of these collusion cases).




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

What does not explain this pattern. Although explaining missing bids is not the goal of this paper, it is useful to clarify what does not explain this pattern. Specifically, we argue

⁸See 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 5: Distribution of bid-difference Δ – cartel cases in national data, before and after JFTC investigation.

that missing bids are not explained by either the granularity of bids, or ex post renegotiation.

Figures 4 and 5 rule out granularity as an explanation for missing bids: if missing bids were a consequence of the granularity of bids, we should see similar patterns both across all bid levels, and before and after the JFTC investigations.⁹

Renegotiation could potentially account for missing bids by making apparent incentive compatibility issues irrelevant. Indeed, some winning firms may seemingly leave money on the table, only to reclaim it through renegotiation ex post. Our national-level data contains information 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, in the auctions we study, the contract that is signed between the awarder and the awardee include renegotiation provisions that greatly reduce firms' incentives to bid aggressively with the hope of renegotiating to a higher price later on. Specifically, the contract stipulates that renegotiated prices be anchored to the level of the initial bid: if the project is estimated to cost y more than initially thought, the renegotiated price is increased by $\frac{\text{initial bid}}{\text{reserve price}} \times y$. This implies that excessively competitive bids that are unprofitable are likely to remain unprofitable after renegotiation.

Our objectives in this paper are: (i) to formalize why the missing mass of bids around zero 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 safe 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.

⁹In addition, there is no missing bid pattern when comparing the second and third lowest bids.




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

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 sealed-bid first-price auction with reserve price r , which we normalize to $r = 1$.

In each period $t \in \mathbb{N}$, a set $\hat{N}_t \subset N$ 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.¹⁰ The set of eligible bidders can vary over time.

Realized costs of production for eligible bidders $i \in \hat{N}_t$ are denoted by $\mathbf{c}_t = (c_{i,t})_{i \in \hat{N}_t}$. Each bidder $i \in \hat{N}_t$ submits a bid $b_{i,t}$. Profiles of bids are denoted by $\mathbf{b}_t = (b_{i,t})_{i \in \hat{N}_t}$. 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

¹⁰See Chassang and Ortner (forthcoming) for a treatment of endogenous participation by cartel members.

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. Ties are broken randomly.

Participants discount future payoffs using common discount factor $\delta < 1$. Bids are publicly revealed at the end of each period.¹¹

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

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

where:¹²

- 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.¹³ We 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 $i \in \hat{N}_t$ obtains profits

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

where $x_{i,t} \in [0, 1]$ is the probability with which i wins the auction at time t . Note that costs include both the direct costs of production and the opportunity cost of backlog.

¹¹For the sake of concision, we do not consider the possibility of transfers at this point. It does not change our analysis.

¹²We stress that cost specification (1) is flexible. For instance, this formulation can accommodate environments in which costs depend on the physical distance between each firm i and the project's location. Indeed, this can be achieved by setting θ_t to be the vector of distances from each firm to the project's location at time t , and α_i to be a vector with all zeros except a 1 in the i -th position.

¹³I.e. given any event E anterior to time t , $\theta_t | \theta_{t-1}, E \sim \theta_t | \theta_{t-1}$.

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.

3.2 Solution Concepts

A public history h_t in period t takes the form $h_t = (\theta_{s-1}, \mathbf{b}_{s-1})_{s \leq t}$. We let $\mathcal{H}$ denote the set of all public histories.

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 emphasize the class of competitive equilibria, which simply corresponds to the class of 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. Competitiveness of equilibrium is a fairly coarse notion. An equilibrium, which could involve many firms, interacting over an extensive timeframe, is either competitive or not. More realistically, 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

We take the perspective of a regulator using bidding data to decide whether or not to launch an investigation. We consider a regulator that seeks to do as little harm to competitive firms as possible, and formalize her decision to intervene as a *safe test*. That is, a test that is robustly passed by competitive firms. We show in Section 6 that preventing harm against competitive firms serves an important strategic purpose: provided the cost of regulatory intervention is high enough (at least for collusive firms), safe tests cannot unwittingly increase a cartel's enforcement capability. As Harrington (2004) highlights, this need not be true for tests that are not safe.

Let H_∞ denote the set of coherent full public histories $(h_{i,t})_{i \in N, t \geq 0}$.¹⁴ A test τ is a mapping from H_∞ to $\{0, 1\}$, where 1 denotes fail and 0 denotes pass.

Let us 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 .

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, $\text{prob}_{\lambda, \sigma}(\tau_i(h) = 0) = 1$.*

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, $\text{prob}_{\lambda, \sigma}(\tau(h) = 0) = 1$.

In words, a test is jointly safe (resp. unilaterally safe) if and only if all competitive industries (resp. firms) pass the test with probability 1.

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 $h_{i,t}$ and all bids $b' \in [0, r]$, player i 's *counterfactual demand* at $h_{i,t}$ is

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

¹⁴I.e. sequences of histories such that $h_{i,t}$ and $h_{i,t+1}$ coincide on events occurring before time t , and $h_{i,t}$ and $h_{j,t+1}$ coincide on public events before t . For each $i \in N$ and t , history $h_{i,t}$ takes the form $h_{i,t} = (h_t, z_{i,t})$.

