

MACHINE-LEARNING BASED APPROACHES TO DETECT CARTELS IN RAILWAY-INFRASTRUCTURE PROCUREMENT

REPORT

Hannes Wallimann*
Philipp Wegelin*+
Silvio Sticher*

December 2022

Project funded by SBB research fund

Acknowledgment: Ramon Hauser, Lea Oberholzer, and Alisha Bühler provided excellent research assistance. Moreover, the authors would like to thank Karin Amberg, Marco Fetz, and Zeno van Uden for helpful comments. Finally, we are grateful to the SBB Research Fund for financial support. All the views expressed in this paper are solely those of the authors.

*University of Applied Sciences and Arts Lucerne, Competence Center for Mobility
+Philipp Wegelin was in the project team until October 1, 2022 (when he joined the Swiss Competition Commission)

Contents

Contents	iii
List of Figures	iv
List of Tables	v
1 Management Summary	1
2 Introduction	2
3 Literature review on screening methods	4
4 Methods	5
4.1 Prediction framework	5
4.2 Centralized approach: An add-on to the digitization of the railways’ infrastructure procurement	6
4.3 Decentralized approach: A category manager’s tool	8
5 Fitting the models	9
5.1 Centralized approach	10
5.2 Decentralized approach	11
6 Ex-ante application to the Swiss railway-infrastructure market	12
6.1 Background	12
6.2 Data and descriptive statistics	13
6.3 Traffic-light rating system and the detection of collusive subsets	15
7 Complementary analysis	19
7.1 Cost estimates of suspicious and non-suspicious tenders	19
7.2 Inclusion of subsidiaries	20
8 Discussion	21
9 Conclusion	22
A Formal introduction of the screens	23
B Decentralized approach: Example	23
C Training the algorithms	24
D Additional Figures and Tables	24

List of Figures

- 1 Centralized approach 6
- 2 Decentralized approach 8
- 3 Influence of the probability threshold when applying the random forest . . 11
- 4 Classification tree of the category manager’s tool 12
- 5 Tenders by canton 14
- C.1 Variable importance plot of the random forest 24
- D.1 Bid values in million Swiss francs 24

List of Tables

- 1 Performance measurements in the centralized approach 10
- 2 Performance measurements of the category manager’s tool 11
- 3 Tenders per year 13
- 4 Descriptive statistics 15
- 5 Absolute and relative amounts of suspicious tenders per algorithm 16
- 6 Screening within clusters (random forest with probability threshold 0.5) . . . 17
- 7 Firm interaction (random forest with probability threshold 0.5) 18
- 8 Differences between cost estimates and lowest bids of suspicious and non-suspicious tenders 19
- 9 Firm interaction, including subsidiaries (random forest with prob. threshold 0.5) 20

1 Management Summary

Public transportation in general and railways in particular exhibit a great extent of scale economies. This renders railway infrastructure a natural monopoly, usually owned and managed by the state or a state-owned infrastructure company. The substantial amount of subsidies raises the need for means to ensure that taxpayers' money is used efficiently. Typically, railway infrastructure construction and maintenance are procured using competitive procedures such as auctions. These auctions fulfill their purpose of allocating public resources efficiently—unless bidders collude.

In the study at hand, we aim to detect collusion by looking at the distribution of bids for a unique series of procurement tenders provided by the Swiss Federal Railways. What makes a distribution of bids suspicious? We let past data answer this question. By courtesy of the Swiss Competition Commission, we have access to structurally very similar procurement data relating to resemblant industries such as road construction and civil engineering. Unless in the case of the railway-infrastructure data, however, we know about the existence of convicted bid-rigging cartels in these instances. That is, we likewise observe bid distributions but can also pin down the tenders where a cartel was involved—and where not. Therefore, we can use this data to feed so-called supervised machine-learning algorithms, which in turn determine the likelihood of the presence of a potential cartel in any set of comparable data.

We employ machine learning in two different ways. In the first instance, which we call the centralized approach, we assume that the railway-infrastructure firm (the procuring company) is in possession of the entirety of the bid-distribution data in question. In the second instance, which we call the decentralized approach, tenders are separately classified by independent so-called category managers directly when receiving the firms' bids.

In both approaches, data requirements are minimal. All we need to observe are the bids submitted by the firms in a procurement auction. The distributions of these bids constitute a set of predictor variables or "screens". Examples are the spread, the difference between the two lowest bids, and many more. As it turns out (and as is expected), the most important of these screens is the coefficient of variation, which decreases in periods when a cartel is active. Using the past data of the Swiss Competition Commission, we apply supervised machine learning where the outcome of interest (the binary variable which determines the presence of a potential cartel) is also known. Hence, we can utilize a portion of the data to assess the performance of our models (by means of the correct-classification rate and the F1 validation score).

In the centralized approach, we make use of established machine-learning algorithms in the literature on algorithm-based cartel detection. Whilst fitting our models, we learn that the "super learner"—a weighted average of multiple machine-learning algorithms—performs the best. However, since the model's performance partially depends on the training and test sets composition, we also perform bootstrap prediction intervals (by randomly drawing a multitude of these sets). We find that the "random forest"—a comparatively simple algorithm derived from repeatedly applying classification trees to balance between bias and variance—is equally well suited in the present context.

In the decentralized approach, we take the everyday perspective of a category manager who typically has insight into only a limited number of bid distributions (and has a plethora of other tasks to manage). Therefore, a simple "category manager's tool" in the shape of

a stable decision tree is helpful. Along the decision tree, the category manager classifies tenders according to benchmarks for the respective screens. As in the centralized approach, we apply machine learning to determine the selection and optimal amount of screens (decision nodes). It happens that the optimal decision tree is not bushy at all: It simply suggests that whenever the coefficient of variation of the bid distribution is below 0.053, a tender should be treated as suspicious to maximize correct classification. Somewhat surprisingly, we find that our basic tool is virtually equally performant as the more opaque algorithms of the centralized approach.

We apply the centralized and decentralized approaches to classify 1,206 railway-infrastructure tenders in Switzerland between 2015 and 2021. All in all, about 10 percent of these tenders are rated "suspicious"; thereof, we classify only a quarter as "very suspicious" (i.e., exceeding a probability threshold of indicating an active cartel of 0.7). We then use additional information to examine subsets of our data in greater detail. For instance, we identify the region of Ticino to be more prone to suspicious bidding behavior—albeit at a moderate level. Regarding the procurement procedure ("open" versus "on invitation"), we do not find any striking differences. Finally, we study clusters of firms that repeatedly coincide in suspicious tenders. We identify two groups of firms whose interactions are indicative of a cartel in more than 20 percent of the instances and, therefore, worth further scrutiny. However, the thing to notice is that even this finding should be taken with a grain of salt, as it represents a small number of suspicious tenders.

In the course of some supplementary analysis, we study the relation between the cost estimates of the Swiss Federal Railways and some key statistics of suspicious and inconspicuous tenders. We do not find any clear-cut patterns. Moreover, we take a deep dive into the interaction of firms. Specifically, we include their subsidiaries into their parent companies, which increases the number of tenders per firm and, thus, the potential to detect collusive behavior. Whereas the interaction between firms remains relatively unsuspecting, we find some firms (including their subsidiaries) which are worthwhile looking at themselves.

In summary, we do not detect any major anomalies indicating bid rigging. However, we recommend that the Swiss Federal Railways pursue screening methods further in the future. For example, from the authors' point of view, our decentralized approach can simply be integrated into Swiss Federal Railway's supplier management. Also, repeating the study in a few years with new data and newly trained models makes sense. Furthermore, regarding the repetition of the study, we recommend that Swiss Federal Railways automatically prepare the data in advance, which can also be the basis for implementing the centralized approach. Finally, a policy recommendation is to communicate this study's results to deter potential cartels actively.

2 Introduction

According to recent OECD data, investment in railway infrastructure in 2019 by 23 of the 27 countries of the European Union amounted to EUR 38,3 bn.¹ France (EUR 10,8 bn) and Germany (EUR 9.6 bn) alone are responsible for around half of this amount ([International Transport Forum, 2022](#)). There are several reasons why investment in

¹See <https://data.oecd.org/transport/infrastructure-investment.htm> (accessed on September 19, 2022)

railway infrastructure (mainly construction but also maintenance) could increase in the near future. For example, the European Union has established a new action plan to boost long-distance and cross-border passenger rail to foster sustainable mobility ([European Commission, 2021](#)). Furthermore, the trend towards intermodal and multi-modal mobility requires corresponding facilities such as transport hubs and IT systems such as digital platforms ([Montero and Finger, 2020](#)).

