Screening Adaptive Cartels

Juan Ortner
Boston University

Sylvain Chassang
Princeton University

Kei Kawai
U.C. Berkeley

Jun Nakabayashi*†
Kindai University

November 19, 2022

Abstract

We propose an equilibrium theory of data-driven antitrust oversight in which regulators launch investigations on the basis of suspicious bidding patterns and cartels can adapt to the statistical screens used by regulators. We emphasize the use of asymptotically safe tests, i.e. tests that are passed with probability approaching one by competitive firms, regardless of the underlying economic environment. Our main result establishes that screening for collusion with safe tests is a robust improvement over laissez-faire. Safe tests do not create new collusive equilibria, and do not hurt competitive industries. In addition, safe tests can have strict bite, including unraveling all collusive equilibria in some settings. We provide evidence that cartel adaptation to regulatory oversight is a real concern.

Keywords: collusion, auctions, bidding rings, cartels, procurement, antitrust.

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

†Some of this work was previously circulated as NBER Working Paper 25654, entitled "Data Driven Regulation: Theory and Application to Missing Bids". We thank John Asker, Ben Brooks, Joe Harrington, Elliot Lipnowski, David McAdams, Michael Riordan, Takuo Sugaya, Alex Wolitzky, as well as seminar audiences at the 2021 BFI Economic Theory Conference, Boston University, Gerzensee, Harvard-MIT, the 2022 IO Theory Conference, HKUST, NYU, Princeton, Seminars in Economic Theory, Stanford, and Yale for helpful comments and suggestions.

1 Introduction

Competition authorities commonly rely on statistical screens to detect and investigate collusion between firms.[1] Even if formal prosecution cannot rely on statistical evidence alone, detection tools can greatly facilitate the work of competition authorities. Such evidence can be used in court to obtain warrants or authorization for a more intrusive investigation, ultimately leading to actionable evidence and convictions (see Imhof et al., 2018, for a concrete example).[2] Statistical evidence may be helpful in convincing cartel members to apply to leniency programs. Stakeholders other than antitrust authorities may also find statistical screens useful. For example, plaintiffs in civil antitrust complaints can use the results of statistical screens to provide sufficient grounds to state a claim under the Sherman Act.[3] [4] However, this growth in the use of statistical screens raises numerous questions: Do firms adapt to the statistical tests implemented by competition authorities? If so, what is the impact of such screens in equilibrium? Can the tests backfire and either strengthen cartels, or hurt competitive firms? Can we find tests that cause no harm, and yet reduce the incentives to form cartels? We provide theory and evidence addressing these questions.

We propose a model of collusion in the shadow of an antitrust authority. A group of firms repeatedly participates in a first-price procurement auction. We allow firms to observe arbitrary signals about one another, and allow bidders' costs to be correlated within periods.[5]

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

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

[3] Federal rules of civil procedure in the U.S. requires a statement showing that the plaintiff is entitled to relief. Failure to provide grounds of entitlement to relief can result in dismissal of the complaint.

[4] Screens for cartels have a wide range of other uses as well. For example, screening can help procurement offices counter suspected bidding rings by more aggressively soliciting new bidders or adopting auction mechanisms that are less susceptible to collusion. Screening may also be helpful for internal auditors and compliance officers of complicit firms to identify collusion and help contain potential legal risks arising from compliance failures.

[5] Our results extend as is if costs are correlated across periods via a publicly observed exogenous time-

At a finite date T, the antitrust authority observes the history of bids placed by firms, and runs a test to screen for noncompetitive behavior. Firms that don't pass the test are further investigated, and may incur penalties if found guilty of bid-rigging. Investigation may also be costly to non-cartel members. We say that the tests used by the antitrust authority are asymptotically safe if and only if competitive bidders pass with probability approaching one as data size becomes large under all environments in the support of beliefs. Our companion papers Chassang et al. (2022) and Kawai et al. (2022b) develop asymptotically safe tests and illustrate their relevance by applying them to procurement data from Japan and the US.

Our main set of results shows that antitrust oversight based on asymptotically safe tests is a robust improvement over laissez-faire. First, we show that regulation based on firm-level asymptotically safe tests does not significantly expand the set of enforceable collusive schemes available to cartels. Hence, asymptotically safe tests don't increase firms' incentives to form a cartel. This addresses a concern raised by Cyrenne (1999) and Harrington (2004) that some natural screens against collusion may backfire and enhance the ability of cartels to collude. Second, we establish by example that asymptotically safe tests can have strict bite. In complete information settings, optimally colluding bidders submit nearly tied bids. Testing for an excessive mass of close bids yields an asymptotically safe test whenever firms face strictly positive bid preparation costs. This safe test causes collusive equilibria to unravel for an open range of discount factors.

Additionally, consistent with the work of Wollmann (2019) and Cunningham et al. (2021), we provide suggestive evidence that cartels do in fact adapt to regulatory screens. One common screen for bid rigging used by regulators is to test for the presence of tied or nearly tied bids.[6] If regulators scrutinize auctions with close bids, a cartel seeking to avoid scrutiny may adapt by instructing its members to refrain from placing close bids. This adaptive response may then lead to "missing bids" in large datasets, where instead of being

varying state.

[6]Variants of this screen include screens proposed by Abrantes-Metz et al. (2005) and Imhof et al. (2016) that flag auctions with low bid dispersion.

excessively frequent, adaptation would lead close bids to become excessively rare. We show that this pattern is present in procurement auctions from Japan. In addition, as we argue in Chassang et al. (2022), testing for missing bids also constitutes a safe test: this bidding pattern cannot be explained by competitive behavior.[7]

Our paper relates to the academic literature on statistical tests of non-competitive behavior.[8] Screens for collusive bidding patterns include those based on the correlation among bids (Bajari and Ye, 2003), variance of bids (Abrantes-Metz et al., 2005), correlation between bids and costs (Porter and Zona, 1993, 1999, Kawai et al., 2022b) or predictions based on the theory of repeated games, such as price wars and other patterns (Porter, 1983, Ellison, 1994, Chassang and Ortner, 2019). Statistical tests of collusion have also been developed for average-price auctions (Conley and Decarolis, 2016) and multi-stage auctions with rebidding (Kawai and Nakabayashi, 2018). Our paper complements this literature by studying the impact of statistical screening in equilibrium when firms can adapt to the regulatory environment, and identifying the role of safe-tests as a robust improvement over laissez-faire.

Previous work that has considered the equilibrium impact of antitrust oversight include Besanko and Spulber (1989), LaCasse (1995), Cyrenne (1999) and Harrington (2004). Besanko and Spulber (1989) and LaCasse (1995) study static models of equilibrium regulation. Closer to our work, Cyrenne (1999) and Harrington (2004) study repeated oligopoly models in which colluding firms might get investigated and fined whenever prices exhibit large and rapid fluctuations. Both papers highlight that antitrust oversight may backfire, allowing cartels to sustain higher equilibrium profits. Intuitively, cartels may use the threat of a regulatory crackdown to discipline their members.[9] We provide evidence that concerns about adaptive cartels are valid, but that they can be addressed by using safe tests.

[7] In Chassang et al. (2022) we also show that missing bids are correlated with plausible markers of collusion: missing bids are more prevalent in auctions with high winning bids, and, for industries that are investigated for bid-rigging, they are more prevalent before the investigation than after.

[8] See Porter (2005) and Harrington (2008) for recent surveys of this literature.

[9] See also McCutcheon (1997), who shows that anti-trust oversight may help sustain collusion by reducing firms' incentives to renegotiate equilibrium play after a deviation.

Our work complements the literature on auction design in the presence of collusion. Abdulkadiroglu and Chung (2003), Che and Kim (2006, 2009) and Pavlov (2008) show that appropriate auction design can limit the cost of collusion when cartel members have deep pockets and can make payments upfront.[10] Che et al. (2018) studies optimal auction design when collusive bidders are cash-constrained. Our paper complements this literature by showing how an antitrust agency can limit the impact of collusion by screening firms using safe tests.

In addition, the notion of safe tests may have practical use for civil antitrust claims in light of two recent U.S. Supreme Court decisions, Bell Atlantic v. Twombly (2007) and Ashcroft v. Iqbal (2009). In Bell Atlantic v. Twombly (2007), the United States Supreme Court ruled that circumstantial evidence of conduct consistent with collusion, such as suspiciously low entry, or simultaneous price increases, is no longer sufficient grounds to hear cases, leading to the dismissal of antitrust cases that would have been heard prior to 2007 under the precedent set by Conley v. Gibson (1957). Because failures of safe tests cannot be explained by innocuous competitive dynamics, they are likely to provide just the type of facts needed to sustain an antitrust complaint under the more stringent requirements.

Lastly, our focus on safe tests is largely motivated by the importance that courts place on the chilling effects of the law (e.g., Kaplow, 2011). Our results show that safe tests have other desirable properties, in addition to not generating undesirable chilling effects on competitive behavior.

The paper is structured as follows. Section 2 sets up our model of collusion in the shadow of investigation. Section 3 presents an example in which a naïve test for collusion ends up strengthening cartels. Section 4 introduces safe tests. Section 5 establishes that safe tests do not create new collusive equilibria and can strictly reduce the payoffs of cartels. Section 6 discusses the policy relevance of our findings by providing evidence that cartels do in fact

[10] See also Marshall and Marx (2009), who study the vulnerability of different auction formats to bidder collusion.

adapt to statistical screens of collusion, and suggesting ways in which known tests of collusion can be made safer. Proofs are collected in Appendix A.

2 Colluding in the Shadow of Antitrust Authorities

We model the interaction between cartel members and antitrust authorities as follows. At each period $t \in \mathbb{N}$, firms participate in a procurement auction. At some fixed time $T \in \mathbb{N}$, the regulator applies tests to the data generated by the players in periods $t \leq T$. If a test comes out against the null hypothesis of competition, one or more firms are investigated. We begin by describing the stage game of the repeated game.

2.1 Repeated Procurement

In each period $t \in \mathbb{N}$, a buyer needs to procure a single project from a finite set $N = \{1, ..., n\}$ of potential suppliers, with $n \geq 2$. The auction format is a sealed-bid first-price auction with reserve price r, which we normalize to $r = 1$. Let $B \subset [0, 1]$ be the set of feasible bids. We assume throughout that $B = [0, 1]$, except in Proposition 2, where we assume B is a finite grid $B = \{0, \nu, 2\nu, ..., 1\}$, with $\nu > 0$ small.

Costs. Firms' period t procurement costs are denoted by $\mathbf{c}_t = (c_{i,t})_{i \in N}$. For simplicity, we assume private values, so that each bidder observes her own cost. Costs are allowed to be correlated across firms within each period, but for simplicity are assumed to be i.i.d. over time. At each time t, $\mathbf{c}_t$ is drawn from distribution $F_C(\cdot)$ supported on set $C \subset [0, 1]^n$.[11]

Information. In each period t, each bidder $i \in N$ privately observes a signal $z_{i,t}$ prior to bidding. Signals $z_{i,t}$ can take arbitrary values, including vectors in $\mathbb{R}^k$. The distribution of signals $\mathbf{z}_t = (z_{i,t})_{i \in N}$ depends only on realized costs $\mathbf{c}_t$: $\mathbf{z}_t$ is drawn from distribution $F_Z(\cdot|\mathbf{c}_t)$,

[11] Our analysis remains unchanged if we allow costs to depend on a public time-varying exogenous state, as in Chassang et al. (2022).

with support contained in some set Z. Signals $(z_{i,t})_{i \in N}$ are allowed to be arbitrary, and may include information about the costs of other bidders. This allows our model to nest many informational environments, including correlated private values, asymmetric bidders, as well as complete information. Since bidders observe their own costs prior to bidding, we assume that firm i's signal $z_{i,t}$ includes firm i's cost $c_{i,t}$.

