



ARTICLE

Cartel Screening and Machine Learning

Joseph E. Harrington, Jr.* & David Imhof**

Abstract. This paper discusses a growing tool of interest for cartel detection: examining market data for evidence of collusion, or what is referred to as cartel screening. Screening identifies collusive patterns in firm conduct such as prices and bids. The first half of the paper describes what to look for in the data, more specifically it features collusive markers, structural breaks, and anomalies. A collusive marker is a pattern in the data more consistent with collusion than competition. A structural break is an abrupt change in the data-generating process that could be due to cartel birth, death, or disruption. An anomaly is a pattern in the data that is inexplicable or inconsistent with competition. The second half of the paper focuses on the recent use of machine learning algorithms to develop more effective screens by extracting the most informative patterns from the data, which then instruct us what to look for in the data. With access to a data set comprising episodes of collusion and competition, supervised learning can identify patterns indicative of collusion. Proof of concept is shown based on work of the Swiss Competition Commission using data from construction cartels in Switzerland. Guidance is provided for other competition authorities to deploy machine learning algorithms, including deep learning, to make cartel screening more effective.

* **Joseph E. Harrington, Jr.** is Patrick T. Harker Professor, Department of Business Economics & Public Policy, The Wharton School, University of Pennsylvania.

** **David Imhof** is an economist at the Swiss Competition Commission.

I. Introduction

When Mark Whitacre informed the FBI of the lysine cartel in November 1992, there was little thought among competition authorities that global cartels were a concern beyond the few that operated openly and legally, such as in crude oil (OPEC) and potash. However, it was not long before the prosecution pipeline was flowing with global cartels.¹ All along, these cartels had been operating beneath the radar of competition authorities. The tremendous growth in cases and convictions since that time is partly attributable to the adoption of leniency programs, the introduction of competition laws in many countries, and a feedback loop whereby the success in convicting cartels motivates competition authorities to invest more resources in finding and prosecuting cartels.

Though the story seems a most heartening one, there is a shadow looming over it: *cartels continue to form*. How many is unknown, for we only know those that are discovered, which tells us nothing about how many go undiscovered.² Still, sufficiently many cartels have been discovered in the last decade to keep competition authorities and plaintiff law firms busy and to remind us that the enforcement of laws against cartels continues to require the attention of practitioners, policymakers, and scholars.

Why do cartels continue to form in spite of the enforcement successes? It could be that some firms have not gotten the memo of heightened enforcement, or perhaps it is that penalties are not severe enough to make collusion an unprofitable activity (which is surely true in some jurisdictions),³ or it could be that cartel detection is not sufficiently likely. Motivated by this last possibility, this paper discusses a growing tool of interest for cartel detection: examining market data for evidence of collusion, or what is referred to as *cartel screening*. While there have been episodic ventures into screening going back more than 50 years, it has only meaningfully been employed by competition authorities in the last decade or so.⁴ To promote the use of screening, we offer an overview and explain how machine learning can increase its effectiveness.

¹ “In 1993, when the lysine investigation was underway, there were only a handful of Division investigations of alleged international cartel activity. By comparison, the Division currently has roughly 30 grand juries investigating suspected international cartel activity. Similarly, at the time of the lysine investigation, less than 2 percent of our corporate defendants were foreign-based, as compared to nearly 50 percent of our corporate defendants last year.” Scott D. Hammond, Director of Criminal Enforcement, Antitrust Division, U.S. Department of Justice, ‘Fighting Cartels - Why and How? Lessons Common to Detecting and Deterring Cartel Activity’ (The 3rd Nordic Competition Policy Conference, 12 September 2000) <https://www.justice.gov/atr/speech/fighting-cartels-why-and-how-lessons-common-detecting-and-deterring-cartel-activity>

² This is a point not understood by many economists and lawyers. See Joseph E. Harrington, Jr. & Yanhao Wei, *What Can the Duration of Discovered Cartels Tell Us About the Duration of all Cartels?*, 127 *ECON. J.* 1983–1984 (2017).

³ Towards enhancing penalties, structural remedies have recently been proposed. See Joseph E. Harrington, Jr., *A Proposal for a Structural Remedy for Illegal Collusion* 82 *ANTITRUST L. J.* 335–359 (2018).

⁴ At the 2016 ICN Chief/Senior Economist Workshop, 27 competition authorities in attendance were surveyed and 15 reported they were doing some screening. Nigel Caesar, René Duplantis & Thomas Ross, ‘Report on ICN Chief/Senior Economists Workshop’ (International Competition Network Chief/Senior Economists Workshop, 12–13 September 2016) https://www.internationalcompetitionnetwork.org/wp-content/uploads/2018/10/AEWG_EconWorkshop2016Report.pdf

Screening is designed to be a low cost method for identifying markets where a cartel may be present. It is not meant to deliver conclusive evidence regarding the presence of a cartel but rather to be the basis for an investigation that conducts an intensive economic analysis and collects non-economic evidence (such as through a dawn raid). Examples of screening discovering cartels include those in cement (South Africa), subway construction (Korea), and retail gasoline (Brazil).⁵ The South African cement cartel is instructive as screening produced evidence to justify a dawn raid which then induced leniency applications and yielded convictions. Any investigation begins with some piece of evidence supporting the hypothesis that a market harbors a cartel. Screening can deliver that evidence.

While screening is largely thought of as a tool for competition authorities, it is also relevant to companies (who may be victims of cartels), law firms, and economic consultancies (who can offer screening services for their clients). Relevant to this point, the German railway company Deutsche Bahn is currently developing a screening program, which is a response to cartels among its input suppliers.⁶ For that reason, we encourage private actors, as well as public enforcers, to read on.⁷

II. Overview of Screening

A. General Approaches to Screening

There are two general approaches to screening.⁸ *Structural screening* identifies markets with structural traits conducive to or associated with collusion. Based on theory and empirical evidence, relevant traits include a small number of firms, homogeneous products, stable demand, and excess capacity, among others. The weakness to structural screening is that we have neither a sufficiently informed theory nor sufficiently rich data to predict where collusion emerges. To appreciate the basis for this pessimistic claim, consider the “ideal” market for collusion: two firms, homogeneous products, stable demand, etc. As that market “checks all of the boxes” according to the usual variables, a structural screening approach would predict that a cartel is very likely when, in practice, probably only a small fraction of such markets have cartels. The problem is one of omitted variables in existing theoretical and empirical models of cartel formation. Until we better understand the determinants of cartel formation and have data on more factors, a structural screening approach is unlikely to be an effective method of cartel detection.

Behavioral screening identifies collusive patterns in firm conduct and outcomes such as prices and market shares. It focuses on the outcome of collusion, as opposed to elements of the environment that facilitate cartel formation as with structural screening. Behavioral screening is in the spirit of methods used to uncover tax evasion (e.g., looking for irregular deductions), accounting fraud (e.g., looking for

⁵ See Ulrich Laitenberger & Kai Hüschelrath, *The Adoption of Screening Tools by Competition Authorities*, CPI ANTITRUST CHRONICLE, Sept. 2, 2011.

⁶ Hannes Beth & Thilo Reimers, *Screening Methods for the Detection of Antitrust Infringements*, COMPLIANCEBUSINESS – DAS ONLINE-MAGAZIN, Oct. 31, 2019.