Public transportation in general and railways in particular exhibit a great extent of scale economies. This renders railway infrastructure a natural monopoly ([Crozet, Nash, and Preston, 2012](#)). Usually, infrastructure is owned and managed by a state-owned infrastructure company or a separate infrastructure division of a state-owned transport company ([Finger, 2014](#)). The accordingly substantial public subsidies involved (see, e.g., [Finger, 2014](#), [Montero and Finger, 2020](#)) raise the need for means to ensure that the taxpayers' money is used efficiently. In the wake of liberalization of the public sector in general and of the railway sector in particular, many infrastructure companies were compelled to start procuring construction and maintenance using competitive procurement procedures such as auctions ([Laffont and Tirole, 1993](#)).

This is supposed to ensure efficient use of (partially) public money. However, as common to procurement auctions, concerns regarding bid-rigging appear ([Lengwiler and Wolfstetter, 2006](#)). Therefore, and because of the large amount of public money involved, infrastructure companies and their regulators have a vital interest in preventing unlawful cooperation of firms applying for projects. Detecting collusive arrangements is an indispensable first step in fighting cartels. In order not to rely heavily on external sources—such as principal witnesses in leniency programs—researchers have proposed statistical methods to screen markets to detect possible patterns of a cartel (see, e.g., [Porter and Zona, 1993](#), [Harrington, 2008](#)).

To the best of our knowledge, none of the existing studies focus on approaches designed for detecting potential cartel behavior in the railway-infrastructure industry. We aim to fill this gap by presenting two complementary approaches to detect potential illegal supplier cartels. Both of these approaches are related to cartel screening using machine learning (for a review see, e.g., [Harrington and Imhof, 2022](#)), where algorithms "learn" to predict the presence or absence of a cartel. First, we directly apply machine-learning based algorithms to a pooled database with many projects previously put to tender (centralized approach). Second, we showcase a tool created specifically for category managers of railway-infrastructure companies (decentralized approach). With the tool, a category manager can directly (live) monitor potential bid-rigging in his or her individual tenders by benchmarking descriptive statistics with critical values. Thereby, we use machine learning to select the relevant statistics and their critical values.

As an illustration for ex-ante cartel screening in the railway-infrastructure industry, we apply both approaches to the Swiss context by using unique procurement data provided by the Swiss Federal Railways, for which a cartel's incidence is unknown. The data set consists of 1,818 tenders, mainly from 2016 to 2021. We classify a (rather low) share of between 2.6% and 10.9% of all tenders as suspicious, depending on which approach and, more specifically, which decision threshold (i.e., probability threshold) of machine-learning algorithms we apply (0.5 and 0.7, respectively). We then present several further investigations of the "suspicious" tenders, an exercise of potentially direct practical use, as railway-infrastructure companies can use our investigation steps to explore their own

markets. Note that, before applying either approach, we develop models using past data from investigations by the Swiss Competition Commission (see also [Wallimann, Imhof, and Huber, 2022](#), [Imhof, 2019](#)). We find that we correctly classify a tender in 80% of the cases. Furthermore, the 95% prediction interval for correct classification (potential collusion or no collusion, respectively) reaches from 70% to 90%.

We proceed as follows. In Section 3, we present the relevant literature on screening methods. In Section 4, we discuss our cartel-detection approaches and their application to railway infrastructure in general. In Section 5, we adapt them to the Swiss railway-infrastructure market. In Section 6, we screen this market for cartel-like bidding patterns using unique procurement data provided by the Swiss Federal Railways. In Section 8, we discuss our results and argue that the external validity is likely to be given—i.e., the market structures in the cases from the Swiss Competition Commission and our application are comparable. In Section 7, we provide some supplementary analysis. In Section 9, we conclude.

3 Literature review on screening methods

Our work is related to studies discussing pro-active statistical methods to detect potential cartels, initially proposed by, for instance, [Harrington \(2008\)](#) and [Porter and Zona \(1993\)](#). These so-called screening methods constitute the first step of a multi-phase process that finally may condemn illegally cooperating firms. From an empirical perspective, screening methods are applied in the presence of "prediction policy problems" ([Kleinberg, Ludwig, Mullainathan, and Obermeyer, 2015](#)). That is, researchers aim to predict the probability of a cartel. As we use bidding patterns in procurement auctions to identify potential cartels, we focus on price screens in this review section. Moreover, our main attention is on behavioral screens, observing firms' behavior in markets. This is opposed to structural screens, which identify markets with typical cartel-conducive traits such as a small number of bidders ([Harrington, 2008](#)). Note that the German Railways use structural screens to identify markets prone to cartels (see, e.g., [Beth and Gannon, 2022](#)).

As we show in Section 5, we find that the coefficient of variation, defined as the standard deviation divided by the arithmetic mean of all bids submitted in an auction, is highly predictive in terms of determining the presence of a potential cartel. This fits the body of literature that uses screens to analyze the variance of bids: Several studies (e.g., [Abrantes-Metz, Froeb, Geweke, and Taylor, 2006](#), [Jimenez and Perdiguero, 2012](#), [Imhof, 2019](#)) point out that the coefficient of variation decreases during periods in which a cartel is active. Other studies find that cartel-involved behavior is often accompanied by asymmetrical bidding distributions. Recent findings from Japan ([Chassang, Kawai, Nakabayashi, and Ortner, 2022](#)), Canada ([Clark, Coviello, Gauthier, and Shneyerov, 2018](#)), and Switzerland ([Imhof, Karagök, and Rutz, 2018](#)) show that winning bids tend to be isolated; that is, the difference between the first- and second-lowest bids of a cartel is relatively high compared to differences among losing bids.

In the past five years, the literature in the context of cartel screens has largely focused on the use of machine learning, a subfield of artificial intelligence ([OECD, 2022](#)). It seems promising in light of the wide-scale use of digitization in procurement processes. The corresponding literature includes papers by [Huber and Imhof \(2019\)](#), and [Silveira, Vasconcelos, Resende, and Cajueiro \(2022\)](#) who apply machine learning algorithms (together

with price screens) to find possible patterns of cartels in Switzerland, Croatia, and Brazil, respectively. Moreover, [Rodríguez, Rodríguez-Montequín, Ballesteros-Pérez, Love, and Signor \(2022\)](#), and [Imhof and Wallimann \(2021\)](#) use machine learning algorithms in multiple countries. The latter paper proposes a method to identify conspicuous groups of firms directly instead of tenders affected by conspiracies. [Huber, Imhof, and Ishii \(2022\)](#) investigate the transnational transferability of machine-learning algorithms. They show that a country’s institutional context matters for a cartel’s influence on the distribution of bids. Therefore, applicants should use algorithms trained with data from other countries only with great reluctance. [Wallimann, Imhof, and Huber \(2022\)](#) detect incomplete cartels, where only some but not all bidders are members of the colluding agreement. Finally worth mentioning is the study of [Silveira, de Moraes, Fiuza, and Cajueiro \(2023\)](#) using unsupervised machine learning for cartel detection.

4 Methods

We recommend two complementary approaches for railway-infrastructure companies to detect potential cartels and, thus, avoid paying excessively high prices. First, we show how machine-learning techniques can be implemented into the wide-scale use of digitization in the railways’ infrastructure-procurement process (centralized approach). Thereby, screens are used as predictor variables to assess whether a cartel is likely to be present in a tender. Second, we propose a tool specifically designed for category managers to classify tenders according to statistical benchmarks (decentralized approach). Here, we define the relevant benchmarks with an intuitive machine-learning algorithm, the classification tree. Both of our approaches require minimal data: We only need the bids submitted by the firms in a procurement auction. This is valuable insofar as obtaining truthful information variables other than price (e.g., business-specific variables such as capacity utilization) is challenging without attracting the attention of a cartel.

Upon establishing the prediction framework we consistently use, we introduce both respective approaches in greater detail in the following.

4.1 Prediction framework

We use supervised machine learning with a set of predictor variables (X) to predict our outcome of interest (Y), the presence of a potential cartel. The outcome variable (also: target variable) is binary—taking on the value 0 if firms compete and the value 1 if they collude. To investigate the performances of machine-learning algorithms, we randomly divide the data set into two parts. The training set consists of 75% of the observations, and we use it to develop predictive models observing both predictor variables and outcomes. The test sample consists of the remaining 25% of observations and is used to evaluate model performance by predicting whether a potential cartel is active or not. Thereby, we can assess the goodness of fit of the models by comparing the predictions with the actual outcomes. To do so, we apply two measures: The correct-classification rate and the F1 score. The correct-classification rate denotes the proportion of correctly classified tenders in the test set. The F1 score is a harmonic mean of the precision and recall (also: sensitivity) in the test set. Precision calculates the proportion of how many cartel classifications are correct (true positives divided by all cases classified as ”suspicious”),

whereas recall measures the proportion of actual cartels that the algorithms correctly classify (true positives divided by all true cartel cases).