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

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

the average counterfactual demand for histories in H , and its sample equivalent

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

Definition 4. *We say that set of histories 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 or location 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 an adapted set of histories H . Under any perfect public Bayesian equilibrium σ , for any $\alpha > 0$,*

$$\text{prob}(|\widehat{D}(\rho|H) - \overline{D}(\rho|H)| \leq \alpha) \geq 1 - 2 \exp(-\alpha^2|H|/2).$$

In particular, with probability 1, $\widehat{D}(\rho|H) - \overline{D}(\rho|H) \rightarrow 0$ as $|H| \rightarrow \infty$.

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 a consequence of the equilibrium assumption, but is implied by optimality conditions weaker than equilibrium. It may fail under sufficiently non-common priors, but will hold if participants use data-driven predictors of demand satisfying no-regret (see for instance Hart and Mas-Colell, 2000).

We highlight that Lemma 1 relies on conservative non-asymptotic concentration bounds. In practice, one may be willing to make additional assumptions on the data generating process leading to tighter estimates. Our approach extends to any probabilistic moment conditions of the form

$$\text{prob} \left(\left(\widehat{D}(\rho|H) - \overline{D}(\rho|H) \right)_{\rho \in R} \in S \right) \geq 1 - \epsilon$$

where R is a finite set of deviations $\rho \in (-1, +\infty)$, and S is a confidence set for demand estimation errors $\widehat{D}(\rho|H) - \overline{D}(\rho|H)$ at different values of $\rho \in R$.¹⁵

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,*

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

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

In words, under any competitive 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. Since the density of Δ_i at 0 is essentially 0, the elasticity of demand is approximately zero at these histories.

¹⁵For instance, if one is willing to impose that $(\theta_t)_{t \in \mathbb{N}}$ is stationary and ergodic, and that σ is a competitive equilibrium, then $\widehat{D}(\rho|H) - \overline{D}(\rho|H)$ is asymptotically Gaussian when set H grows large. One can then derive confidence sets by estimating an asymptotic covariance matrix. Alternatively one may also derive confidence sets using bootstrap. This is the approach we use in our empirical analysis.

Proof. Consider a competitive equilibrium σ . Let

$$V(h_{i,t}) \equiv \mathbb{E}_\sigma \left(\sum_{s \geq t} \delta^{s-t} (b_{i,s} - c_{i,s}) \mathbf{1}_{b_{i,s} < \wedge \mathbf{b}_{-i,s}} \middle| h_{i,t} \right)$$

denote player i 's discounted expected payoff at history $h_{i,t}$. Let b denote the bid that bidder i places at history $h_{i,t}$. Since b is an equilibrium bid, it must be that for all bids $b' > b$,

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

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

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

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. $\blacksquare$

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 corresponding natural upper bound to valuations. 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. Note that, unlike in tests of competitive behavior that presume a particular information structure, the test here is valid under general information structures including private values and common values. It is also valid under any form of unobserved heterogeneity.¹⁶ 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 . 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

¹⁶Since Proposition 1 holds conditional on the information that is available to firms, it must also hold conditional on the possibly coarser information available to the econometrician.

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 histories is such that*

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

Proof. Let $H^{\neg \text{comp}} = H \setminus H^{\text{comp}}$ be the set of non-competitive histories in H . For any $\rho > 0$,

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

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, since $\bar{R}(\rho|H^{\neg \text{comp}}) \leq (1 + \rho) \bar{R}(0|H^{\neg \text{comp}})$,

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

This concludes the proof. ■

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

Proposition 2 suggests a natural estimator. For each $\rho \in (-1, \infty)$, define

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

Note that $\widehat{R}(\rho|H)$ is the sample analog of counterfactual revenue. A result identical to Lemma 1 establishes that $\widehat{R}(\rho|H)$ is a consistent estimator of $\overline{R}(\rho|H)$, whenever set H is adapted. This allows us to use Proposition 2 to empirically estimate the share of competitive histories.

In the extreme case where the density of competing bids is zero just above winning bids, we have that $\overline{R}(\rho|H) - \overline{R}(0|H) \simeq \rho \overline{R}(0|H)$ for ρ small. This implies that $s_{\text{comp}} \leq 1 - \overline{R}(0|H)$.¹⁷

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. We emphasize that our tests allow the regulator or econometrician to include plausible constraints on costs, incomplete information, and the extent of unobserved heterogeneity.

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 $n \in \mathcal{M} = \{-\underline{n}, \dots, \overline{n}\}$, such that $\rho_0 = 0$ and $\rho_n < \rho_{n'}$ for

¹⁷Note that $s_{\text{comp}} \leq 1 - \overline{R}(0|H) \leq 1 - \overline{D}(0|H)$. Hence, the bound on competitive histories in Proposition 2 is, at most, equal to the probability of winning an auction. While this bound is quite weak, we note that it is robust to private values, common values, and any form of unobserved heterogeneity.

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, $d_{h_{i,t},n}$ 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 $\omega_a = (d_{h,n}, c_h)_{h \in a, n \in \mathcal{M}}$.¹⁸ We let $\omega_A = (\omega_a)_{a \in A}$ denote the profile of environments across auctions $a \in A$.