⁷ The case for private actors engaging in cartel screening is put forth in Joseph E. Harrington, Jr., *Cartel Screening is for Companies, Law Firms, and Economic Consultancies, not Just Competition Authorities*, CENTRO COMPETENCIA, Nov. 3, 2021 <https://centrocompetencia.com/harrington-cartel-screening-is-for-companies-law-firms-and-economic-consultancies/>.

⁸ See Joseph E. Harrington, Jr., *Detecting Cartels* in HANDBOOK OF ANTITRUST ECONOMICS (Paolo Buccirossi ed., 2008) 213.

violations of Benford's Law), insider trading (e.g., looking for excessive trading volume prior to a company announcement), and credit card fraud (e.g., looking for anomalous spending patterns).

One way to describe the distinction in screening approaches is that structural screening seeks to identify markets for which it is more likely that a cartel *will form*, while behavioral screening seeks to identify markets for which a cartel *has formed*; the latter is a far easier task. Though this paper focuses on the use of behavioral screening to detect cartels, it is worth noting that structural screening can be used in combination with behavioral screening in that behavioral screens can be applied to those markets flagged by a structural screen.

B. Overview of Behavioral Screening Methods

Behavioral screening can work for two fundamental reasons. First, collusion must mean a change in the data-generating process, which, in principle, can be detected. Second, collusion imposes a unique set of challenges for firms, and their conduct in meeting those challenges leaves an evidentiary trail. Furthermore, screening can work even when cartelists are strategic and try to avoid being detected. For example, cartels in markets for intermediate goods often gradually raise prices, and this is probably so as not to create suspicions among industrial buyers that there is a cartel.⁹ Though raising the price more slowly makes detection less likely, it means foregoing some profits. Thus, a cartel will never go so far as to minimize the probability of detection, for that would mean not raising the price at all! A strategic cartel can reduce, but not eliminate, the power of most screens.

Behavioral screening requires *data* and *knowing what to look for in the data*. The trick to cost effective screening is to use easily available data and apply simple empirical methods that can be automated. In practice, available data generally means prices (or bids in an auction setting) but can also include quantities (and market shares) and some cost information (e.g., input prices that are publicly available). The focus of this paper will be on the use of price and bid data.

In terms of what to look for in the data, there are three approaches: collusive markers, structural breaks, and anomalies. A *collusive marker* is a pattern in the data more consistent with collusion than competition. A *structural break* is an abrupt change in the data-generating process. While there is a multitude of sources of such a change, a structural break could be identifying cartel birth, death, or disruption. Finally, an *anomaly* is a pattern in the data that is inexplicable or inconsistent with competition but may ultimately be found consistent with collusion.

⁹ JOHN M. CONNOR, GLOBAL PRICE FIXING (2008).

C. Collusive Markers

The extensive body of theoretical and empirical work on collusion has produced a wealth of hypotheses for distinguishing collusion from competition.¹⁰ Here we illustrate collusive markers by focusing on a few more practical and useful ones.

The first and foremost implication of collusion is higher prices. An obvious screen is to compare a market's prices with some benchmark competitive price. A setting for which this marker is applicable is when a product is sold in geographically distinct markets, so prices in other (appropriately similar) markets can serve as that benchmark. In that case, a market would be flagged when the price is high relative to the average price (across all markets).¹¹

Collusion does not just entail high prices but also a rich set of pricing dynamics. Some cartels have exhibited a V-shaped pattern to prices where price significantly falls (prior to cartel formation) then rises (just after cartel formation). This pattern is exemplified in Figure 1 for the citric acid cartel. The pre-cartel price decline may, in fact, be the cause of cartel formation. There could have been an intensification of competition due to the collapse of tacit collusion (which has been supplanted with explicit collusion) or a decrease in demand which, with excess capacity and low marginal costs, led to drastically lower prices. Having formed a cartel after this price decline, firms then gradually raise prices as they work their way to a steady-state collusive price. Of course, these pricing dynamics could be driven by other factors (such as dynamics associated with input prices or demand), but then a screen is only asked to flag a market for further investigation. If, in fact, some cost or demand factor is driving such large price movements, it should not be difficult to find that factor and thereby dismiss a change in firm conduct being its source.

¹⁰ For reviews, the reader is referred to Robert H. Porter, *Detecting Collusion*, 26 REV. INDUS. ORG. 147-167 (2005); Margaret C. Levenstein & Valerie Y. Suslow, *What Determines Cartel Success?*, 44 J. ECON. LITERATURE 43-95 (2006); and Harrington, *supra* note 8.

¹¹ That type of exercise was done to determine whether prices were supracompetitive in retail gasoline markets in Italy where the benchmark was the average EU price. See Patrick Andreoli-Versbach & Jens-Uwe Franck, *Endogenous Price Commitment, Sticky and Leadership Pricing: Evidence from the Italian Petrol Market*, 40 INT'L J. INDUS. ORG. 32-48 (2015).

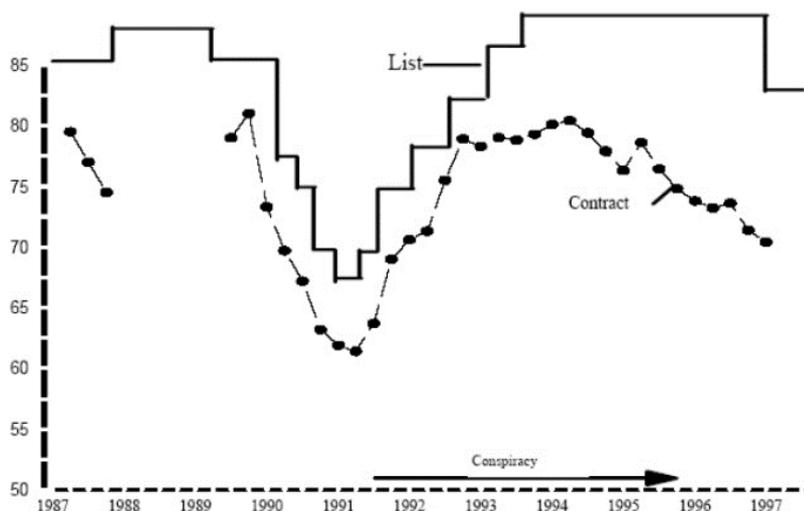


Figure I: V-Shaped Price Pattern for the Citric Acid Cartel¹²

Collusion affects not only the average price but also its variability. After having gone through the post-formation transition phase (as depicted in Figure 1), collusive prices tend to be more stable than competitive prices. Under competition, a firm is likely to change the price in response to all substantive cost and demand shocks; that is, a competitive firm changes its price whenever it is profitable to do so. This process is very different in the case of collusion, as firms are constrained by the cartel. They may only change the price in response to *common* cost and demand shocks when it does not endanger cartel stability and only after cartel members communicated and coordinated. These different rationales and protocols for changing prices can make them less responsive to cost and demand shocks under collusion. Figure 2 shows this property for the urethane cartel, where urethane prices (red line) and the urethane input cost (green line) are plotted. (The blue line is the estimated competitive price and is not relevant to this discussion.) Under competition, price and cost move together, so there is a fair amount of volatility in prices coming from cost variability. In comparison, collusive prices are very stable in spite of wide cost fluctuations.