We exclusively use descriptive statistics of the bids in a tender as predictor variables. These screens describe bidders’ behavior in a tender and capture distributional changes due to a cartel. These exist because cartel members (explicitly or tacitly) coordinate bids in a tender and are aware of their competitors’ bids. For our application, we use nine frequently discussed screens (see, e.g., [Rodríguez, Rodríguez-Montequín, Ballesteros-Pérez, Love, and Signor, 2022](#), [Huber, Imhof, and Ishii, 2022](#), [Wallimann, Imhof, and Huber, 2022](#)): the coefficient of variation (CV), the spread (SPD), the difference between the two lowest bids (DIFFP), the relative distance (RD), the alternative relative distance (ALTRD), the normalized distance (NORMRD), the skewness (SKEW), and the Kolmogorov-Smirnov test (KSTEST). These screens are formally described in [Appendix A](#).

4.2 Centralized approach: An add-on to the digitization of the railways’ infrastructure procurement

Our first approach (depicted in [Figure 1](#)) relies on the instance of an increasingly digitized procurement, which allows to collect bids of firms in a single structured database (optimally in an automated way). Machine-learning algorithms then investigate the database with regard to potential cartels. They classify tenders as suspicious (the stack marked grey, including Tender p) or non-suspicious (the stack kept white, including Tender x). These evaluations are directly made available to the superior category manager of a railway-infrastructure company.




Figure 1: Centralized approach

As all tenders are pooled, the centralized (or cockpit) approach allows for both complete and disaggregated overviews of potentially suspicious clusters, i.e., broken down by regions, time periods, auction formats, and firms involved.

Note that until now, we simplified the classification task by predicting a binary target taking the values 0 or 1, i.e., whether firms collude or not. However, what we really do is carrying out so-called probability predictions. The intuitive default value for machine-learning algorithms is 0.5: If a tender is assigned a probability of 0.5 or higher, the algorithm classifies the outcome value to 1, i.e., "suspicious" in terms of potential cartel participation. Stated differently, this rule puts each tender to the class that is more

likely. Increasing the threshold, e.g., to 0.7, allows applicants (i.e., railway-infrastructure companies) to avoid false-positive results, i.e., wrongly classifying bidders as a cartel. On the other hand, however, it increases the occurrence of false-negative results, i.e., a cartel is wrongly discarded. In the literature, different thresholds are proposed. [Huber and Imhof \(2019\)](#), for instance, suggest a decision threshold between 0.5 and 0.7, while [Silveira, Vasconcelos, Resende, and Cajueiro \(2022\)](#) discuss a threshold between 0.6 and 0.75. We combine these suggestions by labeling tenders with a probability prediction between 0.5 and 0.7 as "suspicious" and those with an even higher value as "very suspicious".

Rail infrastructure companies could use a variety of algorithms to detect potential cartels using this approach, three of which we present in the following. The first is a logistic-regression model ([Berkson, 1944](#)). The second algorithm is the random forest ([Breiman, 2001](#)) which builds on classification trees. The third algorithm is the super learner ([van der Laan, Polley, and Hubbard, 2007](#)), which is an ensembled method averaging several other machine learning algorithms, i.e., random forest, lasso logit regression, gradient boosting, and neural networks.

First, we implement the (binary) logistic regression as baseline model in order to highlight the contribution to the literature by (newer) machine-learning algorithms. The logistic regression has a binary outcome variable. In our case, it takes on the values "cartel" and "competition", respectively. The algorithm converts the outcome variable through a logistic function to assess the probability of a cartel. We apply the *stats* package by the [R Core Team \(2013\)](#) to implement the logistic regression.

Our second algorithm, the random forest, is especially suitable to deal with the bias-variance trade-off: With a single classification tree, we can reduce the bias (i.e., systematically wrongly predicting the outcome variable) when we build a "bushy tree" and thus allow for more splits. By doing so, however, we increase the variance (i.e., how much the results vary when the algorithm is applied to different test sets) in the test set due to many leaves. To tackle potential variance in the test set, the random forest builds many trees (in our case, a thousand) by drawing multiple subsamples from the training set. Moreover, to reduce the correlation of trees across these subsamples and thus make the random forest more flexible, the algorithm considers at each splitting step only a subsample of predictor variables, in our case, the square root of the number of predictors. The final predictions consist of the average over all the classification trees. To implement the random forest, we apply the *randomForest* package for the statistical software R by [Liaw and Wiener \(2018\)](#). In addition to the advantages mentioned above, the random forest is relatively user-friendly because it does not require tuning specific penalty terms ([Athey and Imbens, 2019](#)).

Finally, we implement the super learner. This algorithm is a weighted average of multiple machine learning algorithms. We include the random forest ([Breiman, 2001](#)), lasso logit regression ([Frank and Friedman, 1993](#), [Tibshirani, 1996](#)), gradient boosting ([Friedman, 2001](#)), and neural networks ([McCulloch and Pitts, 1943](#)). The random forest is presented above. The lasso regression is a logistic regression that includes a penalty term, which restricts the number of predictor variables by setting the effect of predictor variables with low predictive power to zero. Similar to the random forest, gradient boosting grows a set of classification trees. In contrast, however, it builds trees in a forward stage-wise manner. The goal is that any additional tree sequentially learns from mistakes made in previous trees. Eventually, a neural network fits a system of nonlinear functions, thereby

allowing flexibility in the association between the outcome and the predictor variables. We apply the *SuperLearner* package by [van der Laan, Polley, and Hubbard \(2007\)](#) together with the *party*, *glmnet*, *xgboost*, and *nnet* packages by [Hothorn, Hornik, Strobl, and Zeileis \(2022\)](#), [Friedman, Hastie, Tibshirani, Narasimhan, Tay, Simon, Qian, and Yang \(2022\)](#), [Yuan \(2022\)](#), and [Ripley and Venables \(2022\)](#) in the statistical software R.

4.3 Decentralized approach: A category manager’s tool

In practical use, continuous (automated) data collection in a single structured database may not be feasible. Instead, monitoring of tenders is typically delegated to multiple category managers—so there is a need to equip them with a cartel-screening tool that works decentrally. Category managers’ tasks in railway-infrastructure companies are usually quite complex. For example, they must orchestrate construction during operation because the trains usually keep running when the building is in progress. Moreover, they are in constant exchange with planners and construction companies. Therefore, if a railway-infrastructure company expects a category manager to screen the market for potential illegal cartels in addition to daily business, a tool must be as simple as possible (that is, reducing effort cost and enhancing interest in the matter as much as possible).

Figure 2 depicts the decentralized application of the tool we propose: Every category manager is responsible for multiple tenders. He or she investigates each assigned tender immediately using the tool. For instance, Category Manager 1 finds tenders *a* and *b* exhibiting suspicious bid patterns (marked grey). He then hands them over to the superior category manager for further investigations. In contrast, the tool does not indicate bid rigging for the tenders of Category Manager 3 and, therefore, no further action is needed.

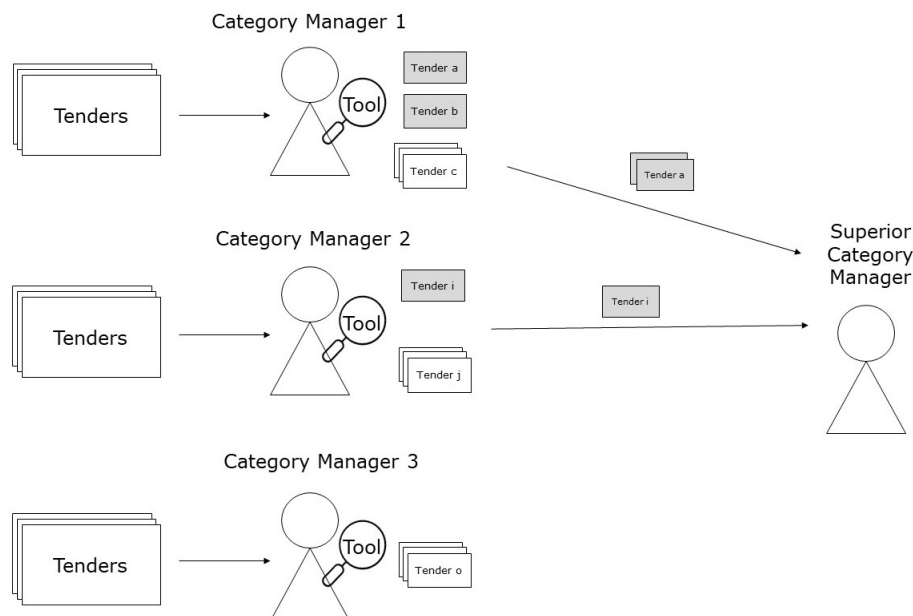


Figure 2: Decentralized approach

We propose a threshold-based approach to screen railway-infrastructure markets. Similar to Section 4.2 and in contrast to theoretically derived benchmarks (see, e.g., [Imhof,](#)

2019), we apply a machine-learning algorithm to define the thresholds. More precisely, we suggest the use of a classification (or decision) tree (Breiman, Friedman, Olshen, and Stone, 1984), a method that is easy to understand (or, at least, easy to utilize) also by non-data-scientists. In our application, the classification tree uses screens to predict whether a tender is suspicious and, thus, a cartel might be active. The classification of a tender starts at a top node with an associated threshold of the screen. Based on the screen’s value of this tender, a category manager can move down the corresponding branch. The agent repeats this procedure until she arrives at a so-called leaf or terminal node.² The terminal node indicates whether the tender is suspicious or not. More technically, we can think of the terminal nodes as non-overlapping regions, i.e., a set of suspicious or non-suspicious tenders. According to a goodness-of-fit criterion, e.g., the Gini coefficient for binary outcomes, the algorithm selects the predictors to maximize the homogeneity of the prediction at each split.