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

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

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

Intuitively, estimator $\widehat{U}$ is a robust upper-bound to the true value $U(\omega_A)$. 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. We refer to this constraint as consistency requirement $(\widehat{CR})$. The following result holds.

¹⁸Note that each history $h \in H$ corresponds to a bid placed by some bidder in an auction $a \in A$. Hence, the histories $h \in a$ associated with auction a correspond to all the bids that were placed in auction a .

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

As we discussed following Lemma 1, any bound of the form

$$\text{prob}((\widehat{D}(\rho_n|H) - \overline{D}(\rho_n|H))_{n \in \mathcal{M}} \in S) \geq 1 - \epsilon$$

yields an extension of Proposition 3 where condition $(\widehat{CR})$ is replaced by $(\widehat{D}(\rho_n|H) - \overline{D}(\rho_n|H))_{n \in \mathcal{M}} \in S$. In that case, estimated metric $\widehat{U}$ is greater than the true metric $U(\omega_A)$ with probability $1 - \epsilon$.

We now show that by applying Proposition 3 to different tuples (u, Ω) , 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 2, and provide a tighter upper bound on the share of competitive histories in H . For expositional purposes, we assume private values and discuss extensions to common values at the end of this section.

Under the private values assumption, at every competitive history $h \in H$ there must exist costs c_h and subjective demands $d_h = (d_{h,n})_{n \in \mathcal{M}}$ satisfying

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

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

Conditions (F), (IC) and $(\widehat{CR})$ exploit some but not all the information content of equilibrium under some information structure. We clarify in Appendix B that this would be the case if we imposed demand consistency requirements conditional on different values of bids

and costs c (corresponding to the bidder's private information at the time of bidding).

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

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

We allow for economic plausibility constraints on environments Ω . We focus on markup constraints, as well as constraints on the informativeness of signals bidders get.

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

where $m \geq 0$ and $M \in (m, +\infty]$ are minimum and maximum markups, and $k \in [0, +\infty)$ provides an upper bound to the information contained in any signal. The focal case of i.i.d. types corresponds to $k = 0$.

For each environment ω_A , define

$$H_{\text{comp}}(\omega_A) \equiv \{h \in H \text{ s.t. } (d_h, c_h) \text{ satisfy (F), (IC) and (EP)}\}$$

to be the set of histories in H satisfying plausibility constraints (EP) that can be rationalized as competitive under ω_A . Program (P) then becomes

$$\begin{aligned} & \max_{\omega_A \in \Omega} \frac{|H_{\text{comp}}(\omega_A)|}{|H|} \quad (\text{P}') \\ & \text{s.t. } \forall n, \quad D_n(\omega_A, H) \in \left[\widehat{D}_n(H) - T(|H|), \widehat{D}_n(H) + T(|H|) \right]. \end{aligned}$$

$\widehat{U}$ 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 $s_{\text{comp}} \in (0, 1]$. Then, with probability at*

least $1 - 2|\mathcal{M}| \exp(-\frac{1}{2}T(|H|)^2|H|)$, $\hat{U} \geq s_{comp}$.

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 tolerance function $T(\cdot)$ 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 again 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, or by using bootstrap methods.

Common values. In the analysis above, the private value assumption implies that each history h is associated with a single cost c_h . Under common values, costs would depend on the deviation ρ chosen by the bidder. An environment at auction a would then take the form $\omega_a = (d_{h,n}, c_{h,n})_{h \in a, n \in \mathcal{M}}$. However, we show in Appendix C that, under the assumption that a bidder's expected costs $c_{h,n}$ are weakly increasing in her opponents' bids (a variant of affiliation), common values cannot relax incentive compatibility constraint (IC). Hence, our results extend without modification to such common value environments.

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 quantities 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, if all of her opponents played competitively, 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.

We define the objective function to be

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

For every environment ω_A , let

$$A_{\text{comp}}(\omega_A) \equiv \{A' \subset A \text{ s.t. } \forall a \in A', \forall h \in a, (d_h, c_h) \text{ satisfy (F) (IC) \& (EP)}\}.$$

be the set of competitive auctions under ω_A . Program (P) then becomes

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

$\hat{U}$ 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. Proposition 3 can be used to derive a lower bound

on the bidders' deviation temptation. Given an environment ω_a , we define objective

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

In this case, with large probability, $-\hat{U}$ 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)), $\hat{U}$ 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, safe tests 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 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 undiscounted penalty of $K \geq 0$ if and only if $\tau_i(h_\infty) = 1$ (i.e., if and only if firm i fails the test).¹⁹

- (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 $\tau(h_\infty) = 1$.

¹⁹Our analysis extends as is if K is the expected cost experienced by colluding firms from being investigated, while non-colluding firms experience a potentially lower cost.

When $K = 0$, under either form of testing the game collapses to the model of Section 3.²⁰

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 $K > \bar{K}$, $\Sigma(K) \subset \Sigma(0)$.*

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 $V_i : \mathbf{b} \mapsto \mathbb{R}$ 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

²⁰Since penalty K is undiscounted, the regulatory game is not continuous at infinity whenever $K > 0$.

shot deviation in a game without testing. The one-shot deviation principle then implies that $\sigma \in \Sigma(0)$.²¹

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 strictly reduce the surplus attainable by collusive bidders. We provide an explicit example in Section 8.