Bids. After privately observing signal $z_{i,t}$, each firm $i \in N$ submits a bid $b_{i,t} \in B \cup \emptyset$, where $\emptyset$ denotes not participating. The procurement contract is allocated to the bidder submitting the lowest bid in B, at a price equal to her bid. Ties are broken randomly. We assume that each bidder $i \in N$ incurs a bid preparation cost $\kappa_i \geq 0$ from submitting a bid in B. Profiles of bids are denoted by $\mathbf{b}_t = (b_{i,t})_{i \in N}$, with $\wedge \mathbf{b}_t$ denoting the lowest bid. We let $\mathbf{b}_{-i,t} \equiv (b_{j,t})_{j \neq i}$ denote bids from firms other than firm i, and define $\wedge \mathbf{b}_{-i,t} \equiv \min_{j \neq i} b_{j,t}$ to be the lowest bid among i's competitors.

We assume that bids are publicly revealed at the end of each period. This matches standard practice in public procurement, where legislation typically requires governments to make bids public. Our main results can be adapted if only the winning bid is made public, or if bidders only observe the identity of the winner.

Overall, firm i's profits in period t are

$$\pi_{i,t} = x_{i,t} \times (b_{i,t} - c_{i,t}) - \kappa_i \mathbf{1}_{b_{i,t} \neq \emptyset},$$

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

2.2 Antitrust Oversight

We assume that the antitrust authority runs screening tests in period T, based on data from the M periods leading up to T.[12]

Fix $T \in \mathbb{N}$ and $M \in \mathbb{N}$, $M \leq T+1$. Let $h_{M,T} = (\mathbf{b}_s)_{s=T+1-M}^{T}$ denote the bids placed during the monitoring phase $t = T+1-M, ..., T$. At the end of period T, after players place their bids and the auction outcome is realized, the antitrust authority runs a vector of tests $(\tau_i)_{i \in N}$, with $\tau_i : h_{M,T} \mapsto \tau_i(h_{M,T}) \in \{0,1\}$.[13] If test τ_i takes value 1, firm i is investigated. This yields an expected penalty $K_T \geq 0$, paid at period T. We allow penalty K_T to grow as the testing date T grows large, for instance, $K_T = \delta^{-T} K$ for some $K > 0$. Hence, the impact of antitrust oversight on firms' payoffs at the start of the game may remain bounded away from zero even as T grows large.

For simplicity, we consider fixed penalties. However, we note that all of our results would continue to hold if penalty K_T was allowed to depend on bidding history $h_{M,T}$. This would allow for penalties that depend on the extent, or impact, of collusion.

Aggregate payoffs to firm i from the perspective of period 0 are as follows:

$$(1-\delta)\left[\sum_{t=0}^{\infty} \delta^t \pi_{i,t}\right] - \tau_i \delta^T K_T.$$

Solution concept. The period-t public history h_t takes the form $h_t = (\mathbf{b}_s)_{s<t}$. Because costs are drawn i.i.d. across periods, past play conveys no information about the private types of other players. As a result we do not need to specify out-of-equilibrium beliefs. A public strategy σ_i is a mapping

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

[12]Our main results continue to hold if the testing date is random and exponentially distributed, or if testing is performed every T periods.

[13]More generally, $h_{M,T}$ may include any data observable to the antitrust authority at the time of running the test.

Under a public strategy, firm i's bid at each time t depends on public history h_t and current signal realization $z_{i,t}$. We focus on perfect public Bayesian equilibria (Athey and Bagwell, 2008, henceforth PPBE); i.e., perfect Bayesian equilibria in public strategies.

3 A Motivating Example

In prior work, Cyrenne (1999) and Harrington (2004) showed that regulatory oversight may backfire, allowing cartels to sustain larger profits. We now illustrate this possibility in the context of our model.

Consider a special instance of our model in which procurement costs are publicly observed among bidders and equal to zero: $\forall i, t$, $c_{i,t} = 0$. Assume also that there are no participation costs: $\forall i$, $\kappa_i = 0$.

Fix $\eta \in (0, 1)$. Suppose the regulator runs the tests at period T, with monitoring length $M = T+1$. Hence, the outcome of the test depends on bidding behavior at periods $t = 0, .., T$. Suppose further that the regulator runs the same test for all firms in the industry: for all $i \in N$, $\tau_i = \tau^{\mathsf{break}}$ with τ^{break} defined by

$$\tau^{\mathsf{break}} \equiv \mathbf{1}\left\{\exists S \leq T \text{ s.t. } |\wedge \mathbf{b}_t - \wedge \mathbf{b}_s| > \eta \text{ for all } t < S, s \in [S, T]\right\}.$$

Test τ^{break} looks for structural breaks in bidding behavior: firms fail the test if there is a discrete jump in the winning bid. Testing for structural breaks is a common way of screening for cartels (e.g., Harrington, 2008): failure to pass test τ^{break} might indicate either that a cartel was formed or that a cartel collapsed.

Assume $\eta < \frac{1}{n}$ and $\delta < \frac{n-1}{n}$. Consider first the case with no regulator; i.e. $K_T = 0$. Since $\delta < \frac{n-1}{n}$, firms are unable to sustain supra-competitive prices. In any equilibrium the winning bid must be equal to zero at all periods, $\wedge \mathbf{b}_t = 0$ for all t.

Consider next the case with a regulator. When K_T is sufficiently large, each player $i \in N$

playing according to the following strategy constitutes an equilibrium of the regulatory game:

- At the initial history h_0, or at any history h_t with $t < T$ and with $b_{j,s} = 1$ for all $j \in N, s < t$, bid $b_{i,t} = 1$;
- At history h_T with $b_{j,s} = 1$ for all $j \in N, s < T$, bid $b_{i,t} = 1 - \eta$;
- At any other history h_t, bid $b_{i,t} = 0$.

Intuitively, play reverts to static Nash following a defection prior to period T, resulting in firms failing the test. When penalty K_T is sufficiently large, the loss from failing the test outweighs any deviation gain, allowing bidders to sustain a collusive equilibrium.[14]

4 Safe Tests

The previous example illustrates that screening for collusion may inadvertently strengthen cartels. This section introduces a class of tests, which we call *asymptotically safe tests.* In words, asymptotically safe tests are tests that are passed with probability close to one by unilaterally competitive firms. Our main results, presented in Section 5, show that preventing harm against competitive firms serves both a direct purpose – it has a vanishingly small impact on firms operating in a competitive market – and an indirect one – it ensures that regulatory screens do not increase a cartel's enforcement capability. Altogether this implies that safe tests do not increase firms' incentives to cartelize. In addition, we show that safe tests can *strictly* reduce the gains from collusion by constraining the equilibrium play of cartel members. In this sense, safe tests are a robust improvement over laissez-faire.

Before introducing safe tests, we need to define what we mean by competitive firms. Following Chassang et al. (2022), we say that a firm is competitive if and only if it plays a stage-game best response at every history on the equilibrium path. Recall that T is the

[14]Note that each firm i can deviate and bid $\eta - \epsilon$ at period 0, without triggering the test. The assumption that $\eta < \frac{1}{n}$ guarantees that such a deviation is not profitable.

period at which the tests are run. Periods $t = T+1-M, ..., T$ correspond to the monitoring phase, and $h_{M,T} = (\mathbf{b}_s)_{s=T+1-M}^{T}$ are the bids placed during the monitoring phase.

Definition 1 (competitive histories)**.** *Fix a common knowledge profile of play σ and a history $h_{i,t}$ of player i. Firm 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}.*

Firm i is competitive during the monitoring phase if it plays competitively at all on-path histories $h_{i,t}$ with $t = T+1-M, ..., T$.

Firm i is competitive if it plays competitively at all histories on the equilibrium path.

The industry is competitive if all firms $i \in N$ play competitively at all histories on the equilibrium path.

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

Asymptotically safe tests. The economic environment corresponds to the tuple $E = (F_C, F_Z, (\kappa_i)_{i\in N})$. We denote by $\mathcal{E}$ the set of environments E that the antitrust authority deems feasible. Set $\mathcal{E}$ captures the subjective restrictions that the antitrust authority is willing to place.

Definition 2 (asymptotically safe tests)**.** *Tests $(\tau_i)_{i\in N}$ are asymptotically safe if and only if for all $i \in N$, all $E \in \mathcal{E}$, and all profiles σ such that firm i is competitive during the monitoring phase, $\textsf{prob}_{E,\sigma}(\tau_i = 0) \geq 1 - G(M)$ for some $G(\cdot) \geq 0$ with $\lim_{S\to\infty} G(S) = 0$.*

In words, tests $(\tau_i)_{i\in N}$ are asymptotically safe if they admit a vanishingly small rate of false positives. This concern over false positives coincides with views expressed by regulators (Imhof et al., 2016). In practice, investigation is a highly disruptive process that is only triggered if sufficient evidence is available.[15] Our companion papers Chassang et al.

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

(2022) and Kawai et al. (2022b) propose various asymptotically safe tests and illustrate their relevance by applying them to Japanese procurement data, as well as data from the Ohio school milk cartel (Porter and Zona, 1999). We emphasize that asymptotically safe tests pass competitive firms with high probability even if other firms do not behave competitively.

Because safe-tests are by nature conservative — they cannot be failed by competitive firms — we believe that failures of safe tests constitute hard-to-dismiss evidence of non-competitive conduct. In this sense, failure of safe tests meet the new requirements set by Bell Atlantic v. Twombly (2007) needed for sustaining federal antitrust cases.

The function $G(\cdot)$ in Definition 2 bounds the rate of false positives of tests $(\tau_i)_{i\in N}$. In Chassang et al. (2022) we propose a family of asymptotically safe tests based on estimated demand, and show that tests in this family have a rate of false positives that is exponentially decreasing in the number of observations: i.e., those tests are asymptotically safe with $G(M) = \beta \exp(-\alpha M)$ for some constants $\alpha, \beta > 0$. In Section 5.3 we present an asymptotically safe test that also has an exponentially decreasing rate of false positives.

Before we study the impact of safe tests on incentives to collude, the following attractive property is worth noting.

Remark 1. *If τ_i and τ_i' are asymptotically safe, then $\max\{\tau_i, \tau_i'\}$ is asymptotically safe. If τ_i is asymptotically safe, then for any test τ_i', $\min\{\tau_i, \tau_i'\}$ is asymptotically safe.*

This means that regulators can combine safe tests with both safe and unsafe tests and still maintain safety. In particular, the combination of two safe tests is safe. In addition, performing a quick and possibly unsafe test τ_i', followed by a safe test τ_i if firm i fails the first test, is also safe.

5 Safe Tests do not Increase Incentives to Cartelize

This section provides normative foundations for safe tests. We proceed in three steps: first, we show that safe tests do not significantly expand the set of (potentially collusive) equilibria;

second, we show that safe tests do not significantly affect competitive industries; third, we show by example that safe tests can have strict bite, including the unraveling of all collusive equilibria.

5.1 Safe Tests do not Create New Collusive Equilibria

We start by showing that asymptotically safe tests do not significantly enlarge the set of equilibria. We formalize this result with two propositions. Our first proposition establishes a bound on the equilibrium payoff set of the game with a regulator that holds for any testing date T and monitoring length M. Our second proposition shows that when testing date T and monitoring length M both grow large, any equilibrium of the game with the regulator must be "close" to an equilibrium of the game without the regulator.

Bounding the equilibrium payoff set. For each environment $E \in \mathcal{E}$, and integers T and $M \leq T+1$, let $\Sigma_{T,M}(E)$ denote the set of PPBE of the game with testing date T and monitoring length M. Let $\Sigma(E)$ denote the set of PPBE of the game without the regulator.

Fix an environment $E \in \mathcal{E}$, a testing date T and a monitoring length $M \leq T+1$. For each strategy profile σ and each history $h_{i,t}$, let $V_i^E(\sigma, h_{i,t})$ denote player i's expected discounted payoff (excluding possible penalties) at history $h_{i,t}$ under strategy profile σ and environment E. Similarly, let $P_i^E(\sigma, h_{i,t}) = \mathbf{1}_{t \leq T} \mathbb{E}_{E,\sigma}\left[\tau_i \delta^{T-t} K_T | h_{i,t}\right]$ denote player i's expected penalty at history $h_{i,t}$ under strategy profile σ and environment E. Player i's total payoff in the regulatory game is $W_i^E(\sigma, h_{i,t}) = V_i^E(\sigma, h_{i,t}) - P_i^E(\sigma, h_{i,t})$.