To prune the tree and to determine the optimal level of tree-complexity, we perform 10-fold cross-validation. For the implementation of the classification tree, we apply the *caret* package for the statistical software R by Kuhn (2019).

Finally, note that—due to the discrete structure of decision trees—probability thresholds should not be freely chosen in the decentralized approach, which is why we only work with a value of 0.5.

5 Fitting the models

To train our machine-learning models, we draw on data from three convicted bid-rigging cartels provided by the Swiss Competition Commission. These cartels are See-Gaster, Grisons, and Ticino, previously introduced by Imhof, Karagök, and Rutz (2018) and Wallimann, Imhof, and Huber (2022).³ The See-Gaster and Grisons cartels were active from 2004 to 2010, while the Ticino cartel was active from 1999 to March 2005. These cartels rigged public and private contracts in construction markets comparable to the railway infrastructure, i.e., road construction, asphaltting, and civil engineering. Wallimann, Imhof, and Huber (2022) classify the procurement process in Switzerland as first-price sealed-bid auctions, where all bidders simultaneously submit bids. The bids are “sealed” because bidders do not know their competitors’ bids. The price is still the most important—albeit not the single—criterion in determining the winner.⁴ Pooling the data of all three cartels leads to 538 competitive (no-cartel) and 519 non-competitive (cartel) tenders. To prevent class imbalances, i.e., preventing the case that one class occurs distinctively more often than the other (a situation where classification methods might not be good at (Békés

²In Appendix B, we show how to apply the decentralized approach.

³See also decisions *Bauleistung See-Gaster: Verfügung vom 8. Juli 2016*, and *Bauleistung Graubünden Strassenbau u.a.: Verfügung vom 19.8.2019* available on the homepage <https://www.weko.admin.ch/weko/de/home/praxis/publizierte-entscheide.html> (accessed on November 29, 2022). For the Ticino cartel see also on this homepage: <https://www.weko.admin.ch/weko/de/home/medien/medieninformationen/nsb-news.msg-id-16109.html> (accessed on November 29, 2022).

⁴For instance, even before the overhaul of the Swiss procurement law in 2021, minimum requirements and eligibility criteria were also considered. For more on this, see <https://www.bkb.admin.ch/bkb/de/home/themen/revision-des-beschaffungsrechts.html> (accessed on February 21, 2023).

and Kézdi, 2021)), we only consider competitive tenders where past (or future) cartel participants submitted.

5.1 Centralized approach

The screens we use in the centralized approach correspond to the predictor variables described in Appendix A as well as their interactions and squares. Table 1 presents the prediction performance in the test set for the three different algorithms (logistic regression, random forest, and super learner). We achieve (decent) correct classification rates from 77.7% to 81.3%. In other words, we correctly classify about four out of five tenders. In this regard, the super learner outperforms the two other algorithms by 2.8 percentage points and more.

Table 1: Performance measurements in the centralized approach

Algorithm	CCR	F1 score	95% PI CCR	95% PI F1 score
Logistic regression	0.777	0.789	[0.717; 0.842]	[0.711; 0.845]
Random forest	0.785	0.803	[0.777; 0.857]	[0.773; 0.860]
Super learner	0.813	0.796	[0.774; 0.853]	[0.766; 0.889]

Note: *PI* and *CCR* denote prediction interval and correct-classification rate, respectively.

Since these performance measures might also depend on the test set’s composition, we additionally calculate bootstrap prediction intervals by randomly drawing 2,000 training and test sets without replacement. In every training sample, we construct a new prediction algorithm and estimate the performance in the test set. The resulting 95%-prediction intervals are relatively wide, ranging from 7.9 percentage points (correct-classification rate of super learner) to 13.4 percentage points (F1 score for logistic regression). Therefore, and as a word of caution, the model’s performance partially depends on the composition of the training and test sets. However, regarding the random forest and the super-learner algorithms, the lower bound of the interval is still higher than three-quarters correct prediction.

Concerning 95%-prediction intervals, we see that the random forest performs almost identically well as the super learner. Therefore, we conclude that both of these algorithms are advisable in the present context. The lower bound of the logistic regression, in turn, is substantially lower (71.7% for the correct-classification rate and 71.1% for the F1 score).

An additional feature of tree-based algorithms is that we can depict the meaningful features. Figure C.1 in Appendix C shows the relative importance of the predictor variables according to the 2,000 bootstrap samples. We notice that the coefficient of variation is the most important predictor variable. This comes at no surprise, as also other studies consider it as among the most important predictor variables for classifying potential cartels (Imhof and Wallimann, 2021).

Finally, applying the random forest, Figure 3 shows how different decision thresholds affect the performance in the test set: By iteratively increasing the thresholds, the overall correct-classification rate is roughly unaffected up to a threshold of around 0.7, while the

false positive predictions substantially decrease (from around 19% to 8%). We conclude that even with the higher threshold, we are still able to identify potential cartels. This finding is in line with the thresholds discussed in the literature.




Figure 3: Influence of the probability threshold when applying the random forest

5.2 Decentralized approach

The category manager’s tool of the decentralized approach consists of a classification tree, which (theoretically, i.e., prior to pruning) utilizes all predictor variables described in Appendix A.

Table 2 presents the accuracies of the classification tree in the test set.

Table 2: Performance measurements of the category manager’s tool

Correct-classification rate	0.800
F1 score	0.806
95% prediction interval correct-classification rate	[0.770; 0.853]
95% prediction interval F1 score	[0.762; 0.855]

The correct-classification rate amounts to 80%, the F1 score is about the same. Analogous to Section 5.1, we calculate bootstrap prediction intervals, constructing a classification tree for each of 2,000 training samples and estimating the performance in the respective

test sets. The resulting 95%-prediction interval ranges from 77.0% to 85.3% for the correct-classification rate. This describes the category manager’s tool as similarly predictive as the more opaque algorithms of the centralized approach.

Figure 4 presents the final classification tree for the direct application by category managers. We can easily observe that our classification tree is far from being bushy and has only two terminal leaves. Moreover, only the coefficient of variation (CV) is relevant for classifying a tender as suspicious or not. A category manager can classify a tender as suspicious whenever the CV is lower than 0.053.




Figure 4: Classification tree of the category manager’s tool

6 Ex-ante application to the Swiss railway-infrastructure market

In this section, we apply both approaches to real-world data from the railway-infrastructure sector. In particular, we draw on data provided by the Swiss Federal Railways. From a data-analytics perspective, we now use the uncovered patterns from the data of the Swiss Competition Commission and use it to predict the target variable (i.e., a potential cartel’s presence or not) in the data of the Swiss Federal Railways. First, we subsequently discuss the institutional background. Then, we present summary statistics of the data set. Finally, we screen the data for collusive behavior using both the centralized and decentralized approaches. We do so by classifying all tenders as “suspicious” and “non-suspicious” in terms of a traffic-light system. The group of suspicious tenders is then analyzed in more detail.

6.1 Background

In 2020, the Swiss railway network amounted to 5,317 km.⁵ Several railway-infrastructure companies are responsible for building, maintaining, and operating infrastructure. These companies are vertically integrated into state-owned public transport companies. The Swiss Rail Infrastructure Fund (BIF), endowed by general-tax revenue as well as fuel and road taxes, funds expenditures to build, maintain, and operate the railway infrastructure. Since the implementation of the BIF in 2016, an annual amount of 4 to 5 billion Swiss francs has been invested in railway infrastructure, roughly 20% in building, 66% in maintaining, and

⁵See <https://www.bfs.admin.ch/bfs/de/home/statistiken/mobilitaet-verkehr/verkehrsinfrastruktur-fahrzeuge/streckenlaenge.html> (accessed on October 12, 2022)

14% in operating railway infrastructure (LITRA, 2022). The federal transport agency and the railway-infrastructure companies conclude four-year agreements about goals and the distribution of funds to specific tasks to achieve these goals. Based on these agreements and the available funds, the infrastructure companies contract specialized private construction companies for construction projects. Usually, the law requires competitive procedures for contracting out.