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 $K > \bar{K}$, $\Sigma_{RP}^+(K) \subset \Sigma_{RP}^+(0)$.*

²¹Note that the game with $K = 0$ is continuous at infinity, and so the one-shot deviation principle holds in such game.

Beyond safe tests. Sufficiently punitive safe tests do not make collusion more profitable, and do 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.

Still, 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 her 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. Details regarding the computation of the tests are collected in Appendix D.

7.1 A Case Study

We first illustrate the mechanics of inference using data from the city of Tsuchiura.²² Specifically, we show how three different deviations $\rho \in \{-.02, -.0005, .0008\}$ affect our estimates of the share of competitive histories. The distribution of Δ and the three deviations that we consider are illustrated in Figure 7.

Figure 8 presents our estimates of the 95% confidence bound on the share of competitive histories as a function of the correlation parameter k in constraint (EP). For these estimates

²²Chassang 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 $\rho \in \{-0.02, -0.0005, 0.0008\}$.

and all the estimates that we present below, we use bootstrap to set tolerance T so that our tests have a confidence level of 5%. We now delineate the mechanics of inference.




Figure 8: Estimated share of competitive auctions, city-level data. Minimum and maximum markups: $m = 0.02$, $M = 0.25$.

An upward deviation. We first consider a small upward deviation $\rho = .0008$, corresponding to the analysis of Section 4. Because the mass of values of Δ falling between 0 and .0008 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. Note that this 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 bidder’s 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 $\rho = -.02$ 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 2% drop in prices leads to a 44% increase in the probability of winning the auction. If costs of production are sufficiently below bids, this is an attractive deviation.

7.2 Consistency between safe tests and proxies for collusion

We now turn to the national data to 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. Figure 4 plots the distribution of Δ for histories with high versus low bid-to-reserve ratios (i.e., bids above/below 80% of the reserve price). The pattern of missing

bids is more striking when we focus on histories at which bidders place relatively high bids. To the extent that missing bids are a marker of non-competitive behavior, Figure 4 suggests that histories at which firms place relatively low bids are more likely to be competitive.

Figure 9 plots our estimates of the 95% confidence bound on the share of competitive histories for the different sets of histories in Figure 4. 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 for histories at which bidders place higher bids.




Figure 9: Estimated share of competitive auctions by bid level, national data. Minimum and maximum markups: $m = 0.02$, $M = 0.5$.

Before and after prosecution. Figure 10 shows our estimates of the 95% confidence bound on the share of competitive histories for the four groups of firms that were investigated by the JFTC (see Section 2, and in particular Figure 5). Our estimates suggests non-competitive behavior in the before period across the four groups of firms, and show a drop in non-competitive behavior following investigation.




Figure 10: Estimated share of competitive auctions, before and after FTC investigation, national-level data. Minimum and maximum markups: $m = 0.02$, $M = 0.5$.

7.3 Zeroing-in on specific firms

Finally we highlight that our tests can be applied to individual firms. We focus on the thirty firms that participate in the most auctions in our national level data. Table 1 shows our estimates of the 95% confidence bound on the share of competitive histories for these thirty firms. Column 4 in the table shows estimates obtained from Proposition 2. Column 5 shows tighter estimates obtained from Corollary 1. We find that the bound from Proposition 2 is less than 1 for only three of the thirty firms, while the bound from Corollary 1 is less than 1 for twenty-five firms.²³

²³Note that the counterfactual demands of individual firms within the same industry turn out to be highly correlated. This leads to the observed correlation in estimates reported in Table 1. Firms exhibiting similar

Ranking	Participation	Share won	Share comp. (Prop. 2)	Share comp. (Coro. 1)
1	4044	0.17	1.00	0.60
2	3854	0.07	1.00	0.68
3	3621	0.12	1.00	0.60
4	2998	0.15	1.00	1.00
5	2919	0.06	1.00	0.68
6	2547	0.08	1.00	0.15
7	2338	0.07	1.00	0.15
8	2333	0.07	1.00	0.18
9	2328	0.04	1.00	0.68
10	2292	0.06	1.00	0.20
11	2237	0.08	0.92	0.68
12	2211	0.03	1.00	0.68
13	2015	0.09	1.00	0.28
14	1984	0.08	1.00	0.28
15	1727	0.07	1.00	1.00
16	1674	0.05	1.00	0.31
17	1661	0.03	1.00	0.68
18	1660	0.08	1.00	0.28
19	1589	0.07	1.00	0.28
20	1427	0.10	1.00	1.00
21	1393	0.06	1.00	0.31
22	1392	0.07	1.00	1.00
23	1370	0.04	1.00	0.56
24	1368	0.14	1.00	1.00
25	1353	0.05	1.00	0.16
26	1342	0.09	1.00	0.91
27	1337	0.04	1.00	0.25
28	1326	0.08	1.00	0.60
29	1291	0.06	0.95	0.31
30	1260	0.06	0.95	0.60

Table 1: Estimated Share of Competitive Auctions, National-Level Data. The table reports our estimates of the share of competitive auctions for the thirty firms that participate in the most number of auctions during the sample period. The first column corresponds to the ranking of the firms and the second column corresponds to the number of auctions in which each firm participates. Column 3 shows the fraction of auctions that each of these firms wins. Columns 4 and 5 present our 95% confidence bound on the share of competitive histories for each firm based on Proposition 2 and Corollary 1, respectively. For our estimates of column five, we use the following parameters: $m = 0.02$, $M = 0.5$ and $k = 1.5$.

shares of non-competitive histories belong to the same industries.