For T and M, define

$$\mathcal{V}_{T,M}(E) \equiv \left\{ \mathbf{V} \in \mathbb{R}^n : \mathbf{V} = \left(\mathbb{E}_E[W_i^E(\sigma, (h_0, z_i))]\right)_{i \in N} \text{ for some } \sigma \in \Sigma_{T,M}(E) \right\},$$

to be the set of period 0 payoffs that can be supported in an equilibrium with testing date

T and monitoring length M under environment E. Similarly, let

$$\mathcal{V}(E) \equiv \left\{ \mathbf{V} \in \mathbb{R}^n : \mathbf{V} = \left(\mathbb{E}_E[V_i^E(\sigma, (h_0, z_i))] \right)_{i \in N} \text{ for some } \sigma \in \Sigma(E) \right\}$$

denote the period 0 equilibrium payoff set of the game without a regulator.

Let $X = B \cup \emptyset$ denote the set of bids and fix $\mathcal{W} \subset \mathbb{R}^n$ a set of possible continuation values. Following Abreu et al. (1990) (henceforth APS), we say that a profile of bidding functions $\beta = (\beta_i)_{i \in N}$, with $\beta_i : z_i \mapsto \beta_i(z_i) \in \Delta(X)$ for each $i \in N$, is enforceable on $\mathcal{W}$ if there exists $W : X^n \to \mathcal{W}$ such that, for all $i \in N$, for all z_i, and all $b_i \in \operatorname{supp} \beta_i(z_i)$,

$$b_i \in \arg\max_b \mathbb{E}_{E,\beta}[(1-\delta)(b - c_i)x_i + \delta W_i(b, \mathbf{b}_{-i}) | z_i],$$

where $x_i \in \{0, 1\}$ denotes whether or not i wins the auction.

Fix a profile of bidding functions β enforced by W on $\mathcal{W}$. For each $i \in N$, define

$$v_i^E(\beta, W) \equiv \mathbb{E}_{E,\beta}[(1-\delta)(b_i - c_i)x_i + \delta W_i(b_i, b_{-i})].$$

For any $\mathcal{W} \subset \mathbb{R}^n$, the APS value-set operator is given by

$$\mathcal{B}_E(\mathcal{W}) \equiv \{\mathbf{V} \in \mathbb{R}^n : \mathbf{V} = (v_i^E(\beta, W))_{i \in N} \text{ for some } \beta \text{ enforced by } W \text{ on } \mathcal{W}\}.$$

Proposition 1 (screening does not increase the value of collusion)**.** *In the game without a regulator,* $\mathcal{V}(E) = \lim_{t \to \infty} \mathcal{B}_E^t([0, 1]^n)$.

Consider a regulator running an asymptotically safe test with a vanishing false positive rate $G(M)$*. Then, for all environments* $E \in \mathcal{E}$*,* $\mathcal{V}_{T,M}(E) \subset \mathcal{B}_E^{T-M}([-G(M)\delta^M K_T, 1]^n)$*.*

In words, we can bound the set of equilibrium values of the game with the regulator by applying the APS operator $T - M$ times to the set $[-G(M)\delta^M K_T, 1]^n$. If the limit $\cap_{t \geq 0} \mathcal{B}_E^t([-\nu, 1]^n)$ depends continuously on ν, and T, M are chosen so that $G(M)\delta^M K_T$ becomes vanishingly small as M and T grow large, then the set of values under screening $\mathcal{V}_{T,M}(E)$ will be included in the set of values without a regulator $\mathcal{V}(E)$ as T and M grow large.

The key intuition behind Proposition 1 is that screening for collusion using tests that are asymptotically safe does not significantly lower firms' min-max values relative to a setting without a regulator. Indeed, following a history $h_{i,t}$ with $t \leq T - M$, a firm can guarantee to pass the test with probability $1 - G(M)$ by playing a static best response to the actions of her opponents during the monitoring phase. Hence, firm i's continuation payoff at $t + 1$ cannot be lower than $-G(M)\delta^{T-t}K_T$. Safety limits the cartel's ability to exploit regulatory screening as an enforcement tool.

Convergence of the equilibrium set. Next, we show that, for all $E \in \mathcal{E}$, the equilibrium set $\Sigma_{T,M}(E)$ converges to $\Sigma(E)$ as T and M grow large when the regulator runs asymptotically safe tests. To establish our result, we make the following assumptions. First, we assume that the set of feasible bids B, the set of possible cost profiles C, and the set of possible signal profiles Z are all finite. This assumption guarantees that the set of mixed strategy profiles is compact under the product topology (by Tychonoff's Theorem).

Second, we assume that the length M of the monitoring phase grows together with T, but at a slower rate, for instance $M_T = \lfloor T^x \rfloor$ for some constant $x \in (0, 1)$. We let $(\tau_i^T)_{i\in N}$ denote the tests that the antitrust agency runs when the testing date is T and the length of the monitoring period is $M = M_T$. Hence, when tests $(\tau_i^T)_{i\in N}$ are asymptotically safe, the probability of a false positive vanishes as T grows large.

Finally, for each testing date T, we assume that firms that fail the test incur a penalty $K_T = \delta^{-T}K$ for some $K > 0$. This implies that antitrust oversight continues to have an impact on firms' overall payoffs as T grows.[16]

[16]Our results continue to hold as long as $K_T \leq \delta^{-T}K$.

Endow the set of strategy profiles with the product topology, and for each $E \in \mathcal{E}$, define

$$\Sigma_\infty(E) \equiv \{\sigma : \exists (T^s)_{s\in\mathbb{N}} \to \infty, (\sigma^s)_{s\in\mathbb{N}} \to \sigma, \text{ with } \sigma^s \in \Sigma_{T^s, M_{T^s}}(E) \text{ for all } s \in \mathbb{N}\}$$

to be the set of limiting PPBE under environment E as T and M grow to infinity.

Proposition 2 (screening does not expand equilibrium strategies)**.** *Suppose the regulator runs asymptotically safe tests. Then, for all $E \in \mathcal{E}$, $\Sigma_\infty(E) \subset \Sigma(E)$.*

Fix $E \in \mathcal{E}$ and a sequence $(T^s)_{s\in\mathbb{N}} \to \infty$. Let $(\sigma^s)_{s\in\mathbb{N}}$ be a sequence satisfying $\sigma^s \in \Sigma_{T^s, M_{T^s}}(E)$ for all s. Then, by Proposition 2, $(\sigma^s)_{s\in\mathbb{N}}$ must approach the equilibrium set $\Sigma(E)$ as $s \to \infty$, even if the $(\sigma^s)_{s\in\mathbb{N}}$ does not converge.[17] Hence, Proposition 2 shows that antitrust oversight based on tests that are approximately safe (i.e., large T) does not significantly enlarge the set of enforceable collusive schemes available to cartels.

The following corollary clarifies that, in any equilibrium in $\Sigma_{T,M}(E)$, players expect to pass the test with high probability whenever T is sufficiently large.

Corollary 1 (tests pass with high probability in equilibrium)**.** *Suppose the regulator runs asymptotically safe tests, with testing date T and a monitoring length M. Then, for all $E \in \mathcal{E}$, all $\sigma \in \Sigma_{T,M}(E)$, all $i \in N$ and all histories $h_{i,T-M}$,*

$$\textsf{prob}_{E,\sigma}(\tau_i = 1 | h_{i,T-M}) \leq \frac{1}{\delta^{-(T-M)}K} + G(M).$$

In particular, if $M = \lfloor T^x \rfloor$ for some $x \in (0,1)$, $\textsf{prob}_{E,\sigma}(\tau_i = 1 | h_{i,T-M}) \to 0$ as $T \to \infty$.

In words, when regulators use safe tests, if cartels persist, they adapt their behavior to ensure that they pass the test with high probability. In equilibrium, safe tests should not

[17] To see why, suppose the result is not true. Then, there exists a subsequence $(\sigma^{s_k})_{k\in\mathbb{N}}$ such that all elements in the subsequence are bounded away from $\Sigma(E)$. Since the set of strategy profiles is compact, $(\sigma^{s_k})_{k\in\mathbb{N}}$ has a convergent subsubsequence $(\sigma^{s_{k_m}})_{m\in\mathbb{N}}$. And by Proposition 2, $(\sigma^{s_{k_m}})_{m\in\mathbb{N}}$ converges to a point in $\Sigma(E)$. But this is a contradiction to $(\sigma^{s_k})_{k\in\mathbb{N}}$ being such that all its elements are bounded away from $\Sigma(E)$.

trigger. Still, we show in Section 5.3 that safe tests can strictly reduce the scope for collusion, even if cartels can adapt.

5.2 Safe Tests do not Affect Competitive Equilibria

Let $\Sigma^{\mathsf{comp}}(E)$ denote the set of competitive equilibria under environment E, i.e., strategy profiles in which all bidders play according to the same Bayes' Nash equilibrium of the stage game at all histories. Assume that $\Sigma^{\mathsf{comp}}(E)$ is non-empty.[18] Then, the game with testing date T and monitoring length M has an equilibrium in which players play according to a Bayes' Nash equilibrium of the stage game at all periods $t \leq T - M$ (regardless of whether the tests that the regulator runs are safe or not). In addition, if the tests are asymptotically safe, we have that for all $E \in \mathcal{E}$, players' payoffs from using strategy profile $\sigma \in \Sigma^{\mathsf{comp}}(E)$ are bounded below by $V_i^E(\sigma, h_0) - G(M)\delta^T K_T$, where for any strategy profile $\hat{\sigma}$, $V_i^E(\hat{\sigma}, h_0)$ denotes player i's payoffs at the start of the game (excluding penalties) when firms play according to $\hat{\sigma}$.

The following Proposition summarizes this discussion.

Proposition 3 (safe tests do not reduce competitive payoffs)**.** *(i) For any tests (τ_i) and every $E \in \mathcal{E}$, the regulatory game has an equilibrium that coincides with an equilibrium in $\Sigma^{\mathsf{comp}}(E)$ for all periods $t \leq T - M$;*

(ii) If tests (τ_i) are asymptotically safe, then for all $E \in \mathcal{E}$ and all $i \in N$, firm i's payoff from any strategy profile $\sigma \in \Sigma^{\mathsf{comp}}(E)$ is bounded below by $V_i^E(\sigma, h_0) - G(M)\delta^T K_T$.

Together Propositions 2 and 3 show that by using safe tests regulators can ensure that they do not increase the incentives to form a cartel: collusive and competitive values respectively increase and decrease by a vanishing amount.

[18] That is, assume that the stage game has a Bayes' Nash equilibrium. Note that the stage game always has a Bayes' Nash equilibrium if sets B, C and Z are all finite.

5.3 Safe Tests Can Have Strict Bite

Using safe tests guarantees that incentives to cartelize are not meaningfully increased by regulation. At the same time, screening for collusion with safe tests will in general restrict firms' equilibrium strategies (and values), making cartels worse off. In this section, we formalize this with an example: we present a setting in which safe tests can strictly reduce incentives to cartelize, and can even cause the set of collusive equilibria to unravel. We stress that, although we focus on a specific test to make our point, Chassang et al. (2022) describes systematic procedures to generate safe tests given a set of environments $\mathcal{E}$.

A safe test of interest. We begin by introducing a safe test closely related to variance tests frequently used by antitrust agencies. We will show that optimal collusive strategies in benchmark environments fail this test, and that in turn, this can cause the set of collusive equilibria to unravel.

We maintain the following assumption throughout this section.