6.2 Data and descriptive statistics

We use a unique data set provided by the Swiss Federal Railways, containing projects of track construction, station construction, and civil engineering projects. The data set contains tenders ranging mainly from 2015 to 2021. Overall, we observe 1,818 tenders. Railway-infrastructure companies in Switzerland can generally choose between two tendering procedures: open procedures and procedures on invitation. Using the latter, a railway-infrastructure company invites a predefined (usually at least three) number of firms. The open procedure does not restrict the participation of firms submitting bids and is thus associated with more severe competition (Wallimann, Imhof, and Huber, 2022).

Note that in order to apply the screens presented above, we must perform some final data wrangling before starting the actual screening. First, since a minimum of three bids is required to calculate the entirety of screens described in Appendix A, we drop tenders with less than three bids. Second, from time to time, firms submit different variants. In such cases, we only consider the lowest bid. Our final sample contains 1,206 tenders altogether.

Table 3 presents the number of tenders (with more than two bids) per year grouped by the two procedures. For example, there were 267 tenders in 2016, of which 68 belonged to the open procedure and 192 to the procedure of invitation, respectively.

Table 3: Tenders per year

	All	Open procedure	Procedure on invitation
older than 2015	5	0	5
2015	47	21	26
2016	267	68	199
2017	155	46	109
2018	162	31	131
2019	185	49	136
2020	224	56	168
2021	158	38	120
2022	1	1	0
Unknown	2	0	2
Total	1,206	310	896

Switzerland consists of 26 member states, the so-called cantons. An interesting feature of our data set is that we obtain data from all over Switzerland. Figure 5 shows the distribution of projects among cantons. We observe, e.g., that the most populous⁶ canton (Zurich) has the most tenders (261), whereby 216 (83%) and 45 (17%) are procedures on invitation and open procedures, respectively. The cantons with the second and third highest number of tenders are Vaud and Ticino with 132 and 125, respectively.

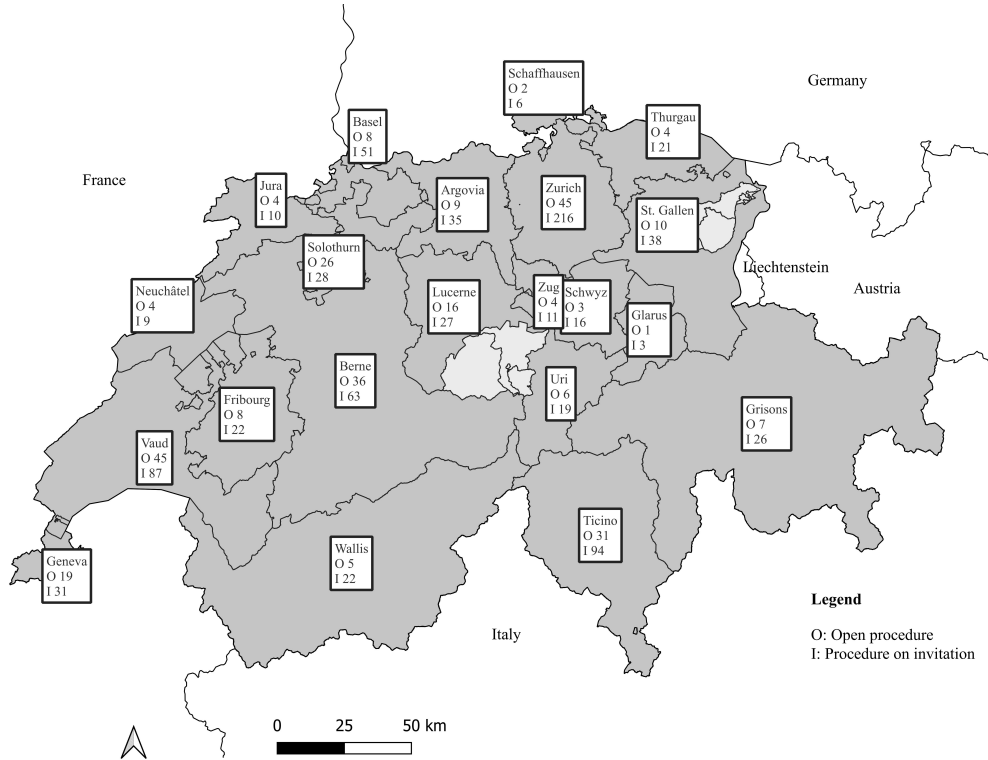


Figure 5: Tenders by canton

Our analysis exclusively rests upon the bids submitted by the firms. Table 4 presents the key descriptive figures of these bids (grouped by tender procedure). We observe, e.g., that the highest bid amounts to approximately 376 million Swiss francs. In comparison, the lowest bid amounts to 2,175 Swiss francs. The mean and median values amount to 2.7 and 0.7 million Swiss francs, respectively, which indicates a heavily right-skewed distribution. We illustrate this distribution also in Figure D.1 in Appendix D. The average bid for projects tendered with the procedure on invitation is lower (0.8 million Swiss francs) than with the open procedures (7.0 million Swiss francs).

Table 4 also shows that the total number of firms in the data set amounts to 2,061. Of these firms, 1,559 and 779 took part at least once in a procedure on invitation and an open procedure, respectively.

⁶With 476km, Zurich is also home to the most expansive railway network in Switzerland. See <https://reporting.sbb.ch/infrastrukturen> (accessed on February 21, 2023).

Table 4: Descriptive statistics

	All	Open procedure	Procedure on invitation
Number of bids			
Mean	3.97	4.68	3.73
Median	4	4	3
Bid values (Swiss francs)*			
Mean	2,688,906	7,021,877	804,103
Median	656,117	3,028,289	408,199
Maximum	375,591,800	375,591,800	36,954,503
Minimum	2,175	11,522	2,175
Standard deviation	11,180,096	19,429,012	1,876,232
Number of bidding firms			
Total	2,061	774	1,559

**Note: We use all bids to calculate the descriptive statistics on the bid values. Therefore, tenders with many bidders have a higher weight.*

In the following, we illustrate the application of the two approaches discussed in this paper using the data of the Swiss Federal Railways. We present the results in two steps: First, we classify all tenders in terms of a traffic-light rating system as red (i.e., suspicious) or green (i.e., non-suspicious). Second, we screen for potential collusion in subsets of the data (regarding geography, auction format, and firm interactions).

6.3 Traffic-light rating system and the detection of collusive subsets

Table 5 presents the results of the two approaches' application to the database of the Swiss Federal railways. Applying the random forest and the super learner of the centralized approach with the decision threshold of 0.5, we classify 8.5% and 5.6% of all tenders as suspicious, respectively. When we look for "very suspicious" tenders by increasing the decision threshold to 0.7, the random forest and the super learner flag only 2.6% and 2.4%, respectively, of all tenders as candidates for further investigations.

Applying the category manager's tool of the decentralized approach, we see that 130 (10.8%) of all tenders are classified as suspicious, whereas the vast majority of tenders (89.2%) again is classified as non-suspicious.

Table 5: Absolute and relative amounts of suspicious tenders per algorithm

Approach	Algorithm	Probability threshold	Suspicious	Non-suspicious
Centralized	Random forest	0.5	102 (8.5%)	1,104 (91.5%)
	Super learner	0.5	68 (5.6%)	1,138 (94.4%)
	Random forest	0.7	31 (2.6%)	1,175 (97.4%)
	Super learner	0.7	29 (2.4%)	1,177 (97.6%)
Decentralized	Classification tree		130 (10.8%)	1,076 (89.2%)

In summary, the results are not indicative of active potential cartels. We conclude that from an overall perspective, procurement of the Swiss Federal Railways is efficient. As expected (regarding the error rates), both approaches classify a subset of several tenders as collusive. This could be due to false-positive results. Nevertheless, if there actually is collusion, it is most likely within these subsets—rendering them still very worthwhile investigating.

Using the 102 tenders classified as suspicious by the random forest, we investigate whether noticeable patterns occur in terms of regions, the procurement process applied, or firm interactions. We use the results of the random forest due to two reasons: First, the centralized approach includes such further analysis. Second, the use of the random forest with the decision threshold of 0.5 has practical reasons, as it identifies more tenders as suspicious than the super learner. Therefore, the chances increase that we find a cluster exhibiting potential cartels. Moreover, 63 of the 68 tenders marked as suspicious by the super learner also classified as suspicious by the random forest. Therefore, we conclude that the findings of the two algorithms are congruent.

We start by geographically allocating the suspicious tenders to cantons. In Table 6, we list all cantons with more than 50 tenders. The strongest penetration of suspicious tenders amounts to 14.4% (present anonymization at the service of the Swiss Federal Railways), which is 5.9 percentage points higher than the nation-wide average. Regarding the two different procedures, we do not see any difference, as for both procedures (open vs. on invitation), the share of suspicious tenders is between 8 and 9 percent (see also Table 6). Finally, differentiating between years, we see that the amount of suspicious tenders is the highest in 2017. However, with 4.4 percentage points, the difference to the overall fraction of suspicious tenders is rather small.