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 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. If this were not the case, the cartel’s pledgeable surplus would have to be spent on keeping the winner from increasing her bid rather than keeping losing bidders from undercutting.

One implication of this is that there exist safe tests that strictly reduce the surplus of cartels. Given a set of auctions A , consider tests that rule out frequent approximately-tied winning bids:

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

where $\kappa, \epsilon > 0$, and $b_{(1)}, b_{(2)}$ are the lowest and second lowest bids. 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. In fact, such tests are likely very highly powered: the probability that the lowest and second-lowest bids are tied for a significant share auctions is likely very low under reasonable priors. This explains why regulators use tied bids as an indicator of collusion (DOJ, 2005). It also suggests an explanation for missing bids, as we now explain.

Missing bids as a side-effect of existing regulatory oversight. It is possible that regulatory oversight itself may be at the origin of the missing bid pattern. Because test $\tau_{\kappa, \epsilon}$ is well powered with relatively small data sets, it makes sense for the regulator to scrutinize auctions in which tied bids occur. If this is the case, 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 may lead to the detectable missing bids pattern highlighted in this paper.

Missing bids as robust coordination. Another 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).

Online Appendix – Not for Publication

A Further Empirics

A.1 Sample Statistics

In this section we report descriptive statistics for our two illustrative datasets. At the auction level, we report the mean and standard deviation of: reserve prices; the lowest bid;²⁴ the lowest bid as a fraction of the reserve price; the number of bidders participating in the auction. At the bidder level, we report the mean and standard deviation of: the number of auctions a bidder participates in; the number of auctions she wins; and her total revenue.

We note that national-level auctions have higher reserve prices and a greater number of participants, but the two datasets are nonetheless broadly comparable. We also note that there is large heterogeneity in reserve prices in both cases, and in participation. Some projects are very large, and some bidders participate very often.

	N	Mean	S.D.
By Auctions			
reserve price (mil. Yen)	78,272	105.121	259.58
lowest first round bid (mil. Yen)	78,272	101.909	252.30
lowest bid / reserve	78,272	0.970	0.10
#bidders	78,272	9.883	2.27
By Bidders			
participation	29,670	26.40	94.61
number of times lowest bidder	29,670	2.64	10.57
total revenue (mil. Yen)	29,670	264.23	1312.77

Table A.1: Sample Statistics – National-Level Data

²⁴This corresponds to the winning bid when the reserve price is public. If the reserve price is secret, then the lowest bid need not be the winning bid.

	N	Mean	S.D.
By Auctions			
reserve price (mil. Yen)	1,469	26.363	65.054
lowest first round bid (mil. Yen)	1,469	24.749	61.745
lowest bid / reserve	1,469	0.936	0.078
#bidders	1,469	4.00	1.84
By Bidders			
participation	269	26.40	94.61
number of times lowest bidder	269	5.803	9.240
total revenue (mil. Yen)	269	167.342	447.014

Table A.2: Sample Statistics – City-Level Data

A.2 Total deviation temptation

In this section, we apply Proposition 3 to the objective function

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

This corresponds to firms' minimum deviation temptation as a share of expected profits. Figure A.1 reports estimates for firms in our city-level data.




Figure A.1: Estimated total deviation temptation as a fraction of profits, city-level data. Minimum and maximum markups: $m = 0.02$, $M = 0.5$.

B Connection with Bayes Correlated Equilibrium

In this section we further extend the class of estimators and tests introduced Section 5 and clarify what would need to be added so that they exploit all implications from equilibrium. This allows us to connect with the work of Bergemann and Morris (2016).

For simplicity we assume that player identities i , bids b and costs c take a fixed finite number of values $(i, b, c) \in I \times B \times C$ that does not grow with sample size. Ties between bids are resolved with uniform probability. Deviations $\rho_n \in (-1, \infty)$ correspond to the ratios of different bids on finite grid B .

We extend problem (P) as follows. As in Section 5.2, the objective function counts whether a history is competitive or not:

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

For any $(i, b, c) \in I \times B \times C$, let us define $H_{i,b,c}(\omega_A) \equiv \{h \in H \mid (i_h, b_h, c_h) = (i, b, c)\}$, histories at which bidder i experiences a cost c and bids b . Note that $H_{i,b,c}$ is adapted to the information of player i . For any tolerance function $T : \mathbb{N} \rightarrow \mathbb{R}^+$ such that

$$\lim_{N \rightarrow \infty} T(N) = 0 \quad \text{and} \quad \lim_{N \rightarrow \infty} \exp\left(-\frac{1}{2}T(N)^2 N\right) = 0$$

we consider inference problem (P'')

$$\begin{aligned} \widehat{U} &= \max_{\omega_A \in \Omega} U(\omega_A) & (P'') \\ \text{s.t. } \forall (i, b, c), \forall n, \quad D_n(\omega_A, H_{i,b,c}(\omega_A)) &\in \left[\widehat{D}_n(H_{i,b,c}(\omega_A)) - T(|H_{i,b,c}(\omega_A)|), \right. \\ &\quad \left. \widehat{D}_n(H_{i,b,c}(\omega_A)) + T(|H_{i,b,c}(\omega_A)|) \right]. \end{aligned}$$