Assumption 1 (costly bid preparation). *For all feasible environments $E = (F_C, F_Z, (\kappa_i)) \in \mathcal{E}$, $\kappa_i \geq \hat{\kappa} > 0$ for all $i \in N$.*

Under all plausible environments, firms face strictly positive bid-preparation costs. We note that bid preparation costs for public works projects can be especially large.[19]

We now introduce asymptotically safe tests (τ_i^{close}). Fix $\rho \in (0,1)$ and $\Delta \in (0,1)$. For each $i \in N$, define test τ_i^{close} as:

$$\tau_i^{\mathsf{close}} \equiv \mathbf{1}\left(\frac{1}{M}\sum_{t=T+1-M}^{T} \mathbf{1}_{\wedge \mathbf{b}_{-i,t} \in (b_{i,t}-\Delta, b_{i,t}]} > \rho\right).$$

Test τ_i^{close} triggers when the mass of bids within Δ of the winning bid is higher than ρ during the monitoring phase. It detects a high frequency of close-to-winning losing bids. Our next

[19]For instance, Krasnokutskaya and Seim (2011) study highway procurement auctions in California, and estimate that bid preparation costs are between 2.2% and 3.9% of the engineer's cost estimate.

result shows that under Assumption 1, test $\tau_i^{\textsf{close}}$ is asymptotically safe.

Proposition 4. *Suppose Assumption 1 holds. Then, tests* $(\tau_i^{\textsf{close}})$ *with* $\rho > \frac{\Delta}{\hat{\kappa}}$ *are asymptotically safe: for all* $E \in \mathcal{E}$*, all* $i \in N$*, and all* σ *such that* i *is competitive during the monitoring phase,* $\textsf{prob}_{E,\sigma}(\tau_i^{\textsf{close}} = 0) \geq 1 - e^{-\frac{\alpha^2}{2}M}$ *for some constant* $\alpha > 0$*.*

In words, when bid preparation costs are strictly positive, nearly-tied bids and winning bids near the reserve price should be infrequent under competitive behavior. We note that the tests $(\tau_i^{\textsf{close}})$ are very similar to variance screens often used by regulators (Abrantes-Metz et al., 2005, Imhof et al., 2016) that flag auctions with bids that are unusually close together. The variance screen focuses on the variance of all of the bids while the tests $(\tau_i^{\textsf{close}})$ focus on the distance between firms' bids and the most competitive other bid. The match is exact when there are only two bidders.

Optimal collusion in a benchmark setting. We now establish two results under a benchmark environment with constant costs: (1) optimal collusive schemes fail tests $\tau_i^{\textsf{close}}$; and (2) this can cause the set of collusive equilibria to unravel for an open set of discount factors.

Consider the repeated procurement model of Section 2 with two firms ($n = 2$), under the benchmark environment $E^{\textsf{bmk}} \in \mathcal{E}$ such that, at each time t, firms share the same procurement cost c, which we normalize to $c = 0$. In addition, each firm $i = 1, 2$ faces bid preparation cost $\kappa_i = \kappa > 0$.

We assume that the set of environments $\mathcal{E}$ that the regulator deems feasible satisfies Assumption 1, and that the antitrust authority uses tests $(\tau_i^{\textsf{close}})$. Lastly, we assume that the auctioneer cancels the auction if there is only one participant.[20] This assumption rules out bidding profiles in which firms alternate in participating.

[20]Many government procurement auctions have a minimum number of required bids. North Carolina, for example, requires at least three competitive bids on all procurement auctions for construction or repairs. We note that our results would continue to hold if the auctioneer runs the auction with probability $q < (1 - 2\kappa)/(1 - \kappa)$ if only one firm participates.

Let $\overline{V}^{\mathsf{bmk}}$ denote the largest cartel equilibrium payoff in the game without a regulator under environment E^{bmk}. For each T, M, let $\overline{V}^{\mathsf{bmk}}_{T,M}$ denote the largest cartel equilibrium payoff (including potential penalties) in the game with testing date T and monitoring phase M. Define $\underline{\delta} \equiv \frac{1}{1-\kappa}\frac{1}{2}$.

Proposition 5. *(i) Suppose there is no regulator. If $\delta \geq \underline{\delta}$, under an optimal collusive equilibrium of the benchmark model both firms submit a bid equal to $r = 1$ at all on-path histories; i.e., $\overline{V}^{\mathsf{bmk}} = 1 - 2\kappa$.*

(ii) Suppose the regulator uses tests $(\tau_i^{\mathsf{close}})$. Then there exists $\mu > 0$ such that, for all $\varepsilon > 0$ and all $\delta \in [\underline{\delta}, \underline{\delta} + \mu]$, $\overline{V}^{\mathsf{bmk}}_{T,M} < \varepsilon$ whenever M, $T - M$ and K are all large enough.

Proposition 5(i) shows that, when there is no regulator and δ is higher than $\underline{\delta}$, optimal collusion involves both firms submitting a bid equal to the reserve price at all histories. Clearly, firms $i = 1, 2$ fail test τ_i^{close} under this bidding profile. Proposition 5(ii) establishes that for an open set of discount factors δ greater than $\underline{\delta}$, firms are unable to sustain supra-competitive profits when the regulator uses tests $(\tau_i^{\mathsf{close}})$. Intuitively, the winning bid must be strictly below the reserve price a positive fraction of periods during the monitoring phase for firms to pass test τ_i^{close}, lowering equilibrium values. For moderate values of δ, reducing equilibrium values precludes firms from sustaining supra-competitive profits.

Altogether this shows that even if firms adapt to regulatory screens, these screens can effectively curb collusion.[21]

6 Discussion: Policy Relevance

6.1 Summary

This paper proposes an equilibrium model of data-driven screening for cartel behavior in which cartel members can adapt their bidding strategies to undermine regulatory oversight.

[21] Of course, if firms do not adapt, regulatory screens may be more effective at detecting cartels.

We emphasize the value of safe tests designed to fail firms whose behavior cannot be explained by any competitive model. We show that such tests cannot help cartels sustain new collusive equilibria, and that they can be freely combined to create new safe tests. Importantly, safe tests can strictly reduce the set of enforceable equilibrium values in some settings. This makes them a robust improvement over laissez-faire.

We conclude by discussing policy relevant aspects of our work. First, we provide suggestive evidence that cartel members do in fact adapt to regulatory scrutiny. Second, we discuss the extent to which known tests of collusion can be made safe.

6.2 Evidence of Adaptation

We now provide suggestive evidence that cartels do in fact adapt to their regulatory environment. In particular, we argue that several puzzling bidding patterns observed in Japanese procurement auctions can be rationalized as adaptations to statistical screens frequently used by antitrust authorities. Interestingly, it turns out that these puzzling patterns themselves can form the basis of additional safe tests. We note that this evidence of adaptation to the regulatory environment is consistent with the work of Wollmann (2019) and Cunningham et al. (2021) who find evidence of bunching of mergers and acquisitions below size thresholds that make them subject to increased regulatory scrutiny.

Consider the class of safe tests $(\tau_i^{\mathsf{close}})_{i \in N}$, which are closely related to variance screens used by antitrust authorities. A simple adaptation to this specific regulatory screen is to systematically avoid close bids. Concretely, bidders may avoid bidding profiles $\mathbf{b} = (b_i)_{i \in N}$ such that $|b_i - \wedge \mathbf{b}_{-i}| < \epsilon$ for some i, as well as bidding profiles that are clustered around the reserve price, r, such that $r - \wedge \mathbf{b} < \epsilon$ for $\epsilon > 0$ small. This ensures that the winning bid and the next closest bid remain far apart.

We now provide suggestive evidence of adaptive bidding behavior described above in a sample of procurement auctions taking place in the city of Tsuchiura, in Ibaraki prefecture, Japan. The data contain approximately 400 auctions taking place between May 2007 and October 2009. The auction format is first-price sealed-bid, with a public reserve price. The median number of bidders is 5, and the median winning bid is approximately USD 98,000. In previous work, Chassang and Ortner (2019) provide evidence of bidder collusion in these auctions.

Formally, for each auction t and each bidder i participating in this auction, let $\Delta_{i,t} \equiv (b_{i,t} - \wedge \mathbf{b}_{-i,t})/r_t$ denote the margin by which i wins or losses the auction normalized by the reserve price r_t. Let $\wedge \mathbf{b}_t \equiv \min_i b_{i,t}$ denote the winning bid in the auction. A cartel that is trying to avoid triggering tests (τ_i^{close}) may generate a sample of bids such that the density of $\Delta_{i,t}$ is close to 0 around $\Delta_{i,t} = 0$, and such that the distribution of winning bids $\wedge \mathbf{b}_t$ has almost no mass close to the reserve price.

Figure 1 illustrates the distribution of normalized bid differences $\Delta_{i,t}$ for the sample of procurement auctions from Tsuchiura.[22] Figure 2 plots the empirical c.d.f. of normalized winning bids $\wedge \mathbf{b}_t / r_t$ (i.e., winning bid divided by reserve price) for this same sample of auctions. As anticipated, firms seek to avoid suspiciously close bids and winning bids close to the reserve price.[23]

Interestingly, as we show in Chassang et al. (2022), this missing mass of bids around $\Delta = 0$ is itself a suspicious pattern that can be turned into a safe test. Recall that, for each bidder i participating in auction t, $\wedge \mathbf{b}_{-i,t} = \min_{j \neq i} b_{j,t}$ is the lowest bid among i's opponents in the auction. For each bidder $i \in N$, integers $T, M \leq T+1$ and constant $\rho \in (-1, \infty)$, let

$$\widehat{D}_i(\rho, T, M) \equiv \frac{1}{M} \sum_{t=T-M+1}^{T} \mathbf{1}_{b_{i,t}(1+\rho) < \wedge \mathbf{b}_{-i,t}},$$

[22] In Chassang et al. (2022), we show that similar bidding patterns also appear in procurement auctions run by the Ministry of Land, Infrastructure and Transportation in Japan, and in procurement auctions run by municipalities located in the Tohoku region in Japan.

[23] Kawai et al. (2022a) proposes an alternative explanation for the bidding patterns in Figure 1. In particular, it shows that such patterns may arise when cartels have access to a mediator that helps them coordinate their bids.

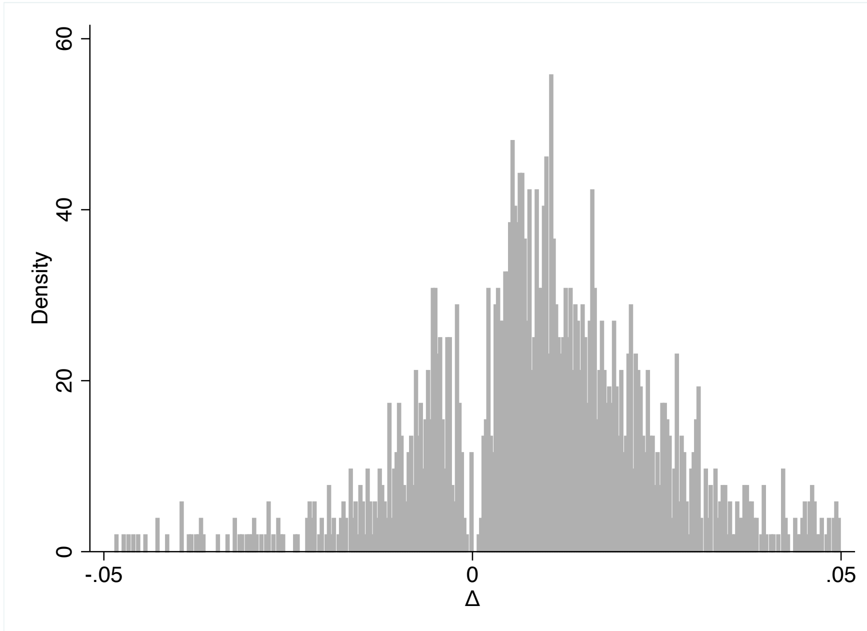

Figure 1: Distribution of bid differences Δ

denote firm i's sample demand; i.e., the empirical probability with which i wins an auction at any given bid during the monitoring phase. Fix $\rho \in (0,1)$. For $\mu > 0$ small, define tests $(\tau_i^{\textsf{missing}})_{i\in N}$ as

$$\tau_i^{\textsf{missing}} \equiv \mathbf{1}\left(\frac{\log \widehat{D}_i(\rho,T,M) - \log \widehat{D}_i(0,T,M)}{\log(1+\rho)} > -1+\mu\right).$$

In words, firm i fails test $\tau_i^{\textsf{missing}}$ if its sample demand $\widehat{D}_i(\rho,T,M)$ is inelastic at $\rho = 0$. We note that the data in Figure 1 fails tests $(\tau_i^{\textsf{missing}})$: when the distribution of bid differences $\Delta_{i,t}$ has no mass at $\Delta = 0$, we have $D_i(\rho,T,M) \approx D_i(0,T,M)$ for all $\rho > 0$ small.