¹² John Connor, *What Can We Learn from the ADM Global Price Conspiracies?* (Purdue Univ., Dep't of Agric. Econ., Working Paper No. 98-14, 1998).

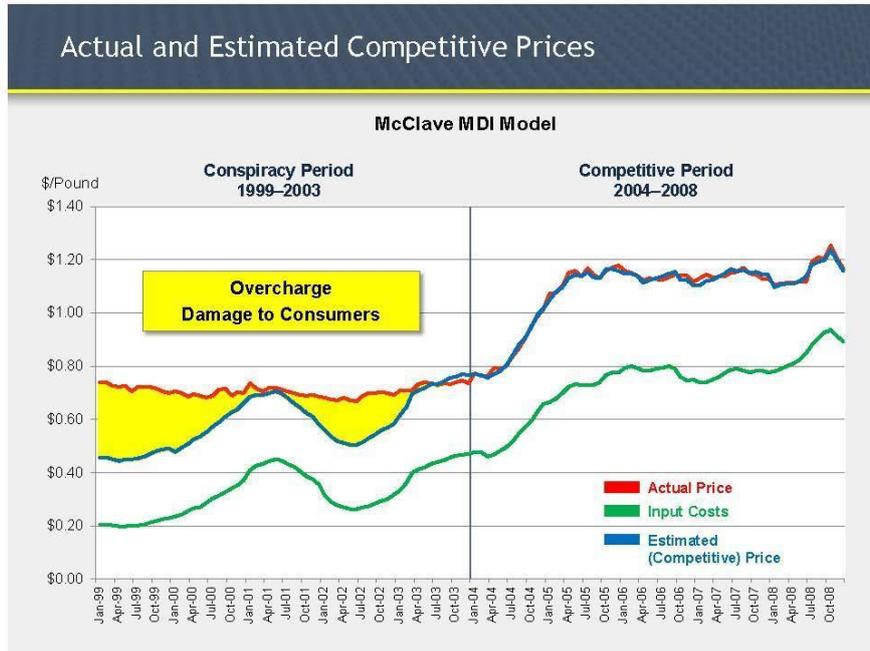


Figure II: Low Price Variability for the Urethane Cartel¹³

As collusion raises the price and lowers price volatility, it is not surprising that a particularly effective collusive marker is the price coefficient of variation (or bids if one is examining auction data). The coefficient of variation (CV) is the ratio of the standard deviation of a random variable (such as price) to its mean.¹⁴ As collusion lowers the numerator and raises the denominator, the price CV is predicted to be lower – perhaps significantly lower – under collusion. Table I shows the stark effect that collusion had on the price CV of the frozen perch cartel. The CV with collusive prices was more than four times smaller than with competitive prices. Furthermore, this differential cannot be explained by cost changes; the cost CV was only modestly lower during the years the cartel was operating, while the price CV was significantly lower.

	Price		Cost	
	Cartel years	Post-cartel years	Cartel years	Post-cartel years
Mean	3.544	2.970	0.722	0.771
Standard deviation	0.078	0.283	0.114	0.173
Coefficient of variation	0.022	0.095	0.158	0.221

Table I: Low Price Coefficient of Variation for the Frozen Perch Cartel¹⁵

¹³ Class Plaintiffs' Response Brief, *In re: Urethane Antitrust Litigation*, 18 (10th Cir. 2014)

¹⁴ Here the price CV is measured across time but it can also be measured across firms. The latter is considered in Section III.

¹⁵ Rosa M. Abrantes-Metz et al., *A Variance Screen for Collusion*, 24 INT'L J. INDUS. ORG. 467-486 (2006).

D. Structural Breaks

A structural break refers to a discrete change in the data-generating process. As a screen, one is looking for a change in the process generating prices or bids (or some other relevant variable) that is attributable to cartel birth, death, or disruption. A cartel could be detected at birth through, for example, higher prices, more stable prices, or a change in how prices respond to cost and demand factors. Though a savvy cartel could manage the transition from competitive to collusive prices in order to make cartel formation less transparent to third parties, it can only partially diminish the power of the screen. Collusion must change the price-generating process in order to be profitable, which means there is, in principle, a detectable structural break. Just as screening can pick up cartel birth, it can also detect cartel death. In fact, cartel death and its often sharply lower prices are a more promising avenue for detection because its collapse means the cartel cannot manage prices to make detection less likely, as it can do with cartel birth. Finally, cartels can be detected through temporary disruptions to collusion. An internal disruption by cartel members could take the form of a temporary price war in response to non-compliance by some cartel members. An external disruption could come from non-cartel members that may cause collusion to temporarily break down or cartel members to engage in aggressive pricing against those firms. With either internal or external disruptions, there is regime switching as prices go from being set collusively to competitively back to collusively. Of course, finding a structural break in the pricing process does not necessarily mean having found collusion for the break could be caused by a persistent change in input prices or demand. Examining these alternative hypotheses would be part of the next phase of the investigation. The role of the screen is to identify a market for closer inspection.

Rather than review the well-established statistical methods for identifying structural breaks, we will offer two episodes for which simply plotting the data is sufficient to detect a cartel. One is associated with cartel death and the other with cartel disruption.¹⁶

In Mexico, the largest public health provider purchased generic drugs by conducting a procurement auction with the supply contract going to the lowest bidder.¹⁷ It turns out the procurer was paying supracompetitive prices, which was revealed upon the cartel's collapse. Figure 3 shows the price paid for two of these drugs: insulin (drug 1) and calcium (drug 2). The vertical line marks the end of collusive bidding, and one can see a striking change in the data. Initially, the price was high and stable across tenders, and then suddenly, it was much lower and more variable. No sophisticated empirical analysis is needed to conclude there has been a radical change in how firms bid and the natural hypothesis is that they had previously coordinated their bids. An investigation would certainly be warranted based on this evidence.

¹⁶ As an example of the use of a structural break statistical test, see Carsten J. Crede, *A Structural Break Cartel Screen for Dating and Detecting Collusion*, 54 REV. INDUS. ORG 543-574 (2019).

¹⁷ Ernesto Estrada & Samuel Vazquez, *Bid Rigging In Public Procurement Of Generic Drugs In Mexico*, 9 COMPETITION POL'Y INT'L J. 100-122 (2013).

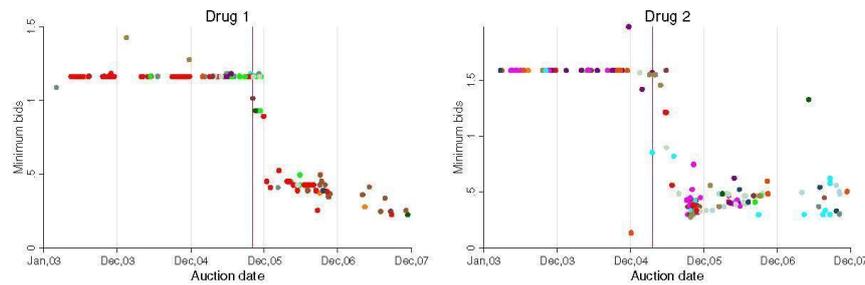


Figure III: Change in Generic Drug Prices After Cartel Death¹⁸

The second case involves local governments conducting procurement auctions for road paving contracts.¹⁹ As part of the auction design, the procurer sets both maximum allowed and minimum allowed bids for each auction. Using data for 139 tenders, Figure 4 reports the winning bid and the minimum allowed bid as a percentage of the maximum allowed bid. In 123 out of the 139 tenders, the winning bid ranged from 91% to 95% of the maximum allowed bid. In the other 16 tenders, the winning bid equaled the minimum bid. Notice the striking gap in the data: Either the winning bid was high (more than 91% of the maximum allowed bid) or low (less than 85% of the maximum allowed bid). The explanation is that there was a cartel that periodically shifted to competition, either because of failure to agree or due to some non-cartel firms having participated, which led the cartel members to bid aggressively (either knowing that they had to compete or to discourage the non-cartel firms from participating in the future). The structural break here takes the form of periodic switching between regimes of collusion and competition.