Problem (P'') differs from (P') by imposing consistency conditions for all triples (i, b, c) . Proposition 3 continues to hold with an identical proof: with probability approaching 1 as $|H|$ goes to ∞ , estimate $\widehat{U}$ will approach 1 whenever all firms are competitive. Imposing consistency requirements conditional on bids and costs lets us establish a converse: if data passes our extended safe tests, then the joint distribution of bids and costs is an ϵ -Bayes correlated equilibrium in the sense of Hart and Mas-Colell (2000).

Consider an environment ω_A solving (P''). Let $\widehat{\mu} \in \Delta([B \times C]^I)$ denote the sample distribution over bids and costs implied by (H, ω_A) .

Proposition B.1. *For any $\epsilon > 0$, for $|H|$ large enough, $\widehat{U} = 1$ implies that $\widehat{\mu}$ is an ϵ -Bayes correlated equilibrium.*

Proof. Consider an environment $(d_{n,h}, c_h)_{h \in H}$ solving Problem (P''), and $\widehat{\mu}$ the corresponding sample distribution over profiles of bids b and costs c .

In order to deal with ties, we denote by $\wedge b_{-i} \succ b_i$ the event “ $\wedge b_{-i} > b_i$, or $\wedge b_{-i} = b_i$ and the tie is broken in favor of bidder i .”

For $|H|$ large enough, we have that for all (i, b, c) and all n ,

$$\frac{1}{|H|} \left| \sum_{h \in H_{i,b,c}} d_{n,h} - \widehat{\mu}(\wedge b_{-i} \succ (1 + \rho_n)b_i | i, b, c) \right| \leq \epsilon. \quad (7)$$

In addition, $\widehat{U} = 1$ implies that (IC) holds at all histories: for all h, n ,

$$d_{n,h}((1 + \rho_n)b_h - c_h) \leq d_{0,h}(b_h - c_h).$$

Summing over histories $h \in H_{i,b,c}$ yields

$$\frac{1}{|H|} \sum_{h \in H_{i,b,c}} d_{n,h}((1 + \rho_n)b_h - c_h) - d_{0,h}(b_h - c_h) \leq 0.$$

Hence for N large enough, for all (b_i, c_i) ,

$$\sum_{b_{-i}, c_{-i}} \widehat{\mu}(b_i, c_i, b_{-i}, c_{-i}) \left(\mathbf{1}_{\wedge b_{-i} \succ (1 + \rho_n)b_i} ((1 + \rho_n)b_i - c_i) - \mathbf{1}_{\wedge b_{-i} \succ b_i} (b_i - c_i) \right) \leq \epsilon.$$

It follows that $\widehat{\mu}$ is an ϵ -Bayes correlated equilibrium in the sense of Hart and Mas-Colell (2000). ■

C Common Values

We now show how to extend the analysis in Section 5.2 to allow for common values. An environment in auction a is now given by $\omega_a = (d_{h,n}, c_{h,n})_{h \in a, n \in \mathcal{M}}$, where for each $n \in \mathcal{M}$, $c_{h,n} = \mathbb{E}[c | h, \wedge \mathbf{b}_{-i,h} > (1 + \rho_n)b_h]$ is the bidder's expected cost at history h conditional on winning at bid $(1 + \rho_n)b_h$.

We make the following assumption on costs.

Assumption C.1. For all histories h and all bids b, b', b'' with $b < b' < b''$, $\mathbb{E}[c|h, \wedge \mathbf{b}_{-i,h} \in (b, b')] \leq \mathbb{E}[c|h, \wedge \mathbf{b}_{-i,h} \in (b', b'')]$.

In words, bidders' expected costs are increasing in opponents' bids.²⁵ This implies that expected costs $c_{h,n}$ are weakly increasing in the deviation ρ_n . We now show that, under these conditions, allowing for common values does not relax the constraints in Program (P').

Note first that, for each deviation n , expected costs $(c_{h,n})_{n \in \mathcal{M}}$ satisfy:

$$\forall n \in \mathcal{M}, \quad d_{h,n}c_{h,n} = d_{h,0}c_{h,0} + (d_{h,n} - d_{h,0})\hat{c}_{h,n}, \quad (8)$$

where $\hat{c}_{h,n} = \mathbb{E}[c|h, \wedge \mathbf{b}_{-i,h} \in (b_h, (1 + \rho_n)b_h)]$.²⁶ Our assumptions on costs imply that $\hat{c}_{h,n}$ is weakly increasing in n .