Table 6: Screening within clusters (random forest with probability threshold 0.5)

	Suspicious	Non-suspicious
All	102 (8.5%)	1,104 (91.5%)
By canton		
Canton <i>A</i>	14.4%	85.6%
Canton <i>B</i>	13.1%	86.9%
Canton <i>C</i>	9.3%	90.7%
Canton <i>D</i>	8.5%	91.5%
Canton <i>E</i>	6.8%	93.2%
Canton <i>F</i>	4.6%	95.4%
Canton <i>G</i>	4.0%	96.0%
By procedure		
Open procedure	27 (8.7%)	283 (91.3%)
On invitation	75 (8.4%)	821 (91.6%)
By year		
2015	5 (10.6%)	42 (89.4%)
2016	15 (5.6%)	252 (94.4%)
2017	20 (12.9%)	135 (87.1%)
2018	12 (7.4%)	150 (92.6%)
2019	14 (7.6%)	171 (92.4%)
2020	21 (9.4%)	203 (90.6%)
2021	15 (9.5%)	143 (90.5%)

To investigate the interaction of firms, we present a matrix quantifying how many times firms participated in suspicious tenders together with other firms in Table 7. Such a matrix is helpful as we assume that a bid-rigging cartel involves regular interaction between firms (Imhof, Karagök, and Rutz, 2018). Therefore, we only include firms with a minimum of three suspicious tenders and at least one interaction with another firm partaking at three or more suspicious tenders. The diagonal of the matrix contains the percentage of suspicious tenders of each individual firm, where the values in brackets refer to suspicious (bolded) and overall tenders the firm in question took part. For instance, regarding firm 20, 22% of the tenders (that is, 7 out of 32 tenders) are classified as suspicious.

For the remaining 12 firms, we analyze all $2^{12} = 4096$ possible clusters (including trivial clusters with no or single firms). For each of these, we compute a "suspicioucy rate",

Table 7: Firm interaction (random forest with probability threshold 0.5)

	20	227	163	23	179	2411	13	70	89	123	180	273
20	22% (7/32)	20% (2/10)		0% (0/1)		25% (1/4)			75% (3/4)	33% (3/9)		
227		23% (7/30)		0% (0/1)		25% (1/4)		67% (2/3)	0% (0/2)	20% (1/5)		
163			7% (5/72)		4% (1/26)		0% (0/3)				0% (0/1)	
23				5% (4/86)			7% (1/15)		0% (0/1)			10% (1/10)
179					7% (4/59)		0% (0/1)					
2411						31% (4/13)		100% (1/1)		0% (0/3)		
13							5% (3/63)					13% (1/8)
70								50% (3/6)				
89									33% (3/9)	100% (1/1)		
123										18% (3/17)		
180											17% (3/18)	
273												9% (3/32)

which we define as the ratio of suspicious to overall tenders involving at least one firm of the cluster. The ranking of firms by suspicious rate depends on whether we include the diagonal elements of Table 7 into our analysis. If we do, the two most "suspicious" clusters (with more than one firm) are {70, 89} and {70, 2411}, even though there is only one interaction between firms 70 and 2411 and none between firms 70 and 89. Even by increasing the cluster size, the suspicious rate is barely driven by the interaction of firms. For example, the most suspicious cluster of size four, {20, 70, 89, 2411}, exhibits a rate of 32%. However, only five out of 22 suspicious tenders relate to interactions.

Once we exclude the diagonal elements, the suspicious rate reads as the ratio of suspicious to overall tenders, each involving at least *two* of the firms of a cluster. Still, the top-ranked clusters exhibit little interaction.⁷ However, further down the list, more interesting cases arise: Cluster {20, 89} consists of four interactions, three of them suspicious. Likewise, two out of three interactions are suspicious in cluster {70, 227}. Fifth in line, cluster {20, 70, 89, 2411} (which we know from above) registers five suspicious interactions out of nine. The thing to notice is that firm 89 never interacts with firms 70

⁷The clusters {70, 2411} and {89, 123} only interact once, reaching a 100% rate, because the single interaction is labeled suspicious.

and 2411, and firms 20 and 70 do not interact. Close next is cluster {20, 89, 123} with a suspicious rate of 50% (seven out of 14).

Since cartels (if stable) often and regularly interact, it would be worthwhile to put the companies of these clusters under increased scrutiny in future tenders, even though, up to today, the results are far from striking.

In summary, we classified between 89.2% and 97.6% of all tenders as non-suspicious. These results do not indicate active large-scale cartels. Furthermore, using the results of the centralized approach with the random forest algorithm and a decision threshold of 0.5, we found no striking bid-rigging patterns regarding cantons, procurement procedures, and years. Nevertheless, looking at firm interactions of firms bidding more frequently in suspicious tenders, we identify potential clusters of firms worth increased scrutiny. However, it should be noted that these results may be due to classification errors and should be treated with caution as they are based on small numbers.

7 Complementary analysis

7.1 Cost estimates of suspicious and non-suspicious tenders

For each tender, a cost estimate is calculated by the Swiss Federal Railways in advance. The difference between the cost estimate and the actual cost is a key variable from a business perspective of the Swiss Federal Railways: A positive number (i.e., actual costs that are lower than the estimate) indicates that competition is playing out. In Table 8, we present the difference between the firms' bids and the cost estimates. Note that we take the lowest bid in each tender for the firms' bids.

Table 8: Differences between cost estimates and lowest bids of suspicious and non-suspicious tenders

Metric	All	Suspicious tenders	Non-suspicious tenders
Mean	547,866	816,561	522,652
Median	89,194	52,311	93,980

Note: 'Suspicious tenders' are tenders classified as suspicious by the random forest with the probability threshold 0.5.

Regarding Table 8, we observe the following. First, all the numbers are positive. This means the estimated cost is consistently higher than the firms' bids. Second, the medians are lower than the means, probably driven by some high-volume tenders (see also Table 4). Finally, the mean is higher for the suspicious bids, whereas their median is lower. We conclude that there are no (clearly visible) differences between the suspicious and non-suspicious tenders in this analysis of the cost estimates.

7.2 Inclusion of subsidiaries

Up to here, we considered each firm’s subsidiary individually (see, e.g., Table 7). This is because they may prepare offers (and bids) independently. In the following, we include the firms’ subsidiaries to the number of tenders. For example, assume that firm X has two subsidiaries, firm Y and firm Z. We now consider these three firms (i.e., X, Y, and Z) as one firm, i.e., firm X.

To investigate the interaction of firms, we again present a matrix quantifying how many times firms participated in suspicious tenders together with other firms in Table 9. But as the inclusion of subsidiaries increases the number of (suspicious) tenders, we only consider firms with a minimum of five suspicious tenders.

Table 9: Firm interaction, including subsidiaries (random forest with prob. threshold 0.5)

	107	1241	20	2022	227	2380	273	36	435
107	12% (5/42)	10% (1/10)						0% (0/1)	
1241		14% (5/37)		33% (2/6)					
20			22% (7/32)		20% (2/10)	25% (1/4)			
2022				17% (6/35)				0% (0/1)	
227					23% (7/30)	25% (1/4)			
2380						8% (14/169)	0% (0/4)	4% (2/45)	3% (1/29)
273							7% (5/67)		
36								7% (8/107)	4% (1/25)
435									6% (5/84)

As in Section 6.3, we compute the suspiciousity rate for each possible cluster. (Here, there are $2^9 = 512$ clusters to analyze.) Including diagonal elements, the most suspicious cluster, $\{20, 227\}$, contains of 72 tenders, 16 of which (i.e., 22%) are classified as suspicious. The second and third suspicious clusters, $\{20, 227, 2022\}$ and $\{227, 2022\}$ also reach suspiciousity rates of 21% and 20%, respectively, but there is no interaction between firm 2022 and the other firms. In general, as above, suspiciousity is again driven by patterns within single firms (including their subsidiaries). Specifically, most noticeable clusters include firm 227 ($7/30 = 23\%$ suspicious tenders), firm 20 ($7/32 = 22\%$), or firm 2022 ($6/35 = 17\%$). Once we exclude the diagonal elements, the most suspicious clusters do not look particularly striking: In $\{1241, 2022\}$, 2 out of 6 (33%) interactions happen in the

context of suspicious tenders. Next on the list is {1241, 2022, 227, 2380} ($3/10 = 30\%$) and the very similar group of firms {1241, 2022, 20, 2380} (also $3/10 = 30\%$.)