¹⁸ Ernesto Estrada and Samuel Vazquez, *Bid Rigging In Public Procurement Of Generic Drugs In Mexico*, draft, undated.

¹⁹ Reiko Ishii, *Collusion in Repeated Procurement Auction: A Study of a Paving Market in Japan* (INST. SOC. ECON. RSCH. Working Paper, Paper No. 710, 2008), https://papers.ssrn.com/sol3/papers.cfm?abstract_id=1148064

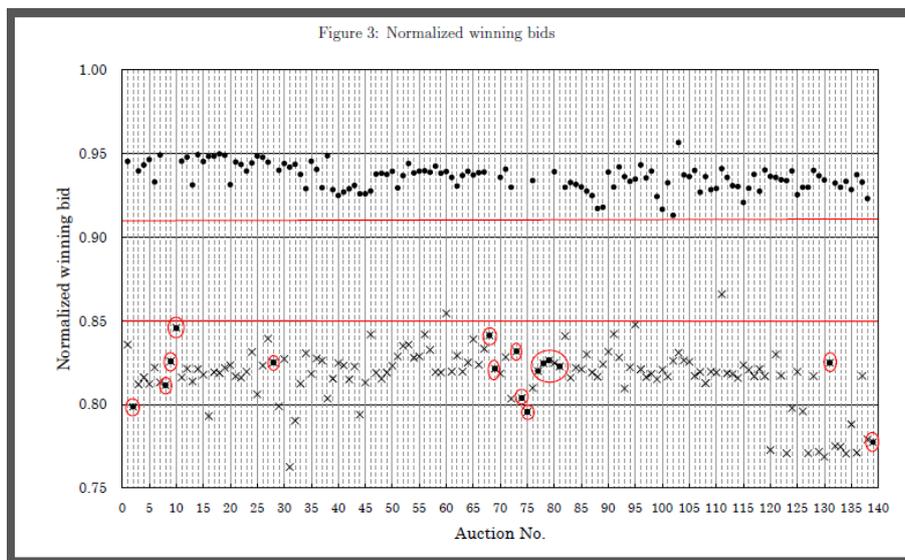


Figure IV: Periodic Shifts in Winning Bids from Collusion to Competition at Road Paving Tenders²⁰

Legend:

- denotes the winning bid
- × denotes the minimum allowed bid
- denotes when winning bid = minimum allowed bid

E. Anomalies

An anomaly is a pattern in the data that is inexplicable or inconsistent with that expected under competition. An example of an inexplicable pattern is firms systematically not charging certain prices such as when Nasdaq market makers avoided bid and ask quotes ending in an odd-eighth.²¹ An example of a pattern inconsistent with competition is charging lower prices when the cost is higher, such as when some milk suppliers submitted lower bids on school milk tenders for districts that were farther away from their plants (and thus had higher transportation costs).²² Of course, some data can be inexplicable or inconsistent with competition but not be consistent with collusion. Nevertheless, finding an anomaly in the data (especially price data) calls for an explanation that may justify an investigation. Where that investigation might lead is unclear though the discovery of a cartel is one possibility. With Nasdaq, it was found that avoidance of odd-eighths was a simple rule for supporting a higher price-cost markup. In the case of school milk tenders, the explanation was that the cartel members colluded at school district tenders close to them (so bids were high) but competed with non-cartel members at more distant school district tenders (so bids were low).

²⁰ *Id.*

²¹ William G. Christie & Paul H. Schultz, *Why do NASDAQ Market Makers Avoid Odd-Eighth Quotes?*, 49 J. FIN. 1813 (1994).

²² Robert H. Porter & J. Douglas Zona, *Ohio School Milk Markets: An Analysis of Bidding*, 30 RAND J. ECON. 263 (1999).

To illustrate this screening approach, let us look at a case in which the data pattern was initially seen as inexplicable, then recognized to be inconsistent with competition, and finally understood to be consistent with collusion.

The dataset is thousands of government procurement auctions (mostly construction contracts) over a five-year period. The auction format is that bidders submit bids, and the tender is awarded to the bidder with the lowest bid (as long as it does not exceed a secret maximum allowed bid). Using all bids from all auctions, Figure 5 plots the histogram for the variable $x_{i,t} = b_{i,t} - \min_{j \neq i} \{b_{j,t}\}$ where $b_{i,t}$ is the bid of bidder i in auction t (divided by the maximum allowed bid) and $\min_{j \neq i} \{b_{j,t}\}$ is the lowest bid of the other bidders in auction t (divided by the maximum allowed bid). Dividing them by the maximum allowed bid makes bids comparable across heterogeneous tenders. When $x_{i,t} > 0$ then this bidder's bid was not the lowest, and when $x_{i,t} < 0$ then it did submit the lowest bid. For each numerical value of $x_{i,t}$ on the horizontal axis, the vertical axis measures how many observations there were with that numerical value.

The most frequent observations are around .02, which means a bidder's bid was higher than the lowest bid of other bidders by 2% (of the maximum allowed bid). Generally, more extreme values are less common: As $x_{i,t}$ rises above .02 or falls below .02, the frequency of observation falls *except when $x_{i,t}$ is around zero*. There are very few instances in which a bidder's bid is close to the minimum bid of the other bidders. These “missing bids” are inexplicable. Upon some reflection, it is clear that these bids are contrary to profit-maximizing conduct on the part of bidders. For suppose a bidder anticipated this property. Conditional on having the lowest bid, the bidder would do better to bid slightly higher; they would still win but at a higher bid. Hence, this anomaly is inconsistent with competition. Finally, it is consistent with a collusive scheme in which there is a designated winner who shares its planned bid with the other bidders, who then submit cover bids that are distinctly above that planned bid in order to ensure the designated firm wins the auction. Those cover bids create a gap between the lowest bid and the next lowest bid.