In Chassang et al. (2022), we show that testing for missing bids is a safe test.[24] The intuition for this result is simple: when firm i's sample demand is inelastic, it is stage-game profitable for the firm to increase its bids. Hence, such bidding profiles cannot arise when firm i is competitive. This implies that adaptive behavior by a cartel to tests $(\tau_i^{\textsf{close}})$ can be

[24]Importantly, the test is safe even if the regulator places no restrictions on the set of economic environments.

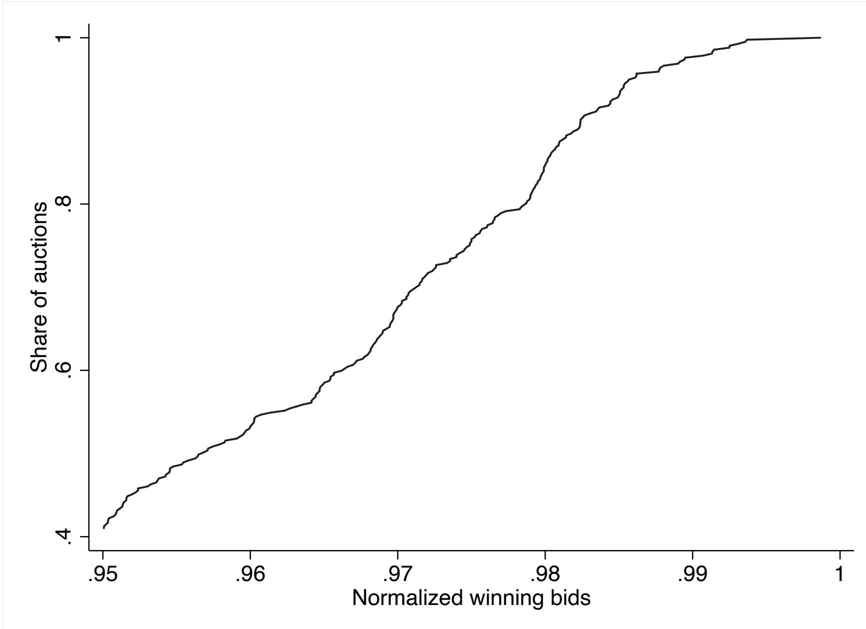

Figure 2: Cumulative distribution function of normalized winning bids

detected using safe tests $(\tau_i^{\text{missing}})_{i \in N}$.

6.3 Assessing the Safety of Known Tests

A well developed literature has identified a number of screens for collusion valid under various assumptions about both the environment and the behavior of firms. This section reviews a selection of prominent tests from the perspective of safety.

We make two remarks before turning to the discussion of specific tests. Our first remark is that safety of a test is a sufficient condition for tests not to create additional collusive equilibria, but it is not a necessary condition.[25] This means that tests that are not unilaterally safe may not in fact be problematic in practice. As we discuss below, the test of Porter and Zona (1993) is an example of an unsafe test that is unlikely to expand the set of collusive equilibria. Our second remark is that the conjunction of an asymptotically safe test and any

[25]For instance, tests $(\tau_i)_{i \in N}$ with the property that, for each σ_{-i}, each firm i has a strategy $\sigma_i(\sigma_{-i})$ available such that (i) firm i passes the test with probability approaching 1 under $(\sigma_i(\sigma_{-i}), \sigma_{-i})$, and (ii) firm i earns non-negative payoffs under $(\sigma_i(\sigma_{-i}), \sigma_{-i})$ won't generate new collusive equilibria.

other test is safe. This means that applying a first screen based on a familiar test (even an unsafe one), and strengthening that initial screen with a safe test yields a compound test that is also safe. Such a compound test allows investigators to rapidly screen evidence using an appealing but unsafe test (e.g. higher than expected prices), and only implement more complex safe tests (such as those of Chassang et al. (2022)) on data likely to be associated with non-competitive behavior.

Conditional independence and exchangeability. Bajari and Ye (2003) propose two tests for competition under the assumption that costs are independent and identically distributed conditional on observable firm characteristics. Under these maintained assumptions, bids from competitive firms should be exchangeable and independent conditional on observables. Bajari and Ye (2003) propose to screen for collusion by regressing each bidder i's bid on bidder characteristics $z_{i,t}$ and testing for the equality of the regression coefficients across bidders, i.e., $\beta_i^z = \beta_j^z$ (exchangeability), and for no correlation of the residuals, i.e., $corr(\epsilon_{i,t}, \epsilon_{j,t}) = 0$ (independence), where β_i^z is the regression coefficient on $z_{i,t}$ and $\epsilon_{i,t}$ is the regression residual. Bidders i and j fail the test if either $\beta_i^z = \beta_j^z$ or $corr(\epsilon_{i,t}, \epsilon_{j,t}) = 0$ is rejected.

Tests of exchangeability need not be unilaterally safe under the maintained assumptions of Bajari and Ye (2003) since a cartel bidder and a competitive bidder that bids against a cartel will not bid in the same way even if their characteristics are similar. Moreover, it seems possible for a cartel to use the possibility of regulatory penalty to punish a deviator by playing strategies that make the deviator fail the test.[26] Hence, tests of exchangeability may increase the set of payoffs available to the cartel. They also raise the possibility that a competitive bidder is incentivized to bid in a similar way as a member of a cartel in order to pass the test, resulting in less competitive bids than in the absence of antitrust scrutiny.

[26]Suppose that bidder i's bid is increasing in $z_{i,t}$ and bidder j's bid is decreasing in $z_{j,t}$. Bidder k will find it impossible to pass the test regardless of how it bids since β_k^z will be different from either β_i^z or β_j^z. Hence bidders i and j can effectively punish bidder k by making sure k fails the test.

On the other hand, under the maintained assumptions in Bajari and Ye (2003), a test of independence is safe. This is because the bid of a competitive bidder who does not communicate with any other bidder will be independent of all other bids.

Bid rotation. Cartels frequently use bid-rotation to allocate cartel rents across its members. Under a bid rotation scheme, bidders take turns winning the auction and designated losers submit complementary bids. Rotation patterns are often considered indicators of collusion[27] and have been used to build tests of collusion (see, e.g., Ishii, 2008, 2009).

An intuitive way to test for bid rotation is to test whether backlog negatively affects the winning probability (Ishii, 2008, 2009), where backlog is defined as the amount of work a firm has won in the recent past. However, testing for the effect of past backlog on the winning probability in this manner is unlikely to be a safe test even if we take the set of environments $\mathcal{E}$ to be such that costs are drawn from a distribution that does not depend on past backlog.[28] Even if costs are drawn i.i.d every period, a cartel can make the bids of its members depend on the backlog of its members or that of non-cartel firms. A competitive firm that best responds to the bidding strategies of such a cartel may fail the test (as its bid will also depend on its own backlog or the backlog of cartel members). This suggests that the test is not safe. Moreover, it seems easy for a cartel to punish a deviator i by making it less likely for bidder i to win when i's backlog is high and vice versa, making bidder i fail the test. Penalty by the regulator can then be exploited to exact harsher punishment on deviators, potentially expanding the set of payoffs.

Kawai et al. (2022b) show, however, that a suitable adjustment of the original insight is in fact a safe test. Under competitive behavior, backlog should not be correlated to whether a firm wins or loses an auction conditioning on being a close winner or loser. This is true even if the environment $\mathcal{E}$ is such that bidder costs are allowed to depend on backlog. If

[27] See, for example, the "Red Flags Of Collusion" report, published by the U.S. DOJ or the pamphlet on bid rigging published by the Canadian Competition Bureau.

[28] If costs depend on past backlog, even a competitive firm would bid more aggressively when its backlog is low and less aggressively when its backlog is high, leading to bid patterns that are likely to fail the test.

backlog is correlated with winning the auction even conditional on close bids, then this is a rejection of competitive behavior and evidence of collusion.

Structural breaks. Building on the observation that cartels may need to go through price-wars to enforce supra-competitive bids, Porter (1983), Ellison (1994), Harrington (2008) suggest that structural breaks in prices time series are a red flag for cartel behavior.

Even if we restrict the set of environments $\mathcal{E}$ to those in which the underlying cost distribution stays constant, the test does not appear safe as we discussed in Section 3. This is because a competitive firm bidding against a cartel that is going through a price-war will also exhibit structural breaks in the time-series of its bids.

Variance screens. Auctions with many similar bids, as well as those with relatively large winning margins are often considered indicative of collusion. For example, Imhof et al. (2016) discuss two tests, one based on the coefficient of variation CV_j, and another based on the winning margin WM_j, both of which were used to launch an investigation against construction firms by the Swiss competition bureau. The test CV_j is defined as $CV_j = \sigma_j/\mu_j$, where σ_j and μ_j are the standard deviation and the mean of the bids in auction j. A low value of CV_j suggests that there is a cluster of similar bids and is deemed indicative of collusion. The test based on the winning margin is defined as $WM_j = \Delta_j/\tilde{\sigma}_j$, where Δ_j and $\tilde{\sigma}_j$ are the winning margin and the standard deviation of the losing bids in auction j, respectively. A high value of WM_j implies that auction j has a relatively large winning margin and is taken as evidence of collusion.

Because the tests CV_j and WM_j take the auction as the unit of analysis, they are unlikely to be safe tests. However, these tests can be made safe by considering τ_i^{close} and $\tau_i^{\mathsf{missing}}$. The test τ_i^{close} focuses on the mass of almost tied bids, similar to CV_j; and the test $\tau_i^{\mathsf{missing}}$ focuses on the winning margin, similar to WM_j. These tests focus on very similar features of the bid distribution to screen for collusion while being safe.

Unresponsive losing bids. Porter and Zona (1993, 1999), Porter (2005) develop tests of cartel bidding based on the observation that collusive bids, and in particular losing bids, may be disconnected to marginal costs. This suggests a test of collusion based on the sensitivity of losing bids with respect to observable cost determinants, including distance, commodity prices, and so on. If bid changes have a low sensitivity to cost changes, this is evidence of collusive behavior.

Under most environments, it seems unlikely that these tests are individually safe. The behavior of a competitive firm bidding against a cartel sustaining prices above marginal cost will be driven by cartel behavior rather than marginal costs.[29] For this reason a firm that behaves competitively against a cartel is not guaranteed to pass such tests. However, given that a competitive firm can always choose to bid as an increasing function of cost shocks (although doing so may not be a best response), the test cannot be used as a punishment against deviating firms. In other words, as we explain in footnote 25, although this test seems not to be safe it does not expand the set of equilibrium payoffs.

Appendix

A Proofs

Proof of Proposition 1. The first statement in the proposition follows from standard arguments (e.g., Abreu et al. (1990)). The second statement in the proposition follows from two observations. First, for every strategy profile $\overline{\sigma}_{-i}$ of i's opponents and every history $h_{i,t}$ with $t \leq T - M$, firm i can guarantee itself a payoff of at least $0 - G(M)\delta^{T-M}K_T$ by playing a static best-response to $\overline{\sigma}_{-i}$ at all periods $s \geq t$. Second, since firms' flow profits are bounded above by $r = 1$, we have that $W^E(\overline{\sigma}, h_{i,t}) \leq 1$ for all i, all $\overline{\sigma}$ and all $h_{i,t}$. Hence, for

[29] For example, suppose that the cartel bids higher when the expected costs of the competitive bidder are low and vice versa. Then the competitive bidder would have incentive to bid relatively high when expected costs are low and vice versa, failing the test.

any $E \in \mathcal{E}$ and any $\sigma \in \Sigma_{T,M}(E)$, firms' payoffs (including expected penalties) at $t < T - M$ must lie $\mathcal{B}_E^{T-M-t}([-G(M)\delta^{T-M}K_T, 1]^n)$. ■

Proof of Proposition 2. Fix an environment $E \in \mathcal{E}$. Recall that for every strategy profile $\overline{\sigma}$, every firm $i \in N$ and every history $h_{i,t}$ of firm i, $V_i^E(\overline{\sigma}, h_{i,t})$ denotes firm i's expected continuation payoff (excluding penalties) at history $h_{i,t}$ under $\overline{\sigma}$, and $P_i^E(\overline{\sigma}, h_{i,t}) = \mathbf{1}_{t \leq T}\mathbb{E}_{E,\overline{\sigma}}\left[\tau_i \delta^{T-t} K_T | h_{i,t}\right]$ denotes firm i's expected discounted penalty at history $h_{i,t}$ under $\overline{\sigma}$. Firm i's total expected payoff at history $h_{i,t}$ under $\overline{\sigma}$ is $W^E(\overline{\sigma}, h_{i,t}) = V_i^E(\overline{\sigma}, h_{i,t}) - P_i^E(\overline{\sigma}, h_{i,t})$.