Consider first downward deviations $\rho_n < 0$ (i.e., $n < 0$). Note that, for such deviations, incentive compatibility constraint (IC) holds if and only if

$$\frac{d_{h,n}(1 + \rho_n)b_{h,n} - d_{h,0}b_{h,0}}{d_{h,n} - d_{h,0}} \leq \hat{c}_{h,n}.$$

Consider upward deviations $\rho_n > 0$ (i.e., $n > 0$). For any such deviation, constraint (IC) becomes

$$\hat{c}_{h,n} \leq \frac{d_{h,0}b_{h,0} - d_{h,n}(1 + \rho_n)b_{h,n}}{d_{h,0} - d_{h,n}}.$$

Since $\hat{c}_{h,\hat{n}}$ is weakly increasing in $\hat{n}$, $\hat{c}_{h,n} \geq \hat{c}_{h,n'}$ for all $n > 0$ and $n' < 0$. Hence there exist costs $(c_{h,n})_{n \in \mathcal{M}}$ satisfying (IC) if and only if

$$\min_{n < 0} \frac{d_{h,n}(1 + \rho_n)b_{h,n} - d_{h,0}b_{h,0}}{d_{h,n} - d_{h,0}} \leq \max_{n > 0} \frac{d_{h,0}b_{h,0} - d_{h,n}(1 + \rho_n)b_{h,n}}{d_{h,0} - d_{h,n}}.$$

Condition (C) implies that there also exists a constant profile of costs $c_{h,n} = c_h$ (i.e. a private value cost), that satisfies (IC).

D Computational Strategy

In this appendix, we discuss computational implementations of the estimates of competitiveness derived in Sections 4 and 5.

²⁵This condition on costs would follow from affiliation if bidders' signals are one-dimensional and bidders use monotone bidding strategies.

²⁶We use the convention that (b, b') is equal to (b', b) in the event that $b' < b$, rather than the empty set.

D.1 Bounds of Section 4

Column 3 of Table 1 reports 95% confidence upper bounds on the share of competitive histories for individual firms using Proposition 2:

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

We take ρ equal to $\{0.01\%, 0.02\%, \dots, 0.3\%\}$ of the reserve price and fix this set in our asymptotics. We estimate the counterfactual revenue using a triangular kernel with bandwidth equal to 0.01% of the reserve price. We use the central limit theorem for Martingale difference sequences to obtain the 95% confidence bound, see e.g., Liu and Yang (2008).

D.2 Bounds of Section 5

Problem (P) is not naturally suited for computational implementation. A solution ω_A is a real vector of dimension $|H| \times (|\mathcal{M}| + 1)$. In addition, objective functions U of interest are typically non-concave, and the set of admissible possible environments Ω need not be convex.

For this reason we study an approximate convexified problem, in the case where objective function $u(\omega_a)$ is separable in histories, i.e. takes the form $u(\omega_a) = \frac{1}{|H|} \sum_{h \in a} v(\omega_h)$, where $\omega_h = (d_{h,n}, c_h)_{n \in \mathcal{M}}$.

In this case, the constrained sets Ω of interest take the form $\Omega = O^H$, where $O \subset \mathbb{R}^{|\mathcal{M}|+1}$. In addition, notice that $\widehat{U}(\omega_A) = \mathbb{E}_{\mu(\omega_A)}[v(\omega_h)]$ where $\mu_{\omega_A} \in \Delta(O)$ is the sample distribution of environments $\omega_h \in O$ induced by profile $\omega_A = (\omega_a)_{a \in A}$ (recall that $\omega_a = (\omega_h)_{h \in a}$).

Hence, it follows that

$$\begin{aligned} \widehat{U} &\leq \max_{\mu \in \Delta(O)} \mathbb{E}_{\mu} v(\omega_h) \\ &\quad \text{under the constraint that} \\ &\quad \mathbb{E}_{\mu} d_{n,h} \in [\widehat{D}_n - T(|H|), \widehat{D}_n + T(|H|)]. \end{aligned}$$

The convexified right-hand side problem is linear, and the dimensionality of a solution is no longer related to sample size $|H|$. The difficulty is that, $\Delta(O)$ is infinite dimensional. Still, because O is compact and finite dimensional it is covered by finitely many balls of radius r

for any $r > 0$. Hence for all N , there exists a finite set $O' \subset O$ such that

$$\widehat{U} - \frac{1}{N} \leq \max_{\mu \in \Delta(O')} |H| \times \mathbb{E}_\mu v(\omega_h) \quad (\text{CVX-P})$$

under the constraint that

$$\mathbb{E}_\mu d_{n,h} \in [\widehat{D}_n - T(|H|), \widehat{D}_n + T(|H|)].$$

The right-hand-side problem is now a well behaved finite dimensional linear problem. Practically we use the following algorithm, which starts with an empty seed $O_0^* = \emptyset$. At each stage $T \geq 1$ of the algorithm, we perform the following operations:

1. Draw a sample S_T of K points in O using a full-support distribution over O .
2. Set $O_T = O_{T-1}^* \cup S_T$, and solve (CVX-P) for O_T ; get solution μ_T .
3. Set $O_T^* = \text{supp } \mu_T$.
4. Iterate until convergence.