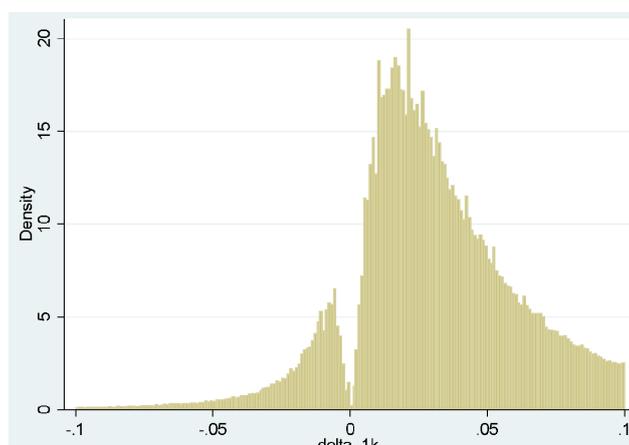


Figure V: Anomalous Missing Bids in Procurement Auctions²³
 Horizontal axis: (own bid minus minimum of other bidders' bids)/maximum allowed bid

III. Detecting Bid-Rigging Cartels with Algorithms

A. Machine Learning and Collusive Markers: A Tender-based Approach

Bid-rigging cartels remain a central issue for many competition agencies.²⁴ The Swiss Competition Commission (hereafter, COMCO) decided in 2008 to develop and implement tools for screening procurement data in order to decrease its dependency on external sources of information.²⁵ In designing these proactive tools, they should be reliable, easy to explain (especially to lawyers and judges), and require only publicly available data. The goal of COMCO consists of building a detection method easily replicable on a large scale while relying on publicly available data (so that data collection does not attract the attention of suspected cartel participants). In the procurement context, publicly available data are typically the bid summaries (also called bid results or official records of the bid opening). The bid summaries provide the following information: bids, bidder identities, contract type, name and location of the contract, date, and tendering procedure. Available in Switzerland, bid summaries are also likely to be accessible in many other jurisdictions.

Since no suitable detection method or statistical tool was directly available, COMCO decided to develop a heuristic detection method based on descriptive statistics (collusive markers) and the observation of collusive behavior in past investigations. They relied on a simple but fundamental hypothesis: Bid rigging

²³ Sylvain Chassang, Kei Kawai, Jun Nakabayashi, & Juan M. Ortner, *Data Driven Regulation: Theory and Application to Missing Bids* (Nat'l Bureau of Econ. Rsch., Working Paper No. 25654, 2019), <https://ideas.repec.org/p/nbr/nberwo/25654.html#>. The published version is Sylvain Chassang, Kei Kawai, Jun Nakabayashi, & Juan Ortner, *Robust Screens for Noncompetitive Bidding in Procurement Auctions*, 90 *ECONOMETRICA* 315 (2022).

²⁴ For a recent survey of competition authorities' use of computational tools for detecting bid rigging, see Thibault Schrepel & Teodora Groza, *The Adoption of Computational Antitrust by Agencies: 2021 Report*, 2 *STAN. COMPUTATIONAL ANTITRUST* 78 (2022).

²⁵ External sources of information include whistleblowers, leniency programs, and complaints from customers and procurement agencies.

affects the distribution of bids. Bid manipulation or bid coordination among cartel participants generally does not reflect the underlying costs of each bidder, while competitive bids do. Because the bid-generating processes differ, the distribution of collusive bids will differ from that of competitive bids. To capture this divergence, COMCO calculated descriptive statistics for the discrete distribution of bids in a tender. For example, the coefficient of variation and the relative distance.²⁶ Other statistics of the bid distribution include kurtosis, skewness, spread, and the percentage difference between the first and second lowest bids, as well as others.²⁷

By comparing the statistics for collusive and competitive tenders, one can identify empirical benchmarks and infer the effects of bid rigging on the bid distribution. Using these statistics along with benchmarks from past investigations, Imhof et al. explain how COMCO flagged a bid-rigging cartel and delivered sufficient evidence for opening an investigation against six firms. Three years after the initiation of the investigation, the firms were fined for bid rigging.

To illustrate the use of a screen, let us examine the relative distance statistic using a real bid-rigging case. The right side of Figure 6 shows the distribution of collusive bids in one tender. As is easily apparent, the difference between the first- and second-best bids (36'800) is much larger than the difference between adjacent losing bids (varying from 3'561 to 6'141). To capture this statistical pattern, we calculate the difference between the first and the second-best bids ($X_2 - X_1$), which is normalized by dividing it by the standard deviation of the losing bids (X_2 to X_6). This suspicion is reinforced if we repeatedly find a similar pattern and observe that bidders in turn successively win tenders. In the case depicted in Figure 6, the relative distance equals 4.97, which leads one to suspect a collusive bidding process.

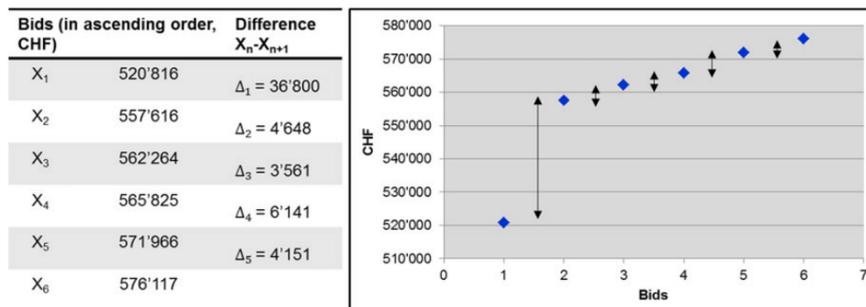


Figure VI: Example of Bid Manipulation Produced by Bid Rigging²⁸

Through its investigations, a competition agency will possess data on competitive and collusive episodes. With that data, it becomes possible to use advanced statistical tools rather than solely relying on markers with benchmarks assessed by human judgment from past investigations. Machine learning

²⁶ David Imhof, Yavuz Karagök, & Samuel Rutz, *Screening for Bid Rigging – Does It Work?*, 14 J. COMPETITION L. & ECON. 235 (2018).

²⁷ See Martin Huber & David Imhof, *Machine Learning with Screens for Detecting Bid-Rigging Cartels*, 65 INT’L J. INDUS. ORG. 277 (2019); David Imhof, *Detecting Bid-Rigging Cartels with Descriptive Statistics*, 15 J. COMPETITION L. & ECON. 427 (2019); and Hannes Wallimann, David Imhof, & Martin Huber, *A Machine Learning Approach for Flagging Incomplete Bid-Rigging Cartels* (arXiv, Working Paper No. 2004.05629v1, 2020), <https://arxiv.org/abs/2004.05629>.

²⁸ Imhof et. al., *supra* note 26.

algorithms, including deep learning, are such advanced tools, and they have been shown to be effective in distinguishing collusive and competitive tenders. The implementation of algorithms for screening could be applied at a large scale and thus reduce the effective costs of screening.

Using machine learning, we predict an outcome Y (also called a response variable) with a set X of predictors as in the following equation: $Y = f(X) + \varepsilon$, where f stands for “function” and ε is the error term reflecting unobserved components. In general, machine learning is not intended to identify a causal relationship between X and Y , but instead, to use X to predict Y . Consequently, high-dimensional data sets and collinearity among predictors are usually not a concern with machine learning. We can therefore use many predictors, which means many collusive markers, even if some might appear highly related. Furthermore, the functional form of f is generally very flexible and (highly) non-linear, which is not a problem since the objective is not to precisely estimate coefficients (that is, measuring the effect of predictors on the outcome) but only to predict the outcome Y . This implies that our approach can remain agnostic about the underlying data-generating process. We do not need to know the effect of the bid coordination among cartel participants to implement an approach with machine learning and markers. All that is required is the maintained hypothesis that the bid-generating process diverges between colluding firms and competing firms.