Fix $T > 0$ and $\sigma = (\sigma_i)_{i \in N} \in \Sigma_{T,M}(E)$. Pick $\epsilon > 0$ and let $\widehat{T}$ be such that $\delta^{\widehat{T}}(1 + \overline{c}) < \frac{\epsilon}{2}$, where $\overline{c}$ is an upper bound to firms' costs (and so firms' flow payoffs from the auction are bounded above by $r = 1$ and bounded below by $-\overline{c}$). Pick a history $h_{i,t}$, and a strategy $\tilde{\sigma}_i \neq \sigma_i$ for player i. Let $\hat{\sigma}_i(\tilde{\sigma}_i)$ be a strategy that coincides with $\tilde{\sigma}_i$ at all histories $h_{i,s}$ of length $s \leq t + \widehat{T}$, and that plays a static best-response to σ_{-i} at all histories $h_{i,s}$ of length $s > t + \widehat{T}$. Note then that

$$|V_i^E((\tilde{\sigma}_i, \sigma_{-i}), h_{i,t}) - V_i^E((\hat{\sigma}_i(\tilde{\sigma}_i)), \sigma_{-i}), h_{i,t})| < \frac{\epsilon}{2}. \tag{1}$$

Since $\sigma \in \Sigma_{T,M}(E)$, it must be that

$$\begin{aligned} & V_i^E(\sigma, h_{i,t}) - P_i^E(\sigma, h_{i,t}) \geq V_i^E((\hat{\sigma}_i(\tilde{\sigma}_i), \sigma_{-i}), h_{i,t}) - P_i^E((\hat{\sigma}_i(\tilde{\sigma}_i), \sigma_{-i}), h_{i,t}) \\ \iff & V_i^E(\sigma, h_{i,t}) - V_i^E((\hat{\sigma}_i(\tilde{\sigma}_i), \sigma_{-i}), h_{i,t}) \geq P_i^E(\sigma, h_{i,t}) - P_i^E((\hat{\sigma}_i(\tilde{\sigma}_i), \sigma_{-i})), h_{i,t}). \end{aligned} \tag{2}$$

Moreover, since tests (τ_i^T) are asymptotically safe, for $T > t + \widehat{T} + M$ it must be that

$$\begin{aligned} P_i^E(\hat{\sigma}_i(\tilde{\sigma}_i), h_{i,t}) &= \mathbb{E}_{E,\sigma}\left[\tau_i^T \delta^{T-t} K_T | h_{i,t}\right] \\ &\leq \delta^{T-t} K_T G(M) = \delta^{-t} K G(M), \end{aligned} \tag{3}$$

where the last equality uses $K_T = \delta^{-T}K$. Since $\lim_{M \to 0} G(M) = 0$ and since $M = \lfloor T^x \rfloor$ for some $x \in (0,1)$, the right-hand side of (3) goes to zero as $T \to \infty$. Let $\overline{T} > 0$ be such that $KG(M) \leq \frac{\epsilon}{2}$ for all $T > \overline{T}$. Then, for all $T > \overline{T}$, we have

$$V_i^E(\sigma, h_{i,t}) \geq V_i^E((\hat{\sigma}_i(\tilde{\sigma}), \sigma_{-i}), h_{i,t}) - \delta^{-t}\frac{\epsilon}{2} \geq V_i^E((\tilde{\sigma}_i, \sigma_{-i}), h_{i,t}) - \delta^{-t}\epsilon, \tag{4}$$

where the first inequality uses (2), (3) and $P_i(\sigma, h_{i,t}) \geq 0$, and the second inequality uses (1) and $\delta^{-t} \geq 1$.

We now use (4) to establish the result. Towards a contradiction, suppose that $\Sigma_\infty(E) \not\subset \Sigma(E)$. Hence, there exists $\sigma^* \notin \Sigma(E)$, and sequences $(T^k) \to \infty$, $(\sigma^k) \to \sigma^*$ with $\sigma^k \in \Sigma_{T^k,M^k}(E)$ for all k.

Since $\sigma^* \notin \Sigma(E)$, there exists $i \in N$, a history $h_{i,t}$, a deviation $\tilde{\sigma}_i \neq \sigma_i^*$ and a scalar $\eta > 0$ such that

$$V_i^E((\tilde{\sigma}_i, \sigma_{-i}^*), h_{i,t}) > V_i^E(\sigma^*, h_{i,t}) + \eta. \tag{5}$$

Pick $\epsilon > 0$ such that $\epsilon\delta^{-t} < \frac{\eta}{4}$. By our arguments above (see equation (4)), for all k large enough it must be that

$$V_i^E(\sigma^k, h_{i,t}) \geq V_i^E((\tilde{\sigma}_i, \sigma_{-i}^k), h_{i,t}) - \delta^{-t}\epsilon > V_i^E((\tilde{\sigma}_i, \sigma_{-i}^k), h_{i,t}) - \frac{\eta}{4}. \tag{6}$$

Since payoffs are continuous in strategies, and since $\sigma^k \to \sigma^*$, for all k large enough it must be that

$$|V_i^E(\sigma^k, h_{i,t}) - V_i^E(\sigma^*, h_{i,t})| \leq \frac{\eta}{4} \tag{7}$$

$$|V_i^E((\tilde{\sigma}_i, \sigma_{-i}^k), h_{i,t}) - V_i^E((\tilde{\sigma}_i, \sigma_{-i}^*), h_{i,t})| \leq \frac{\eta}{4} \tag{8}$$

Combining (6) with (7) and (8) we get

$$V_i^E(\sigma^*, h_{i,t}) - V_i^E((\tilde{\sigma}_i, \sigma_{-i}^*), h_{i,t}) \geq -\frac{3}{4}\eta,$$

which contradicts (5). Hence, it must be that $\Sigma_\infty(E) \subset \Sigma(E)$. Since this holds for all $E \in \mathcal{E}$, this completes the proof. ■

Proof of Corollary 1. Fix $E \in \mathcal{E}$, $\sigma \in \Sigma_{T,M}(E)$, $i \in N$ and a history $h_{i,T-M}$. Firm i's expected payoff at $h_{i,T-M}$ under σ, E is $V_i^E(\sigma, h_{i,T-M}) - P_i^E(\sigma, h_{i,T-M})$. Let $\hat{\sigma}_i$ be a strategy under which firm i plays a static best-response to σ_{-i} at all histories. Firm i's payoff from playing according to $\hat{\sigma}_i$ when her opponents play according to σ_{-i} at history h_{T-M} satisfies

$$V_i^E((\hat{\sigma}_i, \sigma_{-i}), h_{i,T-M}) - P_i^E((\hat{\sigma}_i, \sigma_{-i}), h_{i,T-M}) \geq 0 - G(M)\delta^M K_T = -G(M)\delta^{-(T-M)}K, \quad (9)$$

where the inequality follows since τ_i is an asymptotically safe test, and since firm i's static best response each period must give i a payoff weakly larger than 0 and the equality uses $K_T = \delta^{-T}K$. Since σ is an equilibrium, and since $V_i^E(\sigma, h_{i,T-M}) \leq r = 1$, we have

$$\begin{aligned}
1 - P_i^E(\sigma, h_{i,T-M}) &\geq V_i^E(\sigma, h_{i,T-M}) - P_i^E(\sigma, h_{i,T-M}) \\
&\geq V_i^E((\hat{\sigma}_i, \sigma_{-i}), h_{i,T-M}) - P_i^E((\hat{\sigma}_i, \sigma_{-i}), h_{i,T-M}) \geq -G(M)\delta^{-(T-M)}K \\
&\implies \quad P_i^E(\sigma, h_{i,T-M}) \leq 1 + G(M)\delta^{-(T-M)}K.
\end{aligned}$$

Using $P_i^E(\sigma, h_{i,T-M}) = \delta^M K_T \mathbb{E}_{E,\sigma}[\tau_i | h_{i,T-M}] = \delta^{-(T-M)}K\mathbb{E}_{E,\sigma}[\tau_i | h_{i,T-M}]$, we get

$$\mathbb{E}_{E,\sigma}[\tau_i | h_{i,T-M}] \leq \frac{1}{\delta^{-(T-M)}K} + G(M).$$

Hence, if $M = \lfloor T^x \rfloor$ for some $x \in (0,1)$, $\mathbb{E}_{E,\sigma}[\tau_i | h_{i,T-M}] = \mathsf{prob}_{E,\sigma}[\tau_i = 1 | h_{i,T-M}]$ converges to 0 as $T \to \infty$. This completes the proof. ■

Proof of Proposition 4. Fix $E \in \mathcal{E}$, and let σ be a strategy profile under which firm i is competitive during the monitoring phase. Fix an on-path history $h_{i,t}$ with $t \in [T-M+1, T]$ such that $\sigma_i(h_{i,t}) = b_{i,t} \in [0,1]$.

For each $b \in [0,1]$, define $D_i(b|h_{i,t}) \equiv \mathsf{prob}_{E,\sigma}(\wedge \mathbf{b}_{-i,t} \succ b|h_{i,t})$, where $\wedge \mathbf{b}_{-i,t} \succ b$ denotes the event that either $\wedge \mathbf{b}_{-i,t} > b$, or $\wedge \mathbf{b}_{-i,t} = b$ but ties are broken in favor of bidder i. Hence, $D_i(b|h_{i,t})$ is the probability with which firm i expects to win the auction if she places bid b at history $h_{i,t}$, given strategy profile σ.