8 Discussion

Identifying a potential cartel using statistical methods is a thorny task. A suspicious bid pattern in one tender is no definitive evidence for collusion. Nevertheless, it indicates where to look at more closely, or—if no indication of a cartel is found—increases a railway-infrastructure company’s confidence that procurement is competitive. The means we propose for railway-infrastructure companies to deploy exhibit several advantages. First, data requirements are low. This is important as it comes with low effort (for the category manager in the decentralized approach) and allows for cartel screening in large data sets (using the centralized approach). Second, our approaches are relatively simple to implement from a technical point of view. Third, they complement each other. While the centralized approach allows for a systematic overview and various in-depth evaluation options, the category manager’s tool of the decentralized approach is particularly well-suited for everyday use. Finally, appropriate communication to the industry helps to deter future collusion in the market.

Besides these advantages, one should be aware of some limitations. First, the category manager’s tool could be seen as too simple since it contains only one benchmark. However, the results presented in section 5 demonstrate that it also achieves decent predictions that are almost as accurate as with the centralized approach. Second, being aware of the use of screening tools, firms can try to beat the algorithms. For instance, they could influence the bids in a way that the algorithms classify tenders as “non-suspicious”. However, such efforts can be expensive, especially when it comes to the centralized approach, which is a black box for firms. Moreover, in-depth cooperation requires more communication, potentially increasing evidence that a competition agency can use to condemn firms (see also [Imhof and Wallimann, 2021](#)). Finally, screening tools are just one of several cartel detection methods such as leniency programs. However, increasing the effort to beat algorithms leads to more evidence, making it easier to prosecute potential cartels (see, e.g., [Blatter, Emons, and Sticher, 2018](#)).

Future studies should investigate how other screening methods could complement the approaches presented in this study. An example is the application of econometric models to flag potential cartels (see, e.g., [Bajari and Ye, 2003](#)), as it also takes into account other factors, such as the workload of companies, in order to find potential cartels. One downside of econometric models, however, is their high data requirement, as variables are required that often are intern to firms (such as costs and revenue). Another promising addition could be the application of coalition-based screens ([Imhof and Wallimann, 2021](#)), which screen for groups of suspicious firms instead of suspicious tenders. The main advantage of coalition-based screens is that it directly flags suspicious firms. However, this is not possible by simple means such as the category manager’s tool, which approaches tender by tender. Classifying suspicious groups of firms, in contrast, requires an overview of many tenders in which the same firms submit a bid. Thereby, the coalition-based approach could be promising in identifying collusive subsets when firms are observed frequently interacting in suspicious tenders.

Moreover, similar to the subsample screening we presented above (regarding geography,

different tendering procedures, and firm interactions), practitioners could apply further in-depth analysis—for instance, regarding different types of projects, as some markets could be more susceptible to cartels than others.

Finally, the external validity of our two approaches is of great practical relevance. External validity means the extent to which the correct classifications are similar in new data sets. For instance, the study of [Huber, Imhof, and Ishii \(2022\)](#) investigates the transnational transferability of screening methods and finds that classification rates can go down when training algorithms in one country and testing performance in another. In our study, external validity refers to whether the patterns of association between the target variable (i.e., potential cartel or competition) and the screens found in the road construction, asphaltting, and civil-engineering data apply in railway-infrastructure tenders. We argue that this is the case since the corresponding markets seem comparable due to several reasons: First, the procurement procedures are the same (e.g., [Wallimann, Imhof, and Huber, 2022](#)).⁸ Second, in both markets, the companies need heavy machinery and are therefore limited in their radius of action. Third, in both markets, companies depend on comparable suppliers of, e.g., raw materials. Therefore, natural geographical barriers to entry arise from the latter two points. Fourth, we used training data from the same country, i.e., both data sets are from Switzerland. Fifth, we used a broader set of training data (i.e., three cases) to construct prediction models. Finally, an indication is that several engineering and construction firms are also active in the railway-infrastructure market. However, more research on the existence of general cartel patterns across diverse markets is needed. Furthermore, data of upcoming cases has to be used to develop new models able to detect future conspiracies.

9 Conclusion

This study introduced two complementary approaches allow screening for potential cartels by railway-infrastructure companies. Both approaches rely on machine-learning algorithms, which assess whether tenders show suspicious or non-suspicious patterns in terms of potential collusion among bidders. We draw on previous investigations of the Swiss Competition Commission to uncover such patterns. We used these patterns to predict potential cartel behavior in railway-infrastructure procurement by means of a traffic-light system.

First, we applied several state-of-the-art machine-learning algorithms to flag collusive agreements of firms. The centralized approach of doing so allows a systematic overview of a railway-infrastructure company’s tenders and various in-depth analyses.

Second, we presented a tool tailored to individual category managers responsible for projects put out to tender. When a category manager detects bids exhibiting suspicious patterns, she contacts the superior category manager for further investigations.

Be it by actually disclosing cartels or by deterring their formation, the endeavor of screening tenders is of high practical relevance in light of the vast amounts of (public) money devoted to railway infrastructure. Finally, the thing to notice is that the two approaches can also serve as examples for uncovering potential cartels in other industries.

⁸Note, however, that the Swiss Federal Railways chooses the procedure on invitation relatively more often than it is applied in the data set of the Swiss Competition Commission. Therefore, as that influences the number of bidders, we did not consider this variable as a predictor.

Appendix

A Formal introduction of the screens

To formally introduce the screens note that t denotes some tender t ; \bar{b}_t is the mean of all bids in tender t ; sd_t and $d_{losingbids,t}$ are the standard deviation of all bids and the losing bids not including the lowest one of tender t respectively; n_t is the number of bids submitted in tender t ; $b_{max,t}$ is the highest bid of tender t respectively; $b_{i,t}$ and $b_{j,t}$ are the i -th and the j -th bids of tender t , which are ordered from lowest to highest bids; i_t is the rank of bid i . We calculate the screens using the following formulas:

$$CV_t = \frac{sd_t}{\bar{b}_t}, \quad (\text{A.1})$$

$$SPD_t = \frac{b_{max,t} - b_{1,t}}{b_{1,t}}, \quad (\text{A.2})$$

$$DIFFP_t = \frac{b_{2,t} - b_{1,t}}{b_{1,t}}, \quad (\text{A.3})$$

$$RD_t = \frac{b_{2,t} - b_{1,t}}{sd_{losingbids,t}}, \quad (\text{A.4})$$

$$RDALT_t = \frac{b_{2,t} - b_{1,t}}{\frac{(\sum_{i=2}^{n_t-1} b_{j,t} - b_{i,t})}{n_t - 2}}, \quad (\text{A.5})$$

$$RDNOR_t = \frac{b_{2,t} - b_{1,t}}{\frac{(\sum_{i=1}^{n_t-1} b_{j,t} - b_{i,t})}{n_t - 1}}, \quad (\text{A.6})$$

$$KSTEST_t = \max(D_t^+, D_t^-) \text{ with } D_t^+ = \max_i \left(\frac{b_{i,t}}{sd_t} - \frac{i_t}{n_t + 1} \right), D_t^- = \max_i \left(\frac{i_t}{n_t + 1} - \frac{b_{i,t}}{sd_t} \right) \quad (\text{A.7})$$

B Decentralized approach: Example

In the following, we exemplify the application of the decentralized approach.

Assume we have Tender E with three bids amounting to 100,000, 110,000, and 120,000 Swiss francs. A category manager, aiming to screen for suspicious behavior, can now apply the category manager's tool of the decentralized approach. Specifically, she applies the classification tree presented in Figure 4. To use the classification tree, she computes the coefficient of variation, the standard deviation divided by the mean (see Appendix A). The standard deviation of Tender E amounts to 10,000, and the mean of Tender E to 110,000. Therefore, the coefficient of variation amounts to (approximately) 0.091 (10,000 divided by 110,000). The category manager now follows the left branch of the classification tree, as 0.091 is greater than 0.053 ("yes"). She arrives at the terminal leaf "0": The tool classifies this tender as competitive. Simply put, the observed bidding pattern does not resemble a cartel. In conclusion, the category manager does not have to hand Tender E over to the superior category manager for further investigations (see also Figure 2).