Concretely, the tender-based approach consists of calculating all possible markers to describe the pattern of bids in a tender. In our case, the outcome Y is binary as the tender is either collusive ($Y = 1$) or competitive ($Y = 0$), depending on whether the tender was in the cartel period or in the competitive period (generally, the pre-cartel or post-cartel regimes). To evaluate the performance of a predictive model, we randomly split our data into two samples: a training sample (say, 75% of the observations) and a testing or validation sample (the other 25%). We then use the markers as predictors and estimate (“train”) the models for recognizing collusive and competitive tenders. Once the best predictive model is selected, it is applied to the test data by comparing the predicted outcome \hat{Y} to the actual observed outcome Y to assess the accuracy of the trained model. Since the outcome is binary, efficacy is measured by the fraction of correct classifications in the test sample (that is, how frequent we find $Y = \hat{Y}$). The procedure is repeated many times (say, 100) where, for each repetition, the data is split into training and testing subsamples. In the end, 100 correct classification rates are calculated, which are then averaged.

In different papers applying machine learning with collusive markers, the following algorithms have been used: the Lasso, the random forest, the Support Vector Machine (SVM), and the ensemble method.²⁹ The random forest and the

²⁹ See Huber & Imhof, *supra* note 27; Wallimann et al., *supra* note 27; David Imhof & Hannes Wallimann, *Detecting Bid-Rigging Coalitions in Different Countries and Auction Formats*, 68 INT’L REV. L. & ECON. 106016 (2021); Martin Huber, David Imhof, & Rieko Ishii, *Transnational Machine Learning with Screens for Flagging Bid-Rigging Cartels, Detecting Bid-Rigging Coalitions in Different Countries and Auction Formats* J. ROYAL STAT. SOC’Y SERIES A (2022).

Lasso³⁰ can help in determining the best predictors for classifying the outcome variable Y .³¹ However, the Lasso needs to calculate penalty terms in a cross-validation step before training models, whereas the random forest does not require penalty terms and is, therefore, a simpler algorithm to use. As calculating penalty terms is not required, the calculation time for the random forest is usually lower. The ensemble method, which is a composite algorithm formed with several weighted algorithms, emerges as one of the most powerful machine learners for predicting collusive and competitive tenders.³² However, higher accuracy also means that calculations are more time-consuming, since we must first determine the weights for the different applied algorithms in the ensemble method.

Huber and Imhof apply the Lasso and the ensemble method to a sample of 584 Swiss tenders and find that the algorithms correctly classify 84% of the tenders, as either collusive or competitive.³³ In other words, armed with this estimated screen and given 20 tenders to evaluate, we can expect to correctly classify 17 of the tenders as either collusive or competitive. This result shows the potential of using machine learning to develop effective screens for procurement markets. Finding a large share of collusive tenders – for example, more than 50% – in a specific region or for a specific type of contract or within a specific group of firms would call for further examination, potentially leading to the opening of an investigation. Finding a share of collusive tenders from 20% to 50% would probably necessitate a tightened scrutiny to discard potential cartel candidates.

Huber and Imhof also find that two markers are particularly informative when it comes to prediction: the normalized distance and the coefficient of variation. The normalized distance is the ratio of the difference between the first and second lowest bids to the average difference between all bids.³⁴ For collusive tenders, the coefficient of variation is lower, which means that bid rigging reduces the distances between bids. Moreover, bid manipulation among cartel participants increases the distance between the first and second lowest bids in a tender and reduces the distance between losing bids. We generally observe higher values for the normalized distance (and for the relative distance) in collusive tenders.

It is important to recognize that Huber and Imhof use data involving cartels that encompass all bidders. Given that all bids are manipulated, the statistical pattern produced by bid rigging is easier for the algorithms to recognize. When non-cartel members bid in tenders along with cartel participants, they might alter the statistical pattern produced by bid rigging. This calls for a more robust method for detecting bid rigging.

³⁰ Since the outcome Y is binary, the logit Lasso is used.

³¹ For details on the algorithms, see JAMES GARETH, DANIELA WITTEN, TREVOR HASTIE & ROBERT TIBSHIRANI, AN INTRODUCTION TO STATISTICAL LEARNING WITH APPLICATIONS IN R 219, 319, 349 (2013). The ensemble method uses the “SuperLearner” package for “R”, see Mark J. van der Laan, Erik C. Polley and Allan E. Hubbard, *Super Learner*, 6(1) STAT. APPLICATION IN GENETICS & MOLECULAR BIOLOGY (2007).

³² See Huber, Imhof & Ishii, *supra* note 29; Imhof and Wallimann, *supra* note 29.

³³ The correct classification rate is calculated by comparing the prediction made in the test data with the actual outcome of the test data. Thus, 84% of the tenders are on average correctly predicted as either collusive or competitive according to the actual outcome.

³⁴ The normalized distance differs from the relative distance as the latter divides the difference between the first and second lowest bids in a tender by the standard deviation of the losing bids. For the example in Figure 6, the normalized distance equals 3.33 and the relative distance equals 4.97.

B. A More Robust Approach for Flagging Cartels

Wallimann et al. propose a more robust approach capable of detecting cartels even when some bidders are not part of the cartel. The incompleteness of the cartel could be due to outsiders periodically participating or the cartel involving only a subset of regular bidders. Moreover, a cartel may not always be stable or operating effectively. It may temporarily collapse or there may be periodic deserters to the cartel who do not conform to the collusive bid plan. All of these departures from an all-inclusive cartel operating all of the time will affect the statistical pattern produced by bid rigging.

Wallimann et al. suggest building subgroups of bidders and calculating markers for each subgroup in a tender. In a second step, summary statistics are calculated - the minimum, the maximum, the mean, and the median of those markers across subgroups. These summary statistics are used to create what we call robust screens.

As an illustration, consider Table 2. With eight bids in a tender, there are 56 possible three-firm subgroups and 70 four-firm subgroups. For three-firm subgroups, 56 values of the coefficient of variation are calculated from which we obtain the minimum, the maximum, the median, and the mean for each tender. Other summary statistics could be used, such as the lower and upper quantiles or additional quantiles.

Bids in a tender	Subgroups formed with three bids	Subgroups formed with four bids
4	4	1
5	10	5
6	20	15
7	35	35
8	56	70
9	84	126
10	120	210

Table II: Example of Possible Subgroups for Three and Four Bidders³⁵

The particular context and the extent of available data will determine the choice of the number of bidders for building the subgroups. Of course, many summary statistics require three or more bids, and thus we cannot consider two-bidder subgroups. The tender-based approach of Section III.A should work well if the cartel is almost all inclusive. The real difficulty is when the cartel leaves more than a firm or two outside of it. For that reason, we will consider the challenging case of small cartels by focusing on subgroups of three and four bidders.

The robust screens have the advantage of isolating the effect of outsiders or undisciplined cartel participants. Let us consider again the example of Figure 6. If bidder 6 had submitted a significantly higher bid and thus increased the variance of bids, this effect would affect the statistics calculated for all subgroups encompassing bidder 6. If we calculated subgroups of three bidders, 10 subgroups including bidder 6 will exhibit a higher variance; whereas the variances for the

³⁵ Wallimann, Imhof & Huber, *supra* note 27.

other 10 subgroups that exclude bidder 6 will remain unaffected, and the statistical pattern of bid rigging undistorted. Consequently, the medians and the means of the markers calculated for the subgroups will moderate the effect of competitive bids; and the minimum or the maximum of the statistics will be unaffected. In our example, the coefficient of variation with six colluding bidders is very similar to the minimum of the coefficient of variation for three-bidder subgroups with five colluding bidders and one competitive outsider. In that manner, the screen is robust.