Since firm i is competitive during the monitoring phase under σ, for any $\Delta > 0$ it must be that

$$D_i(b_{i,t} - \Delta|h_{i,t})(b_{i,t} - \Delta - c_{i,t}) \leq D_i(b_{i,t}|h_{i,t})(b_{i,t} - c_{i,t}).$$

Moreover, since bidding is costly, it must be that $D_i(b_{i,t}|h_{i,t})(b_{i,t} - c_{i,t}) \geq \hat{\kappa} > 0$, and so $b_{i,t} - c_{i,t} \geq \hat{\kappa}$. Combining this with the inequality above, we get that for all $\Delta > 0$ small,

$$\begin{aligned} D_i(b_{i,t} - \Delta|h_{i,t}) &\leq D_i(b_{i,t}|h_{i,t}) \frac{b_{i,t} - c_{i,t}}{b_{i,t} - c_{i,t} - \Delta} \leq D_i(b_{i,t}|h_{i,t}) \frac{\hat{\kappa}}{\hat{\kappa} - \Delta} \\ \implies D_i(b_{i,t} - \Delta|h_{i,t}) - D_i(b_{i,t}|h_{i,t}) &\leq \frac{\Delta}{\hat{\kappa}} D_i(b_{i,t} - \Delta|h_{i,t}) \leq \frac{\Delta}{\hat{\kappa}}. \end{aligned} \tag{10}$$

Note that inequality (10) implies $\lim_{\Delta \searrow 0} D_i(b_{i,t} - \Delta|h_{i,t}) = D_i(b_{i,t}|h_{i,t})$; i.e., at history $h_{i,t}$, the distribution of $\wedge \mathbf{b}_{-i,t}$ cannot have a mass point at $b_{i,t}$. Hence, $D_i(b_{i,t}|h_{i,t}) = \mathsf{prob}_{E,\sigma}(\wedge \mathbf{b}_{-i,t} > b_{i,t}|h_{i,t})$, and so (10) implies that

$$\mathsf{prob}_{E,\sigma}(\wedge \mathbf{b}_{-i,t} \in (b_{i,t} - \Delta, b_{i,t}]|h_{i,t}) \leq D_i(b_{i,t} - \Delta|h_{i,t}) - D_i(b_{i,t}|h_{i,t}) \leq \frac{\Delta}{\hat{\kappa}}. \tag{11}$$

We now use (11) to show that tests $(\tau_i^{\mathsf{close}})$ are asymptotically safe when $\rho > \frac{\Delta}{\hat{\kappa}}$. For any t, define

$$\varepsilon_t \equiv \mathbf{1}_{\wedge \mathbf{b}_{-i,t} \in (b_{i,t} - \Delta, b_{i,t}]} - \mathsf{prob}_{E,\sigma}(\wedge \mathbf{b}_{-i,t} \in (b_{i,t} - \Delta, b_{i,t}]|h_{i,t}).$$

For all $t \in [T - M + 1, T]$, let $S_{t,M} = \sum_{s=T-M+1}^{t} \varepsilon_s$. Note that, for all s, $\mathbb{E}_{E,\sigma}[\varepsilon_s|h_{i,s-1}] = 0$. Hence, $S_{t,M} = \sum_{s=T-M+1}^{t} \varepsilon_s$ is a Martingale with respect to $(h_{i,t})$, with the absolute value of its increments bounded above by 1. By the Azuma-Hoeffding inequality, for any $\alpha > 0$

$$\mathsf{prob}_{E,\sigma}(S_{T,M} > \alpha M) \leq \exp\left(-\frac{\alpha^2 M}{2}\right). \tag{12}$$

The probability that firm i fails test τ_i^{close} under E, σ is given by

$$\begin{aligned}
\mathsf{prob}_{E,\sigma}(\tau_i^{\mathsf{close}} = 1) =& \mathsf{prob}_{E,\sigma}\left(\frac{1}{M}\sum_{t=T-M+1}^{T} \mathbf{1}_{\wedge \mathbf{b}_{-i,t} \in (b_{i,t}-\Delta, b_{i,t}]} > \rho\right) \\
=& \mathsf{prob}_{E,\sigma}\left(\frac{1}{M} S_{T,M} > \rho - \frac{1}{M}\sum_{t=T-M+1}^{T} \mathsf{prob}_{E,\sigma}(\wedge \mathbf{b}_{-i,t} \in (b_{i,t}-\Delta, b_{i,t}] | h_{i,t})\right) \\
\leq& \mathsf{prob}_{E,\sigma}\left(\frac{1}{M} S_{T,M} > \rho - \frac{\Delta}{\hat{\kappa}}\right),
\end{aligned} \tag{13}$$

where the last inequality follows from (11). Combining (13) with (12), and letting $\alpha = \rho - \frac{\Delta}{\hat{\kappa}} > 0$, we get that

$$\begin{aligned}
\mathsf{prob}_{E,\sigma}(\tau_i^{\mathsf{close}} = 1) =& \mathsf{prob}_{E,\sigma}\left(\frac{1}{M}\sum_{t=T-M}^{T} \mathbf{1}_{\wedge \mathbf{b}_{-i,t} \in (b_{i,t}-\Delta, b_{i,t}]} > \rho\right) \\
\leq& \mathsf{prob}_{E,\sigma}(S_{T,M} > \alpha M) \leq \exp\left(-\frac{\alpha^2 M}{2}\right)
\end{aligned}$$

This completes the proof. ■

Proof of Proposition 5. Consider first part (i). Since the auction runs only if both firms participate, the cartel's flow payoff in the repeated game is bounded above by $(1-\delta)(r-2\kappa)$. Consider the following strategy σ_i^{coll}:

- at the initial history h_0, or at any history h_t with $b_{j,s} = r$ for $j = 1, 2$, $s < t$, bid $b_{i,t} = r$;
- at any other history h_t, play the static Nash equilibrium.[30]

[30] With $\kappa > 0$, the stage game admits a Nash equilibrium in which both firms always stay out of the auction.

One can verify that, when $\delta \geq \underline{\delta} = \frac{1}{1-\kappa}\frac{1}{2}$, both firms playing according to σ_i^{coll} is an equilibrium of the repeated game without a regulator. Hence, $\overline{V}^{\mathsf{bmk}} = 1 - 2\kappa$.

We now turn to part (ii). We start by providing an upper bound to the cartel's payoffs during the monitoring phase. Let $b_{(1),s}$ denote the winning bid at period s, and for any $t \in [T-M+1, T+1]$ let $B_{t-1} = \sum_{s=T-M+1}^{t-1} b_{(1),s}$ denote the cumulative sum of winning bids during the testing phase prior to t (with $B_{T-M} = 0$). Note that, if both firms participate at all periods $t = T - M + 1, ..., T$, at least one firm $i \in \{1, 2\}$ fails test τ_i^{close} if $\frac{1}{M} B_T > \overline{B} \equiv 1 - (1 - 2\rho)\Delta$. Define $\widehat{W} \equiv \overline{B} - 2\kappa$.

For each history $h_t = (\mathbf{b}_s)_{s<t}$ with $t \in [T - M + 1, T + 1]$ define state $z_t = (t, B_{t-1})$ and let

$$\Phi(z_{T+1}) = 1 - 2\kappa - \delta^{-1} K_T \mathbf{1}_{B_T > M \times \overline{B}}.$$

Note that $\delta\Phi(z_{T+1})$ is an upper bound on the sum of firms' equilibrium payoffs at the end of period T, after bidding is done but before any penalties are paid.

For each $z_t = (t, B_{t-1})$ with $t \in [T + 1 - M, T]$, define

$$\begin{aligned} \Phi(z_t) = \sup_{b_t \sim F \in \Delta([0,1])} \; & \mathbb{E}_F[(1 - \delta)(b_t - 2\kappa) + \delta\Phi(t + 1, B_{t-1} + b_t)] \text{ s.t.} \\ & \forall b_t \in \mathsf{supp}\, F, \quad b_t \leq \frac{\delta}{1 - \delta}\Phi(t + 1, B_{t-1} + b_t). \end{aligned} \tag{14}$$

Two points are worth noting. First, (14) always admits a deterministic solution (i.e., a solution such that winning bid b_t is deterministic for all z_t). Indeed, (14) is a deterministic dynamic programing problem, and hence admits a deterministic solution. Second, whenever penalty K_T is sufficiently large (e.g., larger than $\overline{K} \equiv 1 - 2\kappa$), the solution to (14) will be such that firms pass the test with probability 1: i.e., $B_T \leq M\overline{B}$. From now on, we assume $K_T \geq \overline{K}$. For each $t = T + 1 - M, ..., T + 1$, let $\overline{\Phi}_t$ be the value of Program (14) at t. Since $z_{T+1-M} = (T + 1 - M, 0)$ regardless of the bidding history, the solution to (14) does not depend on bidding behavior prior to $T + 1 - M$.

We now show that, for any equilibrium σ and any history h_{T+1-M} of length $T+1-M$, the cartel's payoff under σ at history h_{T+1-M} is bounded above by $\overline{\Phi}_{T+1-M}$. To establish this result, we show that for any period t during the testing phase with history $h_t = (\mathbf{b}_s)_{s<t}$ and with state $z_t = (t, B_{t-1})$, if history h_t is such that each firm $i \in \{1,2\}$ would pass the test if it were to not participate at any auction $s \geq t$, then the cartel's equilibrium payoff at h_t under σ is bounded above by $\Phi(z_t)$. Note that this implies that $\overline{\Phi}_{T+M-1}$ is an upper bound to equilibrium payoffs at any history h_{T+1-M}.[31]

Note first that, since $\Phi(z_{T+1}) = 1 - 2\kappa$ for any history h_{T+1} such that both firms pass the test, cartel equilibrium payoffs at such a history h_{T+1} are bounded above by $\Phi(z_{T+1})$. Towards an induction, suppose that the result holds for all histories $h_{\hat{s}}$ of length $\hat{s} = t+1, ..., T+1$ with the property that each firm would pass the test if it stopped participating at all $s \geq \hat{s}$. Fix an equilibrium σ and a history $h_t = (\mathbf{b}_s)_{s<t}$ of length t, with t during the monitoring phase, with the property that each firm would pass the test if it stopped participating. Let $b_{(1),t}$ denote the equilibrium winning bid at h_t. Note that, if both firms participate at h_t, then for $i = 1, 2$ we must have

$$(1-\delta)(x_{i,t} b_{(1),t} - \kappa) + \delta W_{i,t+1} \geq (1-\delta)(b_{(1),t} - \kappa), \tag{15}$$

where $x_{i,t}$ denotes the probability with which i wins the auction at time t and $W_{i,t+1}$ denotes i's continuation value given history $h_{t+1} = h_t \sqcup \mathbf{b}_t$ (including expected penalties).[32] Summing across both players, and using $x_{1,t} + x_{2,t} \leq 1$ and $W_{1,t+1} + W_{2,t+1} = W_{t+1}$, we get

$$\begin{gathered}(1-\delta)(b_{(1),t} - 2\kappa) + \delta W_{t+1} \geq 2(1-\delta)(b_{(1),t} - \kappa) \\ \iff \frac{\delta}{1-\delta} W_{t+1} \geq b_{(1),t}.\end{gathered}$$

[31] Indeed, all histories h_{T+1-M} of length $T+1-M$ have $B_{T-M} = 0$ and have the property that each firm would pass test τ_i^{close} if it didn't participate at any auction $s \geq T+1-M$.

[32] To see why (15) must hold under σ, recall that h_t is such that i would pass the test if it were to not participate any longer. Hence, (15) must hold since firm i can obtain a payoff equal to the right-hand side by undercutting bid $b_{(1),t}$ at t and not participating at any future date.

If bids $\mathbf{b}_t = (b_{1,t}, b_{2,t})$ are such that firm $i \in \{1,2\}$ fails the test at history $h_t \sqcup \mathbf{b}_t$ if it were to stop participating, then we have $W_t \leq \Phi(z_t)$ whenever $K \geq \overline{K} = 1 - 2\kappa$. If not, then by the induction hypothesis we have that $W_{t+1} \leq \Phi(z_{t+1})$; and so it must be that

$$W_t = (1-\delta)(b_{(1),t} - 2\kappa) + \delta W_{t+1} \leq \Phi(z_t),$$

where the inequality follows since $\frac{\delta}{1-\delta}\Phi(z_{t+1}) \geq \frac{\delta}{1-\delta} W_{t+1} \geq b_{(1),t}$. Hence, equilibrium payoffs at any history h_{T-M+1} are bounded above by $\overline{\Phi}_{T-M+1}$.

We now show that, for all $\epsilon > 0$, $\overline{\Phi}_t \leq \widehat{W} + \epsilon = \overline{B} - 2\kappa + \epsilon$ for some $t \in [T-M+1, T]$ whenever M is large enough. Since the solution to (14) is deterministic, we have that for all $t = T-M+1, ..., T$, $\overline{\Phi}_t = (1-\delta)(b_t - 2\kappa) + \delta\overline{\Phi}_{t+1}$, or

$$b_t = \frac{1}{1-\delta}\left(\overline{\Phi}_t - \delta\overline{\Phi}_{t+1}\right) + 2\kappa.$$

Summing over periods $t = T-M+1, ..., T$ and dividing by M we get

$$\frac{1}{M}\sum_{t=T-M+1}^{T} b_t = \frac{1}{M}B_T = \frac{1}{M}\left(\sum_{t=T-M+2}^{T}\overline{\Phi}_t + \frac{1}{1-\delta}\left(\overline{\Phi}_{T-M+1} - \delta\overline{\Phi}_T\right)\right) + 2\kappa. \tag{16}$$

Since firms pass the test under the solution to (14), we have $\frac{1}{M}B_T \leq \overline{B} = \widehat{W} + 2\kappa$, and so (16) gives us

$$\frac{1}{M-1}\sum_{t=T-M+2}^{T}\overline{\Phi}_t \leq \frac{M}{M-1}\widehat{W} + \frac{1}{1-\delta}\frac{1}{M-1}\left(\delta\overline{\Phi}_T - \overline{\Phi}_{T-M+1}\right)$$

Since $\overline{\Phi}_t \leq 1 - 2\kappa$ for all t, we have that for all $\epsilon > 0$ there exists $\overline{M}_\epsilon$ such that, for all $M \geq \overline{M}_\epsilon$,

$$\frac{1}{M-1}\sum_{t=T-M+2}^{T}\overline{\Phi}_t \leq \widehat{W} + \epsilon. \tag{17}$$

Equation (17) implies that there exists $t \in [T-M+2, T]$ such that $\overline{\Phi}_t \leq \widehat{W} + \epsilon$ whenever

$M \geq M_\epsilon$.