E Proofs

E.1 Proofs of Section 3

Proof of Lemma 1. Let H be an adapted 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 $s \geq 0$, $\mathbb{E}_\sigma[\varepsilon_{i,t} | h_{j,t-s}] = \mathbb{E}_\sigma[\mathbb{E}_\sigma[\mathbf{1}_{\wedge b_{-i,t} > b_{i,t}(1+\rho)} | h_t, z_{i,t}] - \mathbf{1}_{\wedge b_{-i,t} > b_{i,t}(1+\rho)} | h_{t-s}, z_{j,t-s}] = 0$.²⁷

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 $\varepsilon_k = \varepsilon_{i,t}$,

²⁷This 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 .

so that

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

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 $\alpha > 0$, $\text{prob}(|S_{|H|}| \geq |H|\alpha) \leq 2\exp\{-\alpha^2|H|/2\}$. Therefore, with probability 1, $\frac{1}{|H|}S_{|H|} = \hat{D}(\rho|H) - \overline{D}(\rho|H)$ converges to zero as $|H| \rightarrow \infty$. ■

E.2 Proofs of Section 4

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

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

for each deviation n . It then follows that

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

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

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 $H_{\text{comp}}(\omega_A) = H$, 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 $1 - 2|\mathcal{M}|\exp(-T(|H|)^2|H|/2)$, $\hat{U} \geq s_{\text{comp}} = 1 \geq s_0$. Since $2|\mathcal{M}|\exp(-T(|H|)^2|H|/2) \rightarrow 0$ as $|H| \rightarrow \infty$, firms in this industry pass test τ with probability approaching 1 as $|H| \rightarrow \infty$. ■

E.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 continuation payoff cannot be lower than 0 at any history.

Let $\bar{K} = \frac{\delta}{1-\delta}r = \frac{\delta}{1-\delta}$ (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 public 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 $\frac{\bar{K}}{K} + \epsilon < 1$. Hence, under history $h_t \sqcup h_s$, firm i 's continuation payoff at the end of period $t + s$ (i.e., after the outcome of the auction is determined) is bounded above by $\frac{\delta}{1-\delta} - \left(\frac{\bar{K}}{K} + \epsilon\right)K = -\epsilon K < 0$, 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 $\hat{\sigma}$ at history $h_{i,t}$. Firm i 's total payoff under strategy profile $\hat{\sigma}$ given history $h_{i,t}$ is $V_i(\hat{\sigma}, h_{i,t}) - \mathbb{E}_{\hat{\sigma}}[\tau_i | h_{i,t}]K$.

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:²⁸ 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 $\sigma \in \Sigma(K)$, all 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 . Hence, by the one-shot deviation principle $\sigma \in \Sigma(0)$.²⁹ ■

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

Lemma E.1. *Let σ be a strategy profile with the property that all firms pass the test with*

²⁸Note 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.

²⁹Note that the game with $K = 0$ is continuous at infinity, and so the one-shot deviation principle holds.

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 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 E.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 $\bar{K} \equiv \frac{\delta}{1-\delta}$. 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 $\frac{\bar{K}}{K} + \epsilon < 1$. Hence, under public history $h_t \sqcup h_s$, each firm's continuation payoff at the end of period $t+s$ (i.e., after the outcome of the auction is determined) is bounded above by $\frac{\delta}{1-\delta} - \left(\frac{\bar{K}}{K} + \epsilon\right)K = -\epsilon K < 0$, 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 $h_{i,t} \neq h_{i,\tau}$,

$$\begin{aligned} V_i((\tilde{\sigma}_i, \sigma_{-i}), h_{i,\tau}) - \mathbb{E}_{(\tilde{\sigma}_i, \sigma_{-i})}[\tau_i | h_{i,\tau}]K &\leq V_i(\sigma, h_{i,\tau}) - \mathbb{E}_\sigma[\tau_i | h_{i,\tau}]K \\ \iff V_i((\tilde{\sigma}_i, \sigma_{-i}), h_{i,\tau}) &\leq V_i(\sigma, h_{i,\tau}), \end{aligned} \tag{10}$$

where the second line in (10) follows since, under equilibrium σ , firms pass the test with

probability 1 at every history. By the second line in (10), 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, $\sigma \in \Sigma_{RP}^+(0)$. ■

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.
- 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.
- (2016): “Bayes correlated equilibrium and the comparison of information structures in games,” *Theoretical Economics*, 11, 487–522.
- 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 (2018): “Weak cartels and collusion-proof auctions,” *Journal of Economic Theory*, 178, 398–435.
- CHE, Y.-K. AND J. KIM (2006): “Robustly Collusion-Proof Implementation,” *Econometrica*, 74, 1063–1107.
- CLARK, R. AND J.-F. HOUE (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.
- DOJ, U. (2005): “Price fixing, bid rigging, and market allocation schemes: What they are and what to look for. An Antitrust Primer.” 2005.
- DOVAL, L. AND J. ELY (2019): “Sequential information design,” Tech. rep., Working Paper.
- 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.
- HART, S. AND A. MAS-COLELL (2000): “A simple adaptive procedure leading to correlated equilibrium,” *Econometrica*, 68, 1127–1150.
- 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.
- LIU, R. AND L. YANG (2008): “Kernel estimation of multivariate cumulative distribution function,” *Journal of Nonparametric Statistics*, 20, 661–677.
- 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.

- SCHURTER, K. (2017): “Identification and inference in first-price auctions with collusion,” Tech. rep., working Paper, University of Chicago.
- 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.