C Training the algorithms




Figure C.1: Variable importance plot of the random forest (note: we express the importance of predictor variables relative to the maximum)

D Additional Figures and Tables

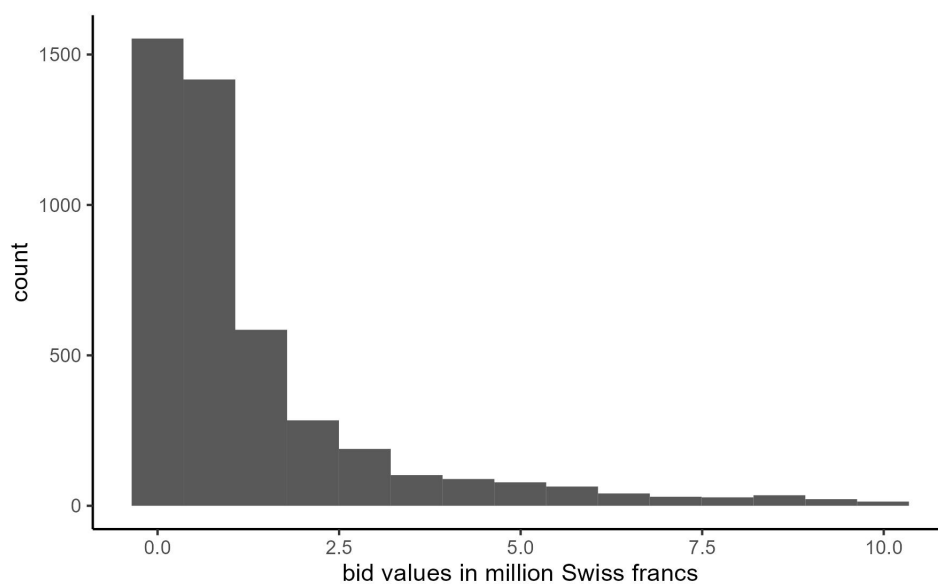


Figure D.1: Bid values in million Swiss francs (note: bid values higher 10 million Swiss francs, i.e., 5.4% of all bids, have been removed for the plot)

Bibliography

- ABRANTES-METZ, R., L. M. FROEB, J. GEWEKE, AND C. T. TAYLOR (2006): “A variance screen for collusion,” *International Journal of Industrial Organization*, 24(3), 467–486.
- ATHEY, S., AND G. W. IMBENS (2019): “Machine learning methods that economists should know about,” *Annual Review of Economics*, 11, 685–725.
- BAJARI, P., AND L. YE (2003): “Deciding between Competition and Collusion,” *The Review of Economics and Statistics*, 85(4), 971–989.
- BÉKÉS, G., AND G. KÉZDI (2021): *Data Analysis for Business, Economics, and Policy*. Cambridge University Press.
- BERKSON, J. (1944): “Application of the logistic function to bio-assay,” *Journal of the American statistical association*, 39(227), 357–365.
- BETH, H., AND O. GANNON (2022): “Cartel screening—can competition authorities and corporations afford not to use big data to detect cartels?,” *Competition Law & Policy Debate*, 7(2), 77–88.
- BLATTER, M., W. EMONS, AND S. STICHER (2018): “Optimal leniency programs when firms have cumulative and asymmetric evidence,” *Review of industrial organization*, 52(3), 403–427.
- BREIMAN, L. (2001): “Random Forests,” *Machine Learning*, 45(1), 5–32.
- BREIMAN, L., J. FRIEDMAN, R. OLSHEN, AND C. STONE (1984): *Classification and Regression Trees*. Wadsworth.
- CHASSANG, S., K. KAWAI, J. NAKABAYASHI, AND J. ORTNER (2022): “Robust screens for noncompetitive bidding in procurement auctions,” *Econometrica*, 90(1), 315–346.
- CLARK, R., D. COVIELLO, J. F. GAUTHIER, AND A. SHNEYEROV (2018): “Bid rigging and entry deterrence in public procurement: Evidence from an investigation into collusion and corruption in Quebec,” *Journal of Law, Economics, and Organization*, 34(3), 301–363.
- CROZET, Y., C. NASH, AND J. PRESTON (2012): “Beyond the quiet life of a natural monopoly: Regulatory challenges ahead for Europe’s rail sector,” *Policy paper CERRE*, 24.

- EUROPEAN COMMISSION (2021): “New Action Plan: boosting long-distance and cross-border passenger rail,” Discussion paper.
- FINGER, M. (2014): “Governance of competition and performance in European railways: An analysis of five cases,” *Utilities Policy*, 31, 278–288.
- FRANK, L. E., AND J. H. FRIEDMAN (1993): “A statistical view of some chemometrics regression tools,” *Technometrics*, 35(2), 109–135.
- FRIEDMAN, J., T. HASTIE, R. TIBSHIRANI, B. NARASIMHAN, K. TAY, N. SIMON, J. QIAN, AND J. YANG (2022): “glmnet: Lasso and Elastic-Net Regularized Generalized Linear Models,” *R package*.
- FRIEDMAN, J. H. (2001): “Greedy function approximation: a gradient boosting machine,” *Annals of statistics*, pp. 1189–1232.
- HARRINGTON, J. E. (2008): *Detecting cartels*.
- HARRINGTON, J. E., AND D. IMHOF (2022): “Cartel Screening and Machine Learning,” *Stanford Computational Antitrust*.
- HOTHORN, A. Z., K. HORNIK, C. STROBL, AND A. ZEILEIS (2022): “party: A Laboratory for Recursive Partytioning,” *R package*.
- HUBER, M., AND D. IMHOF (2019): “Machine learning with screens for detecting bid-rigging cartels,” *International Journal of Industrial Organization*, 65, 277–301.
- HUBER, M., D. IMHOF, AND R. ISHII (2022): “Transnational machine learning with screens for flagging bid-rigging cartels,” *Journal of the Royal Statistical Society: Series A (Statistics in Society)*.
- IMHOF, D. (2019): “Detecting Bid-Rigging Cartels With Descriptive Statistics,” *Journal of Competition Law & Economics*, 15(4), 427–467.
- IMHOF, D., Y. KARAGÖK, AND S. RUTZ (2018): “Screening for bid rigging-does it work?,” *Journal of Competition Law & Economics*, 14(2), 235–261.
- IMHOF, D., AND H. WALLIMANN (2021): “Detecting bid-rigging coalitions in different countries and auction formats,” *International Review of Law and Economics*, 68, 106016.
- INTERNATIONAL TRANSPORT FORUM (2022): “Transport infrastructure investment and maintenance,” *data base*.
- JIMENEZ, J. L., AND J. PERDIGUERO (2012): “Does Rigidity of Price Hide Collusion?,” *Review of Industrial Organization*, 41(3), 223–248.
- KLEINBERG, J., J. LUDWIG, S. MULLAINATHAN, AND Z. OBERMEYER (2015): “Prediction policy problems,” *American Economic Review*, 105(5), 491–95.
- KUHN, M. (2019): “caret: Classification and Regression Training,” *R package*.

- LAFFONT, J.-J., AND J. TIROLE (1993): *A theory of incentives in procurement and regulation*. MIT press.
- LENGWILER, Y., AND E. G. WOLFSTETTER (2006): “Corruption in procurement auctions,” *Handbook of Procurement*.
- LIAW, A., AND M. WIENER (2018): “Package randomForest: Breiman and Cutler’s Random Forest for Classification and Regression,” *R package*.
- LITRA (2022): “Verkehrszahlen - Ausgaben 2021,” Discussion paper.
- MCCULLOCH, W. S., AND W. PITTS (1943): “A logical calculus of the ideas immanent in nervous activity,” *The bulletin of mathematical biophysics*, 5(4), 115–133.
- MONTERO, J., AND M. FINGER (2020): “Railway Regulation: A Comparative analysis of a diverging reality,” in *Handbook on Railway Regulation*, pp. 1–20. Edward Elgar Publishing.
- OECD (2022): “Data Screening Tools in Competition Investigations,” *OECD Competition Policy Roundtable Background Note*.
- PORTER, R. H., AND J. D. ZONA (1993): “Detection of Bid Rigging in Procurement Auctions,” *Journal of Political Economy*, 101(3), 518–538.
- R CORE TEAM (2013): “R: A language and environment for statistical computing,” .
- RIPLEY, B., AND W. VENABLES (2022): “nnet: Feed-Forward Neural Networks and Multinomial Log-Linear Models,” *R package*.
- RODRÍGUEZ, M. J. G., V. RODRÍGUEZ-MONTEQUÍN, P. BALLESTEROS-PÉREZ, P. E. LOVE, AND R. SIGNOR (2022): “Collusion detection in public procurement auctions with machine learning algorithms,” *Automation in Construction*, 133, 104047.
- SILVEIRA, D., L. B. DE MORAES, E. P. FIUZA, AND D. O. CAJUEIRO (2023): “Who Are You? Cartel Detection Using Unlabeled Data,” *International Journal of Industrial Organization*.
- SILVEIRA, D., S. VASCONCELOS, M. RESENDE, AND D. O. CAJUEIRO (2022): “Won’t Get Fooled Again: A supervised machine learning approach for screening gasoline cartels,” *Energy Economics*, 105, 105711.
- TIBSHIRANI, R. (1996): “Regression Shrinkage and Selection via The Lasso,” *Journal of the Royal Statistical Society, Series B*, 58(1), 267–288.
- VAN DER LAAN, M. J., E. C. POLLEY, AND A. E. HUBBARD (2007): “Super Learner,” *Statistical Applications in Genetics and Molecular Biology*, 6.
- WALLIMANN, H., D. IMHOF, AND M. HUBER (2022): “A machine learning approach for flagging incomplete bid-rigging cartels,” *Computational Economics*.
- YUAN, J. (2022): “xgboost: Extreme Gradient Boosting,” *R package*.