Pick $\epsilon > 0$ such that $\widehat{W}+\epsilon < 1-2\kappa$. Note that such an $\epsilon > 0$ exists since $\overline{B} \equiv 1-(1-2\rho)\Delta$ and $\widehat{W} \equiv \overline{B} - 2\kappa$. Assume $M \geq \overline{M} \equiv \overline{M}_\epsilon$. Note then that there exists $\mu > 0$ and $\eta > 0$ such that, for all $V \leq \widehat{W} + \epsilon$, $\frac{\delta}{1-\delta}V \leq V + 2\kappa - \eta$ for all $\delta \in [\underline{\delta}, \underline{\delta} + \mu]$.[33] We now show that, for all $\delta \in [\underline{\delta}, \underline{\delta} + \mu]$, $\overline{\Phi}_{T-M+1} \leq \widehat{W} + \epsilon$. From our arguments above, we know that there exists $t \in [T-M+2, T]$ such that $\overline{\Phi}_t \leq \widehat{W} + \epsilon$. Note then that

$$\begin{aligned} \overline{\Phi}_{t-1} &= (1-\delta)(b_{t-1} - 2\kappa) + \delta\overline{\Phi}_t \\ &\leq (1-\delta)\left(\frac{\delta}{1-\delta}\overline{\Phi}_t - 2\kappa\right) + \delta\overline{\Phi}_t \\ &\leq \widehat{W} + \epsilon - (1-\delta)\eta, \end{aligned}$$

where the first inequality uses $b_{t-1} \leq \frac{\delta}{1-\delta}\overline{\Phi}_t$, and the second inequality uses $\overline{\Phi}_t \leq \widehat{W} + \epsilon$ and $\frac{\delta}{1-\delta}\overline{\Phi}_t \leq \overline{\Phi}_t + 2\kappa - \eta$. If $t-1 = T-M+1$, we are done. Otherwise, we can repeat the same argument until we obtain $\overline{\Phi}_{T-M+1} < \widehat{W} + \epsilon$.

Finally, we show that for all $\varepsilon > 0$ and all $\delta \in [\underline{\delta}, \underline{\delta}+\mu]$, there exists $\overline{T}$ such that $\overline{V}^{\mathsf{bmk}}_{T,M} < \varepsilon$ for all $M > \overline{M}$, $T - M > \overline{T}$.

Fix $\sigma \in \Sigma_{T,M}(E)$ and a history h_t with $t \leq T - M$. Let $b_{(1),t}$ denote the winning bid at h_t. Note that, if both firms participate, then for $i = 1, 2$ we must have

$$(1-\delta)(x_{i,t}b_{(1),t} - \kappa) + \delta W_{i,t+1} \geq (1-\delta)(b_{(1),t} - \kappa),$$

where $x_{i,t}$ denotes the probability with which i wins the auction at time t and $W_{i,t+1}$ denotes

[33] To see why, note that, for all $V \leq \widehat{W} + \epsilon$, we have

$$V + 2\kappa - \frac{\underline{\delta}}{1-\underline{\delta}}V = 2\kappa\left(1 - \frac{V}{1-2\kappa}\right) \geq 2\kappa\left(1 - \frac{\widehat{W}+\epsilon}{1-2\kappa}\right) > 0,$$

where the equality follows since $\frac{\underline{\delta}}{1-\underline{\delta}} = \frac{1}{1-2\kappa}$, and the strict inequality follows since $\widehat{W} + \epsilon < 1 - 2\kappa$. Hence, for $\eta > 0$ and $\mu > 0$ small enough, we have that $\frac{\delta}{1-\delta}V \leq V + 2\kappa - \eta$ for all $\delta \in [\underline{\delta}, \underline{\delta} + \mu]$.

i's continuation value given history $h_{t+1} = h_t \sqcup \mathbf{b}_t$ (including possible penalties). Summing across both players, and using $x_{1,t} + x_{2,t} \le 1$ and $W_{1,t+1} + W_{2,t+1} = W_{t+1}$, we get

$$(1-\delta)(b_{(1),t} - 2\kappa) + \delta W_{t+1} \ge 2(1-\delta)(b_{(1),t} - \kappa)$$
$$\iff \frac{\delta}{1-\delta} W_{t+1} \ge b_{(1),t}.$$

For each value $W \in \mathbb{R}_+$, define operator $\Psi : \mathbb{R}_+ \to \mathbb{R}_+$ as

$$\Psi(W) \equiv (1-\delta)\left(\max\left\{\min\left\{1, \frac{\delta}{1-\delta}W\right\} - 2\kappa, 0\right\}\right) + \delta W.$$

Then, for any history h_t with $t \le T - M$, $\Psi(W)$ is an upper bound to the cartel's payoff at a history h_t whenever the cartel's continuation value at $t+1$ is bounded above by W. Since by our arguments above cartel's continuation payoff at any history h_{T-M+1} is bounded above by $\overline{\Phi}_{T-M+1} \le \widehat{W} + \epsilon$, we have that $\overline{V}^{\mathsf{bmk}}_{T,M} \le \Psi^{T-M}(\widehat{W} + \epsilon)$. Recall that $\frac{\delta}{1-\delta} W \le W + 2\kappa - \eta$ for all $W \le \widehat{W} + \epsilon$. Hence, for all $W \le \widehat{W} + \epsilon$ we have $\Psi(W) \le \max\{W - (1-\delta)\eta, \delta W\}$. And so, for $T - M$ sufficiently large, $\varepsilon > \Psi^{T-M}(\widehat{W} + \epsilon) \ge \overline{V}^{\mathsf{bmk}}_{T,M}$. This completes the proof.

■

References

Abdulkadiroglu, A. and K.-S. Chung (2003): "Auction design with tacit collusion," *Columbia and Northwestern University Working Paper*.

Abrantes-Metz, R., L. Froeb, J. Geweke, and C. Taylor (2005): "A variance screen for collusion. US Federal Trade Commission," Tech. rep., Bureau of Economics Working Paper.

Abreu, D., D. Pearce, and E. Stacchetti (1990): "Toward a theory of discounted repeated games with imperfect monitoring," *Econometrica: Journal of the Econometric Society*, 1041–1063.

Ashcroft v. Iqbal (2009): *United States Reports*.

Athey, S. and K. Bagwell (2008): "Collusion with persistent cost shocks," *Econometrica*, 76, 493–540.

Bajari, P. and L. Ye (2003): "Deciding between competition and collusion," *Review of Economics and Statistics*, 85, 971–989.

Baker, J. and D. Rubinfeld (1999): "Empirical Methods in Antitrust Litigation: Review and Critique," *American Law and Economics Review*, 1, 386–435.

Bell Atlantic v. Twombly (2007): *United States Reports*.

Besanko, D. and D. F. Spulber (1989): "Antitrust enforcement under asymmetric information," *The Economic Journal*, 99, 408–425.

Chassang, S., K. Kawai, J. Nakabayashi, and J. Ortner (2022): "Robust screens for noncompetitive bidding in procurement auctions," *Econometrica*, 90, 315–346.

Chassang, S. and J. Ortner (2019): "Collusion in auctions with constrained bids: Theory and evidence from public procurement," *Journal of Political Economy*, 127, 2269–2300.

Che, Y.-K., D. Condorelli, and J. Kim (2018): "Weak cartels and collusion-proof auctions," *Journal of Economic Theory*, 178, 398–435.

Che, Y.-K. and J. Kim (2006): "Robustly Collusion-Proof Implementation," *Econometrica*, 74, 1063–1107.

——— (2009): "Optimal collusion-proof auctions," *Journal of Economic Theory*, 144, 565–603.

Conley, T. G. and F. Decarolis (2016): "Detecting bidders groups in collusive auctions," *American Economic Journal: Microeconomics*, 8, 1–38.

Conley v. Gibson (1957): *United States Reports*.

Cunningham, C., F. Ederer, and S. Ma (2021): "Killer acquisitions," *Journal of Political Economy*, 129, 649–702.

Cyrenne, P. (1999): "On antitrust enforcement and the deterrence of collusive behaviour," *Review of Industrial Organization*, 14, 257–272.

Ellison, G. (1994): "Theories of cartel stability and the joint executive committee," *The Rand journal of economics*, 37–57.

Harrington, J. E. (2004): "Cartel pricing dynamics in the presence of an antitrust authority," *RAND Journal of Economics*, 651–673.

——— (2008): "Detecting cartels," *Handbook of antitrust economics*, 213, 245.

Imhof, D., Y. Karagök, and S. Rutz (2016): "Screening for bid-rigging-does it work?" Tech. rep., Université de Fribourg.

——— (2018): "Screening for Bid Rigging —- Does It Work?" *Journal of Competition Law & Economics*, 14, 235–261.

Ishii, R. (2008): "Collusion in repeated procurement auction: a study of a paving market in Japan," Tech. rep., ISER Discussion Paper, Institute of Social and Economic Research, Osaka University.

——— (2009): "Favor exchange in collusion: Empirical study of repeated procurement auctions in Japan," *International Journal of Industrial Organization*, 27, 137–144.

Kaplow, L. (2011): "Optimal proof burdens, deterrence, and the chilling of desirable behavior," *American Economic Review*, 101, 277–80.

Kawai, K. and J. Nakabayashi (2018): "Detecting Large-Scale Collusion in Procurement Auctions," *Available at SSRN 2467175*.

Kawai, K., J. Nakabayashi, and J. Ortner (2022a): "The Value of Privacy in Cartels: An Analysis of the Inner Workings of a Bidding Ring," .

Kawai, K., J. Nakabayashi, J. Ortner, and S. Chassang (2022b): "Using Bid Rotation and Incumbency to Detect Collusion: A Regression Discontinuity Approach," *Review of Economic Studies*.

Krasnokutskaya, E. and K. Seim (2011): "Bid preference programs and participation in highway procurement auctions," *The American Economic Review*, 101, 2653–2686.

LaCasse, C. (1995): "Bid rigging and the threat of government prosecution," *The RAND Journal of Economics*, 398–417.

Marshall, R. C. and L. M. Marx (2009): "The vulnerability of auctions to bidder collusion," *The Quarterly Journal of Economics*, 124, 883–910.

McCutcheon, B. (1997): "Do meetings in smoke-filled rooms facilitate collusion?" *Journal of Political Economy*, 105, 330–350.

Mena-Labarthe, C. (2015): "Mexican Experience in Screens for Bid-Rigging," *Competition Policy International*.

OECD (2018): "Summary of the workshop on cartel screening in the digital era," .

Pavlov, G. (2008): "Auction design in the presence of collusion," *Theoretical Economics*, 3, 383–429.

Porter, R. H. (1983): "A study of cartel stability: the Joint Executive Committee, 1880-1886," *The Bell Journal of Economics*, 301–314.

——— (2005): "Detecting collusion," *Review of Industrial Organization*, 26, 147–167.

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.

Wollmann, T. G. (2019): "Stealth consolidation: Evidence from an amendment to the Hart-Scott-Rodino Act," *American Economic Review: Insights*, 1, 77–94.