Wallimann et al. use the Ticino dataset, which includes only complete cartels. Simulating competitive bids and adding them to collusive tenders, they find a potential gain of ten percentage points in correct classification rates with robust screens. This gain translates into a substantial decrease in the error rate (one minus the correct classification rate). For example, when five competitive bids are added to collusive tenders, robust screens increase the correct classification rate from 76% to 86%, which means a 42% reduction in the error rate from 24% to 14%. Using episodes of incomplete cartels with data from two investigations in Switzerland, robust screens improved correct classification rates by 3 to 7.5 percentage points, with a decrease in error rates by around 25%. If we consider the legal consequences and the resources associated with the opening of an investigation, a decrease by one quarter in error rates is a valuable improvement.

The robust screen approach is easy to implement for it requires only bids and does not need the identity of bidders. However, while it is able to flag tenders as conspicuous, without bidder identities it cannot directly infer the participation of a firm in the cartel. It therefore requires additional tests, as for example those suggested by Imhof et al., in order to isolate a suspicious group of firms that may be colluding. This final step is essential if a competition agency wants to open an investigation, especially if one suspects bid-rigging cartels to be incomplete. Another possibility consists of changing the response variable and to move from a tender-based approach to a firm-based or a coalition-based approach, as discussed in the next section.

C. Flagging Collusive Firms

i. The Coalition-based Approach with Screens

The first approach discussed in this section consists of building collusive and competitive coalitions of firms, following the example of Imhof and Wallimann. Instead of calculating markers based on subgroups formed in tenders, we suggest focusing on a group of three firms and using all tenders in which the three firms submitted a bid. Figure 7 illustrates how we form a coalition. Seven firms (F1 to F7) participated in six tenders (T1 to T6). To construct the coalition F1F2F3, we retain all the tenders in which the three selected firms bid (in our case, T1, T2, T3 and T6). We then calculate the markers for each of the four tenders. In the last step, we use the summary statistics for the coalitions to create coalition-based statistics including the median, the mean, the minimum, and the maximum across coalitions for each tender.

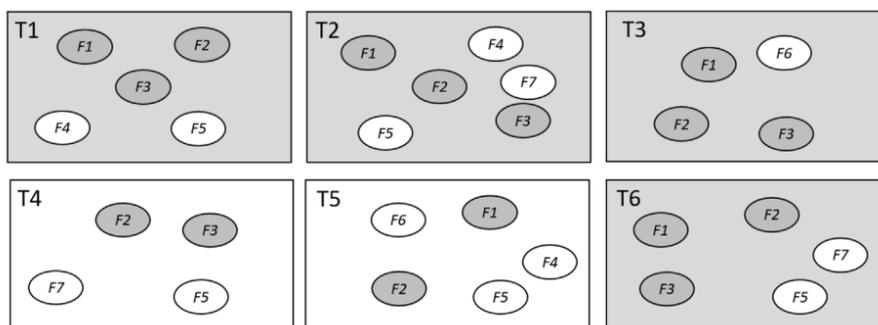


Figure VII: Example of Forming Coalitions³⁶

Like the robust screens, coalition-based screens are made with summary statistics based on collusive markers. Yet, we calculate the robust screens with the markers of all the possible subgroups in one tender, and the robust screens deliver information limited to one tender. In contrast, coalition-based screens deliver information across tenders for a restricted number of firms, the members of the coalition. By summarizing the statistical pattern of bids from the coalition across tenders, we can identify a systematic pattern in a coalition's conduct, which might indicate coordinated bidding across the tenders. In this approach, each observation is a coalition and the outcome Y takes the value of 0 if all coalition members are competing, and the value of 1 if all are colluding.

Imhof and Wallimann use data from Switzerland, Italy, and Japan to identify collusive and competitive coalitions. The best correct classification rates are approximately 90% for the Swiss and Italian coalitions and 94% for the Japanese coalitions. In other words, nine out of ten coalitions are correctly classified as collusive or competitive. The ensemble method appears to be the most accurate algorithm for classifying coalitions. The mechanism is a first-price sealed bid auction in both the Swiss and Japanese data, whereas it is a mean-price sealed bid auction for the Italian data. Furthermore, cartels differ across countries. In Japan and Switzerland, we deal mostly with complete cartels, while in Italy cartels are incomplete. Therefore, the coalition-based approach could have a large field of application since it works for different auction mechanisms and diverse cartels.

Finally, we also note that in all three cases, bid rigging affects the distributional pattern of bids by reducing the variance of bids in collusive coalitions. Table 3 presents the means of the coalitions' medians for the coefficient of variation by country. If bid rigging reduces the variance in the three countries, the magnitude of its effect slightly differs across countries. Competition increases the variance by a factor of three for both Japan and Italy, whereas only by a factor of two in Switzerland. However, we observe that the levels of the screen are noticeably different across countries.

³⁶ Imhof & Wallimann, *supra* note 29.

Country	Collusive Coalitions	Competitive Coalitions
Switzerland	3.38	6.80
Italy	10.13	30.73
Japan	1.06	3.19

Table III: Means of the Coalitions’ Medians for the Coefficient of Variation³⁷

ii. The Firm-based Approach with Image Recognition

A second approach implements convolutional neural networks (CNN), which is a deep learning technique used for image recognition.³⁸ The CNN approach has plots as inputs. To construct plots, we use the bid rotation test suggested by Imhof et al. to compute the interaction of one firm with other firms bidding in the same tenders. In order to compare plots across heterogeneous tenders, bids are normalized with the following min-max transformation:

$$\hat{b}_{it} = \frac{b_{it} - b_{min,t}}{b_{max,t} - b_{min,t}}$$

where $b_{max,t}$ and $b_{min,t}$ are the maximal and minimal bid in tender t , respectively. The normalized bid \hat{b}_{it} lies between 0 and 1; it equals 0 for the lowest bid in a tender and 1 for the highest bid. For each pair of firms, normalized bids lie in $[0,1] \times [0,1]$ space.

Figure 8 plots normalized bids for a hypothetical situation involving firms 1 and 2. Each dot in Figure 8 represents the bids submitted by those firms in a particular tender. A dot along the vertical (horizontal axis) is when firm 1 (2) submitted the lowest bid. Bid rigging is likely to produce more observations in the top right quadrant of Figure 8 and on the upper segment of the vertical and horizontal axes. Those areas are called “non-competitive” because non-lowest bids in the gray regions are quite distant from the lowest bid, which is consistent with cover bids ensuring that the designated cartel member wins the tender. Finding such a pattern could be indicative of possible collusion.

³⁷ *Id.*

³⁸ Martin Huber & David Imhof, *Deep learning for detecting bid rigging: Flagging cartel participants based on convolutional neural networks*, (Working Paper, 2021) https://www.researchgate.net/publication/351063082_Deep_learning_for_detecting_bid_rigging_Flagging_cartel_participants_based_on_convolutional_neural_networks.

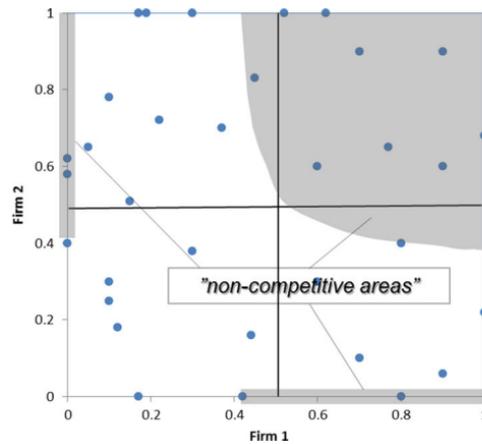


Figure VIII: Normalized Bid Patterns with Two Firms³⁹

Instead of using pairwise plots, Figure 9 plots a firm's normalized bid (firm 1) against the normalized bids of all other firms that submitted bids in the same tenders. The axes in Figure 9 are the same as in Figure 8 but were excluded since they are not relevant for analyzing the bidding pattern of firms. To read Figure 9, consider the vertical sequence of dots on the far left of the left panel. Those dots are to be found on the suppressed vertical axis and they indicate that firm 1's normalized bids are zero (so it submitted the lowest bid in a tender). The height of a dot in this vertical sequence measures the normalized bid of another bidder, and it depicts the distance to the lowest bid in a tender. The left panel of Figure 9 are bids from tenders when the cartel was active and the right panel is when there was competition. Given these images, the CNN identifies distinctive patterns between collusive and competitive plots. For example, consider the two regions highlighted with red rectangles in Figure 9. The absence of dots in the left panel reflects the collusive marker whereby there is a significant gap between the lowest bid and the next lowest bid.

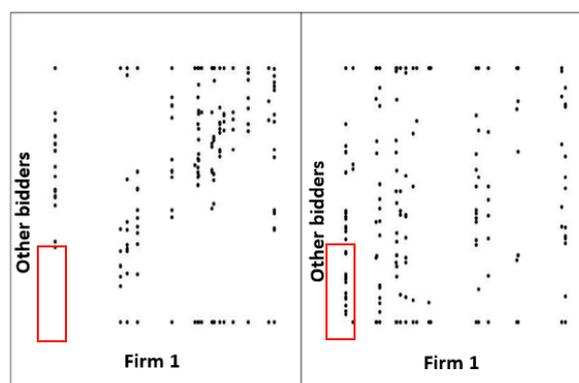


Figure IX: Normalized Bids Under Collusion (Left) and Competition (Right)⁴⁰

³⁹ Imhof, Karagök & Rutz, *supra* note 26.

⁴⁰ Huber & Imhof, *supra* note 38.

With this approach, the outcome Y is 1 if the plot is collusive and 0 if it is competitive. As with other machine learning approaches, we aim to build the best predictive models, which learn from systematic patterns in the data. Contrary to machine learning, the CNNs autonomously learn to recognize specific features, shapes, or colors for training predictive models. Huber and Imhof use 287 plots from Japan and 240 plots from Switzerland for which we know that they were collusive or competitive. CNNs were found to reach an average accuracy of 90% and 91% for predicting collusion in Japan and Switzerland, respectively. In other words, nine out of ten firms are correctly classified as colluding or competing.

D. Transposing Predictive Models for Cartel Screening

We have explained different approaches to detecting the presence of cartels at auctions using bid data. This approach requires data with identified collusive and competitive episodes in order to evaluate the predictive power of models, and we found that combining collusive markers with machine learning is a promising avenue for detecting cartels. The reader may now ask: Is it possible to use a predictive model trained on data in one market or in one country to screen for cartels present in other markets or countries? This is the challenge of transposing models.

When we screen new data to detect cartels, transposing predictive models could affect their effectiveness. The more different the context and the data in which predictive models were trained, the greater the concern. Predictive models with data drawn from the construction sector in Switzerland might well be effective if applied to an industry from the same sector in the same country. However, transposing predictive models might be more problematic if predictive models are trained on Swiss construction sector data and then applied to construction sectors in other countries. How effective is the screen developed in one country when it is applied to another country?

Huber et al. address this question with data from the Okinawa cartel in Japan. The Okinawa bid-rigging cartel is found to affect the bid distribution in the same way as with cartels in similar markets in Switzerland: bid rigging reduces the variance of bids and increases the gap between the first and the second lowest bids relative to the difference between adjacent losing bids. Nevertheless, transposing predictive models produced, in some cases, unsatisfactory results. When training on Japanese data to test Swiss data (or vice versa), classification rates were 57-62% for the random forest and 82-87% for the ensemble method. However, demeaning markers by country improves the results when transposing predictive models.⁴¹ Normalizing screens might then enhance efficacy when screening procurement data in different countries.

The proposed strategy is as follows. First, an algorithm is trained on data from one market or country where its efficacy is established. That will create confidence with various parties – such as competition authorities and judges – that it is a reliable tool for opening investigations and conducting dawn raids. Second, the algorithm is taken to a data set for another market or country in search of possible

⁴¹ Demeaning involves centering a marker by country so that the mean is zero for each country.

cartels. For each observation (which could be a tender, a coalition, or a firm), an estimated probability of collusion is delivered. The observation is labeled as “collusive” when the probability of collusion given by the model exceeds some threshold, such as 0.5. If the competition authority is more concerned with false positives (that is, incorrectly identifying collusion when firms are actually competing) then the threshold can be set higher than 0.5. Of course, a lower chance of starting a wrongful investigation comes along with a higher chance of failing to start an investigation when there actually is a cartel. The approach is flexible in that it allows the competition authority to control the relative likelihood of false positives and false negatives.

IV. Data Availability and Guidance for Developing Screens

In concluding, let us offer some general guidance for constructing a screen for detecting cartels. As one would expect with a data-driven approach, the recommendations are contingent on the available data.

The objective is to search for cartels in data set T (for target). Data set T encompasses prices over some time period or bids for some collection of tenders for a particular market. Our first situation is when there is another data set W for this same market (either prices from a different time period or bids from different tenders) which includes identified episodes of collusion and competition. If data set W has a sufficient number of observations, machine learning is used on data set W to develop a screen which is then applied to data set T. For each observation from data set T, the algorithm reports whether collusion is likely or not. This exercise was described in Sections III.A-III.C.

Now suppose data set W is either too sparse or simply absent but there is another data set Z with prices or bids for a comparable market and, when combined with W, has a sufficient number of observations. In assessing comparability, one will want to check for similar market traits and institutions. A comparable market could be the same product or service but in a different geographic area. For an auction, a comparable market is one that auctions off similar items or contracts using a similar auction format (e.g., sealed bid or oral auctions, public or private reserve price, and so on). It would also be helpful to learn the details of the collusive scheme present in data set Z and whether it is typical. Training data with collusive bids or prices for a rare or idiosyncratic collusive scheme is less likely to produce an effective screen for other markets. Machine learning is applied to data set Z (along with data set W if it is available) to produce a screen which is then used on data set T. An example was discussed in Section III.D where an algorithm was trained on Swiss construction procurement data (data set Z) to screen Japanese construction procurement data (data set T) for cartels.

Finally, we have the case when an algorithm cannot be trained using machine learning because of the absence of adequate data with identified collusive and competitive episodes. Screening would then involve looking for collusive markers, structural breaks, and anomalies, as described in Section II. Depending on the richness of data that is available, there are suitable tools to detect cartels.